CSCI 145 Problem 23: Choose a Horizon Under Persistent Noise

October 27, 2026

Submission Instructions

Please upload your work to Gradescope by **11:59pm on Wednesday, December 9**.

- Turn in a single PDF of *handwritten* solutions. Typed (including LaTeX) solutions will not be accepted.
- A complete, good-faith attempt earns submission credit. The grader may select any recent assigned problem for a brief conversation, so keep your work and be prepared to explain each solution and answer a related question. Scheduling details will be announced in class.

Problem 23: Choose a Horizon Under Persistent Noise

An agent ultimately cares about the far-sighted value

$$V_{\text{target}} = V_{\gamma^\star} = \frac{\mu}{1-\gamma^\star}, \qquad 0 < \gamma^\star < 1, \qquad \mu \neq 0.$$

Across $M \geq 1$ independent episodes, a persistent environmental condition shifts every reward within episode i by the same random amount. Write $r_{i,t}$ for its reward at time t:

$$r_{i,t} = \mu + Z_i, \qquad \mathbb{E}[Z_i] = 0, \qquad \text{Var}(Z_i) = \tau^2.$$

The episode shocks are independent of one another. To reduce estimation noise, the agent may average the shorter returns

$$\overline{G}_{\gamma,M} = \frac{1}{M}\sum_{i=1}^{M}\sum_{t\geq 0}\gamma^t r_{i,t}$$

for any $0 \leq \gamma \leq \gamma^\star$, accepting some bias relative to V_{target}.

1. What value of γ minimizes $\mathbb{E}[(\overline{G}_{\gamma,M} - V_{\text{target}})^2]$, including the exact boundary at $\gamma = 0$?
2. If each persistent Z_i is replaced by noises $Z_{i,t}$ that are independent across time and episodes with the same mean and variance, how does the variance of $\overline{G}_{\gamma,M}$ grow as $\gamma \to 1$ compared with the persistent-shock model?

Hint: You may use the infinite geometric-series identities and the variance rules for independent sums. A change of variable such as $t = (1-\gamma)^{-1}$ is permitted but not required.

Takeaway: *Noise that persists through an episode accumulates much faster in long returns than independent noise does. A shorter horizon can give a more accurate estimate overall, even though it introduces bias. The best horizon depends on both the noise and the number of episodes.*