Probability Theory: Axiomatic Foundations

Mathematics Department

Lecture 1: The Sample Space and The Axiomatic Foundation

What is Probability?

- A mathematical framework for quantifying **uncertainty**.

What is Probability?

- A mathematical framework for quantifying uncertainty.
- Used in almost every field: physics, finance, computer science, biology, engineering, etc.

What is Probability?

- A mathematical framework for quantifying uncertainty.
- Used in almost every field: physics, finance, computer science, biology, engineering, etc.
- Today, we move from *intuitive* notions (like "chance" or "odds") to a rigorous mathematical model.

The Fundamental Concept: The Sample Space

Definition (Sample Space Ω)

The sample space, denoted by Ω (Omega), is the set of all possible outcomes of a random experiment.

Example (Coin Toss)

$\Omega = \{Heads, Tails\}$ or simply $\Omega = \{H, T\}$.

Example (Rolling a Die)

$\Omega = \{1, 2, 3, 4, 5, 6\}$.

More Sample Space Examples

Example (Two Coin Tosses)

What are all possible sequences?

More Sample Space Examples

Example (Two Coin Tosses)

What are all possible sequences? $\Omega = \{HH, HT, TH, TT\}$.

Example (Waiting for a Bus)

You measure the time (in minutes) until the next bus arrives.

More Sample Space Examples

Example (Two Coin Tosses)

What are all possible sequences? $\Omega = \{HH, HT, TH, TT\}$.

Example (Waiting for a Bus)

You measure the time (in minutes) until the next bus arrives.

$\Omega = \{x \in \mathbb{R} \mid x \geq 0\}$. This is an infinite sample space.

Types of Sample Spaces

Sample spaces can be:

- **Finite**: $\Omega = \{\omega_1, \omega_2, ..., \omega_n\}$ (Die roll, coin toss)
- **Countably Infinite**: Outcomes can be listed in a sequence (e.g., $\Omega = \{0, 1, 2, 3, ...\}$ for "number of emails received in an hour")
- **Uncountably Infinite**: Ω is a continuum (e.g., $\Omega = [0, 1] \subset \mathbb{R}$ for "randomly picking a point between 0 and 1")

The type of Ω influences the mathematical tools we need.

From Outcomes to Events

Definition (Event)

An event is a subset of the sample space Ω. We say event A *occurs* if the outcome ω of the experiment is an element of A.

Example (Die Roll)

Let $\Omega = \{1, 2, 3, 4, 5, 6\}$.

- $A = \{2, 4, 6\}$ is the event "roll an even number".
- $B = \{3\}$ is the event "roll a three" (a simple event).
- $C = \{1, 2, 3\}$ is the event "roll a number less than 4".

Set Theory = The Language of Probability

Probability theory is built on set theory.

- **Union** ($A \cup B$): "A or B occurs"
- **Intersection** ($A \cap B$): "A and B occur"
- **Complement** (A^c): "A does not occur"
- **Empty Set** ($\emptyset$): The impossible event (contains no outcomes)
- **Sample Space** (Ω): The certain event (contains all outcomes)

The Goal: Assigning probabilities

We want to assign a number $\mathbb{P}(A)$ to any event A, called the **probability of A**, that represents the "chance" it occurs.

- $\mathbb{P}(A)$ should be between 0 and 1.
- $\mathbb{P}(\Omega)$ should be 1 (something must happen!).
- For disjoint events A and B, $\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B)$.

But can we define $\mathbb{P}(A)$ for *every* subset A of Ω?

A Naive Approach: Discrete Uniform Probability

For a **finite** sample space with N **equally likely** outcomes, we can define:

$$\mathbb{P}(A) = \frac{\text{Number of outcomes in } A}{\text{Total number of outcomes } N} = \frac{|A|}{|\Omega|}$$

Example (Fair Die)

$\Omega = \{1, 2, 3, 4, 5, 6\}$, $|\Omega| = 6$.

- $\mathbb{P}(\text{even}) = \mathbb{P}(\{2, 4, 6\}) = \frac{3}{6} = \frac{1}{2}$.
- $\mathbb{P}(\text{roll a 3}) = \mathbb{P}(\{3\}) = \frac{1}{6}$.

This works perfectly for finite, symmetric experiments. But what about others?

The Need for a General Theory

The naive approach fails for:

- **Infinite Sample Spaces**: "Picking a random point in $[0, 1]$." There are infinitely many points, so the probability of any single point is... 0? But the event is not impossible.
- **Biased Coins**: Outcomes are not equally likely.

We need a more powerful and general system.

Kolmogorov's Axioms (1933)

Andrey Kolmogorov established the modern axiomatic foundation of probability.

The Axioms

A probability measure $\mathbb{P}$ is a function that assigns a number $\mathbb{P}(A)$ to an event A such that:

1. **Non-negativity**: $\mathbb{P}(A) \geq 0$ for any event A.
2. **Unit Measure**: $\mathbb{P}(\Omega) = 1$.
3. **Countable Additivity**: For any sequence of pairwise disjoint events $A_1, A_2, A_3, \ldots$ (i.e., $A_i \cap A_j = \emptyset$ for $i \neq j$),

$$\mathbb{P}\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mathbb{P}(A_i).$$

Immediate Consequences of the Axioms

From the three axioms, we can prove many natural rules:

Theorem

1. $\mathbb{P}(\emptyset) = 0$.
2. $\mathbb{P}(A^c) = 1 - \mathbb{P}(A)$.
3. *If* $A \subseteq B$, *then* $\mathbb{P}(A) \leq \mathbb{P}(B)$.
4. $\mathbb{P}(A \cup B) = \mathbb{P}(A) + \mathbb{P}(B) - \mathbb{P}(A \cap B)$.

These match our intuition! The power of the axioms is that these become theorems.

Example: Using the Axioms and Theorems

In a city, 60% of households have a dog, 45% have a cat, and 30% have both. What is the probability a randomly selected household has...

1. a dog or a cat?

Example: Using the Axioms and Theorems

In a city, 60% of households have a dog, 45% have a cat, and 30% have both. What is the probability a randomly selected household has...

1. a dog or a cat?
   $\mathbb{P}(D \cup C) = \mathbb{P}(D) + \mathbb{P}(C) - \mathbb{P}(D \cap C) = 0.6 + 0.45 - 0.3 = 0.75.$
2. neither?

Example: Using the Axioms and Theorems

In a city, 60% of households have a dog, 45% have a cat, and 30% have both. What is the probability a randomly selected household has...

1. a dog or a cat?
   $\mathbb{P}(D \cup C) = \mathbb{P}(D) + \mathbb{P}(C) - \mathbb{P}(D \cap C) = 0.6 + 0.45 - 0.3 = 0.75.$
2. neither? $\mathbb{P}(\text{neither}) = 1 - \mathbb{P}(D \cup C) = 1 - 0.75 = 0.25.$

Summary: Lecture 1

- **Sample Space** Ω: The set of all possible outcomes.
- **Event**: A subset of Ω.
- **Probability Measure** $\mathbb{P}$: A function satisfying Kolmogorov's three axioms:
  1. Non-negativity
  2. Unit Measure ($\mathbb{P}(\Omega) = 1$)
  3. Countable Additivity
- From these axioms, we can derive all other rules of probability (complement rule, addition rule, etc.).

Next Lecture: What is the domain of the function $\mathbb{P}$? Can we define it on *every* subset?

Lecture 2: σ-algebras and Measurable Spaces

Recap and a New Question

- **Sample Space** Ω: All possible outcomes.
- **Probability Measure** $\mathbb{P}$: Assigns a number to an *event* (a subset of Ω).
- **Axioms**: $\mathbb{P}(A) \geq 0$, $\mathbb{P}(\Omega) = 1$, Countable additivity for disjoint events.

Question: For a given Ω, on which collection of subsets (events) can we define $\mathbb{P}$? **All of them?**

The Problem: Unpleasant Subsets

For $\Omega = [0, 1]$ (e.g., "pick a point at random"), we'd like to define $\mathbb{P}$ for any interval. For example:

$$\mathbb{P}([a, b]) = b - a.$$

This seems natural. But can we define $\mathbb{P}(A)$ for **every possible subset** $A \subset [0, 1]$ and still satisfy the axioms?

The Problem: Unpleasant Subsets

For $\Omega = [0, 1]$ (e.g., "pick a point at random"), we'd like to define $\mathbb{P}$ for any interval. For example:

$$\mathbb{P}([a, b]) = b - a.$$

This seems natural. But can we define $\mathbb{P}(A)$ for **every possible subset** $A \subset [0, 1]$ and still satisfy the axioms?

Vitali's Theorem (1905)

No. It is **impossible** to define a probability measure $\mathbb{P}$ on *all* subsets of $[0, 1]$ such that $\mathbb{P}([a, b]) = b - a$.

We must restrict the domain of $\mathbb{P}$ to a "nice" collection of subsets.

The Solution: σ-algebras

We define $\mathbb{P}$ not on all subsets of Ω, but on a special collection of subsets called a $\sigma - algebra$.

Definition (σ-algebra)

A collection $\mathcal{F}$ of subsets of Ω is called a $\sigma - algebra$ if it satisfies:

1. Contains Sample Space: $\Omega \in \mathcal{F}$.
2. Closed under Complement: If $A \in \mathcal{F}$, then $A^c \in \mathcal{F}$.
3. Closed under Countable Unions: If $A_1, A_2, A_3, \ldots \in \mathcal{F}$, then $\bigcup_{i=1}^{\infty} A_i \in \mathcal{F}$.

The pair $(\Omega, \mathcal{F})$ is called a measurable space.

Why these three rules?

The axioms of a σ-algebra are perfectly matched to the axioms of probability.

- Axiom 1 ($\Omega \in \mathcal{F}$): We need to define $\mathbb{P}(\Omega) = 1$.
- Axiom 2 (Closed under complement): We need to define $\mathbb{P}(A^c)$ if we know $\mathbb{P}(A)$.
- Axiom 3 (Closed under countable unions): We need to define $\mathbb{P}(\bigcup A_i)$ for disjoint events (Axiom 3 of probability).

A σ-algebra is the domain of a probability measure. It is the collection of "measurable events".

Simple Examples of σ-algebras

Example (The Trivial σ-algebra)

For any Ω, the smallest possible σ-algebra is $\mathcal{F} = \{\emptyset, \Omega\}$.

Example (The Power Set)

For any Ω, the largest possible σ-algebra is the power set $\mathcal{P}(\Omega)$, the set of *all* subsets.

For finite or countably infinite Ω, we can usually use $\mathcal{F} = \mathcal{P}(\Omega)$ without issues. The problems arise with uncountable Ω (like $[0, 1]$).

A More Interesting Example

Let $\Omega = \{a, b, c, d\}$ (a finite sample space). Is this a σ-algebra?

$$\mathcal{F} = \{\emptyset, \{a, b\}, \{c, d\}, \Omega\}$$

A More Interesting Example

Let $\Omega = \{a, b, c, d\}$ (a finite sample space). Is this a σ-algebra?

$$\mathcal{F} = \{\emptyset, \{a, b\}, \{c, d\}, \Omega\}$$

Let's check the axioms:

1. $\Omega \in \mathcal{F}$? ✓
2. Closed under complement?
   - $\emptyset^c = \Omega \in \mathcal{F}$ ✓
   - $\{a, b\}^c = \{c, d\} \in \mathcal{F}$ ✓
   - $\{c, d\}^c = \{a, b\} \in \mathcal{F}$ ✓
   - $\Omega^c = \emptyset \in \mathcal{F}$ ✓
3. Closed under countable unions? (e.g., $\{a, b\} \cup \{c, d\} = \Omega \in \mathcal{F}$) ✓

Yes, $\mathcal{F}$ is a σ-algebra. It has "2 events" of size 2.

Generated σ-algebra

Often, we start with a collection of "simple" events we care about (e.g., all intervals in $[0, 1]$) and then generate the smallest σ-algebra that contains them.

Definition (Borel σ-algebra)

For $\Omega = \mathbb{R}$, the Borel σ-algebra $\mathcal{B}(\mathbb{R})$ is the σ-algebra generated by all open intervals (a, b).

- It contains all intervals (open, closed, half-open), single points, countable sets, and much more.
- It excludes the "pathological" sets that cause problems.
- This is the standard σ-algebra used for $\mathbb{R}$ and its subsets.

The Triple $(\Omega, \mathcal{F}, \mathbb{P})$

We now have the three components of a probability model:

Probability Space

A probability space is a triple $(\Omega, \mathcal{F}, \mathbb{P})$ consisting of:

1. A sample space Ω.
2. A σ-algebra $\mathcal{F}$ of subsets of Ω (the measurable events).
3. A probability measure $\mathbb{P}$ on $\mathcal{F}$ satisfying Kolmogorov's axioms.

This is the rigorous mathematical foundation for all of probability theory.

Visualizing a Probability Space

Think of it as a unit mass (total probability $= 1$) spread over the sample space Ω. The σ-algebra $\mathcal{F}$ tells us which chunks of this mass we are allowed to "measure". $\mathbb{P}(A)$ is the mass contained in the event A.

Example 1: Finite Space (Fair Die)

- $\Omega = \{1, 2, 3, 4, 5, 6\}$
- $\mathcal{F} = \mathcal{P}(\Omega)$ (the power set, all 64 subsets)
- For any $A \in \mathcal{F}$, $\mathbb{P}(A) = \frac{|A|}{6}$

Let's verify an axiom. Let $A = \{1\}$ and $B = \{2\}$. They are disjoint.

$$\mathbb{P}(A \cup B) = \mathbb{P}(\{1, 2\}) = \frac{2}{6} = \frac{1}{3}.$$

$$\mathbb{P}(A) + \mathbb{P}(B) = \frac{1}{6} + \frac{1}{6} = \frac{1}{3}. \quad \checkmark$$

$(\Omega, \mathcal{F}, \mathbb{P})$ is a valid probability space.

Example 2: Infinite Space (Uniform on $[0, 1]$)

- $\Omega = [0, 1]$
- $\mathcal{F} = \mathcal{B}([0, 1])$ (the Borel sets on $[0, 1]$)
- For an interval $I = [a, b] \subset [0, 1]$, we define $\mathbb{P}(I) = b - a$.

This definition is extended to all sets in $\mathcal{F}$ in a consistent way that satisfies the countable additivity axiom. We cannot define it for sets outside $\mathcal{F}$.

Why Bother with σ-algebras?

- **Mathematical Necessity**: They resolve the theoretical problem of defining measures on uncountable spaces. They are the "domain" of $\mathbb{P}$.

Why Bother with σ-algebras?

- **Mathematical Necessity**: They resolve the theoretical problem of defining measures on uncountable spaces. They are the "domain" of $\mathbb{P}$.
- **Modeling Information**: Different σ-algebras represent different levels of information. (We will see this in conditional expectation and martingales later).

Why Bother with σ-algebras?

- **Mathematical Necessity**: They resolve the theoretical problem of defining measures on uncountable spaces. They are the "domain" of $\mathbb{P}$.
- **Modeling Information**: Different σ-algebras represent different levels of information. (We will see this in conditional expectation and martingales later).
- **Foundational**: All advanced probability (stochastic processes, modern statistics) relies on this framework.

Summary: Lecture 2

- We **cannot** define a probability measure on *all* subsets of an uncountable Ω.
- The solution is to define $\mathbb{P}$ on a σ-algebra $\mathcal{F}$.
- A σ-algebra is a collection of subsets closed under complement and countable unions.
- The Borel σ-algebra on $\mathbb{R}$ is the most important example.
- A probability space is the triple $(\Omega, \mathcal{F}, \mathbb{P})$.

Next Lecture: How do we actually *construct* and *work with* probability measures?

Lecture 3: Constructing Probability Measures

Recap: The Probability Space

We model randomness with a triple $(\Omega, \mathcal{F}, \mathbb{P})$.

- Ω: Sample Space (what can happen?)
- $\mathcal{F}$: σ-algebra (what questions can we ask?)
- $\mathbb{P}$: Probability Measure (what are the answers?)

Today's Question: How do we actually **define** or **construct** the probability measure $\mathbb{P}$? We can't just list $\mathbb{P}(A)$ for every $A \in \mathcal{F}$.

The Strategy: Start Simple

We often know how to assign probability to "simple" events.

- **Die Roll**: For each simple event $\{i\}$, $\mathbb{P}(\{i\}) = 1/6$.
- **Uniform on** $[0, 1]$: For an interval $[a, b]$, we want $\mathbb{P}([a, b]) = b - a$.

The collection of these "simple" events (e.g., all intervals) is often not a σ-algebra itself. We call it a generator.

From the Generator to the σ-algebra

We have:

1. A collection of "simple" sets $\mathcal{C}$ (the generator) where we know $\mathbb{P}$.
2. The full σ-algebra $\mathcal{F} = \sigma(\mathcal{C})$ generated by $\mathcal{C}$.

Key Problem: Does our definition of $\mathbb{P}$ on the small collection $\mathcal{C}$ extend **uniquely** to the large σ-algebra $\mathcal{F}$?

Carathéodory's Extension Theorem

Yes, under certain conditions. If $\mathbb{P}$ is well-behaved on $\mathcal{C}$ (e.g., it is countably additive on it), then it extends uniquely to a probability measure on $\sigma(\mathcal{C})$.

This theorem is the workhorse behind constructing probability measures.

Construction 1: Discrete Probability Spaces

This is the simplest case. Let Ω be finite or countably infinite: $\Omega = \{\omega_1, \omega_2, \omega_3, \dots\}$.

1. Define a probability mass function (pmf) $p : \Omega \to [0, 1]$ such that $\sum_i p(\omega_i) = 1$.
2. For any subset $A \in \mathcal{F} = \mathcal{P}(\Omega)$, define:

$$\mathbb{P}(A) = \sum_{\omega_i \in A} p(\omega_i)$$

This construction *automatically* satisfies Kolmogorov's axioms (check countable additivity!).

Discrete Example: Poisson Distribution

Model: "Number of emails received in an hour."

- $\Omega = \{0, 1, 2, 3, \ldots\}$
- $\mathcal{F} = \mathcal{P}(\Omega)$ (all subsets)
- For a parameter $\lambda > 0$, the pmf is $p(k) = e^{-\lambda} \frac{\lambda^k}{k!}$ for $k = 0, 1, 2, \ldots$
- For any event A, e.g., $A = \{\text{even numbers}\} = \{0, 2, 4, 6, \ldots\}$,

$$\mathbb{P}(A) = \sum_{k \text{ even}} e^{-\lambda} \frac{\lambda^k}{k!} = e^{-\lambda} \cosh(\lambda).$$

The axioms are satisfied because we've defined $\mathbb{P}$ by a summable series.

Construction 2: Continuous Probability Spaces

Let $\Omega = \mathbb{R}$ (or a subset). We use probability density functions (pdfs).

1. Find a function $f : \mathbb{R} \to [0, \infty)$ such that $\int_{-\infty}^{\infty} f(x)\,dx = 1$.
2. For any event A in the Borel σ-algebra (e.g., an interval, a union of intervals), define:

$$\mathbb{P}(A) = \int_A f(x)\,dx$$

This construction also satisfies the axioms (e.g., countable additivity comes from properties of the Lebesgue integral).

Continuous Example: Uniform Distribution

Model: "Pick a random point between 1 and 5."

- $\Omega = [1, 5]$
- $\mathcal{F} = \mathcal{B}([1, 5])$
- The pdf is:

$$f(x) = \begin{cases} \frac{1}{4} & \text{for } 1 \leq x \leq 5 \\ 0 & \text{otherwise} \end{cases}$$

- What is $\mathbb{P}([2, 3])$? $\mathbb{P}([2, 3]) = \int_2^3 \frac{1}{4}\, dx = \frac{1}{4}(3 - 2) = \frac{1}{4}$.
- What is $\mathbb{P}(\{2.5\})$? $\mathbb{P}(\{2.5\}) = \int_{2.5}^{2.5} f(x) dx = 0$.

Probability zero is not the same as impossible!

A Mix of Both: The CDF

There is a unified way to describe *any* probability measure on $\mathbb{R}$.

Definition (Cumulative Distribution Function (CDF))

The CDF of a probability measure $\mathbb{P}$ is the function $F : \mathbb{R} \to [0, 1]$ defined by:

$$F(x) = \mathbb{P}((-\infty, x])$$

Theorem

A function F is a CDF for some $\mathbb{P}$ if and only if:

1. *F is non-decreasing.*
2. *F is right-continuous.*
3. $\lim_{x\to-\infty} F(x) = 0$ *and* $\lim_{x\to\infty} F(x) = 1$.

CDF Examples

Example (Fair Die (Discrete))

$F(x)$ is a step function.

CDF Examples

Example (Uniform on $[1, 5)$)

$(Continuous)]$F(x)$isapiecewiselinear function.$

Putting It All Together: A Complete Example

Experiment: Toss a fair coin. If Heads, roll a 4-sided die. If Tails, roll a 6-sided die. What is $\mathbb{P}(\text{roll a 3})$?

Putting It All Together: A Complete Example

Experiment: Toss a fair coin. If Heads, roll a 4-sided die. If Tails, roll a 6-sided die. What is $\mathbb{P}(\text{roll a 3})$?

1. **Construct** Ω: $\Omega = \{H1, H2, H3, H4, T1, T2, T3, T4, T5, T6\}$.

Putting It All Together: A Complete Example

Experiment: Toss a fair coin. If Heads, roll a 4-sided die. If Tails, roll a 6-sided die. What is $\mathbb{P}(\text{roll a } 3)$?

1. **Construct** Ω: $\Omega = \{H1, H2, H3, H4, T1, T2, T3, T4, T5, T6\}$.
2. **Construct** $\mathcal{F}$: $\mathcal{F} = \mathcal{P}(\Omega)$.

Putting It All Together: A Complete Example

Experiment: Toss a fair coin. If Heads, roll a 4-sided die. If Tails, roll a 6-sided die. What is $\mathbb{P}$(roll a 3)?

1. **Construct Ω**: $\Omega = \{H1, H2, H3, H4, T1, T2, T3, T4, T5, T6\}$.
2. **Construct $\mathcal{F}$**: $\mathcal{F} = \mathcal{P}(\Omega)$.
3. **Assign probabilities**: The outcomes are **not** equally likely!
   - $\mathbb{P}(\{H_i\}) = \frac{1}{2} \cdot \frac{1}{4} = \frac{1}{8}$ for $i = 1, 2, 3, 4$.
   - $\mathbb{P}(\{T_j\}) = \frac{1}{2} \cdot \frac{1}{6} = \frac{1}{12}$ for $j = 1, ..., 6$.

Putting It All Together: A Complete Example

Experiment: Toss a fair coin. If Heads, roll a 4-sided die. If Tails, roll a 6-sided die. What is $\mathbb{P}(\text{roll a 3})$?

1. **Construct Ω**: $\Omega = \{H1, H2, H3, H4, T1, T2, T3, T4, T5, T6\}$.
2. **Construct $\mathcal{F}$**: $\mathcal{F} = \mathcal{P}(\Omega)$.
3. **Assign probabilities**: The outcomes are **not** equally likely!
   - $\mathbb{P}(\{H_i\}) = \frac{1}{2} \cdot \frac{1}{4} = \frac{1}{8}$ for $i = 1, 2, 3, 4$.
   - $\mathbb{P}(\{T_j\}) = \frac{1}{2} \cdot \frac{1}{6} = \frac{1}{12}$ for $j = 1, ..., 6$.
4. **Define $\mathbb{P}$ for any event**: For $A = \{\text{roll a 3}\} = \{H3, T3\}$,

$$\mathbb{P}(A) = \mathbb{P}(\{H3\}) + \mathbb{P}(\{T3\}) = \frac{1}{8} + \frac{1}{12} = \frac{5}{24}.$$

Summary of the Three Lectures

- **Lecture 1**: We want to assign numbers to events. Kolmogorov's Axioms give the rules. ($\mathbb{P}(A) \geq 0, \mathbb{P}(\Omega) = 1$, countable additivity).
- **Lecture 2**: We can't define $\mathbb{P}$ on all subsets. We restrict its domain to a "σ-algebra" $\mathcal{F}$ of measurable events. This gives us a probability space $(\Omega, \mathcal{F}, \mathbb{P})$.
- **Lecture 3**: We construct $\mathbb{P}$ by first defining it on a simple collection of sets (using pmfs, pdfs, or direct assignment) and then using extension theorems to define it on the entire σ-algebra.

Where to go from here?

You now have the rigorous foundation to explore:

- **Random Variables**: Functions $X : \Omega \to \mathbb{R}$ that translate outcomes into numbers.
- **Distribution**: The probability measure induced by X on $\mathbb{R}$.
- **Expectation**: The "average" value of a random variable.
- **Limit Theorems**: The Law of Large Numbers and the Central Limit Theorem.

Congratulations! You have built the bedrock of modern probability theory.