

Teacher's Corner

You Can Load a Die, But You Can't Bias a Coin

Andrew GELMAN and Deborah NOLAN

Dice can be loaded—that is, one can easily alter a die so that the probabilities of landing on the six sides are dramatically unequal. However, it is *not* possible to bias a coin flip—that is, one cannot, for example, weight a coin so that it is substantially more likely to land “heads” than “tails” when flipped and caught in the hand in the usual manner. Coin tosses can be biased only if the coin is allowed to bounce or be spun rather than simply flipped in the air. We describe a student activity with dice and coins that gives empirical evidence to support this property, and we use this activity when we teach design of experiments and hypothesis testing in our introductory statistics courses. We explain this phenomenon by summarizing a physical argument made in earlier literature.

KEY WORDS: Classroom activity; Experimental design; Fair coin.

“A coin with probability $p > 0$ of turning up heads is tossed . . .”
—Woodroffe (1975, p. 108)

“Suppose a coin having probability 0.7 of coming up heads is tossed . . .”
—Ross (2000, p. 82)

1. INTRODUCTION

The biased coin is the unicorn of probability theory—everybody has heard of it, but it has never been spotted in the flesh. As with the unicorn, you probably have some idea of what the biased coin looks like—perhaps it is slightly lumpy, with a highly nonuniform distribution of weight. In fact, the biased coin does not exist, at least as far as flipping goes.

We have designed classroom demonstrations and student activities around the notion of the biased coin. The simple toss of a coin offers opportunities for learning many lessons in statistics and probability. For example, we ask our students, How is a coin toss random? and What makes a coin fair? This starts a discussion of random and deterministic processes. We can design experiments and collect data to test our assumptions about coin tossing, and because flipping coins is such a simple and familiar

concept, important issues surrounding experimental design and data collection are easy to spot and address.

Gambling and the art of throwing dice have a colorful history. For example, in the eleventh century, King Olaf of Norway wagered the Island of Hising in a game of chance with the King of Sweden (Ekeland 1993). King Olaf beat the Swede's pair of sixes by rolling a thirteen! One die landed six, and the other split in half landing with both a six and a one showing. Jay (2000) has many other interesting stories of the history of biased dice. Ortiz (1984) gave an amusing story of an elaborate confidence game based on a rigged top. What amazes us most about this story is that people are apparently willing to bet with complete strangers in a bar on the outcome of a spinning top.

But for coins, the physical model of coin flipping (see Section 5), which says that the “biased coin,” when flipped properly, should land heads about half the time, may explain why we had trouble finding such stories about biased coins. One exception is in the work of Kerrich (1946). In 1941, while interned in Denmark, he tossed a coin 10,000 times. He describes his method of tossing, “A small coin, balanced on the writer's forefinger, was given a little flip with the thumb so that it spun through the air for about a foot finally landing on a cloth spread out flat over a table . . . if the coin fell heads in one spin it was convenient to balance it head uppermost on the operator's forefinger when preparing for the next, and vice versa.” In addition to tossing this coin (which landed heads 5,067 times), Kerrich also tossed a wooden disc that had one face coated with lead. Calling this face “tails” and the other “heads,” he found the coin landed heads 679 out of 1,000 times. As the coin was tossed only a short distance and was allowed to bounce on the table, a bias was observed.

2. UNFAIR FLIPPING AND EXPERIMENTAL PROTOCOL

We bring a plastic checker to class and affix putty to the crown side, which we also call the heads side. Then we ask the students whether they think the chance the checker lands heads when tossed is $1/2$ or not? Most are positive that the “coin” is biased. We try flipping the checker a few times and varying the way we flip it—high, low, fast spin, no spin, like a frisbee, off kilter, bouncing on the ground, catching it in the air.

Quickly, the students see that to make any sense of this probability statement it is important to specify how to flip the checker. We ask them to come up with a list of rules to follow when flipping the coin to make the flips as similar as possible. For example: begin with the crown side up and parallel to the floor; flip the coin straight up, high in the air so it spins rapidly, and with the spinning axis also parallel to the floor; catch the coin midair in the palm of hand.

We proceed from here with designing an experiment to test the hypothesis that p , the probability the coin lands heads, is

Andrew Gelman is Professor, Department of Statistics, Columbia University, New York, NY 10027. Deborah Nolan is Professor, Department of Statistics, 367 Evans Hall, #3860, University of California, Berkeley, CA 94720 (E-mail: nolan@stat.berkeley.edu). The authors thank Stephen Stigler and an anonymous referee for references and suggestions that led to an improved article. This example is further discussed in Gelman and Nolan (2002).

120 rolls of a loaded die

3	6	3	1	2	1	4	4	1	2	1	4	3	1	5	5	1	1	1	3
1	1	1	1	1	2	3	1	3	4	1	2	3	4	2	2	5	1	5	1
5	5	5	5	3	2	3	5	1	3	5	1	3	1	1	2	1	3	3	1
3	1	2	5	4	3	2	1	2	3	2	2	3	2	3	5	1	3	1	5
5	3	1	5	1	4	1	2	3	2	1	3	2	1	3	5	1	2	1	4
5	1	3	4	1	5	5	1	1	4	3	5	1	5	3	1	3	5	1	3

Figure 1. The results of 120 rolls of a die that has had the edges on the 1-face rounded. In 120 rolls, the 6, which is opposite the 1, showed only once.

1/2. We ask questions such as, how many flips are needed to determine if the coin is not fair? These lead to discussions of hypothesis testing, one- and two-sided alternatives, significance level, and power.

3. CHECKERS, BUBBLE GUM, AND STUDENT ACTIVITIES

Our in-class demonstration continues with a student activity on biased coins and dice. After showing the students our altered checker, and discussing how to flip it fairly, we give them a chance to make their own biased coins and dice. We divide them into pairs and give a plastic checker, wooden die, and piece of sandpaper to each pair. We tell them that they can alter the checkers and dice however they want—for example, they can sand the edges of the die or affix gum to one or both sides of the checker (but then they should let the gum dry before handling the checker). The object is to maximize the probability of tails (or heads) for the checker, and to alter the die so the six sides are not equally probable.

We provide explicit instructions on how to flip and spin the checker and how to roll the die. To roll the die, they must find a smooth level surface and draw a circle on it about one to two feet in diameter. They shake the die in a cup and drop it into the circle. It must remain within the circle when it comes to a rest in order to count as a successful roll. The same circle is used for spinning the checker. The spins are to be contained in the circle, and they must spin quickly before falling. To flip the checker, we follow the rules set up in our earlier discussion. After our in-class demo, the students understand the necessity of closely following the protocol.

We also give specific instructions about the order that they are to work on the dice and checkers. First, we have them modify the die. (Although it is easy enough to do, we were surprised at how much sanding was needed to notice a big difference from what is expected in 120 flips.) Next, they alter the coin. The hitch is that we instruct them to modify the checker until they are satisfied that it is biased when spun (say when tails come up 65 or more times in 100 spins). Then they are to flip the

altered checker, without making any further modifications to it. The students should find that the alterations have essentially no effect on the flips even though they have a large effect on the spins.

We ask the students to bring their modified coins and dice to the next lecture along with a record of their results. They are to roll the die 120 times, and spin and flip the checker 100 times each. Figures 1 and 2 contain the results from one student's efforts to modify her die and checker. After rounding one of the corners on her die, she rolled only one 6 in 120 throws of the die. She also found that her altered checker landed tails 77 times in 100 spins (Figure 2), which has a chance less than one in ten million of occurring with fair spins. But this same altered checker when flipped landed tails 47 times out of 100, a typical result for a "fair coin."

If we have a lot of students in the class, then we would expect a few of the pairs to have significant results for the flipping activity. This is an excellent opportunity to discuss the notion of multiple comparisons. Some students may insist that their checker has a probability greater than 0.5 of landing tails (or heads) when flipped. Since they have brought their checkers to class, we have them turn over their checkers to us for further investigation.

4. SPORTING EVENTS AND QUANTITATIVE LITERACY

After the students, have been "tricked" with the checkers, we discuss the findings of Tomasz Gliszczynski and Wacław Zawadowski, statistics teachers at the Akademia Podlaska in Siedlce, Poland. These statisticians had their students *spin* the Belgian one-Euro coin 250 times, and found that it came up heads 140 times. (The one-Euro coins have a common design on the tails side and a national image on the heads side. Belgium portrays King Albert II on the heads side.) As the introduction of the Euro was the largest currency switch in history, this finding received a lot of press coverage. We present some excerpts from news stories for discussion and critique.

100 spins of the checker	100 spins of the checker
01101001100000101010	11110101010000001111
10001000001010000100	00111000011000110100
00010100010000010000	01111100111001110111
0001000000001001000	00100010111100110010
00100000000000010010	10111100111101100010
(23 heads, 77 tails)	(54 heads, 46 tails)

Figure 2. The results of 100 spins and 100 spins of a plastic checker which has been altered with putty. Heads are denoted by 1 and tails by 0. (We indicate Heads and Tails by 1 and 0, respectively, because "H" and "T" are hard to distinguish visually.)




Figure 3. Angular position of a "tipped coin" as a function of time. Suppose heads was up initially. Then, when the coin is caught, if the angle is between 90 and 270 degrees, it will show "heads" when caught and displayed; otherwise, it will show "tails." The initial condition of the coin is "forgotten" if the uncertainty about when the coin is caught is much greater than the rotation period. In this case, the coin will be in the "heads" region of angular space with probability 1/2, no matter how the coin is weighted.

Memo to all teams playing Belgium in the World Cup this year: don't let them use their own coins for the toss. . . . "It looks very suspicious to me," said Barry Blight, a statistics lecturer at the London School of Economics. "If the coin were unbiased the chance of getting a result as extreme as that would be less than 7%." —Heads, Belgium wins—and wins, *The Guardian*, Jan 4, 2002.

The academics claim Belgium Euro coins have been struck "asymmetrically" and come up tails only 44% of the time. . . . With a French Euro and a limited knowledge of physics—is it best to flip a coin in the air or spin it on a table?—we set to work. . . . The French flip 56% "tree." The French spin 52% "tree." —Hussain's flipping fillip, *BBC Sport*, Jan 4, 2002.

The observation is not to be taken lightly on a sports-mad continent where important decisions can turn on the flip of a coin. . . . Gliszynski says spinning is a more sensitive way of revealing if a coin is weighted than the more usual method of tossing in the air. . . . But Howard Grubb, an applied statistician at the University of Reading, notes that, "with a sample of only 250, anything between 43.8 percent and 56.2 percent on one side or the other cannot be said to be biased." *New Scientist* carried out its own experiments with the Belgian Euro in its Brussels office. Heads came up five percent less often than tails. —Euro coin accused of unfair flipping, *New Scientist*, Jan 4, 2002.

With their experience flipping and spinning their uniquely modified checkers, the students are ready to discuss these news stories. Points that quickly come to surface are: there is confusion between flipping and spinning the coin; two articles report their own experimental results, but they do not supply the number of flips or spins; and the two statisticians quoted do not agree on whether the results are suspicious or not.

5. PICKLE-JAR LIDS AND IDEAL COINS

Deterministic physical laws govern what happens in the flip of a coin and the throw of a die, but we consider these events as random. It's hard to separate the random from the deterministic even in something as simple as the coin flip. What makes a coin toss fair?

The uncertainty of the coin's initial state is the key. A coin tossing is basically deterministic. The coin obeys Newton's laws of motion, with its final state depending on its angular velocity (rate of spin) and time traveled (which in turn depends on the upward velocity with which it is flipped). For tosses where the coin spins rapidly and goes high in the air, the set of initial velocity values that lead to either heads or tails are of equal size.

That is, half of the initial conditions lead to heads and half to tails (see, e.g., Keller 1986; Peterson 1990, 1997). So, uncertainty in the initial state (e.g., a smooth probability distribution on a range of values for the initial state) leads to the coin landing heads half the time.

The law of conservation of angular momentum tells us that once the coin is in the air, it spins at a nearly constant rate (slowing down very slightly due to air resistance). At any rate of spin, it spends half the time with heads facing up and half the time with heads facing down, so when it lands, the two sides are equally likely (with minor corrections due to the nonzero thickness of the edge of the coin); see Figure 3. Jaynes (1996) explained why weighting the coin has no effect here (unless, of course, the coin is so light that it floats like a feather): a lopsided coin spins around an axis that passes through its center of gravity, and although the axis does not go through the geometrical center of the coin, there is no difference in the way the biased and symmetric coins spin about their axes.

Jaynes also described how to add another kind of spin to the coin like the spin when you toss a frisbee, which enables you (if you are good enough at coin flipping) to have the coin, biased or not, always land heads. To prove his point, he tossed the lid of a pickle jar according to three different methods. First he tossed it with a frisbee-type twist and a very slow spin, the lid landed "heads" 99 out of 100 times. (Heads in this case is the inside of the lid.) Then he tossed the lid so that it landed on its edge and spun rapidly on the floor before falling to one side. This time the lid landed heads 0 out of 100 times, because a lopsided coin tends to fall on the side that makes the center of gravity high, and the center of gravity for the lid was closer to the top. (The lid had a diameter of $2\frac{5}{8}$ inches, height of $\frac{3}{8}$ inch, and center of gravity 0.12 inches from the top of the lid.) Finally, when the pickle jar lid was tossed without any bounce or frisbee-spin, it landed heads 54 out of 100 times.

In the first method, the frisbee-style twist on the toss dominates, in the second, the bias takes over, and in the third, we have a "fair coin toss." It does not make sense to say that the coin has a probability p of heads, because it can be completely determined by the manner in which it is tossed—unless it is tossed high in the air with a rapid spin and caught in the air with no bouncing, in which case $p = 1/2$. If we must assign a probability p to a coin, then that probability must be approximately 1/2 (unless it is double-headed or double-tailed), no matter how it is weighted. According to Peterson (1990), "A look at the spread in the way real people flip real coins indicates . . . a slight bias would begin to show up after millions of tosses. The proportion of, say, heads would settle at a number such as 0.503 or 0.497 . . ." But dice can be "loaded" to make some faces more likely because, among other reasons, dice bounce after being thrown, and weighting and beveling can affect the bounces. As we saw with Jaynes's experiment, if a coin is spun, or if it is thrown and allowed to bounce, it can have a stable probability of heads that is not close to 1/2, and it is easy to alter this probability by shaving the edges of the coin to different angles.

6. SUMMARY

Many of us include in our classes activities that use coin flipping or die throwing, most often as a vehicle for sampling. The activities presented here are different in that they study the coin

toss itself. The act of tossing a coin, something with which we are all familiar, is the source of deep questions about randomness. Randomness is an elusive concept, and books on probability rarely even attempt to define it.

Because coin tossing is so simple, we can clearly lay out the experimental protocol. Examples of how others have flipped and tossed coins show the students how essential it is to carefully describe the experimental process. And although it is easy to carry out our experiments, there is a purpose to them and inherent interest on the part of students in the outcome. The flurry of the news stories, and their focus on the use of coins in soccer games makes this point. We use student interest in the topic to engage students in discussion of, for example, the rules for coin flipping, the number of flips needed, and the number of heads required to convince us of a bias.

We have found in developing classroom demonstrations and student activities, that the statistical lesson is most effectively conveyed if there is a surprise ending, or an unexpected twist in the result. That is certainly the case here, when we have the students make the checker display a bias when spun only to find that it remains unbiased when flipped.

The biased coin has long been part of statistical folklore, but it does not exist in the form in which it is imagined. Is this impor-

tant? Certainly not for probability theory. However, it provides an excellent source for a variety of statistics and probability lessons.

[Received July 2001. Revised April 2002.]

REFERENCES

- Ekeland, I. (1993), *The Broken Dice*, Chicago: University of Chicago Press.
- Gelman, A., and Nolan, D. (2002), *Teaching Statistics: A Bag of Tricks*, Oxford: Oxford University Press.
- Jay, R. (2000), "The Story of Dice: Gambling and Death from Ancient Egypt to Los Angeles," *The New Yorker*, Dec. 11, 91–95.
- Jaynes, E. T. (1996), *Probability Theory: The Logic of Science*, pp. 1003–1007; bayes.wustl.edu/etj/prob.html
- Keller, J. B. (1986), "The Probability of Heads," *American Mathematical Monthly*, 93, 191–197.
- Kerrich, J. E. (1946), *An Experimental Introduction to the Theory of Probability*, Copenhagen: J. Jorgensen.
- Ortiz, D. (1984), *Gambling Scams*, New York: Lyle Stuart.
- Peterson, I. (1990), *Islands of Truth*, New York: Freeman.
- (1997), "A Penny Surprise," *MathTrek*, MAA Online, Dec. 15; www.maa.org/mathland/mathtrek_12_15.html.
- Ross, S. M. (2000), *Introduction to Probability Models* (7th ed.), San Diego: Harcourt.
- Woodroffe, M. (1975), *Probability with Applications*, New York: McGraw-Hill.