

CS109 Lecture 14 Worksheet: ANSWER KEY

Beta

Summer 2026

Lecture 14 Worksheet: Answer Key

Instructor / TA copy. Problem-set solutions are hidden in the standard build; run `./convert-to-pdf.sh 14 solutions` for the full instructor/TA edition.

Problem 1 (Review): Five Quick Rolls

$$\binom{5}{2, 2, 1} \left(\frac{1}{6}\right)^2 \left(\frac{1}{6}\right)^2 \left(\frac{1}{6}\right)^1 = \frac{5!}{2! 2! 1!} \left(\frac{1}{6}\right)^5 = \frac{30}{7776} \approx 0.00386.$$

(The three faces we never roll have counts of 0, contributing factors of $0! = 1$ and $(1/6)^0 = 1$.)

Problem 2: The Robot Gripper

- (a) A uniform prior is $\text{Beta}(1, 1)$; add successes and failures: $X \sim \text{Beta}(1 + 6, 1 + 2) = \text{Beta}(7, 3)$.
- (b) $E[X] = \frac{a}{a+b} = \frac{7}{10} = 0.7$. Mode = $\frac{a-1}{a+b-2} = \frac{6}{8} = 0.75$. The posterior is skewed (bounded above by 1), so the mean is pulled below the peak. The mode equals the frequentist estimate $6/8$; the mean 0.7 is between the frequentist estimate and the prior mean 0.5 (the prior acts like one imagined success and one imagined failure).
- (c) A single hump on $[0, 1]$, peaked at 0.75, not symmetric: more probability mass trails off to the left than to the right (it must fit between the peak and 1).

Problem 3: Updating a Beta Prior

- (a) Conjugacy: $\text{Beta}(4 + 3, 2 + 5) = \text{Beta}(7, 7)$. (Symmetric, centered at $E[X] = 1/2$.)
- (b) $\text{Beta}(4, 2)$ acts like having already seen $a - 1 = 3$ imagined successes and $b - 1 = 1$ imagined failure on top of a uniform prior.
- (c) Posterior is $\text{Beta}(2 + s, 2 + (n - s))$, with mode

$$\frac{(2 + s) - 1}{(2 + s) + (2 + n - s) - 2} = \frac{s + 1}{n + 2}.$$

This is the classic Laplace-smoothed estimate: never exactly 0 or 1, even when $s = 0$ or $s = n$.

Problem 4: Puppy Training (Fall 2025 Midterm)

Uniform prior means $P \sim \text{Beta}(1, 1)$; after 3 successes and 1 failure, $P \sim \text{Beta}(4, 2)$. Then

$$\text{Var}(P) = \frac{ab}{(a+b)^2(a+b+1)} = \frac{4 \cdot 2}{6^2 \cdot 7} = \frac{8}{252} = \frac{2}{63} \approx 0.0317.$$

Problem 5: Street Parking (Fall 2023 Final)

A uniform prior between 0 and 1 is $\text{Beta}(a=1, b=1)$. We observe 1 success (open spot) and 9 failures (full spots):

$$X \sim \text{Beta}(2, 10), \quad E[X] = \frac{2}{12} = \frac{1}{6}.$$

Problem 6: Medicine with a Uniform Prior (pset4: warmup_beta)

(Problem set problem)

Problem 7: Thompson Sampling (pset4: learning_while_helping)

- (a) Drug A: $\text{Beta}(1+6, 1+2) = \text{Beta}(7, 3)$, $E = 7/10 = 0.7$. Drug B: $\text{Beta}(1+1, 1+1) = \text{Beta}(2, 2)$, $E = 2/4 = 0.5$.
- (b) Draw one sample θ_A from A's Beta and one sample θ_B from B's Beta; give the patient whichever drug produced the larger sample. Then observe the outcome and add it to that drug's Beta parameters (success adds 1 to a , failure adds 1 to b).
- (c) Because each choice is made from *samples*, a drug whose belief is wide (few observations, high uncertainty) will regularly produce large samples and get chosen, even if its mean is currently lower. This automatically balances exploration (gathering data about uncertain options) with exploitation (using the option that looks best), and the exploration fades naturally as the Betas sharpen with data.

Challenge Problem: Two Beta Facts

- (a) With $a = b = 1$: $x^{a-1}(1-x)^{b-1} = x^0(1-x)^0 = 1$, and $B(1, 1) = \int_0^1 1 dx = 1$. So $f(x) = 1$ on $[0, 1]$, exactly the $\text{Uni}(0, 1)$ PDF.
- (b) By the law of total probability, conditioning on $X = x$:

$$P(\text{heads}) = \int_0^1 P(\text{heads} \mid X = x) f(x) dx = \int_0^1 x f(x) dx = E[X] = \frac{a}{a+b}.$$

So the single best point prediction for the next flip is the mean of your belief.