Approaching the Hot Hand with a Cool Head.

A Comment on Xu and Harvey (*Cognition*, 2014)

Barna Bakó *   Máté Csaba Sándor †

Journal of Comments and Replications in Economics, Volume 5, 2026-12, DOI:10.18718/81781.63

JEL: C15; D91
Keywords: Gambling; Risk-taking; Hot hand fallacy; Gambler's fallacy
Data Availability: The R code and the data to reproduce the results in this comment can be downloaded at JCRE's data archive (DOI: 10.15456/j1.2025161.2021817229).

Please cite as: Bakó, B. and Sándor, M. C. (2026). Approaching the Hot Hand with a Cool Head. A Comment on Xu and Harvey (*Cognition*, 2014). *Journal of Comments and Replications in Economics*, Vol. 5 (2026-12). DOI: 10.18718/81781.63

Abstract

Xu and Harvey (2014) argue that gamblers are more likely to win after winning streaks and more likely to lose after losing streaks. They interpret this pattern as evidence that gamblers create their own hot hand by responding to recent outcomes in a manner consistent with the gambler's fallacy: after wins, they choose bets with higher winning probabilities, and after losses, they choose bets with lower winning probabilities. This comment provides conceptual and empirical evidence that the same aggregate patterns can arise without such streak-dependent behavioral adjustment. Using 916,640 bets placed by 10,963 gamblers on the SatoshiDice platform, we show that aggregate streak-conditioned winning probabilities resemble those reported by Xu and Harvey (2014). However, individual-level analyses do not support the proposed mechanism. After accounting for persistent heterogeneity in betting choices, winning streaks are not followed by systematic shifts toward bets with higher winning probabilities. Our results suggest that aggregate streak-conditioned statistics are insufficient to identify hot-hand-type behavioral adjustment in gambling.

*Corresponding author, Institute of Economics, Corvinus University of Budapest, address: Fővám tér 8 225a, 1093 Budapest, Hungary, tel: +3614825165, barna.bako@uni-corvinus.hu

†Institute of Economics, Corvinus University of Budapest, Fővám tér 8, 1093 Budapest, Hungary. mate.sandor2@stud.uni-corvinus.hu

Acknowledgement: We gratefully acknowledge financial support from the Hungarian National Research, Development and Innovation Office (FK-132343 and K-143276). The present publication is the outcome of the project 'From Talent to Young Researcher project aimed at activities supporting the research career model in higher education', identifier EFOP-3.6.3-VEKOP-16-2017-00007 co-supported by the European Union, Hungary and the European Social Fund.

Received June 9, 2025; Revised May 26, 2026; Accepted June 3, 2026; Published August 21, 2026.

1 Introduction

People often perceive patterns in random sequences. One prominent example is the hot hand fallacy: the belief that recent success increases the probability of future success in a process that is, in fact, governed largely by chance. The idea is especially familiar in sports. Basketball fans and commentators often describe a player who has scored repeatedly as being "on fire" and expect that player to score again with unusually high probability. Similar reasoning appears outside sports. A slot-machine player who has just won twice in succession may continue playing during a perceived lucky streak, while an investor may overweight a stock that has recently generated unusually high returns.

A large literature has examined whether hot-hand effects exist and whether observed streaks reflect genuine changes in performance or cognitive misperceptions. The early literature found little support for the phenomenon and emphasized the hot hand as a cognitive illusion (see, for example, Gilovich et al., 1985; Johnson et al., 2005). More recent work has challenged some of these conclusions. In particular, Miller and Sanjurjo (2018) show that the standard empirical approach used to study sequential performance can suffer from streak-selection bias. Correcting this bias may reverse some earlier conclusions, including those drawn from the influential study by Gilovich et al. (1985). Evidence of hot-hand effects in sports may therefore reflect skill, confidence, or performance dynamics rather than pure randomness.[1]

Outside sports, however, evidence for genuine hot-hand effects remains limited. A notable exception is Xu and Harvey (2014), who claim that hot-hand effects are prevalent in gambling. Using a large sports-betting dataset, they report that gamblers are more likely to win after winning streaks and more likely to lose after losing streaks. Their proposed mechanism is behavioral: gamblers on winning streaks gradually choose bets with higher winning probabilities, whereas gamblers on losing streaks choose bets with lower winning probabilities. In this interpretation, gamblers create their own hot hand by acting on the gambler's fallacy.[2]

These findings were challenged by Demaree et al. (2015), who argued that the relationship between streaks and subsequent winning probabilities may arise mechanically from aggregation. Gamblers who typically choose high-probability bets are naturally more likely to experience longer winning streaks, whereas gamblers who typically choose low-probability bets are more likely to experience longer losing streaks. Conditioning on streaks can therefore select gamblers with different baseline betting tendencies even when no individual changes behavior after wins or losses.

In their reply, Xu and Harvey (2015) argued that selection effects could not fully explain their results and presented additional evidence that they interpreted as supporting streak-dependent behavioral change. However, the central issue remains unresolved. Aggregate streak-conditioned

[1] Yaari and Eisenmann (2011) also find statistical evidence of streakiness in sports, while emphasizing that observed patterns may reflect non-random performance dynamics rather than a simple hot-hand mechanism.

[2] The question studied by Xu and Harvey (2014) is only loosely related to the traditional hot-hand concept. Their analysis does not examine whether success directly raises future performance. Instead, it examines whether bettors choose gambles with different payoff probabilities after streaks. In a basketball analogy, this would be closer to saying that a player has a hot hand because he moves closer to the rim after making several shots. Nevertheless, the behavioral question is important: do gamblers alter risk-taking after winning or losing streaks? This is the mechanism examined in this comment.

statistics cannot by themselves distinguish genuine within-person behavioral adjustment from persistent heterogeneity across gamblers.

This comment contributes to the debate by using a dataset that permits individual-level tracking of gambling behavior. We analyze 916,640 bets placed by 10,963 gamblers on SatoshiDice, an online Bitcoin gambling platform. The data allow us to observe the objective winning probability chosen by each gambler on each bet and to reconstruct complete individual betting histories. We can therefore test the behavioral prediction implied by Xu and Harvey (2014) directly: after a winning streak, the same gambler should choose a subsequent bet with a higher winning probability; after a losing streak, the same gambler should choose a subsequent bet with a lower winning probability.

Our analysis proceeds in two steps. First, we apply the aggregate streak-conditioned procedure used by Xu and Harvey (2014) to the SatoshiDice data. This is not intended as a direct replication of their sports-betting study, since the institutional environment is different. Rather, it is an out-of-sample methodological test in a setting where objective winning probabilities are known and individual histories can be reconstructed. We show that the same aggregate empirical pattern emerges: winning probabilities rise after winning streaks and fall after losing streaks.

Second, we use individual-level data to examine the behavioral mechanism directly. Once we distinguish aggregate streak-conditioned patterns from within-person changes in chosen winning probabilities, the mechanism proposed by Xu and Harvey (2014) receives little support. Gamblers frequently persist with the same betting probability, and when they do change, winning streaks are not followed by systematic shifts toward bets with higher winning probabilities. Indeed, our individual-level estimates point in the opposite direction: conditional on a streak, winners tend to move toward bets with lower winning probabilities, while losers tend to move toward bets with higher winning probabilities.

The appendix provides a simple conceptual explanation for why the aggregate pattern can arise mechanically. If gamblers differ persistently in their chosen winning probabilities and do not change behavior after outcomes, those who choose bets with higher winning probabilities are mechanically overrepresented among long winning streaks, while those who choose bets with lower winning probabilities are overrepresented among long losing streaks. Thus, aggregate streak-conditioned winning probabilities can mimic hot-hand-type patterns without any streak-dependent behavioral adjustment.

2 *Streak-dependent behavior at the population level*

This section applies the aggregate streak-conditioned procedure used by Xu and Harvey (2014) to the SatoshiDice dataset. The purpose is not to claim that SatoshiDice is equivalent to the sports-betting environment studied by Xu and Harvey. Instead, SatoshiDice provides an out-of-sample setting in which objective winning probabilities are known and individual betting histories can be reconstructed. This makes it possible to ask whether the aggregate empirical patterns emphasized by Xu and Harvey (2014) can arise even when they do not reflect systematic within-person adjustment.

SatoshiDice was an online gambling site launched in 2012. During the period studied here, it was one of the most prominent Bitcoin-based gambling platforms and accounted for a substantial share of Bitcoin transactions (see Kondor et al., 2014). The game was simple. The website listed a set of games with a house cut of 1.9%, a minimum and maximum wager size, and a Bitcoin address associated with each game. Bets were evaluated using a random number generated from a combination of a secret hash provided by the website and transaction details.

The SatoshiDice environment differs from the sports-betting setting studied by Xu and Harvey (2014). In SatoshiDice, repeatedly choosing the same winning probability was straightforward, while changing the probability required actively switching to another betting address or multiplier. In sports betting, by contrast, betting opportunities expire and bettors must repeatedly choose among new odds. This difference limits direct behavioral comparability. It does not, however, affect the main methodological question: can aggregate streak-conditioned statistics distinguish genuine within-person behavioral adaptation from persistent heterogeneity across gamblers?

Bitcoin's public ledger records the complete transaction history. Because wagers and payments appear in the ledger, a complete history of bets placed on SatoshiDice can be reconstructed. Each bet can be identified as a transaction sent to one of the game's addresses, and the result can be matched to the return transaction. We constructed user-level betting histories following the method in Kondor et al. (2014).[3]

The dataset covers April 12, 2012, to December 28, 2013, the most active period of SatoshiDice.[4] During this period, the Bitcoin–US dollar exchange rate was highly volatile, ranging from about \$5 to more than \$1,000. To reduce the influence of large price movements, we select five three-week windows with relatively low price volatility. Table 1 reports summary statistics.

Table 1: Summary statistics for the five betting-history windows used in the analysis

	Start date	Number of bets	Number of users	Total bets placed (BTC)	Median bet size (BTC)	Mean daily price (USD/BTC)	Relative daily price deviation (%)
A	2012-05-02	119,399	1002	46,649	0.04	5	1.18
B	2012-09-17	129,265	2114	85,280	0.06	12	2.27
C	2012-12-17	252,301	3405	407,140	0.04	13	0.54
D	2013-05-04	329,155	3432	100,430	0.02	115	4.11
E	2013-09-11	86,520	1400	62,920	0.03	125	1.39

Notes: Exchange rates are from the historical data published by BitCoinCharts (www.bitcoincharts.com). The average relative daily price deviation over 21-day periods between April 21, 2012, and April 21, 2014, was 8.8%, substantially higher than in the selected windows.

During the observed period, 27 games were available on SatoshiDice, three of which were added after the first window. The games covered a wide range of winning probabilities, with many options in the high-risk range below 1%. Each game had a minimum and maximum accepted wager, which changed once during the analyzed period. Table 2 reports the distribution of chosen

[3]The scripts used for this study are publicly available at github.com/sampaat/hot_hand_cold_head.

[4]After 2014, SatoshiDice lost popularity because of competition and changes in crypto-gambling models, including a shift toward site-balance-based or prepayment systems. These changes make it harder to reconstruct individual betting decisions with the same method.

games across the five windows. The dispersion in betting choices is substantial and persistent across periods. The consistency of the results across all five windows strengthens the robustness of our conclusions.[5]

Table 2: Observed distribution of risk choices in the five analyzed periods

P_{win}	Multiplier	A	B	C	D	E
0.002%	64000	0.4%	2.2%	2.2%	2.7%	5.2%
0.003%	31982	0.1%	0.1%	0.0%	0.2%	0.2%
0.006%	15991	0.5%	0.1%	0.1%	0.2%	0.1%
0.01%	7995	0.1%	0.1%	0.1%	0.2%	0.1%
0.02%	3998	0.1%	0.1%	0.2%	0.1%	0.2%
0.05%	1999	0.1%	0.3%	0.2%	0.2%	0.3%
0.1%	999.4	0.3%	0.2%	0.2%	0.2%	0.3%
0.2%	499.7	0.4%	0.2%	0.2%	0.2%	0.2%
0.4%	249.8	0.2%	0.4%	0.2%	0.2%	0.2%
0.8%	124.9	0.4%	0.8%	0.9%	0.4%	0.4%
1.5%	63.97	0.4%	3.0%	4.3%	3.4%	6.0%
2.3%	42.64	0.7%	0.7%	0.2%	0.4%	0.3%
3.1%	31.98	0.3%	0.6%	3.0%	2.1%	8.0%
4.6%	21.32	0.3%	0.7%	0.2%	0.4%	0.4%
6.1%	15.99	0.3%	0.6%	0.6%	0.3%	3.0%
9.2%	10.66	0.9%	0.8%	0.8%	1.5%	1.2%
12%	8.000	0.7%	3.8%	5.1%	5.3%	7.4%
18%	5.335	1.4%	0.9%	3.6%	2.9%	6.1%
24%	4.003	2.0%	8.5%	5.7%	6.5%	7.9%
37%	2.670	5.5%	4.8%	7.8%	5.6%	5.3%
49%	2.004	24%	45%	27%	23%	10%
50%	1.957	47%	14%	16%	25%	22%
73%	1.338	12%	7.3%	11%	9.2%	7.4%
79%	1.235	NA	1.6%	3.6%	2.5%	2.0%
85%	1.147	NA	1.6%	1.8%	2.6%	1.2%
92%	1.071	NA	0.7%	4.4%	3.3%	3.6%
98%	1.004	1.3%	0.2%	0.2%	0.3%	0.3%

Following Xu and Harvey (2014), we identify bets preceded by winning streaks of length $n = 1, \ldots, 6$ and compute their population-wide empirical winning frequency. We then compute the corresponding winning probability among bets not preceded by a winning streak of length n.[6] Figure 1A shows that the population-wide empirical winning frequency generally increases with the length of the preceding winning streak. For bets not preceded by such streaks, the empirical winning frequency remains close to the unconditional population-wide winning probability.

We repeat the same exercise for losing streaks. Figure 1B shows the corresponding inverse pattern: the empirical winning frequency decreases as the length of the preceding losing streak increases, while the non-streak comparison group again remains close to the unconditional

[5]A later version of SatoshiDice offered automated betting strategies. The period studied here predates that feature, so the behavior observed in the data is plausibly human behavior rather than an artifact of platform-provided automation.

[6]We use the term population-wide empirical winning frequency to distinguish this aggregate statistic from the objective winning probability ρ associated with a particular game.

population-wide winning probability. These patterns closely resemble the aggregate results reported by Xu and Harvey (2014) and are robust across the five analyzed periods.

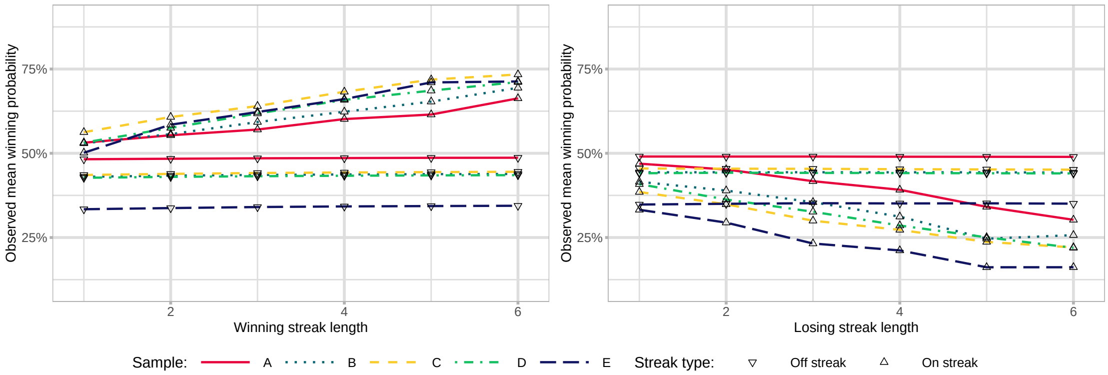

Figure 1: A: Population-wide observed winning probabilities after winning streaks of different lengths and after not obtaining winning streaks of those lengths. B: Population-wide observed winning probabilities after losing streaks of different lengths.

Notes: Colors represent the five observation periods described in Table 1.

We next examine the riskiness of bets chosen after streaks. Instead of odds, we use the implied winning probability, which is bounded between 0 and 100% and allows straightforward comparison across low- and high-risk games. Figure 2 reports the mean implied winning probability chosen after winning and losing streaks of length n. At the aggregate level, mean implied winning probability increases after winning streaks and decreases after losing streaks. Thus, as winning streaks lengthen, the average observed bet becomes safer; as losing streaks lengthen, the average observed bet becomes riskier.

Taken alone, these aggregate results appear consistent with the interpretation in Xu and Harvey (2014). However, the next section shows that the same aggregate pattern can emerge even when individual behavior does not follow the proposed mechanism.

3 *Streak-dependent behavior of individual gamblers*

The aggregate findings in the previous section mirror those of Xu and Harvey (2014). This section uses individual-level data to examine whether the mechanism proposed by Xu and Harvey is actually present. Unlike the aggregate analyses in Xu and Harvey (2014), Demaree et al. (2015), and Xu and Harvey (2015), the SatoshiDice data allow us to observe whether the same gambler changes the winning probability of the next selected bet after a winning or losing streak.

Xu and Harvey (2014) interpret their aggregate results as follows. After a win, gamblers select safer odds; after a loss, they select riskier odds. By doing so, they make themselves more likely to win after winning streaks and more likely to lose after losing streaks. This interpretation has a clear within-person implication: after a winning streak, an individual gambler should increase the winning probability of the next chosen bet; after a losing streak, the same gambler should

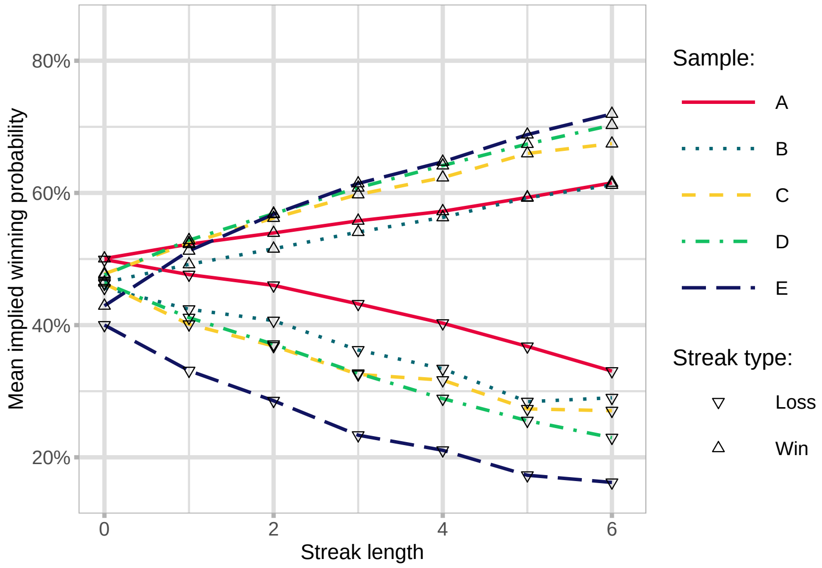

Figure 2: Mean chosen winning probability after winning and losing streaks of length n. Colors represent the five observation periods.

decrease it.

Figure 3 shows that gamblers often do not change their chosen odds at all. After winning streaks, the probability of changing the chosen winning probability is low: gamblers keep the same choice in at least 75% of cases. After losing streaks, gamblers are more likely to change odds, and this tendency rises with streak length. Even then, however, the probability of changing remains below 50% for most streak lengths. Importantly, although only a minority of gamblers actively change the winning probability of subsequent bets, this subgroup alone is approximately three times larger than the entire sample analyzed by Xu and Harvey (2014). Moreover, longer losing streaks are associated with a growing share of gamblers moving toward bets with higher winning probabilities, which is difficult to reconcile with the mechanism proposed by Xu and Harvey (2014).

We then examine the direction and magnitude of changes among observations in which gamblers do change their chosen winning probability. Figure 4 reports the mean relative logarithmic change in the chosen winning probability after winning and losing streaks. Positive values indicate movement toward bets with higher winning probabilities; negative values indicate movement toward bets with lower winning probabilities.

The pattern is again inconsistent with the mechanism proposed by Xu and Harvey (2014). After winning streaks, changes tend to be negative, indicating movement toward bets with lower winning probabilities. After losing streaks, changes tend to be positive, indicating movement toward bets with higher winning probabilities. Although the magnitude varies across periods and streak lengths, the qualitative pattern is the opposite of what the behavioral interpretation in Xu and Harvey (2014) predicts.

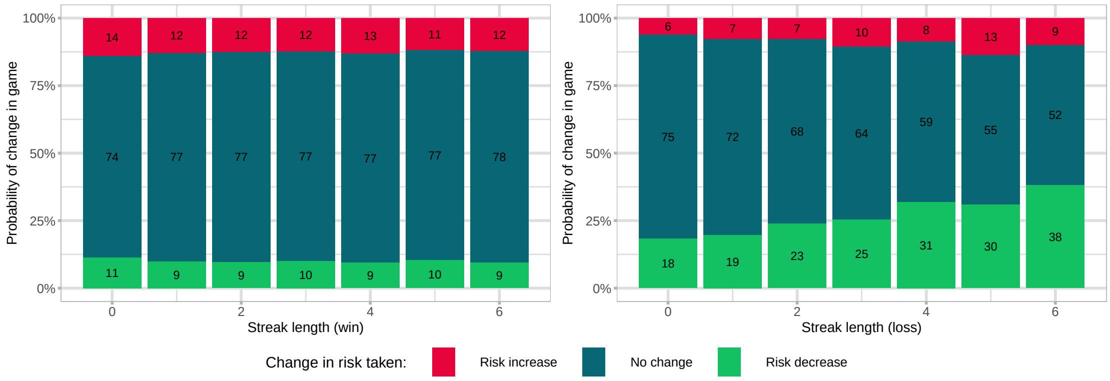

Figure 3: Population-wide probabilities of increasing, decreasing, or holding the chosen winning probability after winning and losing streaks.

Notes: The numbers inside the bars report rounded proportions for each type of risk adjustment by streak length.

The key intuition is straightforward. Gamblers with preference for bets with higher winning probabilities are more likely to generate long winning streaks, while gamblers with preference for bets with lower winning probabilities are more likely to generate long losing streaks. As streak length increases, the streak-conditioned sample becomes increasingly selected. Low-risk gamblers are overrepresented among long winning streaks, and high-risk gamblers are overrepresented among long losing streaks. This produces rising aggregate winning probabilities after winning streaks and falling aggregate winning probabilities after losing streaks, even if individuals do not systematically adjust their betting behavior in the way proposed by Xu and Harvey (2014).

The descriptive evidence therefore suggests that aggregate streak-conditioned patterns may primarily reflect persistent heterogeneity across gamblers rather than within-person behavioral adaptation. However, descriptive comparisons alone cannot fully separate these two mechanisms. To examine whether gamblers systematically adjust the winning probability of subsequent bets after streaks, we next estimate a within-person specification that isolates behavioral changes from time-invariant differences in baseline risk preferences.

3.1 Within-person changes in chosen winning probabilities

The descriptive evidence above suggests that the aggregate patterns documented by Xu and Harvey (2014) may emerge mechanically from persistent differences in gamblers' risk preferences. To examine this more directly, we estimate fixed-effects regressions that exploit within-player variation in subsequent betting choices:

$$\log_2\left(\frac{\rho_{i,t+1}}{\rho_{it}}\right) = \alpha + \beta_w Win5_{it} + \beta_l Lose5_{it} + \mu_i + \varepsilon_{it}. \quad (1)$$

where ρ_{it} denotes the winning probability chosen by player i at time t, while $\rho_{i,t+1}$ denotes the winning probability of the subsequently selected gamble. $Win5_{it}$ and $Lose5_{it}$ are indicator

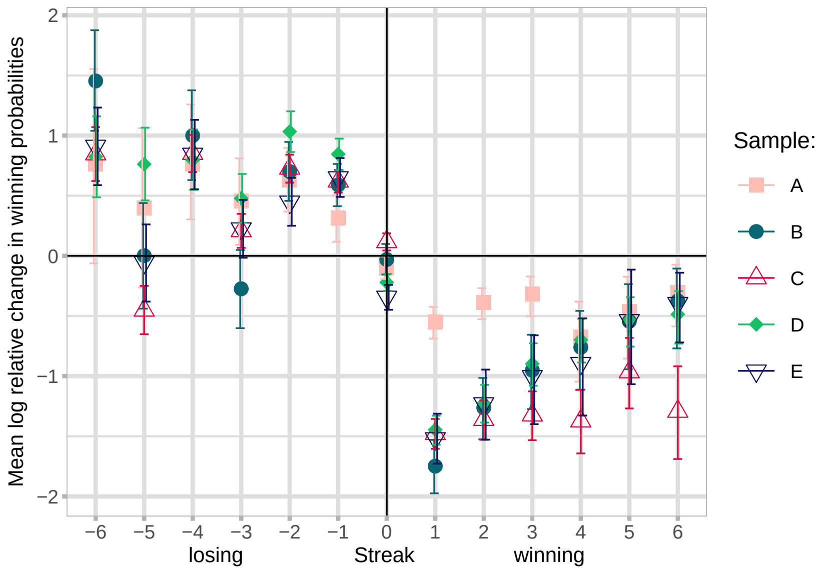

Figure 4: Means of relative logarithmic changes in chosen winning probabilities, $E\left[\log_2(\rho_{i,t+1}/\rho_{it})\right]$, by streak length and observation period.

Notes: Error bars indicate 99% bootstrap confidence intervals. Positive values indicate an increase in the chosen winning probability; negative values indicate a decrease.

variables equal to one if the player is currently experiencing a winning or losing streak of at least five consecutive bets, respectively.[7] Player fixed effects, μ_i, absorb persistent differences in baseline risk preferences. The dependent variable equals zero when the player repeats the same winning probability and differs from zero only when the player switches betting categories.

Table 3: Fixed-effects regressions of within-person changes in chosen winning probabilities

	Coefficient	Clustered SE
$Win5$	-0.131***	(0.025)
$Lose5$	0.263***	(0.028)
Player fixed effects	Yes	
Clustered standard errors	User level	
Observations	785,548	

Notes: Standard errors clustered at the player level are reported in parentheses. *** denotes $p < 0.001$.

The estimates are inconsistent with the behavioral mechanism proposed by Xu and Harvey (2014). After a winning streak of at least five bets, players choose subsequent bets with lower winning probabilities. After a losing streak of at least five bets, they choose subsequent bets with higher winning probabilities. Quantitatively, the coefficient on $Win5$ implies a decrease of approximately 9% in the chosen winning probability, while the coefficient on $Lose5$ implies an increase of approximately 20%. Because the specification includes player fixed effects, these

[7] We focus on streaks of at least five consecutive wins or losses, as Xu and Harvey (2014) argued that any streak-dependent behavioral adjustment should be most pronounced after relatively long winning or losing streaks. Table A1 in Appendix B shows that the results are robust to alternative streak-length thresholds.

estimates reflect within-person changes rather than persistent differences across players in baseline betting tendencies.

We next revisit the selection-bias concern raised by Demaree et al. (2015). In response, Xu and Harvey (2015) used individual-level data to examine whether gamblers with longer winning streaks were generally more cautious. They calculated, for each gambler who reached a given streak length, the average winning probability across that gambler's bets. They then compared these averages across streak lengths and concluded that streak length did not substantially affect betting choices.

This measure is problematic when gamblers differ in how frequently they play. A risky gambler who places many bets and a safer gambler who places few bets receive equal weight in a player-level average, even though they contribute very different numbers of observations to the streak-conditioned data. Moreover, grouping gamblers by the maximum streak length they achieve is highly restrictive. For example, the group with maximum winning streak length zero includes both gamblers who quit after a single loss and gamblers who placed many risky bets without ever winning.

An alternative is to average over individual bets after grouping gamblers by the maximum streak length they achieve. This measure gives each observed bet equal weight and therefore better reflects the composition of the streak-conditioned betting data.[8]

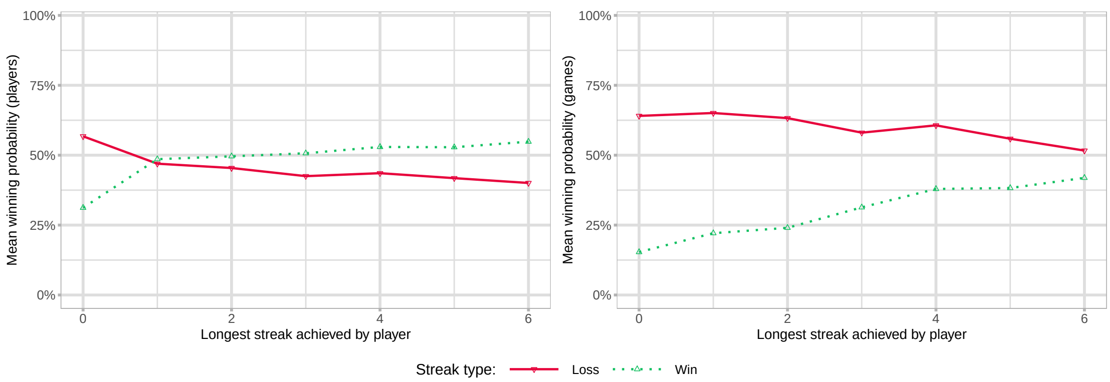

Figure 5: A: Mean winning probability by longest winning or losing streak, using the player-level averaging approach in Xu and Harvey (2015). B: Mean winning probability by longest winning or losing streak, averaging over all bets placed by players with a given maximum streak length.

Figure 5A shows that the player-level averaging approach produces a pattern similar to Xu and Harvey (2015). Figure 5B, however, shows a different pattern once gambling frequencies are taken into account. Players who achieve longer winning streaks tend to place bets with higher winning probabilities, while players who experience longer losing streaks tend to place bets with

[8]Formally, the player-level measure in Figure 5A is $E[\rho(n)]_A = N_{i(n)}^{-1} \sum_{i(n)} N_{j_{i(n)}}^{-1} \sum_{j_{i(n)}} \rho(j_{i(n)})$, while the bet-level measure in Figure 5B is $E[\rho(n)]_B = N_{ij(n)}^{-1} \sum_{i(n)} \sum_{j_{i(n)}} \rho(j_{i(n)})$. Here, $i(n)$ denotes players with maximum streak length n, $j_{i(n)}$ denotes bets placed by player $i(n)$, and N denotes the corresponding number of observations.

lower winning probabilities. This is precisely the selection pattern that can generate aggregate streak-conditioned results without within-person behavioral adjustment.

4 Conclusion

Xu and Harvey (2014) argue that hot-hand effects appear in sports betting because gamblers adjust risk-taking after recent outcomes. According to their proposed mechanism, gamblers on winning streaks choose bets with progressively higher winning probabilities, while gamblers on losing streaks choose bets with progressively lower winning probabilities. This behavior would make subsequent wins more likely after wins and subsequent losses more likely after losses.

This comment shows that such an interpretation is not warranted by aggregate streak-conditioned evidence alone. Applying the same aggregate methodology to SatoshiDice data produces patterns similar to those reported by Xu and Harvey (2014): winning probabilities increase after winning streaks and decrease after losing streaks. However, individual-level analyses do not support the proposed mechanism. Most gamblers often repeat the same betting probability, and when they do change, the direction of change is generally opposite to the one implied by the Xu and Harvey interpretation. Fixed-effects estimates confirm that winning streaks are followed by lower chosen winning probabilities, while losing streaks are followed by higher chosen winning probabilities.

The analysis also clarifies the selection concern raised by Demaree et al. (2015). Gamblers differ persistently in their preferred risk levels. Those who choose bets with higher winning probabilities are mechanically more likely to appear in long winning streaks, while those who choose bets with lower winning probabilities are mechanically more likely to appear in long losing streaks. As a result, aggregate streak-conditioned statistics can produce apparent hot-hand-type patterns even when individuals do not adjust their behavior in the proposed direction.

We do not claim that SatoshiDice is behaviorally identical to the sports-betting environment studied by Xu and Harvey (2014). The institutional settings differ in important ways. SatoshiDice offered objective winning probabilities and allowed gamblers to repeat the same bet immediately, whereas sports bettors must choose among changing opportunities involving events with subjective probabilities. The nature of the data-generating process may matter for how people respond to streaks (see Oskarsson et al., 2009). Some SatoshiDice users may also have been early adopters of Bitcoin, miners, or hobbyists who treated Bitcoin winnings differently from conventional money.

These differences limit direct behavioral generalization. They do not, however, undermine the methodological point. If the aggregate pattern interpreted as a hot-hand effect can arise in a setting where individual behavior is observable and does not support the proposed mechanism, then similar aggregate evidence is insufficient to establish the mechanism in other gambling settings. The existence of hot-hand-type effects in gambling remains an open empirical question, but identifying such effects requires data and methods capable of distinguishing aggregate selection from within-person behavioral adjustment.

Appendix

A Persistent Heterogeneity and the Illusion of a Hot Hand

The main argument of the paper does not depend on a complex theoretical model. The core intuition is simple. Suppose gamblers differ persistently in the winning probabilities they choose, but do not systematically change their behavior after wins or losses. Some gamblers repeatedly choose bets with high winning probabilities, while others repeatedly choose bets with low winning probabilities.

In such a setting, gamblers who repeatedly choose bets with higher winning probabilities are mechanically more likely to experience long winning streaks. Conversely, gamblers who repeatedly choose bets with lower winning probabilities are mechanically more likely to experience long losing streaks. As a result, conditioning observations on streak length gradually changes the composition of the observed sample.

To illustrate the logic, consider a simple example with two types of gamblers:one group repeatedly chooses bets with a 75% winning probability, another group repeatedly chooses bets with a 10% winning probability. Assume that neither group changes behavior after wins or losses.

Even in this simple setting, gamblers with long winning streaks will increasingly consist of players from the first group, while gamblers with long losing streaks will increasingly consist of players from the second group. Consequently, the empirical winning frequency observed after winning streaks rises with streak length, while the empirical winning frequency observed after losing streaks falls with streak length.

Importantly, this pattern emerges even though no gambler changes behavior after outcomes. The result is therefore a mechanical consequence of conditioning on streaks in a heterogeneous population rather than evidence of within-person behavioral adjustment.

To demonstrate this more generally, Figure A1 reports simulation results for three different environments: a continuous uniform distribution of chosen winning probabilities; a binary distribution with 75% and 10% winning probabilities; and the empirical distribution of chosen winning probabilities observed in the SatoshiDice data.

In all three cases, the same qualitative pattern emerges. The empirical winning frequency increases with the length of a winning streak and decreases with the length of a losing streak, even though all players keep their betting behavior fixed over time.

The purpose of Figure A1 is not to draw behavioral conclusions, but to illustrate mechanically how streak-conditioned empirical winning frequencies can emerge under persistent heterogeneity. The figure therefore provides a simple intuition for why aggregate patterns like those reported by Xu and Harvey (2014) do not by themselves establish within-person behavioral adjustment.

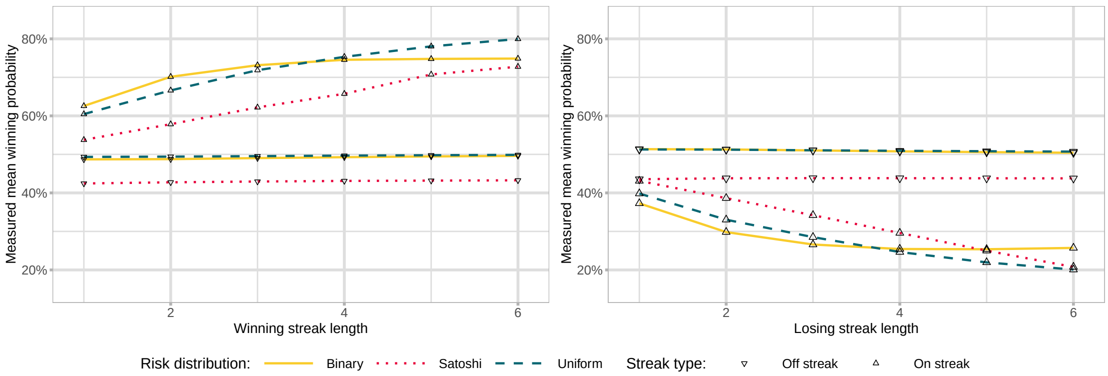

Figure A1: A: Empirical winning frequencies of players with and without winning streaks of different lengths. B: Empirical winning frequencies of players with and without losing streaks of different lengths.

B Robustness to Alternative Streak-Length Thresholds

To examine whether the fixed-effects results depend on the choice of streak-length threshold, we re-estimate Equation (1) using alternative definitions of winning and losing streaks ranging from three to seven consecutive outcomes.

Table A1 reports the results. Across all specifications, winning streaks are associated with statistically significant decreases in the winning probability of subsequently chosen bets, whereas losing streaks are associated with statistically significant increases. The estimated coefficients are stable in sign and similar in magnitude across all threshold definitions, and remain statistically significant at the 0.1% level throughout. The conclusions drawn from the fixed-effects analysis are therefore robust to the choice of streak-length threshold.

Table A1: Robustness of fixed-effects estimates to alternative streak-length thresholds

	Streak-length threshold				
	3	4	5	6	7
Winning streak	-0.198***	-0.165***	-0.131***	-0.115***	-0.093***
Clustered SE	(0.029)	(0.029)	(0.025)	(0.025)	(0.016)
Losing streak	0.281***	0.324***	0.263***	0.308***	0.242***
Clustered SE	(0.029)	(0.032)	(0.028)	(0.028)	(0.024)
Observations	785,548	785,548	785,548	785,548	785,548

Notes: Entries report coefficient estimates from specifications analogous to Equation (1) using alternative definitions of winning and losing streaks. Standard errors clustered at the player level are reported in parentheses. *** denotes $p < 0.001$.

References

Bacidore, J. M., Boquist, J. A., Milbourn, T. T., and Thakor, A. V. (1997). *The Search for the Best Financial Performance Measure*. Financial Analysts Journal, **53**(3), 11–20. DOI: 10.2469/faj.v5

3.n3.2081.

Demaree, H. A., Weaver, J. S., and Juergensen, J. (2015). *A Fallacious "Gambler's Fallacy"? Commentary on Xu and Harvey (2014)*. Cognition, **139**, 168–170. DOI: 10.1016/j.cognition.2014.08.016.

Gilovich, T., Vallone, R., and Tversky, A. (1985). *The Hot Hand in Basketball: On the Misperception of Random Sequences*. Cognitive Psychology, **17**(3), 295–314. DOI: 10.1016/0010-0285(85)90010-6.

Kondor, D., Pósfai, M., Csabai, I., and Vattay, G. (2014). *Do the Rich Get Richer? An Empirical Analysis of the Bitcoin Transaction Network*. PLOS ONE, **9**(2), e86197. DOI: 10.1371/journal.pone.0086197.

Johnson, J., Tellis, G. J., and MacInnis, D. J. (2005). *Losers, Winners, and Biased Trades*. Journal of Consumer Research, **32**(2), 324–329. DOI: 10.1086/432241.

Miller, J. B., and Sanjurjo, A. (2018). *Surprised by the Hot Hand Fallacy? A Truth in the Law of Small Numbers*. Econometrica, **86**(6), 2019–2047. DOI: 10.3982/ECTA14943.

Oskarsson, A. T., Van Boven, L., McClelland, G. H., and Hastie, R. (2009). *What's Next? Judging Sequences of Binary Events*. Psychological Bulletin, **135**(2), 262–285. DOI: 10.1037/a0014821.

Xu, J., and Harvey, N. (2014). *Carry on Winning: The Gamblers' Fallacy Creates Hot Hand Effects in Online Gambling*. Cognition, **131**(2), 173–180. DOI: 10.1016/j.cognition.2014.01.002.

Xu, J., and Harvey, N. (2015). *Carry on Winning: No Selection Effect*. Cognition, **139**, 171–173. DOI: 10.1016/j.cognition.2015.02.008.

Yaari, G., and Eisenmann, S. (2011). *The Hot (Invisible?) Hand: Can Time Sequence Patterns of Success/Failure in Sports Be Modeled as Repeated Random Independent Trials?* PLOS ONE, **6**(10), e24532. DOI: 10.1371/journal.pone.0024532.

Wardrop, R. L. (1995). *Simpson's Paradox and the Hot Hand in Basketball*. The American Statistician, **49**(1), 24–28. DOI: 10.1080/00031305.1995.10476107.

Croson, R., and Sundali, J. (2005). *The Gambler's Fallacy and the Hot Hand: Empirical Data from Casinos*. Journal of Risk and Uncertainty, **30**(3), 195–209. DOI: 10.1007/s11166-005-1153-2.

Kukavica, A. J., and Narayanan, S. (2025). *Rational and Irrational Belief in the Hot Hand: Evidence from "Jeopardy!"*. Stanford University Graduate School of Business, Working Paper No. 4229. DOI: 10.2139/ssrn.5062536.