Prof. Kirstin Hagelskjær Petersen
kirstin@cornell.edu

Fast Robots

Noisy Sensors

- Example: Accelerometer
- Solution?
  - Average over multiple samples
  - mean = -9.97306mg
  - std dev = 7.0318mg
- Normal distributions
  - Described with 2 numbers
    - $[\mu \mp \sigma]$
    - Symmetric
    - Unimodal
    - Sums to unity
- Probabilistic robotics
  - Measurements are uncertain
  - Actions are uncertain
  - States are uncertain

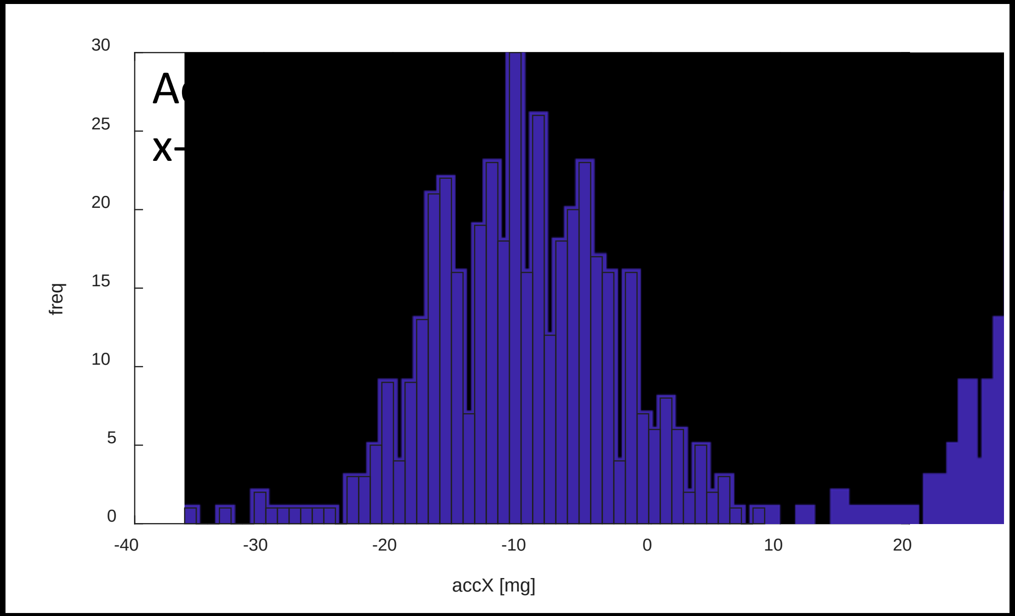

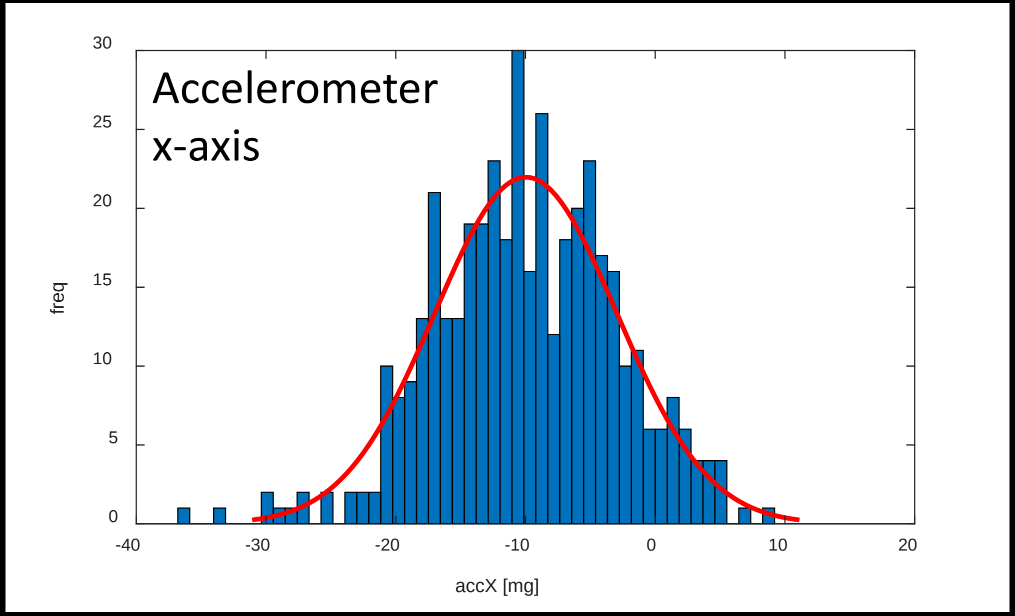

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?

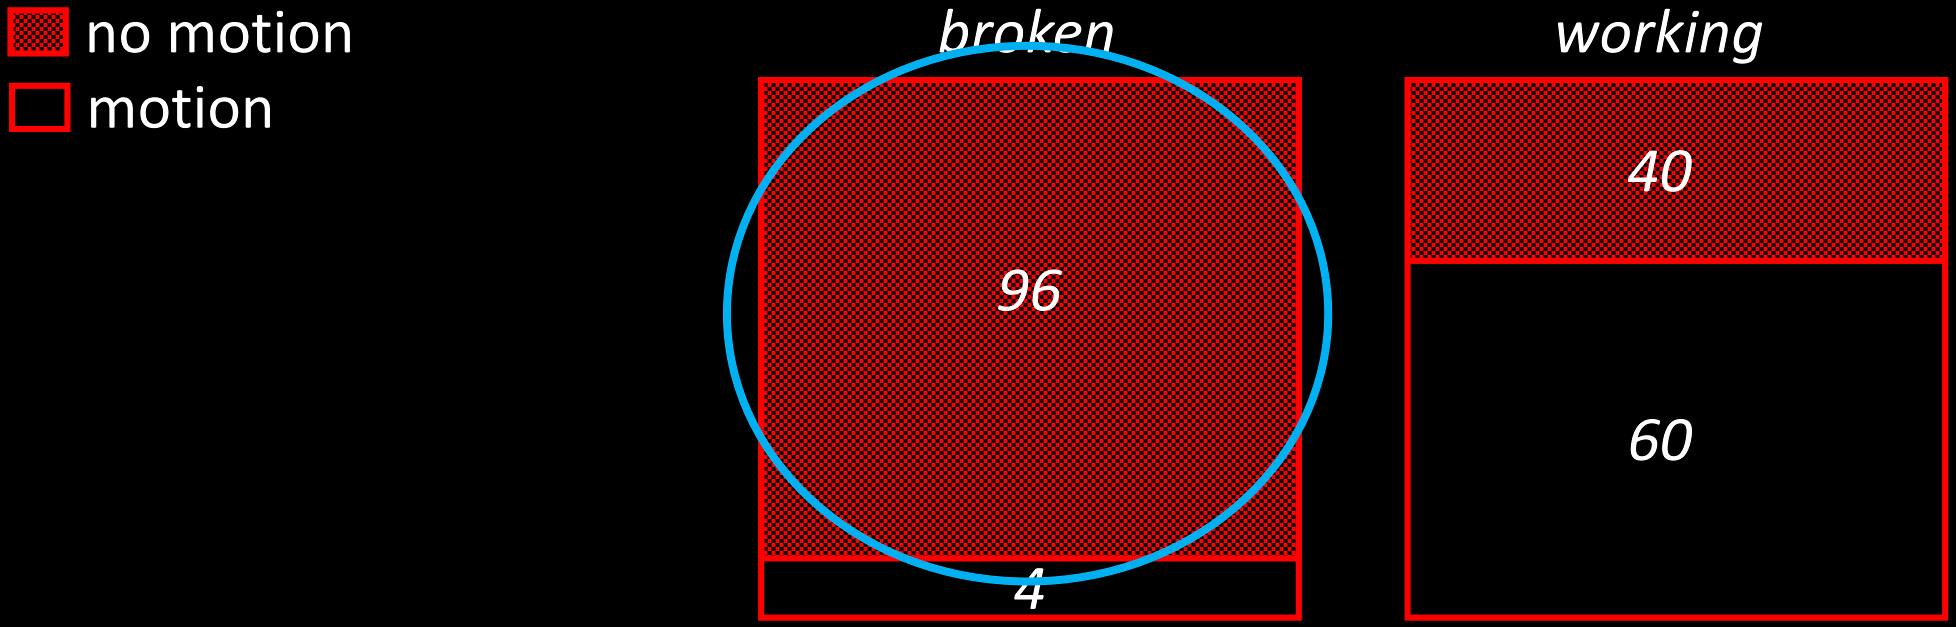

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?

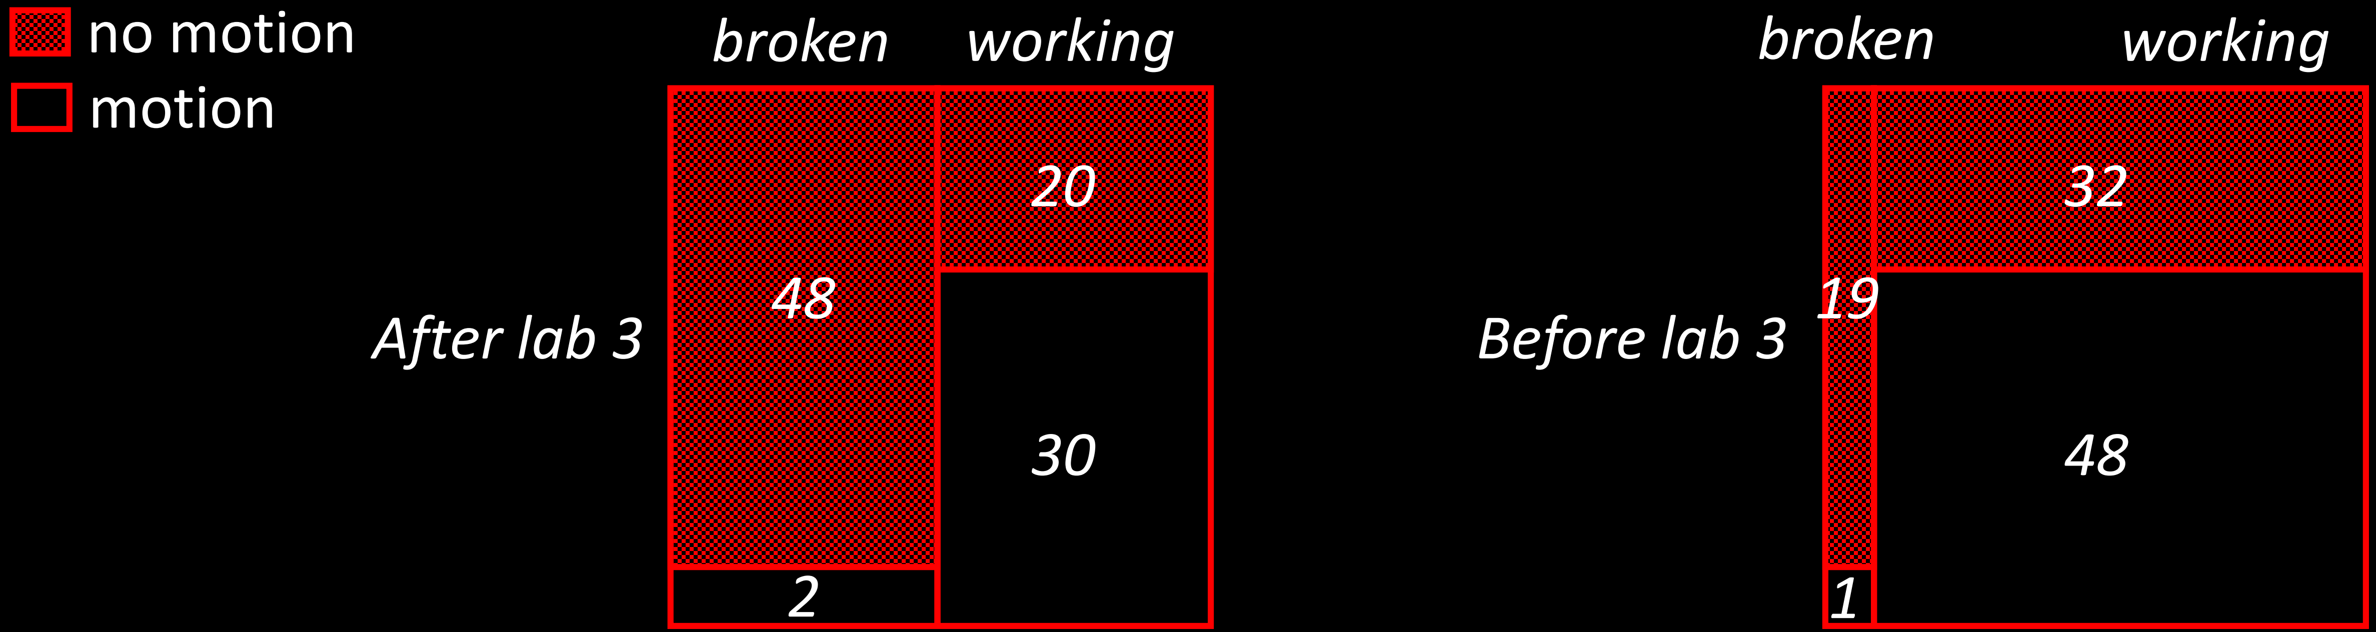

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?
- **Translate to math**
  - P(something) = #something / #everything
  - Before lab 3:
    - P(broken) = #broken / #kits = 20 / 100 = 0.2
    - P(working) = #working / #kits = 80 / 100 = 0.8
  - After lab 3:
    - P(broken) = #broken / #kits = 50 / 100 = 0.5
    - P(working) = #working / #kits = 50 / 100 = 0.5

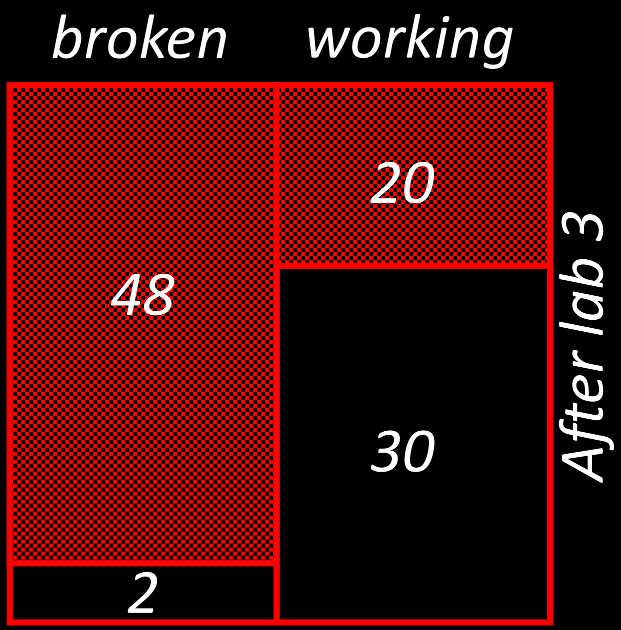

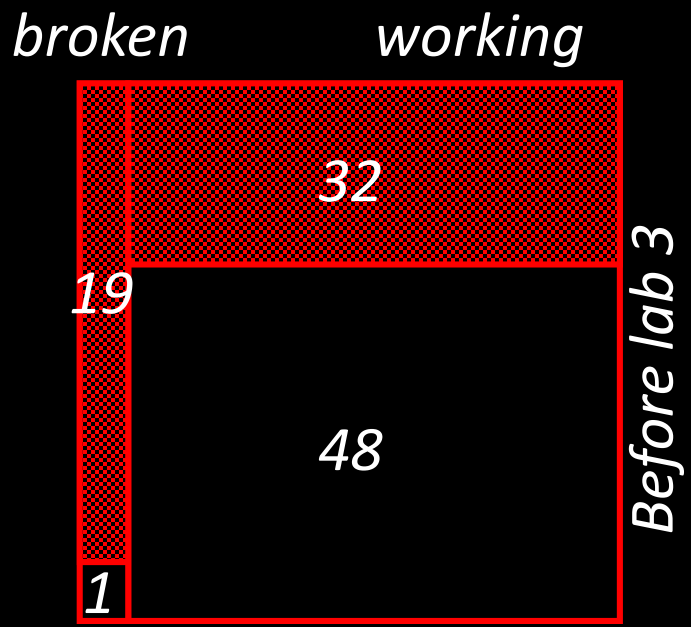

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?

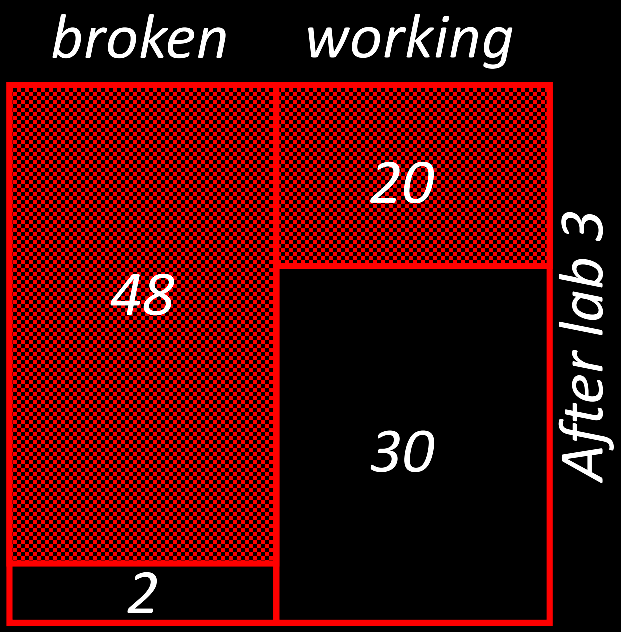

- **Conditional Probability**
  - If you know that the robot is broken, what is the probability that it stopped moving?
  - P(no motion | broken) = #broken and no motion / #broken
  - After lab 3 = 48/50 = 0.96
  - P(no motion | working) = #working and no motion / #working
  - After lab 3 = 20/50 = 0.40

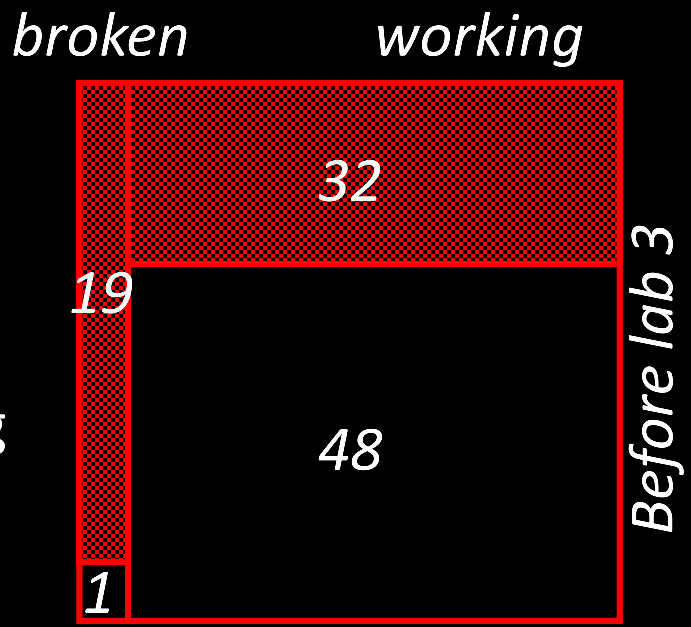

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?

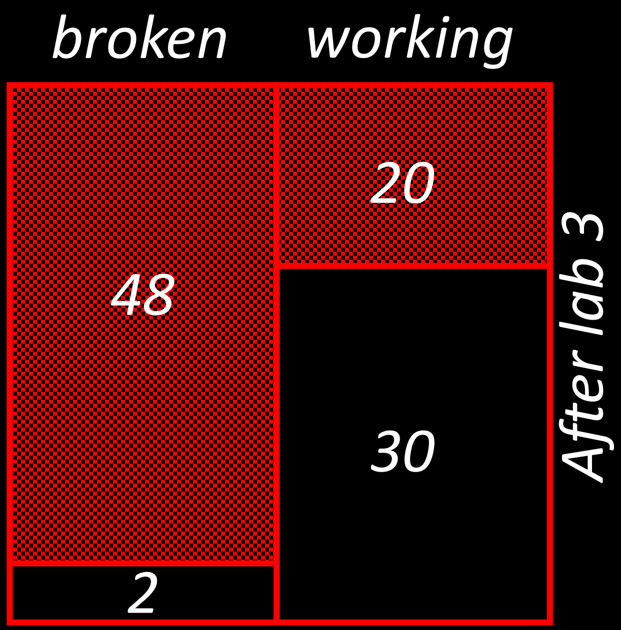

- **Conditional Probability**
  - If you know that the robot is broken, what is the probability that it stopped moving?
  - P(no motion | broken) = #broken and no motion / #broken
  - Before lab 3 = 19/20 = 0.96
  - P(no motion | working) = #working and no motion / #working
  - Before lab 3 = 32/80 = 0.40

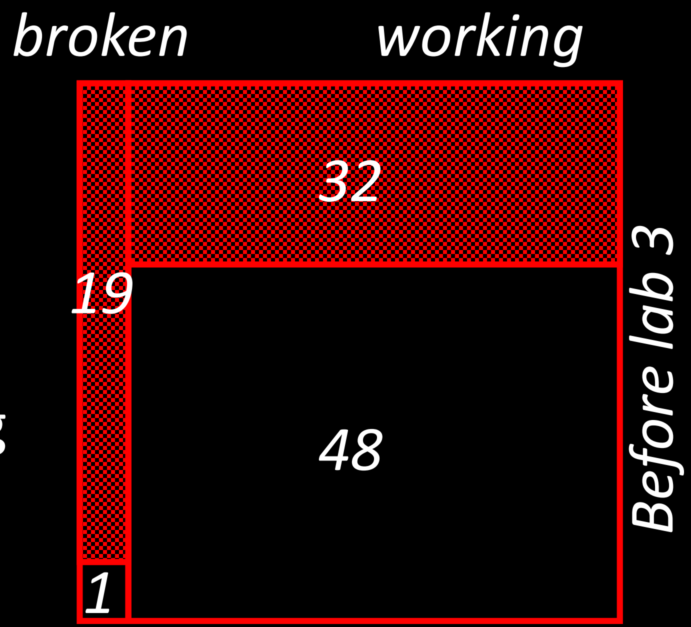

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?

- **Conditional Probability**
  - If you know that the robot is broken, what is the probability that it stopped moving?
  - P(A|B) is the probability of A, given B
  - Note: P(A|B) is not equal to P(B|A)
    - P(cute|puppy) ≠ P(puppy|cute)

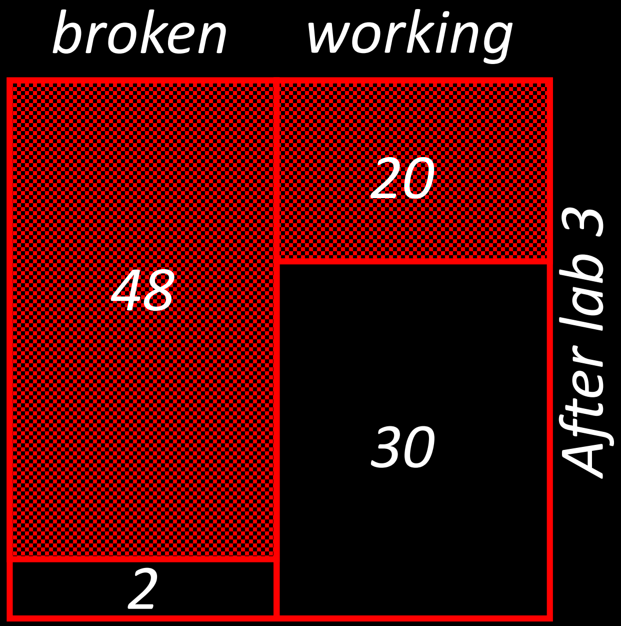

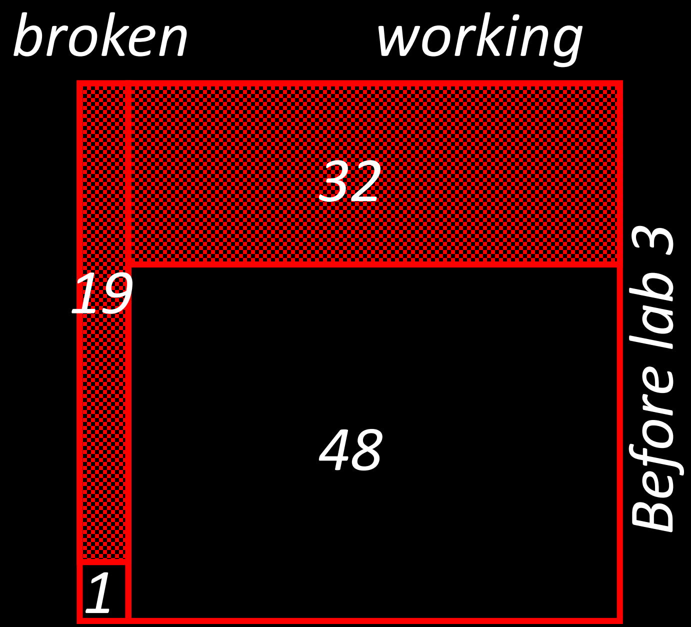

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?

- **Joint Probability**
  - What is the probability that the robot is both broken and not moving?
  - P(broken and not moving)
    = P(broken)*P(not moving | broken)
    = 0.5 * 0.96 = 0.48

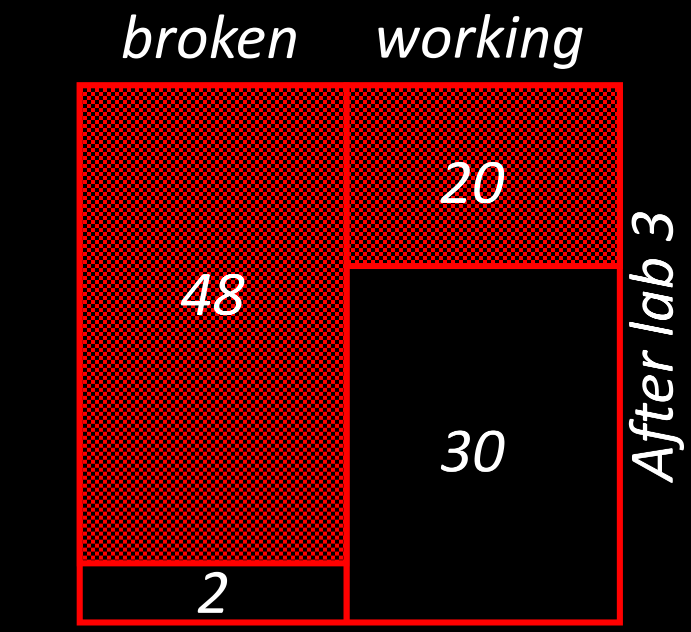

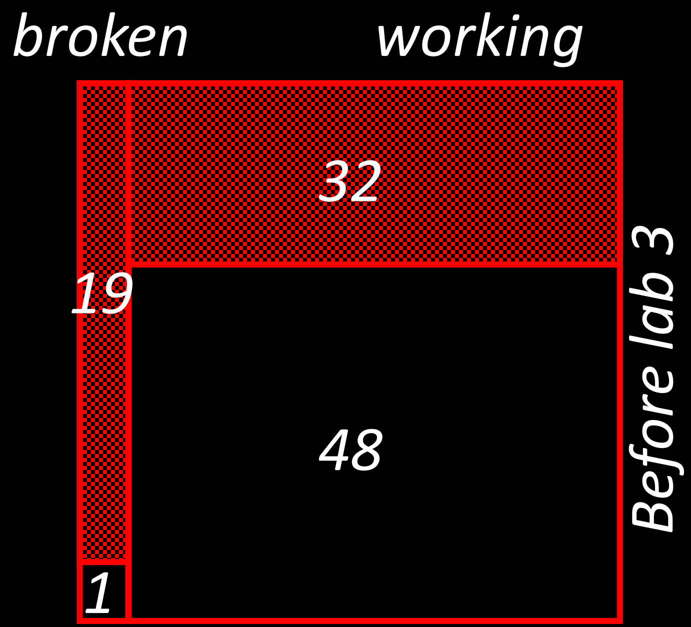

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?
- **Joint Probability**
  - What is the probability that the robot is both broken and not moving?
  - P(broken and not moving)
    = P(broken)*P(not moving | broken)
    = 0.20 * 0.96 = 0.192
  - P(working and moving)
    = P(working)*P(moving | working)
    = 0.80 * 0.60 = 0.48

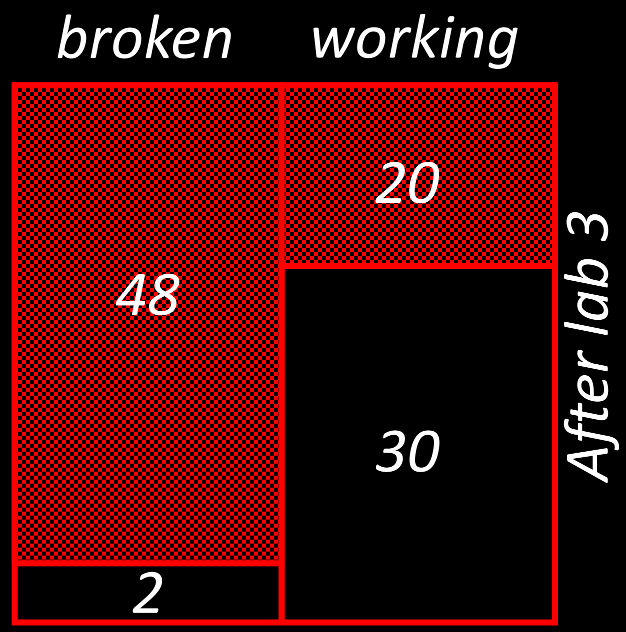

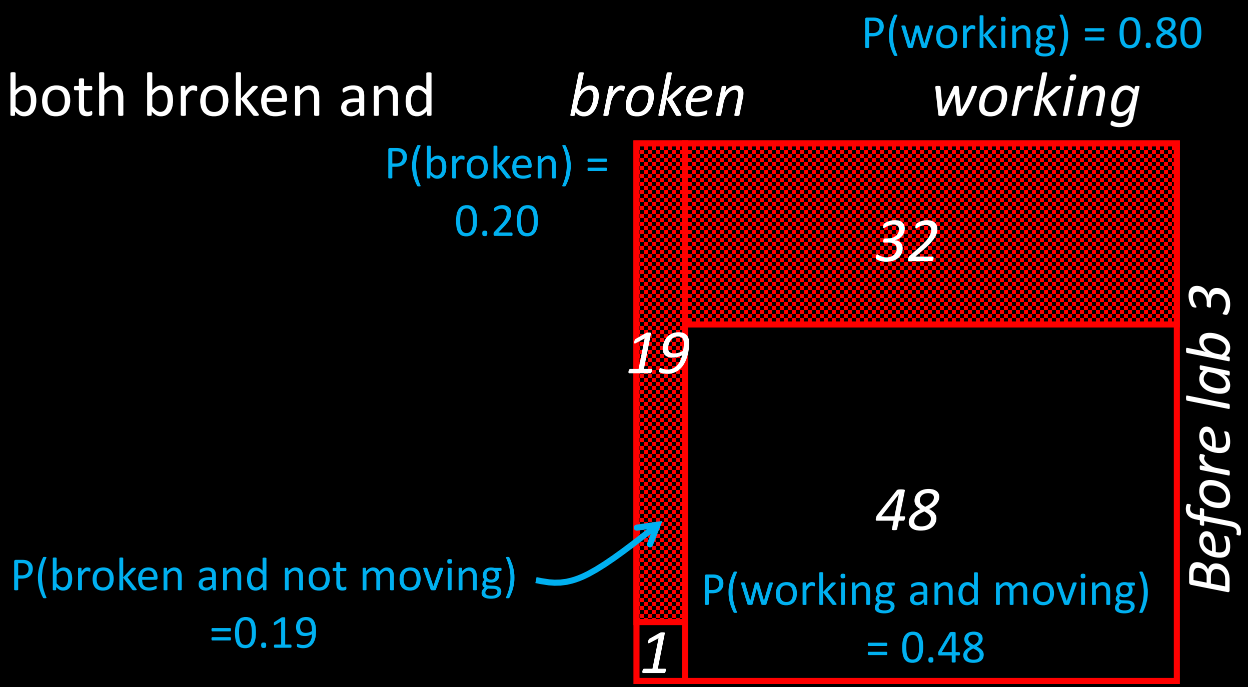

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?
- **Joint Probability**
  - What is the probability that the robot is both broken and not moving?
  - P(A, B) = P(A∩B) = P(A and B)
  - P(A∩B) = P(A)*P(B|A)
  - P(A∩B) = P(B∩A)

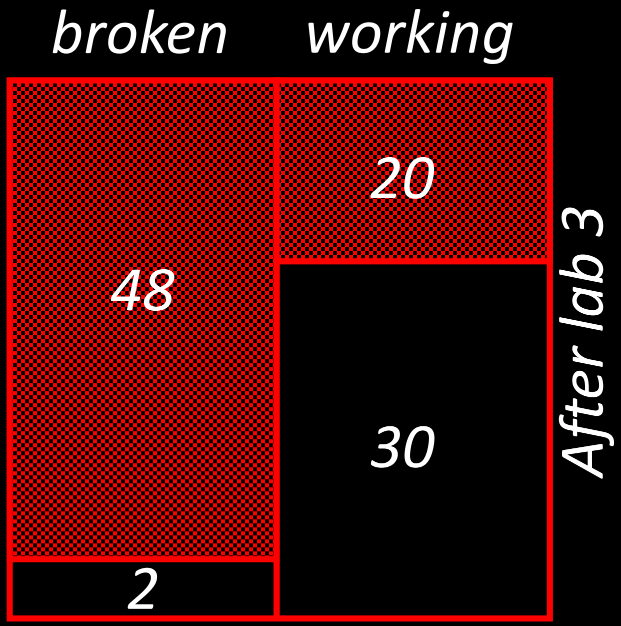

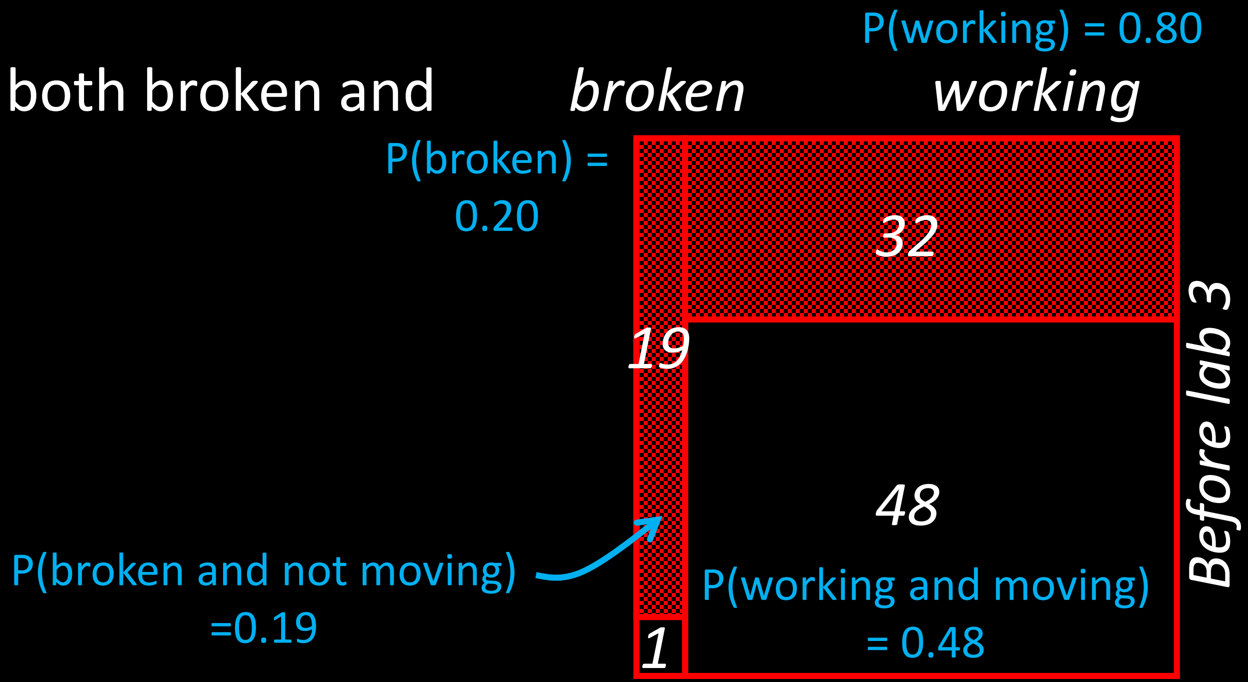

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?

- **Marginal Probability**
  - P(moving)
    = P(broken and moving) + P(working and moving)
    = 1/100 + 48/100 = 0.49
  - P(not moving)
    = 19/100 + 32/100 = 0.51

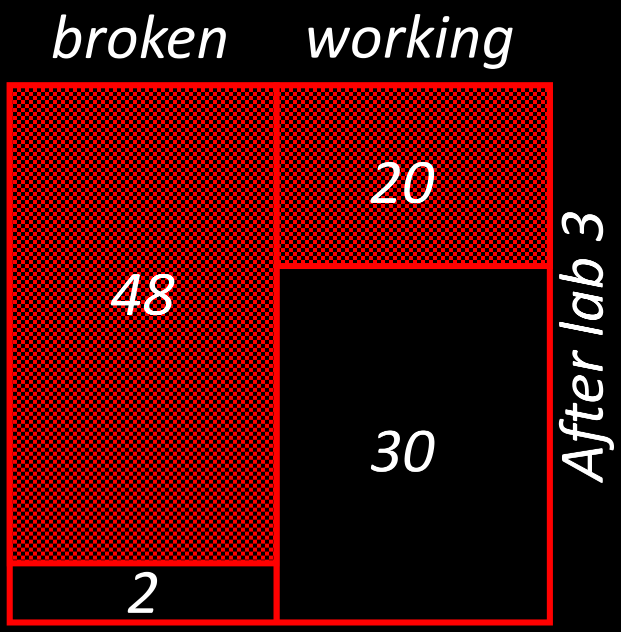

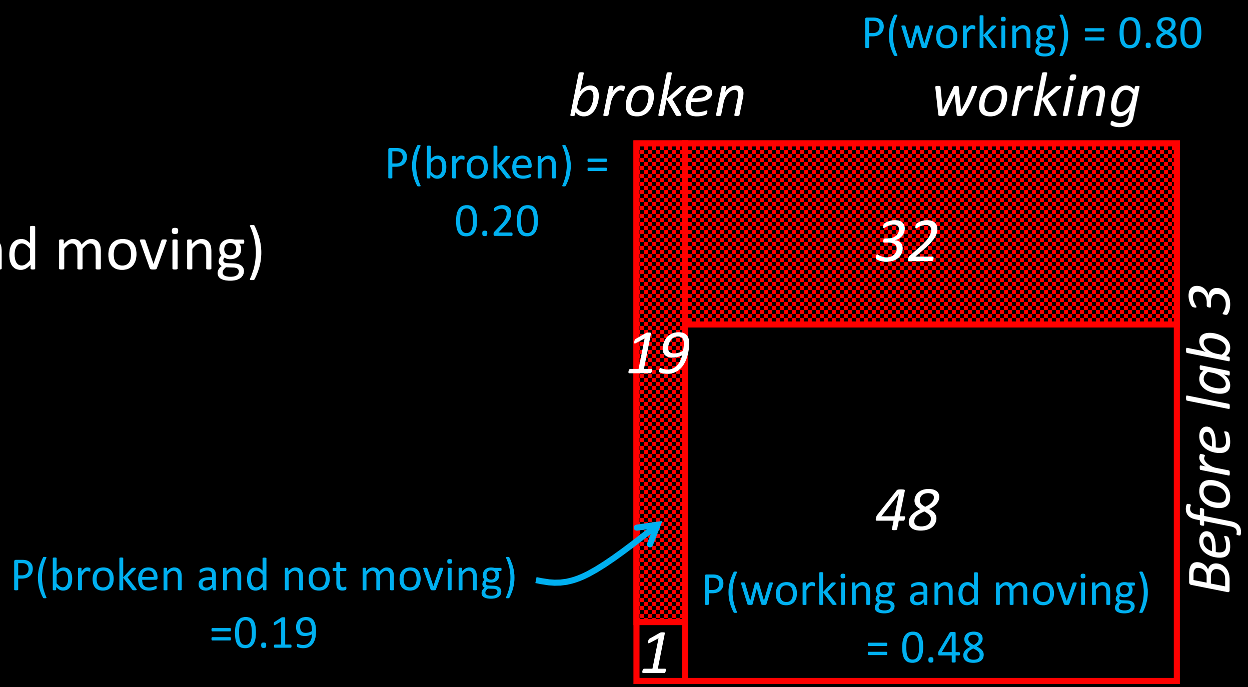

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

- Example
  - *Campuswire:* The robot stopped moving, the hardware is broken, send me new parts
  - What is the probability that the robot is broken, given that it stopped moving?
  - P(broken | not moving) = ???
- P(broken and not moving)
  = P(not moving)*P(broken|not moving)
- P(not moving and broken)
  = P(broken)*P(not moving|broken)
- P(broken|not moving) = $\frac{\text{P(broken)*P(not moving|broken)}}{\text{P(not moving)}}$

- Before lab 3 = 0.2*0.96 / 0.51 = 0.38
- After lab 3 = 0.5*0.96 / 0.68 = 0.71

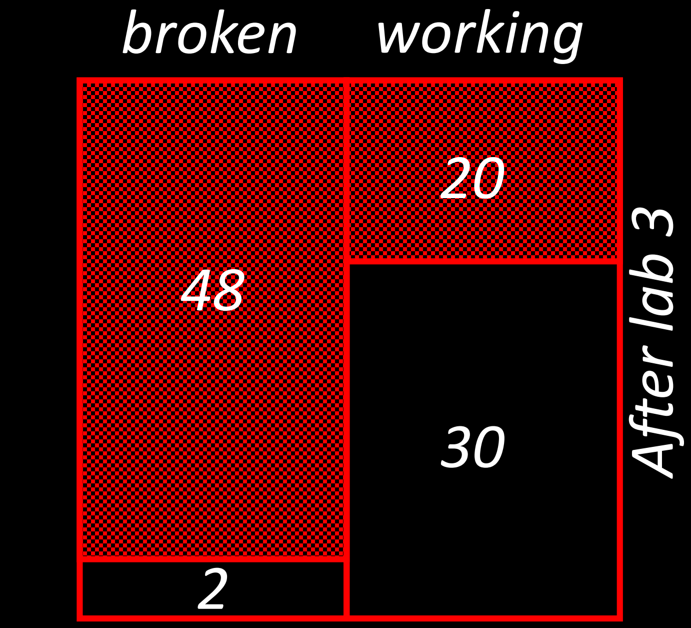

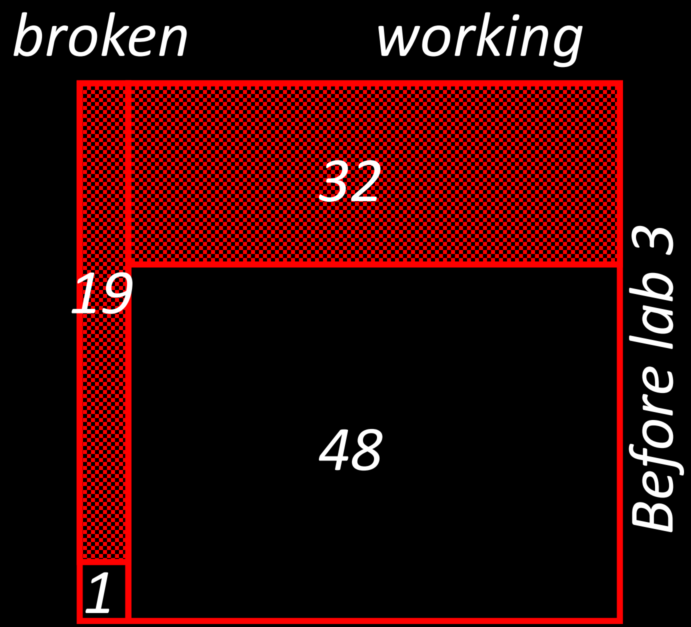

Bayesian Inference

- Bayesian inference = guessing in the style of Bayes

likelihood — prior

$$P(A|B) = \frac{P(B|A)\,P(A)}{P(B)}$$

conditional probability
posterior

marginal likelihood
(constant)

Exercise

- **Conditional probability**
  - You meet a guy, and he says he has a sibling, what is the probability that the sibling is female?
    - guy/girl
    - guy/guy
    - girl/guy
    - girl/girl (<ruled out)
    - 33%
- Independent / dependent variables

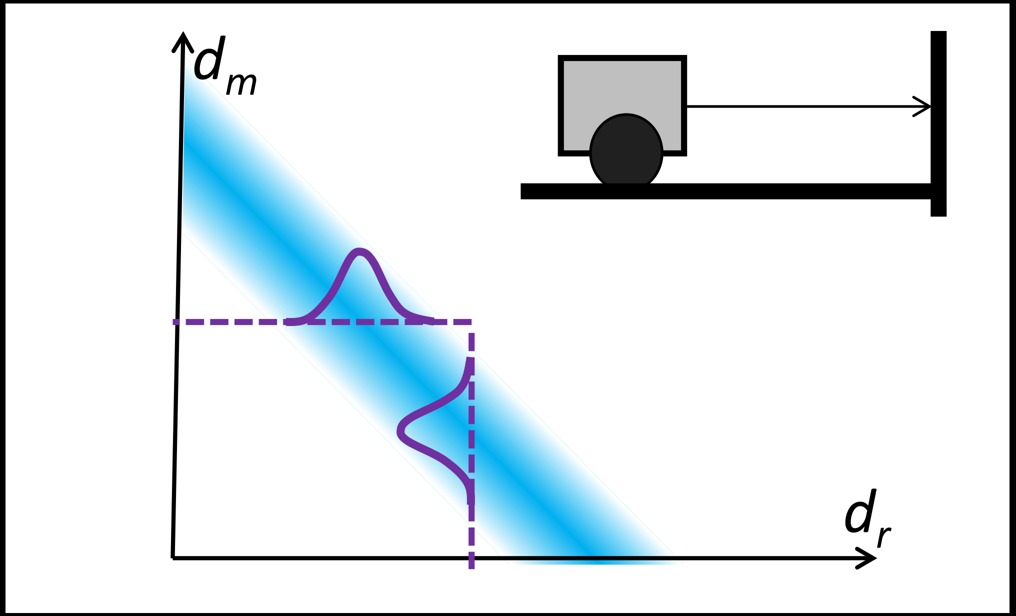

Bayesian Inference

- Unrelated example, borrowed from “Bayes with Beans” by Myriam Hunink, Harvard

$$P(A|B) = \frac{P(B|A)\,P(A)}{P(B)}$$

likelihood

prior

conditional probability
posterior

marginal likelihood
(constant)

Bayesian Inference

- Unrelated example, borrowed from “Bayes with Beans” by Myriam Hunink, Harvard

New info from test

Pre-test info → Post-test info

0 — P(D+) — 1

Bayesian Inference

- Unrelated example, borrowed from "Bayes with Beans" by Myriam Hunink, Harvard

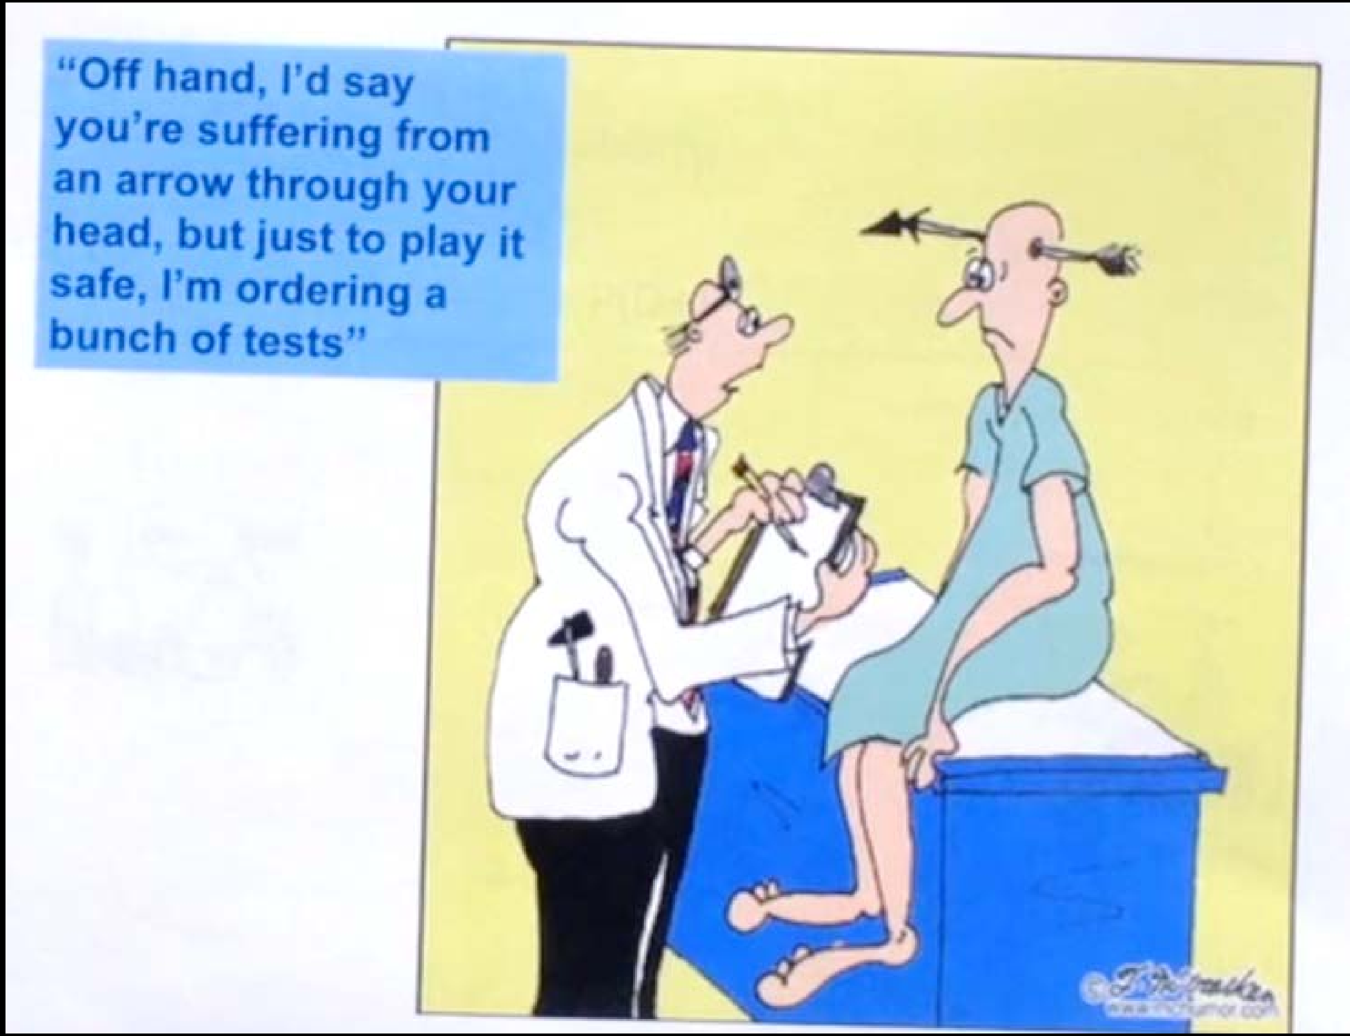

New info from test

Pre-test info → Post-test info

0 P(D+) 1

Bayesian Inference

- Unrelated example, borrowed from “Bayes with Beans” by Myriam Hunink, Harvard

- 100 people go to the doctor with a tick bite, worried they have Lyme disease
- A diagnostic test reveals 16 positives, 84 negative
- Early tests are inaccurate:
  - 40% of sick people will test positive
  - 10% of healthy people will test positive

Bayesian Inference

- Unrelated example, borrowed from “Bayes with Beans” by Myriam Hunink, Harvard

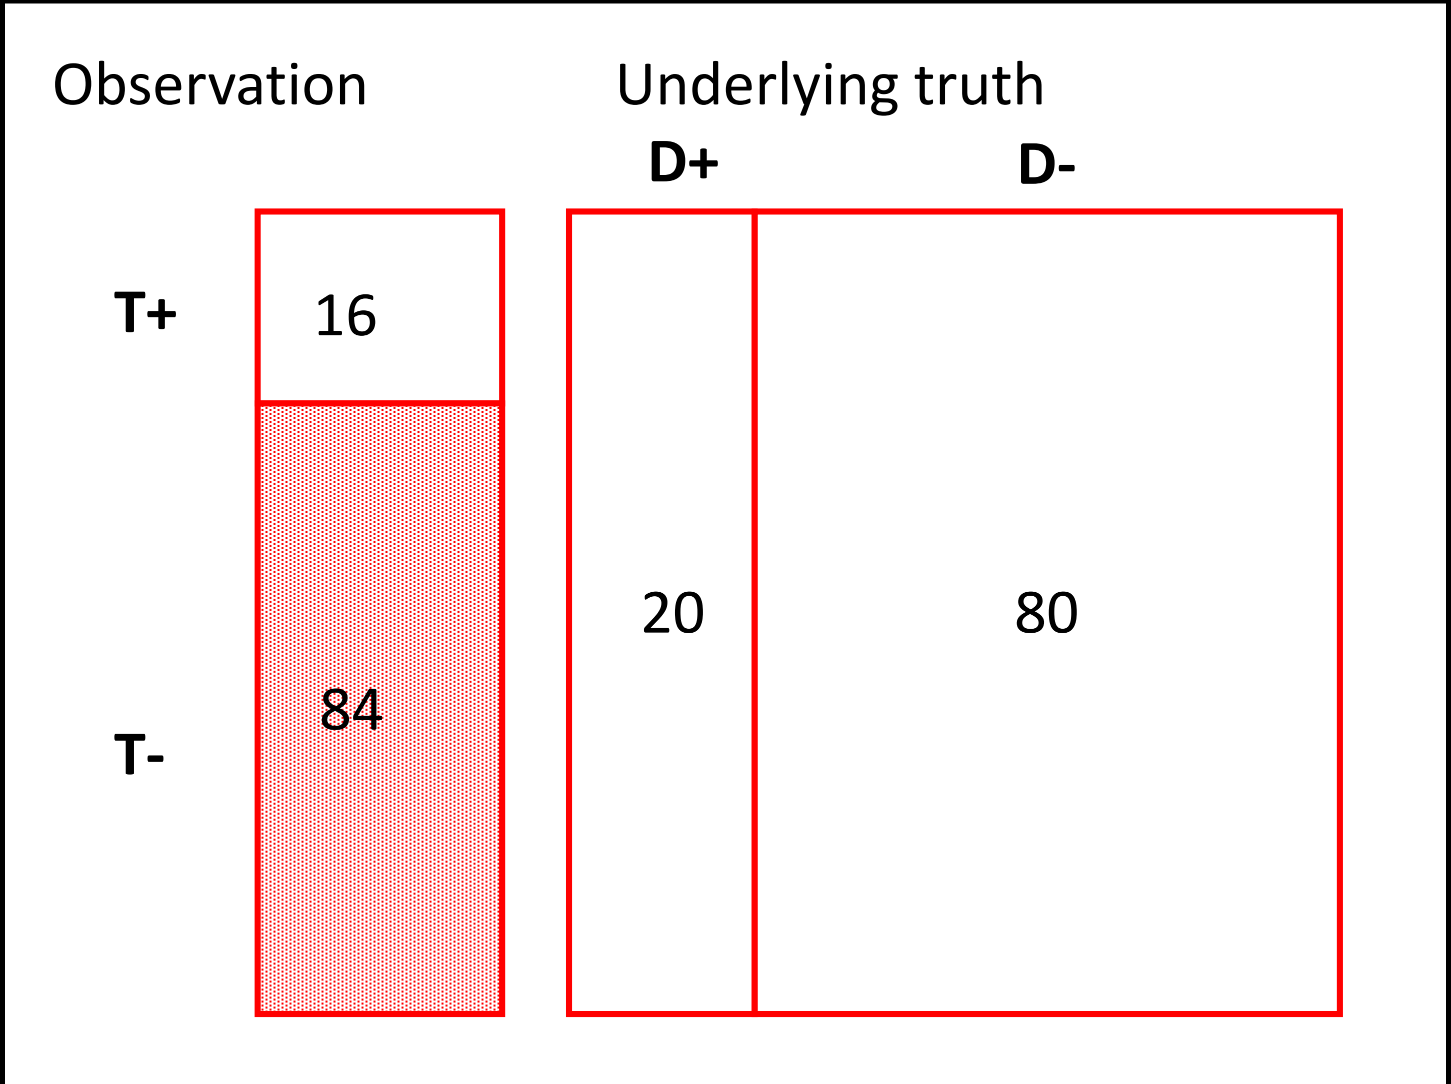

- 100 people go to the doctor with a tick bite, worried they have Lyme disease
- A diagnostic test reveals 16 positives, 84 negative
- Early tests are inaccurate:
  - 40% of sick people will test positive
  - 10% of healthy people will test positive
- Underlying truth:
  - 20% are sick
  - 80% are healthy

Bayesian Inference

- Unrelated example, borrowed from “Bayes with Beans” by Myriam Hunink, Harvard

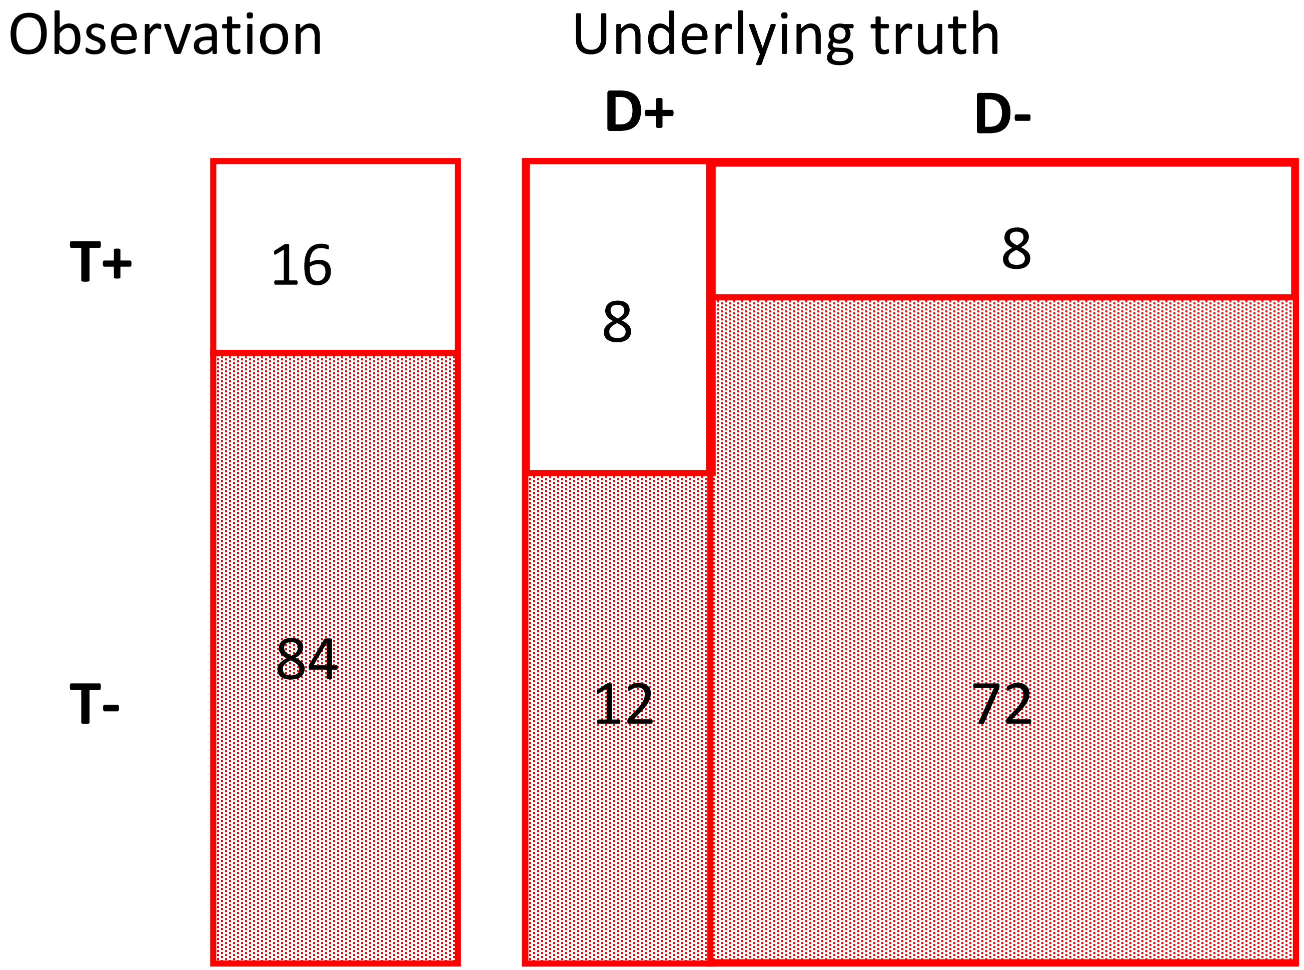

$$= \frac{8}{16} = 50\% = P(D + |T+)$$

$$= \frac{12}{84} = 14\% = P(D + |T-)$$

Bayesian Inference

- Unrelated example, borrowed from “Bayes with Beans” by Myriam Hunink, Harvard

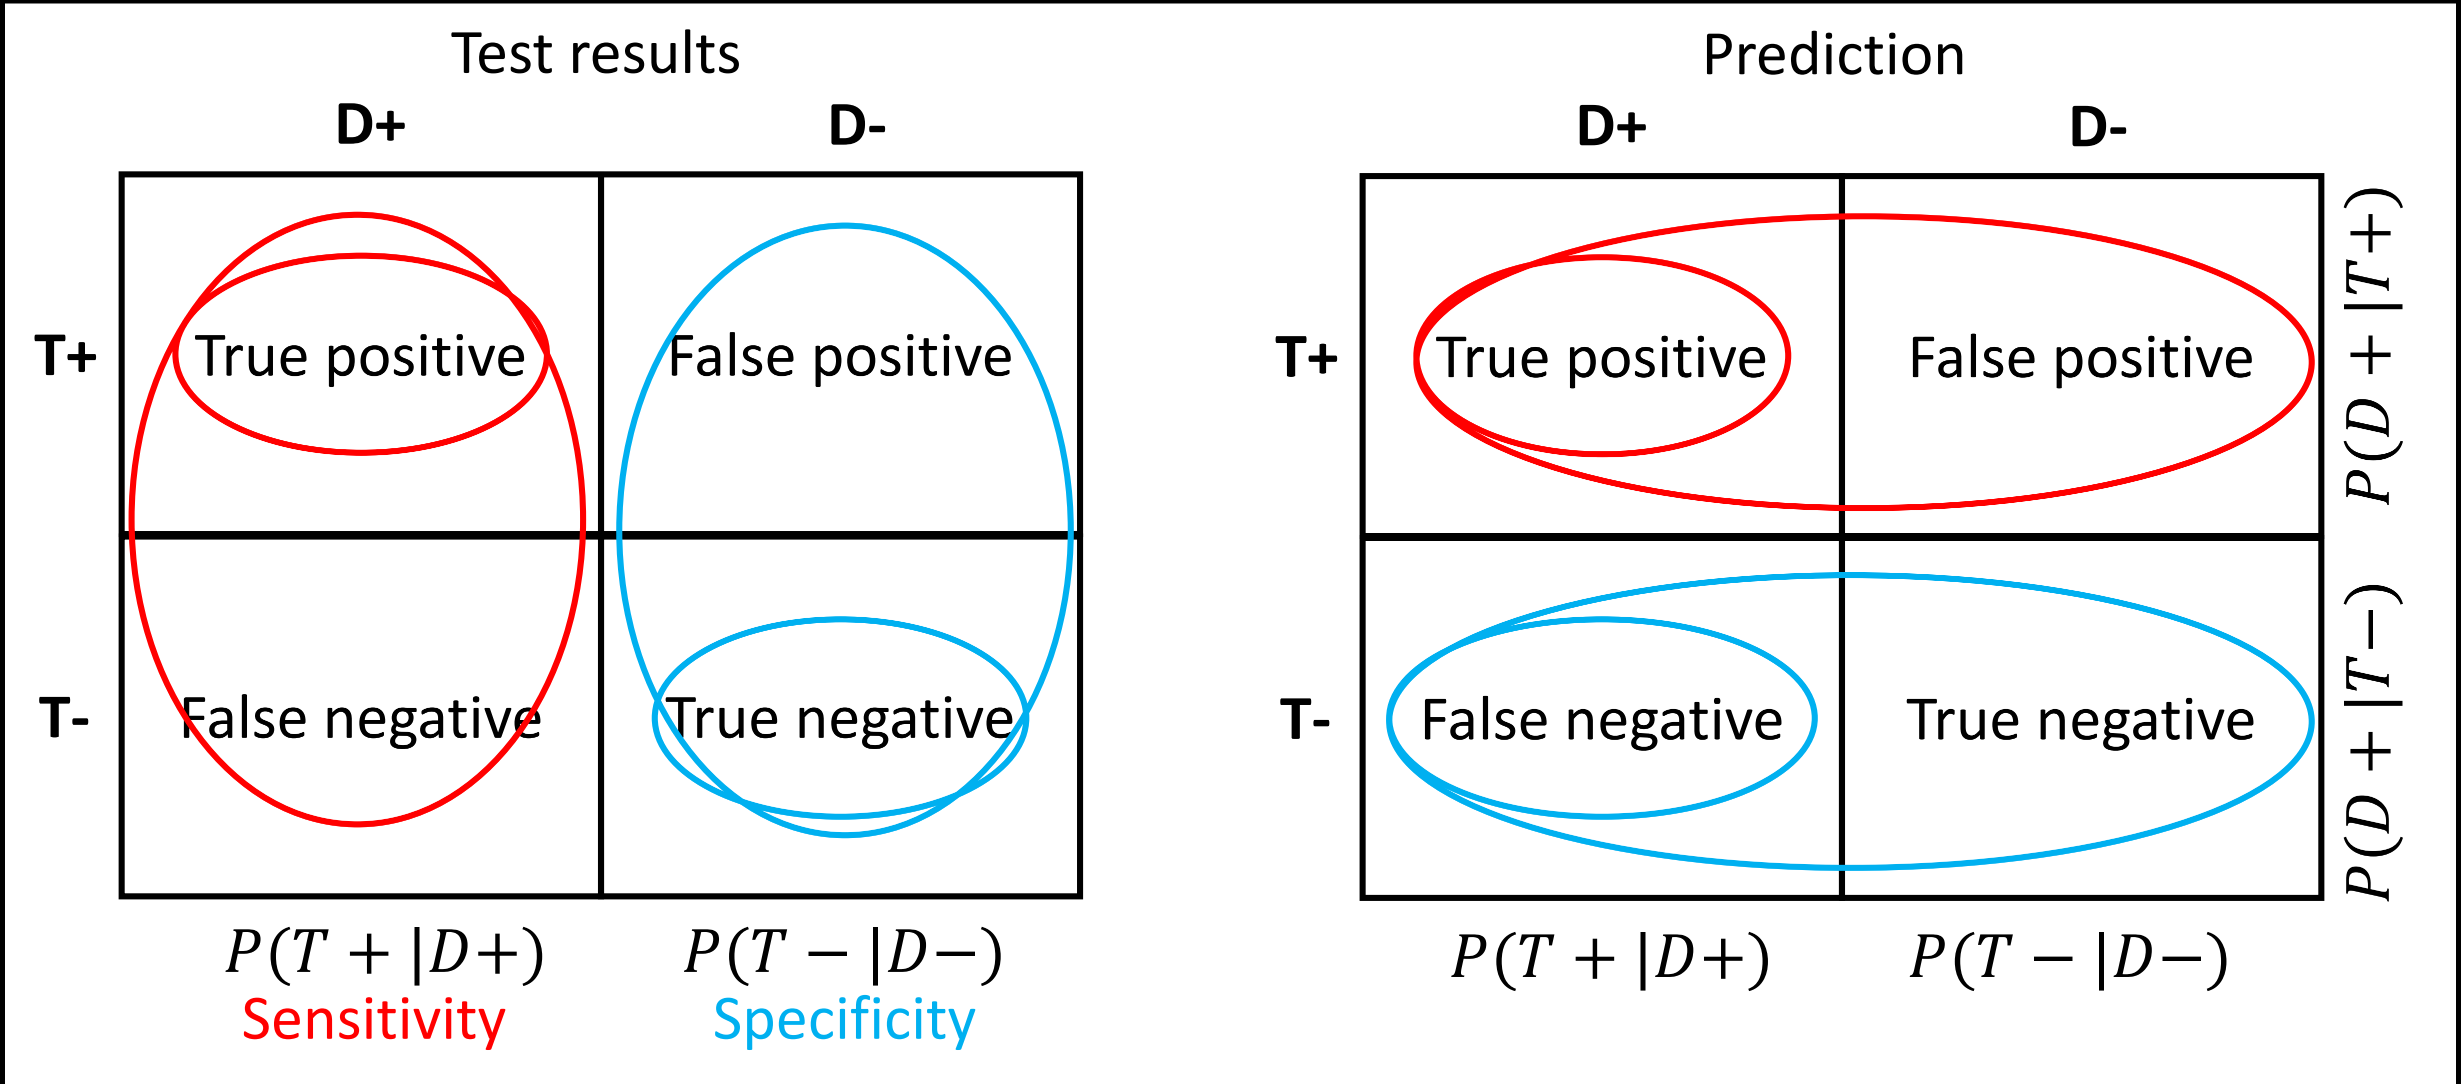

Bayesian Inference

- Unrelated example, borrowed from "Bayes with Beans" by Myriam Hunink, Harvard

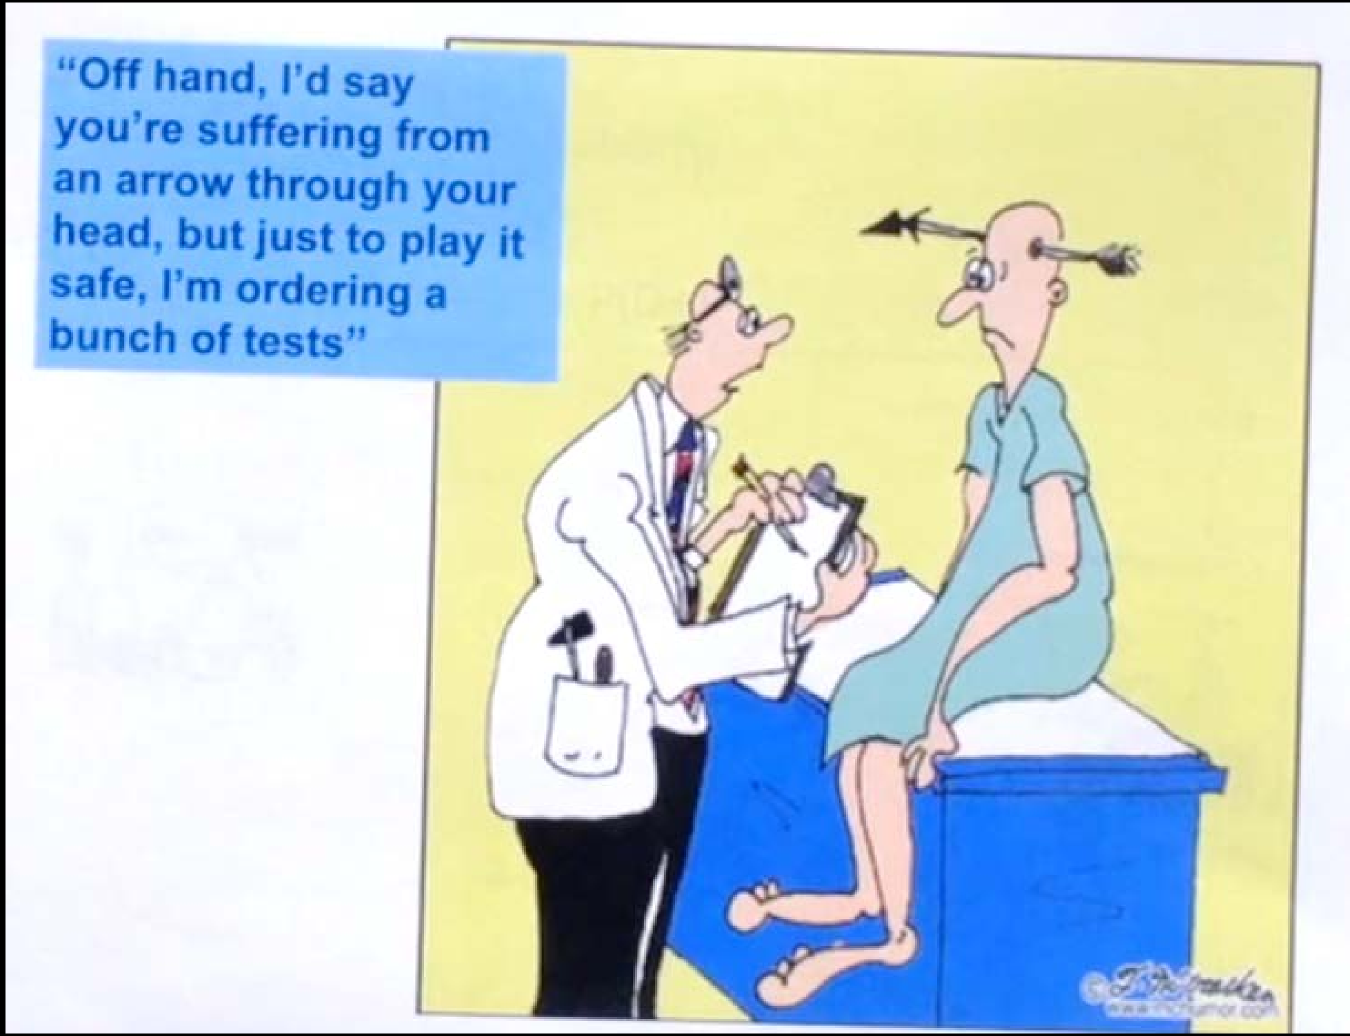

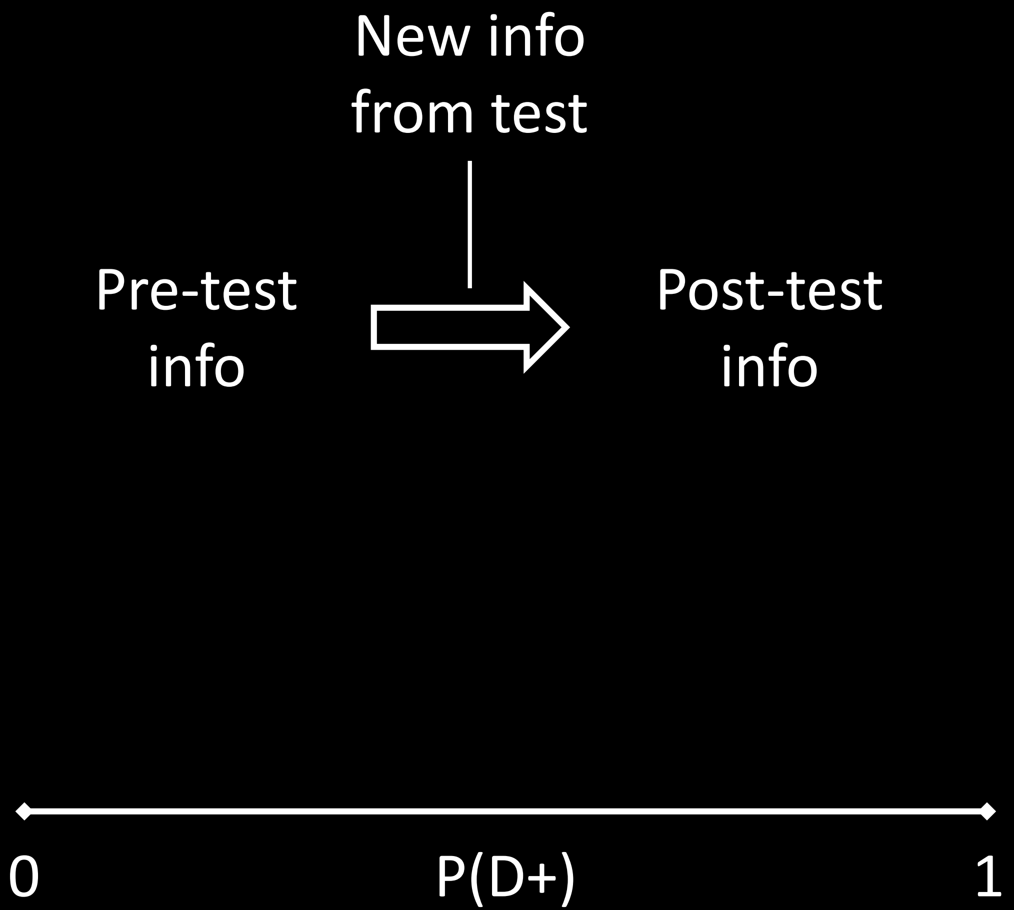

Probability Distribution

- Beliefs

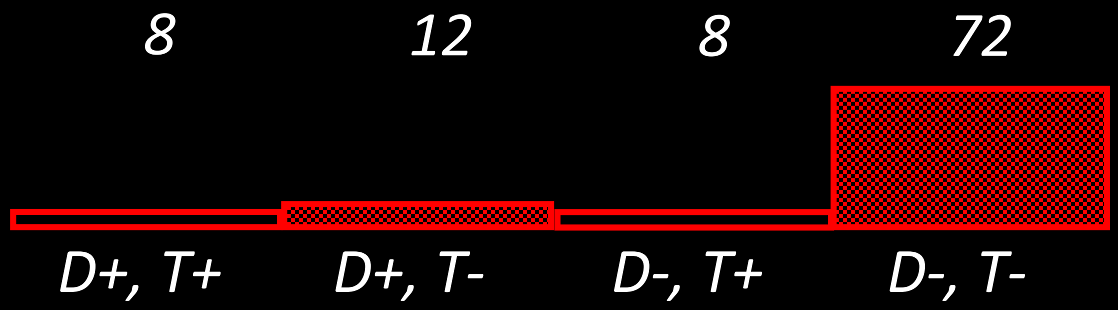

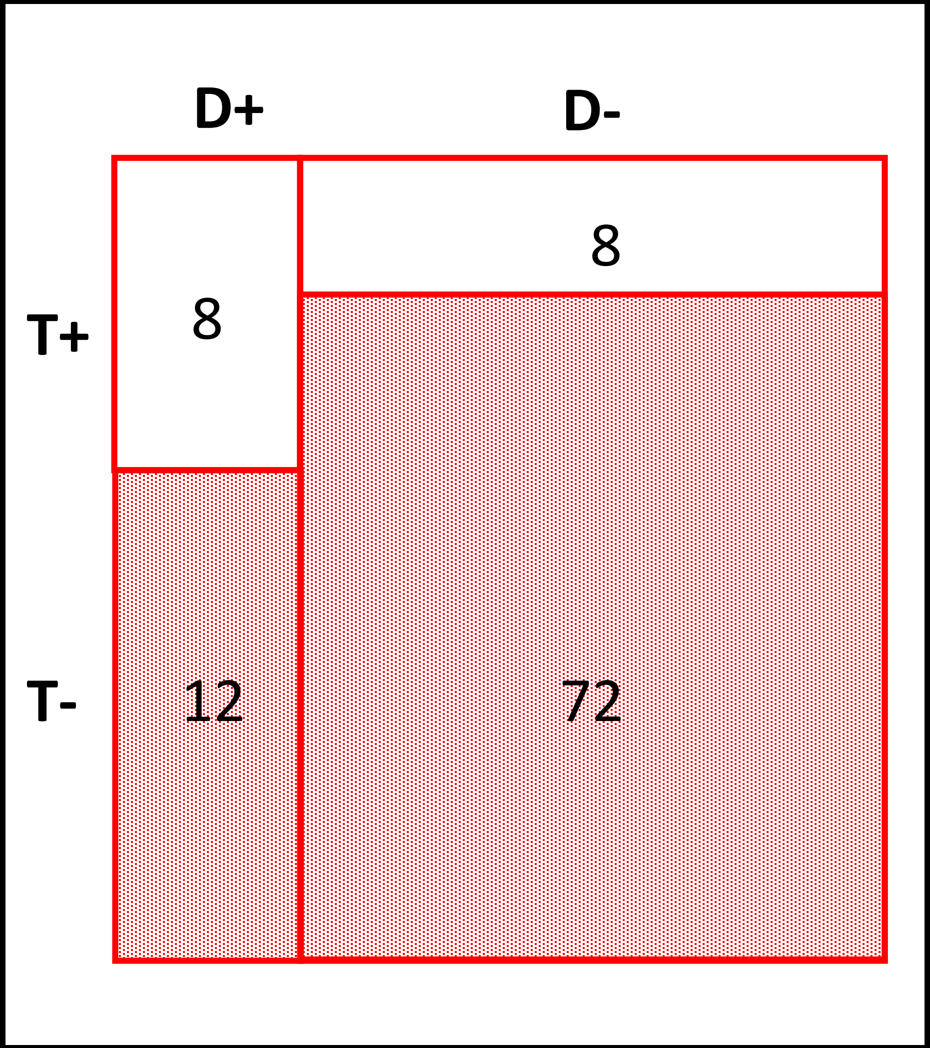

Probability Distribution

- Beliefs

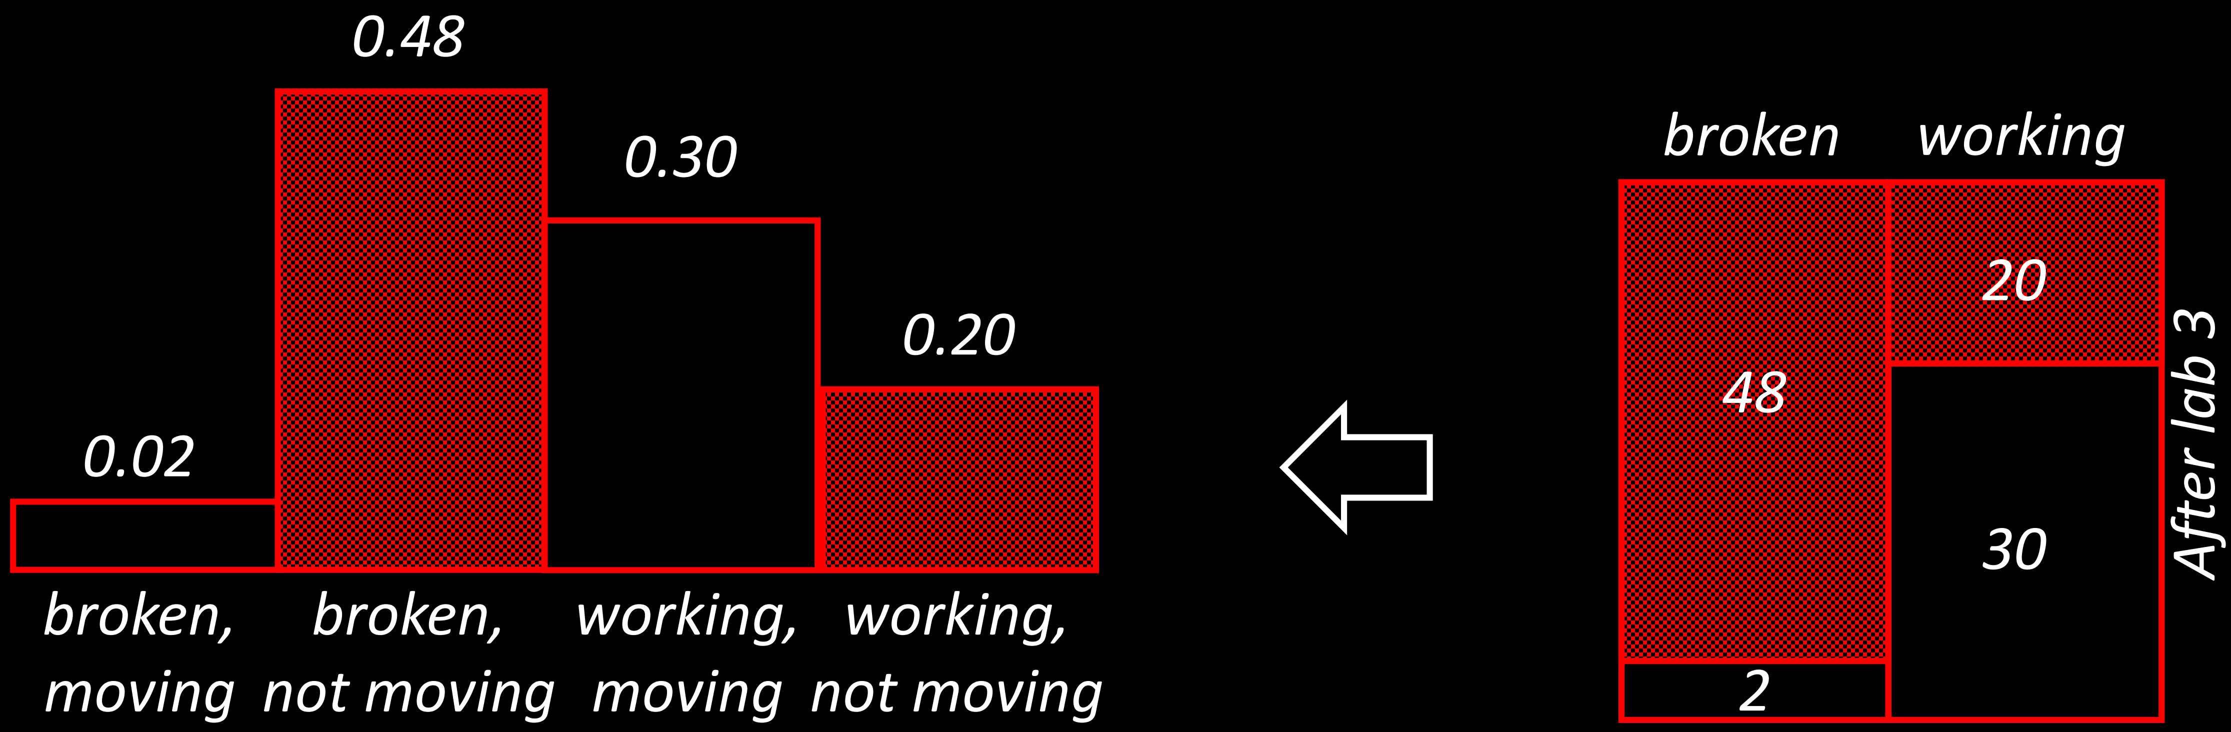

Probability Distribution

- Beliefs

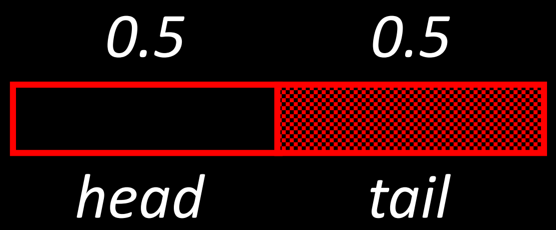

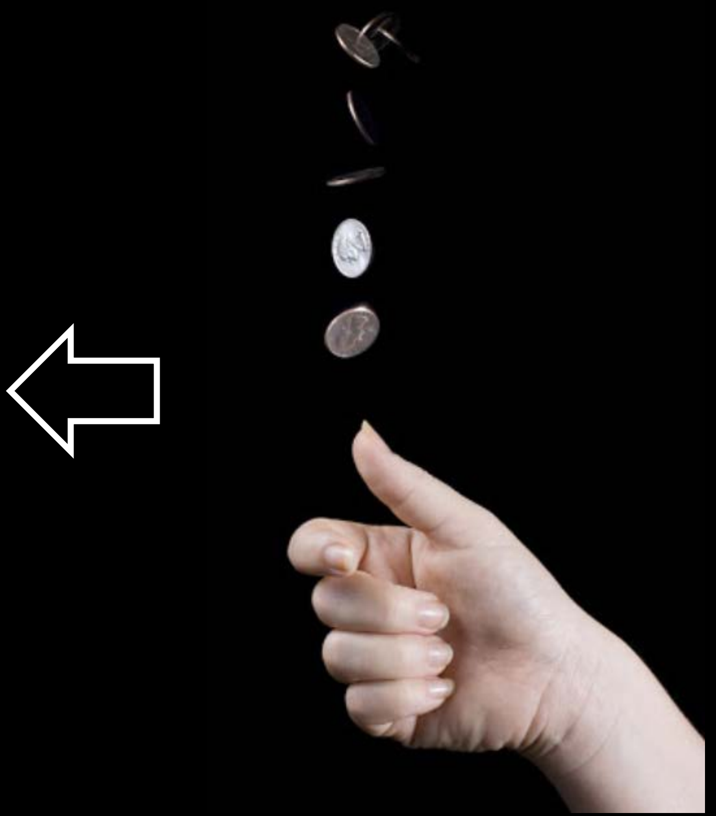

Probability Distribution

- Beliefs

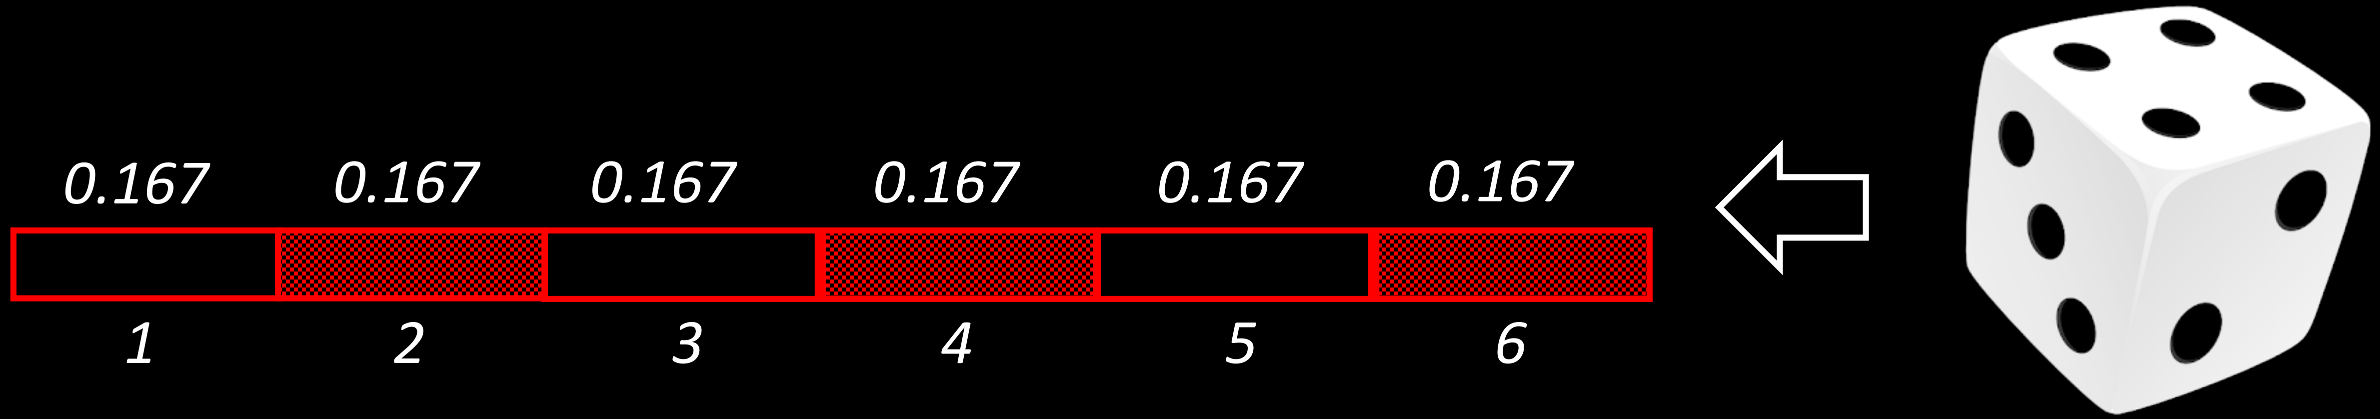

Probability Distribution

- Beliefs

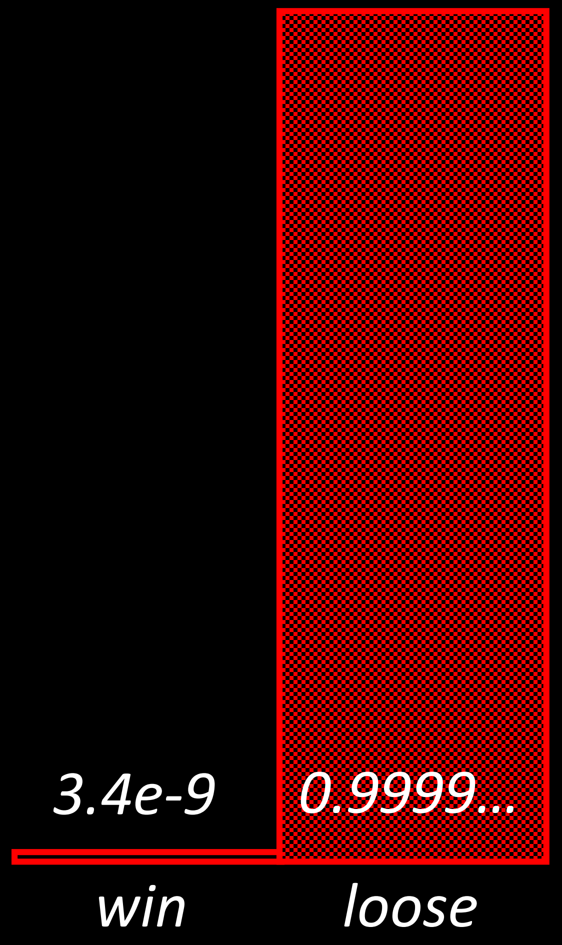

Probability Distribution

- Beliefs
- Discrete -> continuous *probability distribution*
  - Mean, median, most common value, etc.

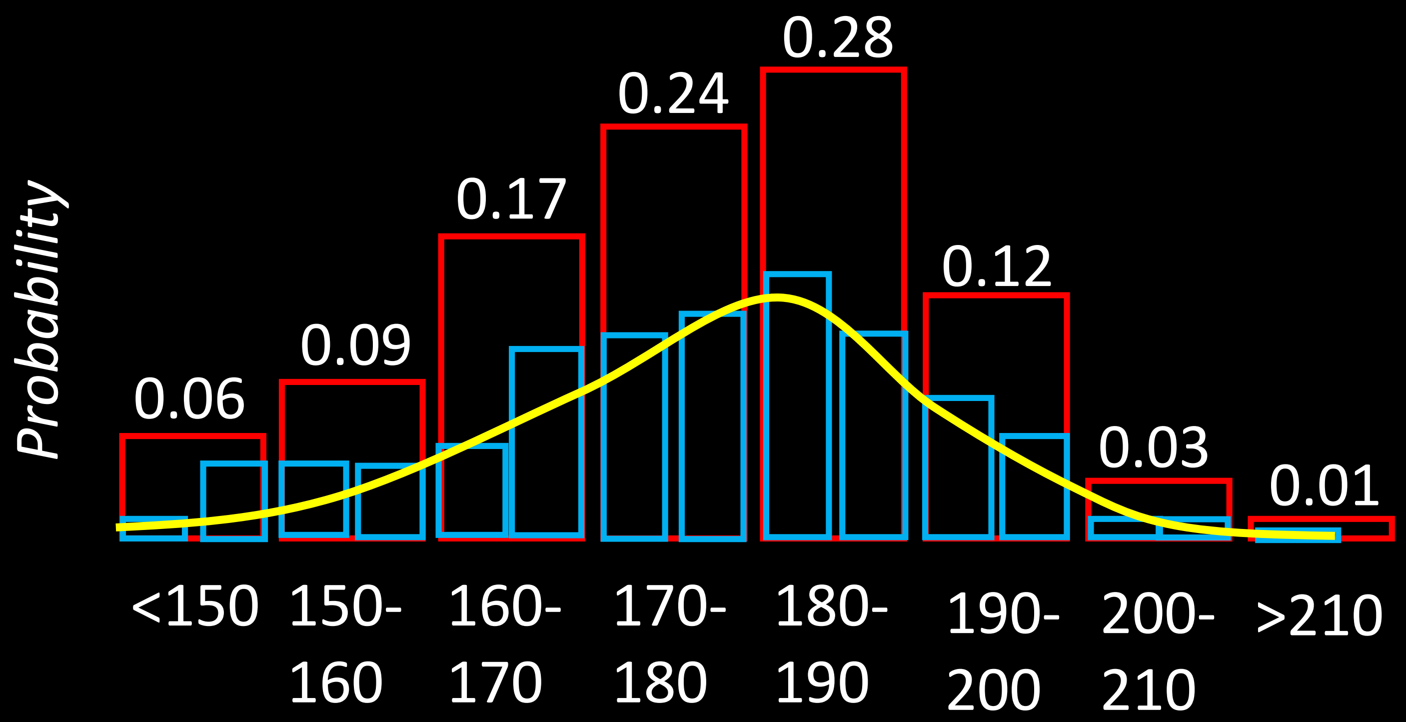

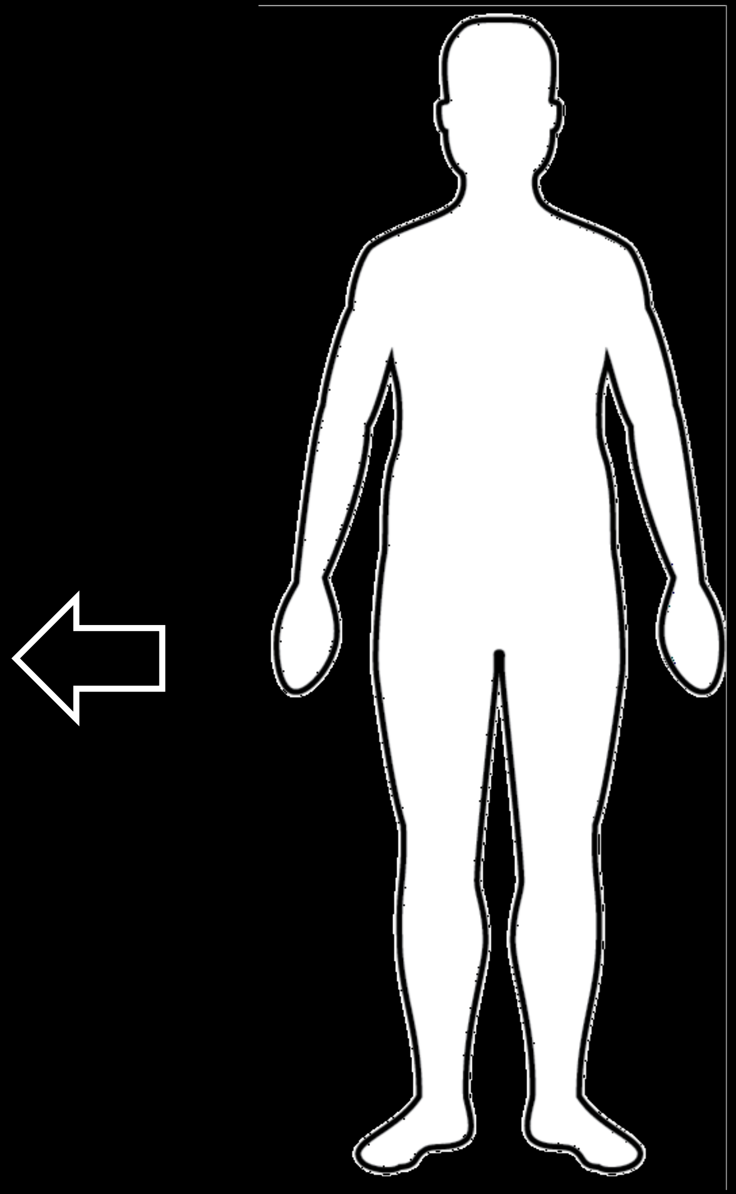

Probability Distributions

- What is the maximum speed of your robot?
  - You weigh 8.8 ft/s, 6.6 ft/s, 8.33 ft/s, but what is the actual value?
- Frequentist Statistics
  - Mean: $\mu = (8.8+6.6+8.33)/3 = 7.91$
  - Variance: $\sigma^2 = ((8.8-7.91)^2 + (6.6-7.92)^2 + (8.33-7.91)^2)/(3-1) = 1.35$
  - Standard deviation: $\sigma = \text{sqrt}(\sigma^2) = 1.16$
  - Standard error: $\sigma / \text{sqrt}(3) = 0.67$
- Bayesian Statistics
  - Probably 7.91ft/s…

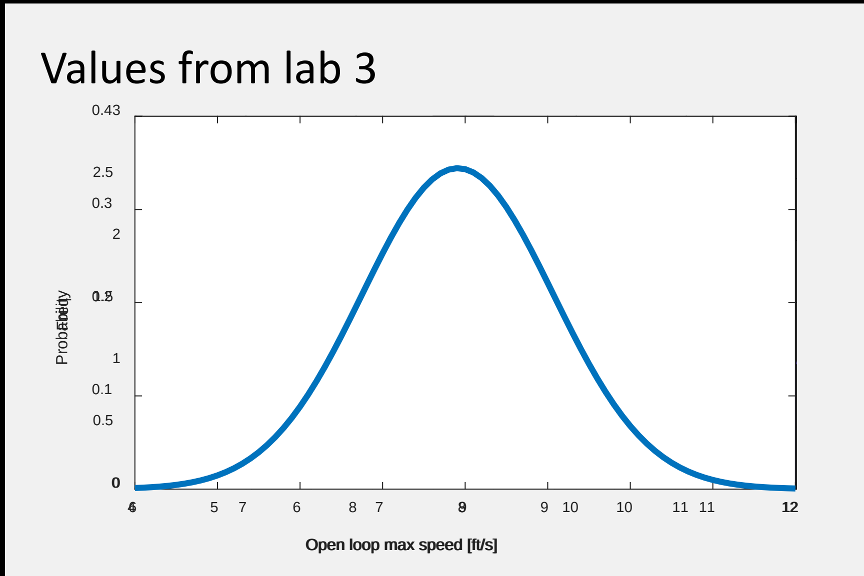

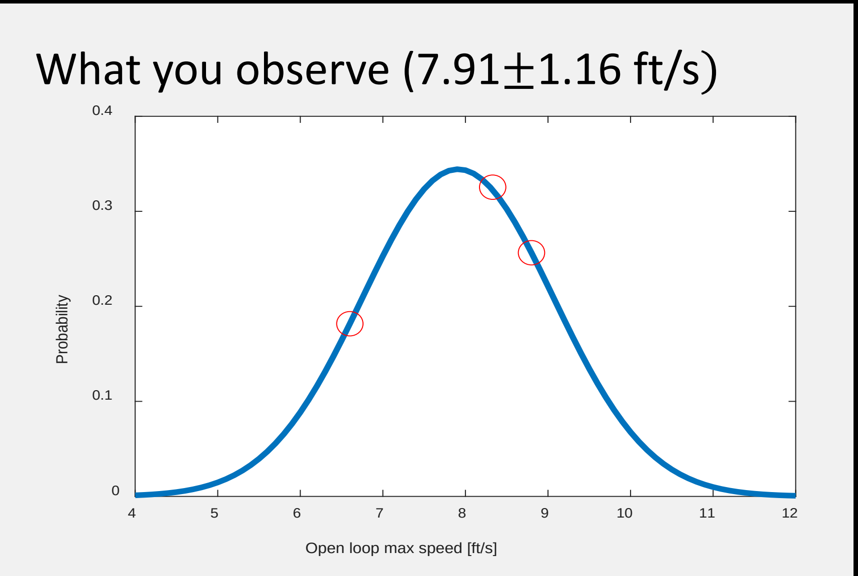

Probability Distributions

$$P(A|B) = \frac{P(B|A)\,P(A)}{P(B)}$$

- Use Bayes theorem
- Instead of A and B
  - Substitute “s” for the actual speed
  - Substitute “m” for the measurements
- P(s) is our prior
- P(m|s) is the likelihood associated with those measurements
- P(s|m) is what we believe about the speed given those measurements
- P(m) is the marginal likelihood
- Procedure:
  - Start with a belief
  - Update it
  - End up with a new belief!

Probability Distributions

$$P(A|B) = \frac{P(B|A)\,P(A)}{P(B)}$$

- Use Bayes theorem
- Start by assuming nothing
  - P(s) = uniform
  - P(s|m) = P(m|s)*c_1/c_2
  - Simplified: P(s|m) = P(m|s)
    - *Guess!* What if the actual max speed is 11 ft/s?
    - P(s=11|m=[6.6,8.33,8.8]) = P(m=[6.6,8.33,8.8] | s=11)
    - P(m = 6.6 | s = 11) * P(m = 8.33 | s = 11) * P(m = 8.8 | s = 11)

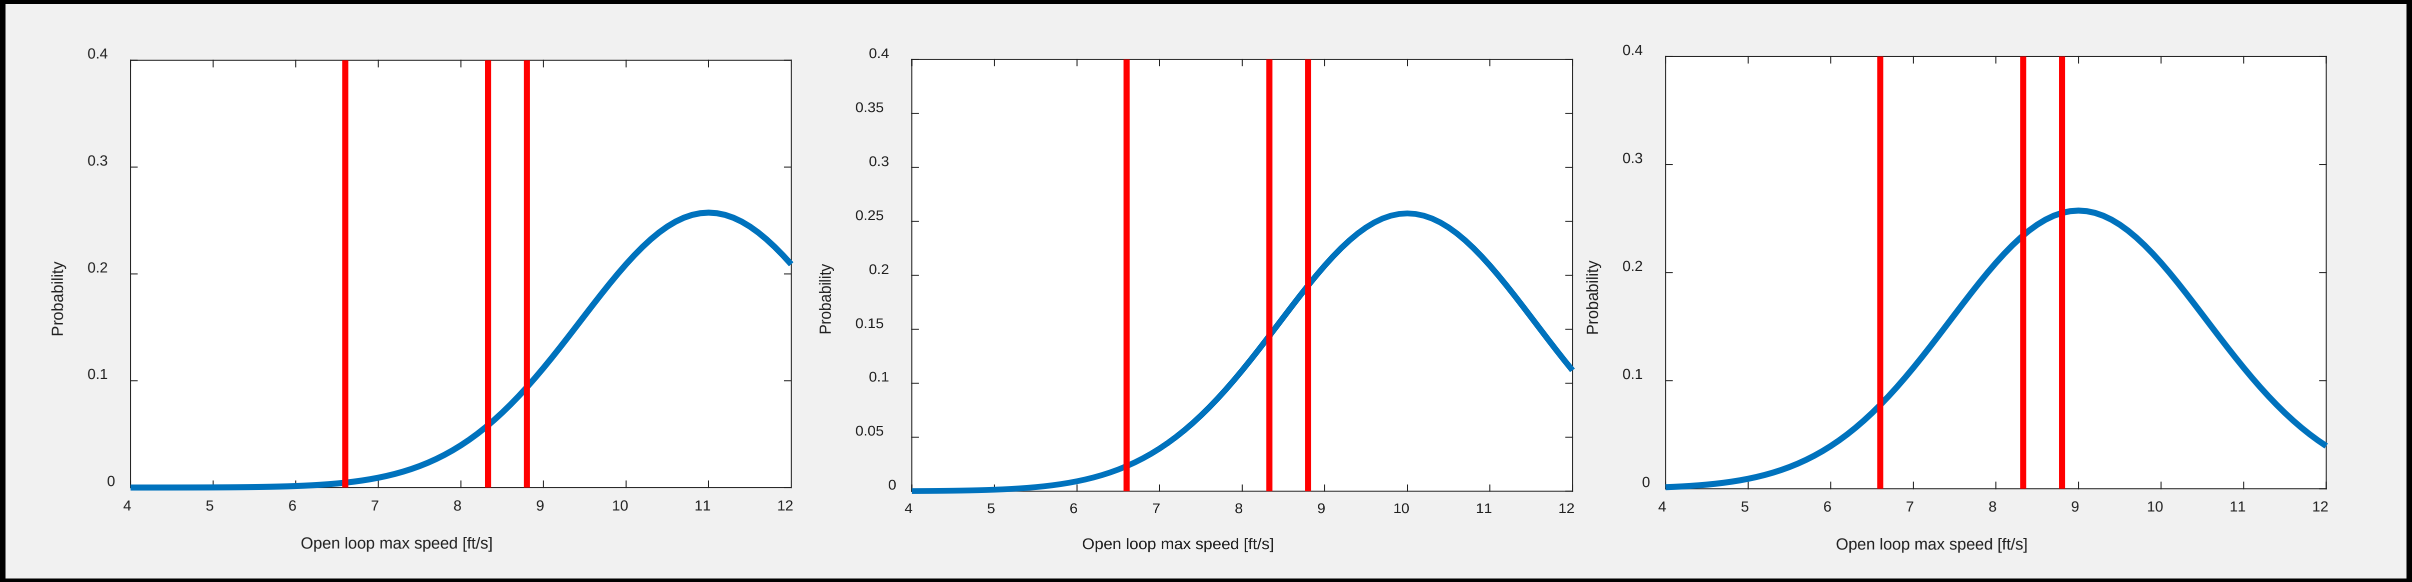

Probability Distributions

$$P(A|B) = \frac{P(B|A)\,P(A)}{P(B)}$$

- Use Bayes theorem
- Start by assuming nothing
  - P(s) = uniform
  - P(s|m) = P(m|s) * c_1/c_2
  - Simplified: P(s|m) = P(m|s)
    - Example, what if the actual max speed is 11 ft/s?
    - P(s = 11 | m = [6.6, 8.33, 8.8]) = P(m = [6.6, 8.33, 8.8] | s = 11)
    - P(m = 6.6 | s = 11) * P(m = 8.33 | s = 11) * P(m = 8.8 | s = 11)

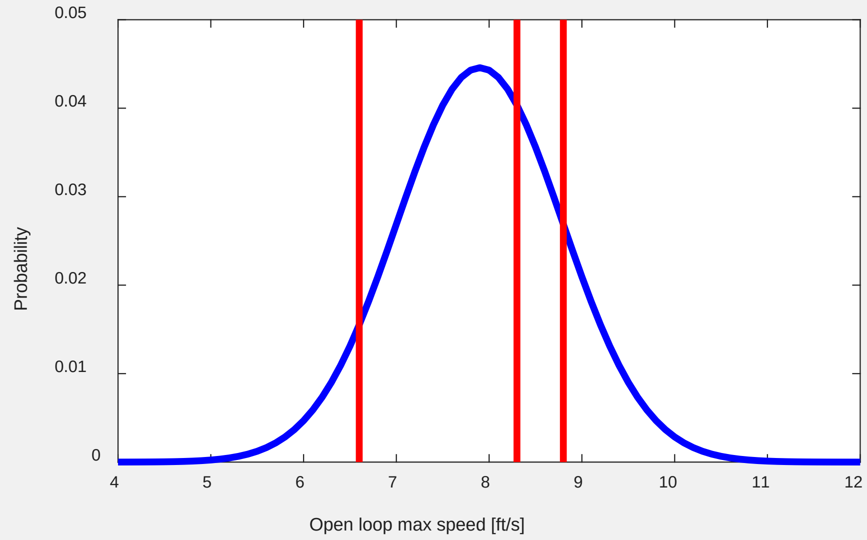

Probability Distributions

$$P(A|B) = \frac{P(B|A)\,P(A)}{P(B)}$$

- Use Bayes theorem
- Add a prior!
  - You know yesterday's speed, and you can kind of judge the current speed by eye
    - Prior: 7.91 ft/s $\pm$ 1.16ft/s
  - P(s = 11 | m = [6.6, 8.33, 8.8]) = P(m = [6.6, 8.33, 8.8] | s = 11) * P(s = 11)
    = P(m=6.6|s=11)*P(s=11) * P(m=8.33|s=11)*P(s=11) * P(m=8.8|s=11)*P(s=11)

Repeat the process!

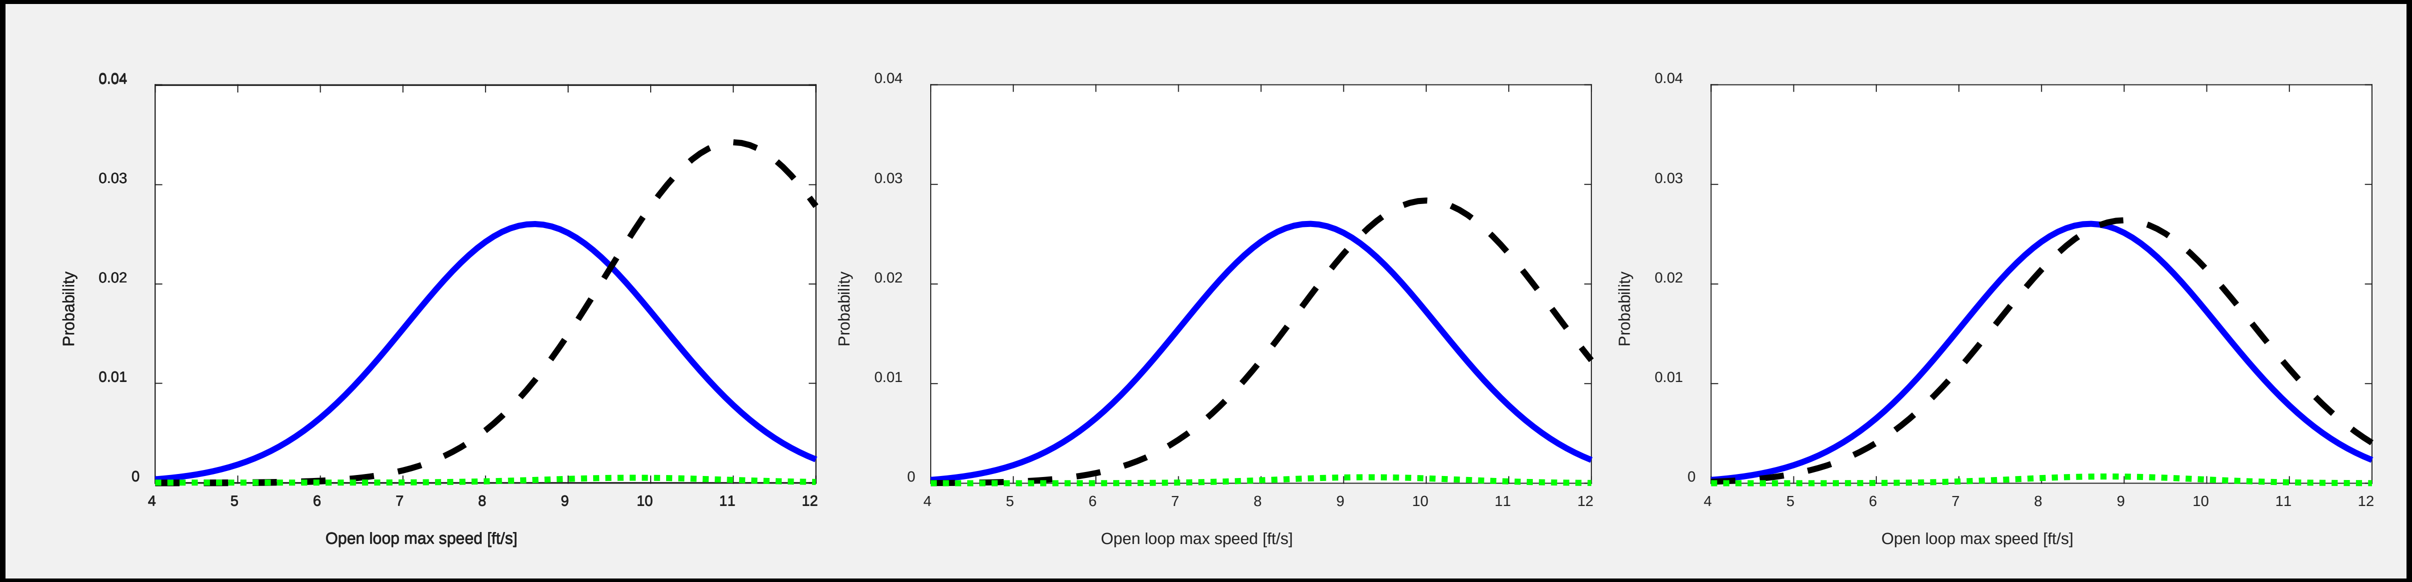

Probability Distributions

$$P(A|B) = \frac{P(B|A)\,P(A)}{P(B)}$$

- Use Bayes theorem
- Add a prior!
  - You know yesterday's speed, and you can kind of judge the current speed by eye
    - Prior: 7.91 ft/s $\pm$ 1.16ft/s
  - P(s = 11 | m = [6.6, 8.33, 8.8]) = P(m = [6.6, 8.33, 8.8] | s = 11) * P(s = 11)
    = P(m=6.6|s=11)*P(s=11) * P(m=8.33|s=11)*P(s=11) * P(m=8.8|s=11)*P(s=11)

Repeat the process!

Add everything up to get the posterior distribution

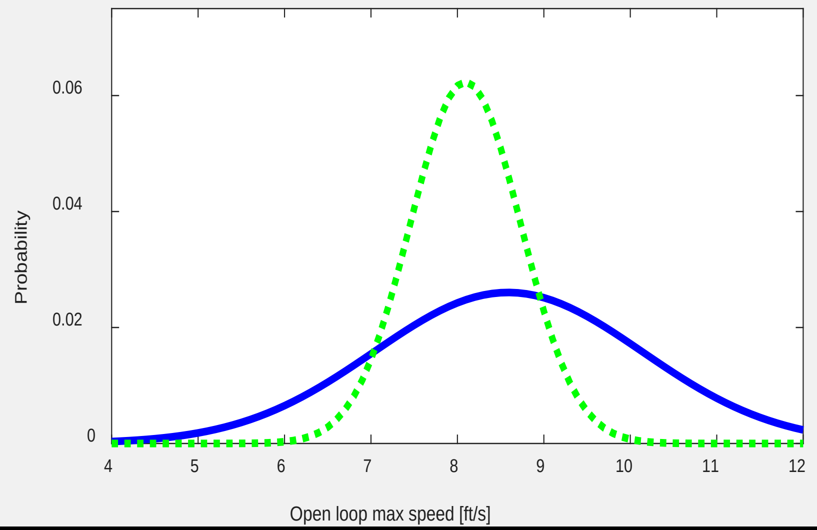

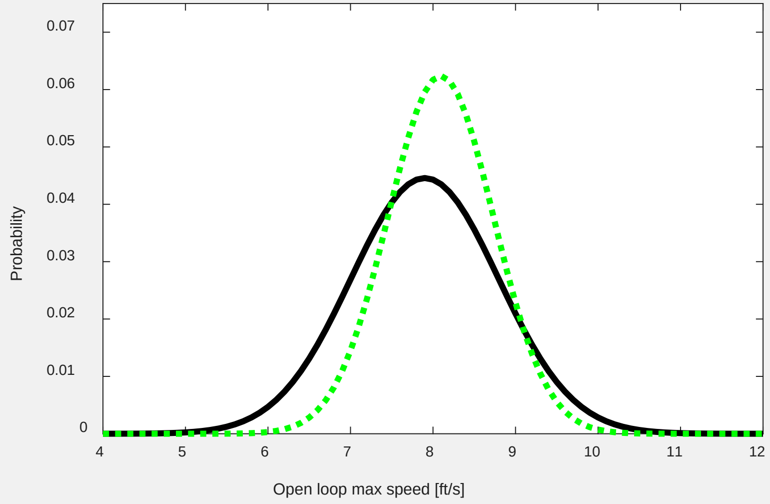

Maximum A Posteriori (MAP)

Probability Distributions

- Always believe the impossible, at least a little bit!
- Leave room for believing the unlikely. Leave a non-zero probability unless you are absolutely certain.
- “It ain’t what you don’t know that gets you into trouble. It’s what you know for sure that just ain’t so.” –Mark Twain
- “When you have excluded the impossible, whatever remains, however improbable, must be true. “ Sherlock Holmes (Sir Arthur Conan Doyle)

Alice’s adventures in wonderland

References

- Probabilistic Robotics, book by *Dieter Fox, Sebastian Thrun, and Wolfram Burgard*
- How Bayes Theorem works (Youtube), by Brandon Rohrer