The Gambler's Fallacy

Kevin Dorst
kmdorst@mit.edu

24.223 Rationality

I. Bayesian Gamblers

Baylee is bored. The Prius has won the front spot 3 days in a row, so she bets the Jeep will get it next. The majority (and the *average*) of her coworkers agree—expecting it to switch. A minority think the Prius is "hot", so its streak will continue.

They've been doing this for months.

In fact, the outcomes of the parking battle are statistically independent.

The majority commit the **gambler's fallacy**—the tendency to think that streaks of (in fact, independent) outcomes are likely to switch.
The minority commit **hot hands fallacy**—the tendency to think that streaks of (in fact, independent) outcomes are likely to continue.

Is this evidence for irrationality? ("Law of small numbers"?)
Claim: No! It's what we should expect from rational Bayesians.

With limited data or memory.

Why? Rule out "streaky" hypotheses quicker than "switchy" ones.
$\rightarrow$ Would exhibit many of the subtle empirical trends found.

So average toward latter.

So these findings are *also* evidence for:

Causal-Uncertainty Hypothesis: The gambler's and hot-hands fallacies are due to causal uncertainty combined with rational responses to limited data and memory.

II. Reasonable Uncertainty

For simplicity, suppose they know the long-run hit rate (= 50%).
Three (classes of) hypotheses that would guarantee that:

Steady: each day, the Prius has a 50% chance of winning.

Eg(?): coin.

Switchy: after Prius wins (a few), Jeep becomes more likely to win—and vice versa.

Eg: a deck of cards (red/black).

Sticky: after Prius wins (a few), Prius is more likely to win again—and vice versa.

E.g. basketball shots

Can draw using *Markov chains*: an arrow labeled t from x to y indicates that in state x, the system is t-likely to transition to state y.

Can vary in both *extremity* (how far deviate from 50%) and *speed* (how quickly they deviate. E.g. here are *2-step-to-70%* chains:

Put the longest miss-streak on the left, shortest streaks in middle, and longest hit-streak on right, e.g. $\langle 00, 10, 01, 11 \rangle$.

2-step-Switchy$_{70\%}$:

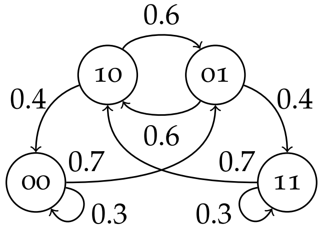

00	10	01	11
0.3	0	0.7	0
0.4	0	0.6	0
0	0.6	0	0.4
0	0.7	0	0.3

2-step-Sticky$_{70\%}$:

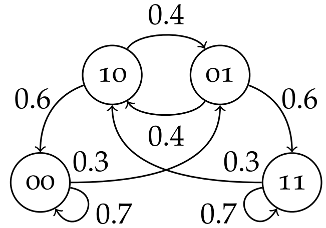

00	10	01	11
0.7	0	0.3	0
0.6	0	0.4	0
0	0.4	0	0.6
0	0.3	0	0.7

For concreteness, focus on *5-step-to-90%* chains (vs. Steady):

5-step-Switchy$_{90\%}$

$$\begin{pmatrix}
.1 & 0 & 0 & 0 & 0 & .9 & 0 & 0 & 0 & 0\\
.18 & 0 & 0 & 0 & 0 & .82 & 0 & 0 & 0 & 0\\
0 & .26 & 0 & 0 & 0 & .74 & 0 & 0 & 0 & 0\\
0 & 0 & .34 & 0 & 0 & .66 & 0 & 0 & 0 & 0\\
0 & 0 & 0 & .42 & 0 & .58 & 0 & 0 & 0 & 0\\
0 & 0 & 0 & 0 & .58 & 0 & .42 & 0 & 0 & 0\\
0 & 0 & 0 & 0 & .66 & 0 & 0 & .34 & 0 & 0\\
0 & 0 & 0 & 0 & .74 & 0 & 0 & 0 & .26 & 0\\
0 & 0 & 0 & 0 & .82 & 0 & 0 & 0 & 0 & .18\\
0 & 0 & 0 & 0 & .9 & 0 & 0 & 0 & 0 & .1
\end{pmatrix}$$

5-step-Sticky$_{90\%}$

$$\begin{pmatrix}
.9 & 0 & 0 & 0 & 0 & .1 & 0 & 0 & 0 & 0\\
.82 & 0 & 0 & 0 & 0 & .18 & 0 & 0 & 0 & 0\\
0 & .74 & 0 & 0 & 0 & .26 & 0 & 0 & 0 & 0\\
0 & 0 & .66 & 0 & 0 & .34 & 0 & 0 & 0 & 0\\
0 & 0 & 0 & .58 & 0 & .42 & 0 & 0 & 0 & 0\\
0 & 0 & 0 & 0 & .42 & 0 & .58 & 0 & 0 & 0\\
0 & 0 & 0 & 0 & .34 & 0 & 0 & .66 & 0 & 0\\
0 & 0 & 0 & 0 & .26 & 0 & 0 & 0 & .74 & 0\\
0 & 0 & 0 & 0 & .18 & 0 & 0 & 0 & 0 & .82\\
0 & 0 & 0 & 0 & .1 & 0 & 0 & 0 & 0 & .9
\end{pmatrix}$$

Since all three have 50% hit rate, suppose Baylee starts out rationally unsure (uniform) over 5-step Switchy$_{90\%}$, Steady, and 5-step Sticky$_{90\%}$.

$P(Swi) = P(Ste) = P(Sti) = \frac{1}{3}$

What about after observing limited data or memory? Updates on E.

Suppose the sequence <u>ends on a run of 0s</u>. Her credence it'll switch to 1 should be a weighted average of the three chance-hypotheses:

Total Probability + Principal Principle

$$\begin{aligned}
P(1|E) &= P(Swi|E)\cdot P(1|Swi) \quad + \quad P(Ste|E)\cdot P(1|Ste) \quad + \quad P(Sti|E)\cdot P(1|Sti)\\
&= P(Swi|E)\cdot(0.5+c) \quad + \quad P(Ste|E)\cdot(0.5) \quad + \quad P(Sti|E)\cdot(0.5-c)
\end{aligned}$$

Where c determined by streak length. If 1, $c = 0.08$, if 2, $c = 0.16$, etc.

→ If $P(Swi|E)$ is high, she should do gambler's reasoning.
→ If $P(Sti|E)$ is high, she should do hot-hands reasoning.
→ Only if $P(Swi|E) = P(Sti|E)$ should she treat tosses independently.

Whenever $P(Swi|E) > P(Sti|E)$, she should *guess* it'll switch.

So *exact* independence will be rare. Still, where's the *asymmetry*?

Why would majority (or *average*) do gambler's reasoning, if data unbiased?

III. Asymmetric Convergence and Memory Limits

Suppose a group of Bayesians start $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ over (*Swi, Ste, Sti*), and then each observe unbiased data from an *in fact* Steady distribution.

Claim: On average, will get evidence against *Switchy or Sticky* (in favor of *Steady*), but will get *more* evidence against *Sticky* than *Switchy*.

Steady likelihoods are closer to Switchy likelihoods than Sticky ones.

Intuition: when deviate from 50%, Sticky wants to *stay* deviated while Switchy wants to *return* toward 50%.

Can see this in (eg) proportion of hits on 20 tosses (left).
Mean posteriors in Swi/Ste/Sti after length-n sequence (right).

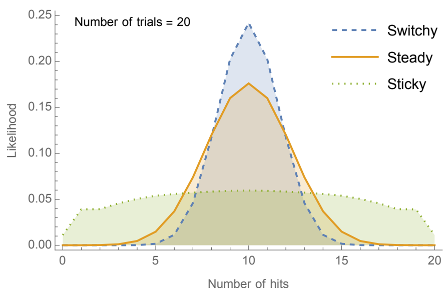

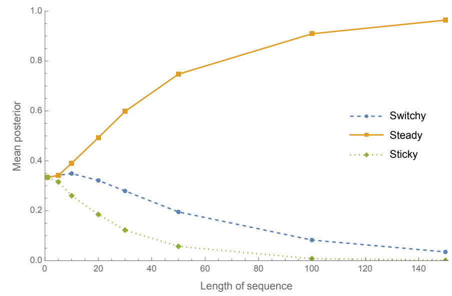

IV. Fallacious Features

1) Nonlinear expectations:

The finding: As streaks grow, people first increasingly exhibit gambler's reasoning—but eventually they exhibit *hot hands* reasoning.

The simulations: Our Bayesians agree. Asymmetric convergence means they start out favoring Switchy. As the streak grows, they know that *if* it's Switchy, it's increasingly primed to switch. But if it keeps going, it eventually provides enough evidence to swamp the asymmetry—they start favoring Sticky over Switchy, so predict continuations.

How? Update on limited data or memory; then present with streaks of varying lengths and look at mean probability of a continuation.

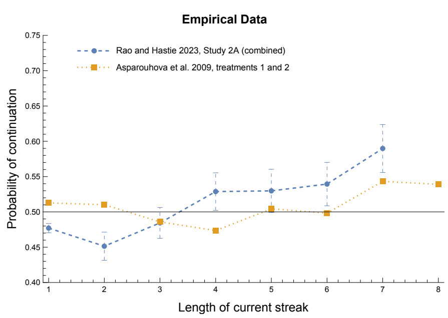

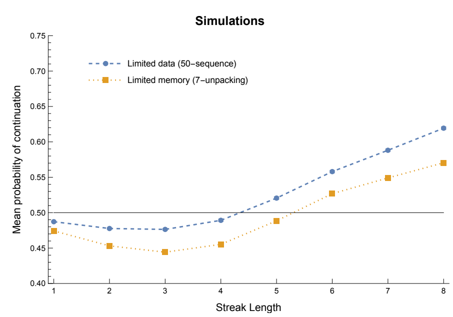

2) Experience-Dependence:

The finding: The degree to which they exhibit gambler's and hot-hands reasoning—and how long of a streak it takes before they transition from the former to the latter—depends on how much experience they have with the causal system.

The simulations: With more data/memory, our Bayesians both (1) are increasingly confident of Steady, and (2) increasingly favor Switchy over Sticky[1]. (1) implies that those with more experience deviate less from 50%; (2) implies that those with more experience will need to see a longer streak before they start favoring Sticky.

[1] $P(Swi|Swi\ or\ Sti) \to 1$

How? As before, but groups have varying amounts of data or memory.

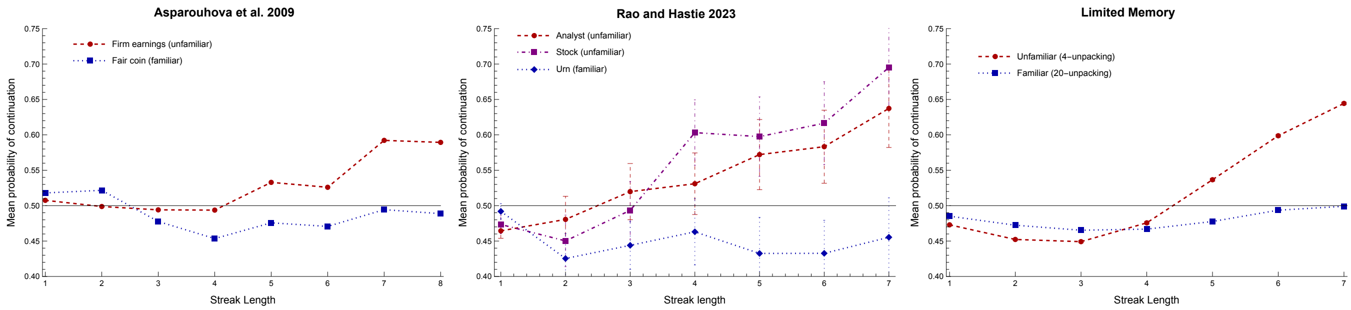

V. Worries

Can this explain other instances of the law of small numbers?
HTHHHH vs HTTHTH? Yes. HTHTHT vs HTTHTH? No.
(4,4,4,4,4) vs (4,4,5,4,3)? No.
Cases with no plausible causal mechanism(?): preserving majority/minority.
Program A: 65% boys, Program B: 45% boys. if 55%, which more likely?

Can this explain other things that the representativeness heuristic can?
Conjunction fallacy? No. Lawyers/engineers? No.