UNIT 5 · BLOCK 1

z-Scores & the Normal Curve

One number that tells you exactly where a score stands.

Why this matters

Is a stress score of 32 "high"? Is a 3.9 GPA "impressive"? You can't say until you know what the rest of the group looks like. A **z-score** answers that in one move — it tells you how far above or below average a score sits, in a language that lets you compare *anything to anything*. It's the last big idea before Exam 1, and it ties together everything we've built with means and standard deviations. Master this and the rest of the course — every test, every *p*-value — is really just z-scores wearing a costume.

By the end you can...

- ✓ Explain what a **z-score** is and why it makes any score interpretable.
- ✓ Calculate a z-score by hand with $\mathbf{z = (X - \mu) \div \sigma}$ and in Google Sheets.
- ✓ Use the **normal curve** and the **68–95–99.7 rule** to picture where scores fall.
- ✓ Convert a z-score into a **percentile** — and back — and compare scores on different scales.

1. The big idea — a score with no context is silent

A raw score by itself is almost meaningless. **32 stress points** tells you nothing until you ask the only question that matters: *compared to whom?* Thirty-two is alarming if the class averages 20 and nearly impossible if it averages 60. The number alone can't speak — it needs the **mean** (where the group centers) and the **standard deviation** (how spread out the group is) before it means anything at all.

A **z-score** supplies exactly that context and folds it into a single number. It rescales any raw score into one universal language: **how many standard deviations away from the mean it is.** Once a score has been turned into a z, it has been **standardized** — stripped of its original units and placed on a common ruler that every other standardized score shares.

If z is...	The score is...
+ (positive)	**above** the mean
− (negative)	**below** the mean
0	**exactly average** (right at the mean)

A z of +2 means "two standard deviations above average." A z of −1 means "one standard deviation below." A z of 0 means you *are* the average. That's the whole idea — the rest of this unit is just learning to compute it, picture it, and read a percentile off of it.

CHECK **A classmate says, "I scored a 47 on the anxiety measure — is that a lot?" What's the one thing you can't answer without?**

The mean and SD of the group. A raw score is silent on its own. Only the group's center (mean) and spread (SD) tell you whether 47 is high, low, or dead average — and those two numbers are exactly what a z-score uses.

2. The normal curve (the "bell curve")

Before we standardize anything, look at the shape most of our scores actually take. Pour together a big pile of natural measurements — height, reaction time, test scores, stress — and they tend to settle into one familiar form: tall in the middle, tailing off symmetrically on both sides. We call it the **normal curve** (or, casually, the **bell curve**). It shows up so often because most outcomes are pushed around by *many small, independent causes* that mostly cancel out — a deep mathematical result that's well beyond us here, but worth knowing isn't an accident.

Three features define every normal curve:

- It's **symmetric** — the left half is a mirror image of the right.
- The **mean, median, and mode all sit in the same spot**: the exact center, where $z = 0$.
- It's **unimodal** — one peak in the middle, thinning steadily as you move toward either tail.

The further you move from the center in either direction, the rarer scores become. That's *why* a z of +2 is impressive and a z of +3 is genuinely unusual — out there, the curve has almost no height left, so almost nobody lands that far out. Drag the sliders below and watch the curve, the score marker, and the percentile move together:

Notice the two controls do different jobs. Sliding the **mean** shifts the whole curve left or right without changing its shape — it just moves where "average" is. Sliding the **SD** changes the spread: a wider SD flattens and stretches the curve, a narrow SD makes it tall and tight. The very same raw score becomes more or less unusual depending on how spread out the group around it is — which is precisely why the z-score formula has to divide by the SD.

3. The 68–95–99.7 rule (the empirical rule)

Here's what makes the normal curve so useful: its proportions are **fixed**. No matter what the mean or SD happen to be, the share of scores within a given number of standard deviations of the mean is always the same. That regularity is called the **empirical rule**, or just the **68–95–99.7 rule**:

Within…	…of the mean fall about	Outside (both tails)
±1 SD (z from −1 to +1)	**68%**	≈ 32% (16% each tail)
±2 SD (z from −2 to +2)	**95%**	≈ 5% (2.5% each tail)
±3 SD (z from −3 to +3)	**99.7%**	≈ 0.3% (rare air)

So almost *everybody* — about 997 of every 1,000 people — lands within 3 SDs of the mean. A z beyond ±3 is genuinely exceptional. The static figure below makes the bands concrete; the shaded region marks a single score (here z = +1) and the area below it:

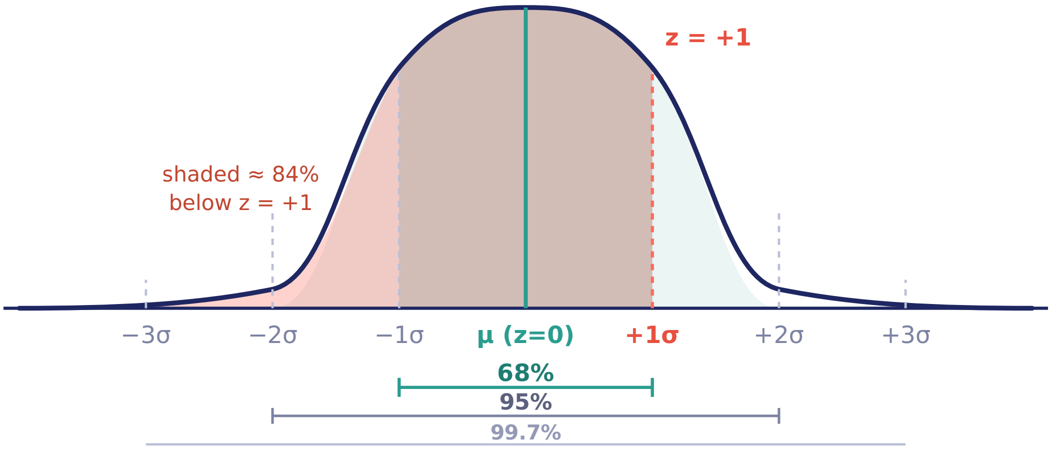

The fixed proportions of a normal curve. The teal core is ±1 SD (68%); the wider band is ±2 SD (95%); the full sweep is ±3 SD (99.7%). The coral region is everything below z = +1 — about 84% of scores.

The cool part — our REAL students follow it

Of the 200 students in the dataset, **66%** fall within ±1 SD of mean stress, **95%** within ±2 SD, and **100%** within ±3 SD. That's almost dead-on the 68–95–99.7 rule — a textbook bell curve, built from actual people sitting in classrooms like ours. The rule isn't a fairy tale; real psychological data really do behave this way.

CHECK **On a normal curve, about what percentage of scores fall *within* ±2 standard deviations of the mean?**

About 95%. The 68 figure is for ±1 SD. Within ±2 SD you capture about 95% — which leaves only ~5% in the two tails combined, or ~2.5% out past +2 on each side.

4. The z-score formula

> **The one formula**
>
> **$z = (X - \mu) \div \sigma$.** In words: take your raw score X, subtract the mean μ, then divide by the standard deviation σ. The subtraction $(X - \mu)$ gives you the *direction and raw distance* from average. The division by σ converts that distance into a count of **standard deviations** — the universal unit.

Two small notes on symbols you'll see all over this course. We write **μ** (mu) for the mean and **σ** (sigma) for the standard deviation when we're treating the numbers as a whole population; you'll also see the sample versions **$\bar{x}$** and **s**. The z-score formula works identically either way — plug in whichever center and spread describe your group.

On OUR data — stress (column J), mean = **20.1**, SD = **5.9**:

- A student with stress **32**: $z = (32 - 20.1) \div 5.9 =$ **+2.02** → *just over two SDs above average — very high stress.*
- A student with stress **14**: $z = (14 - 20.1) \div 5.9 =$ **−1.03** → *about one SD below average — calmer than most.*

On gpa (column K), mean = **3.08**, SD = **0.44**: a **3.9 GPA** gives $z = (3.9 - 3.08) \div 0.44 =$ **+1.87** → *almost two SDs above average — well above the pack.* Notice the GPA scale (0–4) and the stress scale (roughly 0–40) are wildly different in size, yet both produce a z near +2. That is the magic we'll cash in shortly.

Worked example — raw score → z → percentile

A student scores **X = 26** on a measure with mean **μ = 20** and SD **σ = 6**. How unusual is that? Predict each step before revealing it.

STEP 1 — FIND THE DISTANCE FROM THE MEAN

Subtract: 26 − 20 = `6`. The score sits 6 raw points *above* average (positive, so above).

STEP 2 — DIVIDE BY THE SD TO STANDARDIZE

z = 6 ÷ 6 = `+1.00`. The score is exactly one standard deviation above the mean.

STEP 3 — LOCATE Z ON THE CURVE

z = +1.00 sits at the edge of the 68% middle band. By symmetry, the lower half (50%) plus half of the 68% band (34%) lies below it: 50 + 34 = `84%`.

STEP 4 — STATE THE PERCENTILE

The score is at about the `84th percentile` — it beats roughly 84% of the group. (A z-table or `=NORM.S.DIST` gives 84.13%; the 68–95–99.7 rule got us there by eye.)

Try it now — the z-Score calculator

Type a raw score, the mean, and the SD — then read the z-score and percentile straight off. Hit **Load example** to drop in a real stress score from our data, and check it against the worked example above (X = 26, μ = 20, σ = 6).

5. From z-score to percentile (and back)

A **percentile** says what percent of the group scored *below* you. Because the normal curve's proportions are fixed, every z-score maps to exactly one percentile — and that mapping is what a **z-table** (or a spreadsheet) stores. The empirical rule already gives you the famous landmarks by heart:

z-score	Percentile	Meaning
−2	~2nd	near the bottom — only ~2% scored lower
−1	~16th	below average — ~16% scored lower
0	**50th**	exactly average — half the group scored below you
+1	**~84th**	you beat about 84% of the group
+2	~98th	near the top — only ~2% scored higher

Where do the in-between numbers come from? Each is just "50% for the lower half, plus (or minus) half of whatever band you've crossed." For z = +1 you add half of 68% (= 34%) to the bottom 50% and get the **84th**. For z = −1 you subtract: 50 − 34 = the **16th**. So our student with a stress z of +2.02 scored higher than roughly **98%** of classmates. That's a precise, honest statement — and it came from one little formula plus the shape of the curve.

Run it backwards too

The map works in reverse. "Top 5%" means the *highest* 5%, which starts at the **95th percentile** — about $z = +1.65$. "Bottom 10%" sits below roughly $z = -1.28$. Cut-scores, scholarships, and clinical thresholds are usually defined this way: pick a percentile, then read off the z that marks it.

CHECK **A score lands at $z = 0$. What percentile is it, and what does that say about the score?**

50th — exactly average. $z = 0$ means the score equals the mean, which sits dead center on a symmetric curve. Half the group scored below it. $z = 0$ is the *middle*, never "nothing" or "the bottom" — a classic trap.

6. Why z-scores make different scales comparable

This is the payoff. Raw scores live in their own units — stress points, GPA, SAT points, inches — and you cannot compare across them directly. A 1300 SAT and a 29 ACT are measured on scales that don't even overlap. But standardizing strips the units away. Once both scores become z-scores, they live on the **same ruler**, centered at 0 with a standard deviation of 1, and you can lay them side by side. Whoever has the larger z did better *relative to their own group* — and the comparison is fair even though the original tests share nothing.

The same trick answers a question about a single person across two traits. Is our student more unusual in **stress** ($z = +2.02$) or in **GPA** ($z = +1.87$)? Ignore the signs and compare the magnitudes: the bigger $|z|$ is the more extreme. Here stress (2.02) edges out GPA (1.87), so this student is a touch more of an outlier in stress than in grades — a sentence you simply could not write from the raw numbers 32 and 3.9.

CHECK **On a math quiz Mara's z = +1.5; on a writing quiz her z = +0.8. The quizzes used totally different point scales. Relative to her classmates, she did better on…**

Math. Standardizing puts both quizzes on the same ruler, so the raw point scales no longer matter — only the z's do. A larger positive z (+1.5 vs. +0.8) means she stood further above her group in math. That comparison is exactly what raw scores could never give you.

Stat in the Wild — comparing apples to oranges

A friend got a **1300 on the SAT**. Another got a **29 on the ACT**. Who did better? The tests use totally different scales — you can't compare 1300 to 29 directly. But convert each to a **z-score** (using that test's mean and SD) and suddenly they're in the same language. Whoever has the higher z did better *relative to their group*. College admissions offices, percentile rank reports, and "concordance tables" all run on exactly this logic.

The lesson: z-scores let you compare scores measured on **different scales**. They're also the engine under IQ scores (mean 100, SD 15), clinical cutoffs, and — coming soon — every hypothesis test you'll run. When you compute a test statistic and look up a *p*-value, you're standardizing and reading the tail of a curve. Same move, dressed up.

In Google Sheets

You won't hunt through a paper z-table on exam day — Sheets does both halves for you. Open the **SJNY Student Life** dataset and try these. First, get the group's center and spread; then standardize any score:

```
=AVERAGE(J2:J201)                   → the mean (μ) of stress
=STDEV(J2:J201)                     → the standard deviation (σ)
=STANDARDIZE(32, 20.1, 5.9)         → the z-score: (32 − 20.1) ÷ 5.9 ≈ 2.02
```

`=STANDARDIZE(X, mean, SD)` *is* the z formula in one cell — same three inputs, same answer. To turn that z into a percentile (the area of the curve below it), use the standard-normal distribution function:

```
=NORM.S.DIST(2.02, TRUE)            → 0.978  (about the 98th percentile)
=NORM.S.DIST(1, TRUE)               → 0.8413 (the ~84th percentile from our example)
```

The `TRUE` asks for the *cumulative* area — everything below that z. To run it backwards (percentile → z), use `=NORM.S.INV(0.95)`, which returns about **1.645** — the z that marks the top 5%. Between `STANDARDIZE`, `NORM.S.DIST`, and `NORM.S.INV`, you can move freely among raw scores, z-scores, and percentiles without ever touching a table.

Watch out for

Sign matters for direction, not for "extremeness" — a z of −2.5 is just as far out as +2.5, same distance from average, opposite side. **z = 0 is average**, not "zero" or "bad" — it's literally the middle (50th percentile). **Bigger |z| = rarer**: near 0 is common, past ±2 is uncommon, past ±3 is rare air. And the **68–95–99.7 rule assumes a roughly normal shape** — for badly skewed data the percentages are only a rough guide, so picture the histogram first.

Key terms

z-score · Standardize · Raw score (X) · μ (mean)

σ (standard deviation) · Normal curve · Bell curve

Symmetric · 68–95–99.7 rule · Empirical rule

Percentile · z-table · Standard normal

PRACTICE — NEVER GRADED

Check yourself

Ten quick true/false items with instant explanations. (Different from your Canvas quiz — real practice.)

Sigma says: A z-score is just a translation. Subtract the mean to find the distance, divide by the SD to count it in standard-deviation steps. Once a score speaks "z," you can compare it to anything — and read its percentile right off the curve.