BAYESIAN STATISTICAL CONCEPTS

A gentle introduction

Alex Etz
@alxetz ←Twitter (no ‘e’ in alex)
alexanderetz.com ←Blog
November 5th 2015

Why do we do statistics?

- Deal with uncertainty
  - Will it rain today? How much?
  - When will my train arrive?

- Describe phenomena
  - It rained 4cm today
  - My arrived between at 1605

- **Make predictions**
  - It will rain between 3-8 cm today
  - My train will arrive between 1600 and 1615

Prediction is key

- Description is boring
  - Description:
    - On this IQ test, these women averaged 3 points higher than these men
- Prediction is interesting
  - Prediction:
    - On this IQ test, the average woman will score above the average man
- **Quantitative (precise) prediction is gold**
  - Quantitative prediction:
    - On this IQ test, women will score between 1-3 pts higher than men

Evidence is prediction

- Not *just* prediction in isolation
- Competing prediction
  - Statistical evidence is comparative

Candy bags

A

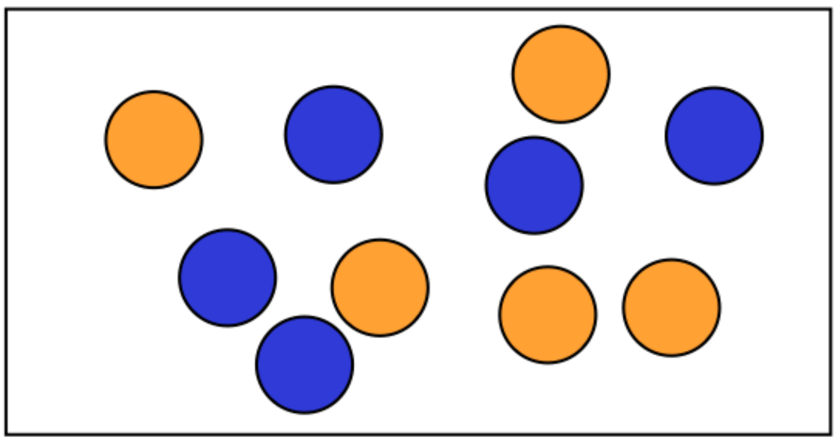

5 orange, 5 blue

B

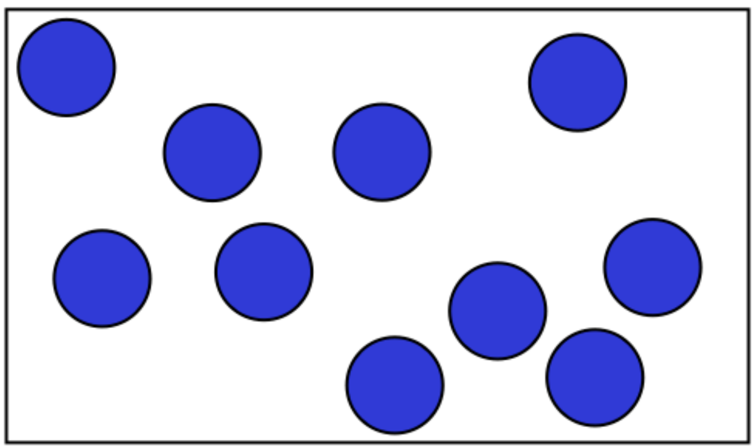

10 blue

Candy bags

- I propose a game
  - Draw a candy from one of the bags
  - You guess which one it came from
  - After each draw (up to 6) you can bet (if you want)

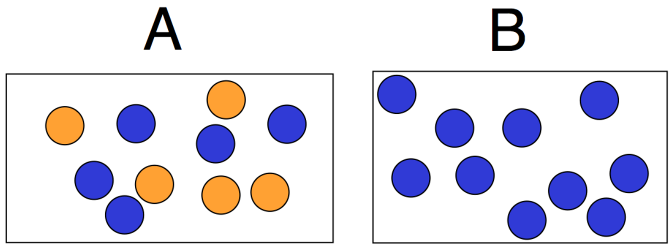

Candy bags

- If **orange**
  - Bag A predicts **orange** with probability .5
  - Bag B predicts **orange** with probability 0

- Given **orange**, there is evidence for A over B
  - How much?
    - Infinity
  - Why?
    - Outcome is impossible for bag B, yet happened
    - Therefore, it **cannot** be bag B

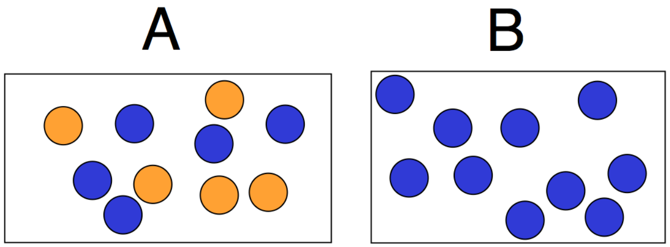

Candy bags

- If **blue**
  - Bag A predicts **blue** with probability .5 (5 out of 10)
  - Bag B predicts **blue** with probability 1.0 (10 out of 10)

- Cannot rule out either bag

- Given **blue**, there is *evidence* for B over A
  - How much?
    - Ratio of their predictions
    - 1.0 divided by .5 = 2 *per draw*

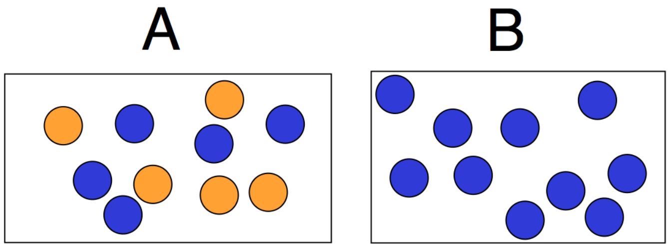

Evidence is prediction

- There is evidence for A over B if:
  - Prob. of observations given by A exceeds that given by B
- Strength of the evidence for A over B:
  - The ratio of the probabilities (very simple!)
  - **This is true for all of Bayesian statistics**
    - More complicated math, but same basic idea
  - This is **not** true of classical statistics

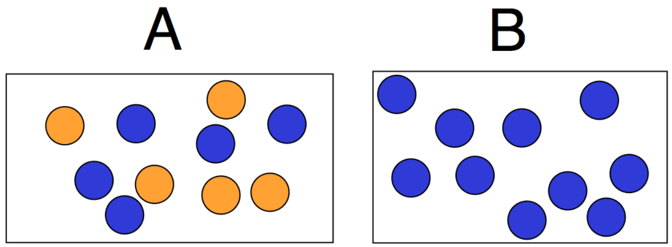

Candy bag and a deck of cards

- Same game, 1 extra step
  - I draw one card from a deck
    - **Red** suit (Heart, Diamond) I draw from bag A
    - **Black** suit (Spade, Club) I draw from bag B
  - Based on the card, draw a candy from one of the bags
  - You guess which one it came from
  - After each draw (up to 6) you can bet
    - (if you want)

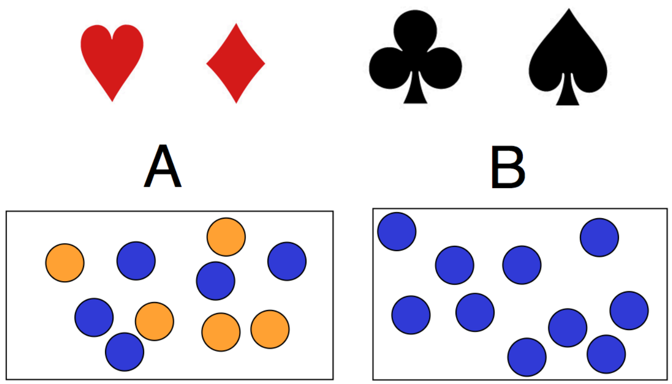

Candy bags and a deck of cards

- If **orange**, it came from bag A 100%. Game ends
- If **blue**
  - Both bags had 50% chance of being selected
  - Bag A predicts **blue** with probability .5 (5 out of 10)
  - Bag B predicts **blue** with probability 1.0 (10 out of 10)
- Evidence for B over A
  - How much?
    - Ratio of their predictions
    - 1.0 divided by .5 = 2 *per* ***blue*** *draw*

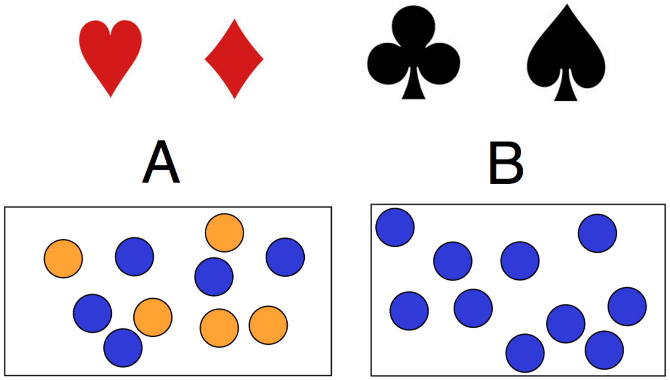

Candy bags and a deck of cards

- Did I add any information by drawing a card?
  - Did it affect your bet at all?

- **If the prior information doesn't affect your conclusion, it adds no information to the evidence**
  - "Non-informative"

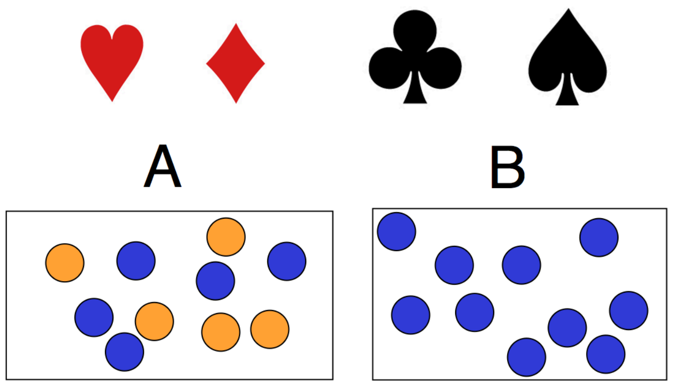

Candy bags and a deck of cards

- Same game, 1 extra step
  - I draw one card from a deck
    - **King of hearts** I draw from bag B
    - ***Any other card*** I draw from bag A
  - I draw a ball from one of the bags
  - You guess which one it came from
  - After each draw (up to 6) you can bet

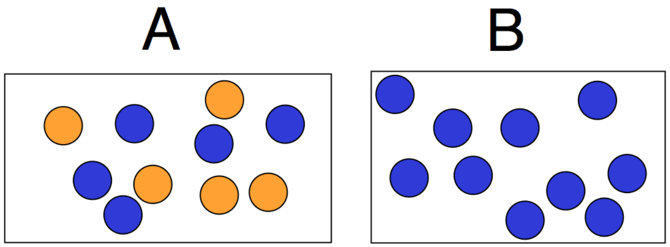

Candy bags and a deck of cards

- Did I add any information by drawing a card?
  - Did it affect your bet at all?

- Observations (evidence) the same
  - But conclusions can differ

- **Evidence is separate from conclusions**

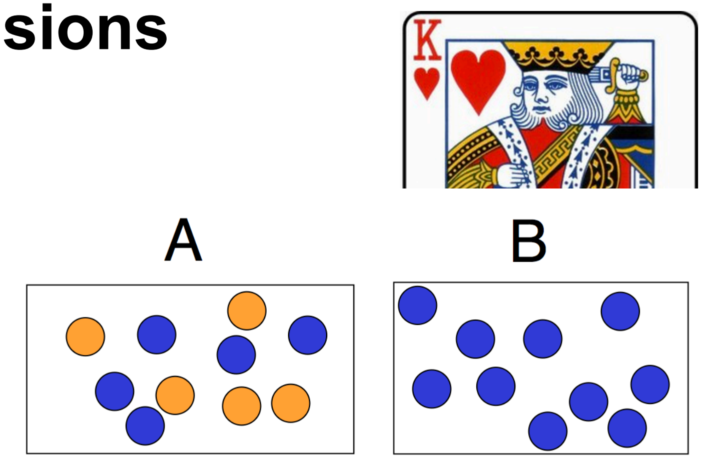

Betting on the odds

- The 1 euro bet
  - If **orange** draw
    - Bet on bag A, you win 100%
    - We have ruled out bag B
  - If **blue** draw
    - Bet on bag A, chance you win is x%
    - Bet on bag B, chance you win is (1-x)%

Betting on the odds

- Depends on:
  - Evidence from sample (candies drawn)
  - Other information (card drawn, etc.)

- A study only provides the evidence contained in the sample

- **You** must provide the outside information
  - Is the hypothesis initially implausible?
  - Is this surprising? Expected?

Betting on the odds

- If initially *fair* odds
  - (Draw **red** suit vs. **black** suit)
  - Same as adding no information
  - Conclusion based only on evidence

- For 1 **blue** draw
  - Initial (prior) odds 1 to 1
  - Evidence 2 to 1 in favor of bag B
  - Final (posterior) odds 2 to 1 in favor of bag B
  - Probability of bag B = 67%

Betting on the odds

- If initially *fair* odds
  - (Draw **red** suit vs. **black** suit)
  - Same as adding no information
  - Conclusion based only on evidence

- For 6 **blue** draws
  - Initial (prior) odds 1 to 1
  - Evidence 64 to 1 in favor of bag B
  - Final (posterior) odds 64 to 1 in favor of bag B
  - Probability of bag B = 98%

Betting on the odds

- If initially *unfair* odds
  - (Draw **King of Hearts** vs. ***any other card***)
  - Adding relevant outside information
  - Conclusion based on evidence combined with outside information

- For 1 **blue** draw
  - Initial (prior) odds 1 to 51 in favor of bag A
  - Evidence 2 to 1 in favor of bag B
  - Final (posterior) odds 1 to 26 in favor of bag A
  - Probability of bag B = 4%

Betting on the odds

- If initially *unfair* odds
  - (Draw **King of Hearts** vs. ***any other card***)
  - Adding relevant outside information
  - Conclusion based on evidence combined with outside information

- For 6 **blue** draws
  - Initial (prior) odds 1 to 51 in favor of bag A
  - Evidence 64 to 1 in favor of bag B
  - Final (posterior) odds 1.3 to 1 in favor of bag B
  - Probability of bag B = 55%

Betting on the odds

- The evidence was the same
  - 2 to 1 in favor of B (1 **blue** draw)
  - 64 to 1 in favor of B (6 **blue** draws)

- Outside information changed conclusion
  - Fair initial odds
    - Initial prob. of bag B = 50%
    - Final prob. of bag B = 67% (98%)

  - Unfair initial odds
    - Initial prob. of bag B = 2%
    - Final prob. of bag B = 4% (55%)

Should you take the bet?

- If I offer you a 1 euro bet:
  - Bet on the bag that has the highest probability
- For other bets, decide based on final odds

Evidence is comparative

- What if I had many more candy bag options?

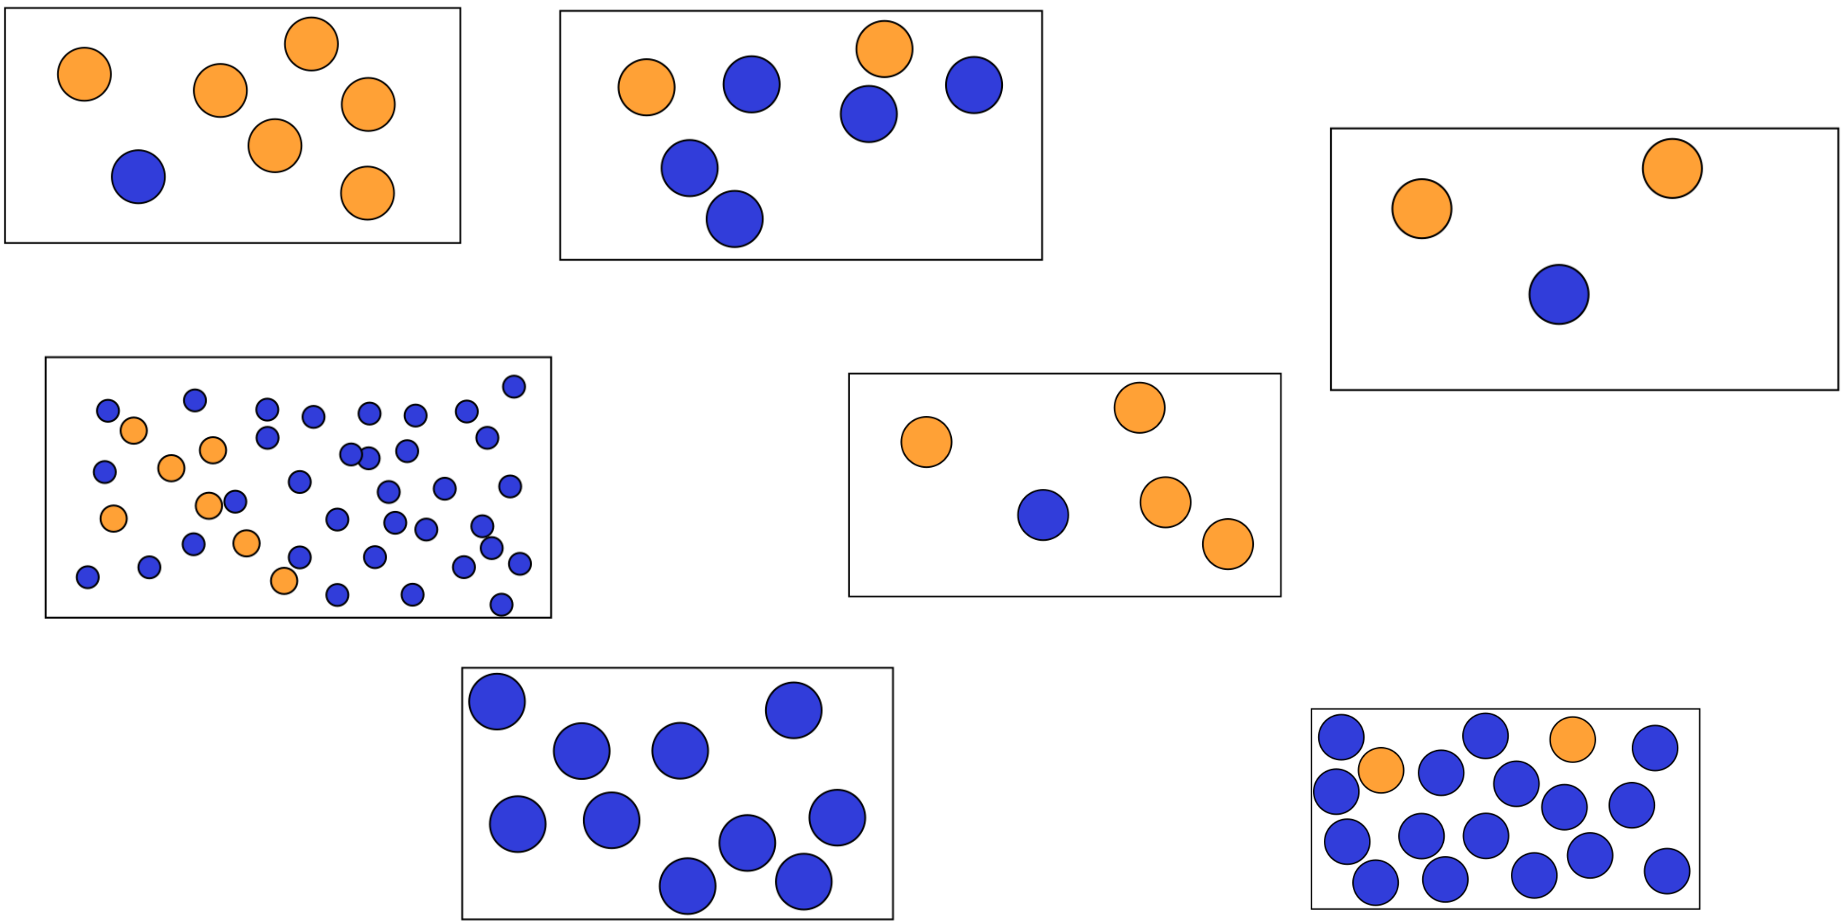

Graphing the evidence

- What if I wanted to compare every possible option at once?
- Graph it!

Graphing the evidence (1 **blue**)

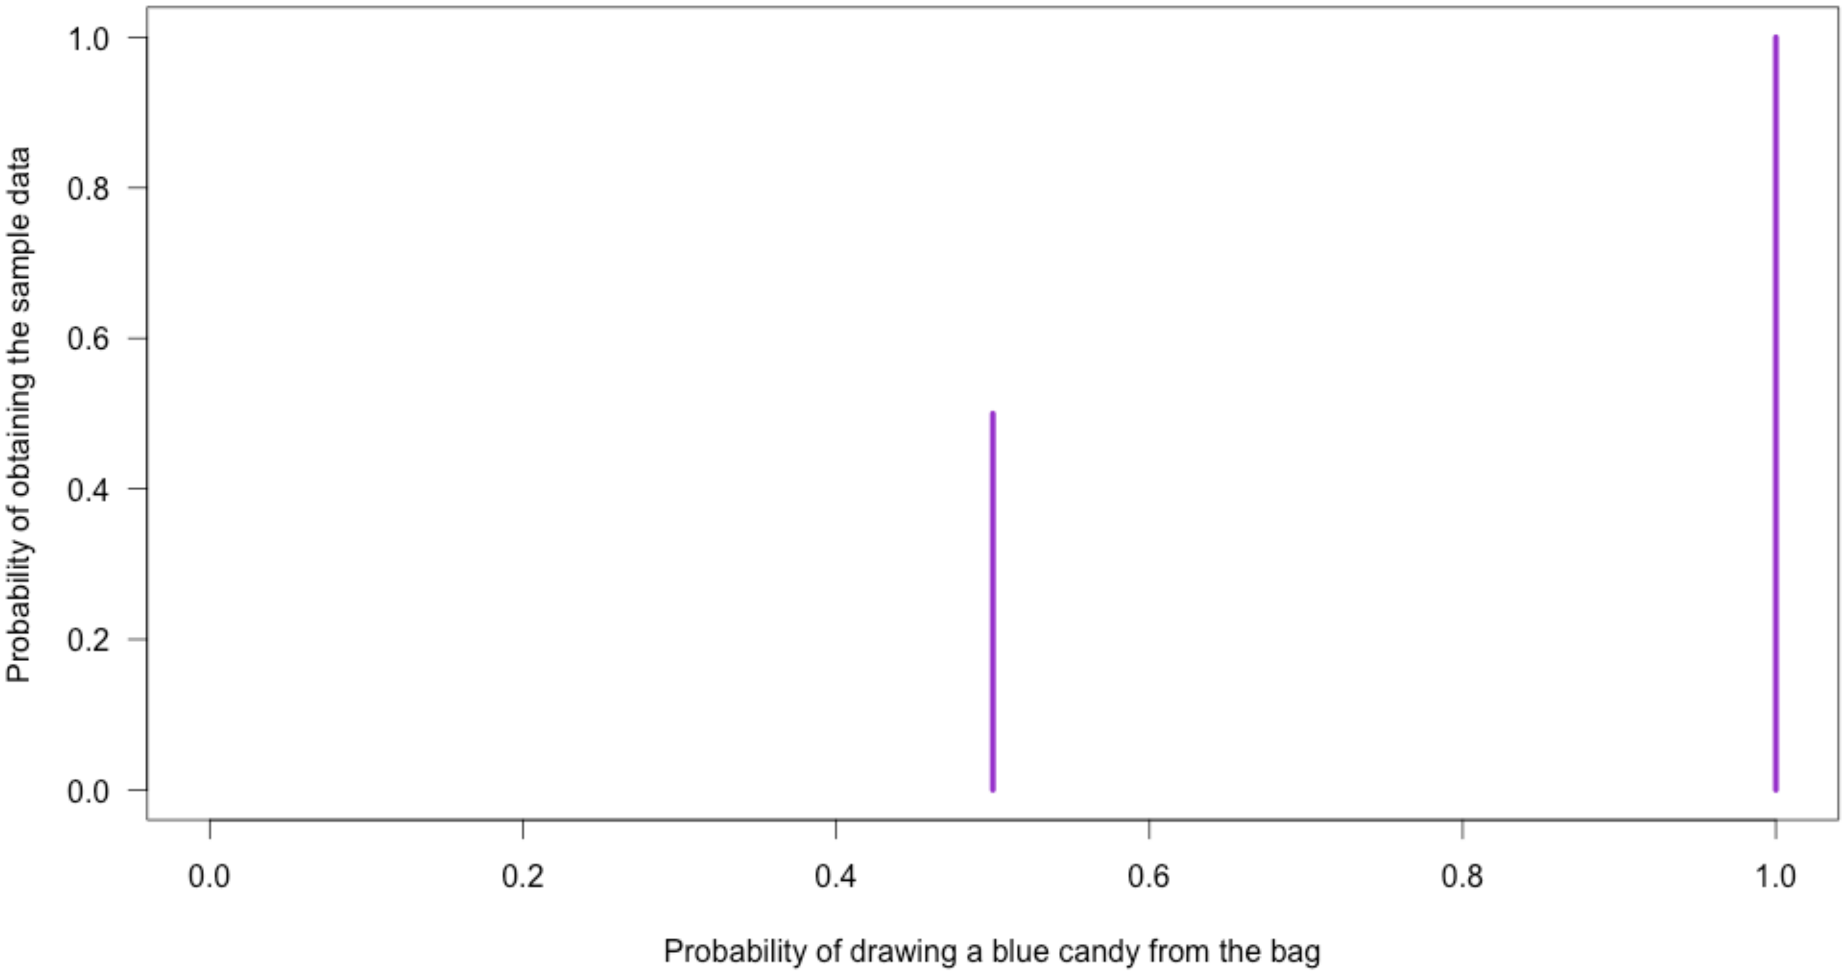

Graphing the evidence (1 **blue**)

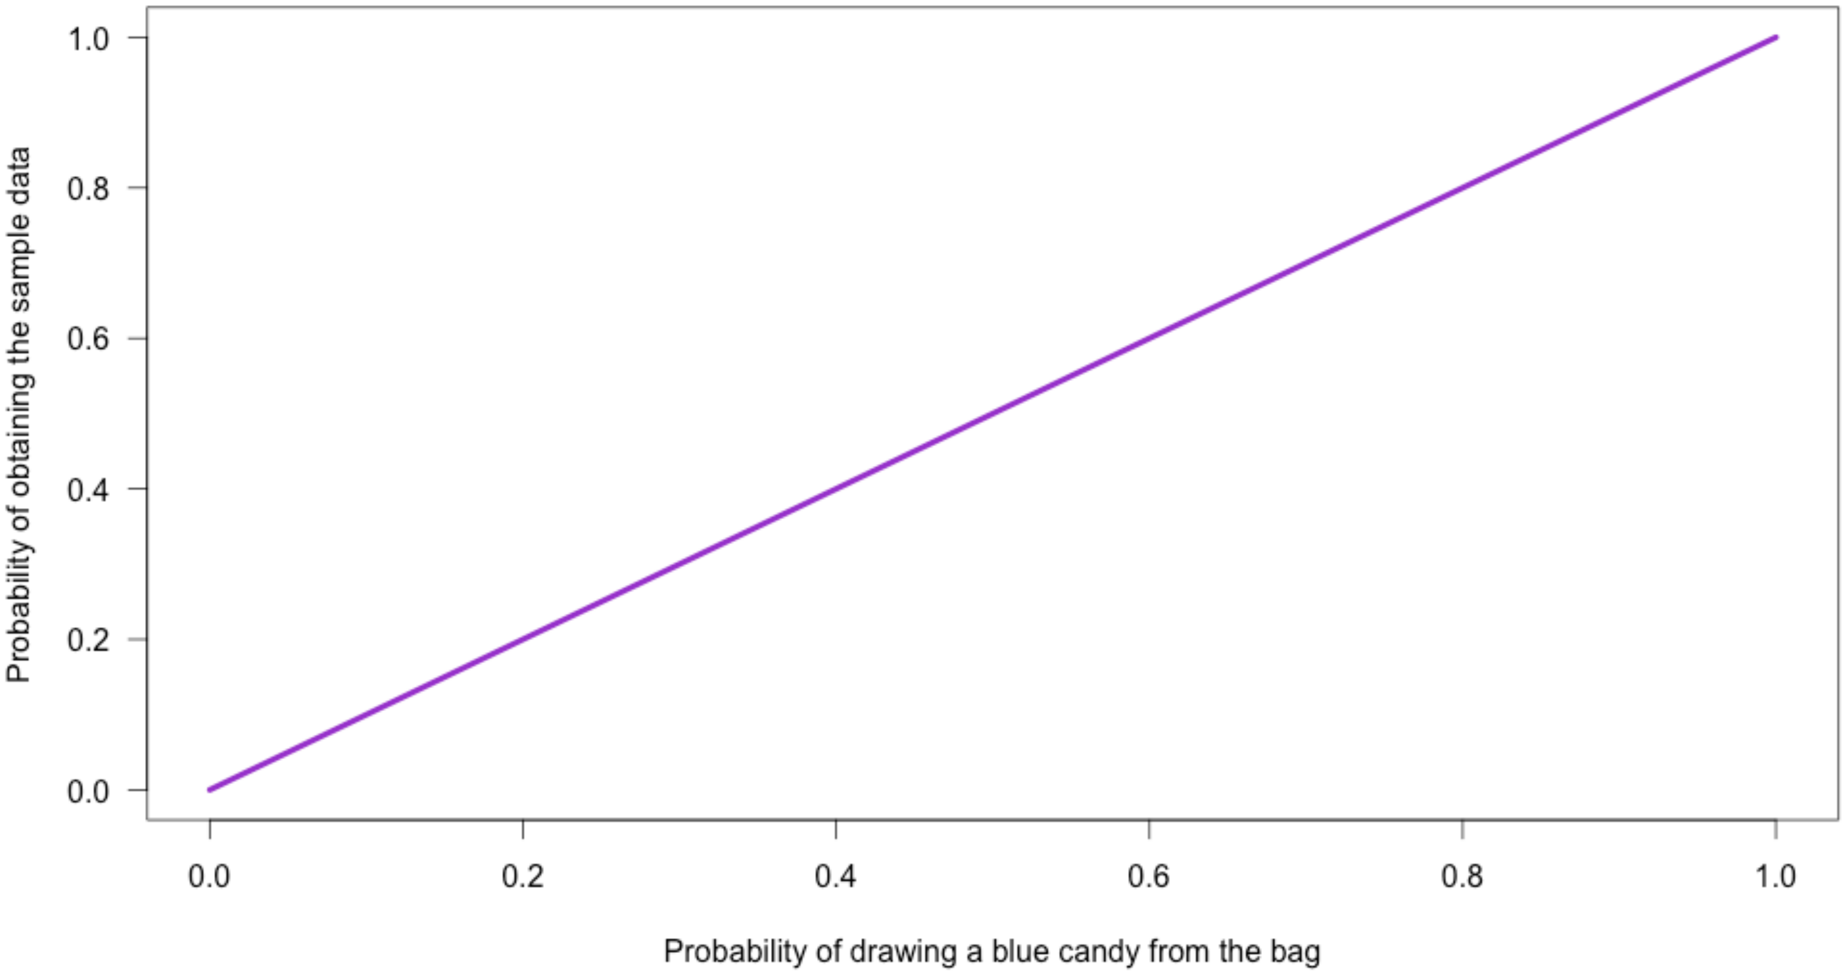

Graphing the evidence (6 **blue**)

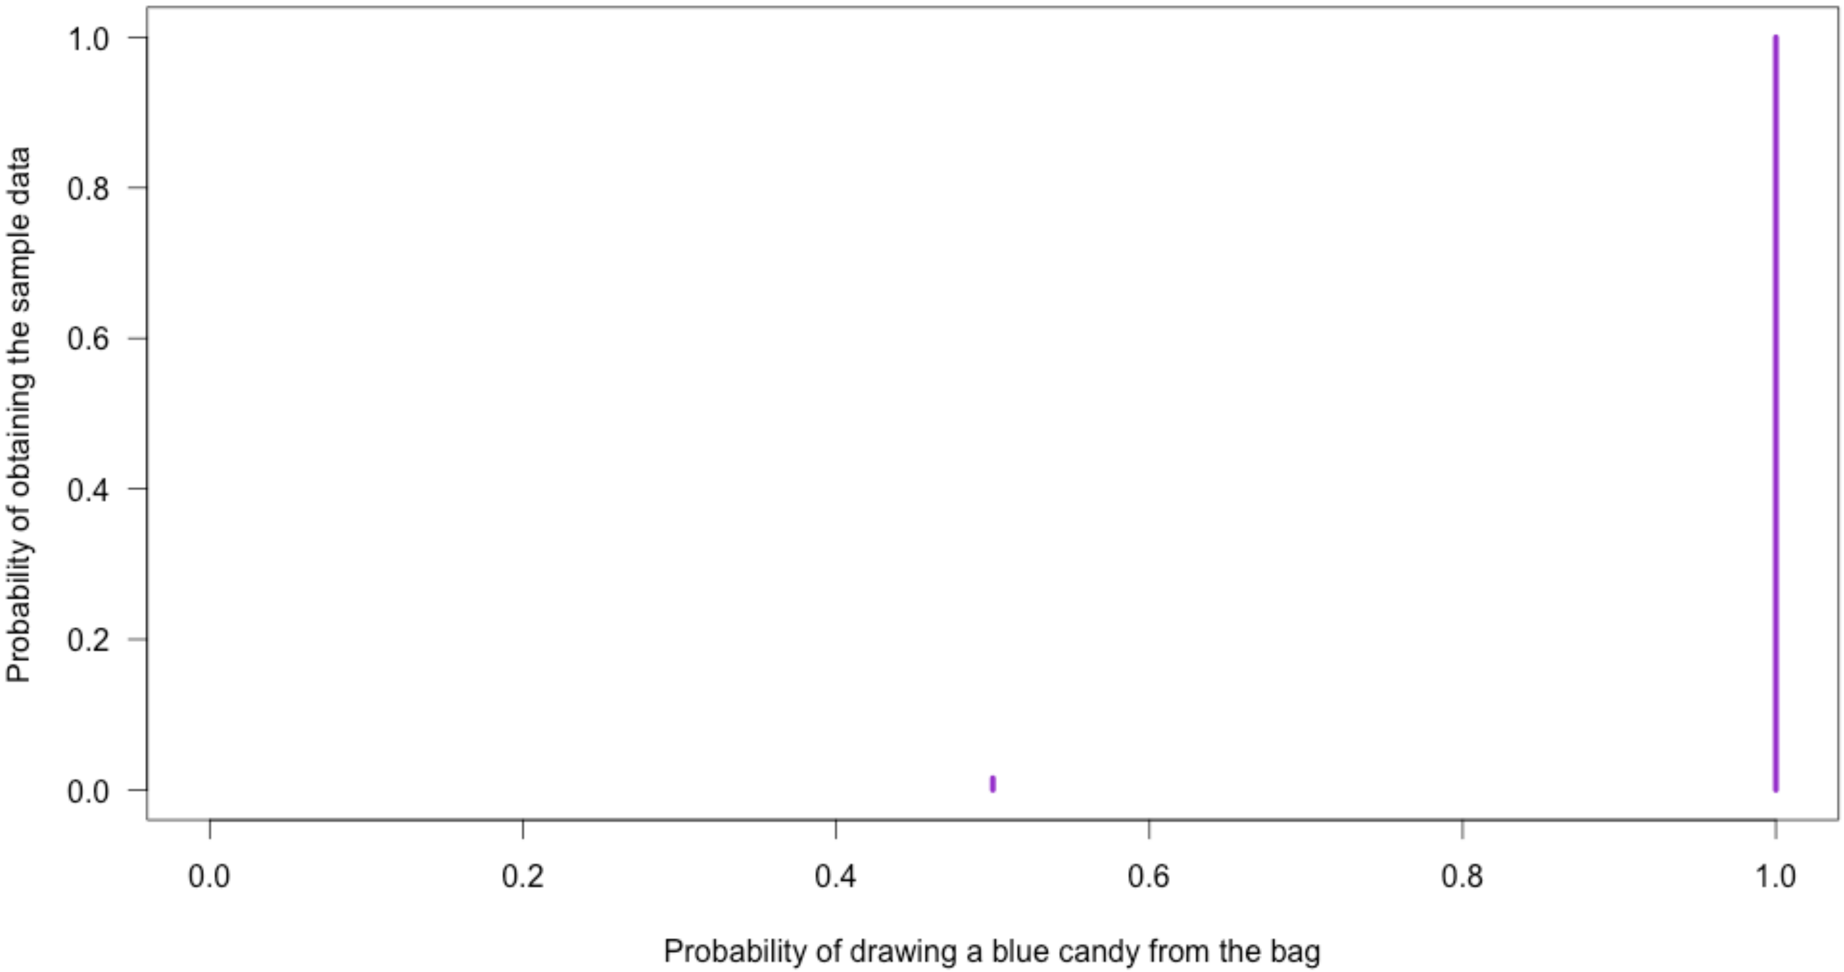

Graphing the evidence (6 **blue**)

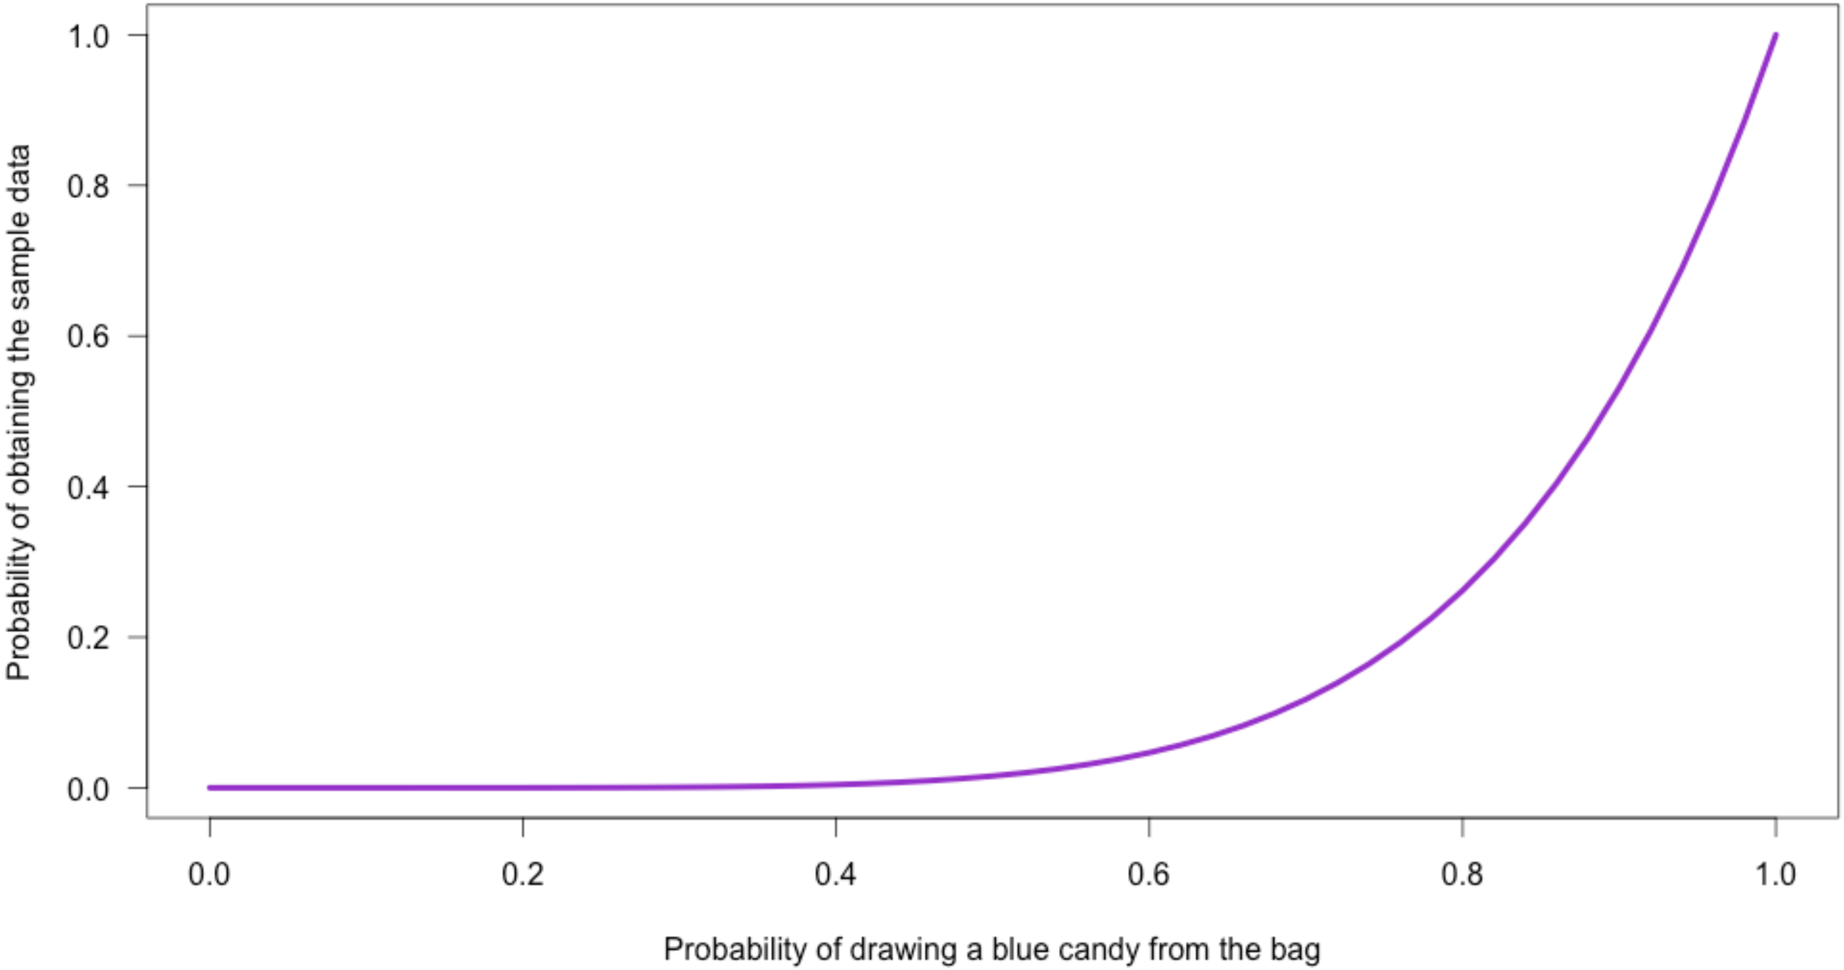

Graphing the evidence

- This is called a *Likelihood function*

- Ranks probability of the observations for all possible candy bag proportions

- Evidence is the ratio of heights on the curve
  - A above B, evidence for A over B

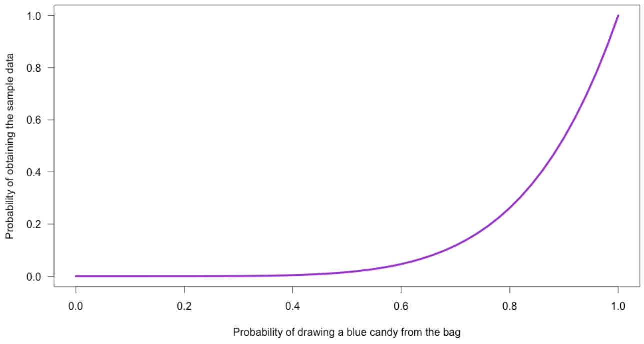

Graphing the evidence

- Where does prior information enter?
  - Prior rankings for each possibility

- Just as it did before
  - But now as a prior distribution

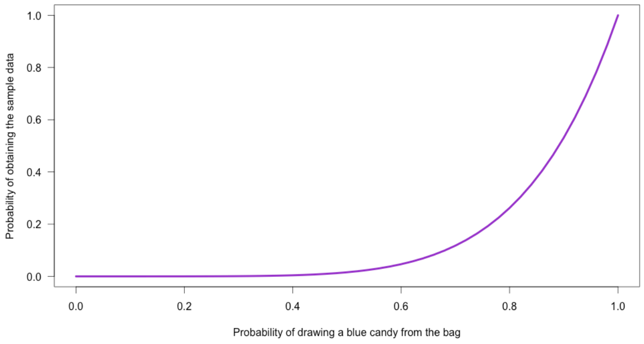

Prior information

- “Non-informative” prior information
  - All possibilities ranked equally
  - i.e. no value preferred over another

- Weak prior information; vague knowledge
  - “The bag has some blue candy, but not all blue candy”
  - After Halloween, for example
    - Saw some blue candy given out, but also other candies

- Strong prior information
  - “Proportion of women in the population is between 40% and 60%”

Non-informative

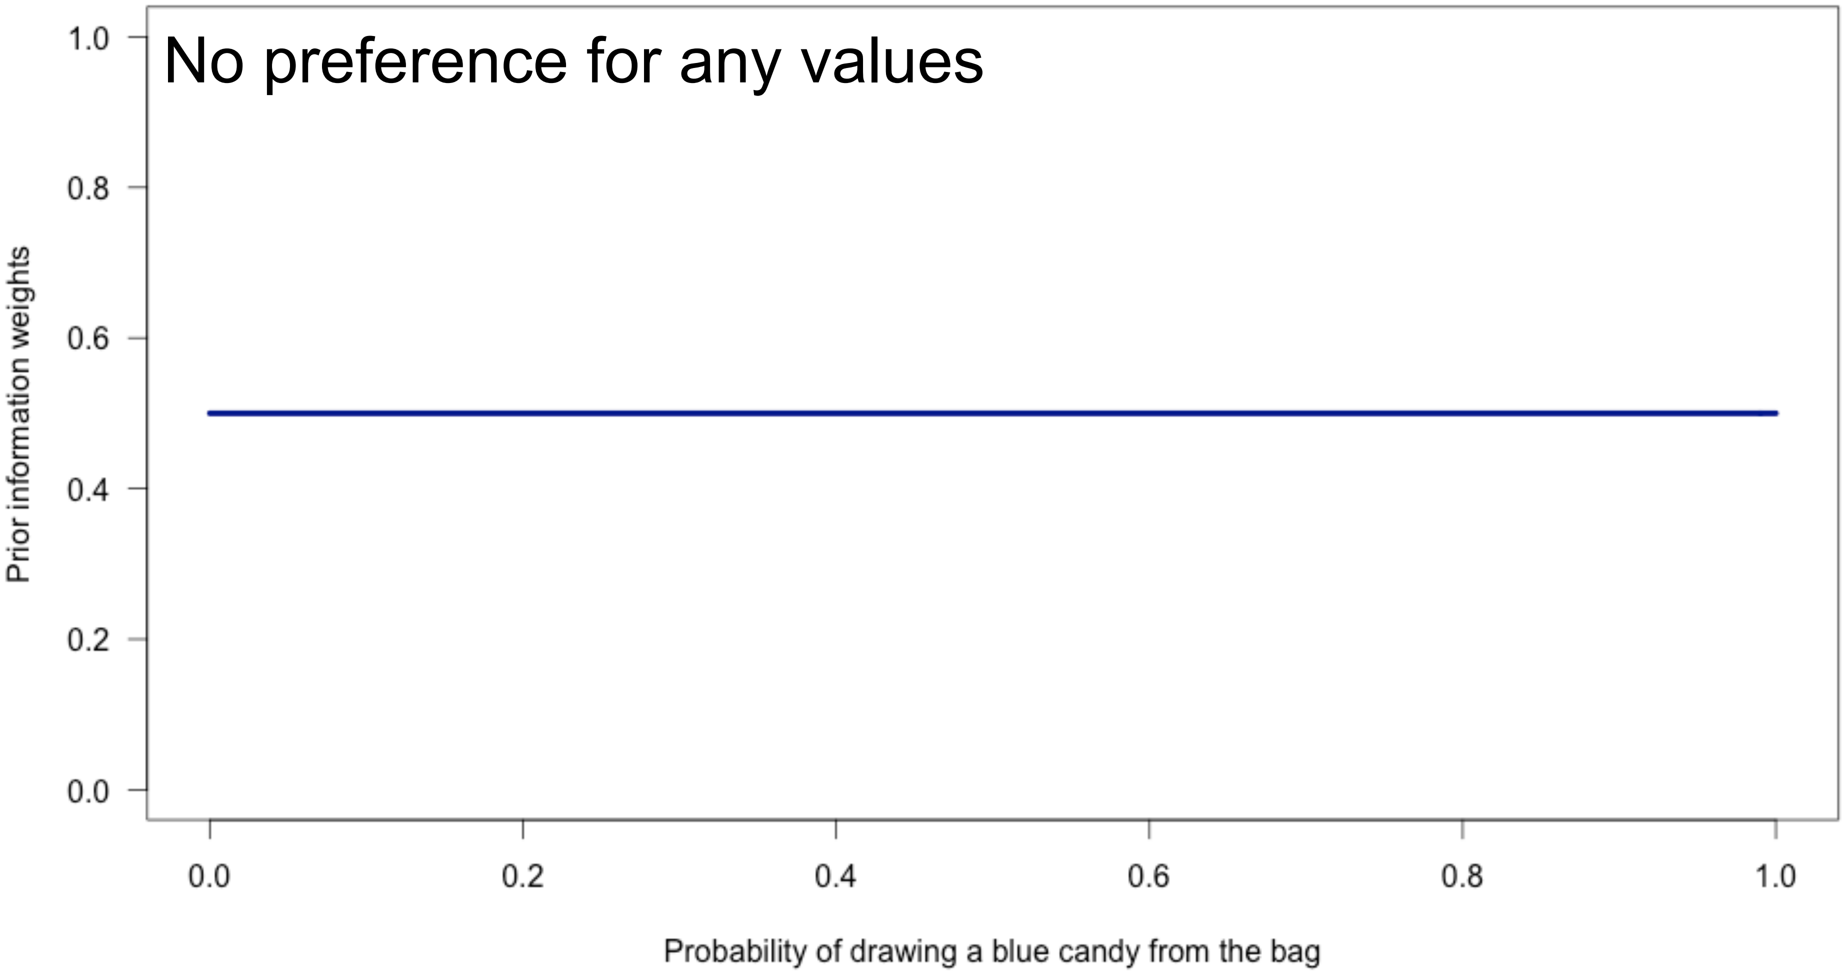

Weakly-informative

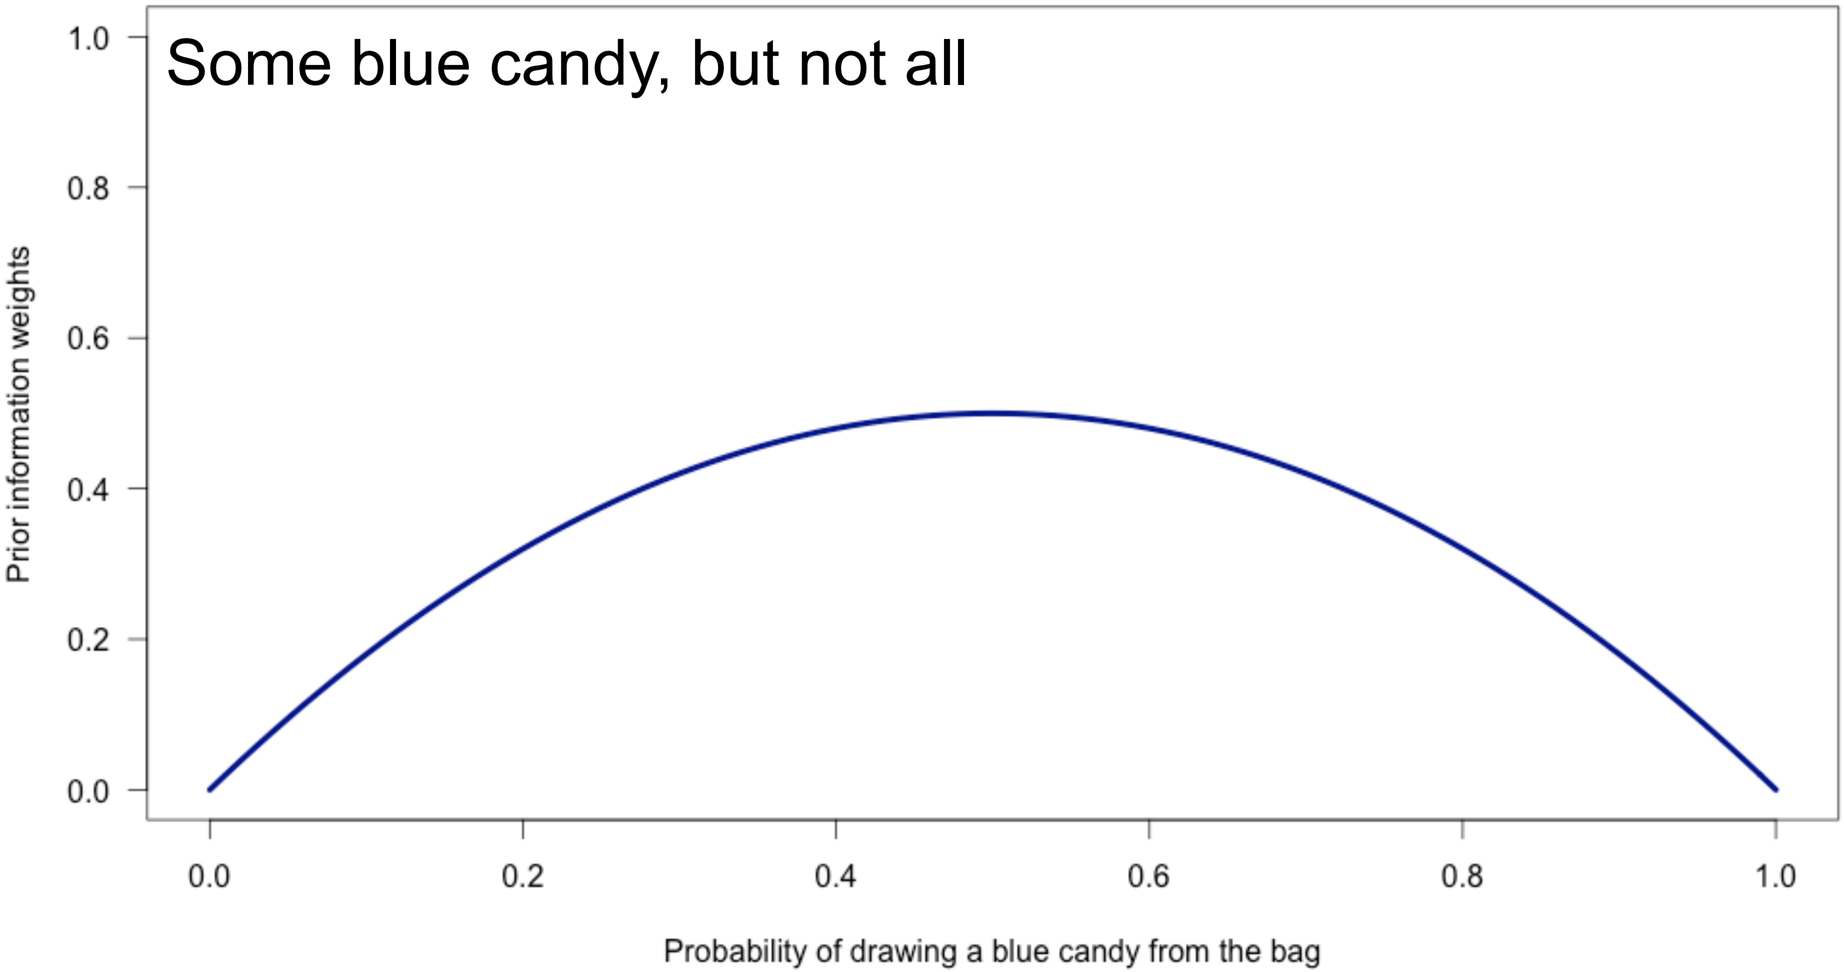

Strongly-informative

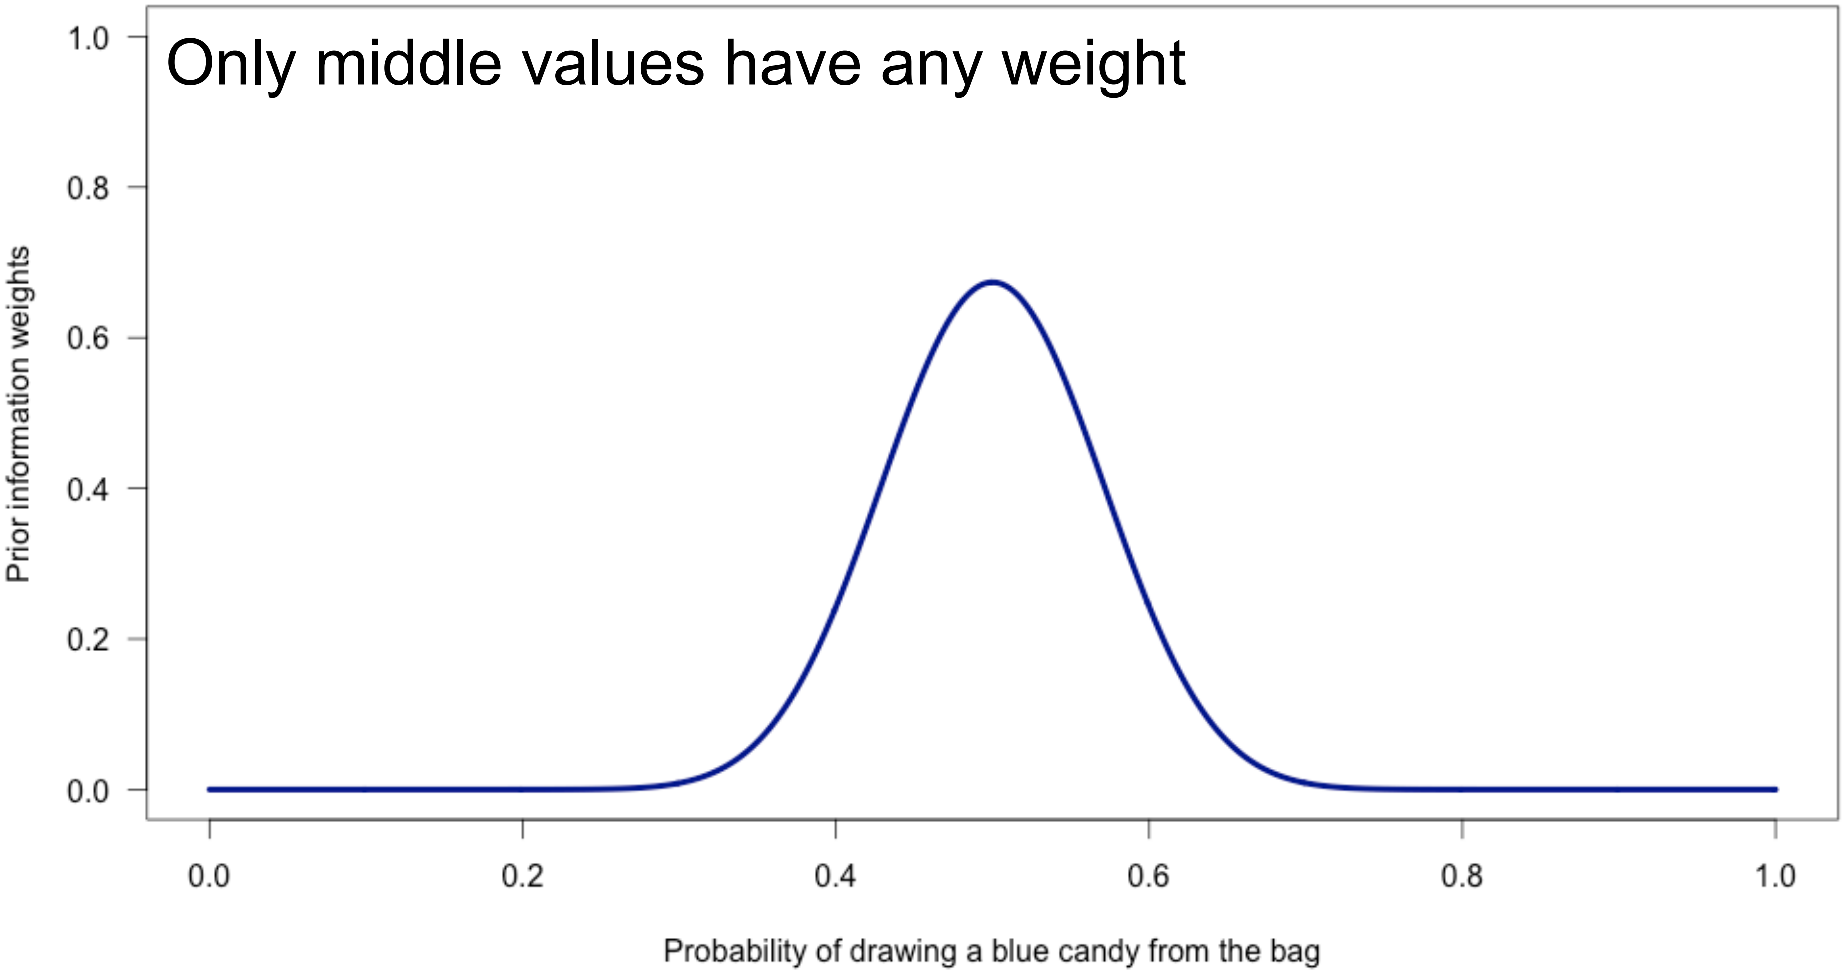

Information and context

- Your prior information depends on context!
  - And depends on what you know!
- Just like drawing cards in the game
  - Just harder to specify
  - Intuitive, personal
- Conclusions must take context into account