The Entropy Method: How to Detect the Hot Hand and Other Dependencies in Data

Ben Johnson-Laird and Phil Johnson-Laird

Athletes are familiar with the 'hot hand' in which they have a streak of performing better than normal. But, a trio of psychologists – Gilovich, Vallone, and Tversky – published an analysis in 1985 showing that two NBA basketball teams seemed never to have hot hands in shooting either field goals or free throws. Later, Tversky and Gilovich debunked the phenomenon in a further article in this journal. Basketball players were not the only ones to be dumbfounded.

Analysts have used a variety of statistical techniques to search for hot hands, including conditional probabilities and correlations. A few years ago, *Putting the 'Hot Hand' on Ice* in this journal used linear regression to take into account various aspects of ice-hockey games, such as the quality of the goalie, in order to show that players in the NHL were less likely to score after having just scored. And a recent study, *Surprised by the hot hand fallacy? A truth in the law of small numbers*, showed that for any finite sequence of hits and misses generated independently, the probability of a hit following a streak is less than the underlying conditional probability of a hit. This proof undermines many of the debunking studies, including the original one, and so the hot hand may be real.

Our aim was to formulate a general analytical procedure for establishing whether or not certain results have a reliable bias equivalent to a hot hand. They could be goals that soccer players score in games, the different sorts of observation that individuals select to test a hypothesis, or the changes in the prices of stocks.

By definition, a hot hand is exceptional. It is a departure from the usual, and therefore harder to predict than a typical performance. Correlations to assess it are dubious – a significant correlation can occur in data that can be split into two subsets in which

neither yields a reliable correlation. Conditional probabilities led one group of investigators astray. And the use of regression models is impossible in the absence of data about relevant factors. A new approach seemed timely.

The Entropy Method

Suppose that the goals (runs, hits, or whatever) that a player makes in a minority of games seem to occur surprisingly often. They could be just chance fluctuations. Yet, some players in the English Premier soccer league showed potential signs of a hot boot. There are 20 teams in this league, and they play each other twice in a season, once at home and once away in a total of 38 games for each team. To search for hot boots, we examined the records for all the players who scored at least 10 goals in the 2018-19 season—fewer than 10 goals was not enough for a hot boot to be likely. There were 32 higher scoring players. One of them, Sergio Aguero of Manchester City, had an extraordinary record: 21 goals in the 33 games he played, which included three hat-tricks: games in which he scored three goals. We suspected that his boot was hot on these occasions, but, is there a way to transform such intuitions into a proper test?

There is: the method of entropy. We first developed it to analyze the sequence of selections of potential evidence that participants in psychological experiments made to test general hypotheses, and we come back to this study below. We adapted the method to examine soccer, and we report this new application to detecting the hot boot.

Entropy is disorder, and therefore unpredictability. It can be measured in bits using Shannon's H , and we present the formula in a moment: the greater its value, the greater the entropy. The 32 goal-scorers in the Premier league played at most four categories of game in the season: games in which they scored zero, one, two, or three goals—no-one scored more than a hat-trick in a game. For each player, the value of H depends on the number of categories of game he played in the season, and their

respective frequencies. With four categories, the maximum value of H equals two bits, and with two categories, the maximum value of H equals one bit. We need to take this difference into account. But, if the distribution of frequencies of a player's different categories of game is biased towards him scoring no more than one goal per game, then H is reduced, and the prediction that he will make such a score in a given game is likely to be correct. H is an additive measure, and so its sum over the four categories of game in a season yields its overall value for the season. We can now describe the main steps in applying the entropy method to scoring in soccer.

The first step is to measure the entropy (H) for a player's actual season based on the probabilities of the different categories of game in his season, i.e., games of 0, 1, 2, or 3 goals. The measure is, of course, the same regardless of the sequence of games in the season: it sums the entropy for each of the categories of game to obtain its value for the season. The second step relies on a Monte Carlo procedure: this distribution of probabilities is used to simulate a large number of seasons, say, ten thousand. The third step is to compute the difference between the actual season's entropy and its mean for the 10,000 simulations. When the actual season's entropy is smaller than the mean of the simulations, the goal scorer tends to play predictable games. When the actual season's entropy is greater than the mean of the simulations, the goal scorer tends to play less predictable games—and it is in his season that instances of the hot boot should occur.

These three steps are repeated for all the goal-scorers, and in theory they can divide the players into those with predictable seasons and those with unpredictable seasons. Given a significant difference between these two sets of players, we can examine those with unpredictable seasons for robust signs of a hot boot. But, given that the difference is not significant, the quest for the hot boot has failed.

The Computation of Normalized Entropy (H)

To compute entropy, we start with a player's number of goals in each game he plays. The highest goal-scorer in the Premier league's 2018-19 season, Mohamed Salah of Liverpool, scored 22 goals in his 38 games. He scored one goal in his first game, no goals in his second game, and so on. In his season, he had:

20 games of 0 goals, 15 games of 1 goal, 2 games of 2 goals, and 1 game of 3 goals.

So, his probability of a one-goal game is $P_1 = 15/38 = .395$. The entropy, H , of a one-goal game equals the product of this probability, P_1 , with its logarithm to base 2:

$$P_1 \log_2 P_1$$

The $\log_2$ of $.395 = -1.341$, and its product with P_1 equals $-.53$, which is the entropy H of Salah's one-goal games. H is an additive measure, and so the entropy of Salah's season is the sum of the values of H for the four different categories of game in his season (from $n = 1$ to 4). The general formula for H is therefore:

$$H = - \sum_{i=1}^n P_i \log_2 P_i$$

The minus sign in the formula eliminates the negative values of logarithms of numbers less than 1. So, Salah's season has the following entropy:

Score of 0 goals in 20/38 games: $P_1 = .526$. $\log_2 P_1 = -0.927$. $P_1 \log_2 P_1 = -.488$

Score of 1 goal in 15/38 games: $P_2 = .395$. $\log_2 P_2 = -1.340$. $P_2 \log_2 P_2 = -.529$

Score of 2 goals in 2/38 games: $P_3 = .053$. $\log_2 P_3 = -4.238$. $P_3 \log_2 P_3 = -.225$

Score of 3 goals in 1/38 game: $P_4 = .026$. $\log_2 P_4 = -5.265$. $P_4 \log_2 P_4 = -.137$

Minus the Sum of $P_i \log_2 P_i = 1.379$

The entropy, H , for Salah's season is thus: 1.379 bits. Does this value show that goals in his games tend to be predictable or unpredictable? The second step tells us.

Monte Carlo Simulations as Resampling Tests of Entropy

The way to find out whether Salah's season tends to have predictable games is to use his actual probabilities of 0, 1, 2, or 3 goal-games to simulate a season of games, and to make a large number of such simulations, say, 10,000. However, it's essential to adjust all the entropies, including the one for his actual season, in order to avoid an artificial bias, and to compare different players. The following example explains why.

Suppose we simulate 10,000 seasons based on a Poisson-like distribution of probabilities for the four different categories of game of 0, 1, 2, and 3 goals. We chose at random one of these simulated seasons containing the four sorts of game, and we treat it as though it were an actual season. Its maximum possible entropy is 2.0 bits. But, even though its generation depended on the same probabilities as did the remaining 9,999 remaining simulations, it is bound to have a higher entropy, H , than their mean entropy. The bias occurs because some of the simulated seasons will have only three categories of game with a maximum possible entropy of only 1.585 bits, and some of the simulated seasons will have only two categories of game with a maximum possible entropy of 1.0 bit. These seasons with fewer categories of game than the actual season reduce the mean entropy for the simulated seasons, so they are bound to be more predictable than the one season that we singled out to treat as an actual season. The frequencies of Salah's actual four categories of game, which were not reliably different from a Poisson distribution, yielded 46% simulated seasons of all four categories of game as in his actual season, 48% simulated seasons of only three categories of game (no three-goal games), and 6% of simulated seasons of only two categories of game (no three- or two-goal games). So, if no correction is made then there is a misleading bias towards simulations yielding unrealistic low values of H .

The same problem occurs in a comparison of different players. We want to determine whether there is a significant division among them between those with predictable seasons and those with unpredictable seasons. But, we have to eliminate a bias resulting from the number of categories of game that different players had in their actual seasons, e.g., Salah had games of all four categories, whereas Aguero had games of only three categories (see Table 1 below). Both sorts of problem are familiar

in analyses of entropy, and they have a standard answer: *normalize* the value of H. We must normalize H for a player's actual season, and for each of its simulations. It eliminates the bias from the simulations, and it ensures that all the players are assessed on the same scale no matter how many categories of game occurred in their actual seasons. As we say, this procedure levels the playing field.

Normalization is straightforward: divide an observed H by its maximum possible value. So, the formula for normalization is:

$$\text{Normalized-H} = \frac{H}{\text{Maximum possible H}}$$

The maximum possible H in soccer depends on the number of categories of game in a season, and therefore equals:

$$-\log_2 P(1/\text{number of categories of game in a season})$$

For Salah's actual season, H equals 1.379, the number of categories of game in his season was 4, and so the maximum possible H equals 2.0. Hence, the normalized-H for his season equals 1.379/2.0, which equals .689. Any normalized-H ranges in value from 0 to 1 no matter how many categories of game occurred in a season. If most of the simulations of Salah's season have a greater normalized entropy (H) than his actual season, his scoring is more predictable in his real season than in his simulated seasons.

The Monte Carlo procedure is straightforward. Turn Salah's probabilities of his four categories of game (see their values above) into cumulative intervals on the unit line [0, 1]. The boundaries are as follows:

Then generate a pseudo-random number from 0 to 1. Its value yields the number of goals in a particular game from 0 to 3 according to the unit line above, e.g., if the value of the random number is .530, then he scored 1 goal in the game. Repeat this procedure 38 times, and the result is a simulation of a single season given that the number of games in each category depends only on their cumulative probabilities in the actual season. Simulate 10,000 seasons in this way. The main datum for each player is the reduction, if any, in the normalized entropy of his actual season in comparison with the mean for his simulated seasons. The greater this reduction, the more predictable were the goals in his games. For Salah, there was a reduction in normalized entropy from the mean of his simulated seasons (.744 bits) to his actual season (.689 bits). So, his actual season tended to be predictable. For Aguero, there was no reduction but an increase in normalized entropy from the mean of his simulated seasons (.820 bits) to his actual season (0.834 bits). So, his actual season tended to be unpredictable.

The predictability of goal-scoring in the English Premier League

Table 1 shows the performance of the 32 top goal-scorers in the English soccer Premier league for the 2018-19 season. It presents (i) the numbers of their games in the four categories (0, 1, 2, or 3 goals in a game), (ii) a statistic showing the distance of these frequencies from a Poisson distribution (the Cramér-von Mises statistic) and whether it is statistically significant, (iii) the normalized entropy, H , for the actual seasons and (iv) for their means of 10,000 simulations, and (v) the difference between these two values. The size of the Cramér-von Mises statistic correlated slightly with the degree to which the players' actual seasons were predictable, i.e., they had a smaller H than the mean of their simulations (Kendall's rank-order correlation, $\tau = 0.27$, $z = 2.17$, $P < 0.025$, two-tailed). Figure 2 is a histogram showing the relative predictability of the 32 players' scoring: the reduction of the normalized entropy of their actual seasons in comparison with the mean of their 10,000 simulations. No reliable

correlation occurred between the total number of goals a player scored and the tendency to have a predictable season (Kendall's, $\tau = -.18$, $P > .9$).

As Table 1 and Figure 1 illustrate, 23 out of the 32 players' had actual seasons that were more predictable than their simulations (a mean difference of 0.048), and 9

Player's name and club	Numbers of 0, 1, 2, & 3 goal games in actual season	Cramér-von Mises statistic	Normalized entropy (H)	Mean normalized entropy (H) in 10K simulations	Reduction in entropy of actual season
1. A. Lacazette, <i>Arsenal</i>	23 11 1 -	0.36	0.674	0.766	0.091
2. R. Jimenez, <i>Wolves</i>	26 11 1 -	0.38	0.650	0.741	0.091
3. A. Barnes, <i>Burnley</i>	26 10 1 -	0.40	0.636	0.721	0.085
4. L. Milivojevic, <i>Crystal Palace</i>	27 10 1 -	0.41	0.628	0.713	0.085
5. S. Rondon, <i>Newcastle</i>	22 9 1 -	0.38	0.658	0.742	0.084
6. A. Martial, <i>Manchester United</i>	18 8 1 -	0.36	0.685	0.763	0.078
7. L. Sane, <i>Manchester City</i>	22 8 1 -	0.40	0.641	0.718	0.077
8. W. Zaha, <i>Crystal Palace</i>	25 8 1 -	0.43	0.610	0.686	0.076
9. R. Firmino, <i>Liverpool</i>	25 7 1 1	0.50*	0.547	0.618	0.071
10. A. Perez, <i>Newcastle</i>	28 7 1 1	0.53*	0.520	0.589	0.069
11. G. Deulofeu, <i>Watford</i>	23 5 1 1	0.52*	0.526	0.583	0.057
12. M. Salah, <i>Liverpool</i>	20 15 2 1	0.35	0.689	0.744	0.055
13. L. Moura, <i>Tottenham</i>	25 5 1 1	0.55*	0.505	0.559	0.054
14. R. Sterling, <i>Manchester City</i>	22 8 3 1	0.38	0.678	0.706	0.028
15. E. Hazard, <i>Chelsea</i>	26 7 3 1	0.42	0.623	0.649	0.026
16. G. Murray, <i>Brighton</i>	27 9 2 -	0.38	0.673	0.686	0.013
17. G. Sigurdsson, <i>Everton</i>	27 9 2 -	0.38	0.673	0.686	0.013
18. H. Kane, <i>Tottenham</i>	13 13 2 -	0.26	0.820	0.832	0.012
19. Richarlison, <i>Everton</i>	24 9 2 -	0.36	0.702	0.713	0.011
20. C. Wilson, <i>Bournemouth</i>	18 10 2 -	0.39	0.777	0.786	0.009
21. Son Heung-min, <i>Tottenham</i>	21 8 2 -	0.35	0.719	0.727	0.008
22. C. Wood, <i>Burnley</i>	30 6 2 -	0.47*	0.576	0.580	0.004
23. M. Arnautovic, <i>West Ham</i>	20 6 2 -	0.37	0.691	0.692	0.001
24. J. Vardy, <i>Leicester</i>	19 12 3 -	0.27	0.826	0.813	-0.013
25. S. Aguero, <i>Manchester City</i>	18 12 0 3	0.34	0.834	0.820	-0.014
26. A. Mitrovic, <i>Fulham</i>	29 5 3 -	0.43	0.605	0.589	-0.016

27. J. King, <i>Bournemouth</i>	26	6	3	-	0.38	0.668	0.651	-0.017
28. R. Lukaku, <i>Manchester United</i>	23	6	3	-	0.36	0.704	0.685	-0.019
29. M. Rashford, <i>Manchester United</i>	23	10	0	-	0.28	0.885	0.864	-0.021
30. P. Pogba, <i>Manchester United</i>	26	5	4	-	0.36	0.680	0.658	-0.022
31. S. Mane, <i>Liverpool</i>	20	10	6	-	0.26	0.893	0.868	-0.025
32. P. Aubameyang, <i>Arsenal</i>	19	12	5	-	0.26	0.890	0.865	-0.025

Table 1. The 32 best goal-scorers in the 2018-19 season of the English Premier soccer league, the numbers of occurrences of their four categories of game (0, 1, 2, and 3 goals), the Cramér–von Mises (omega-squared) measure of the difference between this distribution and a Poisson distribution (* indicates a reliable difference with $P < .05$), the normalized H of their season, its mean over 10,000 simulations of their season based on their distributions of the four categories of game in the actual seasons, and the reduction in entropy of their actual season in comparison with its simulations, where a negative number indicates that a player's actual season had a greater entropy than its simulations, and so it was more unpredictable.

players had actual seasons that were less predictable than their simulations (a mean of -0.019). This difference was highly significant (Mann-Whitney test, $W = 45.0$, $z > 4.3$, $P < .0000025$; two tailed). Could it be a consequence of the caliber of the players or their opposing teams? An inspection of the players' clubs in Table 1 refutes the

Figure 1. The reduction, if any, in the normalized entropy of the actual seasons of the 32 best goal-scorers in the 2018-19 Premier league season in comparison with 10,000 simulations for each player based on the frequencies of their different categories of game (those in which they score 0, 1, 2, or 3 goals). The greater the reduction in H , the more predictable is their season; the players with a negative difference had more unpredictable seasons than those with a positive difference.

possibility. The 23 scorers with predictable seasons played for 15 out of the 20 teams in the Premier league, and 7 out of the 9 scorers with unpredictable seasons were teammates of theirs. The two groups of players did not differ reliably in their teams' overall success in the season as a whole. Teams in the Premier league score 3 points for a win, 1 point for a tie, and 0 points for a loss, and the predictable group played for teams with a mean of 62.0 points for the season, and the unpredictable group played for teams with a mean of 65.1 for the season (Mann-Whitney test $W = 156.0$, $z < .3$, $P > .38$).

In general, goal-scorers differed in the predictability of their seasons. This difference reflected neither their skill in scoring goals nor the caliber of their teams in comparison with their opponents. So, what caused it? The answer is the occurrence of the hot boot, that is, the occurrence of games in which a player scores more than one goal.

The proportion of such games was reliably greater for the 9 players with unpredictable seasons (44%) than for the 23 players with predictable seasons (32%; Mann-Whitney test $W = 92.5$, $z > 2.34$, $p < 0.01$). It is not attributable to departures from a Poisson distribution, which tended to be smaller for unpredictable seasons, or to the caliber of the players' teams or opponents. As we had wondered, one player in the first group was Aguero, who scored three hat-tricks—a real sign of a hot boot—and just over half his goals in one-goal games. One player in the second group was Salah who scored more than two thirds of his goals in one-goal games. In sum, the hot boot did occur for some soccer players in the 2018-19 Premier League season.

Do People Seek Independent Tests of Hypotheses?

The entropy method has widespread applications. It can be used to assess whether sets of data are internally dependent or not. We illustrate an application remote from sports for which Ragni, Kola, and Johnson-Laird first devised the method. In a well-known psychological task, participants have to select potential evidence to test a given hypothesis. The late Peter Wason, who devised the task, aimed to investigate scientific thinking and to examine whether people who are not scientists grasp the relevance of falsification. In his original test, the hypothesis was an abstract general proposition about letters and numbers:

If a card has an E on one side then it has a 4 on the other side.

The participants knew that every card had a letter on one side and a number on the other side. Their task was to make a parsimonious selection of cards that they needed to turn over to find out whether the hypothesis above was true or false about the four cards:

Most people selected the E and 4 cards, some selected E alone, and a few selected E, 4, and 7. What was remarkable was the neglect of 7. Yet, if it has an E on its other side, then the hypothesis is false. This neglect of evidence that could falsify a hypothesis sparked a surge in experiments and theories. At least 16 different theories of the task existed at one time, and the equivalent of the population of a small town – over 14 thousand people – was tested in over 200 experiments. These studies also examined sensible everyday hypotheses, and deontic principles governing permissions. Both these sorts of content elicited many more potentially falsifying selections than Wason's abstract task.

One controversial issue was whether people select each piece of evidence independently from the others, i.e., they decide whether or not to select each of the four cards, E, K, 4, and 7, independently from their other selections. Over half the theories of people's performance implied that they do. To test independence, some investigators examined correlations between the selections of different pairs of cards, but the results were mixed.

The entropy method answered the question. With four pieces of potential evidence, e.g., the cards above, there are 16 possible selections, including a selection of all four cards or of none at all. In fact, the entropy, H , for the 99 experiments that reported the frequencies of all 16 selections was often much smaller than the maximum of 4 bits. Suppose that in an on-line experiment, 244 participants made the following frequencies of selections:

E & 4: 116

E: 62

E & 7: 41

E, 4 & 7: 21

K: 4

The individual probabilities of selecting each of the four cards are: $P(E) = .984$, $P(K) = .016$, $P(4) = .561$, $P(7) = .254$. To simulate independent selections, a program makes each individual selection by deciding whether or not to select each card according to the probabilities in the actual experiment, and therefore it make independent selections. It repeats the process 244 times, once for each participant in the experiment. Next, it simulates the experiment 10,000 times in this way. The entropy, H , of the actual experiment above is: 1.846 bits. Its mean value for the simulations is 2.00, and in 9,765 of them it is larger than that of the actual experiment. It follows with a probability of $P < .025$ that the selections in the actual experiment are dependent on one another, because their entropy is smaller than that of their independent simulations.

In the 99 experiments reporting their complete results for the task, the overall differences between their mean entropies and those of their simulations were highly significant: 2.33 vs. 2.60 for abstract hypotheses such as the one above, 2.43 vs. 2.75 for sensible everyday hypotheses, and 1.57 vs. 1.80 for rules about permissions ($P < .0001$ in all three comparisons). As the smaller entropy in actual performance showed, the participants made dependent selections, because they tended, for instance, not to select the 7 card in the abstract task unless they selected the E card. The entropy method eliminated nine theories of the task.

One unexpected finding was a difference between those experiments carried out face-to-face and those carried out in on-line experiments. The face-to-face experiments had a smaller entropy than that of the on-line experiments. And this difference was highly reliable. So, the selection of evidence in the face-to-face studies was more

predictable than in the on-line studies. Participants tested on line are likely to care less and to pay less attention. Indeed, unlike those in the face-to-face studies, some of them selected no evidence whatsoever. Comparisons between the two ways of carrying out experiments have seldom reported any systematic differences between them. The absence of such reports may be because investigators haven't examined the entropies of the two sets of data.

Conclusion

If you need to test whether or not there's a dependency between events, then the method of entropy, which combines Shannon's measure of entropy, H , with Monte Carlo re-sampling is helpful. It led to the discovery that the hot boot occurs in soccer for some players but not others, that people select potential evidence in dependent ways, and that experiments carried out face-to-face yield less noise than their on-line counterparts.

Further Reading

Gilovich, T., Vallone, R., and Tversky, A. 1985. The hot hand in basketball: On the misperception of random sequences. *Cognitive Psychology*, 17, 295-314.

Miller, J. B., and Sanjurjo, A. 2018. Surprised by the hot hand fallacy? A truth in the law of small numbers. *Econometrica*, 86, 2019-2047.

Ragni, M., Kola, I., and Johnson-Laird, P. N. 2018. On selecting evidence to test hypotheses. *Psychological Bulletin*, 144, 779-796.

Shannon, C. E. 1948. A mathematical theory of communication. *The Bell System Technical Journal*, 27, 379-423.

Tversky, A., & Gilovich, T. (1989). The 'hot hand': Statistical reality or cognitive illusion?. *CHANCE*, 2, 31-34.

Vesper, A. (2015). Putting the 'Hot Hand' on ice. *CHANCE*, 28, 13-18.

Premier League data are from <https://fbref.com>. Its use is allowed from its parent website: https://www.sports-reference.com/data_use.html It states: "As an aside, copyright law is clear that facts cannot be copyrighted, so you are free to reuse facts found on this site in accordance with copyright laws." We thank Peter Ayton, Sangeet Khemlani, and two reviewers, for helpful advice and criticisms.

Appendix

About the authors

Ben Johnson-Laird is a software engineer and data scientist. He received his BSc (Hons) in Computer Science from the University of Kent, and his MPhil in Computer Speech and Language Processing from the University of Cambridge. His chief extra-curricular interest is in listening to vinyl records of classical music and modern jazz.

Phil Johnson-Laird is Stuart Professor of Psychology, Emeritus, Princeton University, and a Visiting Scholar in the Psychology Department of New York University. He is a member of the National Academy of Sciences and the American Philosophical Society, and a fellow of the Royal Society (of London) and the British Academy. His main research is into the role of mental models in reasoning. He too listens to classical music and modern jazz, but not on vinyl.