CSCI 145 Problem 5: The Best Possible Forecast

August 27, 2026

Submission Instructions

Please upload your work to Gradescope by **11:59pm on DATE**.

- Turn in a single PDF of *handwritten* solutions. Typed (including LaTeX) solutions will not be accepted.
- Problems are grouped into three-week **cycles**. Once per cycle, you will meet with the grader and present the solution to *one* problem from that cycle, selected at random when you arrive — so bring your work and be ready to explain any of them.

Problem 5: The Best Possible Forecast — 2 points

Let (X, Y) have a known joint distribution, and score a predictor f by $R(f) = \mathbb{E}[(Y - f(X))^2]$. The regression function is $f^*(x) = \mathbb{E}[Y \mid X = x]$.

Prove the following risk decomposition for every predictor f:

$$R(f) = \mathbb{E}[\mathrm{Var}(Y \mid X)] + \mathbb{E}\big[(f^*(X) - f(X))^2\big].$$

Your proof should identify the term that could prevent the two squares from separating and justify carefully, using conditional expectation, why that term is zero. Use the identity to establish both the best possible predictor and its irreducible error, including the best constant predictor as a special case.

Test the theorem on the lecture's sales example, whose joint probabilities are

	$Y = 10$	$Y = 30$
$X =$ cloudy	0.30	0.10
$X =$ sunny	0.15	0.45

Compute the optimal forecast for each weather condition, the irreducible error, and the excess error paid by the best constant forecast. Verify that these two errors add to the constant predictor's total risk, and use the decomposition to assess a model that reports zero squared error on genuinely noisy test data.

Results you may use. For any integrable random variable Z, $\mathbb{E}[Z] = \mathbb{E}[\mathbb{E}[Z \mid X]]$. A quantity determined by X can be pulled outside an expectation conditional on X, and $\mathrm{Var}(Y \mid X) = \mathbb{E}[(Y - \mathbb{E}[Y \mid X])^2 \mid X]$.