




Figure 3.12

3.4 The cumulative distribution function

Thus far we've studied discrete random variables: random quantities X where the set of possible values is finite or countably infinite. In order to specify the distribution of a discrete random variable, we have two options:

- **(pmf)** for each value $x \in \text{Range}(X)$, specify its individual probability $P(X = x)$. If the range is small enough, these can simply be listed out in a table. Otherwise, this information must be summarized in a *function* called the **probability mass function**:

$$p_X(x) = \begin{cases} P(X = x) & x \in \text{Range}(X) \\ 0 & \text{else.} \end{cases}$$

- **(cdf)** define a function $F_X : \mathbb{R} \rightarrow [0, 1]$ that returns $F_X(x) = P(X \leq x)$ for any $x \in \mathbb{R}$.

Figure 3.12 visualizes these options, and as we saw in Section 3.2, these are fundamentally two sides of the same coin: given one, you can always recover the other, either summing up the pmf to get the cdf, or differencing the cdf to get the pmf. Having said that, we have yet to define a random variable primarily in terms of its cdf. We have worked exclusively with the pmf, and only occasionally visualized the cdf as a curious byproduct that did not see much action. But what was once a bit player will now become the star as we turn our attention to other types of random variables beyond the discrete case.

In order to fully specify any random variable, regardless its type, we must ultimately write down two objects:

- **(range)** the set of all possible values $\text{Range}(X) \subseteq \mathbb{R}$;
- **(distribution)** a probability measure $A \mapsto P(X \in A)$.





Figure 3.13: Valid cdfs come in all shapes and sizes

Figure 3.14: Types of continuity for a function $f : \mathbb{R} \rightarrow \mathbb{R}$

The question then becomes: *how*? Merely as a technical matter, how do we pull this off? Unlike in a first calculus course where we can write down functions $f(x) = x^2$ that take real numbers and return real numbers, it is not straightforward to write down a simple “formula” that summarizes how *any* set $A \subseteq \mathbb{R}$ gets assigned its probability $P(X \in A)$. And even if it were, how do we ensure that the resulting map is a valid probability measure that satisfies all of our axioms? It turns out that the universal answer to the question “how do we write down the distribution of *any* random variable” is “write down its cdf.” All random variables, no matter how pathological, possess a cdf, and the cdf is sufficient for computing the entire distribution, which we will now see.

Recall the four properties of a cumulative distribution function:

- **nondecreasing**: for any $x \leq y$, it must be that $F_X(x) \leq F_X(y)$;
- **right continuous**: F_X has all its limits as you approach *from the right*:

$$\lim_{x \rightarrow a^+} F_X(x) = F_X(a).$$

This is visualized in Figure 3.14.

- **goes to zero**: $\lim_{x \rightarrow -\infty} F_X(x) = 0$;
- **goes to one**: $\lim_{x \rightarrow \infty} F_X(x) = 1$.

Every cdf you will ever encounter for the rest of your life satisfies these four properties...*at least*. Figure 3.13 displays cartoons of what a valid cdf might look like. How is the cdf sufficient for computing the entire distribution of a random variable? That is, given *any* set $A \subseteq \mathbb{R}$, how does the cdf allow us to compute the number $P(X \in A)$? To show this, we start with intervals of the form $A = (a, b]$ for $a < b$:

Theorem 3.7. Let X be any random variable with cdf $F_X(x) = P(X \leq x)$, and let $A = (a, b]$ for any real numbers $a < b$. Then

$$P(X \in A) = P(a < X \leq b) = F_X(b) - F_X(a). \quad (3.15)$$

Proof. Let hilarity ensue:

$$\begin{aligned} P(X \in A) &= P(a < X \leq b) \\ &= P(a < X \text{ and } X \leq b) \\ &= P(a < X) + P(X \leq b) - P(a < X \text{ or } X \leq b) && \text{(inclusion/exclusion)} \\ &= P(a < X) + P(X \leq b) - P(X \in \mathbb{R}) \\ &= 1 - P(X \leq a) + P(X \leq b) - P(X \in \mathbb{R}) && \text{(complement rule)} \\ &= 1 - P(X \leq a) + P(X \leq b) - 1 && \text{(axiom: total measure 1)} \\ &= P(X \leq b) - P(X \leq a) \\ &= F_X(b) - F_X(a). \end{aligned}$$

□

So, computing the distribution of X for a generic interval $A = (a, b]$ can always be collapsed to a cdf calculation. But what about sets of other forms: (a, b) , $(0, \infty)$, $\{c\}$, etc? We will not offer a formal proof, but any “ordinary-looking” subsets of $\mathbb{R}$ like these can ultimately be expressed as limits, unions, intersections, or complements of intervals like $(a, b]$, and so the overall distribution $P(X \in \mathbb{R})$ for any set A can be computed using Theorem 3.7 together with our basic rules from Section 2.3.

In this way then, we see that “the cdf is queen.” It is a universal object. All random variables have one. It fully characterizes the distribution of the random variable. The cdf is sufficient in principle for computing $P(X \in A)$ for any set A . But if it’s so important, why did we hardly ever mention it in Section 3.2? While the cdf is crucially important theoretically, its practical use can be limited. In practice, we would like to have *alternative* characterizations of a distribution that are easier to work with and more interpretable. The precise nature of that alternative characterization depends on the *type* of the random variable. As we saw, when the random variable is of discrete type for instance, we can characterize the distribution with the pmf, and we got a lot of mileage out of this option.

So what other types of random variables are there? We know that the cdf is queen, and it possesses its four universal properties: nondecreasing, *right* continuous, goes to 0 as $x \rightarrow -\infty$, goes to 1 as $x \rightarrow \infty$. Those apply to all random variables, period end of story I don’t care what. As such, we can start to classify random variables according to any *extra* properties their cdf may possess. For example...




Figure 3.15: Cartoons of continuous cdfs

3.5 Continuous random variables

Definition 3.12. A random variable X is **continuous** if its cdf $F_X(x) = P(X \leq x)$ is a continuous function for all $x \in \mathbb{R}$.

So stronger than being merely right-continuous, the cdf is now just plain *continuous*, meaning that the value of F_X at a point a is equal to its limit at that point regardless which direction we approach from:

$$\lim_{x \rightarrow a} F_X(x) = F_X(a).$$

Figure 3.15 displays cartoons of continuous cdfs, which do not have any holes or jumps. One of the consequences of this is that a continuous random variable will now have a range which is an *uncountably* infinite set, like $(0, 1)$, $[0, \infty)$, or all of $\mathbb{R}$. But because of this, we must learn to accept a jarringly uninuitive property:

Theorem 3.8. If X is a continuous random variable, then $P(X = a) = 0$ for any $a \in \mathbb{R}$.

Proof. It suffices to show $P(X \in (-\infty, a)) = P(X \in (-\infty, a])$, since the difference between these two probabilities is precisely $P(X = a)$ (why? check!). Now, take any $b < a$. Since $(-\infty, b] \subset (-\infty, a) \subset (-\infty, a]$, monotonicity of probability gives

$$F(b) = P(X \in (-\infty, b]) \leq P(X \in (-\infty, a)) \leq P(X \in (-\infty, a]) = F(a).$$

Since this is true for any $b < a$, we can take the limit as b approaches a from the left to get that

$$F(a) = \lim_{b \uparrow a} F(b) \leq P(X \in (-\infty, a)) \leq F(a).$$

Therefore, $P(X \in (-\infty, a)) = F(a) = P(X \in (-\infty, a])$. □

Think about a random variable X with $\text{Range}(X) = \mathbb{R}$. There are simply “too many” real numbers for us to assign non-zero probability to each one and have these “add up” to one. So if we wish to place a valid probability measure on all of $\mathbb{R}$, we have to accept Theorem 3.8.

One quick, practical consequence of Theorem 3.8 is that it does not matter whether or not you include the endpoints when you compute the probability that a continuous random variable is in some interval:

Corollary 3.9. If X is a continuous random variable, then

$$P(a < X \leq b) = P(a \leq X < b) = P(a \leq X \leq b) = P(a < X < b).$$




Figure 3.16: The Cantor distribution is an example of a *singular* distribution. This cdf is continuous everywhere, but it is not *absolutely* continuous, and the derivative is equal to zero (almost) everywhere, so it does not have a probability density function.

Partial proof. We won't itemize all the cases, but as an example, consider $[a, b]$. This can be rewritten as a disjoint union: $[a, b] = \{a\} \cup (a, b) \cup \{b\}$. So by countable additivity $P(a \leq X \leq b) = P(X = a) + P(a < X < b) + P(X = b)$. We know from Theorem 3.8 that $P(X = a) = P(X = b) = 0$, so $P(a \leq X \leq b) = P(a < X < b)$. $\square$

Continuity is a nice property to require because it enforces a certain amount of smoothness, but it is ultimately a fairly weak condition. It does not rule out functions like the Cantor distribution displayed in Figure 3.16, which is just plain weird. In practice, when working with continuous RVs, we require an even stronger form of smoothness...




Figure 3.17: Cartoons of what a probability density function might look like. Area under the curve corresponds to probabilities. The total area underneath the pdf must be one.

3.6 Absolutely continuous random variables

Definition 3.13. A random variable X is **absolutely continuous** if its cdf $F_X(x) = P(X \leq x)$ is an absolutely continuous function for all $x \in \mathbb{R}$.

The concept of absolute continuity is beyond the scope of our course. You can look it up if you wish, but all I want you to know is that absolute continuity is a special kind of continuity that is stronger, and it ensures that the cdf is sufficiently smooth as to rule out the pathological cases

By ruling out such cases and requiring the cdf to be extra smooth, we guarantee the existence of a special object, which is the continuous analog of the pmf

Theorem 3.10. If X is an absolutely continuous random variable, then there exists a function $f_X : \mathbb{R} \rightarrow \mathbb{R}$ called the **probability density function (pdf)** with the following properties:

- $f_X(x) \geq 0$ for all $x \in \mathbb{R}$;
- $F'_X(x) = f_X(x)$;
- $F_X(x) = \int_{-\infty}^x f_X(t) dt$;
- $P(a < X < b) = \int_a^b f_X(x) dx = F_X(b) - F_X(a)$ (fundamental theorem of calculus!);
- $\int_{-\infty}^{\infty} f_X(x) dx = P(-\infty < X < \infty) = 1$;




Figure 3.18: $P(a < X < b) = \int_a^b f_X(x) dx = F_X(b) - F_X(a)$

Figure 3.17 displays cartoons of what this function might look like, and Figure 3.18 illustrates how probabilities are calculated.

Remark 3.9. Absolutely continuous random variables are the continuous analog to discrete random variables, and the pdf plays the role of the pmf. In discrete world, we *summed* the pmf to compute probabilities. In continuous world, we *integrate* the pdf. In discrete world, we *differenced* the cdf to compute the pmf. In continuous world, we *differentiate* it.

Remark 3.10. The relationship $F_X(x) = \int_{-\infty}^x f_X(t) dt$ means that the cdf is the **area accumulation function** of the pdf. This is illustrated in Figure 3.19.

Remark 3.11. I hasten to emphasize: not all continuous random variables have a pdf. Random variables whose cdf is continuous but not absolutely continuous are called **singular**, and the Cantor distribution in Figure 3.16 is the most famous example of this. The derivative of the Cantor distribution is zero (almost) everywhere, so a density with the above properties could not possibly exist. When the pdf does exist however, it's super useful, because it allows us to compute probabilities and moments by *integrating*.

Example 3.14. Let X be absolutely continuous with density

$$f_X(x) = \begin{cases} c(4x - 2x^2) & 0 < x < 2 \\ 0 & \text{else.} \end{cases}$$




Figure 3.19: the cdf F is tracking the amount of area (probability) under the pdf f that we accumulate as we move from left to right. When we start on the left with “ $x = -\infty$,” we’ve accumulated none of it, and as $x \rightarrow \infty$, we eventually accumulate all of it (total measure 1).




Figure 3.20

A cartoon is displayed in Figure 3.20. The range of X is the set of real numbers where the density is nonzero. This is called the **support** of the density function:

$$\text{supp}(f_X) = \{x \in \mathbb{R} : f_X(x) > 0\}, \quad (3.16)$$

and so $\text{Range}(X) = \text{supp}(f_X)$ for absolutely continuous random variables, and in this case we see that $\text{Range}(X) = (0, 2)$. And recall that because X is continuous, we need not quibble about whether or not we include the endpoints. The constant c is a **normalizing constant** that serves to ensure that the density integrates to 1, so let's find its value. First, note that

$$\begin{aligned} \int_{-\infty}^{\infty} f_X(x) dx &= \int_{-\infty}^0 f_X(x) dx + \int_0^2 f_X(x) dx + \int_2^{\infty} f_X(x) dx \\ &= \cancel{\int_{-\infty}^0 0 dx} + \int_0^2 c(4x - 2x^2) dx + \cancel{\int_2^{\infty} 0 dx} \\ &= c \int_0^2 (4x - 2x^2) dx \end{aligned}$$

So

$$\begin{aligned} \int_{-\infty}^{\infty} f_X(x) dx = 1 &\implies c \int_0^2 (4x - 2x^2) dx = 1 \\ &\implies c = \frac{1}{\int_0^2 (4x - 2x^2) dx}. \end{aligned}$$

As such:

$$\int_0^2 (4x - 2x^2) dx = \left[2x^2 - \frac{2}{3}x^3 \right]_0^2 = 2 \cdot 4 - \frac{2}{3} \cdot 8 = 8/3.$$

So $c = 1/(8/3) = 3/8$. If we wanted to compute $P(X > 1)$, we can do that by integrating the pdf:

$$\begin{aligned} P(X > 1) &= P(X \in (1, \infty)) \\ &= \int_1^{\infty} f_X(x) dx \\ &= \int_1^2 f_X(x) dx + \int_2^{\infty} f_X(x) dx \\ &= \int_1^2 \frac{3}{8}(4x - 2x^2) dx + \cancel{\int_2^{\infty} 0 dx} \\ &= \frac{3}{8} \left[2x^2 - \frac{2}{3}x^3 \right]_1^2 \\ &= \frac{3}{8} \left[2 \cdot 4 - \frac{2}{3} \cdot 8 - 2 + \frac{2}{3} \right] \\ &= \frac{1}{2}. \end{aligned}$$

Definition 3.14. If X is absolutely continuous with density f , then the expected value of X is

$$\mathbb{E}[X] = \int_{-\infty}^{\infty} x f(x) dx. \quad (3.17)$$

As in the discrete case, the expected value may not exist or be finite.

Example 3.15. Continuing Example 3.14, We see that

$$\begin{aligned} \mathbb{E}[X] &= \int_{-\infty}^{\infty} x f_X(x) dx \\ &= \int_{-\infty}^0 x f_X(x) dx + \int_0^2 x f_X(x) dx + \int_2^{\infty} x f_X(x) dx \\ &= \int_{-\infty}^0 x \cdot 0 dx + \int_0^2 x \frac{3}{8} (4x - 2x^2) dx + \int_2^{\infty} x \cdot 0 dx \\ &= \frac{3}{8} \int_0^2 (4x^2 - 2x^3) dx \\ &= \frac{3}{8} \left[\frac{4}{3} x^3 - \frac{2}{4} x^4 \right]_0^2 \\ &= \frac{3}{8} \left[\frac{4}{3} 8 - \frac{2}{4} 16 \right] \\ &= 1. \end{aligned}$$


It makes sense that this is the mean because the pdf is symmetric about 1.

Theorem 3.11. (LOTUS, again) If X is absolutely continuous with density f_X and $g : \text{Range}(X) \rightarrow \mathbb{R}$ is a transformation, then

$$E[g(X)] = \int_{-\infty}^{\infty} g(x) f_X(x) dx. \quad (3.18)$$

The upshot of this result is the same as before: in order to compute $E[g(X)]$, it is not necessary to first derive the entire density of the transformed variable $g(X)$. We can make due with the original density f_X . All of the results in Section 3.3 about the linearity of expectation and the variance apply unmodified to continuous and absolutely continuous random variables. They apply to all random variables regardless of type, in fact:

- $E(aX + b) = a\mathbb{E}[X] + b$;
- $E\left(\sum_{i=1}^n c_i X_i\right) = \sum_{i=1}^n c_i E(X_i)$, no matter how the X_i are related;
- $\text{var}(X) = E[(X - \mathbb{E}[X])^2] = E(X^2) - \mathbb{E}[X]^2$;
- $\text{var}(aX + b) = a^2 \text{var}(X)$.
- If the X_i are independent, then $\text{var}\left(\sum_{i=1}^n c_i X_i\right) = \sum_{i=1}^n c_i^2 \text{var}(X_i)$.

Figure 3.21: pdf and cdf of $\text{Unif}(a, b)$.

3.6.1 Uniform distribution

Definition 3.15. X has the **continuous uniform distribution** on the interval $\text{Range}(X) = [a, b]$ if the density is

$$f_X(x) = \begin{cases} \frac{1}{b-a} & a \leq x \leq b \\ 0 & \text{else.} \end{cases} \quad (3.19)$$

This density is pictured in Figure 3.21, and we denote this $X \sim \text{Unif}(a, b)$.

Compare this with the discrete uniform distribution from Section 3.2.5. In that case, “uniform” meant every every value in the range had the same probability. That is trivially true for all continuous random variables because they have $P(X = x) = 0$ for all $x \in \mathbb{R}$, so “uniform” has a slightly different meaning in the continuous case. It means that all tiny lil’ sub-intervals of the same length have the same probability. So if $(c, d), (u, v) \subseteq (a, b)$ have the same length $d - c = v - u$, then $P(c < X < d) = P(u < X < v) = (d - c)/(b - a) = (v - u)/(b - a)$. This is illustrated in Figure 3.22

Does the pdf integrate to 1?

It had better. Let’s check:




Figure 3.22: The continuous uniform distribution is “uniform” in the sense that sub-intervals of the same length have the same probability, regardless where they are located.

$$\begin{aligned}
 \int_{-\infty}^{\infty} f_X(x) dx &= \int_{-\infty}^a f_X(x) dx + \int_a^b f_X(x) dx + \int_b^{\infty} f_X(x) dx \\
 &= \int_{-\infty}^a 0 dx + \int_a^b \frac{1}{b-a} dx + \int_b^{\infty} 0 dx \\
 &= \frac{1}{b-a} \int_a^b 1 dx \\
 &= \frac{1}{b-a} [x]_a^b \\
 &= \frac{1}{b-a} (b-a) \\
 &= 1.
 \end{aligned}$$

What is the cdf?

The cdf is $F_X(x) = P(X \leq x) = \int_{-\infty}^x f_X(t) dt$. If $x < a$, then

$$F_X(x) = \int_{-\infty}^x f_X(t) dt = \int_{-\infty}^x 0 dt = 0.$$

If $a \leq x \leq b$, then

$$\begin{aligned}
 F_X(x) &= \int_{-\infty}^x f_X(t) dt \\
 &= \int_{-\infty}^a f_X(t) dt + \int_a^x f_X(t) dt \\
 &= \cancel{\int_{-\infty}^a 0 dt} + \int_a^x \frac{1}{b-a} dt \\
 &= \frac{1}{b-a} [t]_a^x \\
 &= \frac{x-a}{b-a}.
 \end{aligned}$$

If $b < x$, then

$$\begin{aligned}
 F_X(x) &= \int_{-\infty}^x f_X(t) dt \\
 &= \int_{-\infty}^a f_X(t) dt + \int_a^b f_X(t) dt + \int_b^x f_X(t) dt \\
 &= \cancel{\int_{-\infty}^a 0 dt} + \int_a^b \frac{1}{b-a} dt + \cancel{\int_b^x 0 dt} \\
 &= 1.
 \end{aligned}$$

So the cdf of $X \sim \text{Unif}(a, b)$ is

$$F_X(x) = \begin{cases} 0 & x < a \\ \frac{x-a}{b-a} & a \leq x \leq b \\ 1 & b < x. \end{cases} \quad (3.20)$$

This is displayed in Figure 3.21.

What is the mean?

$$\begin{aligned}
 \mathbb{E}[X] &= \int_{-\infty}^{\infty} x f_X(x) dx \\
 &= \int_{-\infty}^a x f_X(x) dx + \int_a^b x f_X(x) dx + \int_b^{\infty} x f_X(x) dx \\
 &= \int_{-\infty}^a x \cdot 0 dx + \int_a^b x \frac{1}{b-a} dx + \int_b^{\infty} x \cdot 0 dx \\
 &= \frac{1}{b-a} \int_a^b x dx \\
 &= \frac{1}{b-a} \left[\frac{1}{2} x^2 \right]_a^b \\
 &= \frac{1}{b-a} \left[\frac{1}{2} b^2 - \frac{1}{2} a^2 \right] \\
 &= \frac{1}{2(b-a)} [b^2 - a^2] \\
 &= \frac{1}{2(b-a)} (b-a)(b+a) \\
 &= \frac{b+a}{2}.
 \end{aligned}$$

So the expected value is the midpoint of the interval $[a, b]$. This makes sense because the uniform distribution on an interval is symmetric around its midpoint.