

The Tea Leaves of Sports Talk

Finding Meaning in Random Sequences

A N D R E W B E R N A R D I N

ALTHOUGH I ENJOY PLAYING BASKETBALL, I cannot claim to have ever experienced a "hot hand." Luke-warm, maybe. Nor have I ever found myself "in the zone." Even if equipped with a G.P.S. and a topological map, I don't know if I could even locate the zone. Perhaps because of this, I will endeavor to deconstruct the zone and plunge the hot hand in a bucket of ice water.

The non-phenomenon of basketball's hot hand has been addressed elsewhere (see Michael Shermer's review of the book, *Intuition: Its Powers and Perils*, in Vol. 10, No. 1 of *SKEPTIC*). In brief, the findings indicate that a player who has hit a number of shots, and is thus considered hot, is no more likely to hit the next shot than his or her overall shooting percentage would predict.

So there's no hot hand. Is this a big deal? I think so. The type of superstitious thinking behind the hot hand operates far beyond the basketball court. And although it may provide entertainment, when unchecked it can lead us seriously astray.

Half a decade ago, tennis champion Goran Ivanisevic played at Wimbledon. That year he went on a winning streak. He advanced further and further toward the finals. Goran was one of the top players in the world, so this wasn't all that surprising. Nonetheless, Goran didn't want to jinx himself. So, game after game, he wore the same hat, which he did not bother to launder. Maybe he was aware of research showing that detergent can wash away benevolent spirits, positive vibrations, and quantum luck tunnels. Why Goran believed his hat was more instrumental to his good play than, say, his underwear, I don't know. With each match the hat got dirtier and sweatier and, no doubt, smellier. Was Goran's hat responsible for his advancing in the tournament? I'm sure Adidas, Nike, or whichever company had their logo on it, wished this were so.

With the help of his hat, Goran made it to the finals. Also making it to the finals was the

world rated number one Goliath of Wimbledon, Pete Sampras. Goran continued to play well in the finals, but his hat was apparently insufficiently powerful to overcome Sampras' dominating service game.

Superstitious behavior, whether on a tennis court, basketball court, football field, or golf course, is relatively easy to recognize. Everyone can plainly see, for example, a baseball player making a sign of the cross before getting into the batter's box (*in the name of the father, the son, and Babe Ruth, help me send the ball into the left-field bleachers*). Another form of sports superstition, finding meaning in random sequences, has a broader following than the quirky rituals and good luck charms of individual players. It permeates the talk of broadcasters, athletes, and fans; it infiltrates the minds of those drawn to games of all types.

Superstitious behavior has been called irrational. Which it is. Yet while many people consider irrationality to be the consequence of faulty logic, much of it can be attributed to a straightforward problem of information misuse and/or neglect. When we are being rational, we properly weigh information to come to some conclusion or decision. Irrationality can thus be seen as the result of one or a number of the following: too little information, biased selection of information, poor quality information, and the improper weighing of information.

Backspin-doctoring Sports

One day in early 2003, I became aware of my own tendency to think irrationally in the domain of sports. I had played basketball and performed poorly. Many more of my shots found the outside of the rim than the inside of the net. I lingered on the court, trying to diagnose what had gone wrong. Why was my shooting "off?" Was I giving the ball enough backspin? I put up a shot with more backspin, and *swish*. Ah-hah. I took three more shots and hit all three. So that was it. *I wasn't putting enough backspin on the ball*. This was the hopeful conclusion I came to after hitting four shots in four attempts. I headed to the shower enthused, imagining that my play would be forever improved.

Fortunately, I am a skeptic. That is why in the following weeks I began to test my conclusion. Over a number of sessions I took 800 free throws (shots from the penalty line): 400 without

concentrating on anything other than trying to get the ball through the hoop, and another 400 while concentrating on giving the ball backspin. I chose 400 because I considered that number sufficiently large to rule out the influence of what some would call luck but what I think of as a non-representative sequence or "turn of events." Because I am not a mathematician, and perhaps because my brain lacks an adequate level of serotonin, I wanted to err, if I did, on the side of too many trials of what I consider to be a pleasingly repetitive behavior.

The results: I hit 278 of the 400 "regular" shots, or 69.5%. Of the "backspin" shots, I hit 267 of 400, or 66.8%. Evidently, adding more backspin to my shooting was not going to transform my game.

What had caused me to come to the faulty conclusion that additional backspin was the difference between poor shooting and good? Too little information. There is not enough data contained in only four trials to develop any sound conclusions. And herein lies the key to superstition in sports: humans have the habit of bringing a quick-draw, deterministic mindset to probabilistic phenomena.

With a deterministic event, there is a strict relation between cause and effect. If I hold a coin in my hand and drop it, the coin will fall to the ground. Each and every time I drop the coin it falls to the ground. Because the relation is deterministic, I can accurately predict the outcome. A probabilistic event, on the other hand, entails a relation between cause(s) and effect that is not one to one. There is, instead, a probability that a particular outcome will follow. In the case of the coin, the effect of it landing heads up cannot be strictly determined. I can only give the probability that it will land heads up: 50%. Try as I might, I cannot determine the outcome of the drop. As importantly—and central to the point—I also cannot, in hindsight, ascribe causal reasoning as to why the outcome was one way or the other. The coin didn't land tails-up because it wanted to, or because I'm unlucky, or because I was leaning on my left foot. The coin simply landed tails-up and that's all I can say about it. As a human I may be wired to find a simple reason for everything, but when it comes to probabilistic events, I must keep that tendency in check. If I don't, I will end up talking gibberish.

Because there isn't a one-to-one correspondence between cause and effect with probabilistic




WWW.SKIPNOC.COM

SKIPNOC

the Humanist

a magazine of
progressive &
rational opinion
on today's
leading concerns



1777 T Street, NW
Washington, DC 20009-7125
202.238.9088 • 800.837.3792
www.theHumanist.org

events, it becomes necessary to gather a lot of information. For instance, imagine you are visiting a foreign planet and you want to know whether a balloon filled with oxygen gas will rise or sink in the windless atmosphere. In this case, you would only have to release the balloon a few times to determine the answer. After a limited number of trials you could predict what would happen. And you could offer a simple explanation as to why it did happen—something about the difference between the weight of the oxygen molecules in the balloon and those in the surrounding atmosphere. In this scenario, the balloon falling or rising would be a deterministic event. You could accurately predict what would happen. You could also offer a straightforward explanation for what had happened.

Suppose you have established that the balloon sinks. Whether or not it will pop upon hitting the uniformly jagged surface of the planet is another question. The answer to this question may belong in the realm of probability. If you dropped the balloon twice and it didn't pop, you might be foolish to predict that the third time it also wouldn't pop. If a dropped balloon popped sometimes, but not others, then you could not determine the outcome of any subsequent drop. You would have to drop the balloon many, many times to establish the probability of what would happen. And once it happened, you could only explain any particular outcome by referring to the probabilities that you established.

The above doesn't mean you are unable to learn or know anything about the balloon and its tendency to pop. You may discover, for example, that imparting a spin to the balloon increases the likelihood it will pop. However, to figure out what the probability is you couldn't rely on a few trials. Any conclusions drawn from a limited number of trials, from too little information, would be meaningless. So you make a great number of trials and discover that imparting a spin to the balloon increases its probability of popping by, say, 10%.

On your next trial, you spin the balloon as you release it, and it pops upon hitting the ground. Can you then conclude that the spin caused it to pop? No. It may have popped regardless of the spin. There's no way for you to know whether the spin played a role in that particular pop or not.

The Precisely Imprecise Nature of Sports

Sports are exciting because we cannot precisely

SKEPTIC

VOLUME 11 NUMBER 2 2004

predict the outcome. If the team to win the Super Bowl could be forecast by adding the weight of the players for each side and comparing them, what fun would that be? Las Vegas, for one, would lose interest.

Sporting events are probabilistic because human behavior is more like the flight of shotgun pellets than laser-guided missiles. Our elbows are not made of titanium. There is always a degree of imprecision—and thus randomness—in how we perform. Even the great Tiger Woods, under ideal conditions, is not going to shoot under par every time he plays golf. His play is imprecise. And when Tiger is imprecise, he is not precisely imprecise. The alternatives to a perfect drive or putt are many and cannot be predicted. Our imprecision generates random outcomes. And, whisperings of golf commentators aside, there is no meaning in the random.

If you situated Tiger Woods 200 yards from the green and told him to hit the ball within 20 feet of the pin, you could not predict whether he would succeed on one shot or the next. After hitting 100 balls, the result would most likely resemble a shotgun blast—white pellets marring a green target, some of the marks inside the designated bulls-eye, some outside. If Tiger hit 60% of his shots within the bulls-eye, what could you say about those 100 shots? Well, from that distance Tiger hit 60% of his shots within 20 feet of the pin.

Now imagine witnessing those 100 shots individually, as each one was made. What could you conclude about a single shot having hit or missed the target? Could you say that Tiger was distracted if he missed, or that he was “feeling it” if he hit? No. The same goes for short strings of hits and misses. Ten hits in a row does not mean that Tiger is on fire. Likewise, if Tiger misses his next 10 out of 15, we cannot rationally speak about some drop in performance.

All-Star George's Hoop Dreams

The game of basketball is particularly well suited to this discussion because the better players hit roughly 50% of their shots. What is the probability that a player is going to score on any given shot? Toss a coin. What is the reason a player hit or missed a specific shot? Unless the shot was deflected by an opponent, we cannot speak rationally about the cause of a single, probabilistic event.

One night, while reading in bed, I wondered

what would happen if I pretended a tossed coin was a basketball player. What would “his” performance look like? Would “he” get as hot and cold as the pros? I fetched a quarter and dubbed it “All Star George.” Because marquee players put up roughly 20 shots per game, I drew a 20 x 5 grid. George would shoot 10 shots in the first half, and 10 in the second, for a 5 game series, resulting in a total of 100 shots. While sitting up in bed—not what you would call controlled laboratory conditions—I began flipping the coin, making “m” (miss) and “h” (hit) marks on the paper.

In the first half of the first game, the sequence looked like this: m-h-m-m-m-h-h-h-m-m. At half-time, George had hit 40% of his shots. As you can see, though George is a 50% shooter, his hits and misses do not predictably alternate. We know that. At least one part of our brain knows that.

For the second half the results were surprising: h-h-h-h-h-h-h-h-m-h. My quarter, George, had “caught fire,” making 90% of his shots. What had inspired George to do this? Combining stats for both halves, in the first game All Star George hit 13 of 20, or 65%. My quarter was obviously trying to impress me.

During the second game, George hit 50% of his first half shots. Would he again catch fire in the second half? The ensuing sequence again surprised me: m-m-m-m-m-h-m-h-h. George went on another streak, but this time it was of misses. He hit only 3 of 10, although he did “pick it up” toward the end of the game. George's shooting percentage for the second game was 40%. Maybe his legs hadn't recovered from the great effort he put into the second half of the first game.

In the middle of the third game I had to bring in another player. George the 1st had rolled under the bed. It was at this time in my *flip, catch, slap, scribble* routine that my wife looked over the top of the book she was trying to read. She told me that if I didn't wrap it up soon, she would cancel my research funding.

Although I knew the probability of George hitting heads was 50%, I could not predict what would happen on each individual flip. Nor could I predict what George's shooting percentage would be for half of a game. When I completed the 5 games, I found that in only 4 of the 10 halves (10 shot sequences) did George hit the expected 50%. For the first halves of all 5 games, George hit 22 of 50 shots, which is 44%. For the

second halves of the 5 games, George hit 28 shots of 50, or 56%. Could I have concluded from these 100 shots that George tended to “step it up” in the second half of each game? No. Numbers clipped from limited sequences are no more meaningful than the pattern left by tea leaves in the bottom of a cup. Once again, for probabilistic events we need *a lot* of information.

If you're quick with math you will have noticed that over the 5-game series George went a perfect 50 for 100—just what one would expect. What are the chances of that? I imagine it's at the peak of a bell curve of probability, but not the only statistically likely outcome.

The Nicks v. the Twenty-Fivers

I had conducted my quarter-flipping test during the NBA playoffs, which followed on the heels of the NCAA tournament games called March Madness. While watching the basketball contests, I had been hearing statements such as: *he's really raised the level of his game* (after a string of hits), *they're looking for someone to get a hot hand* (I kid you not), *they're completely out of sync* (after a string of misses), *that basket will really build his confidence* (same thing happened for George, I think), *he's making some good shot selections* (string of hits), *he's rushing his shot* (he missed), *I couldn't find my rhythm* (he missed a lot), and *he hasn't stepped it up* (after a player went 1 for 6, rather than his average 2.9 for 6).

Many of the NCAA games were so close I wondered if the winner could be said to have been determined by the toss of a coin...by luck (a sequence terminating at a fortuitous point?). Though sports commentators may jabber about the winning team having “brought their A game,” in such a close contest, and with all other things being equal—for instance, one team not having lost their best player to injury—could the commentators come to any rational conclusion about the outcome?

For my next coin-tossing project, I decided to re-enact a game. I would pit one coin against another, alternating tosses like trips up and down a basketball court. I called one team the “Nicks” and used a nickel to represent them. The other I called the “Twenty-Fivers.” Each coin would be tossed 100 times during the game. For each hit, or “heads,” the team would be awarded 2 points. Due to the probability of a coin coming up heads, I figured the final score would be, oh,

near-100 to near-100. But would either coin take a dramatic lead during the contest?

As in basketball, I broke the game into two halves with two periods each half (to avoid confusion, I'll forgo the customary term “quarters”). It began with a coin toss. I must admit I was excited to see who would win. Because I have the sad habit of rooting for the underdog in almost any contest—and am thus frequently disappointed by the outcome—I secretly rooted for the smaller Nicks.

Ten shots into the game, the Nicks were up 14 to 10, having hit 7 of 10. Twelve shots later they were up 16 to 10. At 20 shots they were up by 4, having hit 2 more shots than the Twenty-Fivers. The Nicks increased their lead by the end of the period to 28 to 22. For that period, my Nicks had shot 56%, the opposition, 44%.

During the first half of the second period, each coin went on a mini-run, with the Nick's lead increasing, then falling back to 6 points. Then, during the end of the second period, my team went flat, hitting only 2 of 8. Meanwhile, the Twenty-Fivers went an impressive 7 for 8 and took the lead. The score: 50 to 56. In the second period the Nicks had shot a somewhat wimpy 44%, while the Twenty-Fivers razzle-dazzled their way to a red-hot 68% shooting.

I gave my Nicks a pep talk during half-time. I know, a good scientist tries to eliminate all bias, but I couldn't help myself. “You can do it!” I coached my coin. I was both amused and dismayed by my sincere excitement over a competition involving the tossing of two coins, but this was high drama.

Third period. Although neither team scored above average, and both went 0 for 6 at some point, my Nicks did particularly poorly. For the period—a 25-flip sequence—they shot a dismal 36%. Hadn't they listened to me? Meanwhile, the Twenty-Fivers shot a respectable 48% and increased their lead to double digits. The score after three periods of play: 68 to 80. Could my Nicks rally in the fourth period?

During the final period of play, both teams shot “lights out.” My Nicks put together two impressive strings, hitting 7 for 7—they were in the zone—then later going 5 for 5. Unfortunately, the Nick's 68% shooting in the final period wasn't enough. The Twenty-Fivers shot 64%, winning the game 112 to 102. Apparently the Twenty-Fivers “wanted it more.”

The quarter I tossed landed heads-up 56 times out of 100, 6% greater than expected. Is this unusual? The result, certainly, is somewhat farther out of the fat of the bell curve than a perfect 50% would be. Maybe I could repeat the project 100 times and not get the same result. And that is what makes sports so exciting. Beyond the real human drama, such as raw ability, strategy, injuries, rivalries, etc., sporting events entail a virtual mosh-pit of probabilities.

Science Interrupted

This is not to say that the outcome of sporting events is completely random, the result of a coin toss. Were a pee-wee football team to play the Tampa Bay Buccaneers, the 2003 NFL Super Bowl champions would win. Probably. And that probability is very, very high. That said, if the two teams played a game every day for a million years, one day the pee-wees might win one. Just as if you were to prop a monkey at a keyboard, sooner or later, through random typing, he might string out the entire *Origin of Species*. It may take a billion trillion years, but it could happen. Trying to grasp the meaning, or lack thereof, of probabilities can be like forcing a square peg into a round hole—we aren't accustomed to thinking that way.

Despite the fact that we can rationally say so little about the causes of an individual hit or miss, or even short strings of hits or misses, whether they be in golf, basketball, tennis, baseball or even croquet, human beings tend to jump to conclusions. We like to find meaning for things, and so we do, whether or not this is rational behavior.

In the first game of the 2003 NBA finals, the New Jersey Nets hit their first three shots. They had just come off a one-week layoff, having waited for the Spurs to finish out their series with the Los Angeles Lakers. The television commentators, always in need of something to talk about, had speculated as to whether the time off would help the Nets or hurt them. After going three for three, one of the talking-heads declared that there was "no rust on them." Had the Nets missed their first three, I'm sure I would have heard how they were rusty. Funny, my tossed coins never got rusty, though they did go on streaks.

In a sense, superstitious thought is science interrupted. A relationship is sought and found, but there is a serious lack of follow-through. The conclusions are premature; they are based upon

too little information.

During the closing seconds of game two of the NBA finals, the Spur's Steven Jackson missed a potentially game-winning three-point shot. He had hit his previous two. One of the commentators, an ex-player who really should understand the game better, proposed that Jackson missed the relatively wide-open shot because a Nets player was rushing at him and "when someone is coming at you like that, it can throw your shot off." There had to be a distinct reason why the shot missed.

Maybe, just maybe, Jackson missed that shot because he is a 40% shooter from that distance (give or take 10%). Under the best of circumstances, such as on an empty court while fully concentrating, Jackson will miss shots, for no particular reason. Why? Because a human athlete is imprecise, and this imprecision produces random outcomes. Nothing meaningful can be said about randomness.

During that same game, the star player for New Jersey, Jason Kidd, hit his last 6 free throws to "win the game" for the Nets. (Why the final points in a game, and not the initial points or middle points, are responsible for a win is another issue.) The still-sweating Kidd was interviewed. "How did you do it?" he was asked. Jason replied, "I knew we had to win this one and there wasn't going to be a second chance."

So Kidd put his will power into over-drive? Why wouldn't he *always* give 110%? The answer: because fortuitous sequences eventually come to an end. For a 70% free-throw shooter to hit 6 in a row is not a Herculean feat, as much as fans and announcers would like to view it as such.

In game Five, Jason Kidd hit his first 4 shots. One commentator concluded that Kidd had definitely brought his "A" game. The other concurred: Kidd was "hot." Apparently it is possible to cool off instantly while simultaneously falling into one's "D" game. Kidd hit only 1 of his next 9 shots. Both sports commentators had jumped to a premature conclusion.

Why do irrational beliefs such as the hot hand persist? Perhaps, in part, because the need persists. It is difficult to project human motives and meaning onto raw numbers such as 50%. But with a sequence such as *miss-bit-miss-miss-miss-bit-miss-bit-bit-bit*, a dramatic story can be told about hot hands and cold chokes. Each event has its story, every turn, a reason. ▼

