Introduction and Disclaimer

These mock examination questions span diverse disciplines and are designed for your practice in preparation for the International Research Olympiad (IRO) 2024. Endeavor to answer them to the best of your ability, utilizing this opportunity to enhance your skills and knowledge. For additional practice, it is advisable to engage in extensive reading of various papers; such efforts will contribute to a more comprehensive and nuanced understanding of the subject matter.

Try your best, and good luck! -International Research Olympiad 2024

Fair coins tend to land on the same side they started: Evidence from 350,757 flips

Abstract

Many people have flipped coins but few have stopped to ponder the statistical and physical intricacies of the process. In a preregistered study we collected 350,757 coin flips to test the counterintuitive prediction from a physics model of human coin tossing developed by Diaconis, Holmes, and Montgomery (D-H-M; 2007). The model asserts that when people flip an ordinary coin, it tends to land on the same side it started—D-H-M estimated the probability of a same-side outcome to be about 51%. Our data lend strong support to this precise prediction: the coins landed on the same side more often than not, Pr(same side) = 0.508, 95% credible interval (CI) [0.506, 0.509], $BF_{\text{same-side bias}}$ = 2364. Furthermore, the data revealed considerable between-people variation in the degree of this same-side bias. Our data also confirmed the generic prediction that when people flip an ordinary coin—with the initial side-up randomly determined—it is equally likely to land heads or tails: Pr(heads) = 0.500, 95% CI [0.498, 0.502], $BF_{\text{heads-tails bias}}$ = 0.183. Furthermore, this lack of heads-tails bias does not appear to vary across coins. Our data therefore provide strong evidence that when some (but not all) people flip a fair coin, it tends to land on the same side it started. Our data provide compelling statistical support for D-H-M physics model of coin tossing.

Paper

A coin flip—the act of spinning a coin into the air with your thumb and then catching it in your hand—is often considered the epitome of a chance event. It features as a ubiquitous example in textbooks on probability theory and statistics and constituted a game of chance ('capita aut navia' – 'heads or ships') already in Roman times.

Despite the widespread popularity of coin flipping, few people pause to reflect on the notion that the outcome of a coin flip is anything but random: a coin flip obeys the laws of Newtonian physics in a relatively transparent manner. According to the standard model of coin flipping, the flip is a deterministic process and the perceived randomness originates from small fluctuations in the initial conditions (regarding starting position, configuration, upward force, and angular momentum) combined with narrow boundaries on the outcome space. Therefore the standard model predicts that when people flip a fair coin, the probability of it landing heads is 50% (i.e., there is no 'heads-tails bias'; conversely, if one side of a coin would land on one side more often than the other, we would say there is a 'heads-tails bias').

The standard model of coin flipping was extended by Diaconis, Holmes, and Montgomery who proposed that when people flip a ordinary coin, they introduce a small degree of 'precession' or wobble—a change in the direction of the axis of rotation throughout the coin's trajectory. According to the D-H-M model, precession causes the coin to spend more time in the air with the initial side facing up.

Consequently, the coin has a higher chance of landing on the same side as it started (i.e., 'same-side bias'). Based on a modest number of empirical observations (featuring coins with ribbons attached and high-frame-rate video recordings) D-H-M measured the off-axis rotations in typical human flips. Based on these observations, the D-H-M model predicted that a coin flip should land on the same side as it started with a probability of approximately 51%, just a fraction higher than chance.

A group of 48 people (i.e., all but three of the co-authors) tossed coins of 46 different currencies × denominations and obtained a total number of 350,757 coin flips. The data confirm the prediction from the D-H-M model: the coins landed how they started more often than 50%. Specifically, the data feature 178,078 same-side landings from 350,757 tosses, Pr(same side) = 0.508, 95% CI [0.506, 0.509], which is remarkably close to D-H-M prediction of (approximately) 51%.

In addition, the data show no trace of a heads-tails bias. Specifically, we obtained 175,420 heads out of 350,757 tosses, Pr(heads) = 0.500, 95% CI [0.498, 0.502]

A preregistered Bayesian informed binomial hypothesis test indicates extreme evidence in favor of the same-side bias predicted by the D-H-M model, $BF_{\text{same-side bias}}$ = 1.71 × 1017. A similar (not-preregistered) analysis yields moderate evidence against the presence of a heads-tails bias, $BF_{\text{heads-tails bias}}$ = 0.168.

With the data in hand we realized that the same-side bias was possibly subject to considerable between-people heterogeneity. Therefore we specified a more complex Bayesian hierarchical model that includes both heterogeneity in same-side bias between people and heterogeneity in heads-tails bias between coins; this hierarchical model was then used to estimate the parameters and to test the hypotheses using Bayesian model-averaging and inclusion Bayes factors (for details see the methods section). These analyses were not preregistered. The posterior distribution of the same-side bias is slightly wider than in the simple preregistered analysis, Pr(same side) = 0.5098, 95% CI [0.5049, 0.5147], which is caused by the substantial between-people heterogeneity in the probability of the coin landing on the same side, $sd_{\text{people}}(p_{\text{same-side}})$ = 0.0156, 95% CI [0.0119, 0.0200]. The additional uncertainty lowers the evidence for the same-side bias to $BF_{\text{same- side bias}}$ = 2364, which however remains extreme (e.g., when the hypothesis of a same-side bias has a prior probability of 0.50, a Bayes factor of about 2364 results in a posterior probability of 0.9996). The hierarchical model reveals overwhelming evidence for the presence of between-people heterogeneity in same-side bias, $BF_{\text{people heterogeneity}}$ = 3.04 × 1024 .

We collected 350,757 coin flips and found strong empirical confirmation of the counterintuitive and precise prediction from D-H-M model of human coin tossing: when people flip a coin, it tends to land on the same side as it started. Moreover, the data revealed a substantial degree of between-people variability in same-side bias: some people appear to have little or no same-side bias, whereas others do display a same-side bias, albeit to a varying degree. This variability is consistent with D-H-M model, in which the same-side bias originates from off-axis rotations (i.e., precession or wobbliness), which can reasonably be assumed to vary between people. Future work may attempt to verify whether 'wobbly tossers' show a more pronounced same-side bias than 'stable tossers'. The effort required to test this more detailed hypothesis appears to be excessive, as it would involve detailed analyses of high-speed camera recordings for individual flips.

In order to ensure the quality of the data, we videotaped and audited the data collection procedure (see the method section for details). However, there remains a legitimate concern: at the time when people were flipping the coins they were aware of the main hypothesis under test. Therefore it cannot be excluded that some of the participants were able to manipulate the coin flip outcomes in order to produce the same-side bias. In light of the nature of the coin tossing process, the evidence from the video recordings, and the precise correspondence between the data and the predictions from D-H-M model, we deem this possibility as unlikely; future work is needed to disprove it conclusively (e.g., by concealing the aim of the study).

Could future coin tossers use the same-side bias to their advantage? The magnitude of the observed bias can be illustrated using a betting scenario. If you bet a dollar on the outcome of a coin toss (i.e., paying 1 dollar to enter, and winning either 0 or 2 dollars depending on the outcome) and repeat the bet 1,000 times, knowing the starting position of the coin toss would earn you 19 dollars on average. This is more than the casino advantage for 6 deck blackjack against an optimal-strategy player, where the casino would make 5 dollars on a comparable bet, but less than the casino advantage for single-zero roulette, where the casino would make 27 dollars on average. These considerations lead us to suggest that when coin flips are used for high-stakes decision-making, the starting position of the coin is best concealed.

Paper 5: Statistics

Question 1

Question: What does the D-H-M model assert about flipping a coin?

a) 51 percent of coins land on a side different from that on which they started.

b) 50 percent of coins land on the same side that they started.

c) All coins have a bias towards heads due to engravings shifting coins' mass distribution.

d) When people flip a coin, the axis of rotation changes throughout the coin's trajectory.

Question 2

Question: Does this paper provide statistically significant support for the D-H-M model of coin-flipping?

a) True

b) False

Question 3

Question: What is a problem in the way that data was collected in this research paper?

a) Different denominations of coins consistently produced different results.

b) Coin-flippers were aware that they were testing the fairness of coins.

c) The data collection process was flawed because the researchers did not use a randomized controlled trial to ensure the validity of the coin flips.

d) The data collection was problematic because the researchers did not account for the influence of external factors such as air resistance or surface irregularities on the coin flips.

Question 4

Question: How did the authors account for biases specific to certain coin-flippers in their analysis?

a) They compared each flipper's Gaussian distribution.

b) They used a Bayesian logistic regression model.

c) They used randomized samples. adobe acrobat

d) Coins were put through an automatic flipping machine to remove human bias.