The Expected Value Puzzle

A Note on the Redistricting Court Example

Gov 2001 · Spring 2026

February 9, 2026

The Puzzle

On Thursday, February 5th, a puzzle came up in class. We were working through a redistricting example with a 5-justice state supreme court. Each justice had a different probability of voting to strike down a gerrymandered map:

Justice	1	2	3	4	5
$\mathbb{P}$(strike)	0.70	0.60	0.55	0.40	0.35
$\mathbb{P}$(uphold)	0.30	0.40	0.45	0.60	0.65

Assuming independence, we worked through all the combinations to construct the PMF for X = total votes to strike. (See last Thursday's slides for the full derivation.) The result:

x (votes to strike)	0	1	2	3	4	5
$f(x) = \mathbb{P}(X = x)$	0.02	0.13	0.31	0.34	0.17	0.03

Then we calculated two things:

1. **Probability the map is struck down** (requires 3+ votes):

$$\mathbb{P}(X \geq 3) = 0.34 + 0.17 + 0.03 = \mathbf{0.54}$$

2. **Expected number of votes to strike**:

$$\mathbb{E}[X] = 0(0.02) + 1(0.13) + 2(0.31) + 3(0.34) + 4(0.17) + 5(0.03) = \mathbf{2.60}$$

The puzzle: The expected value is 2.6—*below* the 3-vote threshold. Yet the probability of striking down the map is 54%—the map gets struck down *more often than not*.

How can the *average* outcome be below the threshold while the map is *usually* struck down?

I'm going to explain this three different ways. They all tell the same story, but different visualizations click for different people.

Approach 1: Comparing PMF to Contributions

The PMF shows *where probability sits*. But expected value weights each outcome by its position. Let's see both:

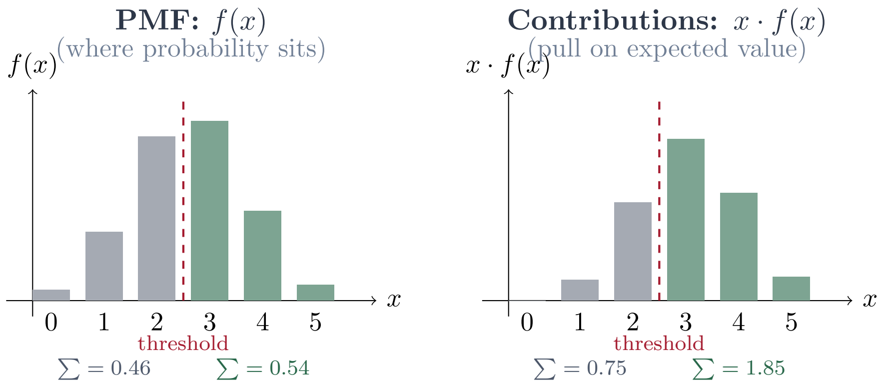

Left chart (PMF): More probability mass sits on the right (54% vs 46%). So the map is *usually* struck down.

Right chart (contributions): The $x = 2$ bar is now substantial—it contributes 0.62 to the expected value. That's a lot of "pull" concentrated just below threshold. The total contribution from the left (0.75) plus the right (1.85) gives us $\mathbb{E}[X] = 2.60$.

Approach 2: The Balance Beam

Think of expected value as a *balance point*. Place weights at each position, where the weight equals $f(x)$. The beam balances at $\mathbb{E}[X]$.

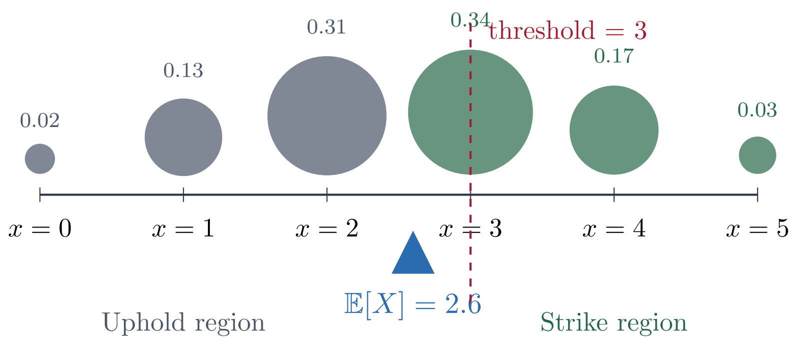

The green weights (strike region) have more total mass: $0.34 + 0.17 + 0.03 = 0.54$. But the gray weights are positioned further from center. The big weight at $x = 2$ sits close to the balance point, pulling it leftward. The beam balances at 2.6—in the uphold region—even though more mass is on the strike side.

Approach 3: The Decomposition Table

Let's decompose $\mathbb{E}[X]$ term by term:

	Uphold ($x < 3$)			**Strike** ($x \geq 3$)		
x	0	1	2	3	4	5
$f(x)$	0.02	0.13	0.31	0.34	0.17	0.03
$x \cdot f(x)$	0.00	0.13	0.62	1.02	0.68	0.15

	Uphold ($x < 3$)	**Strike** ($x \geq 3$)
Total probability $\sum f(x)$	0.46	**0.54**
Contribution to $\mathbb{E}[X]$ $\sum x \cdot f(x)$	0.75	**1.85**

The strike region has more probability (0.54 vs 0.46) *and* contributes more to $\mathbb{E}[X]$ (1.85 vs 0.75). So why is $\mathbb{E}[X] = 0.75 + 1.85 = 2.60$ below threshold?

Because expected value is a *weighted position*, not a comparison of contributions. The right side contributes more, but those contributions come from positions 3, 4, 5. The left side contributes less, but from positions 0, 1, 2. The weighted average of all positions lands at 2.6.

The Bottom Line

All three approaches tell the same story:

Question	**Answer**
"Will the map usually be struck down?"	Yes — $\mathbb{P}(X \geq 3) = 0.54$
"What's the average number of strike votes?"	2.6 — below threshold

These aren't contradictory. They're *different questions*:

- **Probability** asks: which side of the threshold has more mass?
- **Expected value** asks: where does the distribution balance?

When a distribution is asymmetric, these can point in different directions. This happens all the time in political science:

- A candidate can win the popular vote but lose the Electoral College
- A bill can pass with majority support even if the average legislator is lukewarm
- A party can win most seats even if their average vote share is below 50%

The lesson: always be clear about *which question you're asking*. Expected value and probability answer different things.

This note accompanies the Week 3 lecture on discrete distributions. The redistricting example is from the Feb 5 working session.