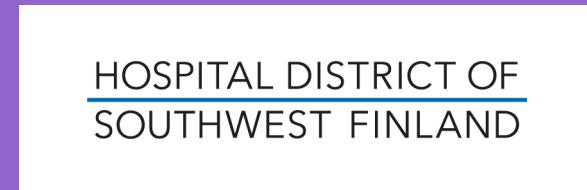


Bayesian statistical inference in neuroimaging

Turku PET Centre Neuroimaging Course, September 2022

Tuulia Malen

PhD student, psychologist



Mission

- Statistical inference
- Bayesian statistical inference
- Bayesian inference in **neuroimaging**
 - Example analysis with PET data

Statistical inference

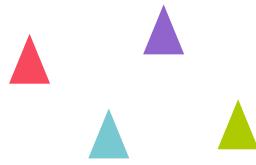
Statistical inference

?

Statistical inference

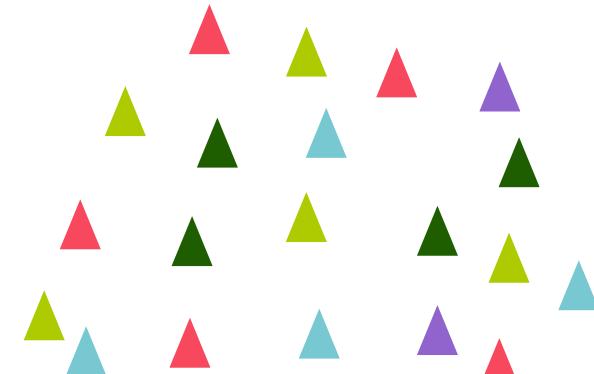


Statistical inference



Sample

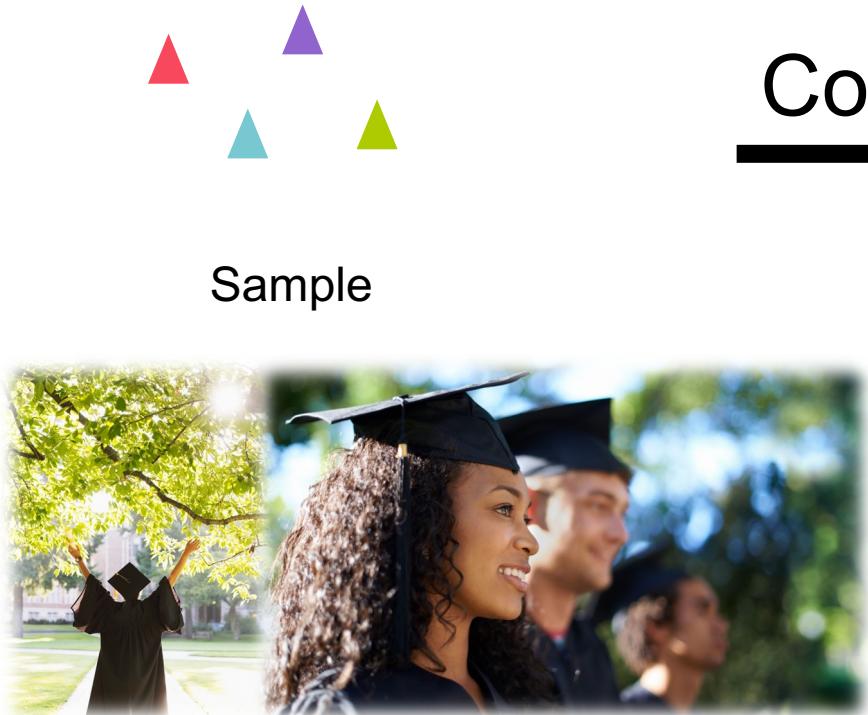
Conclusions



Population

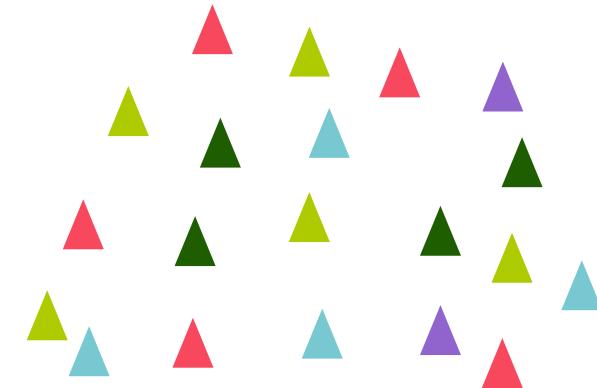


Statistical inference



Sample

Conclusions

A thick black horizontal arrow pointing to the right, indicating the direction of statistical inference from the sample to the population.

Population



Statistical inference: frameworks

Frequentist

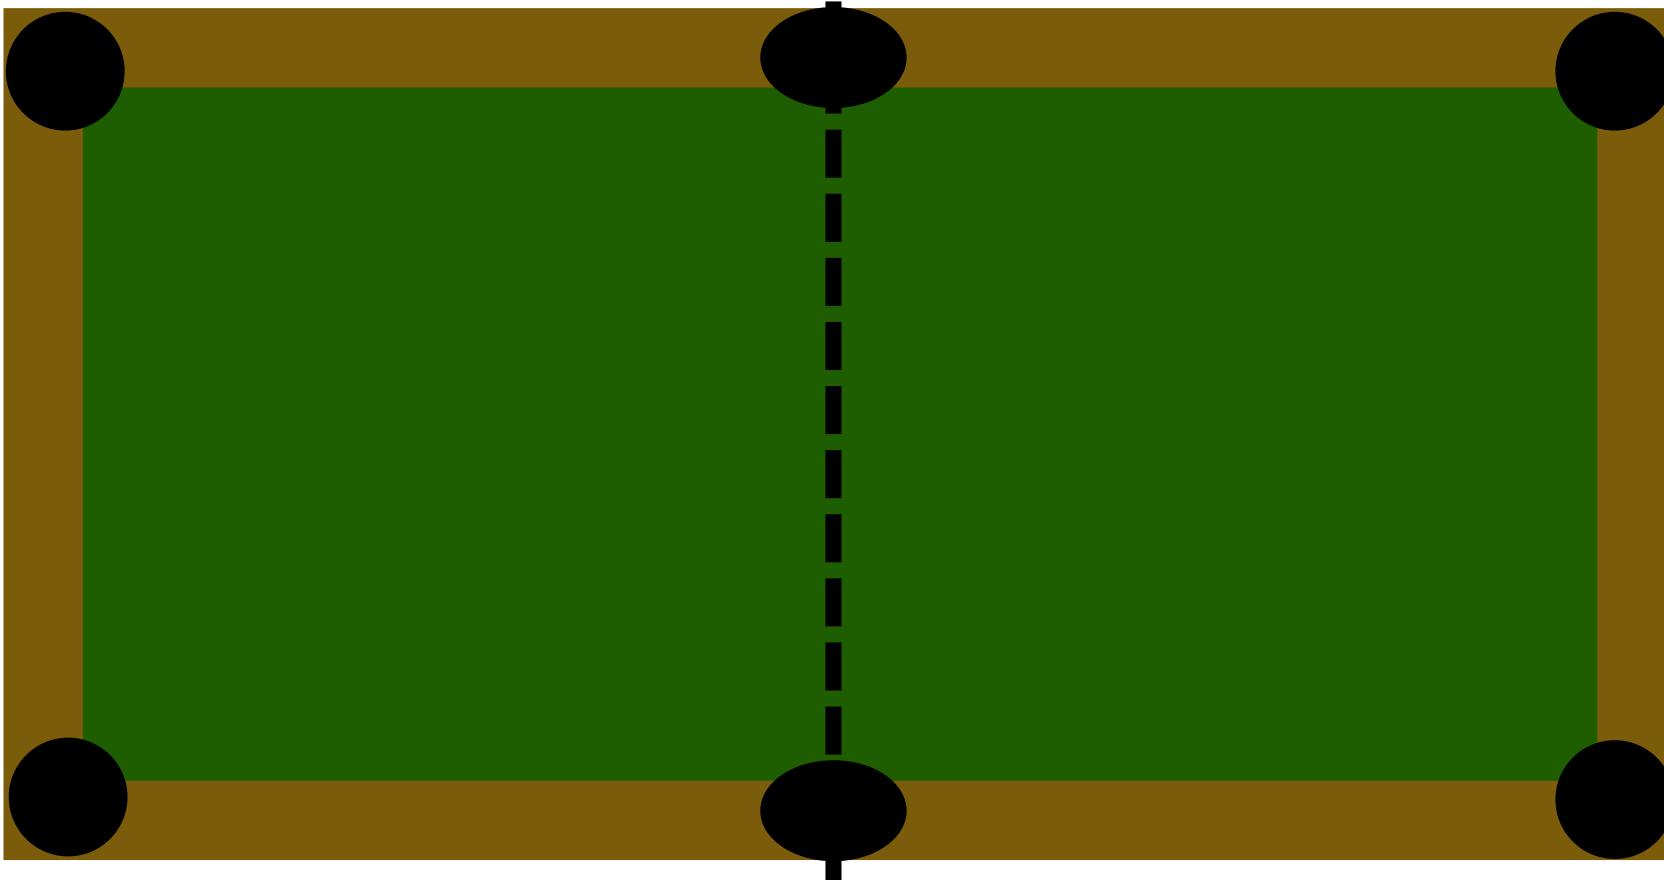
- Null hypothesis testing
- t-test
- p-value
- multiple comparison correction
- confidence interval

Bayesian

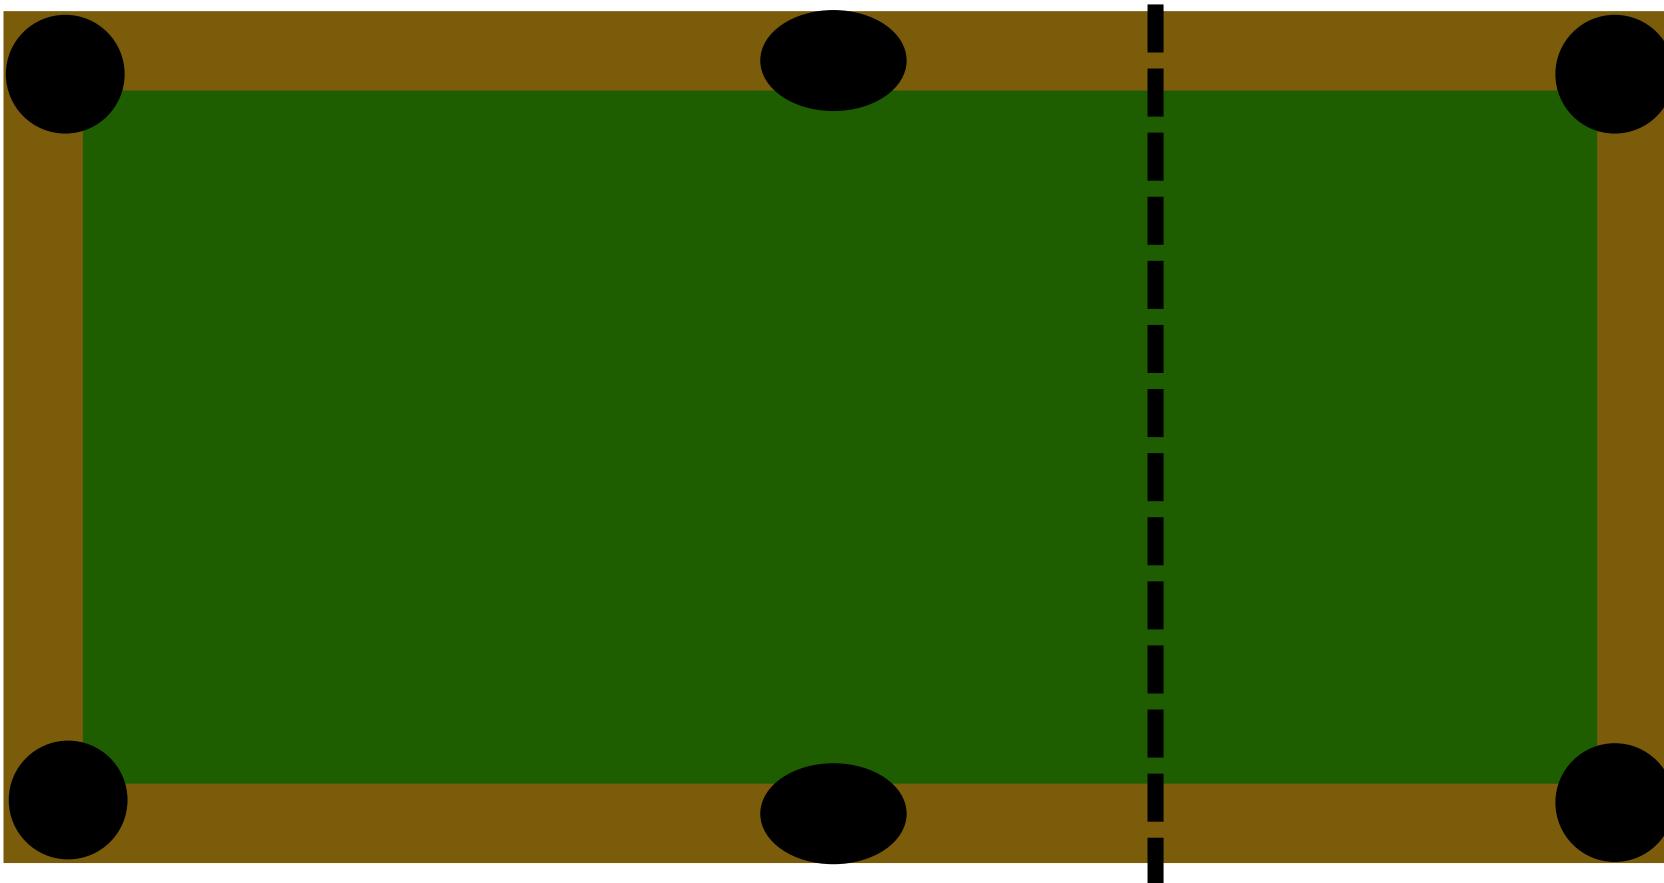
Bayesian statistical inference

Take a mindset of a philosopher...

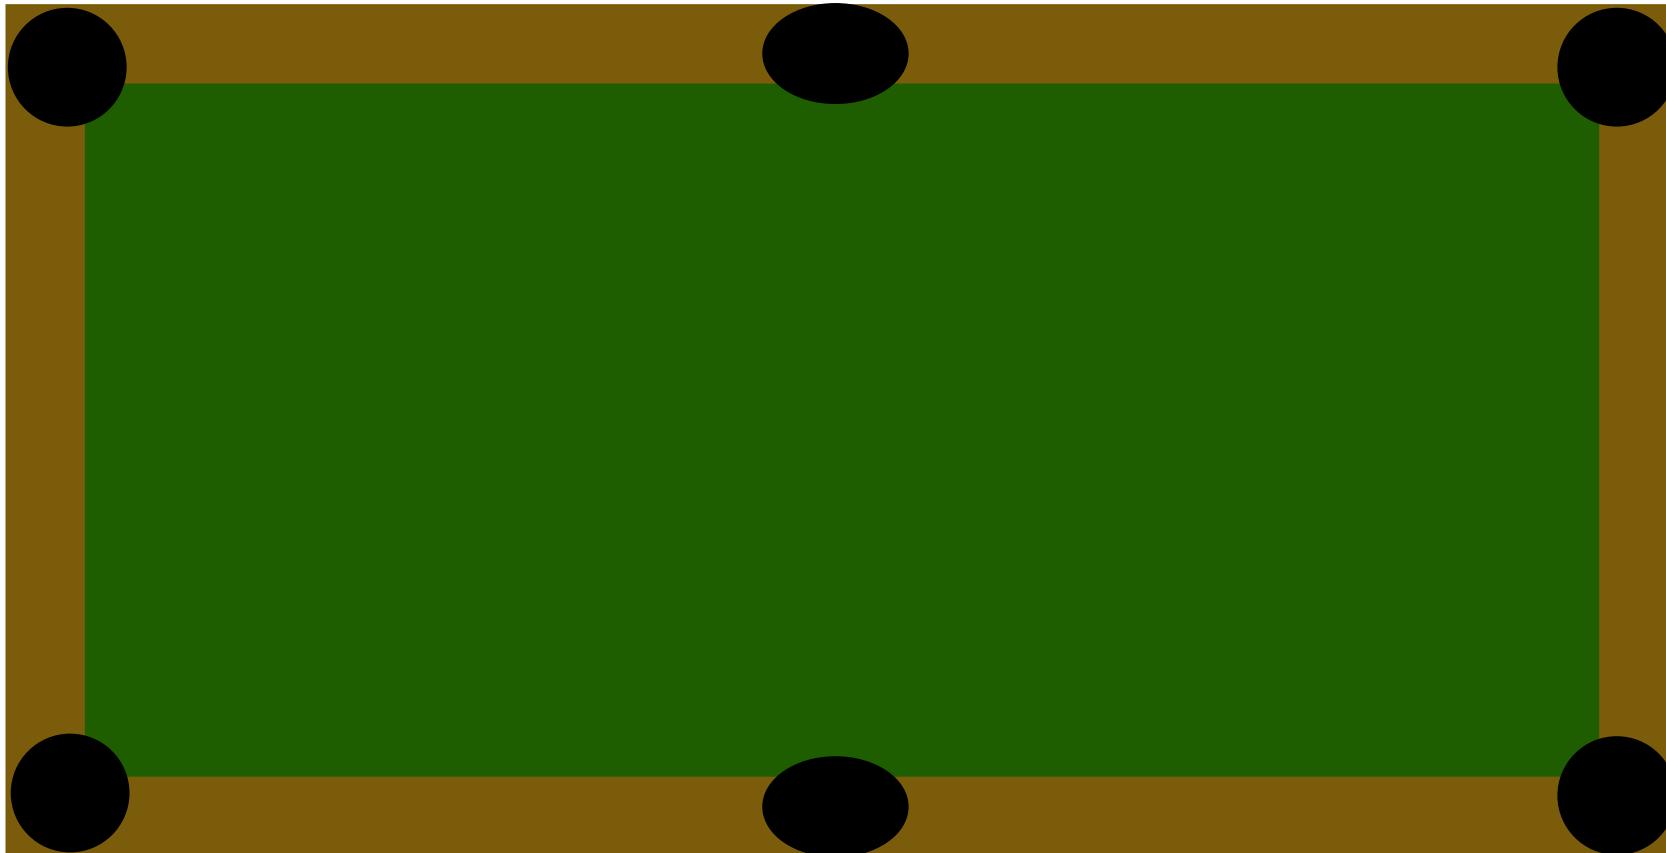
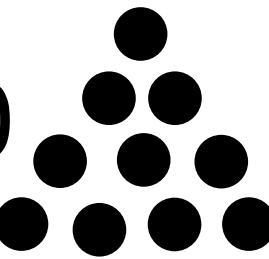
Thomas Bayes



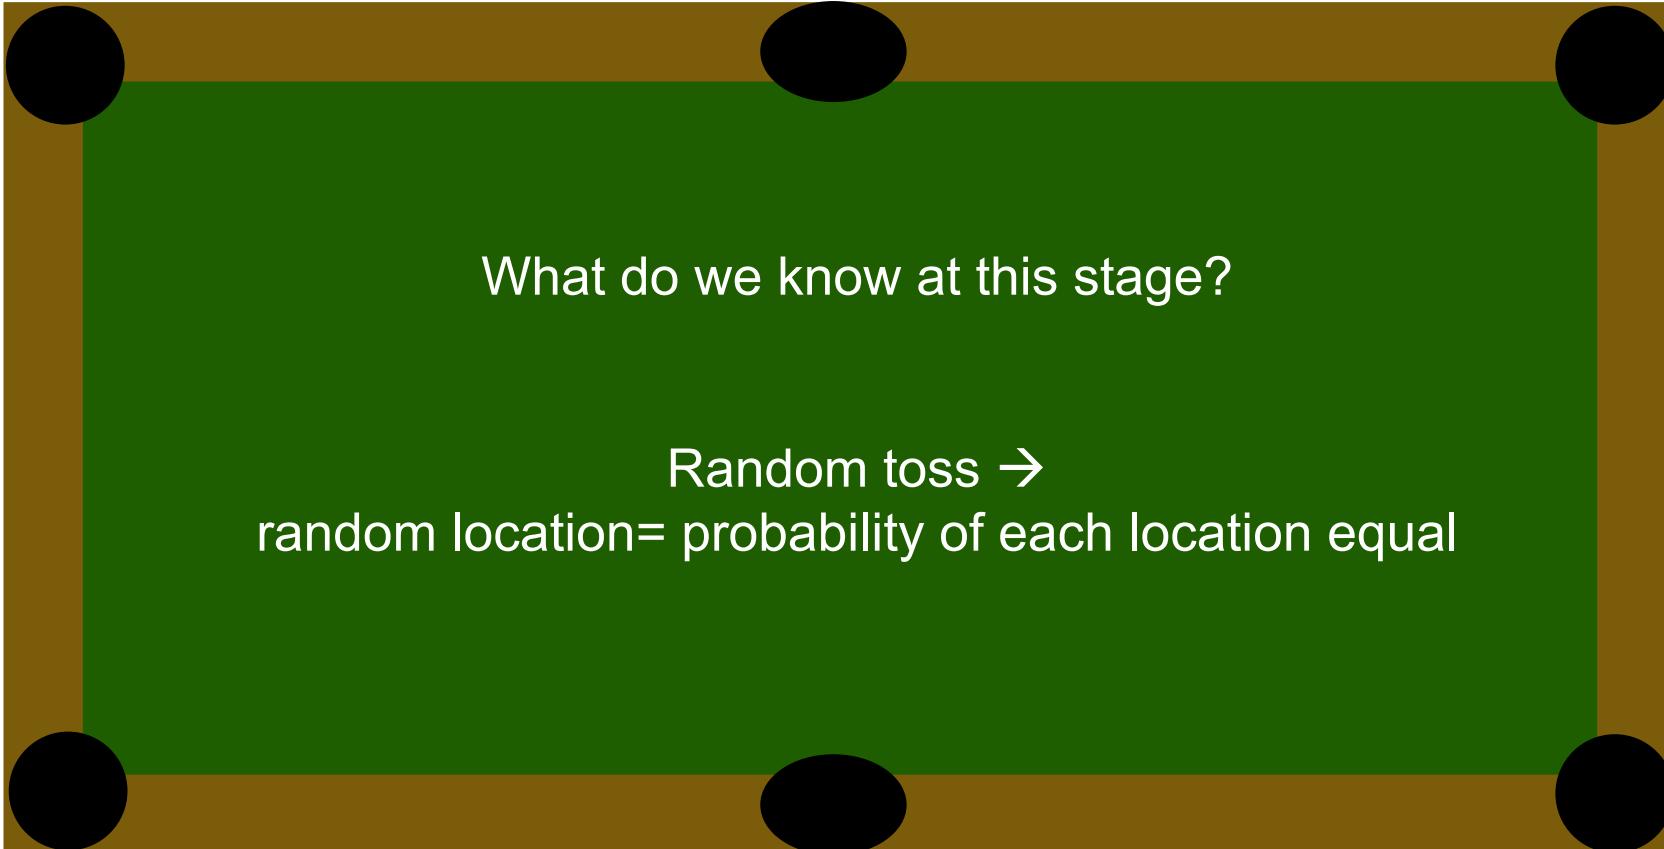
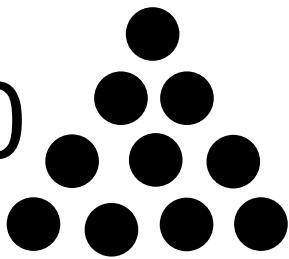
Thomas Bayes

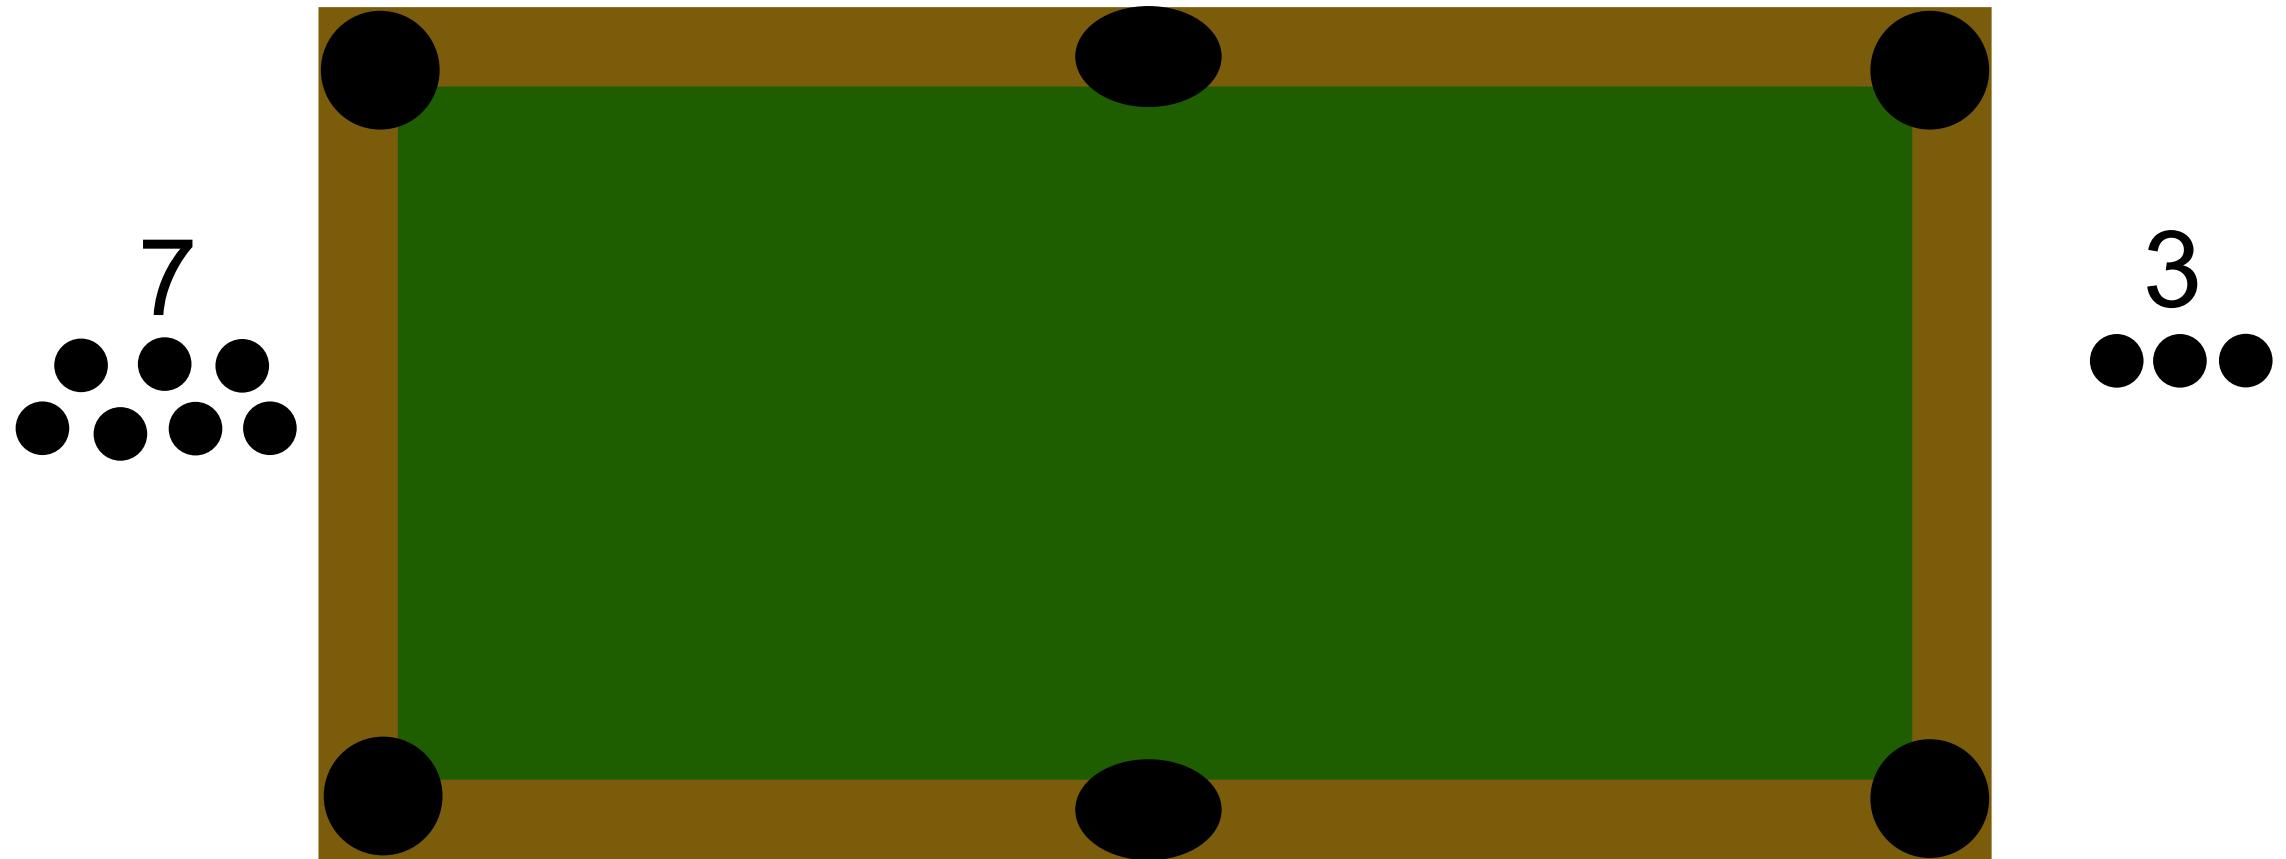


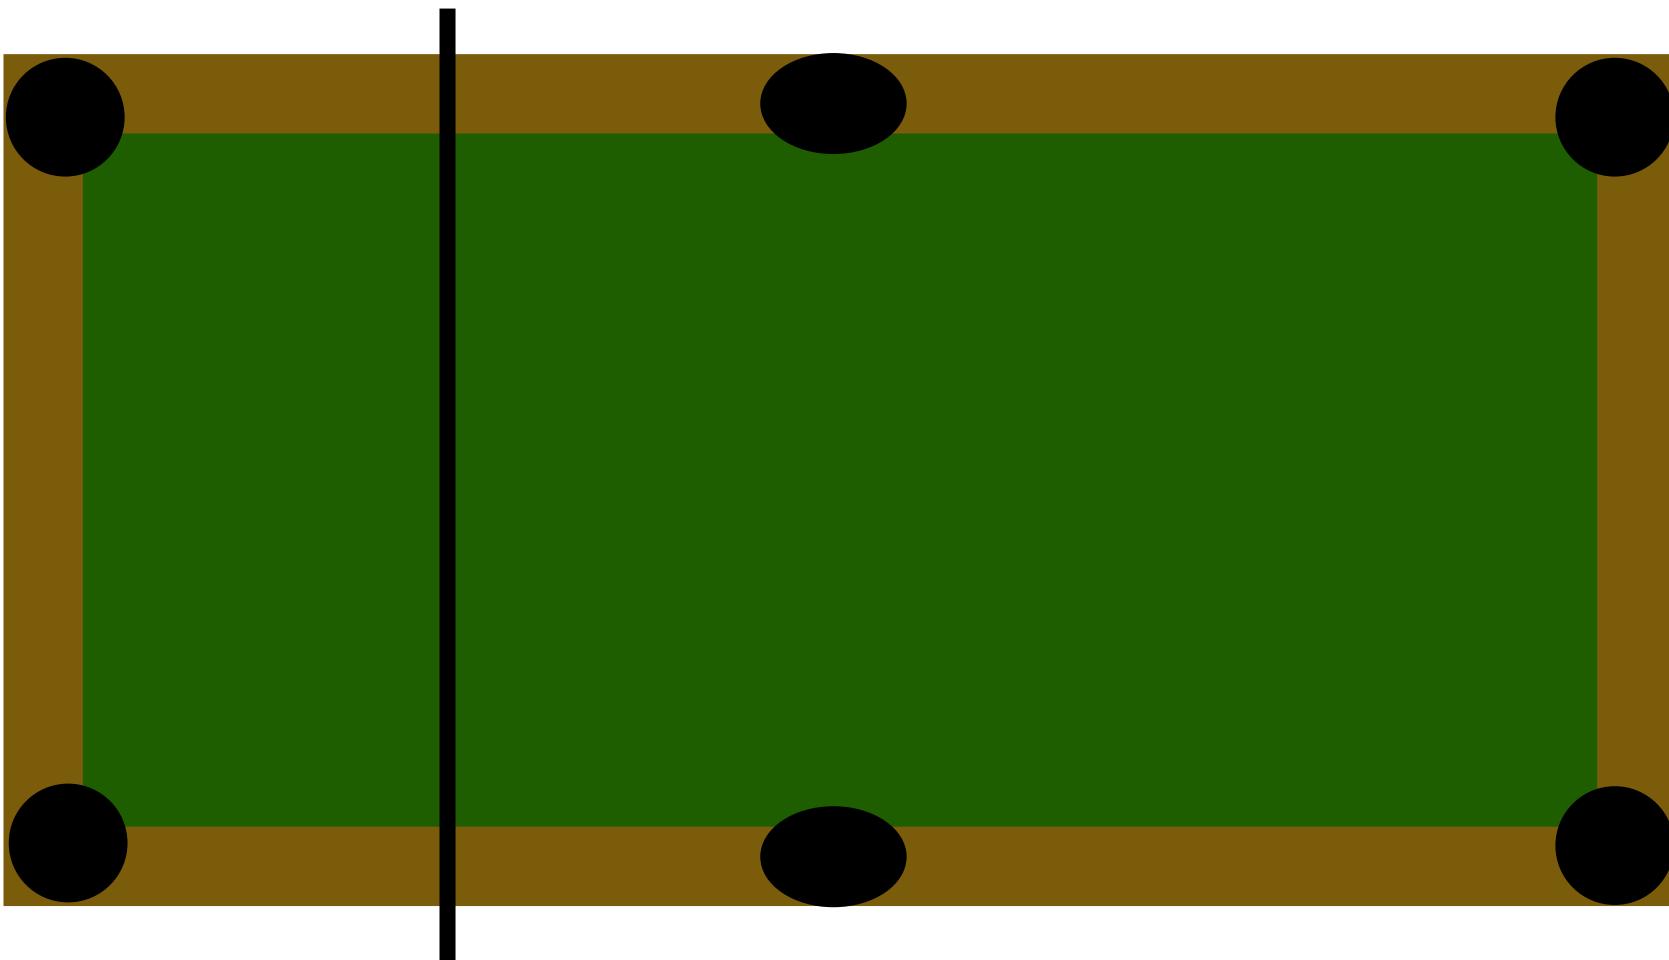
10

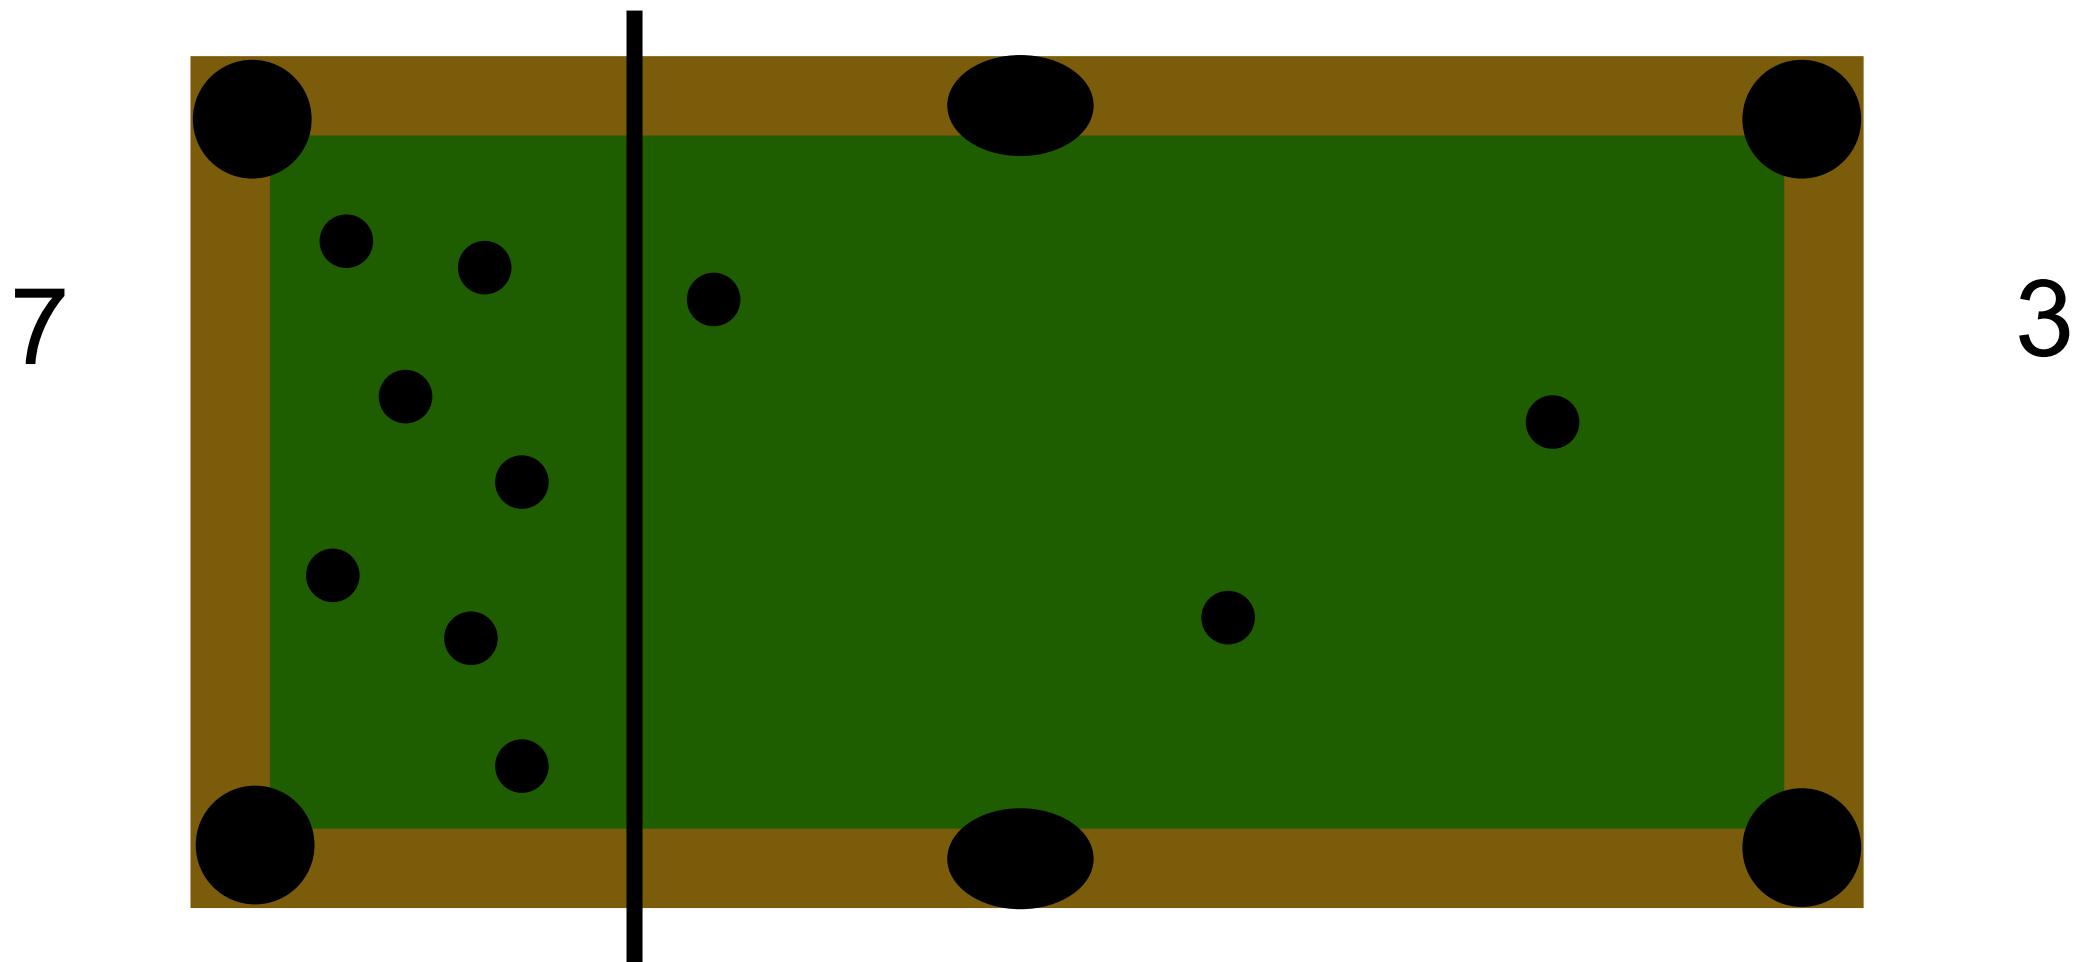


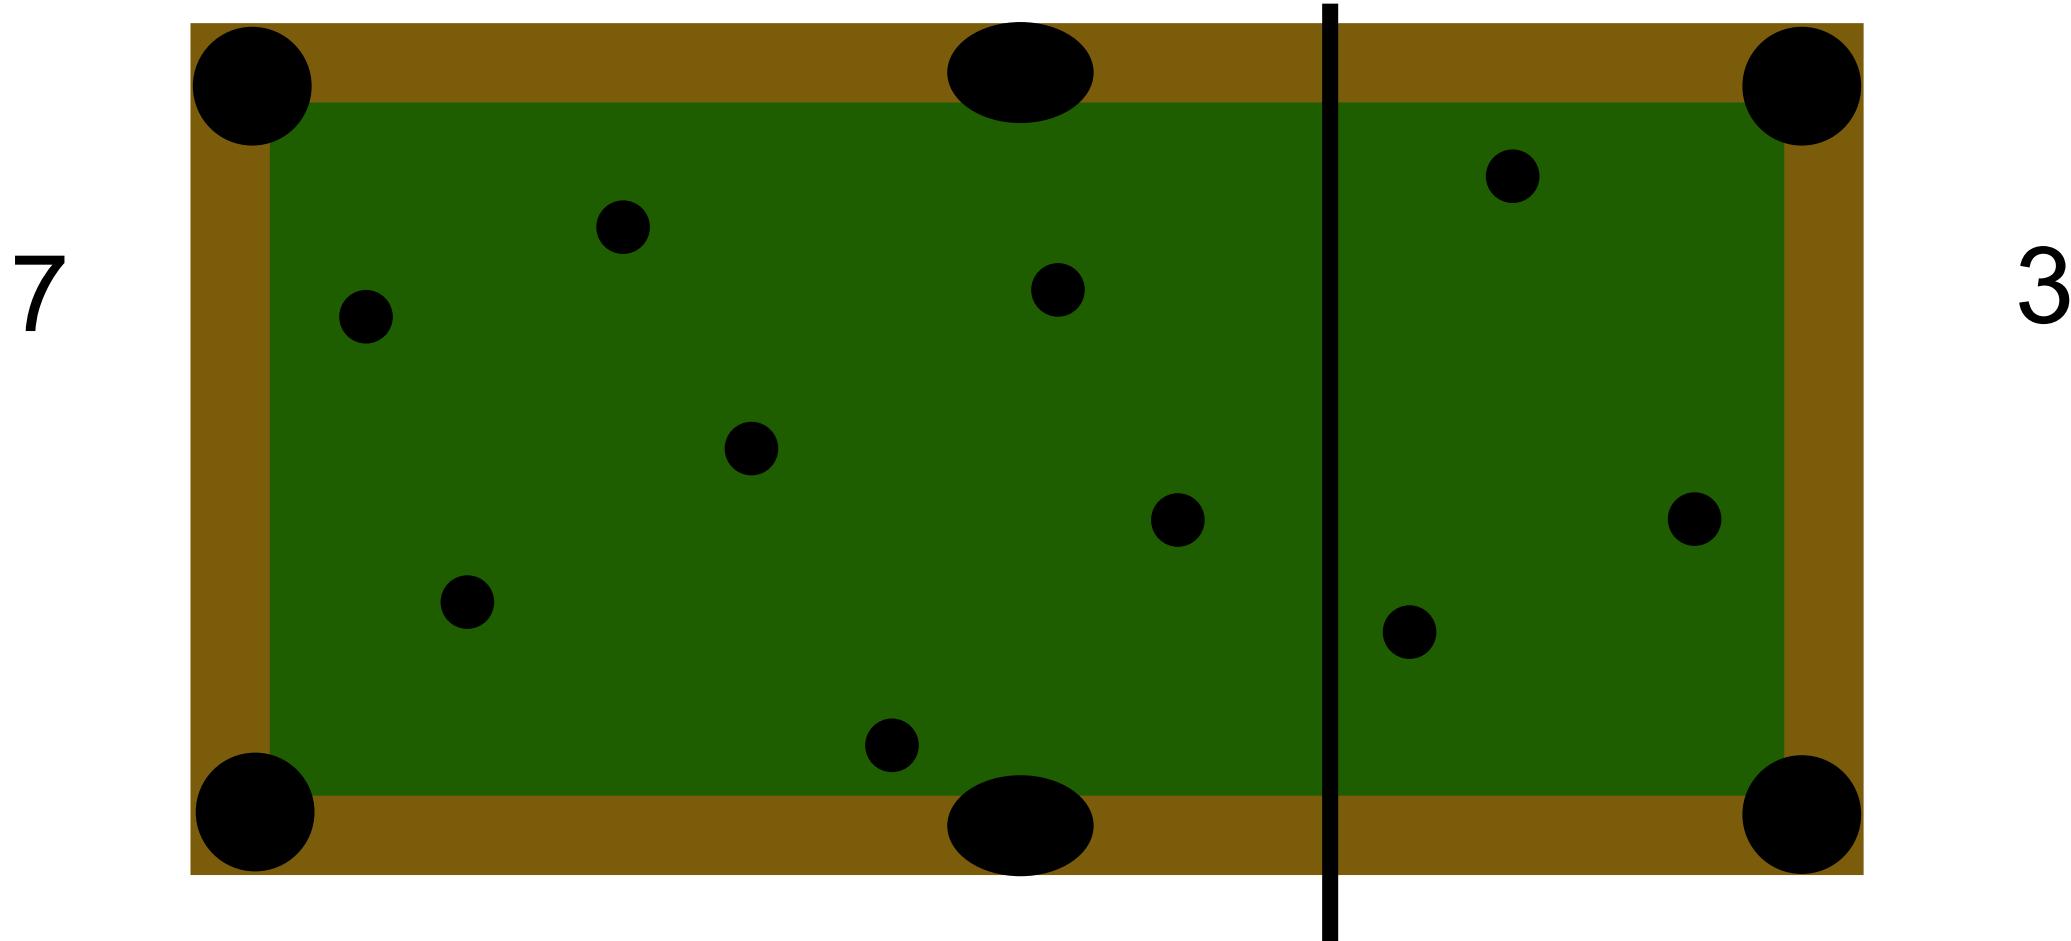
10



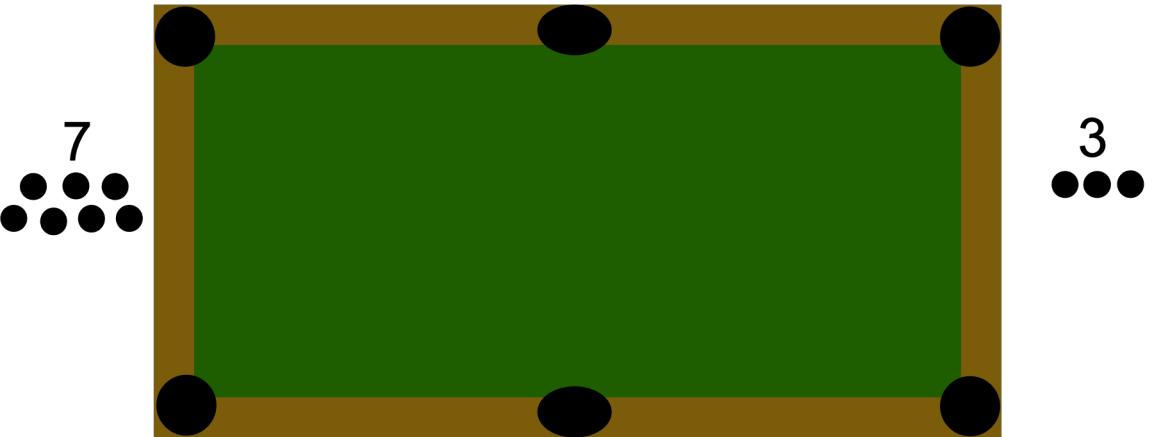








Where the line?



1) What do we know **prior** to data?

- Each ball gets a random position
- All positions equally probable

2) What is our **data**?

- 7 on the left
- 3 on the right

Both the **prior** knowledge and the **data** affect our **posterior** understanding

Key

Prior & data → Conclusions of the parameter

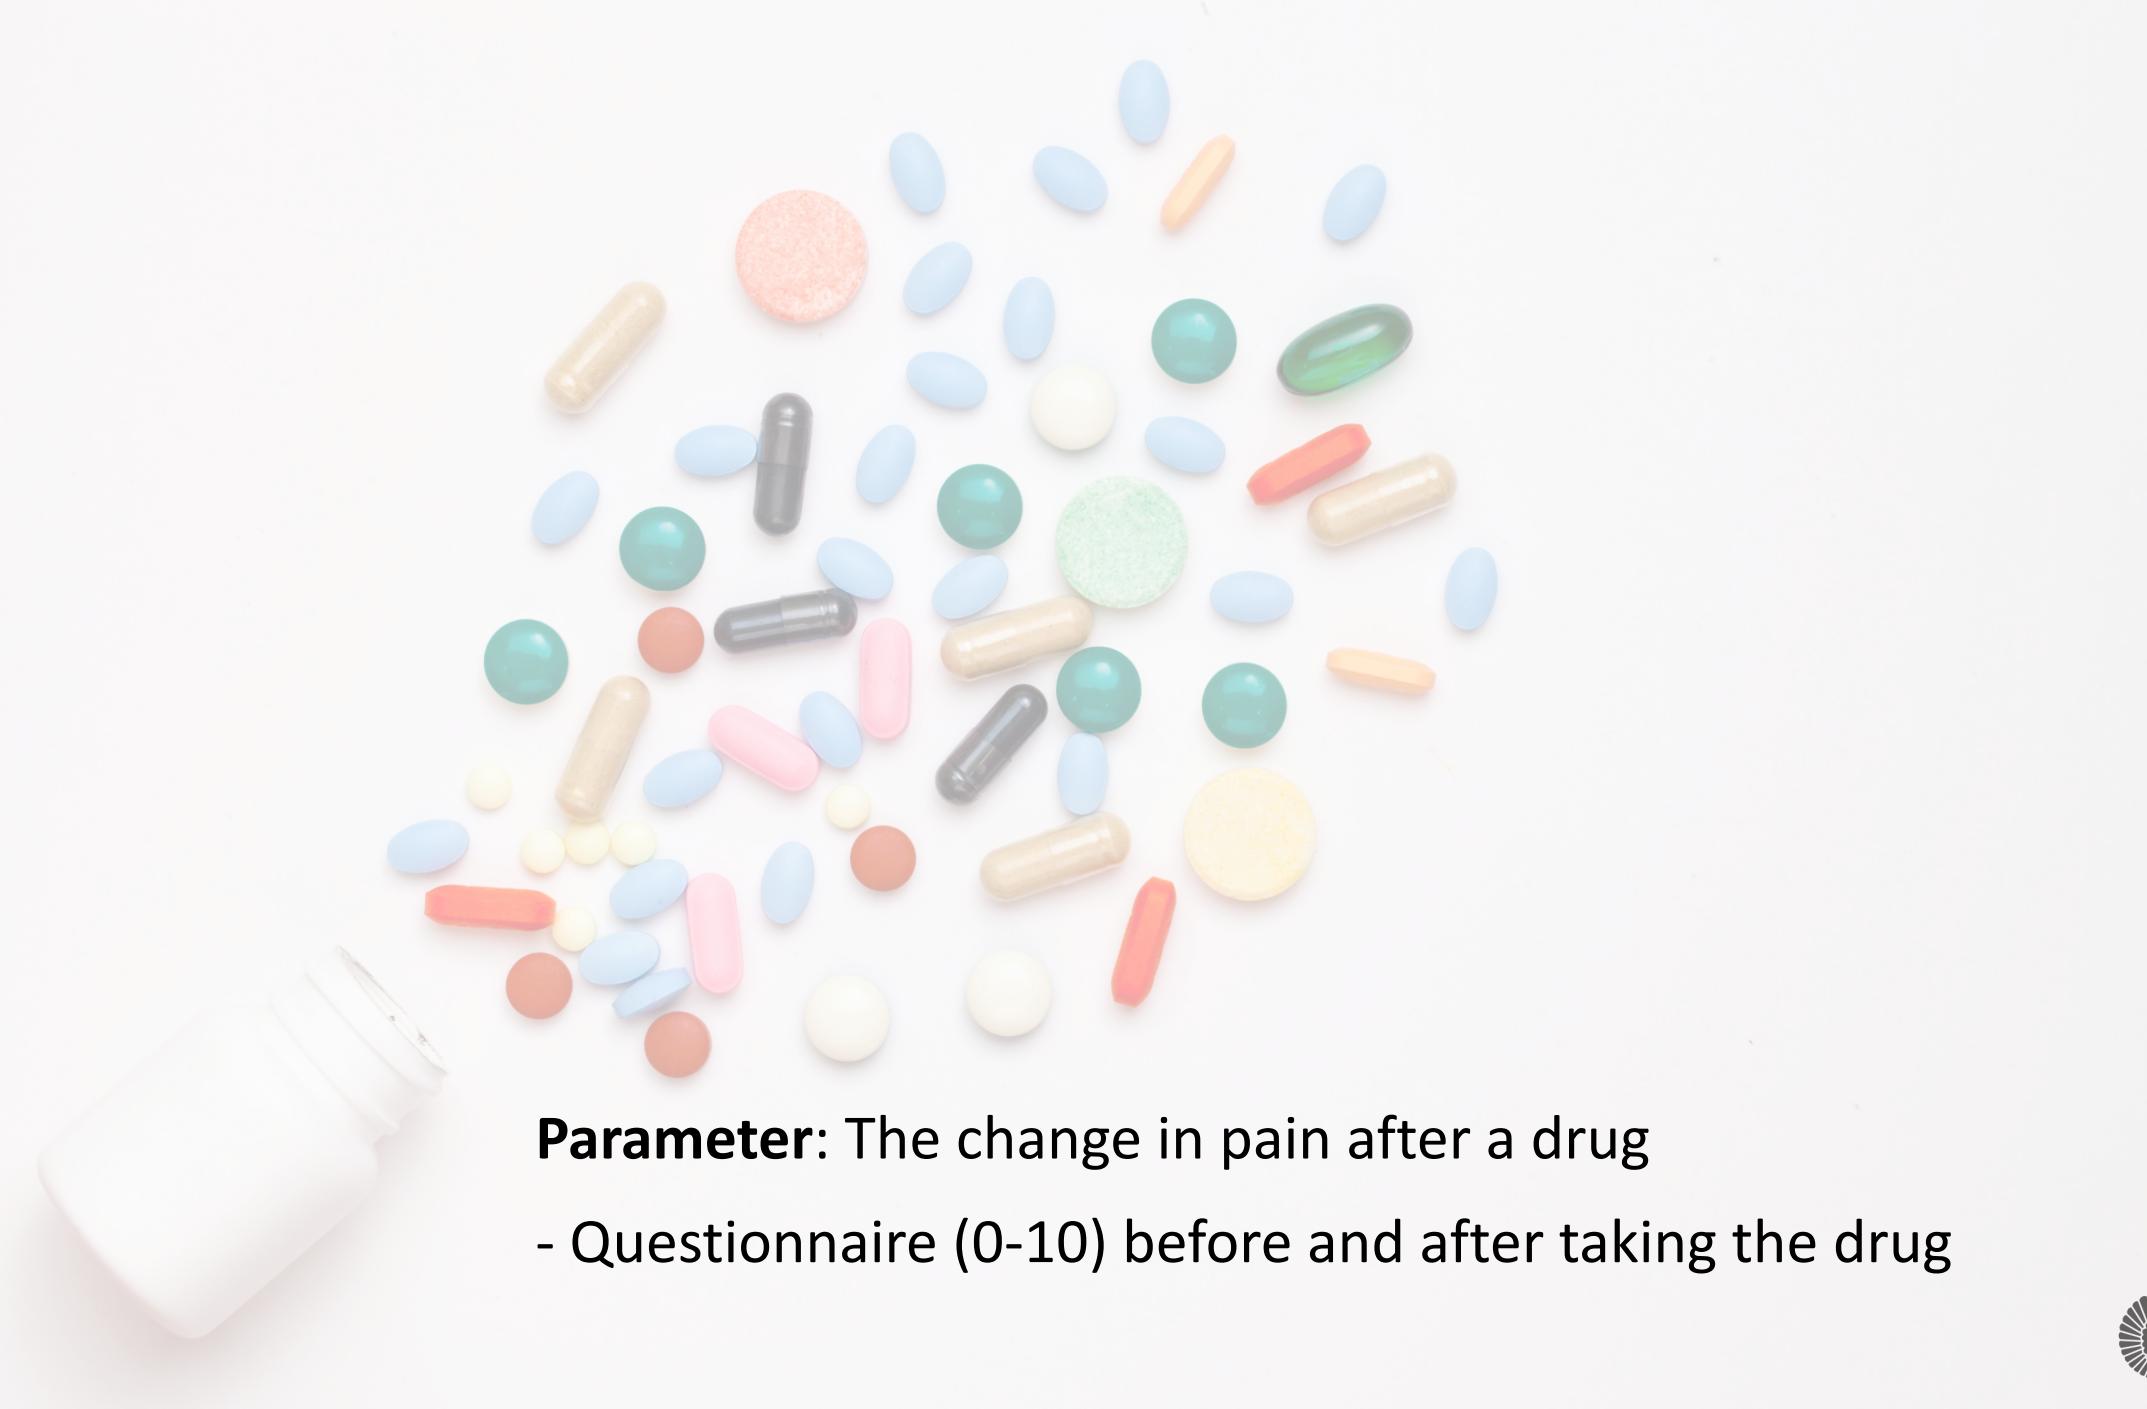
Key

Prior & data → Conclusions of the parameter

Key

Prior & data → Conclusions of the parameter

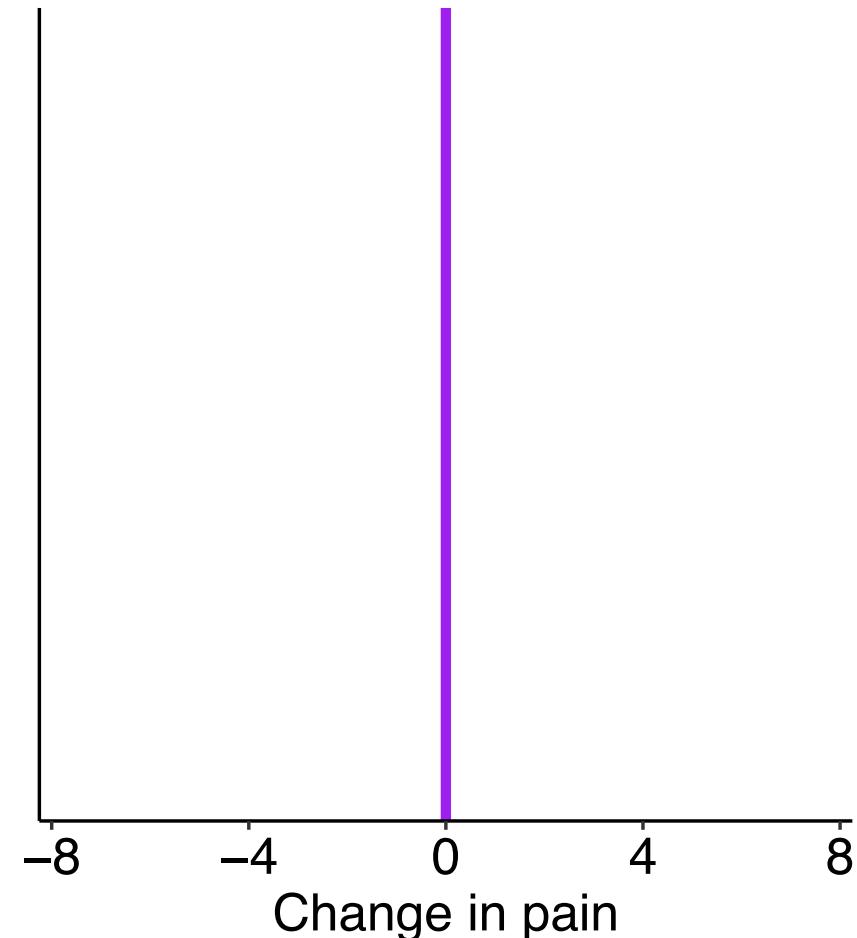




Parameter: The change in pain after a drug
- Questionnaire (0-10) before and after taking the drug

Frequentist

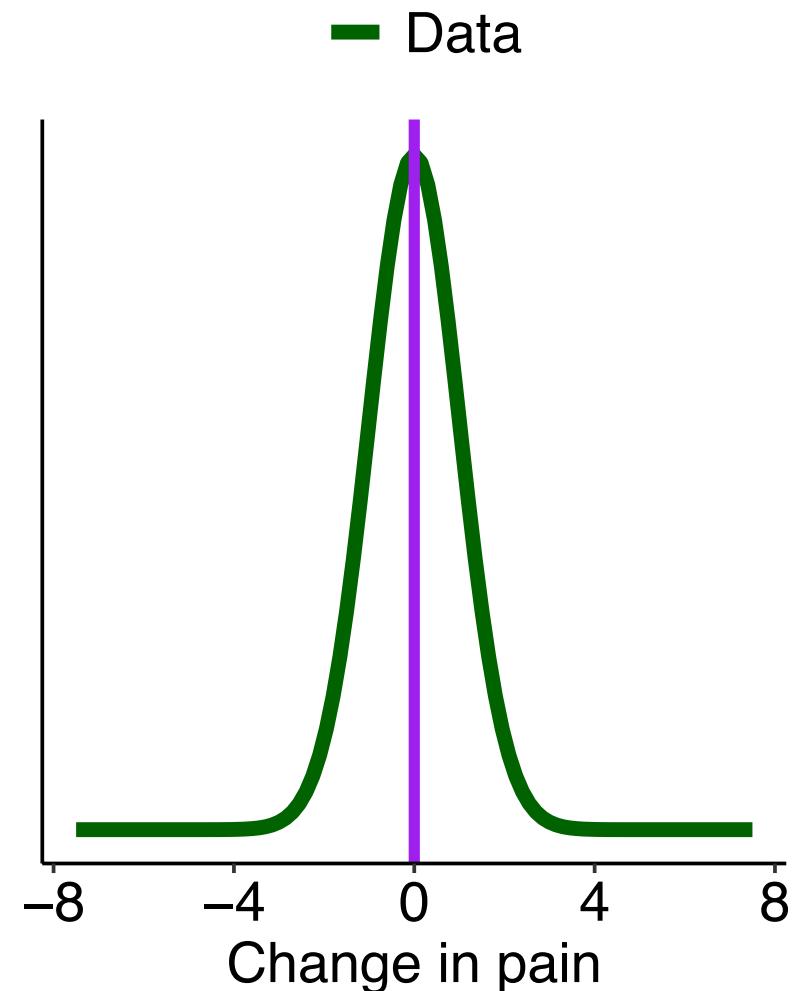
Null-hypothesis (H_0): fixed parameter value



Frequentist

Null-hypothesis (H_0): fixed parameter value

Data: observations from 100 subjects

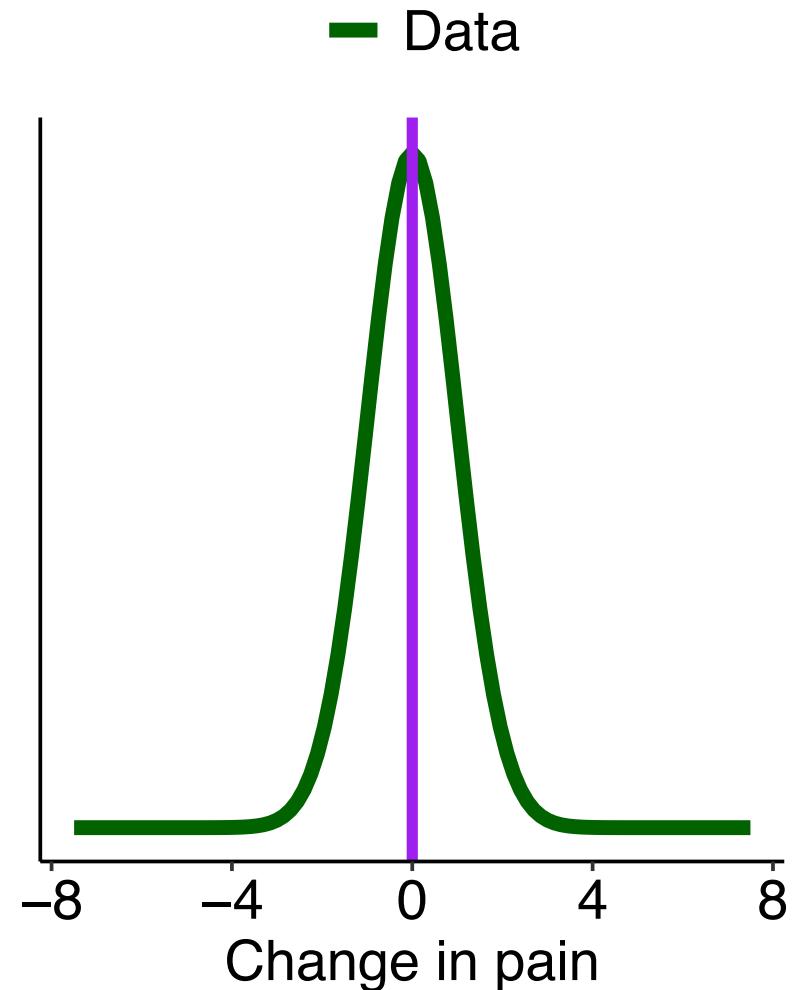


Frequentist

Hypothesis testing:

p-value close to 1= data well aligned with H_0

→ Support for H_0 (no change in the pain)



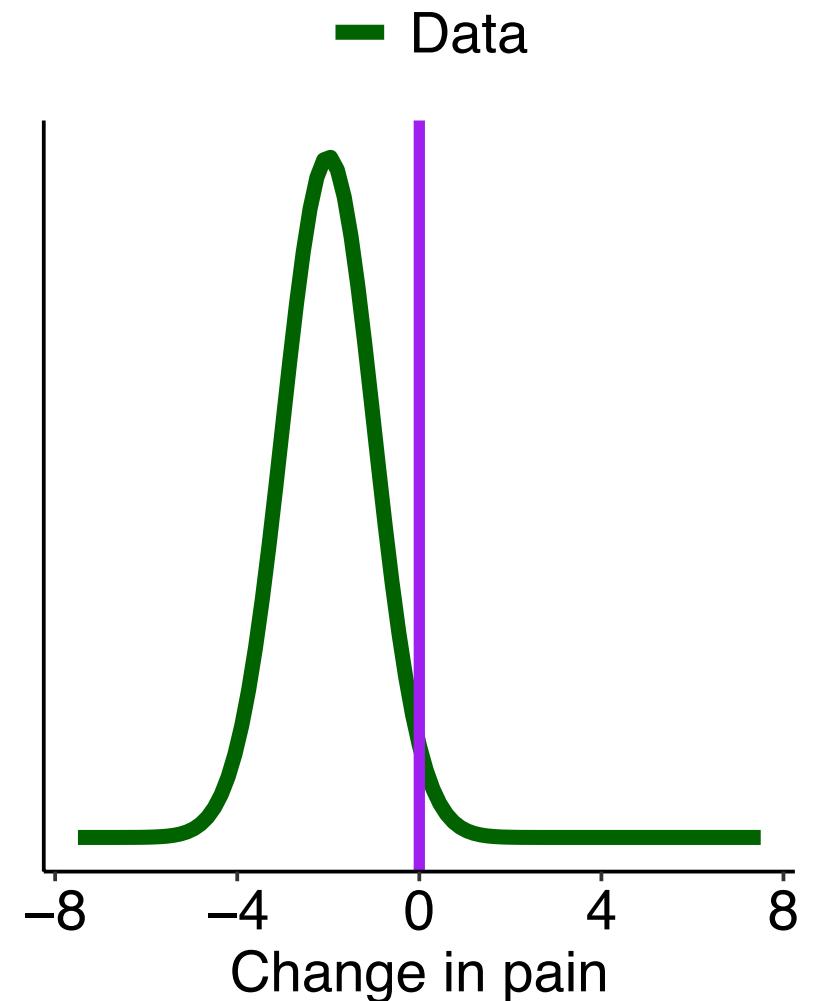
Frequentist

Hypothesis testing:

p-value small, close to 0 = data deviating from H_0

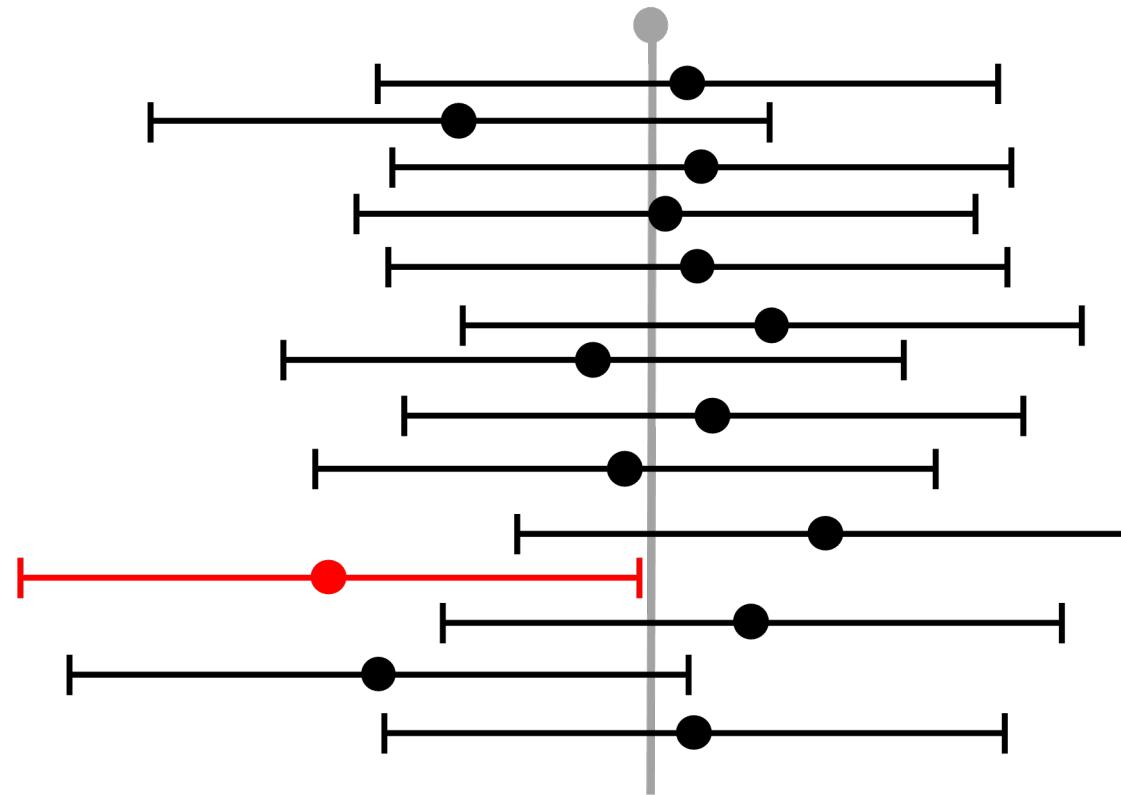
→ If H_0 is true, data rare*

*Sometimes the data is rare just by chance



Frequentist

*Sometimes the data is rare just by chance



Frequentist

How does our *data* look like?



What if...

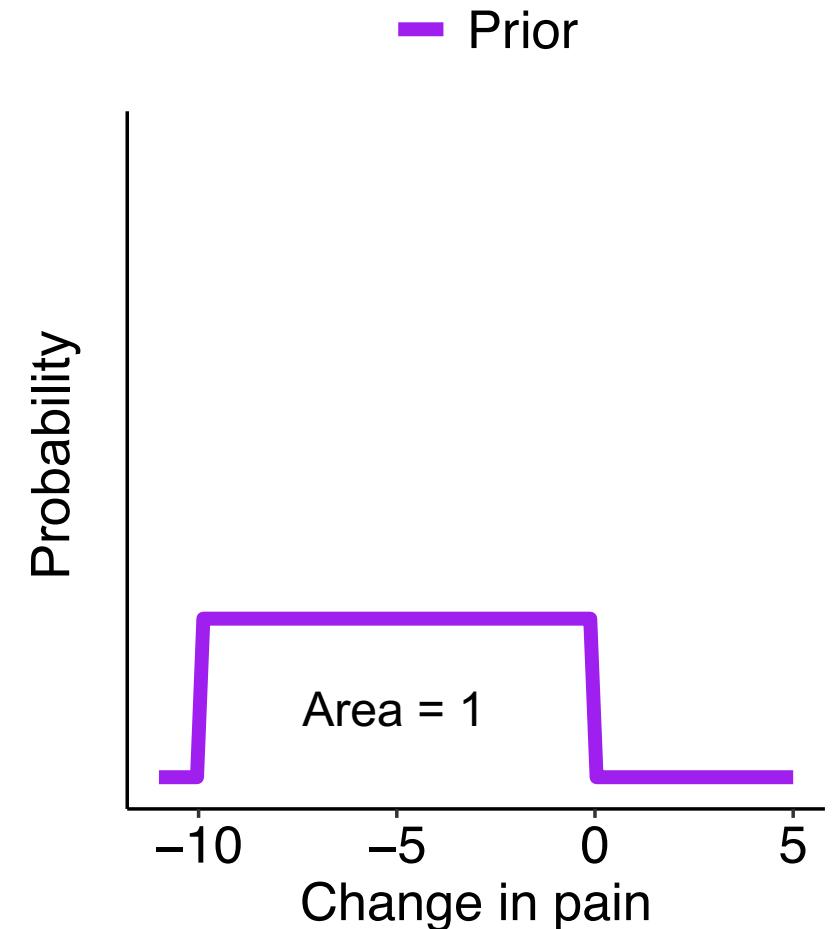
... we did not have to fix the hypothesis (parameter, the change)?

Bayesian:

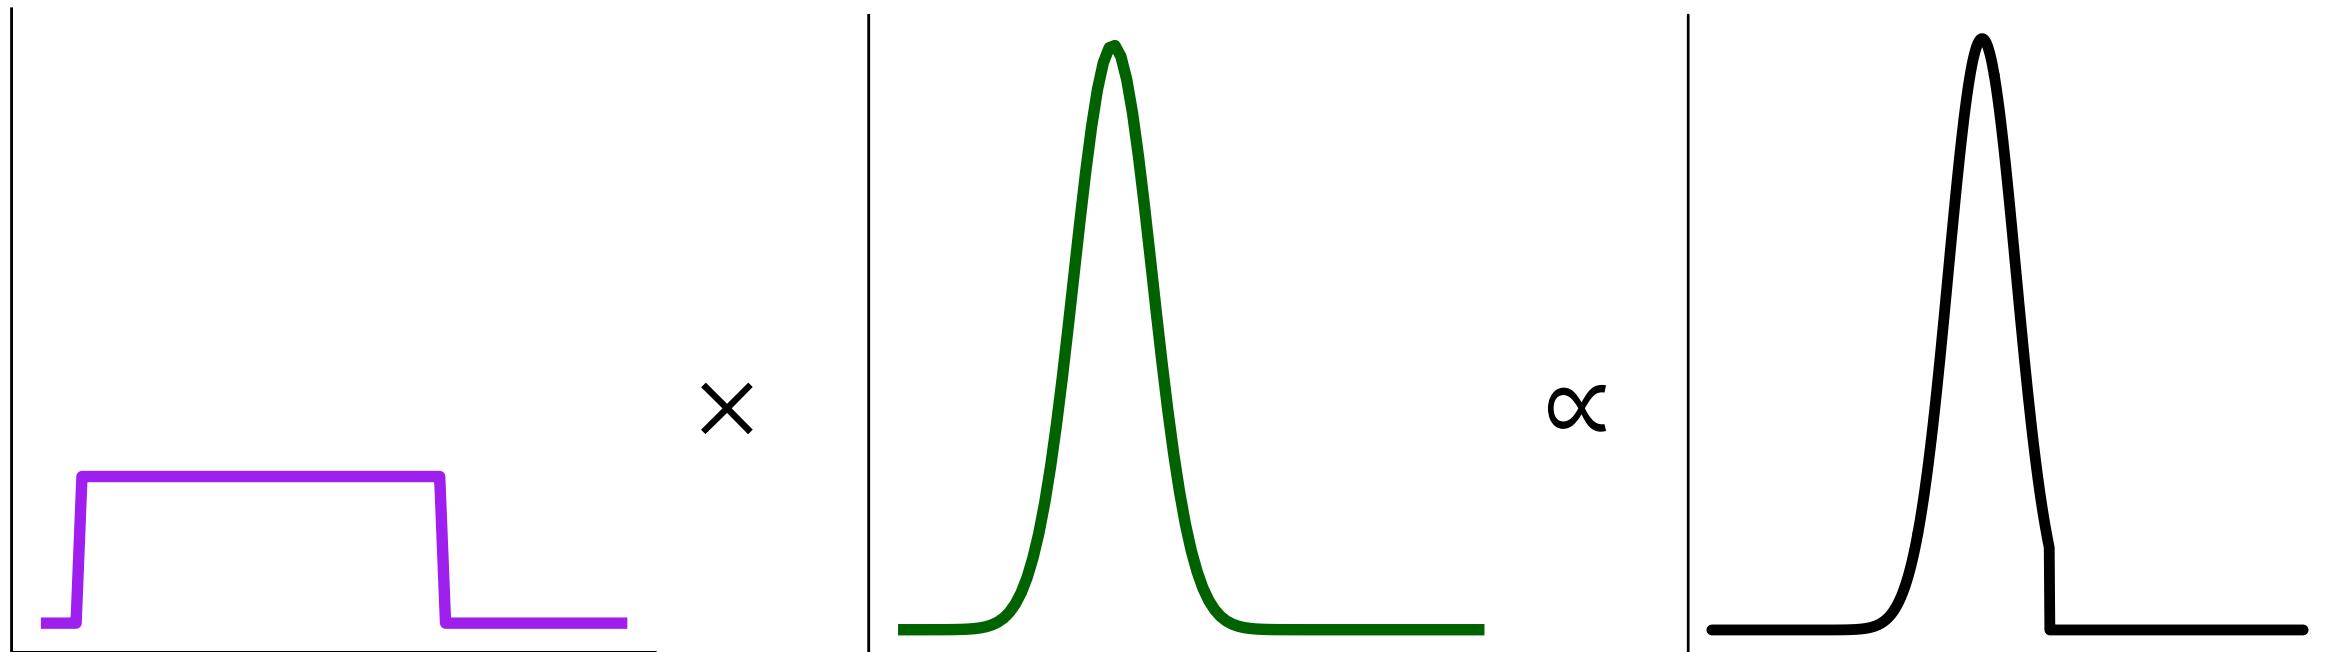
Consider multiple parameter values at once

This is the prior knowledge (in a form of a probability distribution)!

Prior is always given to each parameter in the model



Prior & data form posterior



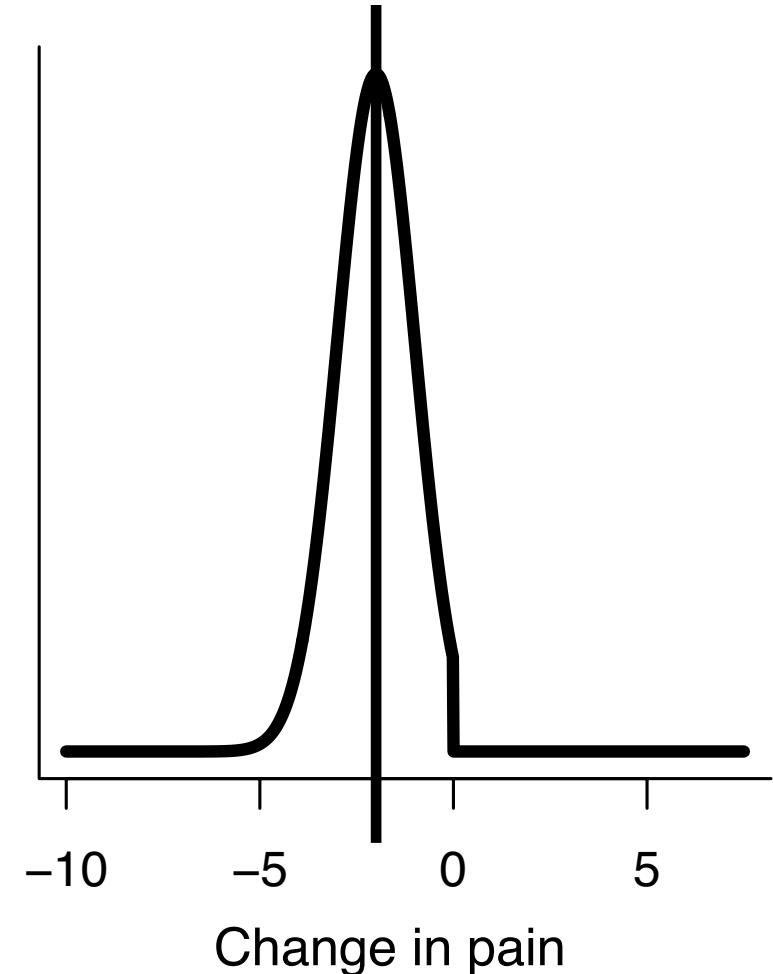
Posterior= prior probability distribution updated with the data, combines information of prior and data

Posterior probabilities

Bayes: Probability for each parameter value

The strongest probability that the pill decreases the pain by 2

Most of the probability mass supports the decrease from 0-5



Frequentist: Assuming the drug does not change the pain, our data is aligned/ rare... (p-value)

The frameworks answer a 'different question'

Frequentist

The probability of our **data**, assuming a fixed parameter value

- *Probability* only in the data, cannot apply to the parameter that is something for sure (we just don't know it)

Bayesian

The probability of different **parameter** values

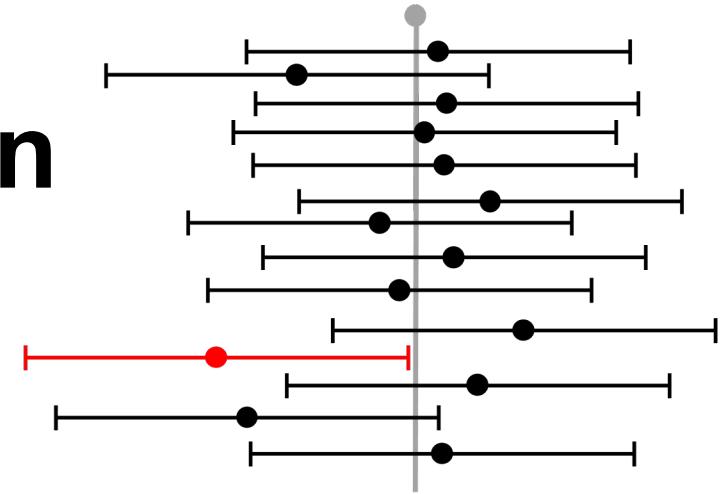
- *Probability* in the data but also in the parameter
- Prior knowledge
- *Subjective* probability

Multiple comparison correction

Multiple comparison= hypothesis testing multiple times

→ increased risk for false null-hypothesis rejection

→ correct the p-value (e.g. Bonferroni p-value*number of comparisons/ tests)



No (p-value based) hypothesis testing in Bayesian modeling, no corrections needed

Bayesian inference in neuroimaging: Example analysis of dopamine receptors using PET data



Dopamine

the molecule of more...and much more!



Emotions (excitement,
reward)



Cognition
(attention, learning)



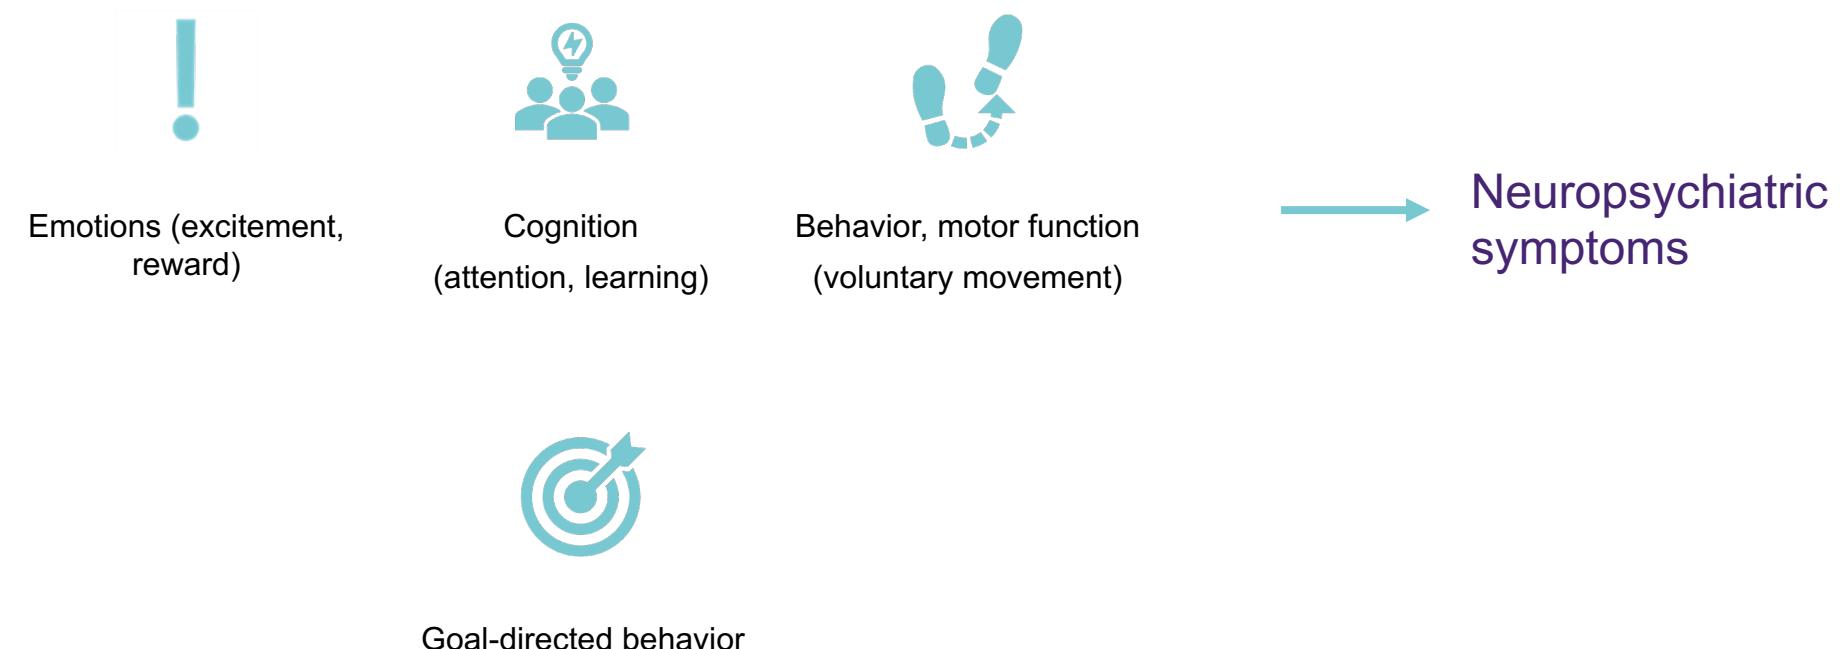
Behavior, motor function
(voluntary movement)



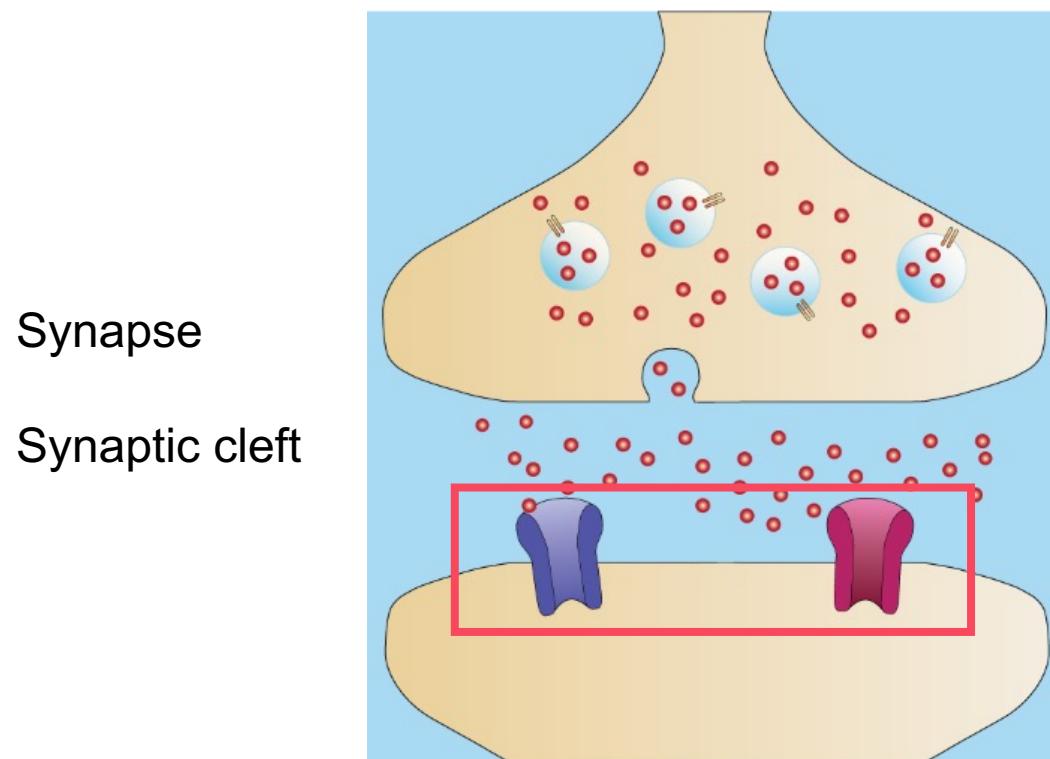
Goal-directed behavior

Dopamine

the molecule of more...and much more!



Dopamine receptors

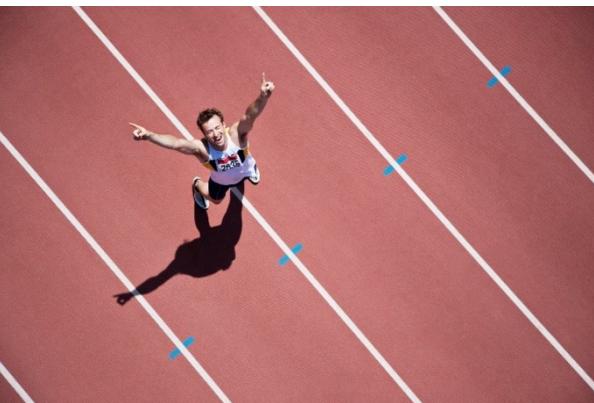




Are we different in our dopamine receptors?

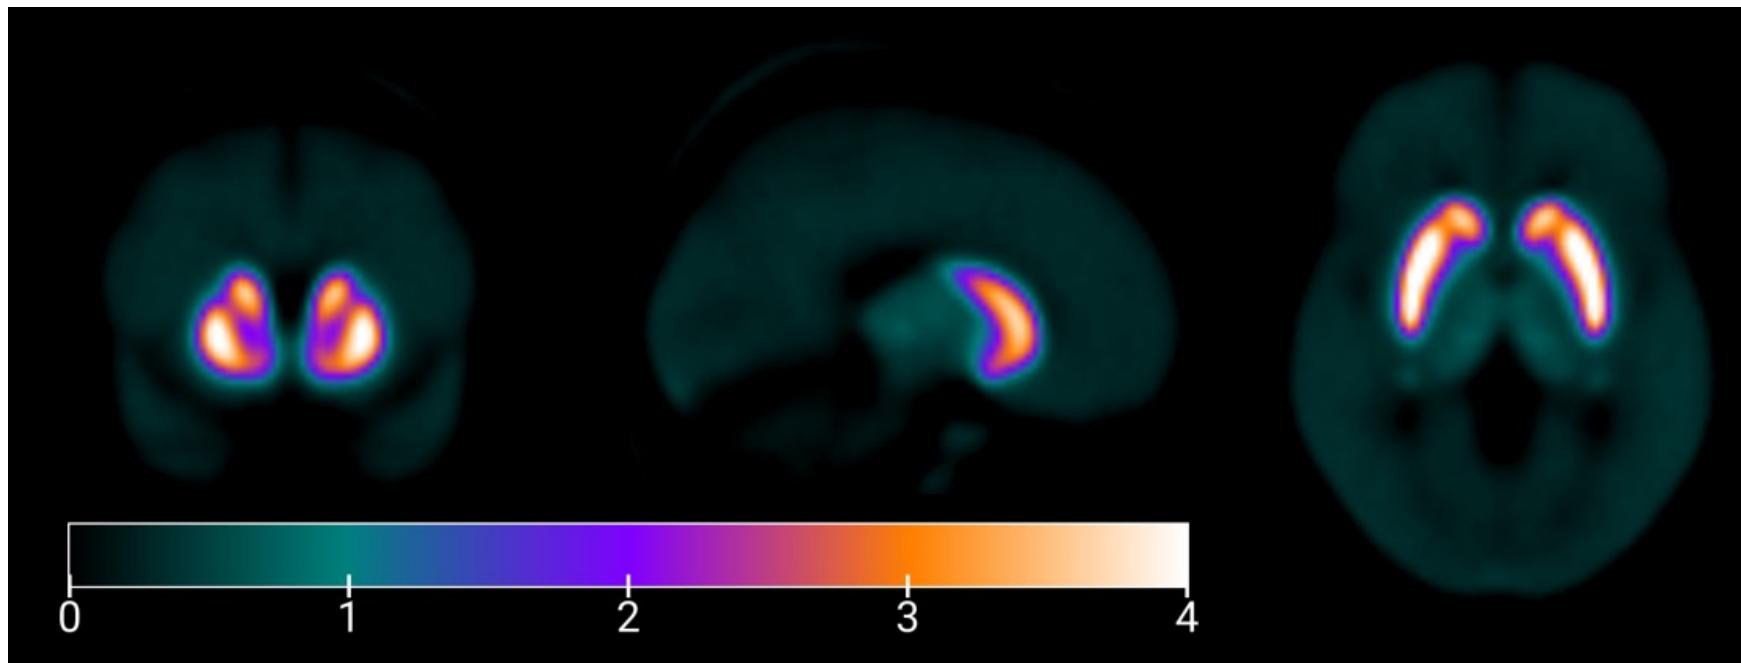
→ Vulnerability to neurological and psychiatric conditions?

- Age
- Sex
- Body mass index



Dopamine receptors with positron emission tomography (PET)

Striatum:
Caudate nucleus



BP_{ND} = binding potential, density and affinity of unoccupied receptors, receptor availability

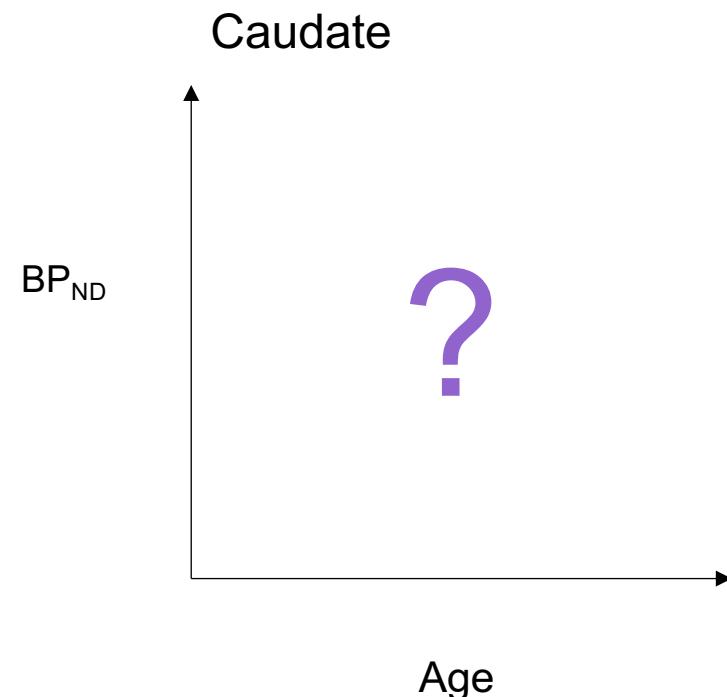
Data

	Dependent variable	Predictor	Predictor	Predictor
Subject	BP _{ND} (caudate)	Age	Sex	Body mass index
1	2.8	20	f	18
2	2.5	23	m	30
3	1.7	35	f	25
4	2.2	50	m	20
...

BP_{ND}= binding potential, density and affinity of unoccupied receptors, receptor availability

Regression: association

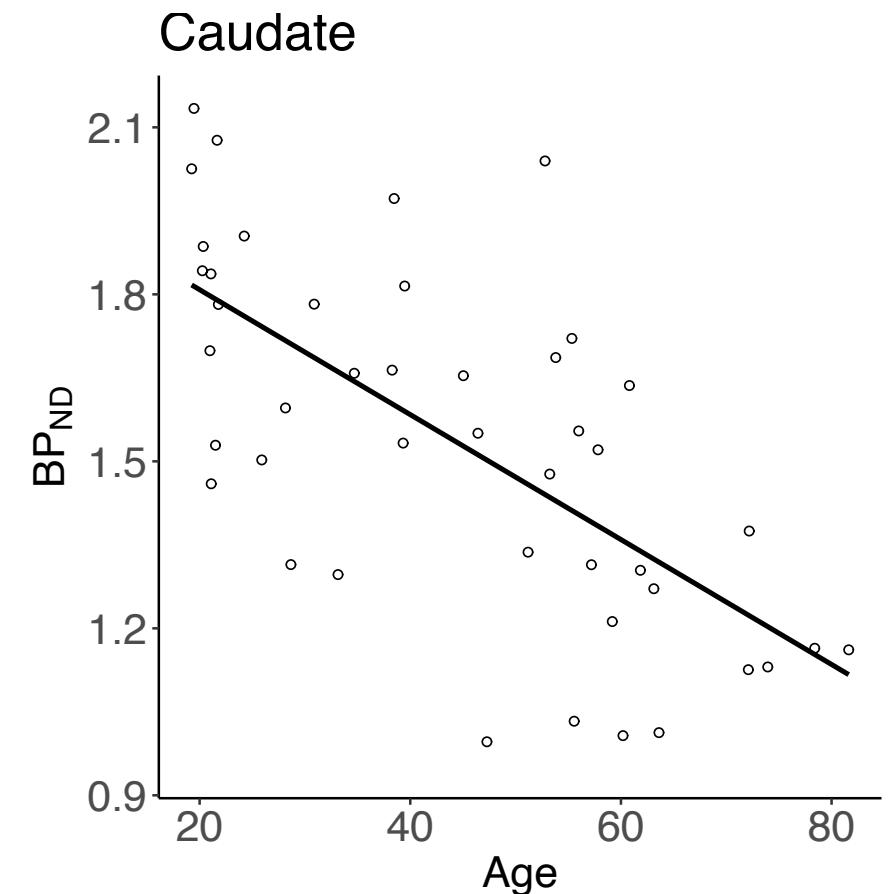
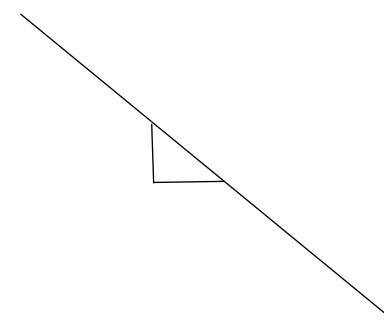
Age —→ BP_{ND} in caudate

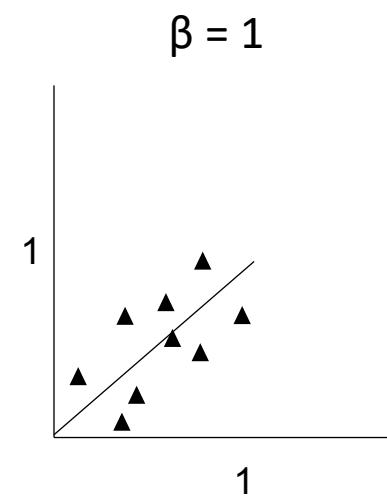
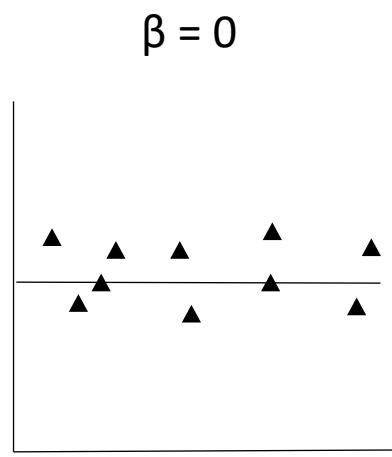
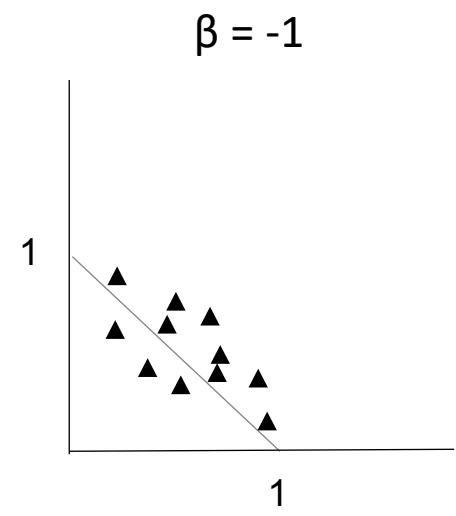


Parameter of interest

β

Regression coefficient
Ageing one unit, the change in BP_{ND}



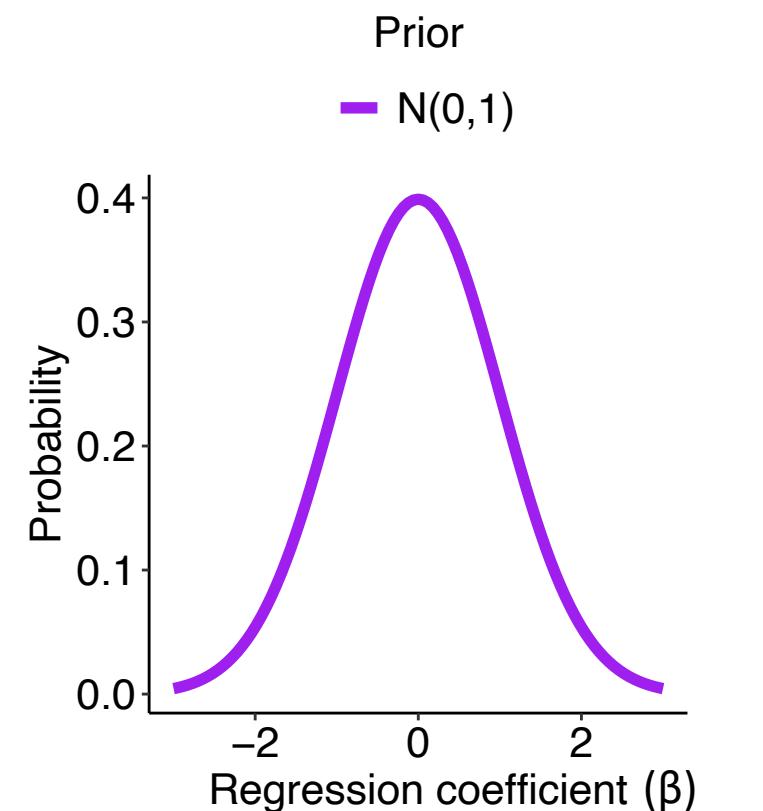


Prior for β

Prior knowledge as a probability distribution

Normal distribution with mean 0 and standard deviation of 1

- Range not limited
- Symmetrical (not weighted on either side)
- Weakly-informative: mean at zero, standard deviation wide

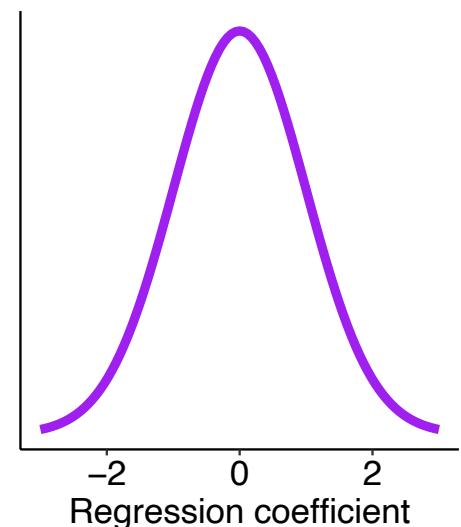
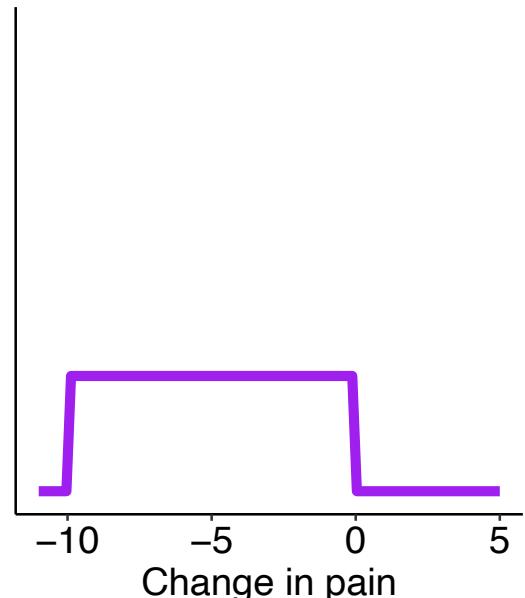


What's great about priors?

Regularize and exclude impossible parameter values

Previous findings to priors

Direct interpretation of the parameter probabilities
(interpretation of the data in frequentist inference)



UNIVERSITY
OF TURKU

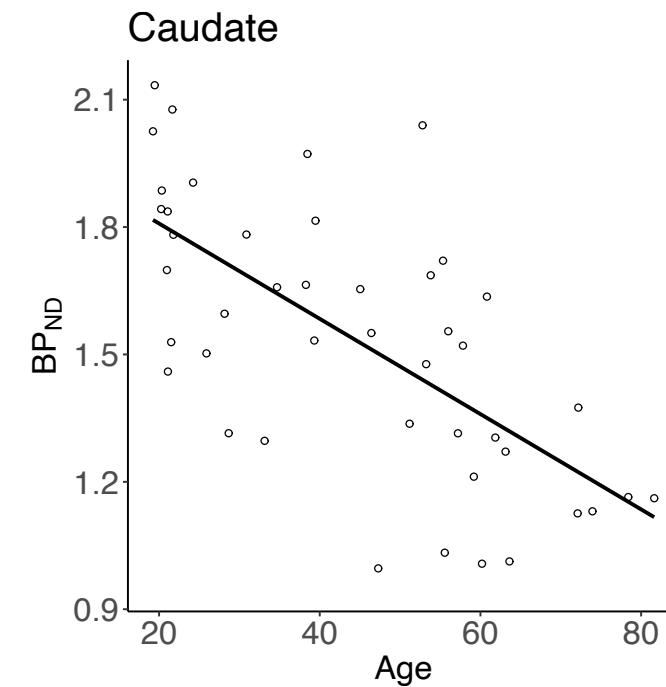
*Nothing Bayesian of frequentist about them

Model: linear regression*

$$BP_{ND} = \text{Intercept} + \beta * \text{age} + \text{error}$$

Intercept= BP_{ND} when age is 0

β = regression coefficient for age



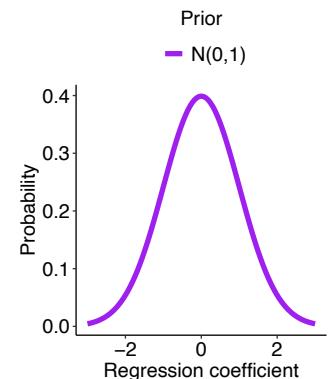
*Nothing Bayesian of frequentist about that

Model: linear regression*

$$BP_{ND} = \text{Intercept} + \beta * \text{age} + \text{error}$$

brms-package in R

```
model <- brm(formula = bf(BPND ~ 1 + age_z),
              data = data,
              prior = set_prior('normal(0,1)', class = 'b', coef = 'age_z'),
              chains = 4,
              iter = 4000,
              warmup = 1000)
```

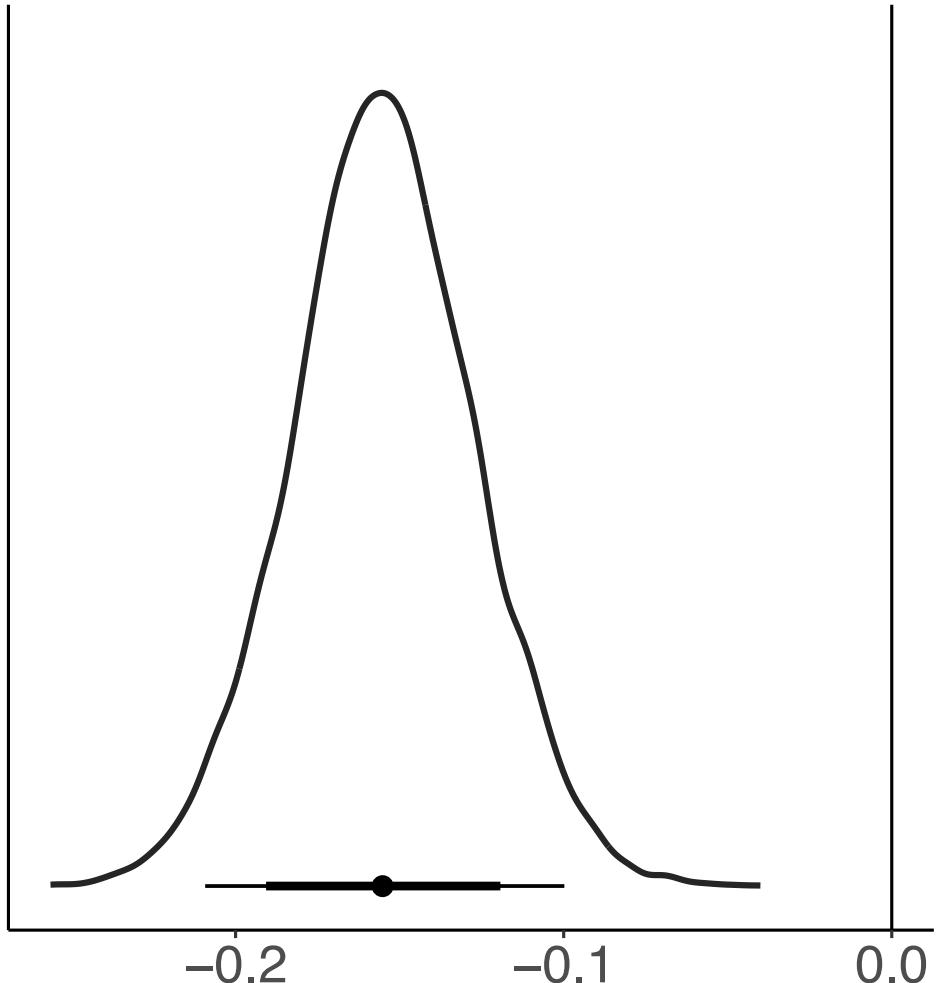


Intercept= BP_{ND} when age is 0

β = regression coefficient for age

(R is doing the tricks)

Posterior age effect (beta)



	Estimate	Est.Error	Q2.5	Q97.5
Intercept	1.6723457	0.04432914	1.5853905	1.75987669
age_z	-0.1549403	0.02783491	-0.2092666	-0.09979254

Ageing 1 unit (~20 years), BP_{ND} decreases 0.1 - 0.2

→ Dopamine receptor availability (density and/ or affinity) decreases through age in caudate

To take home

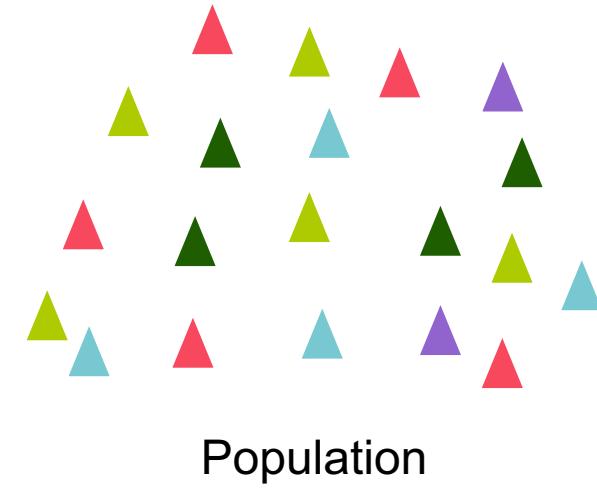
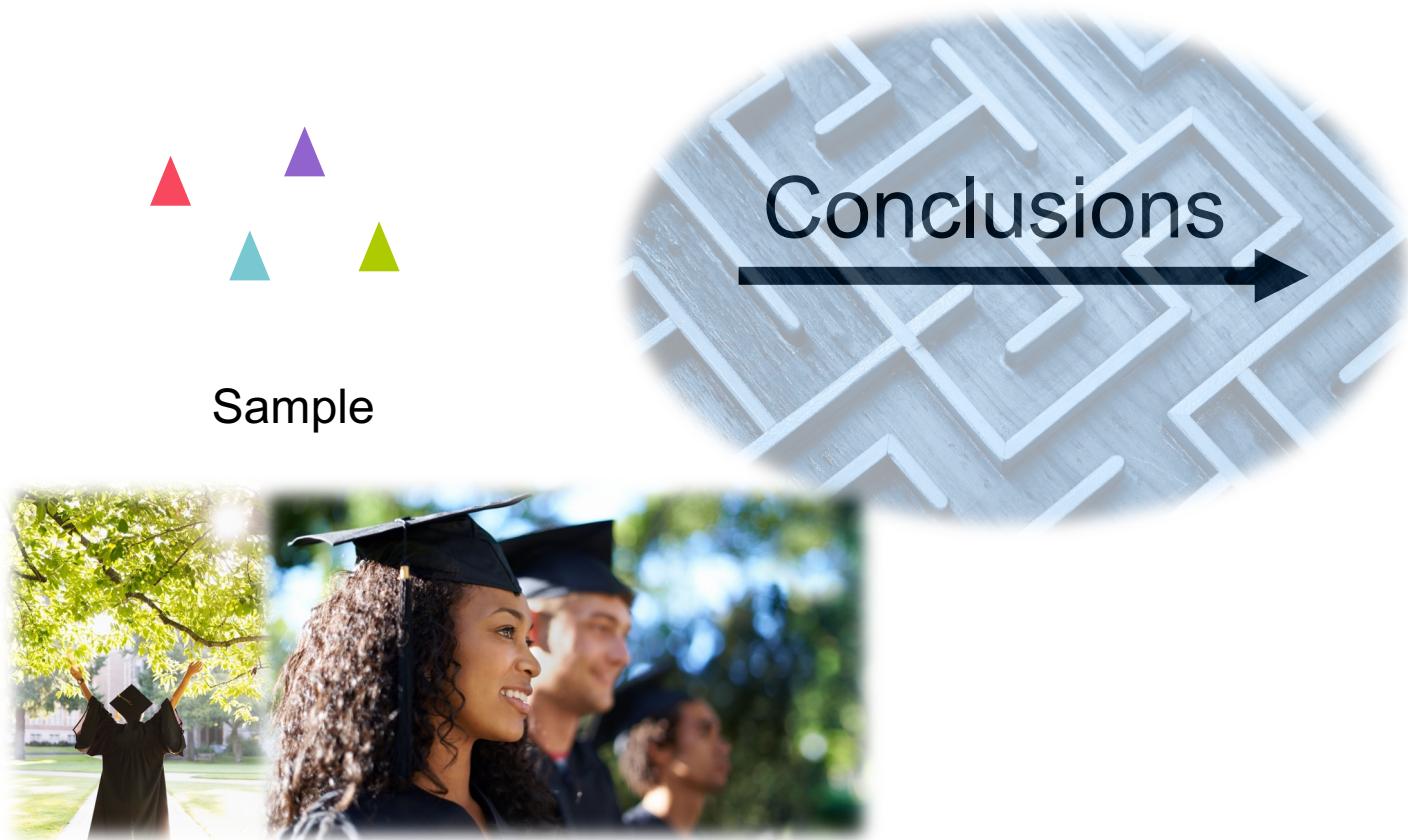
Bayesian statistical inference allows

- Direct inference on the parameter (how probable each parameter value is)
- Regularization with priors: Exclude impossible parameter values, while letting the data talk (objectivity, assumptions also in frequentist)
 - Population of Finland versus world
- Modeling without multiple comparison corrections
- Multiple ways to do inference: know your options and what question they answer to



UNIVERSITY
OF TURKU

Statistical inference



Learning Material

- Bayesian Data Analysis, by Andrew Gelman, John Carlin, Hal Stern, David Dunson, Aki Vehtari, and Donald Rubin
 - <http://www.stat.columbia.edu/~gelman/book/>
- McElreath, R. (2018). *Statistical rethinking: A Bayesian course with examples in R and Stan.* Chapman and Hall/CRC.
 - <https://www.youtube.com/watch?v=4WVeICswXo4>
 - Scripts and datasets



UNIVERSITY
OF TURKU

Thank you!

tukama@utu.fi

emotion.utu.fi

Bonus slide / Fun to know

Mixed effects models*

Fixed, interest in these particular levels of the variable (age), comparison between the levels

Random, random sample of the variable that groups the data (scanner) but we are not interested in the comparison of these particular levels, instead their average difference

