From *SIAM News*, Volume 37, Number 4, May 2004

Tossed Coins and Troubled Marriages: Mathematical Highlights from AAAS 2004

By Barry A. Cipra

Attendees at the annual meeting of the American Association for the Advancement of Science, held February 13–16 in Seattle, got a good dose of the mathematical sciences, with sessions on the changing nature of proof, analyses of the World Wide Web, and phase transitions in computer science, among other topics. As described in this article, the program also included a talk on a surprising new result in the age-old search for randomness and a session on mathematical analysis of the ever-vexing problem of love and marriage.

Unfair Coins

Put a quarter heads up on the tip of your thumb, flick it into the air, and catch it on the way down. What are the odds it'll land heads up? The answer that springs to mind is 50:50. But according to Persi Diaconis of Stanford University, the actual odds may be more like 51:49.

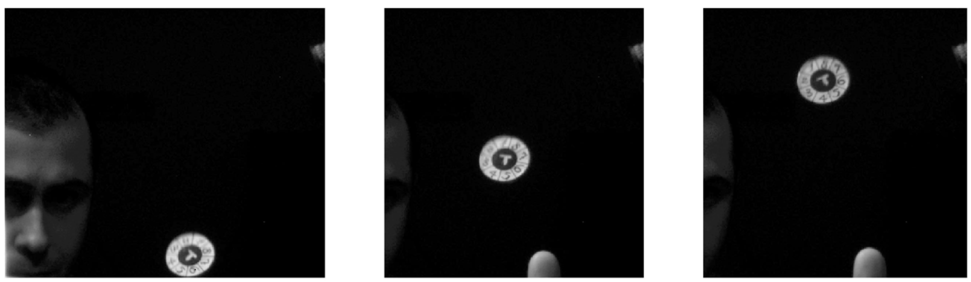

A tossed coin—as shown in three frames (left to right, 48, 68, 88) taken by a high-speed camera—does not simply (or even always) flip end over end, but also precesses. Persi Diaconis, Susan Holmes, and Richard Montgomery have shown that precession introduces a bias into coin tosses, making it essentially impossible to toss a coin fairly.

Diaconis and colleagues Susan Holmes of Stanford and Richard Montgomery of the University of California at Santa Cruz have analyzed what happens when a thin disk spins through space. Their analysis indicates that no matter how vigorously a heads-up coin is launched, it invariably shows a bias toward landing heads up. The results of empirical studies in which a high-speed camera tracked the gyrations of real coins suggest a bias as great as 2%—which exceeds the (theoretical) house advantage in roulette (37:36). Diaconis presented the coin-tossing results in a topical lecture at the AAAS meeting.

Coin tossing has long been the icon of randomness, but in fact there is nothing random about a spinning coin, Diaconis says. When he was at Harvard in the early 1990s, Diaconis had the physics department build a coin-tossing device that launches coins with the exact same initial conditions every time. A quarter placed heads up on the device goes up about a foot, turns over 12 times, and lands heads up every time. In short, a tossed coin is simply a rigid body subject to Newton's laws of motion. Randomness enters via uncertainties in initial conditions, the inability of people to be perfectly precise. Everything else is classical mechanics.

In 1986, Joe Keller of Stanford published an analysis of the physics of tossed coins. Assuming that the coin starts heads up, spins about its diameter with angular velocity ω, and is caught after t seconds (if the coin is tossed upward with velocity v and caught at the height from which it was launched, t is easily seen to equal $2v/g$, $g = 32$ ft/sec^2 being the acceleration due to gravity), its orientation on landing simply depends on whether $\cos(\omega t)$ is positive (heads) or negative (tails). Keller showed that for any (smooth) probability distribution of uncertainties in initial conditions, the integrated odds of heads to tails tends to 50:50 when the distribution is translated far enough out in the (ω,t) plane—that is, coin tossing is fair in the limit of sufficiently vigorous flips as long as there is some uncertainty in initial conditions.

Intuitively, fairness follows from the fact that the curves $\cos(\omega t) = 0$ tend, at the outer reaches of the plane, to look like equispaced straight lines, and the farther outward you go the more closely spaced they become. If, for example, the probability distribution is uniform in a box centered at (ω_0,t_0), then the odds are simply the ratio of the total area for heads to the total area for tails, which tends to 50:50 as ω_0 and t_0 go to infinity.

Keller's analysis would seem to have settled the matter, but he left out a key feature of real coins: precession. A real coin doesn't just spin around a diameter when it's tossed: It also rotates around its center. Many coins, in fact, never turn over at all, but merely wobble on their way up and down, like tiny metallic pizza dough crusts. (In a "low-tech" experiment with a ribbon taped to a coin, Diaconis found that in 4 of 100 tosses the coin never turned over. "Now that's bias!" he says.)

The complete description of a tossed coin has one more parameter: the angle ψ that the angular momentum vector makes with respect to the normal to the coin (assumed to point upward initially). This extra degree of freedom, it turns out, is hugely important. Keller's analysis assumed $\psi = \pi/2$; in a "total cheat" coin—one that wobbles without turning over—ψ is close to 0. The question is, what happens in between?

Diaconis, Holmes, and Montgomery (who is an expert in celestial mechanics—see *SIAM News*, July/August 2001, http://

www.siam.org/siamnews/07-01/dynsys.pdf) began with a simple exercise in freshman physics: deriving a formula for the angle that the normal to the coin makes with respect to the up direction as the coin precesses around the angular momentum vector. If ω is the rate of rotation (angular momentum divided by the coin's moment of inertia, which depends on its mass, radius, and thickness), a matrix calculation shows that the vertical component of the normal at time t is

$$\cos^2\psi + \sin^2\psi\cos(\omega t).$$

This implies that what separates heads-up from heads-down regions in the (ω,t) plane is the equation $\cos(\omega t) = -\cot^2\psi$, from which it follows that the coin never turns over if ψ is between 0 and $\pi/4$ (or between $3\pi/4$ and π). That large range of ψ values explains why it's relatively easy to cheat when tossing a coin—with practice, you can make it look like the coin is spinning like mad, but still have it come up heads almost all the time.

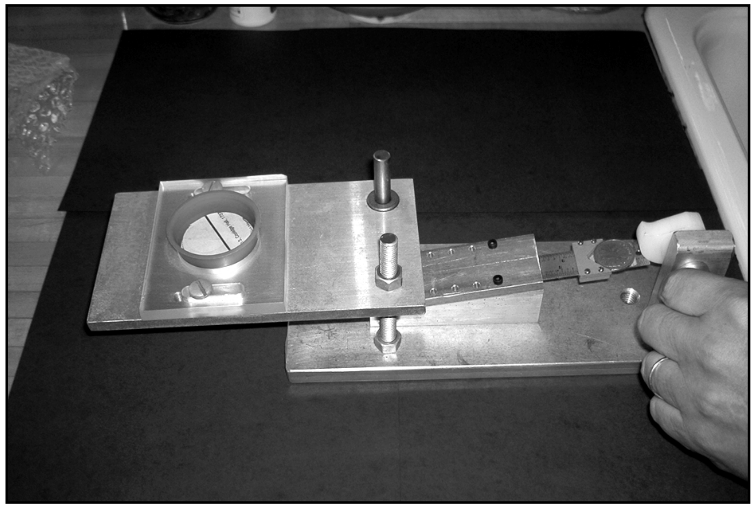

A coin placed heads up in a coin-tossing machine built for Diaconis by Melissa Franklin in the physics department at Harvard University gives the same result—heads up—every time.

When ψ is between $\pi/4$ and $3\pi/4$, regions in the (ω,t) plane do correspond to tails, but they are thinner than those for heads. The main result of Diaconis and his colleagues is that for all (smooth, compactly supported) probability distributions, the limiting probability of heads is identically 1 when ψ is between 0 and $\pi/4$ or $3\pi/4$ and π, and equal to

$$1/2 + (1/\pi)\sin^{-1}(\cot^2\psi)$$

for ψ between $\pi/4$ and $3\pi/4$. In other words, except for Keller's perfectly tossed coin, all coins are biased to land heads up. In short, it is impossible to toss a coin fairly.

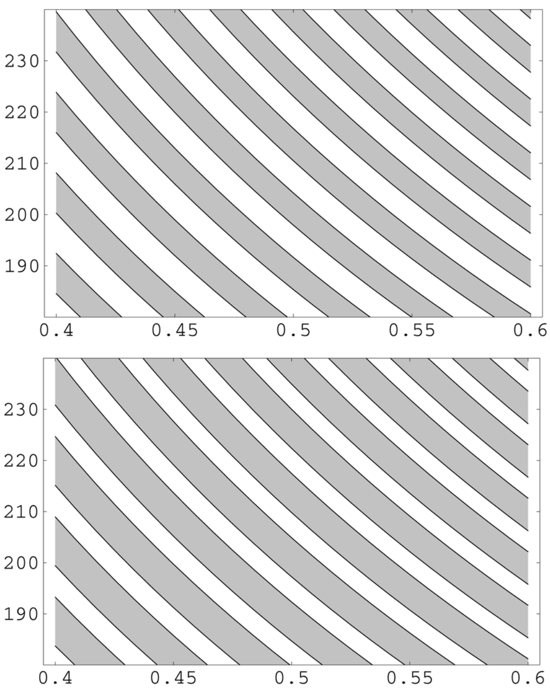

Regions for heads (shaded) and tails (white) in the region $0.4 < t < 0.6$, $180 < \omega < 240$ determined by the equation $\cos(\omega t) = -\cot^2\psi$ for two values of the angular momentum angle ψ: Keller's fair toss, $\psi = \pi/2$ (top), and a highly biased toss, $\psi = \pi/3$ (bottom). Most actual tosses are not nearly so biased.

The new formula is based on a limit of a coin tossed infinitely high with infinitely much angular momentum. Real coin tosses, of course, aren't that vigorous. In most cases the coins ascend only about a foot, which corresponds to an initial upward velocity of 8 feet per second and half a second in the air, and spin at between 35 and 40 revolutions per second. The angle ψ also varies from toss to toss. To get an idea of the actual distribution for people tossing coins, Diaconis and colleagues recorded approximately 50 tosses with a high-speed camera and analyzed the images. In each frame, the coin looks like an ellipse. The normal can be computed from the major and minor axes, and, if hundreds of images are available, the data can be fit to the known equations of motion. These empirical studies have led to the 51:49 estimate.

A true test of that estimate, Diaconis points out, would require a sample of about thirty thousand coin tosses, which is not beyond the realm of possibility. Other factors could affect the outcome, however. One is the actual angle at which the coin leaves the thumb. The model assumes the normal at the start to be straight up; in fact, the coin probably leaves at a range of angles. Another is whether the coin is allowed to bounce on landing. Bouncing introduces a whole new set of equations, with coefficients of restitution and friction, that defy simple analysis. Lingering questions also include the effects of air resistance (totally ignored in Keller's and the current model) and of inhomogeneities in the mass distribution. Finally, in practice, when you take a coin out of your pocket and put it on your thumb, do you bother to notice which side is up?

Can This Equation Be Saved?

Psychologist John Gottman and mathematical biologists James Murray and Kristin Swanson, all of the University of Washington, described a mathematical model Gottman's group has developed to identify—and possibly help—troubled marriages. (The AAAS session, appropriately enough, was held on Valentine's Day.) The model is based on a pair of difference equations that represent the changing attitudes of a husband and wife during an intense, 15-minute conversation:

$$W_{t+1} = b + r_W W_t + I_{HW}(H_t)$$
$$H_{t+1} = a + r_H H_t + I_{WH}(W_{t+1}),$$

where H_t and W_t represent the husband's and the wife's attitudes after the tth interchange. Roughly speaking, the constants a and b represent a base-level attitude that each person brings to the marriage, while r_H and r_W are "coefficients of inertia," i.e., the tendency to maintain the same attitude. Central to the model are the "influence functions" I_{WH} and I_{HW}, which express the extent to which the wife's attitude affects the husband's, and vice versa. (The slight asymmetry in the subscripts t and $t + 1$ reflects the convention that the wife speaks first.) Without the influence functions, the spouses may as well not be aware of each other's existence (which is actually the case for some couples—an individual's influence function can be identically zero).

In principle, the influence functions can be arbitrarily complicated, but the marital modelers have found that piecewise linear functions, with a single bend at the origin (i.e., neutral attitude), give a good fit to the facts. Intuitively, the kink corresponds to the notion that the intensity of a person's response to feedback depends on whether the feedback is positive or negative; some people respond more strongly to negative feedback, others to positive.

Psychologists at Gottman's "Love Lab," as his research group is affectionately called, study videotapes of a couple's conversation and convert what they see and hear into numerical scores, the H_t's and W_t's. The time series is then analyzed to provide estimates for the parameters a, b, r_H, r_W, and the slopes of the lines for the two influence functions. That's where the math kicks in. An elementary analysis of the equations guarantees at least one, and typically two, steady states. If there is a stable equilibrium with both partners in the positive quadrant of the HW-plane, the marriage is in good shape; if the only steady state has one or both partners in a negative quadrant, watch out.

Among the conclusions the researchers have drawn from their studies is that, the opening sentence of *Anna Karenina* notwithstanding, a couple can be happy in more than one way. Conventional counseling wisdom holds that there is just one "healthy" way for couples to interact: Each partner should "validate" the other's point of view, even as each tries to get the other to change her/his mind. But the Love Lab has found two additional styles to be consistent with marital longevity. One is the classic volatile couple, who seem to argue endlessly but also seem to enjoy their rough strife. The other is the confrontation-avoiding couple, who simply agree not to discuss sensitive subjects.

What characterizes marriages headed for disaster is hostility, either overt or detached. (Contempt, Gottman says, is the best predictor of divorce.) Simple as the model is, the equations are uncannily accurate at distinguishing stable from unstable marriages. In four studies of more than 700 couples over the last decade, the model has proved more than 90% accurate in predicting divorce.

Moreover, the analysis of the dynamics offers insight into what goes right or wrong. Indeed, Gottman and colleagues have begun to design interventions based on the model. Roughly speaking, they look at the parameters that describe a couple's current interactions, and then suggest behavioral changes (e.g., "you might want to stop sneering at your spouse") that nudge the parameters into a range that include a positive steady state.

Gottman, Murray, and Swanson, with co-authors Rebecca Tyson and Catherine Swanson, have written a book (*The Mathematics of Marriage*, MIT Press, 2003) summarizing a decade of research. (They are currently working on intervention strategies and are also studying gay and lesbian couples.) The application of mathematics to the study of human emotions may still raise a few eyebrows, but the marriage of math and psychology is one that's likely to last.

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.