CSCI 145 Problem 2: The Optimal Control Variate

July 20, 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 2: The Optimal Control Variate

In lecture we halved a Monte Carlo estimator's error without rolling a single extra die: we subtracted off a surrogate whose average we already knew, scaled by a coefficient guessed out of the air ($c = \frac{1}{2}$). Was the guess lucky, and how much can the trick ever buy us? In Problem 1 correlation was the villain that put a floor under a poll; here it answers both questions.

We want $\mu = \mathbb{E}[f(X)]$, and alongside the target $f(X)$ we have a *surrogate* $g(X)$ whose expectation $\mathbb{E}[g(X)]$ we know exactly. In class we built the adjusted Monte Carlo estimator

$$\hat{\mu}_c = \frac{1}{n}\sum_{i=1}^{n}\Big[f(X^{(i)}) - c\big(g(X^{(i)}) - \mathbb{E}[g(X)]\big)\Big],$$

where $X^{(1)}, \ldots, X^{(n)}$ are independent samples distributed like X, and we proved it unbiased for *every* constant c. So all choices of c are correct on average; they differ only in variance.

Part A: The Variance Is a Parabola

Show that

$$\mathrm{Var}(\hat{\mu}_c) = \frac{1}{n}\Big(\mathrm{Var}(f(X)) - 2c\,\mathrm{Cov}(f(X), g(X)) + c^2\,\mathrm{Var}(g(X))\Big).$$

(Hint: the samples are independent, so the variance of the average is the variance of one adjusted sample divided by n, the $\rho = 0$ case of Problem 1. Then expand the variance of $f(X) - c\,g(X)$ with the variance-of-a-sum identity from class; constants shift nothing.) As a function of c, is this parabola smiling or frowning, and why does that matter for what comes next?

Part B: The Bottom of the Parabola

Find the coefficient $c^\star$ minimizing $\mathrm{Var}(\hat{\mu}_c)$ by setting the derivative in c to zero. Your answer should be

$$c^\star = \frac{\mathrm{Cov}(f(X), g(X))}{\mathrm{Var}(g(X))}.$$

Sanity-check it: what does the formula say when f and g are independent, and does that match what the adjustment *should* do?

Part C: How Much It Saves

Substitute $c^\star$ back into Part A and show that

$$\mathrm{Var}(\hat{\mu}_{c^\star}) = \frac{\mathrm{Var}(f(X))}{n}\left(1 - \rho^2\right), \qquad \text{where} \quad \rho = \frac{\mathrm{Cov}(f(X), g(X))}{\sqrt{\mathrm{Var}(f(X))\,\mathrm{Var}(g(X))}}$$

is the correlation between target and surrogate. Now apply it to the class example, with f the larger of two dice and g their sum, using the exact moments from enumerating the 36 outcomes:

$$\mathrm{Var}(f(X)) = \tfrac{2555}{1296}, \qquad \mathrm{Cov}(f(X), g(X)) = \tfrac{35}{12}, \qquad \mathrm{Var}(g(X)) = \tfrac{35}{6}.$$

Compute $c^\star$ and the reduction factor $1 - \rho^2$, compare $c^\star$ against the value lecture guessed, and say how many games the adjustment saves. Conclude in one sentence: what two properties make a surrogate g worth having, and which of them decides *how much* it is worth?

Where this goes next. Stochastic gradient descent estimates the gradient by Monte Carlo, so everything you proved here applies to it verbatim, and in the final week the variance-minimizing *baseline* an agent subtracts from its returns turns out to be exactly your $c^\star$.