Penney's Game: The Short Version

CS109 Project — Neshanth Anand 2/22/26

What Is Penney's Game?

Pick any sequence of 3 coin flips — say **HHT**. The machine picks its own sequence *after* seeing yours. You flip a coin repeatedly in a stream, and whoever's pattern appears first wins.

It looks fair. You go first. Yet no matter what you pick, the machine wins **66–87% of the time**. Every time. This is not a trick — it is probability theory.

The Best Choices (Read This First)

Before the math: if you want to give yourself the best shot, pick **HTH** or **THT**.

These are the only two patterns that force the machine into its weakest possible counter, where it wins just **66.7%** of the time instead of 75% or 87.5%. You still lose more often than not — but 1 in 3 is a real chance. Every other pattern hands the machine a larger edge.

No pattern gives you the advantage. Non-transitivity means there is no "best" sequence in absolute terms — the optimal counter always exists. But HTH and THT make it as close as it can get.

Why Naive Thinking Fails

Both sequences have probability $(1/2)^3 = 1/8$ of appearing in any isolated window of 3 flips, so the game should be 50/50 — right?

Wrong. We are not looking at isolated windows. We are watching a *continuous stream*, where sequences overlap. The flip that ends one candidate sequence might be the flip that starts another. That overlap structure is deeply asymmetric, and it is exactly what the machine exploits.

The Framework: States and LOTP

At any point in the stream, all that matters is the longest suffix of flips so far that is a prefix of either sequence. That suffix is the **state**. Each sequence is 3 flips, so states have length 0, 1, or 2 — length 3 means someone already won. The maximum number of non-terminal states is:

$$1 + 2 + 4 = 7$$

For each non-terminal state X, define p_X as the probability you win given the current state is X. The Law of Total Probability conditions on the next flip:

$$p_X = \frac{1}{2} \cdot p_{X \xrightarrow{H}} + \frac{1}{2} \cdot p_{X \xrightarrow{T}}$$

with boundary conditions: value = 1 if you win, value = 0 if machine wins. One equation per state. Solve the linear system. Done.

Death States

A **death state** is any state where one player wins with certainty — formally an absorbing state. Once entered, there is no escape.

- $p_X = 0$: you are trapped. Every path leads to machine winning.
- $p_X = 1$: machine is trapped. You win with certainty.

The machine's counter is engineered so its death trap for you is easy to reach — often on the very first flip — while your death trap for the machine requires a lucky corridor.

The Three Outcomes

There are 8 possible sequences. By H/T symmetry there are 4 structurally distinct ones, which collapse into exactly **3 distinct win rates**:

87.5% — All same (HHH, TTT). Machine's trap is reached on flip 1 with probability 1/2 and is perfectly absorbing. Your best state gives only a 50% chance even after reaching it. The machine's structural advantage is overwhelming.

75% — One repeated letter (HHT, HTT, THH, TTH). Your death state for the machine becomes perfect ($p = 1$) — a real improvement. But the machine's trap is still reachable on flip 1 with probability 1/2, while yours requires two consecutive matching flips (probability 1/4). That reachability gap is what keeps the machine at 75%.

66.7% — Fully alternating (HTH, THT). Both sequences share a first letter, so the machine cannot engineer a first-flip trap. Both death states are reached with equal probability from the same state — but the machine's trap is perfectly absorbing while yours is leaky (you escape back to Start with probability 1/2). That asymmetry alone is enough to keep the machine above 50%.

The Formula

Given your sequence XYZ, the machine picks $\bar{Y}$XY — the opposite of your middle letter, then your first two letters. For HHT: middle letter H, opposite T, first two letters HH → machine picks **THH**.

The middle letter is the load-bearing letter of your sequence. It sits between your required buildup (XY) and your win. The machine copies your buildup as its ending (stealing your progress) and opens with the opposite of Y (locking the trap before you can build anything).

This weakens for HTH and THT because both sequences start with the same letter as the machine's counter — early flips advance both players simultaneously, diluting the trap.

Summary

Your Sequence	Machine's Counter	Machine Wins
HHH	THH	**87.5%**
TTT	HTT	**87.5%**
HHT	THH	**75.0%**
HTT	HHT	**75.0%**
THH	TTH	**75.0%**
TTH	HTT	**75.0%**
HTH	HHT	**66.7%**
THT	TTH	**66.7%**

Pick HTH or THT. Make it close.

Note: The state framework used here is an instance of absorbing Markov chain theory. A complete treatment would use transition matrices — but everything above follows from LOTP alone.

Products Used: GPT-4 for LaTeX equations and table formatting. UPMATH Markdown & LaTeX editor.