On August 18th, 1913, the ball of a roulette wheel at Le Grand Casino of Monte Carlo landed on black 26 consecutive times. During this streak a betting frenzy ensued as a growing crowd staked increasingly large amounts on red. They believed that the ball was "due" to land on red. They lost and the casino reaped a large reward, thanks to Gambler's Fallacy (GF). GF is well-known and yet it remains a powerful intuition for most of us. Why? Where does its seductive power come from?

To begin with, what exactly is the GF? Let's take a look at some of the readily available definitions.

Wikipedia's definition: *"The Gambler's fallacy…is the belief that if deviations from expected behaviour are observed in repeated independent trials of some random process then these deviations are likely to be evened out by opposite deviations in the future."*

Not bad. The crucial premise is a sequence of *independent* events, i.e., events whose outcomes do not depend on previous events. Actually, they don't have to be random. All that is required is that the next event not be predictable by the previous ones.

Skeptic's Dictionary definition: *"The gambler's fallacy is the mistaken notion that the odds for something with a fixed probability increase or decrease depending upon recent occurrences."*

This is a special case of GF. The probability doesn't have to be constant. Again, all that's needed is that the next event not be predictable by the previous ones. That goes for the probability of a particular outcome in the sequence too.

The Investopedia definition (also reproduced here): *"When an individual erroneously believes that the onset of a certain random event is less likely to happen following an event or a series of events."*

Not even close. Can an event change the probability of a future event? Of course it can. But if it does, that means the future event depends on the earlier one.

The Internet Encyclopedia of Philosophy definition: *"This fallacy occurs when the gambler falsely assumes that the history of outcomes will affect future outcomes."*

Perfectly circular. The GF is a fallacy—why? Because it's a false assumption. There are plenty of circumstances under which previous outcomes can affect future ones, so all you have to do to avoid the GF is assume this when it's true…?

So, at the most general level, the key components of GF are:

1. We're observing a sequence of events that are independent of one another.
2. Our minds tell us that they aren't independent.
3. We therefore make predictions about the sequence based on its earlier events.

The specifics are what could give us some clues about why this happens. Our minds tell us the events aren't independent, so how do we think they're related? And why would our minds tell us they're related in that particular way?

A widely-cited explanation given in 1974 by Amos Tversky and Daniel Kahneman attempted to answer both questions. They claimed that people use a representativeness heuristic, whereby people judge the probability of an event by its resemblance to the data available to them. Thus, people expect that a short run of random events should have the same properties that it has "in the long run." If we believe a coin is fair and therefore the probability of it landing Heads is 1/2, we think that a run of 9 consecutive tosses landing Tails is not representative of the coin's long-run properties and so we expect the 10th toss to land Heads, just like what happened to the misguided gamblers in Le Grand Casino.

The problem with the representativeness heuristic explanation for GF is that it begs the "why" question. It may well be that we expect a short run of random events to share its long-run properties, but where does this expectation come from?

Alex Bellos, the author of a fascinating book Adventures in Numberland, makes a stab at the "why" question. He thinks it's a motivational matter of trying to establish control over events by reading patterns into them. That's plausible, but if so then it's largely unconscious. Also, the intuition behind GF seems sufficiently powerful to not require a motivational explanation.

More problems with the representativeness heuristic explanation arise when we look at a related fallacy, namely the belief in the hot hand. Like the GF, you can find somewhat different definitions in various places.

Wikipedia definition: *"… the hot hand fallacy is the idea that a streak of positive successes are likely to continue."*

Fallacy Files definition: *"A gambler has had a streak of luck. Therefore, the gambler is either "hot" or "cold", depending on whether the luck is good or bad, and the good or bad luck will continue at a probability greater than chance."*

Changingminds.org definition: *"The Hot Hand Phenomenon occurs where people believe that 'success breeds success' such that when a person succeeds at something then they are more likely to succeed in subsequent attempts, whereas the truth is that they are still governed by the laws of chanced [sic]."*

So, a hot hand belief is a special case of a belief that a sequence of identical (or similar) events will continue to produce similar events. It's a fallacious belief when the events in the sequence are independent of one another, so that the occurrence of one does not predict anything about its successors. A hot hand belief is a belief in a *positive* correlation between earlier and later events in a sequence, whereas the GF is a belief in a *negative* correlation between earlier and later events. The representativeness heuristic has been used to explain both of these, but that begs the question of when or why people would find positively correlated sequences more "representative" in some situations and negatively correlated sequences more so in other situations.

James Sundali and Rachel Croson (2006) studied real betting and belief patterns in games of roulette. They pointed out that there are two kinds of beliefs that players bring to bear: Beliefs about players and beliefs about outcomes. They found that players' betting patterns fell into two categories. Those who acted consistently with GF (betting on numbers that hadn't appeared previously) were more likely to also act as if they believed in the hot hand (increasing the number of bets after a win). But players who acted consistently with a belief in the hot *outcome* (betting on numbers that had already appeared) were more likely to act consistently with a kind of GF belief about players' luck (decreasing the number of bets after a win because one's stock of luck was due to run out).

Sundali and Croson speculated that locus of control could account for this pattern of individual variability. Players with an "external" locus of control believe in luck and are more likely to believe in lucky streaks (hot outcomes) and a stock of luck that can run out. Players with an "internal" locus of control believe they can learn the underlying principles by which, for instance, a roulette wheel operates (whence GF) and also are more likely to have hot hand beliefs.

Ayton and Fischer (2004) came up with yet another explanation, simply based on learning from experience: “(1) The hot hand fallacy arises from the experience of characteristic positive recency in serial fluctuations in human performance. (2) The gambler’s fallacy results from the experience of characteristic negative recency in sequences of natural events, akin to sampling without replacement.”

Andreas Wilke and H. Clark Barrett (2009) take this experiential explanation for hot hand beliefs one step further. They point out that humans didn’t evolve in casinos but instead in environments where resources tend to be “clumped” as is the case for food resources in foraging environments, and they claim the hot hand intuition evolved from successful foragers being more likely to pass their genetic material on to succeeding generations.

In my 1997 paper, "Judgment Under Chaos," I presented both the learning-from-experience and evolutionary explanations a decade before Wilke and Barrett, and applied them to both GF and hot hand beliefs. My evidence was a demonstration that humans are good at short-range predictions of chaotic processes. To do this, I created a chaotic attractor with a parameter that controlled the extent to which the resulting sequence was anti-persistent (i.e., frequently oscillating as would be expected by GF beliefs) or persistent (infrequently oscillating as expected by hot hand beliefs). Examples of both kinds are shown in the figure below.

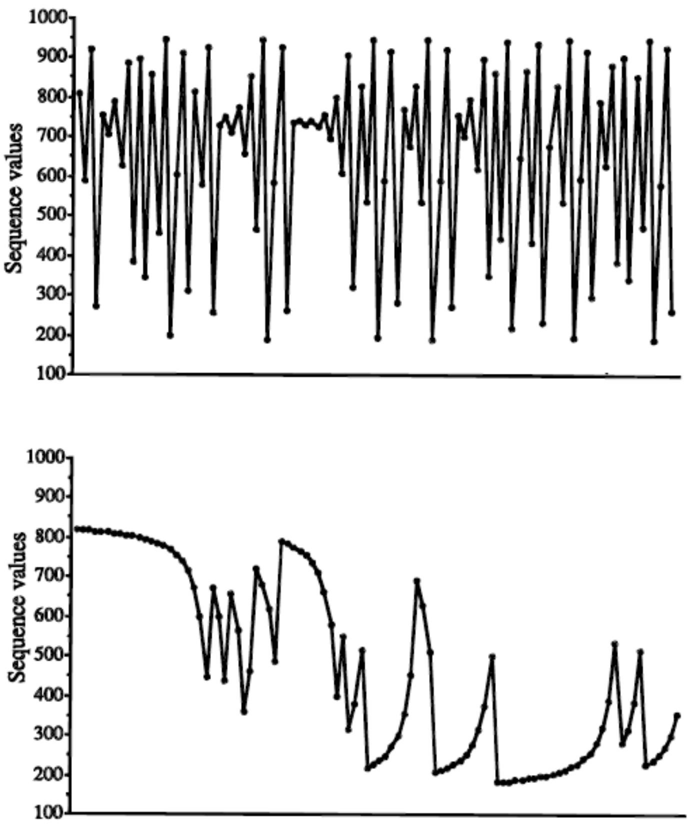

FIG. 1. Anti-persistent (top) and Persistent (bottom) sequences.

Participants in the study were randomly assigned to either tracking a persistent or anti-persistent process, and after seeing 100 iterations of their sequence they were asked to predict its next value for 50 more iterations. They were also asked to do the same for a random sequence containing exactly the same values as their chaotic sequence (but randomly re-ordered).

Of course, no one was any good at predicting the random sequence. I suspect that when we see a random sequence like coin-tosses, we misread it as an anti-persistent process and therefore fall for GF. But everyone was quite good at predicting the chaotic one (the average correlation between their predictions and the sequence itself was .68, and some people's correlations were around .9). Moreover, they were equally good at predicting persistent and anti-persistent processes. So, Ayton and Fischer's learning-from-experience proposition received solid support before they wrote their paper. The Wilke-Barrett evolutionary explanation gets a circumstantial evidential boost too. We didn't evolve in casinos; we evolved in environments full of anti-persistent (like prey-populations) and persistent (like the weather) processes. We expect the world around us to have “memory,” and for good reason: It usually does.