Two Tales, One Insight: Today as the Safest Bet for Tomorrow

The Python Quants with ChatGPT

June 14, 2025

Contents

1 “Tomorrow Looks a Lot Like Today”

Maya loved two things about her small seaside town of Breezewood—the old lighthouse where her grandmother once studied the weather, and the little café where her brother Leo watched the stock market between serving lattes.

One summer afternoon their ten-year-old cousin, Sam, burst in with a homework riddle:

> “Explain why, if you may pick only *one* number, the smartest way to guess tomorrow’s stock price **and** tomorrow’s temperature is to use today’s value.”

Sam flopped into a booth. “Weather and money have nothing in common,” he sighed.

1.1 A Lighthouse Lesson

Maya led him up the spiral stairs of the lighthouse. At the top, a brass thermometer read `22 °C`.

“If I know nothing else,” she said, “my safest single-number guess for tomorrow is `22 °C` again. Some days I’ll be off, but *on average* that guess keeps the squared error—the ‘oops distance’—as small as possible.”

1.2 The Café Counter

Back in the café, Leo slid a hot chocolate toward Sam and pointed at today’s closing price of HarborTech:

$$P_t = \texttt{€47.10}.$$

“Same trick here,” Leo said. “With no fresh information, my best forecast for tomorrow is $\hat{P}_{t+1} = P_t$. Finance folks call that the *random-walk hypothesis*.”

1.3 The *A-ha* Moment

Sam’s eyebrows rose. “So both the sea breeze and the stock screen are like a marble nudged by random taps?”

“Exactly,” Maya and Leo chimed. Sam scribbled his answer:

> When you’re allowed just one number and no other clues, yesterday’s value is the fairest bet for both weather and stocks. It minimises the mean-squared forecast error.

He raced home, already picturing the teacher’s smile.

2 "The One-Step Forecast"

Sasha, a second-year finance student, hurried across campus clutching a print-out of HarborTech's price history. She found Professor Alvarez, whiteboard marker in hand, waiting under Finley Hall's colonnade.

2.1 A Stock Chart in Chalk

Alvarez wrote the simplest of rules:

$$\hat{P}_{t+1} = P_t. \tag{1}$$

"This is the random-walk hypothesis," he said. "With no new information, the conditional expectation of tomorrow's price, given everything we know today, is today's price. In least-squares terms, yesterday wins the contest."

2.2 Turning to Temperature

Flipping the board, he produced a coastal weather log and wrote the analogous rule:

$$\hat{T}_{t+1} = T_t. \tag{2}$$

"If you must rely on a single datum," Alvarez continued, "today's thermometer reading is your best least-squares predictor for tomorrow. Measurements from one, two, or four weeks ago add essentially no extra punch for a *one-day* horizon—just as ancient prices add nothing once you see P_t."

2.3 A Shared Logic

Alvarez summarised:

$$\begin{aligned} \textbf{Stocks:} \quad & E[P_{t+1} \mid P_t, P_{t-1}, \dots] = P_t, \\ \textbf{Weather:} \quad & E[T_{t+1} \mid T_t, T_{t-1}, \dots] = T_t. \end{aligned}$$

"Both systems are *Markovian* for the next step," he said. "Whatever last month knew is already encoded in today's value."

Sasha exhaled. "My fancy models were squeezing juice from fruit that's already been pressed."

Alvarez laughed. "Nicely put. Of course, if an earnings surprise or a cold front arrives, you update. But absent that, yesterday is your statistically safest bet."