2. Kolmogorov Extension Theorem

We will now address the following question: Do the homogeneous Poisson process and Brownian motion exist? And if so, in what sense are they unique? This will turn out to be an application of an "extension theorem" due to A.N. Kolmogorov. Credit is sometimes given also to P.J. Daniell for his development of infinite-dimensional integration theory.

2.1 Finite-dimensional distributions.

We start by recalling the notion of a *product measure space*: Given two measure spaces $(\mathscr{X},\Sigma)$ and $(\mathscr{X}',\Sigma')$, their product is a measure space $(\mathscr{X}\times\mathscr{X}',\Sigma\otimes\Sigma')$ where

$$\Sigma\otimes\Sigma' := \sigma\Bigl(\bigl\{A\times A'\colon A\in\Sigma,\, A'\in\Sigma'\bigr\}\Bigr) \tag{2.1}$$

Here $\sigma(\mathcal{A})$ denotes the least σ-algebra containing the collection of sets $\mathcal{A}$. As is readily checked, the formation of product spaces is an associative operation. This allows us to define powers $\Sigma^{\otimes n}$ inductively as $\Sigma\otimes\Sigma^{\otimes(n-1)}$. (A direct definition is possible as well.)

A stochastic process $\{X_t\colon t\in T\}$ gives us access to the law of a single variable X_t, but also the joint law of X_t and X_s, etc. This naturally leads to:

Definition 2.1 *Let $(\mathscr{X},\Sigma)$ be a measurable space and let $\{X_t\colon t\in T\}$ be an $\mathscr{X}$-valued stochastic process. The measures indexed by $n\geqslant 1$ and $(t_1,\dots,t_n)\in T$ and defined by*

$$\forall A\in\Sigma^{\otimes n}\colon\quad \mu_{(t_1,\dots,t_n)}(A) := P\bigl((X_{t_1},\dots,X_{t_n})\in A\bigr) \tag{2.2}$$

form a family of finite-dimensional distributions *of X.*

A moment's thought reveals that not every family of probability measures on products of $(\mathscr{X},\Sigma)$ indexed by elements of T can arise from a stochastic process. Indeed, as is readily checked, (2.2) obeys both (1) and (2) in:

Definition 2.2 (Consistent family of measures) *Let $(\mathscr{X},\Sigma)$ be a measure space and let $\mathcal{S}_n$ denote the set of all permutations of $\{1,\dots,n\}$. A family*

$$\bigl\{\mu_{\underline{t}}\colon \underline{t}=(t_1,\dots,t_n)\in T^n,\, n\geqslant 1\bigr\}, \tag{2.3}$$

where $\mu_{(t_1,\dots,t_n)}$ is a probability measure on $(\mathscr{X}^n,\Sigma^{\otimes n})$, is said to be consistent *if*

(1) $\forall n\geqslant 1\,\forall\sigma\in\mathcal{S}_n\,\forall t_1,\dots,t_n\in T\,\forall A_1,\dots,A_n\in\Sigma$:

$$\mu_{(t_{\sigma(1)},\dots,t_{\sigma(n)})}\bigl(A_{\sigma(1)}\times\dots\times A_{\sigma(n)}\bigr) = \mu_{(t_1,\dots,t_n)}(A_1\times\dots\times A_n) \tag{2.4}$$

(2) $\forall n\geqslant 1\,\forall t_1,\dots,t_n\in T\,\forall A_1,\dots,A_{n-1}\in\Sigma$:

$$\mu_{(t_1,\dots,t_n)}(A_1\times\dots\times A_{n-1}\times\mathscr{X}) = \mu_{(t_1,\dots,t_{n-1})}(A_1\times\dots\times A_{n-1}) \tag{2.5}$$

Consistency is thus necessary for a family of measures to be the finite-dimensional distributions of process. A deep, and perhaps surprising, fact is that consistency is also sufficient for the existence of a process, provided the space $(\mathscr{X},\Sigma)$ is "nice" — meaning, conforming to the following definition:

Definition 2.3 (Standard Borel space) *A standard Borel space is a measure space $(\mathscr{X},\Sigma)$ where $\mathscr{X}$ is a completely-metrizable second-countable topological space and Σ is the σ-algebra of Borel sets — i.e., the smallest σ-algebra containing all open sets.*

We remark that, while the above definition talks about topological spaces, its requirements simply state that, for a suitable metric, $\mathscr{X}$ is a complete and separable metric space. Besides $\mathbb{R}$ or $\mathbb{R}^d$ endowed with the Euclidean norm, the class of standard Borel spaces thus includes all finite and countable sets endowed with the discrete metric, the Banach space $C([0,1])$ of continuous functions on $[0,1]$ endowed with the supremum norm, the Banach spaces $L^p(\mathbb{R})$ and $\ell^p(\mathbb{N})$ for all $p \in [1,\infty)$, the set of Borel probability measures on $\mathbb{R}$ endowed with the topology of weak convergence, etc.

From a certain perspective, Definition 2.3 may appear too restrictive. Indeed, the celebrated Kuratowski Theorem states that any uncountable standard Borel space $(\mathscr{X},\Sigma)$ is bi-measurably isomorphic to the interval $[0,1]$ endowed with the Borel sets of Euclidean topology. On the other hand, since the isomorphism is only required to be bi-measurable (which is less than homeomorphic), the class of standard Borel spaces is actually so rich that we rarely need to worry about anything else.

We are ready for the main result of this lecture:

Theorem 2.4 (Kolmogorov Extension Theorem) *Let $(\mathscr{X},\Sigma)$ be a standard Borel space and let T be a non-empty set. Then for any consistent family*

$$\{\mu_{\underline{t}} \colon \underline{t} = (t_1,\dots,t_n) \in T^n,\, n \geqslant 1\} \tag{2.6}$$

where $\mu_{(t_1,\dots,t_n)}$ is a measure on $(\mathscr{X}^n,\Sigma^{\otimes n})$, there exists a probability space $(\Omega,\mathcal{F},P)$ and an $\mathscr{X}$-valued stochastic process $\{X_t \colon t \in T\}$ on $(\Omega,\mathcal{F},P)$ such that

$$\forall n \geqslant 1\, \forall t_1,\dots,t_n \in T\, \forall A \in \Sigma^{\otimes n} \colon\; P\big((X_{t_1},\dots,X_{t_n}) \in A\big) = \mu_{(t_1,\dots,t_n)}(A) \tag{2.7}$$

In short, the family (2.6) *are the finite-dimensional distributions of X.*

Before we delve into the proof, let us remark that this theorem, along with the problem of disintegration of measure, are the main reasons why we sometimes have to assume an underlying topological structure of our probability space. There have in fact been attempts (e.g., by D. Blackwell) to cast the foundations of probability in terms of topological spaces. However, this turns out to be an overkill: The additional structure is often needed just tangentially and putting it into the foundations carries a lot of unnecessary technical overhead that can, for the most part, be avoided.

2.2 Proof of Kolmogorov Extension Theorem.

As is not hard to guess, the proof reduces to a measure-extension problem. For this we recall that an *algebra* on Ω is a collection of subsets of Ω containing $\emptyset$ and Ω which is closed under finite unions and complements. We aim to invoke the following standard measure-extension tool:

Theorem 2.5 (Hahn-Kolmogorov) *Let $\mathcal{A}$ be an algebra of sets on Ω and let $\mu \colon \mathcal{A} \to [0,\infty]$ be a set function that is*

(1) *finitely additive, meaning*

$$\forall A, B \in \mathcal{A}\colon \quad A \cap B = \emptyset \;\Rightarrow\; \mu(A \cup B) = \mu(A) + \mu(B) \tag{2.8}$$

(2) *countably subadditive, meaning*

$$\forall \{A_n\}_{n\geqslant 1} \subseteq \mathcal{A}\colon \quad \bigcup_{n\geqslant 1} A_n \in \mathcal{A} \;\Rightarrow\; \mu\Bigl(\bigcup_{n\geqslant 1} A_n\Bigr) \leqslant \sum_{n\geqslant 1} \mu(A_n) \tag{2.9}$$

Then there exists a measure $\bar\mu$ *on* $(\Omega, \sigma(\mathcal{A}))$ *such that*

$$\forall A \in \mathcal{A}\colon \quad \bar\mu(A) = \mu(A). \tag{2.10}$$

Equipped with this result, we are ready to start constructing the underlying probability space. For the sample space we use the set

$$\Omega := \mathscr{X}^T := \bigl\{ \text{ functions } T \to \mathscr{X} \bigr\}. \tag{2.11}$$

This set can naturally be endowed with the product σ-algebra

$$\Sigma^{\otimes T} := \sigma\biggl(\Bigl\{ \bigtimes_{t\in T} A_t\colon \{A_t\colon t \in T\} \subseteq \Sigma \,\wedge\, \{t \in T\colon A_t \neq \mathscr{X}\} \text{ is finite}\Bigr\}\biggr) \tag{2.12}$$

We will also need the notation for events that depend only on a given (typically finite) number of coordinates. Given any $S \subseteq T$, for this we let

$$\mathcal{F}_S := \sigma\biggl(\Bigl\{ \bigtimes_{t\in T} A_t\colon \{A_t\colon t \in T\} \subseteq \Sigma \,\wedge\, \{t \in T\colon A_t \neq \mathscr{X}\} \subseteq S\Bigr\}\biggr) \tag{2.13}$$

Let $\pi_S\colon \Omega \to \mathscr{X}^S$ be the projection on the coordinates in S. We now observe:

Lemma 2.6 *Using the above notation,*

$$\mathcal{A} := \bigcup_{\substack{S\subseteq T\\ \text{finite}}} \mathcal{F}_S \tag{2.14}$$

is an algebra on Ω*. Moreover, given a consistent family* (2.6)*, for any* $\mathcal{A} \in \mathcal{A}$ *and any* $S = \{t_1, \dots, t_n\}$ *such that* $A \in \mathcal{F}_S$*, the expression*

$$P(A) := \mu_{(t_1,\dots,t_n)}\bigl(\pi_S(A)\bigr) \tag{2.15}$$

does not depend on S (as long as $A \in \mathcal{F}_S$*) and thus defines a map* $P\colon \mathcal{A} \to [0,1]$*. This map is finitely additive on* $\mathcal{A}$*.*

Proof. That $\mathcal{A}$ is an algebra follows from the fact that $A \in \mathcal{F}_S$ and $B \in \mathcal{F}_{S'}$ imply $A, B \in \mathcal{F}_{S\cup S'}$ and noting that $\mathcal{F}_{S\cup S'}$ is even a σ-algebra. The independence of (2.15) on S is a consequence of the consistency. The finite additivity is again proved by noting that $\mu_{(t_1,\dots,t_n)}$ is even countably additive on $\pi_S(\mathcal{F}_S)$. □

As is directly checked by comparing (2.14) and (2.12),

$$\mathcal{F} = \sigma(\mathcal{A}). \tag{2.16}$$

Motivated by Theorem 2.5, our task is to prove that P is countably subadditive on $\mathcal{A}$. It is here where we will make use of the standard Borel property of $(\mathscr{X},\Sigma)$. There are two facts we will need; these come in the next two lemmas:

Lemma 2.7 (Inner regularity) *Let μ be a finite measure on a standard Borel space $(\mathscr{X},\Sigma)$. Then*

$$\forall A\in\Sigma\colon\quad \mu(A)=\sup\bigl\{\mu(C)\colon C\subseteq A \text{ compact or empty}\bigr\} \tag{2.17}$$

In short, μ is inner regular.

Lemma 2.8 *Let $(\mathscr{X},\Sigma)$ be a standard Borel space. Then so is $(\mathscr{X}^n,\Sigma^{\otimes n})$ for all $n\geqslant 1$.*

Leaving the proof of these lemmas to homework, we now use these to check:

Proposition 2.9 *Using above notation,*

$$\forall\{B_n\}_{n\geqslant 1}\in\mathcal{A}^{\mathbb{N}}\colon\quad B_n\downarrow\varnothing\ \Rightarrow\ P(B_n)\downarrow 0 \tag{2.18}$$

Proof. We will proceed by proving the contrapositive, so let us assume that $\{B_n\}_{n\geqslant 1}\subseteq\mathcal{A}$ is a non-increasing sequence such that

$$\epsilon:=\inf_{n\geqslant 1}P(B_n)>0 \tag{2.19}$$

If we are sloppy about details, the argument is actually very simple: We use the standard Borel property and inner regularity to find, for each $n\geqslant 1$, a set $C_n\subseteq B_n$ such that $\bigcap_{k=1}^n C_k$ has probability very close to that of $\bigcap_{k=1}^n B_k$ and is thus non-empty. It follows that $\bigcap_{k=1}^n C_k$ contains a "point" $x^{(n)}\in\mathscr{X}^T$ for each $n\geqslant 1$. Moreover, C_n will be compact in the coordinates that B_n depends on which will permit us to conclude that $\{x^{(n)}\}_{n\geqslant 1}$ contains a subsequence whose every coordinate converges. The limit point then lies in $\bigcap_{n\geqslant 1}C_n$ and thus in $\bigcap_{n\geqslant 1}B_n$.

The actual argument is notationally quite a bit more complicated because need to move carefully between sets in $\mathscr{X}^T$ and their finite dimensional projections. Note first that, since each B_n belongs to $\mathcal{F}_S$ for some finite $S\subseteq T$, repeating terms if necessary we may assume that for a sequence $\{t_n\}_{n\geqslant 1}$ of elements of T,

$$\forall n\geqslant 1\colon\quad B_n\in\mathcal{F}_{S_n} \tag{2.20}$$

holds with

$$S_n:=\{t_1,\dots,t_n\} \tag{2.21}$$

Next note that $\mu_{(t_1,\dots,t_n)}$ is a measure on $(\mathscr{X}^n,\Sigma^{\otimes n})$ which by Lemma 2.8 is standard Borel. Lemma 2.7 then ensures the existence of a compact set $C_n'\subseteq\pi_{S_n}(B_n)$ such that

$$\mu_{(t_1,\dots,t_n)}\bigl(\pi_{S_n}(B_n)\smallsetminus C_n'\bigr)<\epsilon 2^{-n-1} \tag{2.22}$$

Writing $C_n:=\pi_{S_n}^{-1}(C_n')$, which we note belongs to $\mathcal{F}_{S_n}$, we thus get

$$\forall n\geqslant 1\,\exists C_n\in\mathcal{F}_{S_n}\colon\quad C_n\subseteq B_n\ \wedge\ \pi_{S_n}(C_n)\text{ is compact }\ \wedge\ P(B_n\smallsetminus C_n)<\epsilon 2^{-n-1}. \tag{2.23}$$

Since, for each $n \geqslant 1$,

$$P\Big(\bigcap_{k=1}^n C_k\Big) \geqslant P\Big(\bigcap_{k=1}^n B_k\Big) - \sum_{k=1}^n P(B_k \smallsetminus C_k) \geqslant \epsilon - \sum_{k=1}^n \epsilon 2^{-k-1} \geqslant \frac{\epsilon}{2} > 0 \tag{2.24}$$

we get that $\bigcap_{k=1}^n C_n \neq \emptyset$ for each $n \geqslant 1$. In light of $C_n \in \mathcal{F}_{S_n}$, for each $n \geqslant 1$ this implies existence of $\{x_k^{(n)}\}_{1\leqslant k\leqslant n}$ such that

$$\forall n \geqslant 1\colon \quad (x_1^{(n)}, \dots, x_n^{(n)}) \in \pi_{S_n}\Big(\bigcap_{k=1}^n C_k\Big) \tag{2.25}$$

Thanks to the product structure, for each $m \geqslant 1$, the sequence $\{(x_1^{(n)}, \dots, x_m^{(n)})\}_{n\geqslant m}$ belongs to the compact set $\bigcap_{k=1}^m \pi_{S_m}(C_k)$. It follows that this sequence contains a subsequence that converges in $\mathscr{X}^m$. As this can be done for all $m \geqslant 1$, the Cantor diagonal argument thus produces a subsequence $\{n_j\}_{j\geqslant 1}$ such that

$$\forall k \geqslant 1\colon \; \bar{x}_{t_k} := \lim_{j\to\infty} x_k^{(n_j)} \text{ exists} \tag{2.26}$$

with, again by compactness,

$$\forall n \geqslant 1\colon \quad (\bar{x}_{t_1}, \dots, \bar{x}_{t_n}) \in \pi_{S_n}\Big(\bigcap_{k=1}^n C_k\Big). \tag{2.27}$$

Define x_t as above for $t \in \{t_n : n \geqslant 1\}$ and by $\bar{x}_t := \bar{x}_{t_1}$ else. The "point" $\bar{x} = \{\bar{x}_t\}_{t\in T}$ then obeys $\bar{x} \in \bigcap_{k\geqslant 1}^n C_k$ for all $n \geqslant 1$ and thus $\bar{x} \in \bigcap_{n\geqslant 1} C_n$. As the latter is a subset of $\bigcap_{n\geqslant 1} B_n$, we conclude that $\bigcap_{n\geqslant 1} B_n \neq \emptyset$, thus proving the contrapositive of (2.18). □

We now ready for:

Proof of Theorem 2.4. We start by showing that P is countably additive on $\mathcal{A}$. Let $\{A_n\}_{n\in\mathbb{N}} \in \mathcal{A}^{\mathbb{N}}$ be disjoint (meaning that $A_n \cap A_m \neq \emptyset$ implies $m = n$) and such that $\bigcup_{n\geqslant 1} A_n \in \mathcal{A}$. Writing $B_n := \bigcap_{k>n} A_k$, we then get

$$\bigcup_{n\geqslant 1} A_n = B_n \cup \bigcup_{k=1}^n A_k \tag{2.28}$$

and, by disjointness, $B_n = \bigcup_{n\geqslant 1} A_n \smallsetminus \bigcup_{k=1}^n A_k$ implying $B_n \in \mathcal{A}$. The finite additivity of P on $\mathcal{A}$ then gives

$$P\Big(\bigcup_{n\geqslant 1} A_n\Big) = P(B_n) + \sum_{k=1}^n P(A_k) \tag{2.29}$$

As $B_n \downarrow \emptyset$ thanks to disjointness of $\{A_n\}_{n\geqslant 1}$, we have $P(B_n) \to 0$ as $n \to \infty$ by (2.18). Taking $n \to \infty$ we thus get

$$P\Big(\bigcup_{n\geqslant 1} A_n\Big) = \sum_{k\geqslant 1} P(A_k) \tag{2.30}$$

which is the desired countable additivity.

The Hahn-Kolmogorov Theorem shows that P extends to a measure on $(\Omega, \mathcal{F})$, where $\Omega := \mathscr{X}^T$ and $\mathcal{F} := \Sigma^{\otimes T}$. For each $\omega = \{\omega_t : t \in T\} \in \Omega$, set

$$X_t(\omega) := \omega_t. \tag{2.31}$$

The product structure of $\mathcal{F}$ ensures that X_t is $\mathcal{F}/\Sigma$-measurable and so $\{X_t : t \in T\}$ is a stochastic process on $(\Omega, \mathcal{F}, P)$. The condition (2.7) then follows from (2.15). □

While the Kolmogorov Extension Theorem is a very elegant tool, it is possible to avoid in many cases. This is because the key conclusion (2.18) can be verified by other means with the help of the particular structure of the finite-dimensional distribution. In particular, this will work for *all* product (probability) measures (regardless of the index set) and thus all independent random variables, but also measures with a specific dependency structure such as Markov chains. All that is needed is that the finite-dimensional measures are not just consistent but also disintegrate with respect to one another.

With this being said, we warn the reader that some conditions on the underlying measure space cannot be avoided in general. Indeed, E.S. Andersen and B. Jessen in "On the introduction of measures in infinite product spaces" (Danske Vid. Selsk. Mat.-Fys. Medd., vol.25, no.4 (1948), 8pp) constructed an example of a consistent sequence of probability measures on finite product spaces with no extension to the infinite product space. However, these issues have little bearing on "practical" questions involving probability where we invariably deal with spaces that are standard Borel.

2.3 Uniqueness.

With the tools to construct stochastic processes available, the next question is that of *uniqueness*. For the measure extension problem, this will follow from Dynkin's π/λ-Theorem. Recall the following definitions:

Definition 2.10 *Given a set $\Omega \neq \emptyset$, a collection $\mathcal{P}$ of subsets of Ω is a π-*system *if*

$$\forall A, B \in \mathcal{P}: \ A \cap B \in \mathcal{P} \tag{2.32}$$

*A collection $\mathcal{L}$ of subsets of Ω is a λ-*system*, if*

(1) $\emptyset, \Omega \in \mathcal{L}$
(2) $\forall A, B \in \mathcal{L}: \ A \subseteq B \ \Rightarrow \ B \smallsetminus A \in \mathcal{L}$
(3) $\forall \{A_n\}_{n \geqslant 1} \subseteq \mathcal{L} \, \forall A \in \mathcal{L}: \ A_n \uparrow A \ \Rightarrow \ A \in \mathcal{L}$

Every σ-algebra is a λ-system but there are λ-systems that are not σ-algebras (in general). The example relevant for us comes in:

Lemma 2.11 *Let μ and ν be two probability measures on $(\Omega, \mathcal{F})$. Then*

$$\mathcal{L} := \big\{A \in \mathcal{F} : \ \mu(A) = \nu(A)\big\} \tag{2.33}$$

is a λ-system.

Proof. This follows from the additivity and continuity of measure along with the fact that both measures have equal, and finite, total mass. □

The previous lemma shows that λ-systems naturally appear when studying equality of two measures. The vehicle that provides the needed abstract-nonsense part of the argument is then:

Theorem 2.12 (Dynkin's π/λ-Theorem) *Given a non-emptyset Ω, let $\mathcal{L}$ be a λ-system on Ω and let $\mathcal{P}$ be a π-system on Ω. Then*

$$\mathcal{P} \subseteq \mathcal{L} \;\Rightarrow\; \sigma(\mathcal{P}) \subseteq \mathcal{L} \tag{2.34}$$

Proof. Let $\mathcal{P}$ be a π-system contained in a λ-system $\mathcal{L}$. We start by noting that it suffices to prove the claim with

$$\mathfrak{m}(\mathcal{P}) := \bigcap\bigl\{\mathcal{L}' \colon \lambda\text{-system} \wedge \mathcal{P} \subseteq \mathcal{L}'\bigr\} \tag{2.35}$$

instead of $\mathcal{L}$. Indeed, as the intersection of any number of λ-systems is a λ-system, $\mathfrak{m}(\mathcal{P})$ is a λ-system containing $\mathcal{P}$. Our argument will actually show that $\mathfrak{m}(\mathcal{P})$ is also a π-system and thus a σ-algebra. This gives $\sigma(\mathcal{P}) \subseteq \mathfrak{m}(\mathcal{P}) \subseteq \mathcal{L}$, proving the claim.

For any $A \subseteq \Omega$, denote

$$\mathcal{G}_A := \bigl\{B \in \mathfrak{m}(\mathcal{P}) \colon A \cap B \in \mathfrak{m}(\mathcal{P})\bigr\}. \tag{2.36}$$

Using that $\mathfrak{m}(\mathcal{P})$ is a λ-system, we readily verify that so is $\mathcal{G}_A$, regardless of A. Now observe that $\mathcal{P} \subseteq \mathfrak{m}(\mathcal{P})$ along with $\mathcal{P}$ being a π-system imply $\mathcal{P} \subseteq \mathcal{G}_A$ for each $A \in \mathcal{P}$. The minimality of $\mathfrak{m}(\mathcal{P})$ then gives $\mathfrak{m}(\mathcal{P}) \subseteq \mathcal{G}_A$ for each $A \in \mathcal{P}$. This translates into $\mathcal{P} \subseteq \mathcal{G}_B$ for each $B \in \mathfrak{m}(\mathcal{P})$ and the minimality of $\mathfrak{m}(\mathcal{P})$ then yields $\mathfrak{m}(\mathcal{P}) \subseteq \mathcal{G}_B$ for each $B \in \mathfrak{m}(\mathcal{P})$. It follows that $\mathfrak{m}(\mathcal{P})$ is a π-system, as desired. □

Using Theorem 2.12 we get:

Corollary 2.13 *The extension on $P\colon \mathcal{A} \to [0,1]$ from* (2.15) *to $\Sigma^{\otimes T}$ is unique.*

Proof. The collection of sets $\mathcal{A}$ from (2.14) is an algebra and so it is a π-system. If P and P' are two σ-additive extensions of P to $\sigma(\mathcal{A})$, then they are equal on a λ-system containing $\mathcal{A}$. By (2.34) they are thus equal on $\Sigma^{\otimes T} = \sigma(\mathcal{A})$. □

It is worth noting that, while the measure defined by finite-dimensional distributions on the infinite product space is unique, the induced stochastic process $\{X_t \colon t \in T\}$ is not in general. An obvious obstacle for uniqueness is the fact that, in probability, we do not make a distinction between different realizations of a random variable, as long as these induce the same distribution. A less obvious obstacle comes in the following example: Consider a real-valued stochastic process $X = \{X_t \colon t \in [0,1]\}$ and let U be a Uniform$([0,1])$ and independent of X. Set

$$Y_t := \begin{cases} X_t + 1, & \text{if } t = U, \\ X_t, & \text{else.} \end{cases} \tag{2.37}$$

As is then checked readily from Fubini-Tonelli, $Y_t = X_t$ a.s. for each $t \in [0,1]$. It follows that Y has the same finite dimensional distributions as X, and thus induces the same measure on $(\mathbb{R}^{[0,1]}, \mathcal{B}(\mathbb{R})^{\otimes[0,1]})$. Yet the two processes have clearly distinct sample paths — for instance, if X continuous then Y is not.

Examples of this sort do not (and cannot) arise unless the index space T is uncountable which is why we have been able to avoid dealing with them until now. But 275D is primarily concerned with continuum index sets and so the uniqueness issue deserves a closer look. This is what we will do in the next lecture.