

RESEARCH

How random is the toss of a coin?

Matthew P.A. Clark MBBS, Brian D. Westerberg MD

ABSTRACT

Background: The toss of a coin has been a method used to determine random outcomes for centuries. It is still used in some research studies as a method of randomization, although it has largely been discredited as a valid randomization method. We sought to provide evidence that the toss of a coin can be manipulated.

Methods: We performed a prospective experiment involving otolaryngology residents in Vancouver, Canada. The main outcome was the proportion of “heads” coin tosses achieved (out of 300 attempts) by each participant. Each of the participants attempted to flip the coin so as to achieve a heads result.

Results: All participants achieved more heads than tails results, with 7 of the 13 participants having significantly more heads results ($p \leq 0.05$). The highest proportion of heads achieved was 0.68 (95% confidence interval 0.62–0.73, $p < 0.001$).

Interpretation: Certain people are able to successfully manipulate the toss of a coin. This throws into doubt the validity of using a coin toss to determine a chance result.

The toss or flip of a coin to randomly assign a decision traditionally involves throwing a coin into the air and seeing which side lands facing up. This method may be used to resolve a dispute, see who goes first in a game or determine which type of treatment a patient receives in a clinical trial. There are only 2 possible outcomes, “heads” or “tails,” although, in theory, landing on an edge is possible. (Research suggests that when the coin is allowed to fall onto a hard surface, the chance of this happening is in the order of 1 in 6000 tosses.¹)

When a coin is flipped into the air, it is supposedly made to rotate about an axis parallel to its flat surfaces. The coin is initially placed on a bent forefinger, and the thumb is released from under the coin surface, where it has been held under tension. The thumbnail strikes the part of the coin unsupported by the index finger, sending it rotating upward. All this is done with an upward movement of the hand and forearm. The coin may be allowed to fall to the floor or other surface or it may be caught by the “tossler” and sometimes turned onto the back of



© 2009 Jupiterimages Corp.

the opposite hand and then revealed. The catching method should not matter, provided it is consistent for each toss. The opponent often calls the toss when the coin is airborne, although in the case of randomization for clinical trials, this is unnecessary because one is simply looking for an outcome.

From St. Paul's Rotary Hearing Clinic, University of British Columbia, Vancouver, BC

Cite as *CMAJ* 2009. DOI:10.1503/cmaj.091733

DOI:10.1503/cmaj.091733

The appeal of the coin toss that it is a simple, seemingly unbiased, method of deciding between 2 options. Although the outcome of a coin toss should be at even odds, the outcome may well not be. Historically, the toss of a coin before a duel reputedly decided which person had his back to the sun — an obvious advantage when taking aim! In medical trials, a simple statistical manipulation can have a dramatic effect on the treatment a patient receives. Our hypothesis is that with minimal training, the outcome of the toss can be weighted heavily to the call of the tosser, thus abolishing the 50:50 chance result that is expected and allowing for manipulation of an apparently random event.

Methods

We included 13 otolaryngology residents from the University of British Columbia who verbally consented to participate. We excluded any residents who had a hand injury that would prevent them from tossing a coin with their dominant hand.

The participants were told about the purpose of the study several weeks before the trial. The investigators instructed the residents and demonstrated the toss. The residents were allowed to practise tossing the coin in a consistent manner for a couple of minutes. Each resident then performed 300 coin tosses in which they tried to achieve a heads result each time. The results were recorded by an observer to avoid the possibility of cheating. The 2 participants who achieved the greatest manipulation of the results (i.e., the highest proportion of heads) were given an incentive (i.e., \$20 and \$10 coffee vouchers). We felt that the use of an incentive would reproduce real-life situations.

Statistical analysis

Because our null hypothesis was that the coin-weighting was unbiased, we assumed that 50% of the tosses would result in heads. We used 2-sided binomial testing because we could not assume that if we were trying to increase the proportion of heads, we would not instead reduce this proportion. For a change of at least 10%, we calculated that 263 coin tosses per participant would be required to achieve 90% power to detect this difference with a significance level of 0.05.

We performed a comparison of proportions. The statistical calculations were performed on the basis of how many coin tosses an individual would have to perform to show a significant manipulation. Therefore, any participant who achieved this level would have a significant result. We did not use group statistics for this reason.

Results

Each of the 13 participants tossed a coin 300 times. All participants tried to achieve a heads result. Each participant successfully achieved more heads than tails results; this difference was statistically significant for 7 participants (Table 1). The participant who was most successful at manipulating the outcome achieved a proportion of heads of 0.68 (95% confidence interval 0.62–0.73; $p < 0.001$).

Interpretation

This study shows that when participants are given simple instructions about how to manipulate the toss of a coin and only a few minutes to practise this technique, more than half can significantly manipulate the outcome. With devoted training, more participants would probably be able to achieve this figure, and the magnitude of the manipulation would probably be increased.

With respect to the use of a coin toss to randomly assign patients to a treatment in a clinical trial, our results could be considered clinically significant if only 1 participant in this study had achieved a nonrandom result. We have shown that a person tossing a coin may have the ability to manipulate the toss and significantly bias the results to their liking. Given that we would never know the manipulation skills or motivation of the person tossing the coin, this method seems unsuitable for randomization procedures in experiments in which bias needs to be minimized.

Research by Diaconis and colleagues² has suggested a dynamic bias to coin tosses. They suggest that, for natural flips, the chance of the coin landing as it started is about 0.51, with a number of assumptions or required conditions. Diaconis states that, for a fair toss, the coin must be caught in the palm of the hand and not allowed to land on a surface and bounce, because the latter often incorporates a degree of spinning on the coin's edge. Because catching a coin in one's palm and turning it onto the back of the opposite hand allows for manipulation with sleight of hand, the coin must be allowed to just land in the palm of the hand.

Diaconis and colleagues² assumed that when tossed, the

Table 1: Results of coin tosses in which 13 participants attempted to achieve a "heads" result

Participant	Outcome <i>n</i> = 300		Proportion of heads (95% CI)	<i>p</i> value*
	Heads	Tails		
1	162	138	0.54 (0.48–0.60)	0.18
2	175	125	0.58 (0.53–0.64)	0.005
3	159	141	0.53 (0.47–0.59)	0.37
4	179	121	0.60 (0.54–0.65)	0.001
5	203	97	0.68 (0.62–0.73)	< 0.001
6	168	132	0.56 (0.50–0.62)	0.043
7	170	130	0.57 (0.51–0.62)	0.024
8	160	140	0.53 (0.48–0.59)	0.27
9	192	108	0.64 (0.58–0.69)	< 0.001
10	167	133	0.56 (0.50–0.61)	0.06
11	154	146	0.51 (0.46–0.57)	0.69
12	153	147	0.51 (0.45–0.57)	0.77
13	176	124	0.59 (0.53–0.64)	0.003

Note: CI = confidence interval.

*Determined by use of a binomial test.

uppermost side of the coin is known by the tosser (and caller). If this is not the case, the true 50:50 probability of the result prevails. But if it is known, the side of the coin that starts off face up is more likely to end that way up because that side spends more time facing up during the flight than does the opposite side. However, because of the large number of tosses required to detect this difference (shown to be 250 000 tosses), this apparent difference is generally irrelevant.

Coins have different raised profiles on each side. In theory, one side could be more weighted, thus making the coin toss unfair. This could be achieved by tampering with a coin. With the introduction of the Euro, Polish statisticians claimed that the 1 € coin (from Belgium), when spun on a surface, came up heads more often than tails.³ This report resulted in a tongue-in-cheek warning in the British press to teams playing against Belgium in the forthcoming soccer World Cup.⁴ Of 250 spins, 56% came up heads. However, independent statistical analysis showed that random variation could produce such scatter even with an unbiased coin given this number of spins (acceptable range 43.8%–56.2%). Other research suggests that even if a grossly weighted coin is used (1 side lead, the other balsa wood), no significant bias shows up.³

The usual method of tossing a coin should not be taken for granted. Gary Kosnitzky, a Las Vegas magician adept in the art of coin magic, can reputedly manipulate coin tosses to give a predictable outcome by use of a method that gives the illusion of a spinning coin. When his coin “spins” in the air, it is not rotating about its axis but is instead fluttering or wobbling and not turning over. This gives the illusion of a normally spinning coin, yet gives a predictable result if caught in the palm of the hand.⁵ Admittedly, this is probably an art more practised by the gambler than the researcher, but it serves as a reminder that an apparently fair situation can be manipulated.

The most compelling finding that raises concerns about the validity of the coin toss come from the use of mechanical coin flippers, which can be made to impart exactly the same initial conditions for every toss, namely the starting position, velocity and force. In these cases, the outcome can be highly, if not entirely, predictable. Coin tossing becomes physics rather

than a random event. It is the human element that makes the process random in that each toss tends to be at a different speed, sent to a different height, launched at a different angle or caught in a different manner. Therefore, the possibility of practising the task to reduce these differing elements can be considered. If you try to toss the coin the same way each time, you should be able to make the outcome significantly different than 50:50.

A limitation of our study is the uncertainty about whether the individual results are repeatable. However, we conclude that the validity of using the toss of a coin to provide a random 50:50 outcome is thrown into doubt, both in medicine and every day life.

Competing interests: None declared.

Contributors: Both authors contributed to the conception of the study design, the acquisition and analysis of the data and writing and revising the manuscript. Both authors approved the final version submitted for publication. Matthew Clark accepts full responsibility for the work and the conduct of the study, had access to the data and controlled the decision to publish.

Acknowledgements: The authors thank Hong Qian for her advice on the statistical analysis.

REFERENCES

1. Murray DB, Teare SW. Probability of a tossed coin falling on its edge. *Phys Rev E Stat Phys Plasmas Fluids Relat Interdiscip Topics* 1993;48:2547-52.
2. Diaconis P, Holmes S, Montgomery R. Dynamical bias in the coin toss. *SIAM Rev* 2007;49:211-35.
3. MacKenzie D. Euro coin accused of unfair flipping. *New Sci* 2002 Jan. 4. Available: www.newscientist.com/article/dn1748-euro-coin-accused-of-unfair-flipping.html (accessed 2009 Oct. 22).
4. Denny C, Dennis S. Heads, Belgium wins — and wins. *The Guardian* [UK] 2002 Jan. 4. Available: www.guardian.co.uk/world/2002/jan/04/euro.eu2 (accessed 2009 Oct. 22).
5. Kosnitzky G. Heads or tails. *Murphy's Magic Supplies*. Rancho Cordova (CA); 2006.

Correspondence to: Matthew P.A. Clark, Consultant Otolaryngologist, Department of Otolaryngology, Gloucestershire Royal Hospital, Gloucester, UK GL1 3NN; matthew.clark@glos.nhs.uk