

Unbiasedness of the Sample Mean: A Structured Proof Note

Basis

March 6, 2026

Abstract

This note gives a fully rigorous proof that the sample mean is an unbiased estimator of the population mean for a finite sample of real-valued random variables under a single essential assumption: integrability (finite expectation). In the i.i.d. setting, we show that if X_1, \dots, X_n are integrable with common mean $\mu = \mathbb{E}[X_1] \in \mathbb{R}$, then the sample mean $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$ satisfies $\mathbb{E}[\bar{X}] = \mu$ by linearity of expectation. We also isolate the minimal condition needed for this identity—namely, that each X_i is integrable and shares the same mean—highlighting that independence plays no role in unbiasedness for fixed n . Finally, we emphasize a concrete way in which “unbiasedness” can fail to be a meaningful statement: for heavy-tailed laws such as the standard Cauchy, the expectations $\mathbb{E}[X_i]$ and $\mathbb{E}[\bar{X}]$ are undefined. As a partial remedy in such settings, we briefly discuss truncation and bounded odd transforms, which restore well-defined expectations and yield unbiasedness for the transformed targets under symmetry.

1 Introduction

The sample mean is among the most fundamental estimators in probability and statistics. A basic and widely used claim is that, when data are sampled from a distribution with a well-defined (finite) mean, the sample mean is an unbiased estimator of that population mean. While the underlying calculation is short, a fully rigorous statement requires spelling out the precise assumption under which the expectation of the sample mean exists and the linearity step is justified.

The main purpose of this paper is to present a clean, measure-theoretic proof of unbiasedness for a finite sample X_1, \dots, X_n of real-valued random variables. The key assumption is integrability: $\mathbb{E}[X_1] \in \mathbb{R}$. Under this condition the sample mean \bar{X} is itself integrable, and linearity of expectation yields $\mathbb{E}[\bar{X}] = \mathbb{E}[X_1]$. Although the classical formulation often includes independence and identical distribution, the argument for unbiasedness at a fixed sample size does not use independence; what matters is that the relevant expectations exist and that the X_i share a common mean.

Beyond establishing the i.i.d. result, we clarify two points that are sometimes glossed over. First, we record a minimal hypothesis guaranteeing unbiasedness: if each X_i is integrable and satisfies $\mathbb{E}[X_i] = \mu$, then $\mathbb{E}[\bar{X}] = \mu$ regardless of dependence. Second, we note that unbiasedness is not merely false but can be ill-posed when expectations do not exist as finite real numbers. A concrete example is provided by the standard Cauchy distribution, for which $\mathbb{E}[X_i]$ and $\mathbb{E}[\bar{X}]$ are undefined; in that case the statement “ \bar{X} is unbiased for μ ” has no mathematical content.

Organization. Section 1 states the setup and the integrability assumption. Section 2 proves the main unbiasedness theorem for i.i.d. integrable variables. Section 3 formulates a minimal condition for unbiasedness that drops independence. Section 4 gives an explicit failure mode in which expectations are undefined (the Cauchy case). Section 5 briefly discusses a partial remedy for heavy-tailed distributions via truncation and, more generally, bounded odd transforms that retain well-defined expectations.

2 Setup and Assumptions

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let X_1, \dots, X_n be real-valued random variables.

Assumption (i.i.d. with finite mean). Assume X_1, \dots, X_n are independent and identically distributed, and integrable:

$$\mu := \mathbb{E}[X_1] \in \mathbb{R}.$$

Define the sample mean

$$\bar{X} := \frac{1}{n} \sum_{i=1}^n X_i.$$

Under this assumption, each X_i is integrable with $\mathbb{E}[X_i] = \mu$, and any finite linear combination of the X_i (in particular \bar{X}) is integrable.

3 Main Result

Theorem 3.1 (Unbiasedness of the sample mean). *Under Assumption 2, the sample mean \bar{X} is an unbiased estimator of μ , i.e.,*

$$\mathbb{E}[\bar{X}] = \mu.$$

Proof. By Assumption 2, each X_i is integrable and $\mathbb{E}[X_i] = \mu$. Since finite sums of integrable random variables are integrable, $\sum_{i=1}^n X_i$ and \bar{X} are integrable. By linearity of expectation for integrable random variables,

$$\mathbb{E}[\bar{X}] = \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n X_i\right] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[X_i] = \frac{1}{n} \sum_{i=1}^n \mu = \mu.$$

□

Remark (independence is not needed for unbiasedness). The conclusion $\mathbb{E}[\bar{X}] = \mu$ uses only integrability and the identities $\mathbb{E}[X_i] = \mu$ for all i , together with linearity of expectation; no independence assumption is required for this equality.

4 Beyond i.i.d.: a minimal condition for unbiasedness

This section records a minimal hypothesis under which the same unbiasedness identity holds, emphasizing which parts of Assumption 2 are actually used in the proof of Theorem 3.1.

Proposition 4.1 (Unbiasedness under identical means). *Let X_1, \dots, X_n be integrable random variables (not necessarily independent) such that $\mathbb{E}[X_i] = \mu$ for each i . Then $\mathbb{E}[\bar{X}] = \mu$.*

Proof. Since each X_i is integrable, \bar{X} is integrable. Linearity of expectation gives

$$\mathbb{E}[\bar{X}] = \mathbb{E}\left[\frac{1}{n} \sum_{i=1}^n X_i\right] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[X_i] = \frac{1}{n} \sum_{i=1}^n \mu = \mu.$$

□

Remark (existence of $\mathbb{E}[\bar{X}]$ is a separate issue). In non-i.i.d. settings it can happen that some individual expectations exist while $\mathbb{E}[\bar{X}]$ does not. For example, with $n = 2$, let $X_1 \equiv 0$ (integrable) and let X_2 be standard Cauchy (not integrable). Then $\mathbb{E}[X_1] = 0$ exists, but

$$\bar{X} = \frac{X_1 + X_2}{2} = \frac{X_2}{2}$$

is again Cauchy-distributed, hence $\mathbb{E}[\bar{X}]$ is undefined. A simple sufficient condition ensuring that $\mathbb{E}[\bar{X}]$ exists and that linearity applies is: each X_i is integrable (as assumed in Proposition 4.1). If one studies a sequence $(\bar{X}_n)_{n \geq 1}$ as n varies, additional conditions (e.g. uniform integrability) are often imposed to control existence and limit-interchange issues, but for a fixed finite n integrability of the summands suffices.

5 A Concrete Case Where Unbiasedness Fails (Expectation Undefined)

Unbiasedness is only defined when the relevant expectations exist as finite real numbers. A standard concrete failure occurs when $\mathbb{E}[X_1]$ does not exist.

Let X_1, \dots, X_n be i.i.d. standard Cauchy random variables. Recall that for a real-valued random variable X one may write $X = X^+ - X^-$ where $X^+ := \max\{X, 0\}$ and $X^- := \max\{-X, 0\}$, and $\mathbb{E}[X]$ is defined (as a finite number) only if $\mathbb{E}[X^+] < \infty$ and $\mathbb{E}[X^-] < \infty$, in which case $\mathbb{E}[X] = \mathbb{E}[X^+] - \mathbb{E}[X^-]$.

For the standard Cauchy distribution one has $\mathbb{E}[X_1^+] = \infty$ and $\mathbb{E}[X_1^-] = \infty$, hence $\mathbb{E}[X_1]$ is undefined (does not exist as a finite real number). Consequently, the parameter μ in Assumption 2 cannot be formed. Moreover, the sample mean \bar{X} is also Cauchy-distributed (stability of the Cauchy law), so $\mathbb{E}[\bar{X}]$ is likewise undefined. In this setting, the statement “ \bar{X} is unbiased for μ ” is not meaningful because neither side is a well-defined finite expectation.

6 A partial remedy for heavy tails: truncation and bounded transforms

When $\mathbb{E}[X_i]$ does not exist, a common workaround is to replace X_i by a bounded (hence integrable) transform. One simple choice is truncation at level $M > 0$:

$$X_i^{(M)} := X_i \mathbf{1}\{|X_i| \leq M\}, \quad \bar{X}^{(M)} := \frac{1}{n} \sum_{i=1}^n X_i^{(M)}.$$

Then $|X_i^{(M)}| \leq M$, so $X_i^{(M)}$ and $\bar{X}^{(M)}$ are integrable for every $M > 0$, regardless of whether X_i is integrable.

Proposition 6.1 (Truncated mean for symmetric heavy tails). *Assume X_1, \dots, X_n are identically distributed and symmetric about 0 (i.e. $X_1 \stackrel{d}{=} -X_1$). Then for every $M > 0$,*

$$\mathbb{E}[\bar{X}^{(M)}] = 0.$$

In particular, if X_1, \dots, X_n are i.i.d. standard Cauchy, then $\mathbb{E}[\bar{X}^{(M)}] = 0$ for every $M > 0$ even though $\mathbb{E}[\bar{X}]$ is undefined.

Proof. For fixed $M > 0$, the map $g_M(x) := x \mathbf{1}\{|x| \leq M\}$ is bounded and odd: $g_M(-x) = -g_M(x)$. If $X_1 \stackrel{d}{=} -X_1$, then $\mathbb{E}[g_M(X_1)]$ exists (boundedness) and

$$\mathbb{E}[g_M(X_1)] = \mathbb{E}[g_M(-X_1)] = \mathbb{E}[-g_M(X_1)] = -\mathbb{E}[g_M(X_1)],$$

so $\mathbb{E}[g_M(X_1)] = 0$. By linearity of expectation for integrable random variables,

$$\mathbb{E}[\bar{X}^{(M)}] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}[X_i^{(M)}] = \mathbb{E}[X_1^{(M)}] = 0.$$

□

Remark (a generalized “unbiasedness” viewpoint). More generally, for any bounded odd function ψ and any X symmetric about 0, one has $\mathbb{E}[\psi(X)] = 0$ (the expectation exists by boundedness). This suggests replacing the undefined condition $\mathbb{E}[X - \theta] = 0$ by the well-defined estimating equation $\mathbb{E}[\psi(X - \theta)] = 0$ (with a chosen bounded odd ψ), a standard robust alternative in heavy-tailed settings.

7 Conclusion and Outlook

Under the integrability assumption $\mathbb{E}[X_1] \in \mathbb{R}$, we proved that the sample mean $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$ is an unbiased estimator of the population mean in the i.i.d. setting, i.e., $\mathbb{E}[\bar{X}] = \mu$. The proof is entirely driven by two facts: finite sums of integrable random variables are integrable, and expectation is linear on integrable random variables. We also isolated the minimal requirement for unbiasedness at fixed n : it suffices that each X_i be integrable with the common mean μ ; independence is not needed for the identity itself.

We further emphasized a concrete sense in which unbiasedness can “fail”: if the underlying mean does not exist as a finite real number, then unbiasedness is not a well-posed property. For i.i.d. standard Cauchy variables, neither $\mathbb{E}[X_i]$ nor $\mathbb{E}[\bar{X}]$ exists, so there is no meaningful target mean μ for \bar{X} to estimate without bias.

As an outlook, for heavy-tailed settings where integrability is doubtful or false, a practical next step is to replace X_i by an integrable transform, such as truncation $X_i^{(M)} = X_i \mathbf{1}\{|X_i| \leq M\}$ or, more generally, a bounded odd score function $\psi(X_i)$. These choices guarantee existence of expectations and can support well-defined estimating equations even when $\mathbb{E}[X_i]$ is undefined, at the cost of changing the estimand and introducing a tuning parameter or modeling choice that must be handled explicitly.