



HAL
open science

A Formalization of Doob's Martingale Convergence Theorems in mathlib

Kexing Ying, Rémy Degenne

► **To cite this version:**

Kexing Ying, Rémy Degenne. A Formalization of Doob's Martingale Convergence Theorems in mathlib. 12th ACM SIGPLAN International Conference on Certified Programs and Proofs, Jan 2023, Boston (Massachusetts), United States. 10.1145/3573105.3575675 . hal-04337785

HAL Id: hal-04337785

<https://inria.hal.science/hal-04337785>

Submitted on 12 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

A Formalization of Doob’s Martingale Convergence Theorems in mathlib

Kexing Ying
University of Cambridge
Cambridge, UK
ky316@cam.ac.uk

Rémy Degenne
Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9198-CRISTAL
Lille, France
remy.degenne@inria.fr

Abstract

We present the formalization of Doob’s martingale convergence theorems in the mathlib library for the Lean theorem prover. These theorems give conditions under which (sub)martingales converge, almost everywhere or in L^1 . In order to formalize those results, we build a definition of the conditional expectation in Banach spaces and develop the theory of stochastic processes, stopping times and martingales. As an application of the convergence theorems, we also present the formalization of Lévy’s generalized Borel-Cantelli lemma. This work on martingale theory is one of the first developments of probability theory in mathlib, and it builds upon diverse parts of that library such as topology, analysis and most importantly measure theory.

CCS Concepts: • Mathematics of computing → Probability and statistics; • Theory of computation → Logic and verification.

Keywords: probability, martingale, Lean, formal mathematics, proof assistant, mathlib

ACM Reference Format:

Kexing Ying and Rémy Degenne. 2023. A Formalization of Doob’s Martingale Convergence Theorems in mathlib. In *Proceedings of the 12th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP ’23)*, January 16–17, 2023, Boston, MA, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3573105.3575675>

1 Introduction

Martingales are a fundamental object of study in modern probability theory. First introduced by Lévy in 1934, martingale theory in the mathematical sense was developed by Doob in the 1940’s. Today, martingales have formed the foundations of stochastic calculus and have found applications in economics, finance and statistics among many other fields.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

CPP ’23, January 16–17, 2023, Boston, MA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0026-2/23/01...\$15.00

<https://doi.org/10.1145/3573105.3575675>

Informally, a martingale models a series of fair games, while the related sub/super-martingales model the outcomes of a series of games biased towards/against the player. The martingale convergence theorems are a powerful result in martingale theory. Developed by Doob, they show that a sub/super-martingale converges provided some boundedness condition. It has a wide range of consequences, most notably the strong law of large numbers, the Radon-Nikodym theorem and Lévy’s generalized Borel-Cantelli lemma.

We describe our formalization of martingale theory in the Lean theorem prover and provide the formalization of the almost everywhere martingale convergence theorem and the L^1 martingale convergence theorem. To achieve this, we use and develop Lean’s mathematical library mathlib [12]. This entire project consists of approximately 9200 lines of code, all of which has been merged into mathlib. The source code for this project can be found in the mathlib GitHub repository¹. mathlib is a constantly evolving library: this paper discusses the code as it was at the time of the conference.

As an overview, this project involved the development and formalization of

- the definitions of conditional expectation, stopping time, hitting time and martingale;
- theory of uniform integrability and the Vitali convergence theorem;
- the optional stopping theorem/fair game theorem;
- Doob’s maximal inequality;
- the a.e. and L^1 martingale convergence theorems;
- Lévy’s upward theorem;
- Doob’s decomposition theorem;
- the one-sided martingale bound and
- Lévy’s generalized Borel-Cantelli lemma.

A proof of the almost everywhere martingale convergence theorem has previously been formalized by [18] earlier this year for the proof of the Dvoretzky’s stochastic approximation theorem in Coq. However, in contrast to their formalization, this project provides a more general framework to work with probability theory overall. In particular, our project contains a definition of the conditional expectation which applies to general Banach space-valued random variables, and thus, allowing us to define Banach space-valued martingales instead of only real-valued ones. This generality is useful and is required for the theory of UMD-spaces (c.f.

¹<https://github.com/leanprover-community/mathlib>

[8]). Furthermore, this project provides the first formalization of the L^1 martingale convergence theorem.

Other than the aforementioned project, [6] defines real-valued martingales in Isabelle/HOL for the purpose of formalizing the Cox-Rubinstein model in financial mathematics. Similarly, the HOL4 library [16] also formalizes the definition of the martingale property however, does so without showing the existence of conditional expectations. These two libraries do not contain any significant results on martingales beyond their definition.

More broadly, parts of probability theory have been formalized in many systems. Most notably, the Isabelle/HOL library contains a general formulation of measure-theoretic probability theory including a proof of the central limit theorem [2], the laws of large numbers [5] and several definitions central to stochastic processes, e.g. filtration and stopping time [9]. Similarly, the HOL4 library contains a general framework for probability theory including the formalization of the laws of large numbers along side results such as the Borel-Cantelli lemmas and Kolmogorov's 0-1 law [14]. Parts of probability theory has also been formalized in Mizar [15] and Coq [1, 17] focusing on discrete/finite sample spaces. Probability theory in Lean as with other theories is consolidated in `mathlib`. Built upon the measure theory part of the library, `mathlib` contains probability results ranging from the strong law of large numbers to parts of the portmanteau theorem alongside theorems introduced through this project (see the `mathlib` overview page² for more details).

1.1 Martingales and the convergence theorems

Measurable spaces and measures. This work concerns objects defined in a measure space $(\Omega, \mathcal{A}, \mu)$, that is a type (or set) Ω , a σ -algebra \mathcal{A} and a measure on that σ -algebra μ . A σ -algebra is a collection of subsets of Ω closed under complement, countable unions, and countable intersections. The sets in \mathcal{A} are called measurable sets. A measure is a map from \mathcal{A} to $\mathbb{R}_{\geq 0} \cup \{\infty\}$ with value zero on the empty set and which verifies a countable additivity assumption. For a measure μ , there is a standard way to give a value $\mu(s)$ to any subset s of Ω , not only the sets in \mathcal{A} (through outer measures, which we will not define).

We will say that a property $p : \Omega \rightarrow Prop$ is true almost everywhere (a.e.) if $\mu(\{\omega \mid \neg p(\omega)\}) = 0$, and denote this in `mathlib` by $\forall^m \omega \partial \mu, p \omega$. Similarly, we define almost everywhere equality of functions (or inequalities): $f =^m[\mu] g$ means that f and g are almost everywhere equal.

A measure μ is said to be finite if Ω has finite measure: $\mu(\Omega) < \infty$. It is σ -finite if there is an increasing sequence of measurable sets $(S_n)_{n \in \mathbb{N}}$ with finite measure such that $\bigcup_{n \in \mathbb{N}} S_n = \Omega$. If $\mu(\Omega) = 1$, then μ is said to be a probability measure. This paper is about objects like stochastic

processes, which are most often used in the context of probability theory. However none of our results will require that μ is a probability measure. We will try to state results in the greatest possible generality. Following that principle, many lemmas will not use any assumption on the measure, while other will require it to be σ -finite, or even finite. We never use probability measures because such a strong assumption is nowhere needed. This approach is very different from the one of [18] for example, in which one of the design choice of the probability library is to only consider probability measures.

Random variables and spaces of functions. A random variable $f : \Omega \rightarrow E$ is a function from Ω to another space E . In most places, we will suppose that E is a Banach space (complete real normed vector space), because that assumption allows us to define a notion of integral called the Bochner integral, which will be the main theory of integration we use in this paper.

`mathlib` contains several notions of measurability, and we now explain the ones which are most relevant to our work: `strongly_measurable` and `ae_strongly_measurable`. A simple function is defined as a piecewise constant function taking finitely many values, such that the preimage of each value is a measurable set. We denote the set of such functions by $\mathcal{S}(E, \mathcal{A})$. A function is said to be strongly measurable if it is the limit of a sequence of simple functions. It is said to be a.e. strongly measurable if it is a.e. equal to a strongly measurable function.

In second-countable metrizable spaces with the Borel σ -algebra, “strongly measurable” and the more commonly used “measurable” are equivalent, hence a reader accustomed to that second notion who would be mostly interested in applications to spaces like \mathbb{R}^d can safely ignore the difference.

We denote by $L^0(E, \mu)$ the space of a.e. strongly measurable functions, quotiented by the a.e. equal relation. The objects of that space are rigorously speaking classes of functions, but we still use the word function. For $p \geq 1$, we denote by $L^p(E, \mu)$ the subset of $L^0(E, \mu)$ containing the functions whose norm raised to the power p has finite Lebesgue integral (denoted by \int^-), that is

$$L^p(E, \mu) = \left\{ f \in L^0(E, \mu) \mid \int_x^- \|f(x)\|^p d\mu(x) < \infty \right\}.$$

The functions in $L^1(E, \mu)$ are said to be integrable. On this L^1 space the Bochner integral is well defined and we can talk about the integral of f , $\int_x f(x) d\mu(x)$, for which we also introduce the notation $\mu[f]$ in Lean. When μ is a probability measure, this integral is also called expectation of f .

Stochastic processes, filtrations and martingales. A stochastic process $(f_i)_{i \in \iota}$ is an ι -indexed family of random variables, where the index set is ordered and usually represents time. Most often $\iota = \mathbb{N}$ or $\iota = \mathbb{R}_{\geq 0}$.

²<https://leanprover-community.github.io/mathlib-overview.html>

Let us consider a simple example of a stochastic process: the amount of money of a player who bets on coin flips. At time $t \in \mathbb{N}$, a coin is flipped and the player gains 1 if the result of the flip is heads, and loses 1 if it is tails. The amount of money of the player over time is a stochastic process $(f_n)_{n \in \mathbb{N}}$, where each f_n is a real random variable $\Omega \rightarrow \mathbb{R}$.

A filtration $(\mathcal{F}_i)_{i \in \iota}$ is a non-decreasing sequence of sub- σ -algebras, and often represents information available up to some time. In our coin flip example, let X_i be the random variable describing the coin flip at time i , and let $\sigma(X_i)$ be the σ -algebra generated by X_i . We can consider the filtration defined by $\mathcal{F}_i = \mathcal{F}_{i-1} \vee \sigma(X_i)$ (i.e. the smallest σ -algebra generated by \mathcal{F}_{i-1} and X_i) with \mathcal{F}_0 the trivial σ -algebra, which encodes the observations available after flip i . A stochastic process is said to be adapted to a filtration if for all $i \in \iota$, f_i is \mathcal{F}_i -strongly measurable (strongly measurable with respect to the σ -algebra \mathcal{F}_i). Intuitively, the capital of the player after i coin flips can be computed from the knowledge of all the results up to time i .

While the expectation $\mu[f_n]$ for a probability measure μ can be interpreted as the expected value of f_n , which equals the initial capital f_0 in our example for fair coins since we expect heads and tails to happen equally often; the conditional expectation $\mu[f_n \mid \mathcal{F}_i]$ is the value of f_n we expect once we have seen the i first coin tosses. If the coins are fair, the process describing the amount of money of the player verifies $\mu[f_n \mid \mathcal{F}_i] = f_i$ almost everywhere for all $i \leq n$: such a process is called a martingale.

The martingale convergence theorems. The two martingale convergence theorems we formalize give conditions under which a (sub)martingale $(f_n)_{n \in \mathbb{N}}$ has a limit function f as n tends to infinity. The first one gives almost everywhere convergence to that limit.

Theorem (Almost everywhere martingale convergence). *Given a submartingale $(f_n)_{n \in \mathbb{N}}$ adapted to the filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$, if there exists some $R \in \mathbb{R}$ such that $\|f_n\|_1 := \int |f_n| \leq R$ for all $n \in \mathbb{N}$, then there exists some $\bigvee_{n \in \mathbb{N}} \mathcal{F}_n$ -measurable function f which is integrable such that $f_n \xrightarrow{\text{a.e.}} f$.*

We will for shorthand denote $\bigvee_{n \in \mathbb{N}} \mathcal{F}_n$ by \mathcal{F}_∞ .

The second result we prove characterizes the L^1 convergence of martingales.

Theorem (L^1 martingale convergence). *Given $(f_n)_{n \in \mathbb{N}}$ a martingale adapted to the filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$, the following are equivalent:*

- (i) (f_n) is uniformly integrable (in the probability sense);
- (ii) (f_n) converges in L^1 to some \mathcal{F}_∞ -measurable, integrable function f ;
- (iii) there exists some \mathcal{F}_∞ -measurable and integrable f such that for all $n \in \mathbb{N}$, $\mu[f \mid \mathcal{F}_n] = f_n$ a.e. .

Those two theorems are at the core of the study of martingales, which are themselves very versatile tools in probability

theory. They can for example be used as in [19] to prove Kolmogorov's 0-1 law, the strong law of large numbers or the generalized Borel-Cantelli lemma.

1.2 Lean and mathlib

All the code produced for this paper was contributed to mathlib [12]: the mathematical library for the Lean theorem prover [13]. This library contains formalized definitions and theorems from many domains of mathematics, in a way that ensures that all those definitions are usable together without compatibility issues. We will see in the section on the definition of the conditional expectation that this feature was essential to our work. Probability theory is built on a large corpus of results in analysis, topology, descriptive set theory and measure theory, and its formalization beyond discrete spaces requires a library which is well developed in all those areas.

New code submitted to mathlib goes through a review process supervised by the mathlib maintainers, helped by automatic checks (for example checking for unused arguments in theorems) as described in [4]. Ensuring that the new code is compatible with and general enough to be used by other parts of the library is a central part of that review process.

Our paper can be understood without prior knowledge of Lean syntax, especially since it contains only Lean statements and no proofs. These statements however contain three distinct types of parameters, which we now explain on an example. In the remainder of the paper we assume that a measure space $(\Omega, \mathcal{A}, \mu)$ is given, as well as an index set ι and a Banach space E for the values of the random processes. To obtain the same effect as that sentence in Lean, we would add the following line at the top of the file:

```
variables { $\Omega \iota E : \text{Type}^*$ } { $\mathcal{A} : \text{measurable\_space } \Omega$ }
{ $\mu : \text{measure } \Omega$ } [normed\_add\_comm\_group  $E$ ]
[normed\_space  $\mathbb{R} E$ ] [complete\_space  $E$ ]
```

Then those parameters don't need to be given to every lemma and definition, and we will also not add them to the Lean code we write below. The parameters enclosed in curly brackets like $\{\mu : \text{measure } \Omega\}$ are implicit. Lean is supposed to infer them from other hypotheses: the measure will appear in other hypotheses of the theorems and Lean should guess it from there. An explicit parameter would be given in parentheses, for example $(\mu : \text{measure } \Omega)$, and the user would have to provide it when they want to apply the theorem. In the above variables definition, we also see [complete_space E], which is an instance parameter: Lean uses typeclass resolution to deduce those parameters from the definition of E and other instances in the hypotheses. If for example $E = \mathbb{R}$, we won't need to supply a proof that \mathbb{R} is complete, since typeclass resolution can build it automatically. For a detailed discussion of instance parameters in mathlib, see [3].

2 Stopping times and martingales: definitions

In order to prove the martingale convergence theorems, we first need to formalize the objects they use, namely conditional expectations, uniform integrability, stochastic processes, stopping times and martingales. This section is dedicated to our development of these basic objects of probability theory.

2.1 Conditional expectation

The definition of martingales requires a notion of conditional expectation, which is a fundamental tool in probability.

Definition 2.1. Let $(\Omega, \mathcal{A}, \mu)$ be a measure space and let \mathcal{B} be a sub- σ -algebra of \mathcal{A} . Let E be a Banach space (complete normed vector space) and $f : \Omega \rightarrow E$ be an integrable function (with respect to μ). Then a function $g : \Omega \rightarrow E$ is said to be a *conditional expectation* of f with respect to \mathcal{B} if it is integrable, \mathcal{B} -strongly measurable and for all \mathcal{B} -measurable sets s ,

$$\int_{x \in s} g(x) d\mu(x) = \int_{x \in s} f(x) d\mu(x).$$

If the measure μ restricted to the σ -algebra \mathcal{B} is σ -finite, then there exists a conditional expectation and it is unique up to almost everywhere equality.

Our formalization broadly follows the procedure described in [8] and leads to a definition with the following type.

```
def condexp (B : measurable_space Ω)
  {A : measurable_space Ω}
  (μ : measure Ω) (f : Ω → E) : Ω → E :=
if hm : B ≤ A
  then if h : sigma_finite (μ.trim hm) ∧ integrable f μ
    then if strongly_measurable[B] f then f
    else (@ae_strongly_measurable' _condexp_L1 _ _ _ _ B
      A μ hm h.1 _).mk (@condexp_L1 _ _ _ _ _ hm μ h.1 f)
  else 0
else 0
```

We will always try to make the use of the library as close as possible to paper proofs, and to that purpose we define a notation $\mu[f | \mathcal{B}]$ for $\text{condexp } \mathcal{B} \mu f$.

```
notation μ \[ f \] \[ B \] := measure_theory.condexp B μ f
```

Let us parse the definition. Ω is a measurable space with σ -algebra \mathcal{A} on which we have a measure μ and E is a Banach space. Then \mathcal{B} is the σ -algebra with respect to which we are computing the conditional expectation and f is the function we are computing it for. What we obtain is a function $\Omega \rightarrow E$.

In contrast to definition 2.1, the Lean definition does not require \mathcal{B} to be a sub- σ -algebra of \mathcal{A} , nor does it require f to be integrable or the measure to be σ -finite on \mathcal{B} . Instead, we define $\mu[f | \mathcal{B}]$ as 0 if those conditions are not verified. Our function `condexp` is then a true conditional expectation only under additional assumptions, but defining it this way allows us to write $\mu[f | \mathcal{B}]$ without having to add to that

notation proofs of the three required properties. Defining such global functions with default values is a common strategy in theorem provers to make the definitions easier to use. Of course the hypotheses we omitted are then required for most lemmas describing the properties of `condexp`.

The `condexp_L1` auxiliary function is the result of the construction detailed in the next section. We also check whether f is \mathcal{B} -strongly measurable: in that case, any \mathcal{B} -strongly measurable function a.e. equal to f is a conditional expectation, and `condexp_L1` will be one of those, but we choose to enforce everywhere equality instead of a.e. equality to ease the use of the function. Indeed, while the properties of the conditional expectation we will be interested in are invariant to a.e. equality and thus the choice of a particular function changes nothing mathematically, having exact equality to f in the strongly measurable case makes it easier to replace $\mu[f | \mathcal{B}]$ by f in the code.

Among the many properties of `condexp`, we most notably prove the equality of integrals which characterizes conditional expectations:

```
theorem set_integral_condexp (hB : B ≤ A)
  [sigma_finite (μ.trim hB)] (hf : integrable f μ)
  (hs : measurable_set[B] s) :
  ∫ x in s, μ[f | B] × ∂μ = ∫ x in s, f × ∂μ
```

This proves that $\mu[f | \mathcal{B}]$ is indeed the conditional expectation of f . The generality in which we define the conditional expectation is close to the greatest possible one: the only condition which is more relaxed in [8] compared to our implementation is the integrability. We require the function f to be integrable to have a non-zero $\mu[f | \mathcal{B}]$, while [8] define a meaningful conditional expectation for the larger class of σ -integrable functions, which means that there is an increasing sequence of sets with union equal to the whole space such that the function is integrable on each individual set. σ -integrability is not defined in `mathlib` and we chose not to add it since it is rarely used in the literature.

Defining the conditional expectation for Banach space and σ -finite measure can be useful in other fields than probability, for example in harmonic analysis where the measure is often the Lebesgue measure, which is not finite.

Construction of the conditional expectation. The main goal of our construction is to obtain the conditional expectation as a continuous linear map from $L^1(E, \mu)$ to itself. The final definition is then obtained from that map by wrapping it in a few if...then...else checks as explained above.

The first step in that construction is to define a conditional expectation for indicators of measurable sets with finite measure. An indicator of a measurable set with finite measure is a function in $L^2(E, \mu)$, so we can focus on those. We want the conditional expectation to be a \mathcal{B} -strongly measurable function, and those functions form a linear subspace of $L^2(E, \mu)$. We can then obtain a conditional expectation for any function in $L^2(E, \mu)$ by projecting orthogonally on the subspace

of \mathcal{B} -strongly measurable functions. The existence of such an orthogonal projection, the linearity properties and the fact that the projection verifies the conditional expectation characterization are all consequences of properties of projections in inner product spaces which were already available in the analysis part of `mathlib` and directly applicable. The nature of `mathlib` as a monolithic library containing all kinds of mathematical domains made the following quick definition possible:

```
def condexp_L2 (hm :  $\mathcal{B} \leq \mathcal{A}$ ) :
  ( $\Omega \rightarrow_2 [\mu] E$ )  $\rightarrow$  L[ $\mathbb{R}$ ] (Lp_meas E  $\mathbb{R}$   $\mathcal{B}$  2  $\mu$ ) :=
  @orthogonal_projection  $\mathbb{R}$  ( $\Omega \rightarrow_2 [\mu] E$ ) _ _
  (Lp_meas E  $\mathbb{R}$   $\mathcal{B}$  2  $\mu$ )
  (by { haveI : fact ( $\mathcal{B} \leq \mathcal{A}$ ) := (hm), apply_instance, })
```

That `condexp_L2` is a linear map from $L^2(E, \mu)$ (denoted by $\Omega \rightarrow_2 [\mu] E$) to the type `Lp_meas E \mathbb{R} \mathcal{B} 2 μ` , which is the subspace of $L^2(E, \mu)$ which contains a.e. \mathcal{B} -strongly measurable functions. It is defined by using `orthogonal_projection`. The last argument is a proof that the subspace is complete: that proof first makes the fact that \mathcal{B} is a sub- σ -algebra of \mathcal{A} available to the type class resolution mechanism, then calls the `apply_instance` tactic, which uses that mechanism to generate a proof.

We now explain how we can extend that definition which works in particular for indicators of measurable sets with finite measure, to obtain a linear map from $L^1(E, \mu)$ to itself, which will give us the full conditional expectation.

Let E be a normed real vector space and G be a Banach space. Let $(\Omega, \mathcal{A}, \mu)$ be a measure space. Suppose that we have a map T from the subsets of Ω to the continuous linear maps from E to G which verifies the following properties:

1. for any two disjoint measurable sets with finite measure s, t , $T(s \cup t) = T(s) + T(t)$,
2. there exists $C \geq 0$ such that for any measurable set s with finite measure, $\|T(s)\| \leq C\mu(s)$.

Then we can define a continuous linear map T^{L^1} from $L^1(E, \mu)$ to G , such that for all measurable sets with finite measure s and all $x \in E$, T^{L^1} applied to the indicator of the set s with value x is equal to $T(s, x)$.

The extension process is done in several successive phases. First, the function can naturally be defined on indicators of measurable sets with finite measure by $T(s, x)$. Then thanks to the linearity property of T it can be extended to integrable simple functions in $\mathcal{S}(E, \mathcal{A})$, which are linear combinations of such indicators. It can then be extended to functions which are almost everywhere equal to an integrable simple function f_s by taking the value of T on f_s . The second property of T (dominated norm) is used to prove that the resulting map from a.e. classes of integrable simple functions to G is a continuous linear map. Since that map is continuous, we can give a unique value to functions which are a.e. limits of sequences of integrable simple functions. Those are the a.e.

strongly measurable and integrable functions, e.g. functions in $L^1(E, \mu)$, to which we can then finally extend our map.

This extension process is added to `mathlib` and is used for two applications:

- For $T(s, x)$ the conditional expectation of the indicator of set s with value x , with value in $G = L^1(E, \mu)$, we get for T^{L^1} the conditional expectation of a function in $L^1(E, \mu)$, as a linear map from $L^1(E, \mu)$ to $L^1(E, \mu)$.
- For $T(s, x) = \mu(s) \cdot x$ with value in $G = E$, T^{L^1} maps a function in $L^1(E, \mu)$ to its Bochner integral.

A special case of the extension process was previously implemented for the definition of the Bochner integral, but we refactored the library to isolate this general construction and apply it to both the definition of the integral and the conditional expectation.

With this refactor we were able to reuse many proofs first written for integrals, up to appropriate generalization, and obtain relatively quickly the corresponding properties for `condexp`. Like many parts of `mathlib`, the integration files were subject to frequent changes in the last year and the integral was slightly altered or generalized several times to allow an easier use in other parts of the library. See [7] for an overview of the other recent integral refactors.

2.2 Uniform integrability

We now define uniform integrability, which will feature in the L^1 convergence theorem. In the literature, there are two non-equivalent formulations of uniform integrability used by analysts and probabilists respectively.

Definition 2.2 (Uniform integrability, analyst's definition). A sequence of functions $(f_n) \subseteq L^1(E, \mu)$ is said to be uniformly integrable if for all $\epsilon > 0$, there exists some $\delta > 0$ such that for all $A \in \mathcal{A}$ with $\mu(A) < \delta$, we have for all n ,

$$\|f_n 1_A\|_1 < \epsilon.$$

Where we recall $\|f\|_1 = \int_x \|f(x)\| d\mu(x)$.

More generally, for $p \geq 1$ we define the norm $\|f\|_p = (\int_x \|f(x)\|^p d\mu(x))^{1/p}$ on $L^p(E, \mu)$. We say that a sequence of $L^p(E, \mu)$ functions is said to converge in L^p if it converges with respect to the topology induced by this norm.

Definition 2.3 (Uniform integrability, probabilist's definition). A sequence of functions $(f_n) \subseteq L^1(E, \mu)$ is said to be uniformly integrable if

$$\lim_{C \rightarrow \infty} \sup_n \|f_n 1_{\{|f_n| \geq C\}}\|_1 = 0.$$

Uniform integrability establishes a relation between convergence in measure and convergence in L^p known as the Vitali convergence theorem. This is the main result required to prove the L^1 martingale convergence theorem.

Theorem 2.4 (Vitali convergence theorem). *A sequence of functions $(f_n) \subseteq L^p(E, \mu)$ converges in L^p if and only if (f_n)*

converges in measure and $(|f_n|^p)$ is uniformly integrable in the analyst's sense.

The Vitali convergence theorem also holds for uniform integrability in the probabilist's sense as the probabilist's definition is strictly stronger than the analyst's definition. This is made clear by the following proposition.

Proposition 2.5. *If μ is a finite measure, then $(f_n) \subseteq L^1(E, \mu)$ is uniformly integrable in the probabilist's sense if and only if it is uniformly bounded in L^1 and is uniformly integrable in the analyst's sense.*

Hence, it is equivalent to define probabilist's uniform integrability with the characterization of proposition 2.5. This alternative characterization has the advantage of directly relating the two definitions and so, allowing us to only make one set of API without needing too many translation lemmas.

In Lean, these two definitions are formulated as,

```
def unif_integrable { $\mathcal{A}$  : measurable_space  $\alpha$ }
  (f :  $\iota \rightarrow \alpha \rightarrow \beta$ ) (p :  $\mathbb{R} \geq 0 \infty$ ) ( $\mu$  : measure  $\alpha$ ) :=
   $\forall (\epsilon : \mathbb{R}) (h\epsilon : 0 < \epsilon), \exists (\delta : \mathbb{R}) (h\delta : 0 < \delta),$ 
   $\forall i s, \text{measurable\_set } s \rightarrow \mu s \leq \text{ennreal.of\_real } \delta \rightarrow$ 
   $\text{snorm } (s.\text{indicator } (f i)) p \mu \leq \text{ennreal.of\_real } \epsilon$ 
```

```
def uniform_integrable { $\mathcal{A}$  : measurable_space  $\alpha$ }
  (f :  $\iota \rightarrow \alpha \rightarrow \beta$ ) (p :  $\mathbb{R} \geq 0 \infty$ ) ( $\mu$  : measure  $\alpha$ ) :=
  ( $\forall i, \text{ae\_strongly\_measurable } (f i)$ )
   $\wedge$  unif_integrable f p  $\mu$ 
   $\wedge \exists C : \mathbb{R} \geq 0, \forall i, \text{snorm } (f i) p \mu \leq C$ 
```

for the analyst's and probabilist's definition respectively where $\text{snorm } f p \mu$ is the function evaluating to $\|f\|_p$ for some suitable f with respect to the measure μ .

We note that the Lean formulation of both definitions has an extra parameter $p : \mathbb{R} \geq 0 \infty$ for which our original definition is recovered if $p = 1$. This allows us to talk about uniform integrability in the case (f_n) is not necessarily L^1 . In the case that the measure μ is finite, by Hölder's inequality, we have

```
p  $\leq$  q  $\rightarrow$ 
  uniform_integrable f q  $\mu \rightarrow$  uniform_integrable f p  $\mu$ 
```

To recover the original definition of uniform integrability in the probabilist's sense, we proved the following:

```
lemma uniform_integrable_iff
  [is_finite_measure  $\mu$ ] (hp :  $1 \leq p$ ) (hp' :  $p \neq \infty$ ) :
  uniform_integrable f p  $\mu$ 
   $\leftrightarrow$  ( $\forall i, \text{ae\_strongly\_measurable } (f i)$ )
   $\wedge \forall \epsilon : \mathbb{R}, 0 < \epsilon \rightarrow \exists C : \mathbb{R} \geq 0,$ 
   $\forall i, \text{snorm } (\{x \mid C \leq \|f i x\|_+\}. \text{indicator } (f i)) p \mu$ 
   $\leq \text{ennreal.of\_real } \epsilon$ 
```

where $\|x\|_+$ denotes the non-negative real valued norm of any suitable x .

As we included the extra parameter p in our formalization of uniform integrability, in contrast to the formulation of the Vitali convergence theorem above, we do not need to say " $(|f_n|^p)$ is uniformly integrable" but simply " (f_n) is

p -uniformly integrable" resulting in the following formalization of the Vitali convergence theorem,

```
lemma tendsto_in_measure_iff_tendsto_Lp
  [is_finite_measure  $\mu$ ] (hp :  $1 \leq p$ ) (hp' :  $p \neq \infty$ )
  (hf :  $\forall n, \text{mem\_Lp } (f n) p \mu$ ) (hg :  $\text{mem\_Lp } g p \mu$ ) :
  tendsto_in_measure  $\mu$  f at_top g
   $\wedge$  unif_integrable f p  $\mu$ 
   $\leftrightarrow$  tendsto ( $\lambda n, \text{snorm } (f n - g) p \mu$ ) at_top ( $N 0$ )
```

with $\text{tendsto_in_measure } \mu f \text{ at_top } g$ being the formalization of " (f_n) converges in measure along μ to g as $n \rightarrow \infty$ ", $\text{tendsto } (\lambda n, \text{snorm } (f n - g) p \mu) \text{ at_top } (N 0)$ the formalization of " $\|f_n - g\|_p$ converges to 0 as $n \rightarrow \infty$ ", and $\text{mem_Lp } g p \mu$ the assertion that g is a.e.-measurable and $\|g\|_p < \infty$.

Our first implementation of uniform integrability in the probability sense required that the functions are strongly measurable instead of a.e. strongly measurable. The initial consensus among interested mathlib community members was that the probability part of the library could focus solely on functions which are strongly measurable, and did not need to use the a.e. strongly measurable notion. However, later work on the L^p -version of the strong law of large numbers (SLLN) indicated that this definition of uniform integrable is too strong.

The a.e.-version of the SLLN states that, given a sequence of independent and identically distributed (i.i.d.) random variables $(f_n)_{n \in \mathbb{N}}$ with $\|f_0\|_1 < \infty$, $S_n := \frac{1}{n} \sum_{i=1}^n f_n$ converges to $\mu[f_0]$ almost everywhere. Hence, as a sequence of identically distributed random variables is uniformly integrable, and as the averaging of a uniformly integrable sequence is also uniformly integrable, the L^p -version of the SLLN follows by the Vitali convergence theorem, i.e. for i.i.d. random variables $(f_n)_{n \in \mathbb{N}}$ with $\|f_0\|_p < \infty$, S_n converges to $\mu[f_0]$ in L^p .

In the outlined proof above, we use the following lemma.

```
lemma uniform_integrable_of_ident_distrib {f :  $\iota \rightarrow \alpha \rightarrow E$ }
  {j :  $\iota$ } {p :  $\mathbb{R} \geq 0 \infty$ } (hp :  $1 \leq p$ ) (hp' :  $p \neq \infty$ )
  (hLp :  $\text{mem\_Lp } (f j) p \mu$ )
  (hf :  $\forall i, \text{ident\_distrib } (f i) (f j) \mu$ ) :
  uniform_integrable f p  $\mu$ 
```

We note that, in the case uniform integrability only requires a.e. strongly measurability, the above assumptions are sufficient. Indeed, as we have asserted that f_j is L^p , f_j is a.e. strongly measurable by the definition of mem_Lp . Hence, by the properties of identically distributed functions, f_i is also a.e. strongly measurable for all $i \in \iota$. In contrast, should uniform integrability require strongly measurable, as a function identically distributed with a strongly measurable functions is not automatically strongly measurable, we will in addition require f_i to be strongly measurable for all i rather than only imposing measurability of f_j .

In essence, we are manipulating properties of L^p classes and any predicate P on them should apply to all elements of

said class; i.e. P should factor through the following triangle:

$$\begin{array}{ccc} \mathcal{L}^p := \{f \text{ a.e. measurable} \mid \|f\|_p < \infty\} & \xrightarrow{P} & \text{Prop} \\ \downarrow Q & \nearrow \bar{P} & \\ L^p := \mathcal{L}^p / \text{a.e. equality} & & \end{array}$$

with Q being the quotient map. With this in mind, we will allow a.e. strongly measurable functions in our discourse.

2.3 Stochastic processes

A stochastic process is a sequence of random variables evolving in time. While this project mostly deals with discrete time processes, we will in general not be too restrictive on the time indices while making definitions, e.g. we will allow the time index to be an arbitrary set equipped with some ordering in most cases.

Given a measurable space (Ω, \mathcal{A}) , a filtration on this measurable space is a sequence of increasing sub- σ -algebras of \mathcal{A} : $(\mathcal{F}_i)_{i \in \iota}$ with some time index set ι (most commonly \mathbb{N} or $\mathbb{R}_{\geq 0}$).

```
structure filtration {Ω : Type*} (ι : Type*) [preorder ι]
  (A : measurable_space Ω) :=
  (seq : ι → measurable_space Ω)
  (mono' : monotone seq)
  (le' : ∀ i : ι, seq i ≤ A)
```

In Lean, we define filtrations as a structure and prove that it forms a complete lattice with respect to the order inherited from the order on σ -algebras.

```
instance : complete_lattice (filtration ι A)
```

With this, given a filtration $(\mathcal{F}_i)_{i \in \iota}$ and a stochastic process $(f_i)_{i \in \iota}$, we say (f_i) is adapted to (\mathcal{F}_i) if f_i is \mathcal{F}_i -strongly measurable for all $i \in \iota$.

```
def adapted (F : filtration ι A) (f : ι → Ω → β) : Prop :=
  ∀ i : ι, strongly_measurable[F i] (f i)
```

In stochastic processes, it is useful to be able to talk about not just deterministic times but also random times as well. To achieve this we introduce the notion of stopping times. Stopping times are a useful definition to characterize the times at which a process exhibits some given special property. To prove the martingale convergence theorems, we develop a theory of stopping times in Lean and define an important class of stopping time known as hitting time³.

Given some ι -indexed filtration $(\mathcal{F}_i)_{i \in \iota}$, we say a function $\tau : \Omega \rightarrow \iota$ is a stopping time if for all $i \in \iota$, $\{\tau \leq i\}$ is \mathcal{F}_i -measurable.

```
def is_stopping_time [preorder ι] (F : filtration ι A)
  (τ : Ω → ι) :=
  ∀ i : ι, measurable_set[F i] {ω | τ ω ≤ i}
```

³While hitting times are not in general stopping times, they are in the case that the time index is discrete which is the setting we are working under.

Given some E -valued process $(f_n)_{n \in \mathbb{N}}$ and a set $A \subseteq E$, the hitting time is thought of as the first time (f_n) enters A . Namely, it is defined to be

$$H^f(A) : \Omega \rightarrow \mathbb{N} \cup \{\infty\} : \omega \mapsto \inf\{n \mid f_n(\omega) \in A\}$$

where we define $\inf \emptyset = \infty$. Given this definition, it is important in our proof of the martingale convergence theorems that whenever $\{n \mid f_n(\omega) \in A\} \neq \emptyset$ we observe

1. $f_{H^f(A)} \in A$ and,
2. $\forall n < H^f(A), f_n \notin A$.

In particular, taking the stopped process defined by $N \in \mathbb{N} \mapsto f_{H^f(A) \wedge N}$ where $N \wedge \infty = N$, we have $f_{H^f(A)} \in A$ and for all $n < H^f(A)$, $f_n \notin A$ if $H^f(A) < N$. This is however not possible in a typed system such as Lean since $f_{H^f(A)}$ does not type check as (f_n) is indexed by \mathbb{N} while $H^f(A)$ has type $\mathbb{N} \cup \{\infty\}$ (in contrast to on paper where we prove the definition is well-defined after the fact by showing $N \wedge \infty \in \mathbb{N}$ for all $N \in \mathbb{N}$). Thus, an alternative formulation of the hitting time is required.

A naïve approach is to directly take the range of $H^f(A)$ to be \mathbb{N} where $\inf \emptyset$ is defined to be some natural number n (in Lean n is taken to be 0). This approach however, does not allow us to reobtain the two properties required for the martingale convergence theorem since for $N > n$, $\inf \emptyset \wedge N = n < N$ and hence, is not suitable for our purpose.

Instead, we would like to incorporate the stopped process notion into our definition of the hitting time. Namely, by introducing an ending time N , we define the hitting time to be the first time (f_n) enters A before time N . Furthermore, in order for our definition to also work in the case that the time index is unbounded from below, it is also useful to introduce a starting time, resulting in the following definition

```
def hitting [preorder ι] [has_Inf ι]
  (f : ι → α → β) (s : set β) (n m : ι) : α → ι :=
  λ x, if ∃ j ∈ set.Icc n m, f j x ∈ s then
    Inf (set.Icc n m ∩ {i : ι | f i x ∈ s}) else m
```

where $\text{set.Icc } n \ m$ denotes the closed interval $[n, m]$. We introduce the notation $H_{N,M}^f(A)$ for the hitting time of f in A with starting time N and ending time M .

The lemma that discrete hitting times are stopping times is formalized as

```
lemma hitting_is_stopping_time
  [topological_space β] [pseudo_metrizable_space β]
  [measurable_space β] [borel_space β]
  {f : filtration ℕ A} {u : ℕ → Ω → β} {s : set β} {n n' : ℕ}
  (hu : adapted f u) (hs : measurable_set s) :
  is_stopping_time f (hitting u s n n')
```

2.4 Martingales

Definition 2.6 ([8], definition 3.1.1). A family of integrable functions $(f_i)_{i \in \iota}$ with values in a Banach space E is called a martingale with respect to a σ -finite filtration $(\mathcal{F}_i)_{i \in \iota}$ if it is

adapted to (\mathcal{F}_i) and for all $i, j \in \iota$ with $i \leq j$, $\mu[f_j | \mathcal{F}_i] = f_i$ almost everywhere.

If there is an order \leq defined on E , then we can also define *submartingales* and *supermartingales*, with the same definition except that the equality of the conditional expectation is replaced respectively by $\mu[f_j | \mathcal{F}_i] \geq f_i$ and $\mu[f_j | \mathcal{F}_i] \leq f_i$. Martingales are both submartingales and supermartingales. Our Lean definition of a martingale is

```
def martingale [preorder ι] (f : ι → Ω → E)
  (F : filtration ι ℳ) (μ : measure Ω) : Prop :=
  adapted F f ∧ ∀ i j, i ≤ j → μ[f j | F i] =m[μ] f i
```

submartingale and supermartingale are defined similarly, under the additional hypothesis `[has_le E]` which requires that E has an order.

This definition inherits the generality of the conditional expectation construction: E is any Banach space. The index set can also be any preorder, which covers in particular the two usual cases of discrete and continuous time indexes (\mathbb{N} and \mathbb{R}). The function sequence (f_i) is not required to contain integrable functions (in contrast to Definition 2.6), but if f_j is not integrable then $\mu[f_j | \mathcal{F}_i] = 0$ (by definition) and the sequence has to be zero for all $i \leq j$. Similarly to the definition of the conditional expectation, the martingale property is defined more generally than for integrable functions for usability reasons, but will of course have interesting properties only for integrable functions.

An important example is the martingale generated by a function f , defined by $(\mu[f | \mathcal{F}_i])_{i \in \iota}$. A consequence of the L^1 martingale convergence theorem is that any real uniformly integrable martingale indexed by \mathbb{N} is of that form for some function f , as we will see in the following sections.

3 Martingale convergence theorems

We will in this section describe our formalization of the almost everywhere and the L^1 martingale convergence theorems.

In the remainder of the paper, the time index for stochastic processes will be \mathbb{N} and these processes will be real-valued (although a few results are stated for more general indexes), since Doob's convergence theorems are results about such processes. The hypotheses of the theorems also use the term *measurable* instead of the *strongly measurable* property explained in the introduction, but these two notions coincide for real valued random variables.

3.1 The almost everywhere martingale convergence theorem

We recall the statement of the almost everywhere martingale convergence theorem as mentioned in the introduction.

Theorem 3.1 (Almost everywhere martingale convergence). *Given a submartingale $(f_n)_{n \in \mathbb{N}}$ adapted to the filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$, if there exists some $R \in \mathbb{R}$ such that $\|f_n\|_1 \leq R$ for all $n \in \mathbb{N}$,*

then there exists some $\bigvee_{n \in \mathbb{N}} \mathcal{F}_n$ -measurable function f which is integrable such that $f_n \xrightarrow{a.e.} f$.

We will for shorthand denote $\bigvee_{n \in \mathbb{N}} \mathcal{F}_n$ by \mathcal{F}_∞ .

The proof of the theorem relies on the following crucial observation on the convergence of real sequences. Namely, a real sequence (x_n) converges if and only if

- (a) $(|x_n|)$ is bounded above;
- (b) and for all $a < b \in \mathbb{Q}$, (x_n) does not visit below a and above b infinitely often.

Thus, for a submartingale (f_n) satisfying the assumptions of the theorem, it suffices to show that (f_n) satisfies the conditions (a) and (b) almost everywhere. To show this, we will introduce Doob's upcrossing estimate which provides a bound on the expected number of times (f_n) can cross from below a to above b .

3.2 Doob's upcrossing estimate

Given real values $a < b$, we would like to define a notion of "upcrossings" of (f_n) across the band $[a, b]$ which counts the number of times any realization of (f_n) crosses from below a to above b . To make this heuristic rigorous, we introduce the following stopping times mutually inductively defined using hitting times,

$$\sigma_0 := 0, \tau_n := H_{\sigma_n, N}^f(-\infty, a], \text{ and } \sigma_{n+1} := H_{\tau_{n+1}, N}^f[b, \infty),$$

for some given $N \in \mathbb{N}$. Then, we may define the upcrossings of (f_n) before time N to be

$$U_N^f(a, b) := \sup\{n \mid \sigma_n < N\}.$$

In Lean, the stopping times $(\sigma_n)_{n \in \mathbb{N}}$ and $(\tau_n)_{n \in \mathbb{N}}$ are defined as

```
def upper_crossing [preorder ι] [order_bot ι] [has_inf ι]
  (a b : ℝ) (f : ι → Ω → ℝ) (N : ι) : ℕ → Ω → ι
| 0 := ⊥
| (n + 1) := λ x, hitting f (set.Ici b)
  (lower_crossing_aux a f (upper_crossing n x) N x)
```

and

```
def lower_crossing
  [preorder ι] [order_bot ι] [has_inf ι]
  (a b : ℝ) (f : ι → Ω → ℝ) (N : ι) (n : ℕ) : Ω → ι :=
λ x, hitting f (set.Iic a) (upper_crossing a b f N n x) N x
```

respectively, where `lower_crossing_aux` is defined by

```
def lower_crossing_aux [preorder ι] [has_inf ι]
  (a : ℝ) (f : ι → Ω → ℝ) (c N : ι) : Ω → ι :=
hitting f (set.Iic a) c N
```

As an example, suppose we have the following realization of a process crossing the band $[a, b]$ in which we take N to be 13. Then, by construction, $\sigma_0 = 0$ and as τ_0 is the first time the process hits below a , $\tau_0 = 1$. Similarly, as σ_1 is the first time the process hits above b after τ_0 , $\sigma_1 = 5$. By the same rationale, $\tau_1 = 7$ and $\sigma_2 = 10$. Finally, as the realization does not cross another band before 13, $\sigma_n, \tau_n = 13$ for all $n > 2$.




Figure 1. An example of realization of a stochastic process with σ_n, τ_n labeled.

With this, the number of upcrossings $U_N^f(a, b)$ is

```
def upcrossing_before
  [preorder ι] [order_bot ι] [has_Inf ι]
  (a b : ℝ) (f : ι → α → ℝ) (N : ι) (x : α) : ℕ :=
  Sup {n | upper_crossing a b f N n x < N}
```

We are allowing arbitrary indices in contrast to just natural number indices as we would eventually like to use the same definition for the discretization of continuous martingales.

We may now formulate Doob's upcrossing estimate.

Theorem 3.2 (Doob's upcrossing estimate). *For all $a < b \in \mathbb{R}$ and $N \in \mathbb{N}$, we have*

$$(b - a)\mu[U_N^f(a, b)] \leq \mu[(f_N - a)^+], \quad (1)$$

where g^+ denotes the positive part of the function g . Furthermore, denoting $U^f(a, b) := \sup_{N \in \mathbb{N}} U_N^f(a, b)$, we have

$$(b - a)\mu[U^f(a, b)] \leq \sup_{N \in \mathbb{N}} \mu[(f_N - a)^+]. \quad (2)$$

Our formalization closely follows the proof presented in [10] while making some small adjustments. In contrast to proof from [10] in which the theorem is reduced to the case where $a = 0$ and $f_n \geq 0$ for all n , we make a weaker reduction where the theorem is reduced to the case where $0 \leq f_0$ and $a \leq f_N$ using the equality

$$U_N^{(f-a)^+}(0, b-a) = U_N^f(a, b).$$

In particular, fixing $a < b \in \mathbb{R}$ and $N \in \mathbb{N}$, it suffices to prove

$$(b - a)\mu[U_N^f(a, b)] \leq \mu[f_N]$$

in the special case where $0 \leq f_0$ and $a \leq f_N$. Indeed, in the general case, we may apply the above inequality to $((f_n - a)^+)_n$ where $((f_n - a)^+)_n$ is also a submartingale by a simple computation using the monotonicity of the conditional expectation. The reason for this weaker reduction will be made clear once we have presented the proof.

With the above reduction, by constructing a suitable predictable process satisfying the following lemma, the inequality follows by considering the submartingale property of its discrete stochastic integral with some predictable process.

Definition 3.3. A stochastic process $(c_n)_{n \in \mathbb{N}}$ is said to be predictable with respect to the filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$ if for all n , c_{n+1} is \mathcal{F}_n -measurable.

Lemma 3.4. Given a processes $(f_n)_{n \in \mathbb{N}}$ and $(c_n)_{n \in \mathbb{N}}$, let us denote the discrete stochastic integral of c with f the process

$$(c \cdot f)_n := \sum_{k=0}^{n-1} c_{k+1}(f_{k+1} - f_k).$$

Then, if $(f_n)_{n \in \mathbb{N}}$ is a submartingale adapted to the filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$, and $(c_n)_{n \in \mathbb{N}}$ is predictable, non-negative and bounded, $((c \cdot f)_n)$ is also a submartingale.

We omit the proof of this lemma as it follows by a simple computation.

```
lemma submartingale.sum_mul_sub [is_finite_measure μ]
  {R : ℝ} {ξ f : ℕ → Ω → ℝ}
  (hf : submartingale f ℱ μ) (hξ : adapted ℱ ξ)
  (hbdd : ∀ n ω, ξ n ω ≤ R) (hnonneg : ∀ n ω, 0 ≤ ξ n ω) :
  submartingale
  (λ n, ∑ k in finset.range n, ξ k * (f (k + 1) - f k)) ℱ μ
```

With the above lemma, we will now construct a predictable process satisfying the above lemma and use the submartingale property to obtain an inequality bounding the discrete stochastic integral. Define $c_n := \sum_{k < N} 1_{[\sigma_k, \tau_{k+1})}(n)$. It is clear that $(1 - c_n)$ is non-negative, bounded and predictable, and hence, $(1 - c) \cdot f$ is again a submartingale. Then, by the very definition of a submartingale, $0 \leq \mu[((1 - c) \cdot f)_0] \leq \mu[((1 - c) \cdot f)_N]$ implying

$$\mu[(c \cdot f)_N] \leq \mu[(1 \cdot f)_N] = \mu[f_N] - \mu[f_0] \leq \mu[f_N].$$

Now, by noting

$$\begin{aligned} (c \cdot f)_N &= \sum_{n \leq N} \sum_{k \leq N} 1_{[\sigma_k, \tau_{k+1})}(n)(f_{n+1} - f_n) \\ &= \sum_{k \leq N} \sum_{n \leq N} 1_{[\sigma_k, \tau_{k+1})}(n)(f_{n+1} - f_n) \\ &= \sum_{k \leq N} (f_{\sigma_{k+1}} - f_{\sigma_k} + f_{\sigma_{k+2}} - \dots + f_{\tau_{k+1}} - f_{\tau_{k+1}-1}) \\ &= \sum_{k \leq N} (f_{\tau_{k+1}} - f_{\sigma_k}) \geq \sum_{k < U_N(a, b)} (b - a) \\ &= (b - a)U_N(a, b) \end{aligned}$$

where the inequality follows since

- for $k < U_N(a, b)$, $f_{\tau_{k+1}} - f_{\sigma_k} \geq b - a$;
- for $k > U_N(a, b)$, $f_{\tau_{k+1}} = f_{\sigma_k} = f_N$ implying

$$f_{\tau_{k+1}} - f_{\sigma_k} = 0;$$

- for $k = U_N(a, b)$,

$$f_{\tau_{k+1}} - f_{\sigma_k} = f_{\tau_{U_N(a, b)+1}} - f_{\sigma_{U_N(a, b)}} = f_N - a \geq 0.$$

We have

$$(b - a)\mu[U_N(a, b)] \leq \mu[(c \cdot f)_N] \leq \mu[f_N],$$

as required.

The overall proof strategy is conducted as described above with a large part of the formalization effort spent building a sufficiently large API for the definitions `upper_crossing` and `lower_crossing`. Indeed, due to their definitions, many seemingly obvious statements regarding (σ_n) and (τ_n) often involve long and tedious mutual inductions. As an example, while the equality

$$U_N^{(f-a)^+}(0, b-a) = U_N^f(a, b),$$

is rather intuitive: the number of upcrossings is invariant under translations and reflections provided we also translate the bands by the same amount; the formalization of this fact spanned almost 100 lines of code. As a result, we sought to minimize proof steps requiring properties of (σ_n) and (τ_n) resulting in the aforementioned weaker reduction compared to that is presented in [10].

With this, the first inequality of Doob's upcrossing estimate is formalized following the above outlined proof in Lean as

```
lemma upcrossing_estimate [is_finite_measure μ]
  (a b : ℝ) (hf : submartingale f ℱ μ) (N : ℕ) :
  (b - a) * μ[upcrossing_before a b f N] ≤ μ[λ x, (f N x - a)^+]
```

We note that in contrast to the theorem statement, our formulation does not require $a < b$. Indeed, this assumption is not necessarily as in the case $a \geq b$, $b - a < 0$ implying the left hand side is non-positive while the right hand side is non-negative. This remark may appear inconsequential, but it means that we dispense the user of that lemma from having to prove trivial inequalities like $a < b$, which is very desirable in a formal library.

To formalize inequality (2) of Doob's upcrossing estimate, we will introduce a new definition known as `upcrossings`. While we have defined $U_N^f(a, b)$ to take value in the natural numbers, $U^f(a, b)$ cannot, as it is not necessarily finite since $(U_N^f(a, b))_N$ might not be bounded above in N . In the case that $(U_N^f(a, b))_N$ is not bounded above, we say on paper that

$$\sup_N U_N^f(a, b) = \infty.$$

This does not translate directly into Lean as the supremum of an unbounded set of natural numbers is defined to be 0. As a result, to define $U^f(a, b)$ so that $U^f(a, b) = \infty$ makes sense, we need to first coerce $U_N^f(a, b)$ to take value in the extended non-negative reals $\mathbb{R}_{\geq 0} \cup \{\infty\}$.

```
def upcrossings [preorder ι] [order_bot ι] [has_Inf ι]
  (a b : ℝ) (f : ι → Ω → ℝ) (ω : Ω) : ℝ≥0∞ :=
  ⌊ N, (upcrossings_before a b f N ω : ℝ≥0∞)
```

We note that $U_N^f(a, b)$ is coerced to $\mathbb{R}_{\geq 0} \cup \{\infty\}$ rather than $\mathbb{N} \cup \{\infty\}$ since as we intend to integrate $U_N^f(a, b)$, if $U_N^f(a, b)$ is coerced into $\mathbb{N} \cup \{\infty\}$, the integration will introduce another coercion into $\mathbb{R}_{\geq 0} \cup \{\infty\}$, making the original coercion redundant.

With this definition, we may now phrase inequality (2) as

```
lemma mul_lintegral_upcrossings_le_lintegral_pos_part
  [is_finite_measure μ] (a b : ℝ) (hf : submartingale f ℱ μ) :
  ennreal.of_real (b - a) * ∫- ω, upcrossings a b f ω ∂μ ≤
  ⌊ N, ∫- ω, ennreal.of_real ((f N ω - a)^+) ∂μ
```

In contrast to the first inequality in Doob's upcrossing estimate where the expectation is formulated using the Bochner integral, as `upcrossings` takes value in the non-negative extended reals, the target space is no longer a Banach space and so we may not use the Bochner integral. Instead, the expectation is formulated using the lower Lebesgue integral (known as `lintegral` in Lean).

The proof of inequality (2) is rather simple. Indeed, by taking the supremum of both sides of inequality (1), inequality (2) follows by the monotone convergence theorem where we notice $(U_N^f(a, b))_N$ is monotone in N . However, it is important to note that, to apply the monotone convergence theorem, we need to make sure $(U_N^f(a, b))_N$ is both measurable and integrable for all N . Integrability is easy to check as $U_N^f(a, b)$ is bounded above by N , and so it suffices to show measurability. Indeed, by noting that

$$U_N^f(a, b) = \sum_{i=1}^N 1_{\{n \in \mathbb{N} \mid \sigma_n < N\}}(i)$$

$U_N^f(a, b)$ is measurable as $\{n \mid \sigma_n < N\}$ is a measurable set since (σ_n) is a stopping time.

With the Doob's upcrossing estimate, we may now finish the proof of the almost everywhere martingale convergence theorem. We recall that a real sequence (x_n) converges if and only if

- (a) $(|x_n|)$ is bounded above;
- (b) and for all $a < b \in \mathbb{Q}$, (x_n) does not visit below a and above b infinitely often.

Thus, for a submartingale (f_n) adapted to the filtration (\mathcal{F}_n) such that for all n , $\|f_n\|_1 \leq R$ for some $R \in \mathbb{R}$, we will show that (f_n) satisfies the above two conditions almost everywhere.

Firstly, as (f_n) is uniformly bounded in L^1 by R , by inequality (2) of Doob's upcrossing estimate, we have

$$(b - a)\mu[U^f(a, b)] \leq \sup_{N \in \mathbb{N}} \mu[(f_N - a)^+] \leq R + |a| \cdot \mu(\Omega),$$

implying $\mu[U^f(a, b)]$ is finite and so, $U^f(a, b)$ is finite almost everywhere. Thus, as $U^f(a, b) = \infty$ if and only if (f_n) does not visit below a and above b infinitely often, (f_n) satisfies condition (b) almost everywhere.

On the other hand, $(|f_n|)$ is bounded almost everywhere. Indeed, in the case that $(|f_n(\omega)|)$ is not bounded above, either

- one of $\liminf_n f_n(\omega)$ or $\limsup_n f_n(\omega)$ equals $-\infty, +\infty$ respectively with the other being finite,
- or $f_n(\omega) \rightarrow \pm\infty$.

The first case is a subset of $\{U^f(a, b) = \infty\}$ and so can only happen with measure 0 as shown above. The second case is

a subset of $\{\liminf_n |f_n| = \infty\}$. By Fatou's lemma, we have

$$\mu[\limsup_{n \rightarrow \infty} |f_n|] \leq \limsup_{n \rightarrow \infty} \mu[|f_n|] < \infty$$

implying $\liminf_n |f_n| < \infty$ almost everywhere and so, the second case can also only happen with measure 0.

With this, we conclude that (f_n) converges almost everywhere and it remains to show that it has an a.e. limit which is integrable and measurable with respect to \mathcal{F}_∞ . The definition of a.e. measurable becomes especially useful here. Traditionally, one obtains measurability by constructing the a.e. limit to be $\liminf_n f_n$ which is \mathcal{F}_∞ -measurable. Instead, as an a.e. pointwise limit of a sequence of measurable functions is at least a.e. measurable, since f_n is \mathcal{F}_∞ -measurable for all n , if $f_n \xrightarrow{a.e.} f$ then f is \mathcal{F}_∞ -a.e. measurable, namely, there exists some \mathcal{F}_∞ -measurable f' equal to f almost everywhere. Thus, $f_n \xrightarrow{a.e.} f'$ and hence, choosing our limiting function to be f' suffices.

Finally, integrability of the limiting function is clear by considering the following corollary of Fatou's lemma.

Lemma 3.5. *If (f_n) is a sequence of measurable functions such that $f_n \xrightarrow{a.e.} f$, then*

$$\|f\|_p \leq \liminf_{n \rightarrow \infty} \|f_n\|_p,$$

for all $p \in \mathbb{R}_{\geq 0} \cup \{\infty\}$.

We note that our proof is slightly more roundabout than the usual proof found in literature where we conclude the almost everywhere convergence directly after noticing that a real sequence (x_n) converges if and only if $\liminf_n |x_n| < \infty$ and for all $a < b \in \mathbb{Q}$, (x_n) crosses the band $[a, b]$ only finitely many times. The reason choosing the indirect route is that the key observation in our argument where we characterize the convergence of real sequences was independently proven by Gouëzel for the purpose of proving the Lebesgue differentiation theorem [7]. Thus, by changing our proof to utilize Gouëzel's work, we are able to avoid redundant work.

Defining a limit function instead of stating existence.

Our formalization of the a.e. martingale convergence theorem follows the proof outlined above with one notable implementation difference. While we had stated the a.e. martingale convergence theorem as the existence of a limiting random variable satisfying some properties, since applications of the theorem often require us to manipulate this limit, it is easier to simply extract the limit of the process into its own definition. In particular, we define the `limit_process` of a sequence of random variables (f_n) with respect to the filtration (\mathcal{F}_n) and the measure μ as the \mathcal{F}_∞ -measurable, μ -a.e. limit of (f_n) if it exists, otherwise we define it as 0.

```
def limit_process (f : ι → Ω → E) (F : filtration ι A)
  (μ : measure Ω) :=
if h : ∃ g : Ω → E, strongly_measurable[[]] n, F n g ∧
  ∀m ω ∂μ, tendsto (λ n, f n ω) at_top (N (g ω))
then classical.some h else 0
```

Here, using the axiom of choice, `classical.some h` for $h : \exists x : \alpha, p x$ where p is a predicate on the type α provides a term of α which satisfies p .

`limit_process` is a useful definition. By extracting the definition, not only do we avoid needing to deconstruct the a.e. martingale convergence theorem every time we need to manipulate the limiting random variable, we can also directly reason about the definition itself. This is exemplified by the following lemmas in which f is not required to be a submartingale.

```
lemma strongly_measurable_limit_process :
  strongly_measurable[[]] n, F n (limit_process f F μ)
```

```
lemma mem_Lp_limit_process_of_snorm_bdd
  (hfm : ∀ n, ae_strongly_measurable (f n) μ)
  (hbdd : ∀ n, snorm (f n) p μ ≤ R) :
  mem_Lp (limit_process f F μ) p μ
```

With this definition, the a.e. martingale convergence theorem is phrased and formalized as

```
lemma submartingale.ae_tendsto_limit_process
  [is_finite_measure μ] (hf : submartingale f F μ)
  (hbdd : ∀ n, snorm (f n) 1 μ ≤ R) :
  ∀m ω ∂μ, tendsto (λ n, f n ω) at_top
  (N (F.limit_process f μ ω))
```

3.3 L^1 martingale convergence theorem

Continuing, we now formulate and prove the L^1 martingale convergence theorem. The L^1 martingale convergence theorem is commonly phrased as the following equivalence.

Theorem 3.6 (L^1 martingale convergence). *Given $(f_n)_{n \in \mathbb{N}}$ a martingale adapted to the filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$, the following are equivalent:*

- (i) (f_n) is uniformly integrable (in the probability sense);
- (ii) (f_n) converges in L^1 to some \mathcal{F}_∞ -measurable, integrable function f ;
- (iii) there exists some \mathcal{F}_∞ -measurable and integrable f such that for all $n \in \mathbb{N}$, $\mu[f | \mathcal{F}_n] = f_n$ a.e. .

However, while this theorem is stated for martingales, several implications from the proof of this theorem also hold for submartingales. So, rather than proving the theorem as stated above, we extract these implications into their own statements and shall instead consider the following slightly more general theorem.

Theorem 3.7. *Given some filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$,*

- (a) *if $(f_n)_{n \in \mathbb{N}}$ is a submartingale adapted to (\mathcal{F}_n) and is uniformly integrable in the probability sense, it converges in L^1 to some \mathcal{F}_∞ -measurable, integrable function f ;*
- (b) *if furthermore, (f_n) is a martingale, then for all $n \in \mathbb{N}$,*

$$\mu[f | \mathcal{F}_n] = f_n \text{ a.e. ;}$$

- (c) *finally, if g is some integrable, \mathcal{F}_∞ -measurable function, then $(\mu[g | \mathcal{F}_n])_{n \in \mathbb{N}}$ converges to g a.e. and in L^1 .*

Indeed, these three statements implies the L^1 martingale convergence theorem directly where the implication (i) \implies (ii) follows from (a), the implication (ii) \implies (iii) follows from (b), and the implication (iii) \implies (i) follows from (c).

As we have already constructed a sufficiently large amount of API for uniform integrability, the proof of theorem 3.7 is rather simple. By recalling that uniformly integrable in the probability sense implies uniform boundedness in L^1 , (a) follows. Indeed, as (f_n) is bounded in L^1 , by the a.e. martingale convergence theorem we have (f_n) converges almost everywhere to some \mathcal{F}_∞ -measurable, integrable function f . Thus, as convergence a.e. implies convergence in measure, we have by the Vitali convergence theorem that (f_n) converges in L^1 to f as required.

Now, taking (f_n) to be a martingale and denoting its L^1 limit by f (which exists by (a)), (b) follows since for all n , we have for all $m \geq n$,

$$\|f_n - \mu[f \mid \mathcal{F}_n]\|_1 = \|\mu[f_m - f \mid \mathcal{F}_n]\|_1 \leq \|f_m - f\|_1.$$

Thus, $\|f_n - \mu[f \mid \mathcal{F}_n]\|_1 = 0$ by taking $m \rightarrow \infty$ implying $f_n = \mu[f \mid \mathcal{F}_n]$ a.e.

Finally, we notice on one hand that any class of the form $\{\mu[g \mid \mathcal{G}_i] \mid i \in \iota\}$ for some ι -indexed sub- σ -algebras $(\mathcal{G}_i)_{i \in \iota}$ is uniformly integrable, and on the other hand that $(\mu[g \mid \mathcal{F}_n])_{n \in \mathbb{N}}$ is a martingale, hence by (b) it converges a.e. and in L^1 to some function g' . Furthermore, $\mu[g \mid \mathcal{F}_n] = \mu[g' \mid \mathcal{F}_n]$ almost everywhere for all n . Thus, by the law of iterative expectation, for all $A \in \mathcal{F}_n$, we have

$$\mu[g 1_A] = \mu[\mu[g \mid \mathcal{F}_n] 1_A] = \mu[\mu[g' \mid \mathcal{F}_n] 1_A] = \mu[g' 1_A]$$

for all n . Hence, as $\bigcup_n \mathcal{F}_n$ is a π -system generating \mathcal{F}_∞ , by Dynkin's theorem, the equality also holds for all $A \in \mathcal{F}_\infty$. Thus, as both g, g' are \mathcal{F}_∞ measurable, we conclude that $g = g'$ almost everywhere and so, $\mu[g \mid \mathcal{F}_n] \rightarrow g$ a.e. and in L^1 as required.

This theorem in Lean is stated and proved as described above in three parts. Similar to the statement of the a.e. martingale convergence theorem, instead of showing the existence of an L^1 -limit, by noting that the L^1 limit coincides with the a.e. limit, we may alternatively phrase the theorem using `limit_process`.

```
lemma submartingale.tendsto_snorm_one_limit_process
(hf : submartingale f  $\mathcal{F}$   $\mu$ )
(hunif : uniform_integrable f 1  $\mu$ ) :
tendsto ( $\lambda$  n, snorm (f n -  $\mathcal{F}$ .limit_process f  $\mu$ ) 1  $\mu$ )
at_top (N 0)
```

```
lemma martingale.ae_eq_condexp_limit_process
(hf : martingale f  $\mathcal{F}$   $\mu$ )
(hbdd : uniform_integrable f 1  $\mu$ ) (n :  $\mathbb{N}$ ) :
f n =m  $\mu$  [ $\mathcal{F}$ .limit_process f  $\mu$  |  $\mathcal{F}$  n]
```

```
lemma integrable.tendsto_snorm_condexp
(hg : integrable g  $\mu$ )
(hgmeas : strongly_measurable ([ ] n,  $\mathcal{F}$  n) g) :
tendsto ( $\lambda$  n, snorm ( $\mu$ [g |  $\mathcal{F}$  n] - g) 1  $\mu$ ) at_top (N 0)
```

As the proof of this theorem only uses standard analysis arguments, the formalization was straightforward with the help of existing lemmas in `mathlib`.

3.4 Lévy's generalization of the Borel-Cantelli lemmas

The martingale convergence theorems has numerous applications in probability theory and provides alternative proofs to the SLLN, the Radon-Nikodym theorem, the Borel-Cantelli lemmas etc. While the SLLN and the Radon-Nikodym theorem have already been formalized in `mathlib` via other methods, the second Borel-Cantelli lemma is yet to be formalized. Thus, as an application the martingale convergence theorems, we formalize a more general version of the second Borel-Cantelli lemma known as Lévy's generalized Borel-Cantelli lemma.

Theorem 3.8 (Lévy's generalized Borel-Cantelli lemma). *Given a filtration $(\mathcal{F}_n)_{n \in \mathbb{N}}$ and a sequence of sets $(S_n)_{n \in \mathbb{N}}$ such that $S_n \in \mathcal{F}_n$ for all n , then*

$$\limsup_{n \rightarrow \infty} S_n = \left\{ \sum_{n=0}^{\infty} \mu[1_{S_{n+1}} \mid \mathcal{F}_n] = \infty \right\}$$

almost everywhere.

This is formalized in Lean as follows:

```
theorem ae_mem_limsup_at_top_iff [is_finite_measure  $\mu$ ]
{s :  $\mathbb{N}$   $\rightarrow$  set  $\Omega$ ] (hs :  $\forall$  n, measurable_set [ $\mathcal{F}$  n] (s n)) :
 $\forall^m \omega \partial \mu, \omega \in \text{limsup at\_top s} \iff$ 
tendsto ( $\lambda$  n,  $\sum$  k in finset.range n,
 $\mu$ [(s (k + 1)).indicator (1 :  $\Omega \rightarrow \mathbb{R}$ ) |  $\mathcal{F}$  k]  $\omega$ )
at_top at_top
```

We shall omit the details of the proof here and refer the reader to [10]. The proof of this theorem illustrated the usage Doob's decomposition theorem to construct martingales. In particular, to prove Lévy's generalized Borel-Cantelli lemma, we need to construct the martingale defined by

$$f_n := \sum_{k < n} (1_{S_{k+1}} - \mu[1_{S_{k+1}} \mid \mathcal{F}_k]).$$

It is easy to check that this is indeed a martingale by using the law of iterated expectations. However, f_n is in fact the martingale part from Doob's decomposition of the process

$$\tilde{f}_n := \sum_{k < n} 1_{S_{k+1}}.$$

Hence, rather than checking (f_n) is a martingale explicitly and since that notion is more generally useful, we chose to implement Doob's decomposition of a process into a martingale part and a predictable part. To complete our proof, it then suffices to check that (f_n) equals the martingale part of some adapted process (\tilde{f}_n) .

We remark that an alternative approach is to prove that a process of the form $f_n := \sum_{k \leq n} g_k$ is a martingale if (g_n) is an adapted process and $\mu[g_n] = 0$ for all n .

4 Concluding thoughts

We added to `mathlib` definitions of the conditional expectation, filtrations, stopping times and martingales and we formalized Doob’s convergence theorems about those objects. Formalizing theorems about the new definitions was essential, since it is very hard to know if a definition will be adequate until we actually use it, as exemplified by our change of measurability conditions for `uniform_integrable`. The analysis and measure theory sections of `mathlib` have now matured enough for probability theory to be developed and we see no obstacle to the formalization of other fundamental tools of probability. We saw with the betting example in the introduction that martingales can be used to model the interaction of an player with a stochastic environment. Suppose now that the player can choose among several coins which might be differently biased, and makes that decision adaptively based on the result of the previously observed coin flips⁴. Then the distribution of the next reward is random since it depends on past observations. To model such interactions we need Markov kernels and regular conditional probability, which we see as a natural next step in our formalization effort.

Retrospection of a design decision. The formalization of Lévy’s generalized Borel-Cantelli illustrates an annoyance with our current formulation of the theory of stopping times. Our current formulation of stopping times is similar to the Isabelle/HOL [9] formulation in the sense that we let the range of a stopping time to be the same type as the time index of its associated filtration: if the filtration is indexed by \mathbb{N} , the stopping time must be finite. As demonstrated in the definition of hitting times, this is slightly different to what we write on paper where it is allowed to be infinity. The same is true for stopping times. Indeed, for a stopping time σ , on paper we normally allow σ to take value infinity and then work with $\sigma \wedge n$. This is not possible with our current set up. This discrepancy is not problematic in our formalization of the martingale convergence theorems, however we observe that it introduces code duplication in the formalization of the Lévy’s generalized Borel-Cantelli. In particular, Lévy’s generalized Borel-Cantelli requires a lemma known as the one-sided martingale bound which proof starts by introducing a stopping time σ which can take value infinity and argue that the stopped process $(f_{\sigma \wedge n})$ of a submartingale (f_n) is also a submartingale. We however cannot proceed like this since we cannot have a stopping time with possibly infinite value. Instead, we introduce an object which is almost a stopped process but not exactly since we don’t have a stopping time, and again reprove a few lemmas specifically for it. A possible way to avoid this duplication is to change the definition of stopping time to allow for infinity in the range, namely change the definition to

```
def is_stopping_time [preorder ι] (F : filtration ι A)
  (τ : Ω → with_top ι) :=
  ∀ i, measurable_set[F i] {ω | τ ω ≤ i}
```

`with_top ι` is the type ι together with an infinity value. Not having the same type for the filtration index and the stopping time range could however be cumbersome since we then can’t directly index by a stopping time, and only our upcoming refactor attempt will tell which of the definitions is easier to use.

Acknowledgements

The authors would like to thank Sébastien Gouézel and Kalle Kytölä for their help and suggestions during the formalization process. They also thank the other `mathlib` reviewers and contributors shaping `mathlib` without which this project would not have been possible. They would also like to thank the conference reviewers for their helpful comments.

References

- [1] Reynald Affeldt and Manabu Hagiwara. 2012. Formalization of Shannon’s Theorems in SSReflect-Coq. In *3rd Conference on Interactive Theorem Proving (ITP 2012)*, Princeton, New Jersey, USA, August 13–15, 2012 (Lecture Notes in Computer Science, Vol. 7406). Springer, 233–249.
- [2] Jeremy Avigad, Johannes Hölzl, and Luke Serafin. 2017. A formally verified proof of the central limit theorem. *Journal of Automated Reasoning* 59, 4 (2017), 389–423.
- [3] Anne Baanen. 2022. Use and Abuse of Instance Parameters in the Lean Mathematical Library. In *13th International Conference on Interactive Theorem Proving (ITP 2022) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 237)*, June Andronick and Leonardo de Moura (Eds.), Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 4:1–4:20. <https://doi.org/10.4230/LIPIcs.ITP.2022.4>
- [4] Floris van Doorn, Gabriel Ebner, and Robert Y Lewis. 2020. Maintaining a library of formal mathematics. In *International Conference on Intelligent Computer Mathematics*. Springer, 251–267.
- [5] Manuel Eberl. 2021. The Laws of Large Numbers. *Archive of Formal Proofs* (February 2021). https://isa-afp.org/entries/Laws_of_Large_Numbers.html, Formal proof development.
- [6] Mnacho Echenim, Hervé Guiol, and Nicolas Peltier. 2020. Formalizing the Cox–Ross–Rubinstein Pricing of European Derivatives in Isabelle/HOL. *Journal of Automated Reasoning* 64, 4 (April 2020), 737–765. <https://doi.org/10.1007/s10817-019-09528-w>
- [7] Sébastien Gouézel. 2022. A formalization of the change of variables formula for integrals in mathlib. <https://doi.org/10.48550/ARXIV.2207.12742>
- [8] Tuomas Hytönen, Jan Van Neerven, Mark Veraar, and Lutz Weis. 2016. *Analysis in Banach spaces*. Vol. 12. Springer.
- [9] Johannes Hölzl, Andreas Lochbihler, and Dmitriy Traytel. 2015. A Zoo of Probabilistic Systems. *Archive of Formal Proofs* (May 2015). https://isa-afp.org/entries/Probabilistic_System_Zoo.html, Formal proof development.
- [10] Olav Kallenberg. 2021. *Foundations of Modern Probability* (3rd ed.). Springer. 196–197 pages.
- [11] Tor Lattimore and Csaba Szepesvári. 2020. *Bandit algorithms*. Cambridge University Press.
- [12] The mathlib Community. 2020. The Lean Mathematical Library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (New Orleans, LA, USA) (CPP 2020)*. Association for Computing Machinery, New York, NY, USA, 367–381. <https://doi.org/10.1145/3372885.3373824>

⁴that interaction is called a stochastic bandit [11]

- [13] Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. 2015. The Lean theorem prover (system description). In *International Conference on Automated Deduction*. Springer, 378–388.
- [14] Michael Norrish and Konrad Slind. 2022. *The HOL System DESCRIPTION* (3 ed.). 206–218 pages.
- [15] Hiroyuki Okazaki, Yuichi Futa, and Yasunari Shidama. 2013. Formal definition of probability on finite and discrete sample space for proving security of cryptographic systems using Mizar. *Artificial Intelligence Research* 2 (05 2013). <https://doi.org/10.5430/air.v2n4p37>
- [16] Konrad Slind and Michael Norrish. 2008. A brief overview of HOL4. In *International Conference on Theorem Proving in Higher Order Logics*. Springer, 28–32.
- [17] Joseph Tassarotti and Robert Harper. 2019. A separation logic for concurrent randomized programs. *Proceedings of the ACM on Programming Languages* 3, POPL (2019), 1–30.
- [18] Koundinya Vajjha, Barry Trager, Avraham Shinnar, and Vasily Pestun. 2022. Formalization of a Stochastic Approximation Theorem. In *13th International Conference on Interactive Theorem Proving (ITP 2022) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 237)*, June Andronick and Leonardo de Moura (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 31:1–31:18. <https://doi.org/10.4230/LIPIcs.ITP.2022.31>
- [19] David Williams. 1991. *Probability with martingales*. Cambridge university press.

Received 2022-09-21; accepted 2022-11-21