

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Subjective Randomness in a Non-cooperative Game

Permalink

<https://escholarship.org/uc/item/8hw3k32v>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 41(0)

Authors

Payton, Michael P.

Zemla, Jeffrey C.

Austerweil, Joseph L.

Publication Date

2019

Peer reviewed

Subjective Randomness in a Non-cooperative Game

Michael Payton (mjpayton@wisc.edu)

Jeffrey C. Zemla (zemla@wisc.edu)

Joseph L. Austerweil (austerweil@wisc.edu)

Department of Psychology, University of Wisconsin Madison
1202 W. Johnson Street
Madison, WI 53706 USA

Abstract

Rock, Paper, Scissors (RPS) is a competitive game. There are three actions: rock, paper, and scissors. The game's rules are simple: scissors beats paper, rock beats scissors and paper beats rock (all signs stalemate against themselves). Over multiple games with the same opponent, optimal play according to a Nash Equilibrium requires subjects to play with genuine randomness. To examine randomness judgments in the context of competition, we tested subjects with identical sequences in two conditions: one produced from a dice roll, one from someone playing rock, paper, scissors. We compared these findings to models of subjective randomness from Falk and Konold (1997) and from Griffiths and Tenenbaum (2001), which explain assessments of randomness as a function of algorithmic complexity and statistical inference, respectively. In both conditions the models fail to adequately describe subjective randomness judgements of ternary outcomes. We also observe that context influences perceptions of randomness such that some isomorphic sequences produced from intentional play are perceived as less random than dice rolls. We discuss this finding in terms of the relation between patterns and opponent modeling.

Keywords: Randomness, pattern recognition, opponent modeling

Introduction

Humans often detect patterns in everyday life—so much so that they often attribute spurious patterns (generated by a random mechanism) to intentional actions. For example, in the *gambler's fallacy* (Kahneman & Tversky, 1972), people believe that “a Red is due” after observing several Black rolls from a roulette wheel, despite the rolls being independent from one another. Conversely, according to the *hot hand effect*, if a player has scored several free throws in a row in basketball practice, people believe that player is more likely to score again, despite this not being true empirically (Gilovich, Vallone, & Tversky, 1985). Why are the patterns that people detect sensitive to their context?

Psychologists have approached these phenomena in terms of *subjective randomness*, or the perceived randomness of observations. Previous literature has shown that in some circumstances, people tend to judge sequences as random even when the underlying pattern is systematic. In a classic example, people tend to believe that a sequence of coin flips will have more alternations (e.g., heads followed by tails) and fewer streaks (e.g., several heads in a row) than is likely to occur in a sequence produced by flipping an unbiased coin repeatedly (Falk & Konold, 1997). To date, much of the literature on subjective randomness has focused on binary

sequences or grids generated from a truly random mechanism, such as a coin flip, an animate mechanism (Ayton & Fischer, 2004), or a human or other intentional agent (Burns & Corpus, 2004; Caruso, Waytz, & Epley, 2010).

In this article, we compare the subjective randomness of sequences generated from a die roll to comparable sequences generated by a player who is in direct competition with another player in a non-cooperative game: Rock, Paper, Scissors (RPS; also called RoShamBo). We had two hypotheses: (1) sequences generated from a random mechanism (a die roll) would be judged as more random than equivalent sequences generated from a human playing RPS, and (2) “complex” sequences generated from a human playing RPS would be perceived even less random due to opponent modeling in a competitive context.

The rules of RPS are straightforward. Two players simultaneously present one of three hand signs, “rock”, “paper”, or “scissors”. The scoring of the game is also simple: scissors beats paper, rock beats scissors, and paper beats rock (all signs stalemate against themselves). From a game-theoretic perspective, the Nash equilibrium (Nash, 1950) is generating signs uniformly at random. Thus, if people played according to the Nash equilibrium, they would be required to produce truly random sequences. This strategy would prevent any player from gaining advantage over another after playing repeated games over time. However, this is unlikely as people are poor at producing random sequences (Baddeley, 1966, Towse, 1998).

Although RPS may appear to be a simple game, it is actually much more complex than one might first think. For instance, while one might expect that the winners of RPS are determined by luck or chance, there are genuine RPS masters. RPS tournaments have been held throughout the world where experienced RPS players will consistently outpace novices (Hegen, 2004). One might at first think that this is simply due to extraneous factors. For example, perhaps one player produces their sign slightly before the other and the other player uses that information to change their play (note that this is illegal in tournament play). However, it is reported that a player in 2001 was allowed to bring a random number generator to inform his sequence generation. He failed to even make the qualifying rounds in the regional tournament (Hegan, 2004).

Further, there have also been machine RPS competitions (Billings, 1999), where researchers submitted automated RPS agents or “bots” to play against each other. Even when only bots compete against each other (no human players), certain

strategies are far more advantageous than others (Billings, 2000). In fact, a bot playing the Nash equilibrium using a random number generator tends to score poorly in these competitions. The results of human and machine tournaments naturally lead to the question of how opponent modeling of intentionally produced sequences might affect subjective judgments of randomness.

In a recent study of multiple repeated games of RPS between human players, Wang (2014) found that a Nash Equilibrium was never obtained by any subset of the population of participants. Rather, successful players often employ a ‘win-stay, lose switch’ strategy which is beneficial in identifying patterns in another’s strategy, exploiting them and also retaining a fail-safe strategy which prevents repeated losses. This is notable as win-stay lose-switch has also been proposed as an explanation of human category learning (Restle, 1962), and recently has been shown to approximate Bayesian inference in some cases (Bonawitz Denison, Gopnik, & Griffiths, 2014).

Below we outline two cognitive models of subjective randomness and introduce an experiment to test their robustness in explaining ternary sequences in competitive and non-competitive environments. We close by discussing the implications and limitations of the experiment in furthering our understanding of subjective randomness.

Models of Subjective Randomness

People’s judgments of randomness notoriously deviate from the prescriptions of formal probability theory in systematic ways. In experimental settings, subjects are less likely to agree that a set of coin flips that come up HHHHHHHHHH are random when compared to a set of flips that came up HTTTHHTHT. Yet, both sets of results are exactly as likely as the other given a fair coin. Even after learning probability theory, it is hard for people to escape the intuition that the latter *feels* more random than the former. How do we explain this intuition?

One popular way to model human deviations from a straightforward probabilistic account is to assume randomness judgments are a function of how difficult it is to encode a sequence or its “complexity”. In these models, psychologists try to identify an encoding process or measure of sequence complexity by specifying a theoretically motivated model that is correlated with subjective ratings. Below, we discuss two prominent models from the literature.

Falk and Konold (1997). Building on an intuition from Kahneman and Tversky (1972), Falk and Konold (1997) proposed that people ‘chunk’ a sequence into smaller subsequences which are easier to encode and remember. The perceived randomness of the sequence is inversely related to the ease with which humans can divide sequences into fewer, more manageable subsequences. To quantify this process, Falk and Konold (1997) developed their model, the Difficulty Predictor (DP), to define the complexity of a sequence to be a function of the number of runs (subsequences with the same outcome) and alternations (subsequences which switch

between two outcomes repeatedly). For example, the sequence “XXOXOX” can be described as “X twice, OX twice.” Each sequence can be encoded in terms of runs and alternations. The DP of a sequence is the sum of the number of runs and two times the number of alternations. For example, the above sequence would assign one point for a sequence length of one for the first subsequence (“X”) and a score of two for a sequence length of two in the second subsequence (“OX”). The DP for this sequence is three.

Only the smallest repeating unit is needed to calculate the score for a subsequence. As such “OXOX” and “OXOXOXOX” are both given a score of two points. Any given sequence can be apportioned many different ways: for example, the sequence “XOXOO” can be described a “XO twice, O once” (DP of 3) or “XOX once, O twice” (DP of 4). DP is the minimal score over possible encodings of a sequence.

This formalization of subjective randomness instantiates the concept of Kolmogorov complexity (Kolmogorov, 1965), which states that the complexity of an object is the length of the shortest program that can be used to generate that object. Previous work in psychology suggests that humans are adept at finding patterns in data and encoding them in a way consistent with Kolmogorov complexity (Chater, 1996, 1999). In fact, Griffiths et al. (2018) showed that DP is a special case of the complexity producing the sequence on a finite state machine with four motifs: all Hs, all Ts, alternating HTs, and alternating THs. The machine is biased to stay in its current motif.

Griffiths and Tenenbaum (2001) Griffiths and Tenenbaum (2001) propose an alternative model of subjective randomness by realizing that randomness judgments are not made in a vacuum: The randomness of a sequence is its relative likelihood of having been generated from a random rather than a *regular* process. Thus,

$$\text{random}(x) = \frac{P(\text{random}|x)}{P(\text{regular}|x)} = \frac{P(x|\text{random})P(\text{random})}{P(x|\text{regular})P(\text{regular})}$$

Their Bayesian model then differs from the standard normative account, which only considers the likelihood of the sequence assuming a random generating process, or $\text{random}(x) = P(x|\text{random})$. They implement a Bayesian model that captures the likelihood of a sequence being generated by a random process compared to a regular (non-random) process. Here, we generalize their model from binary sequences to the ternary sequences that we use in our experiment.

The probability of a ternary sequence x of length N being generated by a random process is:

$$P(x|\text{random}) = (1/3)^N$$

In contrast to a random sequence, we define a regular sequence as one that is generated by a systematic process in which each token in a sequence is generated by a multinomial

process with parameter vector $\vec{\theta}$. $\vec{\theta}$ gives the probability of each type, and so θ_1 would be the probability of the first type. Because the multinomial parameters are unknown for “regular processes”, an ideal observer should consider all possible parameter combinations:

$$P(x|\text{regular}, \vec{\alpha}) = \int P(x|\vec{\theta})P(\vec{\theta}|\vec{\alpha})d\vec{\theta}$$

where $\vec{\alpha}$ represents the parameters for the prior distribution. We use the Dirichlet distribution due to its conjugacy with the Multinomial distribution. Integrating over all possible values for $\vec{\theta}$, we find:

$$P(x|\text{regular}, \vec{\alpha}) = \frac{\Gamma(A)}{\Gamma(N + A)} \prod_{k=1}^K \frac{\Gamma(n_k + \alpha_k)}{\Gamma(\alpha_k)}$$

where $\Gamma(x)$ is the Gamma function evaluated at x , n_k denotes the number of tokens of type k in a sequence and $A = \sum \alpha_k$. Assuming equal prior odds, $P(x|\text{random}) = P(x|\text{regular})$, the complexity of a sequence can then be defined as the log-likelihood that it was generated by a random process, as opposed to a regular process:

$$\text{LR} = \log \frac{P(x|\text{random})}{P(x|\text{regular}, \vec{\alpha})}$$

Sequences with a LR greater than zero are more likely to have been generated by a random process, whereas sequences with a LR less than zero are more likely to have been generated by a regular process than a random process. Prior odds can be included in the model to shift the boundary between regular and random to a value other than zero. See Williams and Griffiths (2013) for additional empirical support of this model in capturing human randomness judgments for binary sequences.

Experiment

In previous work and both models, randomness judgments are not made in the context of two intentional human agents directly competing where the result of their competition is based on their joint decisions. Motivated by these considerations, we ask: how does a sequence being generated within a competitive context affect its perceived randomness?

In the present study we examine how well these models explain ternary sequences, rather than binary ones. We also manipulate conditions of how the sequence is assumed to be generated: either by a person playing RPS (competitive) or by the roll of a die (neutral).

Materials and Methods

We collected data from 148 subjects on Amazon Mechanical Turk. We excluded 42 subjects (28%) who had a mean response time of less than 800ms in either condition. This minimum average response time was based on an estimate of how long a subject would need to view the

sequence, encode any perceived patterns and make a motor response. The data presented here reflect the remaining 106 subjects (mean age 37.3, 53 male, 52 female, 1 unknown). Each subject saw 100 sequences (sequentially) in each of the two conditions: die (neutral context) and RPS (competitive context).

In the die condition, subjects were told that a friend was playing a board game with a six-sided die that had two blue faces, two yellow faces, and two red faces. On each trial, the subject observed a sequence of seven rolls from that die.

In the RPS condition, subjects were told that they were watching two friends play a game of rock, paper, scissors. On each trial, the subject observed a sequence of seven hand gestures from the game. (See Figure 1.)




Figure 1. (A) An example sequence from the die condition. (B) An example sequence from the RPS condition that is conceptually identical to the die sequence.

There are 2,187 possible ternary sequences of length seven. We assumed that the perceived randomness of individual classes was irrelevant. e.g. that a die sequence BLUE YELLOW YELLOW is perceived as equally random as YELLOW BLUE BLUE. This reduces the pool of sequences to 729. For each subject, sequences were randomly selected without replacement from the 729 possible sequences. Images were assigned randomly to the three types so that all 2,187 sequences were observed.

The two conditions were blocked so that subjects saw 100 trials from one condition, followed by 100 trials from the other condition. The starting condition (RPS or die) was counterbalanced between subjects. The order of trials within a condition was random, but identical for both conditions for each subject (e.g., if a subject saw sequence A from Figure 1 as trial 1, they might see the isomorphic sequence B from Figure 1 as trial 101).

Subjects rated each sequence on a Likert scale from 1 (“Not random at all”) to 10 (“Very random”), with midpoint label of “Somewhat random.” Following the experiment, subjects completed a brief demographic survey that also included questions about their level of education, experience playing rock paper scissors, and whether they had taken a statistics or probability course.

Results

We began with two hypotheses: (1) sequences generated from a random mechanism (a die roll) would be judged as more random than equivalent sequences generated from a human playing RPS, and (2) high alternation sequences generated by an intentional agent in the context of a game

would be perceived even less random due to opponent modeling in a competitive context.

To test the effect of sequence production on a subject's randomness judgments, we first compared the scores of reported randomness between conditions. We found a significant effect, with sequences produced from a die being considered more random than sequences produced by games of rock, paper, scissors, $M_{\text{die}} = 5.78$, $M_{\text{RPS}} = 5.47$, $t(105) = 2.93$, $p = .004$. This confirms our first hypothesis that participants perceive outcomes produced by people to be more random than those produced by a die.

Though individual subjects varied greatly in their mean scores of subjective randomness, there is a clear trend towards evaluating sequences from the RPS condition as more random than the sequences produced from the dice condition. See Figure 2.

We expanded this analysis further by examining whether randomness judgments are partially explained by the difficulty of encoding a sequence to memory (as they were in Falk and Konold, 1997). Using response time as a proxy for encoding difficulty, we found that there was a significant effect for reaction time between conditions, $M_{\text{die}} = 2372\text{ms}$, $M_{\text{RPS}} = 3091\text{ms}$, $t(105) = 4.45$, $p < .001$. This means subjects took longer to respond to randomness judgements in the RPS condition compared to the dice condition.




Figure 2. Each point denotes an individual subject's mean randomness judgment for the die condition (x-axis) and RPS condition (y-axis). The identity line is shown for comparison.

We calculated the complexity of each sequence according to three measures: Falk and Konold (1997)'s difficulty predictor (DP), Griffiths and Tenenbaum (2001)'s Likelihood Ratio (LR), and the probability of alternation. Overall, all three measures were highly correlated with subjective randomness judgments (see Figure 3). We found a

difference in subjective randomness judgments by condition that was moderated by the complexity of the sequence. Sequences of low complexity (as judged by either DP, LR, or probability of alternation) were judged to be equally non-random regardless of whether the sequence was in the die or RPS condition. However sequences of high complexity were judged to be more random in the die condition compared to the RPS condition. We discuss this further later in the article.

Discussion

Classic studies on subjective randomness have had subjects judge binary (e.g., heads and tails, black and white tiles), or digit sequences. The pattern of performance described by Falk and Konold (1997) is that subjects will overestimate the number of alternations that would need to be present in a truly random sequence. In a sequence of tosses from a fair coin, we expect that a genuinely random sequence has a probability of alternation of 0.5. While subjects studying coin flips might overestimate the number of alternations in a given sequence, a fully alternating sequence would not be seen as random but following a predictable alternating pattern.

This contrasts sharply with the current findings, where a truly random sequence would have a probability of 0.67. Subjects continue to overestimate the number of alternations within a random sequence, but they do so without showing a decline towards less randomness at higher alternation values. This is because using the probability of alternation is not as useful as a measure in the case of ternary sequences. For example, the sequence RPRPRP has the same probability of alternation as the sequence RPSPSR, though the latter appears more random. In a binary sequence, an "alternation" implies what the next item in the sequence will be, but this is not true for ternary sequences. This highlights a limitation of using probability of alternation as a proxy for subjective randomness judgments.

We found that on average, a sequence of die rolls was judged to be more random than an equivalent sequence of rock paper scissors throws. This effect seems to be driven by higher judgments of randomness for high-complexity sequences in the die condition compared to the RPS condition. Currently, no model adequately describes why this difference between conditions might occur, or why the differences between conditions should be primarily observed in high complexity sequences.

There are several potential explanations for these trends. One possibility is that that randomness judgments are primarily influenced by the mechanism that generates that sequence, rather than the sequence itself.

The fact that RPS throws are the product of intentional action, while die outcomes are generated by chance is a promising hypothesis. Caruso, Waytz, and Epley (2010) explored this type of intentional action as a possible explanation for differences between the hot hand effect and the gambler's fallacy. They found that participants who were told to focus on the intentions of a coin tosser were more likely to expect a coin toss streak to continue compared to

participants who were told to focus on the motor actions of the tosser. However participants were never asked explicitly to judge the randomness of sequences and it was not a directly competitive context. Similarly, Ayton and Fischer (2004) tested whether differences in gambler's fallacy and hot hand might be accounted for by animacy in the generation process. Neither of these mechanisms alone explain the results observed in our study where only high complexity sequences appear to show differences in randomness ratings.

A related explanation is that subjects may be reluctant to use the upper end of the randomness scale in the RPS condition because they are explicitly told that the sequences were generated by a human, and their belief that humans cannot (or do not) produce truly random sequences. This interpretation is anticipated by Burns and Corpus (2004) who found that subjects expect streaks to continue if they are generated by a non-random process ie: a human player. Therefore subjects might perceive sequences with high rates of alternation as less likely.

There is some counterevidence to this hypothesis in our results: z-scoring each participants' ratings does not eliminate the lower randomness ratings specific to more complex sequences..

A second hypothesis is that in the context of playing a game of RPS, subjects expect to see more complex sequences. A competent rock, paper, scissors player should try to make each throw as unpredictable as possible in order to beat his or her opponent. Therefore, we should expect a player to generate complex sequences intentionally. Subjects may have judged highly complex sequence as less random in the RPS condition because they believe a player planned that sequence in order to fool their opponent. Somewhat paradoxically, this means that sequences that are descriptively more random are seen as less random, due to the fact that they are unsurprising in the context of the game. This distinction between descriptive complexity and observed complexity has been used to explain, for instance, why descriptively simple lottery results (such as 1-2-3-4-5) are seen as more surprising (Dessalles, 2017).

Thirdly, RPS presents a sequence in a two-player game. This may lead subjects to underestimate randomness by urging them to look more closely for possible subtle patterns in the sequences generated by opponent modeling. In the die condition, each roll is assumed to be independent of the previous roll. But in the RPS condition, each throw may be conditionally dependent not only on the player's previous throw, but also the opponent's previous throw. This naturally leads to a larger hypothesis space from which subjects may be inferring potential patterns. This expansion of the hypothesis space could disproportionately affect more complex sequences and therefore explain the observed differences between low and high complexity sequences.

Hypotheses 2 and 3 operate under the assumption that expectations based on both context and generation method contribute to perceptions of randomness. While previous studies have shown that generation method plays a role, they do not explicitly contrast between a sequence produced in a directly competitive vs. non-competitive contexts. An aim of future research will be to understand how generation process, context and the complexity of a sequence may interact in order to explain the current results.

One limitation of the current study is that subjects may have a biased prior belief that die rolls are more random than RPS sequences, independently of the likelihood of a given sequence. Future studies may explicitly equate these priors. For instance, subjects could be shown two sequences of die rolls and informed that one sequence was generated by a fair die (random) and the other by a weighted die. Identifying the "cheater" in this case depends only on the likelihood, as the experiment can be designed so that the prior probability of each die is equal (0.5).

Another limitation is that our stimuli consist only of sequences of length 7, and each unit in the sequence can only be one of three possible types. It is not clear whether our results extend to longer sequences, and to multinomial sequences beyond three types. We also do not account for perceptual similarity in our stimuli: the die images in our experiment are similar to each other (except for color), and




Figure 3. Average subjective randomness ratings were correlated with Falk and Konold (1997)'s Difficulty Predictor (left), Griffiths and Tenenbaum (2001)'s Likelihood Ratio (center) and the probability of alternation (right).

participants may be sensitive to perceptual similarity when assessing sequences.

Humans are notoriously poor at inferring randomness from sequences. This cognitive error seems to be exacerbated in competitive contexts. However, this might just as easily be reframed in a different light: People are more attuned to possible patterns of behavior when they are inspecting it within a competitive context. This may lead them to be less likely to write off certain patterns as ‘mere luck’ when they might carry valuable adaptive information for future planning and strategizing.

Acknowledgements

We would like to thank Amber Nomani for providing illustrations of the RPS hand signs and 3 anonymous reviewers for their thoughtful comments on an earlier draft of this paper. This work was funded by the Office of the VCRGE at University of Wisconsin-Madison with funding from the WARF

References

- Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler’s fallacy: Two faces of subjective randomness?. *Memory & cognition*, 32(8), 1369-1378.
- Baddeley, A. D. (1966). The capacity of generating information by randomization. *Quarterly Journal of Experimental Psychology*, 18, 119-129.
- Billings, D. (1999). The first international RoShamBo programming competition. *ICGA Journal*, 23(1), 42-50.
- Billings, D. (2000). Thoughts on roshambo. *ICGA Journal*, 23(1), 3-8.
- Bonawitz, E., Denison, S., Gopnik, A., & Griffiths, T. L. (2014). Win-Stay, Lose-Sample: A simple sequential algorithm for approximating Bayesian inference. *Cognitive Psychology*, 74, 35-65.
- Burns, B. D., & Corpus, B. (2004). Randomness and inductions from streaks: “Gambler’s fallacy” versus” hot hand “. *Psychonomic Bulletin & Review*, 11(1), 179-184.
- Caruso, E. M., Waytz, A., & Epley, N. (2010). The intentional mind and the hot hand: Perceiving intentions makes streaks seem likely to continue. *Cognition*, 116(1), 149-153.
- Chater, N. (1996). Reconciling simplicity and likelihood principles in perceptual organization. *Psychological Review*, 103(3), 566-581.
- Chater, N. (1999). The search for simplicity: A fundamental cognitive principle? *The Quarterly Journal of Experimental Psychology: Section A*, 52(2), 273-302.
- Billings, D. (1999). The first international RoShamBo programming competition. *ICGA Journal*, 23(1), 42-50.
- Dessalles, J. L. (2017, July). Conversational topic connectedness predicted by Simplicity Theory. In *Proceedings of the 39th Annual Conference of the Cognitive Science Society* (pp. 1914-1919).
- Gilovitch, T., Vallone, R., & Tversky, A. (1985). The hot hand in basketball: On misperception of random sequences. *Cognitive Psychology*, 17, 295-314.
- Griffiths, T. L., & Tenenbaum, J. B. (2001). Randomness and coincidences: Reconciling intuition and probability theory. In *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 370-375). Edinburgh: University of Edinburgh.
- Griffiths, T. L., Daniels, D., Austerweil, J. L., & Tenenbaum, J. B. (2018). Subjective randomness as statistical inference. *Cognitive Psychology*, 103, 85-109.
- Hegan, K. (2004). Hand to Hand Combat. *Rolling Stone*, (940), 38-38.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3, 430-454.
- Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1, 1-7.
- Nash, J. F. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1), 48-49.
- Restle, F. (1962). The selection of strategies in cue learning. *Psychological Review*, 69, 329-343.
- Towse, J. N., & Neil, D. (1998). Analyzing human random generation behavior: A review of methods used and a computer program for describing performance. *Behavior Research Methods, Instruments, & Computers*, 30(4), 583-591.
- Walker, D., & Walker, G. (2004). *The official rock paper scissors strategy guide*. Simon and Schuster.
- Wang, Z., Xu, B., & Zhou, H. J. (2014). Social cycling and conditional responses in the Rock-Paper-Scissors game. *Scientific reports*, 4, 5830.
- Williams, J. J., & Griffiths, T. L. (2013). Why Are People Bad at Detecting Randomness? A Statistical Argument. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(5), 1473-1490.