

1. Independence

1.1. Events

A collection of events $\{A_i\} \subset \mathcal{F}$ in some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ are called *independent* if

$$\mathbb{P}[\cap_{i \in I} A_i] = \prod_{i \in I} \mathbb{P}[A_i]$$

for each finite set I of indices. This is a stronger requirement than “pairwise independence,” the requirement merely that

$$\mathbb{P}[A_i \cap A_j] = \mathbb{P}[A_i]\mathbb{P}[A_j]$$

for each $i \neq j$; for a simple counter-example, toss two fair coins and let $A_1 = \{\text{first coin shows Heads}\}$, $A_2 = \{\text{second coin shows Heads}\}$, $A_3 = \{\text{coins disagree}\}$; then each $\mathbb{P}[A_i] = 1/2$ and each $\mathbb{P}[A_i \cap A_j] = 1/4$ for $i \neq j$, but $\cap A_i = \emptyset$ has probability zero.

1.2. Classes of Events

Classes $\{\mathcal{C}_i\}$ of events are called *independent* if

$$\mathbb{P}[\cap_{i \in I} A_i] = \prod_{i \in I} \mathbb{P}[A_i]$$

whenever each $A_i \in \mathcal{C}_i$. An important tool for simplifying the proof of independence is

Theorem 1.1 (Basic Criterion). *If classes $\{\mathcal{C}_i\}$ of events are independent and if each \mathcal{C}_i is a π -system, then $\{\sigma(\mathcal{C}_i)\}$ are independent too.*

Proof. By induction it is enough to consider independent π -systems $\mathcal{C}_1, \mathcal{C}_2$. Fix any $A_2 \in \mathcal{C}_2$ and set

$$\mathcal{L} := \{B \in \mathcal{F} : \mathbb{P}[B \cap A_2] = \mathbb{P}[B] \cdot \mathbb{P}[A_2]\}.$$

Then

- $\Omega \in \mathcal{L}$, obviously;
- $B \in \mathcal{L} \Rightarrow B^c \in \mathcal{L}$, quick computation;
- $B_n \in \mathcal{L}$ and $\{B_n\}$ disjoint $\Rightarrow \cup B_n \in \mathcal{L}$, quick computation.

Thus \mathcal{L} is a λ -system containing \mathcal{C}_1 , and by Dynkin's π - λ theorem it contains $\sigma(\mathcal{C}_1)$. Thus $\sigma(\mathcal{C}_1) \perp\!\!\!\perp \mathcal{C}_2$; the same argument now shows $\{\sigma(\mathcal{C}_1), \sigma(\mathcal{C}_2)\}$ are independent and induction completes the proof. \square

1.3. Random Variables

A collection of random variables $\{X_i\}$ on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ are called *independent* if

$$\mathbb{P}(\cap_{i \in I} [X_i \in B_i]) = \prod_{i \in I} \mathbb{P}[X_i \in B_i]$$

for each finite set I of indices and each collection of Borel sets $\{B_i \in \mathcal{B}(\mathbb{R})\}$. This is just the same as the requirement that the σ -algebras $\mathcal{F}_i := \sigma(X_i) = X_i^{-1}(\mathcal{B})$ be independent; by the Basic Criterion it is enough to check that the joint CDF functions factor, *i.e.*, that

$$\mathbb{P}(\cap_{i \in I} [X_i \leq x_i]) = \prod_{i \in I} F_i(x_i)$$

for each $x \in \mathbb{R}^I$ (or just a dense set of them). For jointly continuous random variables this is equivalent to requiring that the joint density function factor as the product of marginal density functions, while for discrete random variables it's equivalent to the usual factorization criterion for the joint p.m.f.

2. Zero-One Laws

2.1. Borel-Cantelli and Borel's Zero-One Law

Our earlier proof of the Strong Law of Large Numbers for *i.i.d.* bounded random variables relied on the almost-trivial:

Lemma 2.1 (Borel-Cantelli). *Let $\{A_n\}$ be events on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ that satisfy*

$$\sum_{n=1}^{\infty} \mathbb{P}[A_n] < \infty.$$

Then the event that infinitely-many of the $\{A_n\}$ occur (the limit supremum) has probability zero.

Proof.

$$\mathbb{P}\left[\bigcap_{n=1}^{\infty} \bigcup_{m=n}^{\infty} A_n\right] \leq \mathbb{P}\left[\bigcup_{m=n}^{\infty} A_n\right] \leq \sum_{m=n}^{\infty} \mathbb{P}[A_n] \rightarrow 0$$

□

This result does *not* require independence of the $\{A_n\}$, but its partial converse does:

Proposition 2.2 (Borel Zero-One Law). *Let $\{A_n\}$ be independent events on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ that satisfy*

$$\sum_{n=1}^{\infty} \mathbb{P}[A_n] = \infty.$$

Then the event that infinitely-many of the $\{A_n\}$ occur (the limit supremum) has probability one.

Proof. First recall that $1 + x \leq e^x$ for all real $x \in \mathbb{R}$, positive or not. For each pair of integers $1 \leq n \leq N < \infty$,

$$\begin{aligned} \mathbb{P}\left[\bigcap_{m=n}^N A_m^c\right] &= \prod_{m=n}^N (1 - \mathbb{P}[A_m]) \\ &\leq \prod_{m=n}^N e^{-\mathbb{P}[A_m]} = \exp\left(-\sum_{m=n}^N \mathbb{P}[A_m]\right) \\ &\rightarrow \exp\left(-\sum_{m=n}^{\infty} \mathbb{P}[A_m]\right) = e^{-\infty} = 0 \end{aligned}$$

as $N \rightarrow \infty$; thus $\bigcap_{m=n}^{\infty} A_m^c$ is a null set, and

$$\begin{aligned} \mathbb{P}\left[\bigcap_{n=1}^{\infty} \bigcup_{m=n}^{\infty} A_n\right] &= 1 - \mathbb{P}\left[\bigcup_{n=1}^{\infty} \bigcap_{m=n}^{\infty} A_m^c\right] \\ &\geq 1 - \sum_{n=1}^{\infty} \mathbb{P}\left(\bigcap_{m=n}^{\infty} A_m^c\right) = 1 \end{aligned}$$

□

Together these two results comprise a “zero-one law”— for independent events $\{A_n\}$, the limsup $A := \limsup A_n$ has probability $\mathbb{P}(A) = 0$ or $\mathbb{P}(A) = 1$, depending on whether the sum $\sum \mathbb{P}(A_n)$ is finite or not.

2.2. Kolmogorov's Zero-One Law

For any collection $\{X_n\}$ of random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, define two sequences of σ -algebras by:

$$\mathcal{F}_n := \sigma\{X_i : i \leq n\} \quad \mathcal{F}'_n := \sigma\{X_i : i \geq n+1\}$$

and, from them, construct the π -system \mathcal{P} and σ -algebra \mathcal{T} by

$$\mathcal{P} := \bigcup_{n=1}^{\infty} \mathcal{F}_n \quad \mathcal{T} := \bigcap_{n=1}^{\infty} \mathcal{F}'_n.$$

In general \mathcal{P} will not be a σ -algebra, because will not be closed under countable unions and intersections, but it is a π -system, and generates the σ -algebra $\sigma(\mathcal{P}) = \vee \mathcal{F}_n \subseteq \mathcal{F}$.

The class \mathcal{T} is called the *tail* σ -field; it includes such events as “ X_n converges” or “ $\limsup X_n \leq 1$ ” or, with $S_n := \sum_1^n X_j$, “ $\frac{1}{n}S_n$ Converges” or “ $\frac{1}{n}S_n \rightarrow 0$.”

Theorem 2.3 (Kolmogorov's Zero-One Law). *For independent events X_n , the tail σ -field \mathcal{T} is “almost trivial”—that is, every event $\Lambda \in \mathcal{T}$ has probability $\mathbb{P}[\Lambda] = 0$ or $\mathbb{P}[\Lambda] = 1$.*

Proof. Let $A \in \mathcal{P} = \cup \mathcal{F}_n$, and $\Lambda \in \mathcal{T}$. Then for some $n \in \mathbb{N}$, $A \in \mathcal{F}_n$ and $\Lambda \in \mathcal{F}'_n$, so $A \perp\!\!\!\perp \Lambda$. Thus \mathcal{P} and \mathcal{T} are independent; since \mathcal{P} is a π -system, it follows from the Basic Criterion that $\sigma(\mathcal{P})$ and \mathcal{T} are independent. But each X_n is $\sigma(\mathcal{P})$ -measurable, so $\mathcal{T} \subset \sigma(\mathcal{P})$ and each $\Lambda \in \mathcal{T}$ must also be in $\sigma(\mathcal{P}) \perp\!\!\!\perp \mathcal{T}$; thus

$$\mathbb{P}[\Lambda] = \mathbb{P}[\Lambda \cap \Lambda] = \mathbb{P}[\Lambda]\mathbb{P}[\Lambda] = \mathbb{P}[\Lambda]^2,$$

so $0 = \mathbb{P}[\Lambda](1 - \mathbb{P}[\Lambda])$ proving the theorem. \square

3. Fubini

Do independent random variables exist, with arbitrary (marginal) distributions? How can they be constructed?

One way is to build product probability spaces; let's see how to do that.