Kan extensions in probability theory

Ruben Van Belle

Marunouchi Quantitative Finance seminar 73, March 2023

Overview

1. Category theory

Overview

1. Category theory
   - In particular, Kan extensions:

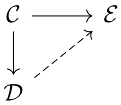

Overview

1. Category theory
   - In particular, Kan extensions:

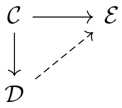

2. Probability monads as codensity monads

Overview

1. Category theory
   - In particular, Kan extensions:

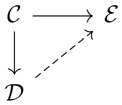

2. Probability monads as codensity monads
3. A categorical proof of the Radon-Nikodym theorem

Overview

1. Category theory
   - In particular, Kan extensions:

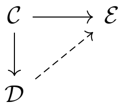

2. Probability monads as codensity monads
3. A categorical proof of the Radon-Nikodym theorem
   - Martingale convergence theorem

1.1 Category theory

Categories, functors and natural transformations

Definition: A **category** $\mathcal{C}$ is a collection of *objects* and for every pair of objects (A, B) a collection $\mathcal{C}(A, B)$ of *morphisms* between them together with a composition operation.

Categories, functors and natural transformations

Definition: A **category** $\mathcal{C}$ is a collection of *objects* and for every pair of objects (A, B) a collection $\mathcal{C}(A, B)$ of *morphisms* between them together with a composition operation.

Examples:

- **Set** is the category of *sets* and *functions*.

Categories, functors and natural transformations

Definition: A **category** $\mathcal{C}$ is a collection of *objects* and for every pair of objects (A, B) a collection $\mathcal{C}(A, B)$ of *morphisms* between them together with a composition operation.
Examples:

- **Set** is the category of *sets* and *functions*.
- **Grp** is the category of *groups* and *group homomorphisms*.

Categories, functors and natural transformations

Definition: A **category** $\mathcal{C}$ is a collection of *objects* and for every pair of objects (A, B) a collection $\mathcal{C}(A, B)$ of *morphisms* between them together with a composition operation.

Examples:

- **Set** is the category of *sets* and *functions*.
- **Grp** is the category of *groups* and *group homomorphisms*.
- **Mble** is the category of *measurable spaces* and *measurable maps*.

Categories, functors and natural transformations

Definition: A **category** $\mathcal{C}$ is a collection of *objects* and for every pair of objects (A, B) a collection $\mathcal{C}(A, B)$ of *morphisms* between them together with a composition operation.

Examples:

- **Set** is the category of *sets* and *functions*.
- **Grp** is the category of *groups* and *group homomorphisms*.
- **Mble** is the category of *measurable spaces* and *measurable maps*.
- **Prob** is the category of *probability spaces* and *measure-preserving maps*

Categories, functors and natural transformations

Definition: A **category** $\mathcal{C}$ is a collection of *objects* and for every pair of objects (A, B) a collection $\mathcal{C}(A, B)$ of *morphisms* between them together with a composition operation.

Examples:

- **Set** is the category of *sets* and *functions*.
- **Grp** is the category of *groups* and *group homomorphisms*.
- **Mble** is the category of *measurable spaces* and *measurable maps*.
- **Prob** is the category of *probability spaces* and *measure-preserving maps*
- A category that has precisely *one* object is a *monoid*.

Categories, functors and natural transformations

Definition: A **category** $\mathcal{C}$ is a collection of *objects* and for every pair of objects (A, B) a collection $\mathcal{C}(A, B)$ of *morphisms* between them together with a composition operation.

Examples:

- **Set** is the category of *sets* and *functions*.
- **Grp** is the category of *groups* and *group homomorphisms*.
- **Mble** is the category of *measurable spaces* and *measurable maps*.
- **Prob** is the category of *probability spaces* and *measure-preserving maps*
- A category that has precisely *one* object is a *monoid*.
- Every *parially ordered set* is a category
- **Cat** is the category of *categories* and *functors*.
- The *functor category* $[\mathcal{C}, \mathcal{D}]$ of *functors* $\mathcal{C} \to \mathcal{D}$ and *natural transformations*.

Categories, functors and natural transformations

Let $\mathcal{C}$ and $\mathcal{D}$ be categories.

Categories, functors and natural transformations

Let $\mathcal{C}$ and $\mathcal{D}$ be categories.
Definition: A **functor** $F : \mathcal{C} \to \mathcal{D}$ is an assignment on objects and a composition-preserving assignment on morphisms.

Categories, functors and natural transformations

Let $\mathcal{C}$ and $\mathcal{D}$ be categories.
Definition: A **functor** $F : \mathcal{C} \to \mathcal{D}$ is an assignment on objects and a composition-preserving assignment on morphisms.
Examples:

- The *forgetful* functor U : **Prob** $\to$ **Mble** sends a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ to its underlying measurable space $(\Omega, \mathcal{F})$.
- The functor **Top** $\to$ **Mble** that sends a topological space $(X, \mathcal{T})$ to the $(X, \sigma(\mathcal{T}))$.

Categories, functors and natural transformations

Let $\mathcal{C}$ and $\mathcal{D}$ be categories.
Definition: A **functor** $F : \mathcal{C} \to \mathcal{D}$ is an assignment on objects and a composition-preserving assignment on morphisms.
Examples:

- The *forgetful* functor U : **Prob** $\to$ **Mble** sends a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ to its underlying measurable space $(\Omega, \mathcal{F})$.
- The functor **Top** $\to$ **Mble** that sends a topological space $(X, \mathcal{T})$ to the $(X, \sigma(\mathcal{T}))$.
- The functor **Top** $\to$ **Grpd** that sends a topological space X to its fundamental groupoid $\pi_1(X)$.
- The *Giry functor* $\mathcal{G}$: **Mble** $\to$ **Mble** that sends a measurable space $(\Omega, \mathcal{F})$ to the space of all probability measures on $(\Omega, \mathcal{F})$.

Categories, functors and natural transformations

Let $\mathcal{C}$ and $\mathcal{D}$ be categories.
Definition: A **functor** $F : \mathcal{C} \to \mathcal{D}$ is an assignment on objects and a composition-preserving assignment on morphisms.
Examples:

- The *forgetful* functor U : **Prob** $\to$ **Mble** sends a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ to its underlying measurable space $(\Omega, \mathcal{F})$.
- The functor **Top** $\to$ **Mble** that sends a topological space $(X, \mathcal{T})$ to the $(X, \sigma(\mathcal{T}))$.
- The functor **Top** $\to$ **Grpd** that sends a topological space X to its fundamental groupoid $\pi_1(X)$.
- The *Giry functor* $\mathcal{G}$: **Mble** $\to$ **Mble** that sends a measurable space $(\Omega, \mathcal{F})$ to the space of all probability measures on $(\Omega, \mathcal{F})$.
- For an object A in $\mathcal{C}$, there is a functor $\mathcal{C}(A, -) : \mathcal{C} \to$ **Set** that sends B to the set of morphisms from A to B.

Categories, functors and natural transformations

Let F and G be functors between categories $\mathcal{C}$ and $\mathcal{D}$.

Categories, functors and natural transformations

Let F and G be functors between categories $\mathcal{C}$ and $\mathcal{D}$.
<u>Definition</u>: A **natural transformation** $\tau : F \to G$ is a collection of morphisms $(\tau_C : Fc \to Gc)_{c\in\mathcal{C}}$ such that

$$\begin{array}{ccc} Fc & \xrightarrow{\tau_c} & Gc \\ {\scriptstyle Ff}\big\downarrow & & \big\downarrow{\scriptstyle Gf} \\ Fd & \xrightarrow[\tau_d]{} & Gd \end{array}$$

commutes for all morphisms $f : c \to d$.

Categories, functors and natural transformations

Examples:

- Let $1 : \mathbf{Mble} \to \mathbf{Mble}$ be the *identity functor* and let $\mathcal{G} : \mathbf{Mble} \to \mathbf{Mble}$ be the *Giry functor*.

Categories, functors and natural transformations

Examples:

- Let $1 : \mathbf{Mble} \to \mathbf{Mble}$ be the *identity functor* and let $\mathcal{G} : \mathbf{Mble} \to \mathbf{Mble}$ be the *Giry functor*. For a measurable space X, there is a map

$$\eta_X : X \to \mathcal{G}X : x \mapsto \delta_x.$$

Categories, functors and natural transformations

Examples:

- Let $1 : \mathbf{Mble} \to \mathbf{Mble}$ be the *identity functor* and let $\mathcal{G} : \mathbf{Mble} \to \mathbf{Mble}$ be the *Giry functor*. For a measurable space X, there is a map

$$\eta_X : X \to \mathcal{G}X : x \mapsto \delta_x.$$

  These form a natural transformation $\eta : 1 \to \mathcal{G}$.

- Let $g : A_2 \to A_1$ be a morphims in $\mathcal{C}$.

Categories, functors and natural transformations

Examples:

- Let $1 : \mathbf{Mble} \to \mathbf{Mble}$ be the *identity functor* and let $\mathcal{G} : \mathbf{Mble} \to \mathbf{Mble}$ be the *Giry functor*. For a measurable space X, there is a map

$$\eta_X : X \to \mathcal{G}X : x \mapsto \delta_x.$$

  These form a natural transformation $\eta : 1 \to \mathcal{G}$.

- Let $g : A_2 \to A_1$ be a morphims in $\mathcal{C}$. For every object B there are functions

$$\mathrm{Hom}(A_1, B) \to \mathrm{Hom}(A_2, B) : f \mapsto f \circ g.$$

Categories, functors and natural transformations

Examples:

- Let $1 : \mathbf{Mble} \to \mathbf{Mble}$ be the *identity functor* and let $\mathcal{G} : \mathbf{Mble} \to \mathbf{Mble}$ be the *Giry functor*. For a measurable space X, there is a map

$$\eta_X : X \to \mathcal{G}X : x \mapsto \delta_x.$$

  These form a natural transformation $\eta : 1 \to \mathcal{G}$.

- Let $g : A_2 \to A_1$ be a morphims in $\mathcal{C}$. For every object B there are functions

$$\mathrm{Hom}(A_1, B) \to \mathrm{Hom}(A_2, B) : f \mapsto f \circ g.$$

  These form a natural transformation $\mathcal{C}(A_1, -) \to \mathcal{C}(A_2, -)$.

Categories, functors and natural transformations

Examples:

- Let $1 : \mathbf{Mble} \to \mathbf{Mble}$ be the *identity functor* and let $\mathcal{G} : \mathbf{Mble} \to \mathbf{Mble}$ be the *Giry functor*. For a measurable space X, there is a map

$$\eta_X : X \to \mathcal{G}X : x \mapsto \delta_x.$$

  These form a natural transformation $\eta : 1 \to \mathcal{G}$.
- Let $g : A_2 \to A_1$ be a morphims in $\mathcal{C}$. For every object B there are functions

$$\mathrm{Hom}(A_1, B) \to \mathrm{Hom}(A_2, B) : f \mapsto f \circ g.$$

  These form a natural transformation $\mathcal{C}(A_1, -) \to \mathcal{C}(A_2, -)$. By the Yoneda lemma, every such natural transformation is of this form.

Categories, functors and natural transformations

$\mathcal{C}$ $\mathcal{E}$

$\mathcal{D}$

- Categories

Categories, functors and natural transformations

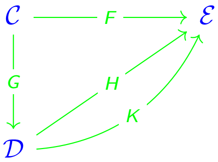

- Categories
- Functors

Categories, functors and natural transformations

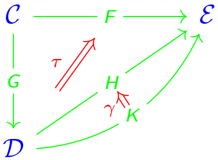

- Categories
- Functors
- Natural transformations

Universal constructions

Universal constructions are objects that are *optimal* in some sense.

Universal constructions

Universal constructions are objects that are *optimal* in some sense. We can think them as *maximal* or *minimal* approximations. Their maximality/minimality implies *uniqueness up to isomorphism*.

Universal constructions

Universal constructions are objects that are *optimal* in some sense. We can think them as *maximal* or *minimal* approximations. Their maximality/minimality implies *uniqueness up to isomorphism*. These are usually *limits* and *colimits* (*ends* and *coends*).

Examples:

- **Products** in **Set**: For sets A and B, there are functions $\pi_1 : A \times B \to A$ and $\pi_2 : A \times B \to B$.

Universal constructions

Universal constructions are objects that are *optimal* in some sense. We can think them as *maximal* or *minimal* approximations. Their maximality/minimality implies *uniqueness up to isomorphism*. These are usually *limits* and *colimits* (*ends* and *coends*).

Examples:

- **Products** in **Set**: For sets A and B, there are functions $\pi_1 : A \times B \to A$ and $\pi_2 : A \times B \to B$. For any other object C with morphims $p_1 : C \to A$ and $p_2 : C \to B$, *there is a unique morphism* $C \to A \times B$ such that

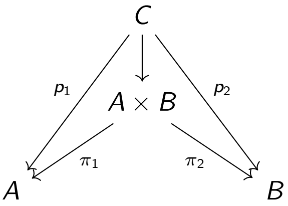

commutes.

Universal constructions

- **Suprema (colimits)**: Let $S \subset \mathbb{R}$ be a bounded, non-empty subset.

Universal constructions

- **Suprema (colimits)**: Let $S \subset \mathbb{R}$ be a bounded, non-empty subset. We have

$$s \leq \sup S$$

for all s.

Universal constructions

- **Suprema (colimits)**: Let $S \subset \mathbb{R}$ be a bounded, non-empty subset. We have

$$s \leq \sup S$$

for all s.
Suppose $s \leq r$ for all s, then

$$\sup S \leq r$$

.

Universal constructions

- **Filtrations (filtered limits)**: Let $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n\in\mathbb{N}}, \mathbb{P}))$ be a filtered probability space.

Universal constructions

- **Filtrations (filtered limits)**: Let $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n\in\mathbb{N}}, \mathbb{P}))$ be a filtered probability space.

$$(\Omega, \mathcal{F}_1, \mathbb{P}\mid_{\mathcal{F}_1}) \longleftarrow (\Omega, \mathcal{F}_2, \mathbb{P}\mid_{\mathcal{F}_2}) \longleftarrow \dots \longleftarrow (\Omega, \mathcal{F}, \mathbb{P})$$

Universal constructions

- **Filtrations (filtered limits)**: Let $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n\in\mathbb{N}}, \mathbb{P}))$ be a filtered probability space. Suppose

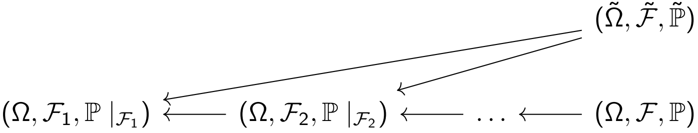

Universal constructions

- **Filtrations (filtered limits)**: Let $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n\in\mathbb{N}}, \mathbb{P}))$ be a filtered probability space. Then there exists a unique

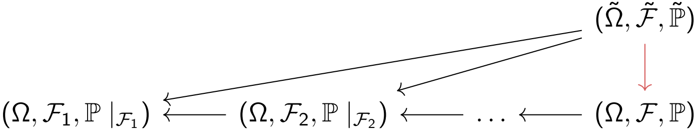

1.2 Kan extensions

Universal property

Let $F : \mathcal{C} \to \mathcal{E}$ and $G : \mathcal{C} \to \mathcal{D}$ be functors.

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{F} & \mathcal{E} \\ {\scriptstyle G}\big\downarrow & & \\ \mathcal{D} & & \end{array}$$

Universal property

Let $F : \mathcal{C} \to \mathcal{E}$ and $G : \mathcal{C} \to \mathcal{D}$ be functors.
The **right Kan extension of** F **along** G is a functor $H : \mathcal{D} \to \mathcal{E}$ together with a natural transformation $\epsilon : H \circ G \Rightarrow F$.

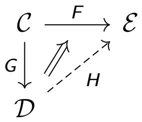

Universal property

Let $F : \mathcal{C} \to \mathcal{E}$ and $G : \mathcal{C} \to \mathcal{D}$ be functors.
The **right Kan extension of** F **along** G is a functor $\text{Ran}_G F : \mathcal{D} \to \mathcal{E}$ together with a natural transformation $\epsilon : \text{Ran}_G F \circ G \Rightarrow F$.
Such that for every other functor $\tilde{H} : \mathcal{D} \to \mathcal{E}$ with a natural transformation $\tilde{\epsilon} : \tilde{H} \circ G \Rightarrow F$,

Universal property

Let $F : \mathcal{C} \to \mathcal{E}$ and $G : \mathcal{C} \to \mathcal{D}$ be functors.
The **right Kan extension of** F **along** G is a functor $\mathrm{Ran}_G F : \mathcal{D} \to \mathcal{E}$ together with a natural transformation $\epsilon : \mathrm{Ran}_G F \circ G \Rightarrow F$.
Such that for every other functor $\tilde{H} : \mathcal{D} \to \mathcal{E}$ with a natural transformation $\tilde{\epsilon} : \tilde{H} \circ G \Rightarrow F$, *there exists a unique* natural transformation $\gamma : \mathrm{Ran}_G F \to \tilde{H}$ such that

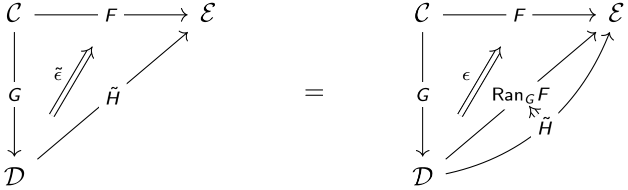

Kan extensions using limits and ends

For functors $F : \mathcal{C} \to \mathcal{E}$ and $G : \mathcal{C} \to \mathcal{D}$.

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{F} & \mathcal{E} \\ {\scriptstyle G}\big\downarrow & & \\ \mathcal{D} & & \end{array}$$

Kan extensions using limits and ends

For functors $F : \mathcal{C} \to \mathcal{E}$ and $G : \mathcal{C} \to \mathcal{D}$.

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{F} & \mathcal{E} \\ {\scriptstyle G}\downarrow & & \\ \mathcal{D} & & \end{array}$$

Let d be an object in $\mathcal{D}$,

Kan extensions using limits and ends

For functors $F : \mathcal{C} \to \mathcal{E}$ and $G : \mathcal{C} \to \mathcal{D}$.

$$\begin{array}{ccc} \mathcal{C} & \xrightarrow{F} & \mathcal{E} \\ {\scriptstyle G}\downarrow & & \\ \mathcal{D} & & \end{array}$$

Let d be an object in $\mathcal{D}$,

$$\begin{aligned} \mathrm{Ran}_G F(d) &= \lim(d \downarrow G \to \mathcal{C} \xrightarrow{F} \mathcal{E}) \\ &= \int_{c \in \mathcal{C}} [\mathcal{C}(d, Gc), Fc] \end{aligned}$$

Example

Let A and B be *partially ordered sets*. Let $f : A \to \mathbb{R}$ and $g : B \to \mathbb{R}$ be *order-preserving maps*.

$$\begin{array}{ccc} A & \xrightarrow{f} & \mathbb{R} \\ {\scriptstyle g}\downarrow & & \\ B & & \end{array}$$

Example

Let A and B be *partially ordered sets*. Let $f : A \to \mathbb{R}$ and $g : B \to \mathbb{R}$ be *order-preserving maps*.

$$\begin{array}{ccc} A & \xrightarrow{f} & \mathbb{R} \\ {\scriptstyle g}\big\downarrow & & \\ B & & \end{array}$$

Then for $b \in B$,

$$\begin{aligned} \mathrm{Ran}_g f(b) &= \lim(b \downarrow g \to A \xrightarrow{f} \mathbb{R}) \\ &= \inf\{f(a) \mid b \leq g(a)\} \end{aligned}$$

2 The Giry functor as a Kan extensions

The Giry functor

Let X be a measurable space.

The Giry functor

Let X be a measurable space. Define

$$\mathcal{G}X := \{\text{probability measures on } X\}$$

together with the σ-algebra generated by the evaluation maps

$$\mathsf{ev} : \mathcal{G}X \to [0,1] : \mathbb{P} \mapsto \mathbb{P}(A).$$

The Giry functor

Let X be a measurable space. Define

$$\mathcal{G}X := \{\text{probability measures on } X\}$$

together with the σ-algebra generated by the evaluation maps

$$\text{ev} : \mathcal{G}X \to [0,1] : \mathbb{P} \mapsto \mathbb{P}(A).$$

For a measurable map $f : X \to Y$, define $\mathcal{G}f : \mathcal{G}X \to \mathcal{G}Y$ by the assignment

$$\mathbb{P} \mapsto \mathbb{P} \circ f^{-1}.$$

The Giry functor

Let X be a measurable space. Define

$$\mathcal{G}X := \{\text{probability measures on } X\}$$

together with the σ-algebra generated by the evaluation maps

$$\text{ev} : \mathcal{G}X \to [0,1] : \mathbb{P} \mapsto \mathbb{P}(A).$$

For a measurable map $f : X \to Y$, define $\mathcal{G}f : \mathcal{G}X \to \mathcal{G}Y$ by the assignment

$$\mathbb{P} \mapsto \mathbb{P} \circ f^{-1}.$$

This defines a functor $\mathcal{G}$: **Mble** $\to$ **Mble**, and is called the **Giry functor**.

The Giry functor

Let X be a measurable space. Define

$$\mathcal{G}X := \{\text{probability measures on } X\}$$

together with the σ-algebra generated by the evaluation maps

$$\text{ev} : \mathcal{G}X \to [0,1] : \mathbb{P} \mapsto \mathbb{P}(A).$$

For a measurable map $f : X \to Y$, define $\mathcal{G}f : \mathcal{G}X \to \mathcal{G}Y$ by the assignment

$$\mathbb{P} \mapsto \mathbb{P} \circ f^{-1}.$$

This defines a functor $\mathcal{G}$: **Mble** $\to$ **Mble**, and is called the **Giry functor**.
Remark: This endofunctor is the underlying functor of the *Giry monad*.

The Giry functor as Kan extension

Define the functor g as

$$\mathbf{Set}_c \to \mathbf{Mble} \xrightarrow{\mathcal{G}} \mathbf{Mble}.$$

Theorem

The Giry functor $\mathcal{G}$ is the right Kan extension of g along itself.

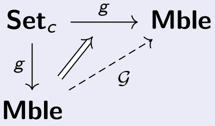

The Giry functor as Kan extension

Define the functor g as

$$\mathbf{Set}_c \to \mathbf{Mble} \xrightarrow{\mathcal{G}} \mathbf{Mble}.$$

Theorem

The Giry functor $\mathcal{G}$ is the right Kan extension of g along itself.

$$\begin{array}{ccc} \mathbf{Set}_c & \xrightarrow{\;g\;} & \mathbf{Mble} \\ {\scriptstyle g}\downarrow & \Uparrow & \nearrow {\scriptstyle \mathcal{G}} \\ \mathbf{Mble} & & \end{array}$$

Meaning: *Probability measures arise naturally as the categorical extension of the more intuitive probability measures on countable sets.*

Codensity monads

Any right Kan extension along itself can be given a canonical monad structure, these are **codensity monads**.

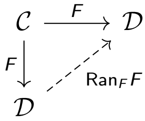

Codensity monads

Any right Kan extension along itself can be given a canonical monad structure, these are **codensity monads**.

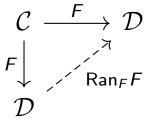

The Giry monad arises as a codensity monad, by the previous result.

3.1 Radon-Nikodym theorem

Radon-Nikodym theorem: finite version

We will give a proof for a special case.

Radon-Nikodym theorem: finite version

We will give a proof for a special case. Let A be a finite set and $(p_a)_{a \in A}$ a probability measure on A.

Radon-Nikodym theorem: finite version

We will give a proof for a special case. Let A be a finite set and $(p_a)_{a \in A}$ a probability measure on A. Let q be a measure on A such that $q \ll p$.

Radon-Nikodym theorem: finite version

We will give a proof for a special case. Let A be a finite set and $(p_a)_{a \in A}$ a probability measure on A. Let q be a measure on A such that $q \ll p$. Define a map $f : A \to \mathbb{R}$ by

$$a \mapsto \begin{cases} \frac{q_a}{p_a} \text{ if } p_a \neq 0 \\ 0 \text{ otherwise.} \end{cases}$$

Radon-Nikodym theorem: finite version

We will give a proof for a special case. Let A be a finite set and $(p_a)_{a \in A}$ a probability measure on A. Let q be a measure on A such that $q \ll p$.
Define a map $f : A \to \mathbb{R}$ by

$$a \mapsto \begin{cases} \frac{q_a}{p_a} \text{ if } p_a \neq 0 \\ 0 \text{ otherwise.} \end{cases}$$

It can be checked that f is the Radon-Nikodym derivative of q with respect to p.

Categorically extending the finite version

Let **Prob** be the category of probabilty spaces and measure preserving maps.

Categorically extending the finite version

Let **Prob** be the category of probabilty spaces and measure preserving maps. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space.

Categorically extending the finite version

Let **Prob** be the category of probabilty spaces and measure preserving maps. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space.

- Define $M_n(\Omega, \mathcal{F}, \mathbb{P})$ as the set

$$\{\mu \mid \mu \leq n\mathbb{P}\},$$

  together with the *total variation metric*.
- Define $RV_n(\Omega, \mathcal{F}, \mathbb{P})$ as the set

$$\mathbf{Mble}(\Omega, [0, n]) / =_{\mathbb{P}},$$

  together with the L^1-*metric* (multiplied by a factor 1/2).

Categorically extending the finite version

Let **Prob** be the category of probabilty spaces and measure preserving maps. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space.

- Define $M_n(\Omega, \mathcal{F}, \mathbb{P})$ as the set

$$\{\mu \mid \mu \leq n\mathbb{P}\},$$

  together with the *total variation metric*.
- Define $RV_n(\Omega, \mathcal{F}, \mathbb{P})$ as the set

$$\mathbf{Mble}(\Omega, [0, n])/ =_{\mathbb{P}},$$

  together with the L^1-*metric* (multiplied by a factor 1/2).

These are *complete* metric spaces (Riesz-Fischer).

Categorically extending the finite version

Let $\mathbf{Prob}_f$ be the full subcategory of **Prob** of **finite** probability spaces.

Categorically extending the finite version

Let $\mathbf{Prob}_f$ be the full subcategory of **Prob** of **finite** probability spaces. Let $s : (A, p) \to (B, q)$ be a measure-preserving map of finite probability spaces.

- Define $M_n^f(s) : M_n(A, p) \to M_n(B, q)$ by the assignment

$$r \mapsto r \circ s^{-1}.$$

Categorically extending the finite version

Let $\mathbf{Prob}_f$ be the full subcategory of **Prob** of **finite** probability spaces. Let $s : (A, p) \to (B, q)$ be a measure-preserving map of finite probability spaces.

- Define $M_n^f(s) : M_n(A, p) \to M_n(B, q)$ by the assignment

$$r \mapsto r \circ s^{-1}.$$

- Define $RV_n^f(s) : RV_n(A, p) \to RV_n(B, q)$ by sending a map $g : A \to [0, n]$ to the map $B \to [0, n]$, which is defined by

$$b \mapsto \begin{cases} \frac{1}{q_b} \sum_{s(a)=b} p_a g(a) & \text{if } q_b \neq 0 \\ 0 & \text{otherwise.} \end{cases}$$

Categorically extending the finite version

Let $\mathbf{Prob}_f$ be the full subcategory of **Prob** of **finite** probability spaces. Let $s : (A, p) \to (B, q)$ be a measure-preserving map of finite probability spaces.

- Define $M_n^f(s) : M_n(A, p) \to M_n(B, q)$ by the assignment

$$r \mapsto r \circ s^{-1}.$$

- Define $RV_n^f(s) : RV_n(A, p) \to RV_n(B, q)$ by sending a map $g : A \to [0, n]$ to the map $B \to [0, n]$, which is defined by

$$b \mapsto \begin{cases} \frac{1}{q_b} \sum_{s(a)=b} p_a g(a) & \text{if } q_b \neq 0 \\ 0 & \text{otherwise.} \end{cases}$$

These are 1-*Lipschitz maps*.

Categorically extending the finite version

Let $\mathbf{CMet}_1$ be the category of complete metric spaces and 1-Lipschitz maps.

Categorically extending the finite version

Let $\mathbf{CMet}_1$ be the category of complete metric spaces and 1-Lipschitz maps. We have two functors:

$$\mathbf{Prob}^f \underset{RV_n^f}{\overset{M_n^f}{\rightrightarrows}} \mathbf{CMet}_1.$$

Categorically extending the finite version

Let $\mathbf{CMet}_1$ be the category of complete metric spaces and 1-Lipschitz maps.
We have two functors:

$$\mathbf{Prob}^f \underset{RV_n^f}{\overset{M_n^f}{\rightrightarrows}} \mathbf{CMet}_1.$$

By the finite Radon-Nikodym theorem, we see that

$$\cong$$

$$\mathbf{Prob}^f \underset{RV_n^f}{\overset{M_n^f}{\rightrightarrows}} \mathbf{CMet}_1.$$

Categorically extending the finite version

It follows that also the right Kan extensions along $i : \mathbf{Prob}_f \to \mathbf{Prob}$ are isomorphic.

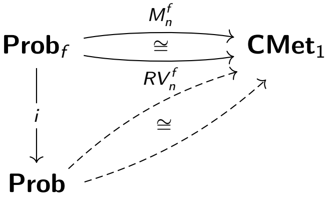

What do these Kan extensions look like?

Proposition

For a probability space $\mathbf{\Omega} := (\Omega, \mathcal{F}, \mathbb{P})$, we have for all $n \geq 1$ that

$$M_n(\mathbf{\Omega}) \to (\mathrm{Ran}_i M_n^f)(\mathbf{\Omega}),$$

is an isomorphism.

What do these Kan extensions look like?

Proposition

For a probability space $\mathbf{\Omega} := (\Omega, \mathcal{F}, \mathbb{P})$, we have for all $n \geq 1$ that

$$M_n(\mathbf{\Omega}) \to (\mathrm{Ran}_i M_n^f)(\mathbf{\Omega}),$$

is an isomorphism.

Proof (sketch): Let $\mathbf{\Omega} := (\Omega, \mathcal{F}, \mathbb{P})$ be a probability space.

$$\mathrm{Ran}_i M_n^f(\mathbf{\Omega}) \cong \int_{\mathbf{A} \in \mathbf{Prob}_f} [\mathbf{Prob}(\mathbf{\Omega}, i\mathbf{A}), M_n^f(\mathbf{A})]$$

What do these Kan extensions look like?

Proposition

For a probability space $\mathbf{\Omega}$, we have for all $n \geq 1$ that

$$(\mathrm{Ran}_i RV_n^f)(\mathbf{\Omega}) \cong RV_n(\mathbf{\Omega}).$$

The proof for this results requires some measure theory.

Radon-Nikodym theorem

Combining everything gives a *bounded* Radon-Nikodym theorem, namely

$$\begin{aligned}\{\mu \mid \mu \leq n\mathbb{P}\} = M_n(\mathbf{\Omega}) &\cong \mathrm{Ran}_i M_n^f((\mathbf{\Omega}) \\ &\cong \mathrm{Ran}_i RV_n^f(\mathbf{\Omega}) \\ &\cong RV_n(\mathbf{\Omega}) = \mathbf{Mble}(\Omega, [0, n])/ =_{\mathbb{P}}\end{aligned}$$

Radon-Nikodym theorem

Combining everything gives a *bounded* Radon-Nikodym theorem, namely

$$\begin{aligned}\{\mu \mid \mu \leq n\mathbb{P}\} = M_n(\boldsymbol{\Omega}) &\cong \mathrm{Ran}_i M_n^f((\boldsymbol{\Omega}) \\ &\cong \mathrm{Ran}_i RV_n^f(\boldsymbol{\Omega}) \\ &\cong RV_n(\boldsymbol{\Omega}) = \mathbf{Mble}(\Omega, [0, n])/ =_{\mathbb{P}}\end{aligned}$$

We can look at the colimit over all $n \geq 1$,

$$\begin{array}{ccccccccc} M_1\boldsymbol{\Omega} & \longrightarrow & M_2\boldsymbol{\Omega} & \longrightarrow & \dots & \longrightarrow & M_n\boldsymbol{\Omega} & \longrightarrow & \dots \\ \wr\Vert & & \wr\Vert & & & & \wr\Vert & & \\ RV_1\boldsymbol{\Omega} & \longrightarrow & RV_2\boldsymbol{\Omega} & \longrightarrow & \dots & \longrightarrow & RV_n\boldsymbol{\Omega} & \longrightarrow & \dots \end{array}$$

This gives us

$$\{\mu \mid \mu \ll \mathbb{P}\} \cong \{f : \Omega \to [0, \infty) \mid f \text{ is integrable}\} \; / =_{\mathbb{P}} .$$

Remark on conditional expectation

For a probability space $\mathbf{\Omega}$, we know what $(\mathrm{Ran}_i M_n^f)(\mathbf{\Omega})$ and $(\mathrm{Ran}_i RV_n^f)(\mathbf{\Omega})$ look like.

Remark on conditional expectation

For a probability space $\mathbf{\Omega}$, we know what $(\mathrm{Ran}_i M_n^f)(\mathbf{\Omega})$ and $(\mathrm{Ran}_i RV_n^f)(\mathbf{\Omega})$ look like.
What can we say about $M_n(g) := (\mathrm{Ran}_i M_n^f)(g)$ and $RV_n(g) := (\mathrm{Ran}_i RV_n^f)(g)$ for $g : \mathbf{\Omega}_1 \to \mathbf{\Omega}_2$?

Remark on conditional expectation

For a probability space $\mathbf{\Omega}$, we know what $(\mathrm{Ran}_i M_n^f)(\mathbf{\Omega})$ and $(\mathrm{Ran}_i RV_n^f)(\mathbf{\Omega})$ look like.
What can we say about $M_n(g) := (\mathrm{Ran}_i M_n^f)(g)$ and $RV_n(g) := (\mathrm{Ran}_i RV_n^f)(g)$ for $g : \mathbf{\Omega}_1 \to \mathbf{\Omega}_2$?
They are the unique morphisms such that

$$\begin{array}{ccc} M_n\mathbf{\Omega}_1 & \xrightarrow{M_n(g)} & M_n\mathbf{\Omega}_2 \\ & \searrow \quad \swarrow & \\ & M_n^f\mathbf{A} & \end{array} \qquad \begin{array}{ccc} RV_n\mathbf{\Omega}_1 & \xrightarrow{RV_n(g)} & RV_n\mathbf{\Omega}_2 \\ & \searrow \quad \swarrow & \\ & RV_n^f\mathbf{A} & \end{array}$$

commute for morphisms $\mathbf{\Omega}_2 \to \mathbf{A}$.

Remark on conditional expectation

In particular, these commute for all $1_E : \mathbf{\Omega_2} \to \mathbf{2}_E$.

Remark on conditional expectation

In particular, these commute for all $1_E : \mathbf{\Omega_2} \to \mathbf{2}_E$. We conclude that for all $E \in \mathcal{F}_2$

$$M_n(g)(\mu) \circ 1_E^{-1} = \mu \circ 1_{g^{-1}(E)}^{-1},$$

and

$$\int_E RV_n(g)(f) \mathrm{d}\mathbb{P}_2 = \int_{g^{-1}(E)} f \mathrm{d}\mathbb{P}_1$$

Remark on conditional expectation

In particular, these commute for all $1_E : \mathbf{\Omega_2} \to \mathbf{2}_E$. We conclude that for all $E \in \mathcal{F}_2$

$$M_n(g)(\mu) \circ 1_E^{-1} = \mu \circ 1_{g^{-1}(E)}^{-1},$$

and

$$\int_E RV_n(g)(f) \mathrm{d}\mathbb{P}_2 = \int_{g^{-1}(E)} f \mathrm{d}\mathbb{P}_1$$

This means that

$$M_n(g)(\mu) = \mu \circ g^{-1} \quad \text{and} \quad RV_n(g)(f) = \mathbb{E}[f \mid g].$$

Summary

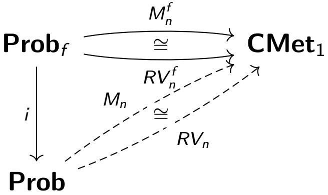

- (Bounded) Radon-Nikodym theorem:

$$M_n(\mathbf{\Omega}) = \{\mu \mid \mu \leq n\mathbb{P}\} \quad RV_n(\mathbf{\Omega}) = \mathbf{Mble}(\Omega, [0, n])/ =_{\mathbb{P}} .$$

- Conditional expectation:

$$RV_n(g)(X) = \mathbb{E}[X \mid f].$$

3.2 Martingales

Martingale convergence

Consider a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n\in\mathbb{N}}, \mathbb{P})$.

Martingale convergence

Consider a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n\in\mathbb{N}}, \mathbb{P})$.
The space $\mathbf{\Omega} := (\Omega, \mathcal{F}, \mathbb{P})$ is the limit of

$$\mathbf{\Omega}_1 \xleftarrow[s_{21}]{} \mathbf{\Omega}_2 \xleftarrow[s_{32}]{} \mathbf{\Omega}_3 \longleftarrow \dots \longleftarrow \mathbf{\Omega}_m \longleftarrow \dots$$

in **Prob**, where $\mathbf{\Omega}_m := (\Omega, \mathcal{F}_m, \mathbb{P}\,|_{\mathcal{F}_m})$.

Martingale convergence

Consider a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n\in\mathbb{N}}, \mathbb{P})$.
The space $\mathbf{\Omega} := (\Omega, \mathcal{F}, \mathbb{P})$ is the limit of

$$\mathbf{\Omega}_1 \xleftarrow[s_{21}]{} \mathbf{\Omega}_2 \xleftarrow[s_{32}]{} \mathbf{\Omega}_3 \longleftarrow \dots \longleftarrow \mathbf{\Omega}_m \longleftarrow \dots$$

in **Prob**, where $\mathbf{\Omega}_m := (\Omega, \mathcal{F}_m, \mathbb{P}\,|_{\mathcal{F}_m})$.
Suppose that $RV_n : \mathbf{Prob} \to \mathbf{CMet}_1$ preserves this limit, then

$$\begin{aligned} RV_n(\mathbf{\Omega}) &\cong \lim_m RV_n(\mathbf{\Omega}_m) \\ &\cong \{(X_m)_m \mid RV_n(s_{m_1 m_2})(X_{m_1}) = X_{m_2} \text{ for } m_2 \le m_1\} \\ &\cong \{(X_m)_m \mid \mathbb{E}[X_{m_1} \mid \mathcal{F}_{n_2}] = X_{m_2} \text{ for } m_2 \le m_1\} \\ &\cong \{(X_m)_m \mid \text{ martingale }\} \end{aligned}$$

Martingale convergence

Consider a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n\in\mathbb{N}}, \mathbb{P})$.
The space $\mathbf{\Omega} := (\Omega, \mathcal{F}, \mathbb{P})$ is the limit of

$$\mathbf{\Omega}_1 \xleftarrow[s_{21}]{} \mathbf{\Omega}_2 \xleftarrow[s_{32}]{} \mathbf{\Omega}_3 \longleftarrow \dots \longleftarrow \mathbf{\Omega}_m \longleftarrow \dots$$

in **Prob**, where $\mathbf{\Omega}_m := (\Omega, \mathcal{F}_m, \mathbb{P}\mid_{\mathcal{F}_m})$.
Suppose that $RV_n : \mathbf{Prob} \to \mathbf{CMet}_1$ preserves this limit, then

$$\begin{aligned} RV_n(\mathbf{\Omega}) &\cong \lim_m RV_n(\mathbf{\Omega}_m) \\ &\cong \{(X_m)_m \mid RV_n(s_{m_1 m_2})(X_{m_1}) = X_{m_2} \text{ for } m_2 \leq m_1\} \\ &\cong \{(X_m)_m \mid \mathbb{E}[X_{m_1} \mid \mathcal{F}_{n_2}] = X_{m_2} \text{ for } m_2 \leq m_1\} \\ &\cong \{(X_m)_m \mid \text{ martingale } \} \end{aligned}$$

It follows that for every martingale $(X_m)_m$ such that $X_m \leq n$ for all m, there exists a random variable $X : (\Omega, \mathcal{F}) \to [0, n]$ such that for all m,

$$\mathbf{E}[X \mid \mathcal{F}_m] = X_m.$$

Enriched categories

Ordinary categories: A *set* of morphisms between two objects.

Enriched categories

Ordinary categories: A *set* of morphisms between two objects.
$\mathcal{V}$-enriched categories: An *object in* $\mathcal{V}$ of morphisms between two objects.

Enrichment over $\mathbf{CMet}_1$

Everything from the first part still works when everything is *enriched* over $\mathbf{CMet}_1$.

Enrichment over $\mathbf{CMet}_1$

Everything from the first part still works when everything is *enriched* over $\mathbf{CMet}_1$.

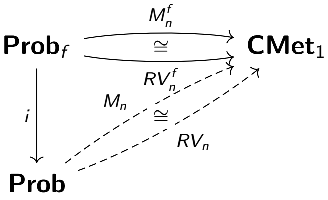

Enrichment over $\mathbf{CMet}_1$

Everything from the first part still works when everything is *enriched* over $\mathbf{CMet}_1$.

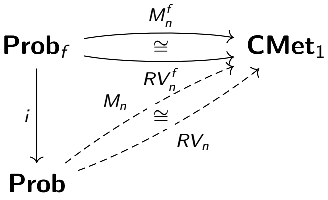

How is **Prob** enriched over $\mathbf{CMet}_1$?

Enrichment over $\mathbf{CMet}_1$

Everything from the first part still works when everything is *enriched* over $\mathbf{CMet}_1$.

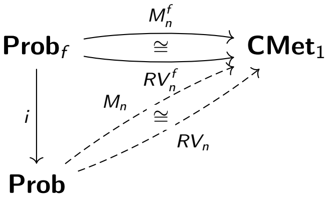

How is **Prob** enriched over $\mathbf{CMet}_1$?
Answer: $\mathbf{Prob}(\mathbf{\Omega}_1, \mathbf{\Omega}_2)$ is the *completion* of

$$\{f : \mathbf{\Omega}_1 \to \mathbf{\Omega}_2 \mid \text{ measure preserving}\}$$

with the pseudometric

$$d(f_1, f_2) := \sup\left\{\mathbb{P}_1(f_1^{-1}(A)\Delta f_2^{-1}(A)) \mid A \in \mathcal{F}_2\right\}.$$

RV_n preserves cofiltered limits

For any finite probability space $\mathbf{A}$, we always have a map

$$\text{colim}_i \mathbf{Prob}(\mathbf{\Omega}_i, \mathbf{A}) \to \mathbf{Prob}(\mathbf{\Omega}, \mathbf{A}).$$

RV_n preserves cofiltered limits

For any finite probability space $\mathbf{A}$, we always have a map

$$\text{colim}_i \mathbf{Prob}(\boldsymbol{\Omega}_i, \mathbf{A}) \to \mathbf{Prob}(\boldsymbol{\Omega}, \mathbf{A}).$$

Since $\{f : \boldsymbol{\Omega} \to \mathbf{A} \mid f \text{ is } \mathcal{F}_i\text{-measurable for some } i\}$ is dense in $\mathbf{Prob}(\boldsymbol{\Omega}, \mathbf{A})$, this is an isomorphism.

RV_n preserves cofiltered limits

For any finite probability space $\mathbf{A}$, we always have a map

$$\mathrm{colim}_i \mathbf{Prob}(\mathbf{\Omega}_i, \mathbf{A}) \to \mathbf{Prob}(\mathbf{\Omega}, \mathbf{A}).$$

Since $\{f : \mathbf{\Omega} \to \mathbf{A} \mid f \text{ is } \mathcal{F}_i\text{-measurable for some } i\}$ is dense in $\mathbf{Prob}(\mathbf{\Omega}, \mathbf{A})$, this is an isomorphism. We can now conclude:

$$\begin{aligned}
RV_n(\mathbf{\Omega}) &\cong \int_{\mathbf{A}} [\mathbf{Prob}(\mathbf{\Omega}, \mathbf{A}), RV_n^f(\mathbf{A})] \\
&\cong \int_{\mathbf{A}} [\mathrm{colim}_i \mathbf{Prob}(\mathbf{\Omega}_i, \mathbf{A}), RV_n^f(\mathbf{A})] \\
&\cong \int_{\mathbf{A}} \lim_i [\mathbf{Prob}(\mathbf{\Omega}_i, \mathbf{A}), RV_n^f(\mathbf{A})] \\
&\cong \lim_i \int_{\mathbf{A}} [\mathbf{Prob}(\mathbf{\Omega}_i, \mathbf{A}), RV_n^f(\mathbf{A})] \cong \lim_i RV_n(\mathbf{\Omega}_i)
\end{aligned}$$

RV_n preserves cofiltered limits

For any finite probability space $\mathbf{A}$, we always have a map

$$\mathrm{colim}_i \mathbf{Prob}(\mathbf{\Omega}_i, \mathbf{A}) \to \mathbf{Prob}(\mathbf{\Omega}, \mathbf{A}).$$

Since $\{f : \mathbf{\Omega} \to \mathbf{A} \mid f \text{ is } \mathcal{F}_i\text{-measurable for some } i\}$ is dense in $\mathbf{Prob}(\mathbf{\Omega}, \mathbf{A})$, this is an isomorphism. We can now conclude:

$$\begin{aligned} RV_n(\mathbf{\Omega}) &\cong \int_{\mathbf{A}} [\mathbf{Prob}(\mathbf{\Omega}, \mathbf{A}), RV_n^f(\mathbf{A})] \\ &\cong \int_{\mathbf{A}} [\mathrm{colim}_i \mathbf{Prob}(\mathbf{\Omega}_i, \mathbf{A}), RV_n^f(\mathbf{A})] \\ &\cong \int_{\mathbf{A}} \lim_i [\mathbf{Prob}(\mathbf{\Omega}_i, \mathbf{A}), RV_n^f(\mathbf{A})] \\ &\cong \lim_i \int_{\mathbf{A}} [\mathbf{Prob}(\mathbf{\Omega}_i, \mathbf{A}), RV_n^f(\mathbf{A})] \cong \lim_i RV_n(\mathbf{\Omega}_i) \end{aligned}$$

Remark: We did not use anything about RV_n^f.

Summary

Enriched version of

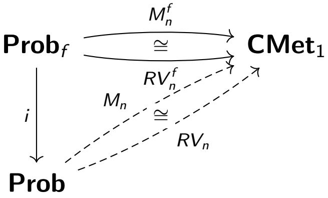

- (Bounded) Radon-Nikodym theorem:

$$M_n(\boldsymbol{\Omega}) = \{\mu \mid \mu \leq n\mathbb{P}\} \quad RV_n(\boldsymbol{\Omega}) = \mathbf{Mble}(\Omega, [0, n])/ =_{\mathbb{P}} .$$

- Conditional expectation:

$$RV_n(g)(X) = \mathbb{E}[X \mid f].$$

- Martingale convergence: RV_n preserves cofiltered limits.
- Weaker Kolmogorov extension theorem : M_n preserves cofiltered limits.

What about left Kan extensions?

Let $H : \mathbf{Prob}_f \to \mathbf{CMet}_1$ be a functor. Suppose that $\mathbf{\Omega}$ is a probability space that is **not** *essentially finite*.
Then $\mathbf{Prob}(\mathbf{A}, \mathbf{\Omega}) = \emptyset$ for all finite probability spaces $\mathbf{A}$ and

$$\mathrm{Lan}_i H(\mathbf{\Omega}) = \int^{\mathbf{A}} \mathbf{Prob}(\mathbf{A}, \mathbf{\Omega}) \times H\mathbf{A} = \emptyset.$$