Statistical Plot Twists

A 5-minute journey through statistical illusions

Frequentist vs Bayesian Inference Coin Problem

A coin is flipped 14 times and shows 10 heads. We must decide whether to bet on the next two flips being both heads (HH) or not.

Frequentist Approach

Estimate: p^ = 10/14 = 5/7

Predictive probability: P(HH) (5/7)^2 = 25/49 = 0.510204...

Decision: Since P(HH) > 0.5 bet FOR HH

Bayesian Approach (Uniform Prior)

Prior: p ~ Beta(1,1) (uniform prior)

After 10 heads, 4 tails Posterior: p ~ Beta(11,5)

Predictive probability: E[p^2] = (1112)/(1617) = 132/272 = 33/68 = 0.4853

Decision: Since P(HH) < 0.5 bet AGAINST HH

- Earlier, the Bayesian and the Frequentist just had a friendly argument. But give them a giant dataset — and they go to war.

- This is **Lindley's Paradox**, where a Frequentist shouts 'Highly significant!' and the Bayesian calmly replies, 'There's no evidence at all.

- n = 10,000 flips, 5,098 heads Is the coin fair?

Frequentist:

n = 10,000 flips, 5,098 heads

z = 1.96, p = 0.05

Reject the null, the coin is not fair.

Not Fair

Bayesian:

Prior Beta(1,1)

Bayes factor 11.7 : 1 for H0

Posterior(H0) = 92% → coin probably fair

Probably fair

- Sample size is huge

- Observed effect is tiny

- Alternative model (H_1) uses a diffuse prior

- Frequentist says “significant!”, Bayesian says “meh, still probably null.”

Simpson’s Paradox, 2 Batsmen, K vs R

Year 2023	Runs	Innings	Average
K	825	15	55
R	1620	30	54
Year 2024			
K	900	30	30
R	290	10	29

Total average
K = 1725/45 = 38.33
R = 1910/40 = 47.75

“K > R in each year, but R > K overall — weighting (innings) changed the verdict. That’s Simpson’s paradox: aggregation with unequal group sizes can flip conclusions**.**

Mathematical Insight: This is due to confounding by the subgroup proportions

Formally, let $p_1 = \frac{a}{b}$, $p_2 = \frac{c}{d}$for subgroups, but combined $p = \frac{a+c}{b+d}$, and if $b \ll d$but $p_1 > p_2$, the weights can flip.

Aggregate improperly and the story flips.

Moral: always check subgroup sizes before trusting an overall average.

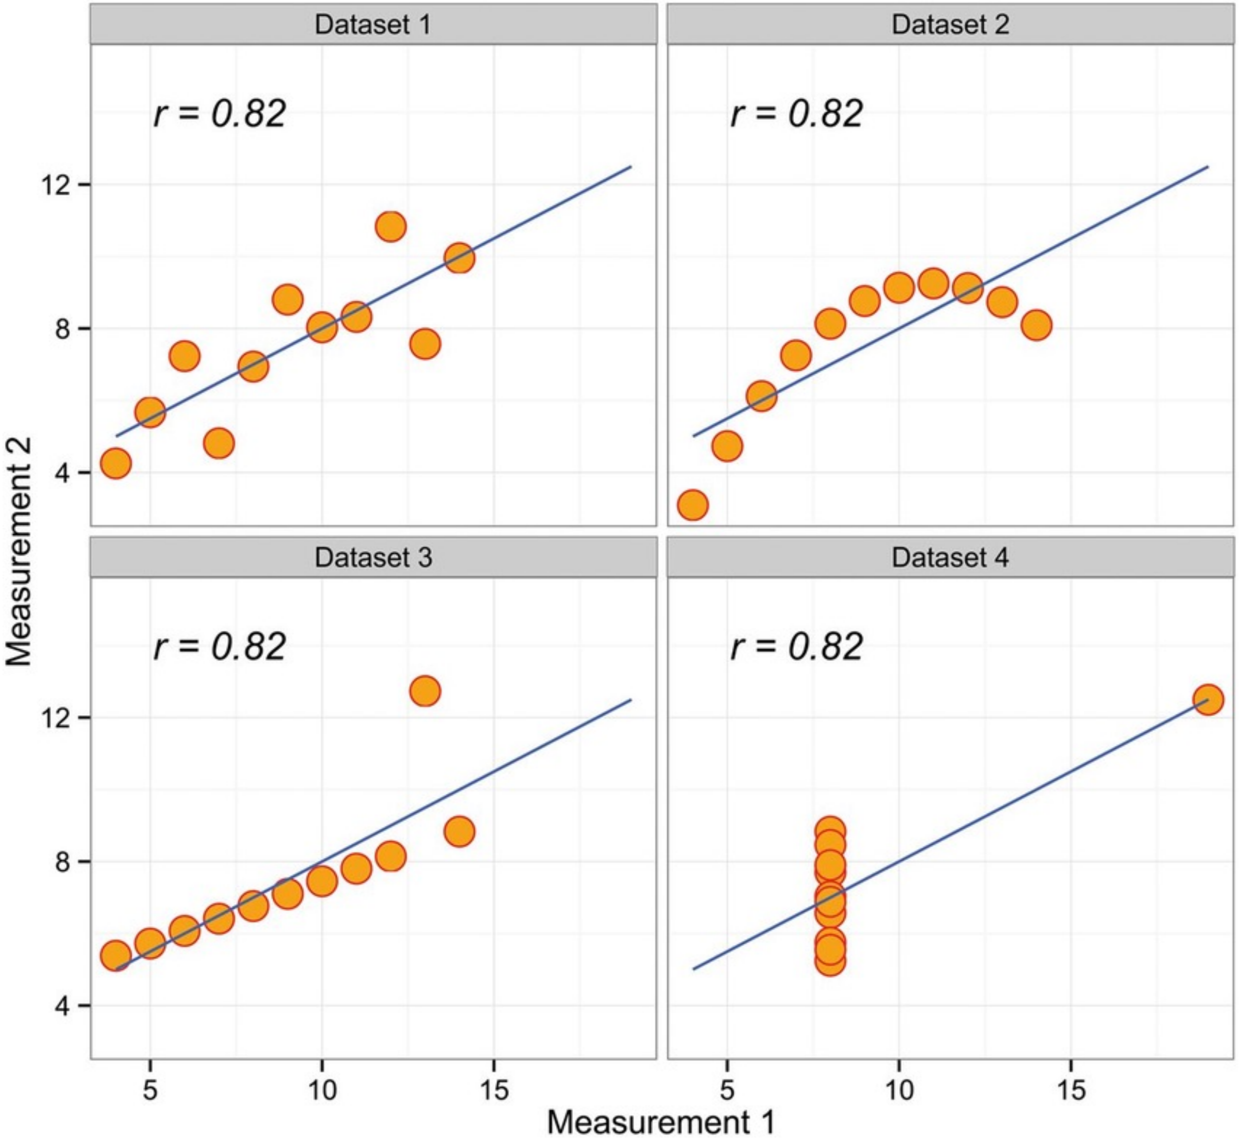

Anscombe’s Quartet — When Averages Lie

Same mean of x = 9.0

Same mean of y = 7.5

Same variance of x = 11

Same variance of y = 4.127

Same correlation (r) = 0.816

Same regression line:

y=3+0.5xy=3+0.5x

“Always plot your data before you trust your stats.”

Ex : 13 Datasets of Datasaurus Dozen.

Berkson's Paradox :The Hospital Illusion

Why diabetics in hospital seem less likely to have cancer

In the general population, diabetes and cancer are independent.

Both increase the chance of hospital admission.

Among hospital patients, people with diabetes are less likely to also have cancer —because either one is enough to get admitted.

The negative correlation is an illusion — caused by conditioning on “being in hospital.”

Conditioning on a common effect creates false relationships.

Beyond the paradox

- **Lindley:** Significance isn’t evidence.
- **Anscombe:** Summary isn’t the story.
- **Simpson:** Aggregates can lie.
- **Berkson:** Conditioning can mislead.
- **Together:** “It’s not the data — it’s the viewpoint.”