On the Effect of Perceived Patterns in Decisions from Sampling

Doron Cohen, Kinneret Teodorescu

University of Basel; Technion

© 2021, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: 10.1037/dec0000159

Author Note

Doron Cohen, Center for Economic Psychology, University of Basel, Switzerland; Kinneret Teodorescu, Faculty of Industrial Engineering and Management, Technion, Israel.

Some of this work was done when Doron Cohen was at the Max Wertheimer Minerva Center for Cognitive Studies, Technion. This research was supported by the I-CORE program of the Planning and Budgeting Committee and the Israel Science Foundation (grant no. 1821/12). We thank Andrei Teodorescu for his valuable comments.

Correspondence concerning this article should be addressed to Doron Cohen, Center for Economic Psychology, Faculty of Psychology, University of Basel, Missionsstrasse 62a, 4055 Basel, Switzerland. E-mail: doronco30@gmail.com.

Abstract

Many real-life choices are based on previous experiences. Research devoted to these decisions from experience has typically employed static settings, where the probability of a given outcome is constant across trials. However, recent studies of repeated choice suggest that people tend to follow perceived patterns of outcomes even when true patterns do not exist (i.e., in static settings). Here we examine whether the tendency to follow perceived patterns above and beyond external incentives also characterize decisions from sampling. To this aim, we modified the static sampling paradigm to include a conspicuous sequence of outcomes while the incentive dictated disregarding the sequence. In two studies we found a strong tendency to follow the fixed pattern of outcomes. This tendency was evident not only in sampling choices where following the pattern required additional effort and did not provide additional information. The same tendency was also evident in participant's final consequential choices, where following the pattern impaired financial returns. The results were replicated after ensuring comprehension of the task, doubling the expected payoffs and also under partial feedback design. Overall, our results suggest that decisions from sampling, like repeated consequential choice, reflect a strong tendency to follow perceived environmental regularities. Our results are consistent with the assumption that during free sampling and during consequential choice, most participants respond to *when* (i.e., on which specific trials) each of the options is better rather than to *which* option is better overall (i.e., implies a more attractive distribution of outcomes).

Keywords: Sampling; Static environments; Dynamic environments; Patterns; Sequential dependencies.

On the Effect of Perceived Patterns in Decisions from Sampling

Many real-life situations require choosing repeatedly between options without prior knowledge of the option's underlying payoff distributions. In such environments people can only rely on their previous experiences (Hertwig & Erev, 2009). Research devoted to these "decisions from experience" (DfE) has typically distinguished between two types of settings. In one type, financially consequential choices are preceded by a sampling stage in which the available options can be examined without financial consequences. In another setting, the product of each choice is financially realized (i.e., free sampling is unavailable). Two influential paradigms were developed to test these types of DfE – the sampling paradigm (e.g., Hertwig et al., 2004; Weber et al., 2004) and the repeated choice paradigm (e.g., Barron & Erev, 2003), respectively.

In a sampling experiment, a participant can sample the outcomes of the choice alternatives for as long as she wants. Once the participant feels ready, she makes a final choice based on her sampling experiences, and only this last choice implies real financial consequences (e.g., buying a garment after sampling different options at the clothing store; see Wulff et al., 2018). A common assumption is that this paradigm completely separates between two aspects of behavior: Information gathering when sampling the outcomes of each option (i.e., exploration of the environment), and exploitation of the obtained information when making the final choice (Hills & Hertwig, 2010; Wulff et al., 2018; Wulff & Hertwig, 2019). Accordingly, choices during the sampling stage are assumed to reflect processes that maximize the usefulness of the information gathered (see Dubey & Griffiths, 2019 for a similar idea), while final choices reflect an effort to maximize returns (Hills & Hertwig, 2012; but see also Gonzalez & Dutt, 2011, 2012).

In contrast, in a repeated-choice experiment, each choice implies real financial consequences for the participant (e.g., trying out different brands of shampoo; see Erev et al., 2017). Behavior in the repeated choice paradigm is usually assumed to reflect an ongoing tradeoff between acquiring information (regarding the possible outcomes, e.g., trying new products) and acquiring specific values (e.g., buying again a favored product). Thus, the two paradigms differ with respect to the temporal dynamics of the realization of choice. While each choice in the repeated choice paradigm is financially realized, only the final choice in the sampling paradigm implies any financial consequence. The repeated choice paradigm can also be seen as a simplified version of the Iowa Gambling Task (Bechara et al., 1997, but note the IGT typically focuses on choice between a larger set of options).

Notwithstanding their important dissimilarities, the sampling and repeated choice paradigms have been argued to share similar underlying cognitive processes. This similarity is made evident by the fact that the two paradigms tend to yield remarkably similar results (e.g., Erev et al., 2010; Gonzalez & Dutt, 2011; Hertwig et al., 2004; Hertwig & Erev, 2009). For example, robust DfE phenomena such as the Description-Experience gap[1] and underweighting of rare events[2] were found with both the sampling and repeated choice paradigms (for example, see Barron & Ursino, 2013; Ungemach et al., 2009).

---

[1] The Description-Experience gap refers to the observation that choice tends to be markedly different when it is made based only on a description of the option's possible outcomes and probabilities, or whether it is made based on previous experiences with the same options (Hertwig & Erev, 2009).

[2] Underweighting of rare events refers to the robust observation that with experience, people tend to behave "as if" rare events (outcomes that appear with low probability) are less probable than objectively warranted. For example, behaving as if one believes that extreme negative events "won't happen to me" (e.g., Barkan et al., 1998; Erev et al., 2008).

Furthermore, Gonzalez and Dutt (2011, 2012, 2016) found that in both paradigms, alternation rates between options tend to decrease as the number of samples increase (even when controlling for sample size; see Gonzalez & Dutt, 2012; see also Teodorescu & Erev, 2014). Assuming alternation rates represent a measure of exploration of the environment, this decrease suggests that the way people sample from the available options may reflect more than just exploration: It may also reflect a process of exploitation of the available options, already during sampling. Gonzalez and Dutt (2016) also find that participant's consequential choices positively correlate with the option they sampled more frequently and more recently (Gonzalez & Dutt, 2016b, 2016a). In addition, recently sampled options with more positive (negative) outcomes were more (less) likely to be chosen by participants. This positive recency effect (see Hertwig et al., 2004) is markedly similar to the trend often observed in repeated choice experiments (e.g., see Erev & Barron, 2005; Nevo & Erev, 2012). Specifically, in such repeated tasks people often exploit the option that gave the best outcomes on preceding trials.

Taken together, the findings reviewed above support the assumption that people base their choices on a similar underlying process in the two paradigms (Gonzalez & Dutt, 2012). This assertion of similarity contrasts with other studies that assume the two paradigms evoke different cognitive processes. For example, one common assumption implies that non-realized sampling choices solely reflect a process of exploration (e.g. see Camilleri & Newell, 2011; Hills & Hertwig, 2012).

Most relevant to the current work, the two paradigms share an additional aspect that has yet to be fully explored. In both paradigms, each choice option usually represents a distribution of outcomes, and at each trial, one of the option's outcomes is randomly drawn (with replacement) from the corresponding distribution. That is, in both paradigms, trials are typically independent from one another, and the probability of observing each outcome remains the same on every trial[3]. Accordingly, most DfE studies rely on the implicit assumption that participants understand and behave in accordance with the task's underlying static nature. Even when not explicitly informed about the nature of the payoff structure, participants are assumed to learn with experience that each option represents a static distribution of outcomes. This assumption (i.e., that participants understand, believe and act upon the task's static nature) implies that arbitrary sequences of outcomes (i.e., patterns) are uninformative and are therefore ignored by participants.

Yet, several findings question the assumption that people's behavior follows the static and independent nature of the choice environment. For example, studies that focus on judgment of randomness in static environments found that many participants tend to (mistakenly) judge certain sequences of outcomes as more probable than others (this is also known as the Gambler's fallacy, Jarvik, 1951). Interestingly, participant's judgments tend to reflect sensitivity to sequential dependencies even when the same participants explicitly report a correct understanding of the static outcome generating process (i.e., that outcomes are not sequentially dependent, see Farmer et al., 2017)[4]. Yet, despite this and other long-standing evidence for an effect of perceived patterns on behavior (e.g., Barberis et al., 1998; Jarvik, 1951; Skinner, 1948),

[3] This static setting was probably developed to enable the important comparison to decisions from description, studied extensively in the Prospect theory literature (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992).

[4] These results suggest a bottom-up rather than a top-down process as the main driver of the effect of patterns on behavior (see Farmer et al., 2017). Yet, we regard the important question (of whether the effect of pattern is mainly a product of bottom-up or top-down process) outside the scope of the current paper (see Plonsky & Erev, 2017 for further discussion of this question). Instead, we focus on observed behavior, and examine whether choices in the sampling paradigm imply exploitation of perceived patterns.

only in recent years has this possibility been systematically investigated in the repeated choice paradigm.

To investigate the effect of perceived patterns on repeated choice, Plonsky, Teodorescu and Erev (2015) developed a unique sequential dependencies analysis. This analysis showed that participants seem to behave as if they attempt to exploit perceived regularities in the sequences of outcomes they observe. That is, participant's suboptimal behavior reflects a (false) belief that the environment they encounter in the repeated choice paradigm is dynamic and trials are not independent (see Navarro et al., 2016; Schulze et al., 2020 for similar observations in related tasks). Specifically, Plonsky et al., (2015) calculated the impact of a given outcome on subsequent choices in repeated choice experiments. Under the assumptions that participants (correctly) believe the payoff distributions are static (i.e., the probability of a given outcome remains the same across trials), observing a string of outcomes that appear to reflect some temporal regularity should not change one's preferences in subsequent trials. Instead, each observed outcome should be used to update the decision maker's estimation of the payoff distribution. Accordingly, with sufficient experience participants should develop an accurate representation of the payoff distributions and form the correct belief that the distributions are independent and static. Thus, the impact curve of a given outcome on subsequent choices should be flat. Alternatively, assuming learning processes give more weight to recent experiences (e.g., (Bramley et al., 2017; Gonzalez & Dutt, 2011), the impact of observed outcomes should gradually decrease over time (i.e., more recent outcomes have a larger impact on subsequent choice).

In contrast, Plonsky et al., (2015) found that experiencing rare outcomes produces a sharp transitory decrease in the rare outcome's impact on immediate subsequent decisions, reaching a nadir of negative impact after a few trials. The nadir is immediately followed by a gradual increase of the rare outcome's impact on later choices. This non-trivial wavy recency effect suggests participants behave "as if" they (falsely) believe that the probability of the rare event changes dynamically across trials. Together with the short-term negative impact, the subsequent gradual increase in the impact of the rare event provides clear evidence for a negative recency effect. This negative recency effect cannot be accounted for by learning models that do not assume some form of reaction to perceived pattern (Plonsky et al., 2015). The wavy recency effect was also observed independently in later studies that relied on the repeated choice paradigm (e.g., Szollosi et al., 2019). These findings point to the robustness of the suboptimal tendency to follow perceived patterns in static repeated choice tasks.

The main aim of the current work is to examine whether the tendency to follow perceived patterns above and beyond external incentives also characterizes sampling decisions. Finding that participants follow perceived patterns against their explicit incentives (in decisions from sampling) will support the assumption that a similar mechanism underlies choices in both the sampling and the repeated choice paradigms. It will also question the common assumption that participants detect, understand, and follow the static nature of the sampling paradigm. Alternatively, observing no effect of perceived patterns on choices in the sampling paradigm will highlight previously overlooked differences between the processes that underlie choice in the sampling and repeated choice paradigms.

Yet an examination of the tendency to follow perceived patterns in the sampling paradigm entails some challenges that cannot be easily overcome. Even if participants choose based on the patterns they perceive, each participant might perceive and follow a different spurious pattern. In repeated choice tasks, the tendency to exploit perceived patterns can be traced due to the availability of many data points following a specific, rare event. These data points are then aggregated so that the effects of several spurious patterns converge to one impact curve, demonstrating the wavy recency effect. However, the impact curve analysis suggested by Plonsky et al. (2015) cannot be used in a typical sampling experiment, as these tend to include only a small number of consequential observations (e.g., see Wulff et al., 2018).

To clarify the challenge, consider an option with a static distribution implying some loss of L or some gain of G with equal probability. Repeated choices (i.e., on consecutive experimental trials) of this option may generate different sequences of outcomes for different participants: For example, one participant choosing this option for 8 consecutive trials may observe a sequence of GGLLGGLL (reflecting the 8 consecutive L and G outcomes she observes). A second participant may observe a sequence of LLGGLLGG. Assuming participants understand and follow the static nature of the task, they would ignore spurious temporal patterns and focus on the overall mean value of the option's payoff distribution (which implies the two participants have the same experience with that option). Alternatively, participants might (wrongly) perceive the environment as dynamic and expect that outcomes can be predicted according to temporal regularities (e.g., a fixed number of losses tends to follow a fixed number of gains). If the latter is true, the participants in our example are likely to respond to the specific spurious pattern each encountered. On the 9th trial, the first participant is likely to choose the same option again, since she is assuming previous outcomes follow a pattern and therefore the next outcome is expected to be a G. Yet the second participant is likely to avoid this option, as the next outcome in the pattern she observed is expected to be an L. Under these circumstances, it would be difficult to reveal the underlying tendency to follow patterns. Instead, behavior will appear to be random (noisy). With an additional assumption that different participants might perceive patterns of different lengths, any analysis of behavior in these settings will reflect responses to numerous different spurious patterns. These observations will cancel each other out, making detection of the underlying tendency practically impossible.

Therefore, in the current work we employed a different approach. To be able to infer which specific pattern each participant followed (if any), we embedded in our experiments a simple and easily identifiable pattern of outcomes. That is, instead of the outcomes being drawn from a static distribution (which can generate many different spurious patterns), participants in our study experienced a short sequence of outcomes that repeated itself in an easily predictable pattern. To ensure that all participants perceive the exact same pattern of outcomes, we provided full feedback that includes both the obtained and the forgone outcomes (we relax this condition in Study 2). Thus, irrespective of individual sequences of choices, if participants tend to follow perceived patterns, they will most likely exploit the same intended pattern. Most importantly, we incentivized participants to ignore the pattern, in line with the typical static experimental settings in which following perceived (spurious) patterns is counterproductive.

The focus of the current paper is to test the impact of observing a specific pattern of outcomes on sampling choices. Theoretically, we aim to shed light on possible similarities between sampling decisions vs. repeated choices (and the underlying cognitive processes each implies) and refine the assumptions commonly used in modeling these two types of DfE. Empirically, we aim to evaluate the impact of perceived patterns on DfE. We focus on an environment in which following the perceived pattern leads to suboptimal choice and impairs payoff maximization.

Study 1

Method

Participants. Four-hundred participants were recruited using Prolific Academic (https://prolific.ac). As two participants abandoned the experiment prematurely, their data was omitted before any analysis took place and the final sample size consisted of 398 participants (161 female, $M_{\text{age}} = 29.4$, $SD_{age} = 9.9$, $\text{Range}_{\text{age}} = [18, 65]$). Participants were informed they will earn a fixed show-up fee of 0.45£ (about 0.62$) and will also receive a bonus based on their choices (the payoff rule differed between conditions, see Procedure). Mean bonus was about 0.16£ (about 0.19$). The experimental session lasted 6.34 minutes on average.

Procedure. All our participants faced the two problems appearing in Figure 1 in random order. Each problem was played for 27 trials and presented two unmarked buttons on-screen at the beginning of each trial. At each trial, participants had to choose one of the buttons (by clicking on it). Immediately after each choice was made, participants were presented with the outcomes both from the option they clicked on and from the unclicked option (i.e., the forgone outcome; see Camilleri & Newell, 2013 for a similar design). Each option implied a different fixed underlying sequence of outcomes. In both problems, one option yielded an outcome of 1 point on every trial, while the outcome of the other, "sequence", option differed as a function of the problem and trial number. In Problem 1, the "sequence" option consisted of the outcomes -9 and +11, presented in a fixed sequence (i.e., -9, -9, +11, +11, -9, -9, +11, +11… -9, -9), while in

Problem 2 the "sequence" option consisted of the outcomes +10 and -8, also presented in a fixed sequence (i.e., +10, +10, -8, -8, +10, +10, -8, -8…+10, +10).

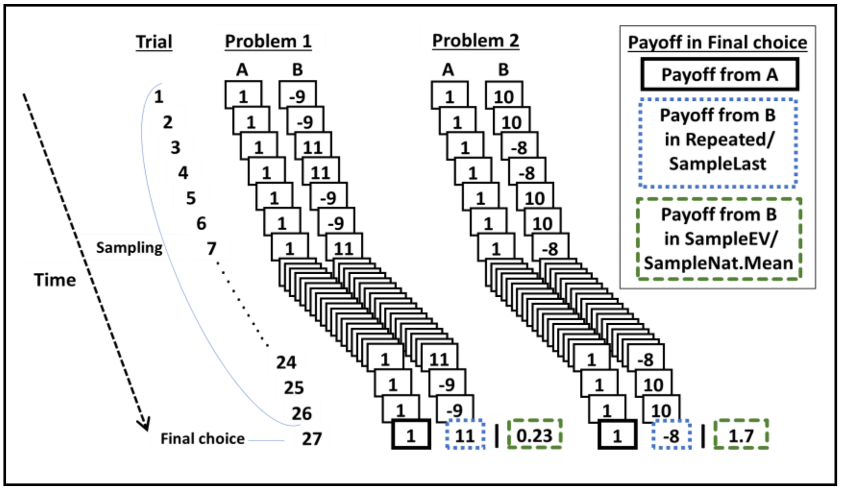

Figure 1. Participants faced the two sequences in random order for 27 trials each. In the Repeated Condition, each choice implied real financial consequences. In the three sampling Conditions (SampleLast, SampleEV and SampleNat.Mean) participants freely sampled the sequence for 26 trials, and then made a final, financially consequential choice. In Condition SampleLast participants received the last outcome of the sequence if they chose option B (blue square). In Conditions SampleEV and SampleNat.Mean, participants received the mean outcome of the sequence if they chose option B (green square).

Participants were randomly allocated to one of four conditions. In the two experimental conditions "SampleEV" and "SampleNat.Mean", participants were required to sample from the two options for the first 26 trials, and then were asked to choose the alternative with either the higher expected value overall, or the higher natural mean over the previous 26 trials, respectively (see Appendix A for full instructions). In these two experimental conditions, the final payoff on trial 27 in problems 1 and 2 was equal to the natural mean (or experienced EV) of the preceding 26 trials of each option[5] (0.23 and 1.7, respectively, with a conversion rate of 1 point = 0.15£).

---

[5] We used both the "SampleEV" and "SampleNat.Mean" versions of the instructions (which in the current experiment are numerically identical) to reduce the likelihood of confusion by participants with regards to the payoff rule. In the SampleEV Condition, participants were told their payoff would be

We also ran two control conditions, in which participants are incentivized to follow the sequence. In the control condition “SampleLast”, participants were required to sample from the two options during the first 26 trials, in any order they chose. The participants were informed that only their final choice on trial 27 will determine their final payoff (with a conversion rate of 1 point = 0.01£ + an endowment of 0.15£, to offset any losses). In this condition, participants were notified in advance that their payoff will be determined according to their outcome on the last trial. This contrasted with the two experimental sampling conditions, where participant’s payoffs were determined according to the average of the distribution chosen on the last trial. Thus, in this control condition, there is a clear incentive to follow the pattern only on the last, 27th trial. The second control condition employed a repeated choice design, rather than a sampling setting. In condition “Repeated”, participants were required to repeatedly choose on each trial between the two options appearing on screen. The participants were informed their payoff will be determined according to the outcome of their choice in one randomly selected trial (with a conversion rate of 1 point = 0.01£ + an endowment of 0.15£). Thus, in this control condition, participants are incentivized to follow the sequence on each of the 27 trials.

Participants were informed of the current trial number and were presented with instructions relevant to the stages of the task. On trial 27, a pop-up message either reminded them that only this last choice will determine their payoff (in the three sampling Conditions), or that this was the last trial for that problem (in the repeated choice Condition). The full

---

equal to the expected value of the key they choose, which is equal to the average outcome they would expect to get from that key had the choice been repeated an infinite number of times. In the SampleNat.Mean Condition, we refined the instructions: Participants were told their payoff would be equal to the key’s average value at the beginning of trial #27, i.e., the sum of the outcomes observed in each key divided by 26.

instructions used in each condition appear in Appendix A. We had 100, 98, 101 and 99 participants in Conditions SampleEV, SampleNat.Mean, Repeated and SampleLast, respectively. The mean bonuses were 0.18£, 0.17£, 0.16£ and 0.14£ in each condition, respectively.

Results

The left panel of Figure 2 presents choice rates of the option that offered the better outcome on each trial, given the set sequence of outcomes in each problem (hence, pattern-accuracy rates), aggregated across the two problems[6]. As can be seen on Figure 2, the pattern accuracy rates start around 50% and quickly increase to 70-90% (in all four conditions). We first focus on participant's pattern accuracy rates during the sampling stages of the task (i.e., trials 1-26). We use a linear mixed effects model with random intercepts for participants and fixed effects for condition (four levels) and trial (trials 1–26, treated as continuous variable), as well as the interaction between them. We used Satterthwaite approximation to estimate degrees of freedom, restricted maximum likelihood estimations are reported. This analysis revealed that the "pattern-accuracy" rates, across trials 1-26 (conforming to the sampling stages) did not differ significantly in the four conditions, main effect for Condition: $X^2(3) = 5.00$, $p = .172$; Bayesian F test, BF_{10} = 5.57E-5. In these trials, average pattern-accuracy rates were $M_{Repeated} =$ .77, 95% CI [.74, .80], $M_{SampleLast}$ = .72, 95% CI [.69, .76], $M_{SampleEV}$ = .72, 95% CI [.69, .76] and $M_{SampleNat.Mean}$ = .74, 95% CI [.70, .79]. We also found a significant main effect for Trial, $X^2(1) = 1022.6$, $p < .001$; Bayesian F test, BF_{10} = 1.88e+190, and a significant interaction between Condition and Trial, $X^2(1) = 12.4$, $p = .006$; Bayesian F test, BF_{10} = 278.7. A post-hoc analysis showed that pattern

[6] There were no significant differences between pattern-accuracy rates in Problem 1 (across the four conditions), $M_{Problem\ 1} =$ .74, 95% CI [.73, .76] and Problem 2, $M_{Problem\ 2}$= .73, 95% CI [.71, .75], $t(397) = 1.2$, $p = .228$.

accuracy rates significantly increased with trial number in all four conditions, but this increase was more pronounced in Condition Repeated[7].

Pattern-accuracy rates in trial 27 (conforming to the final choice stages) significantly differed between the conditions, $F(3,394) = 8.64$, $p < .001$; Bayesian F test, $BF_{10} = 144.3$. A post-hoc analysis showed that pattern-accuracy rates in the last trial were significantly higher in Condition Repeated, $M_{Repeated} = .82$, 95% CI [.77, .88], compared with Conditions SampleLast, $M_{SampleLast} = .69$, 95% CI [.63, .75], $t(394) = -3.10$, $p = .002$; Bayesian t test, $BF_{10} = 14.62$, SampleEV, $M_{SampleEV} = .65$, 95% CI [.59, .72], $t(394) = -3.93$, $p < .001$; Bayesian t test, $BF_{10} = 206.91$, and SampleNat.Mean, $M_{SampleNat.Mean} = .62$, 95% CI [.55, .68], $t(394) = -4.76$, $p < .001$; Bayesian t test, $BF_{10} = 6900$. There was no significant difference between Condition SampleLast and Conditions SampleEV, $t(394) = -0.80$, p = .421; Bayesian t test, $BF_{10} = 0.145$ and SampleNat.Mean, $t(394) = -1.64$, p = .100; Bayesian t test, $BF_{10} = 0.380$.

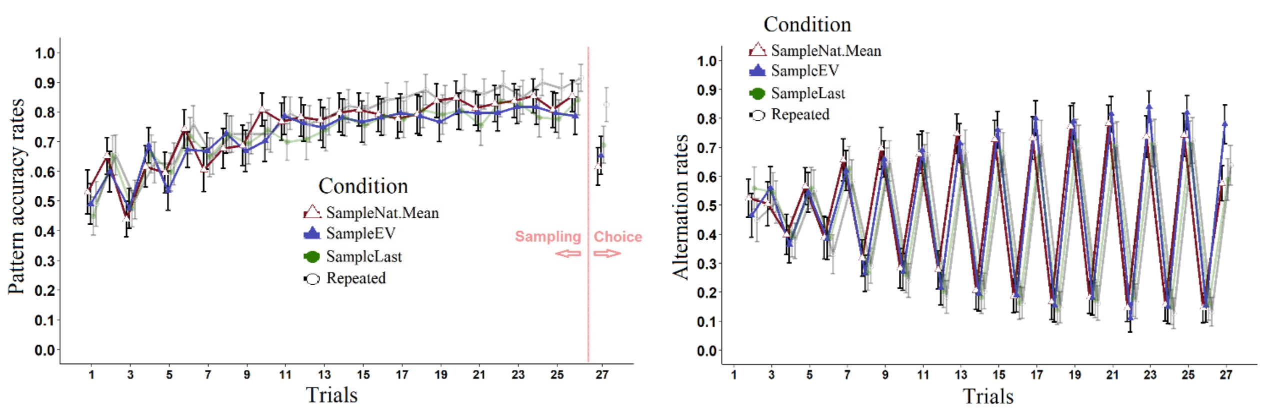

Figure 2. Left panel: Choice rates of the better-outcome alternative ("pattern-accuracy" rates) in each trial, as a function of trial number across the two problems. Right panel: Alternation rates between alternatives as a function of trials, across the two problems.

[7] In Condition Repeated, $b = .0140$, 95% CI [.0126, .0154], $t(20293) = 19.0$, $p < .001$. In Condition SampleLast, $b = .0108$, 95% CI [.0093, .0122], $t(20293) = 14.5$, $p < .001$. In Condition SampleEV, $b = .0109$, 95% CI [.0095, .0123], $t(20293) = 14.9$, $p < .001$. In Condition SampleNat.Mean, $b = .0121$, 95% CI [.0107, .0136], $t(20293) = 16.4$, $p < .001$.

Note that in the experimental conditions SampleEV and SampleNat.Mean, choosing according to the underlying sequence on trial 27 harmed the participants' payoff (see Figure 1). Nevertheless, the rate of choosing according to the pattern was above 50%. Specifically, in Condition SampleEV, the pattern accuracy rate in trial 27 was .65 (95% CI [.59, .72]), significantly higher than random choice, $t(100) = 4.79$, $p < .001$; Bayesian *t* test, $BF_{10} = 2612.6$. Similarly, in Condition SampleNat.Mean, the pattern accuracy rate on trial 27 was .62 (95% CI [.55, .68]) also significantly higher than random choice, $t(99) = 3.63$, $p < .001$; Bayesian *t* test, $BF_{10} = 46.5$.

The right panel of Figure 2 presents alternation rates between the two alternatives (also aggregated across the two problems) as a function of trial number. There were no significant differences in alternation rates between the four Conditions (see Appendix B for statistical analysis).

One possible explanation for the similarities across the four conditions, despite differences in instructions and incentives, is that participants did not read, understand, or believe the instructions of the task (but see Cohen et al., 2019). Thus, it is possible that participants in Conditions SampleEV and SampleNat.Mean did not realize that the sequence is irrelevant. In such a case, deviation from maximization should be higher before participants receive feedback (i.e., the first problem played) compared to after participants receive feedback (i.e., on the second problem). To explore this potential order effect, Figure 3 compares pattern accuracy and alternation rates in the two problems when faced first or second in order. Aggregate pattern-accuracy rates (top panels) across trials 1-26 and across the four conditions were significantly lower when the problems were first played (top left panel), $M_{1st} = .68$, 95% CI [.66, .69], compared to when they were faced second (top right panel), M_{2nd} = .80, 95% CI [.78, .82], t(397) = -16.09, p < .001, Bayesian t test, BF_{10} = 5.863e+41. A weaker, yet similar effect emerged when comparing choice on trial 27, as pattern-accuracy rates were significantly lower when the problems were first played, M_{1st} = .66, 95% CI [.61, .70], compared to when they were faced second, M_{2nd} = .73, 95% CI [.69, .78], t(394) = -2.43, p = .016, Bayesian t test, BF_{10} = 1.03. The same pattern of results was obtained when analyzing alternation rates (see Appendix C for a detailed analysis).

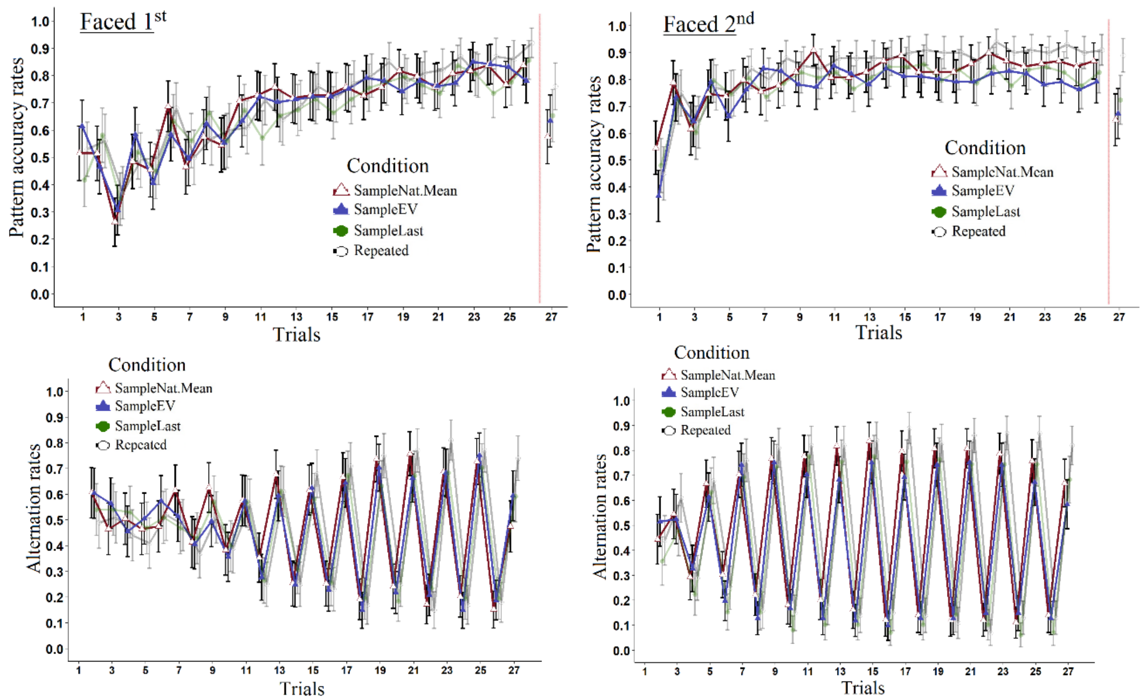

Figure 3. Top panels: Pattern accuracy rates as a function of trial number, across the two problems when faced first (left) and second (right). Bottom panels: Alternation rates between alternatives as a function of trials, across the two problems when faced first (left) and second (right).

Discussion

Overall, Study 1's results demonstrate that when the environment includes an easily identifiable pattern of outcomes, participants quickly learn to match their behavior to that pattern. The tendency to follow the observed pattern was the dominant behavior in all four conditions, regardless of the instructions and payoff rules. Moreover, participants continued to follow the pattern even when it required additional effort (i.e., during sampling with full feedback in the 3 sampling conditions) and at the expense of monetary gains (final choice in conditions Sample.EV and SampleNat.Mean). In addition, participants in the two experimental conditions continued to follow the pattern even after receiving feedback (from the previous problem) showing that following the pattern leads to lower payoffs.

The current findings can be criticized on several grounds. First, it is possible that participants did not read, understand or believe the experimental instructions (but see Cohen et al., 2019). This possibility implies that our findings could have been driven by participants' misunderstanding of the task instructions. Second, it is possible that the performance-based payoffs were too low in Study 1 (bonus payments had an expected value of 0.21 GBP and 0.12 GBP for the value maximizing and non-maximizing options, respectively), implying lower engagement with the task. This explanation frames the results as an effect of low motivation to find the best option. The two explanations suggested above raise the possibility that confusion or lack of motivation are confounded with a tendency to follow the pattern. For example, it is possible that a fundamental misunderstanding of the task led participants to attempt signaling the experimenters that they identified the obvious sequence (e.g., see Plonsky & Teodorescu, 2020). That is, it is not that participants find following the pattern more appealing than avoiding additional effort and earning higher monetary gains. Rather, they are either unaware of the consequences that such behavior entails, or they are motivated by a different, internal goal, stronger than the small external incentives.

A third, and deeper concern, is the possibility that the observed tendency to follow patterns in Study 1 is solely the consequence of the full feedback design we employed. Yet, recent findings suggest that at least in repeated choice tasks, reaction to perceived patterns is also observed when decisions are made under partial feedback (Plonsky & Erev, 2017). However, in partial-feedback settings, each discrete decision determines not only the outcome of that decision, but also the information one can rely on in the future. Thus, under partial feedback settings, different participants may observe different outcomes, and the number of possible patterns significantly increases. That is, any “sequence” of outcomes does not depend solely on the underlying distribution of outcomes in partial feedback settings, but also on the string of choices participants make.

Thus, we assume that under an incentive structure with a fixed sequence design (similar to Study 1), but without exposure to the forgone outcome, aggregate behavior will exhibit a weaker effect for the fixed sequence (while still assuming a tendency to exploit perceived patterns of outcomes). Specifically, we hypothesized that it would take participants longer to converge to the embedded sequence, compared to participants that are also shown the forgone outcome.

Study 2 was designed to test these three main critiques. First, we introduced a short, mandatory comprehension quiz immediately after reading the instructions and before starting the main task. Participants were not allowed to start the task before providing the correct answers. This constraint provided participants who were not sufficiently attentive to the instructions with the opportunity to learn the most important features of the task. We recorded participant’s number of attempts to correctly answer the quiz as a measure of initial comprehension of the instructions. Second, to increase participants engagement and motivation we double the expected monetary earnings (i.e., bonus payoffs). Third, we focus on comparing pattern accuracy rates when participants observe the forgone outcome (i.e., Full Feedback Condition) and when they are only presented with the outcomes of their choice (Partial Feedback Condition). As noted, partial-feedback settings inherently allow for perceiving many possible patterns. Since our ability to detect exploitation of perceived patterns is limited to the intended sequence, we added more trials to allow participants to find the intended pattern. This allows an analysis of the impact of forgone outcomes on behavior in our task. Importantly, in both conditions, participants were incentivized to ignore the pattern.

Study 2

Method

Participants. One-hundred and ninety-eight participants (119 female, M_{age} = 33.1, SD_{age} = 11.42, $\text{Range}_{\text{age}}$ = [18, 65]) were recruited using Prolific Academic (https://prolific.ac). Participants were informed they will earn a fixed show-up fee of 0.45£ (about 0.62$) and will also receive a bonus based on their choices (the payoff rule differed between conditions). Mean bonus was about 0.29£ (about 0.4$). The experimental session lasted 6.17 minutes on average.

Procedure. Study 2 used a similar experimental design to that used in Study 1 (see Figure 1), with the following variations. First, participants in Study 2 faced only one choice problem for a total of 43 trials. Second, we simplified the fixed sequence underlying the task's incentive structure. In each trial, one of the options presented to participants always yielded an outcome of 0 (i.e., with probability 1). The outcomes of the other, "sequence" option changed according to a fixed schedule as a function of trial number. Specifically, the "sequence" option presented a repeated sequence consisting of the outcomes +6, -3 and -4. These repeated in the same fixed order until the end of the task (i.e., a sequence of +6, -3, -4, +6, -3, -4, +6, -3, -4 … +6, -3, -4, appearing in tandem for 42 trials). Note that the first 42 experimental trials (representing the sampling stage) imply this 3-outcome sequence repeats 14 times. According to this sequence, on the final choice trial (i.e., trial 43), the +6 outcome in the local 3-outcome sequence should appear.

Participants faced one of two variations of the SampleNat.Mean Condition used in Study 1, in a between-subject design. In both conditions, participants were required to sample the outcomes of the two options (appearing on screen) for the first 42 trials. Then, at trial 43, participants were asked to choose the alternative with the higher natural mean (running average) over the previous 42 trials (see Appendix A for full instructions). Participants in both conditions were told that their bonus payment will comprise of an endowment of 0.5£ (to offset possible losses) + the natural mean (over the 42 trials) of the option they chose on trial 43. In both conditions, the value maximizing choice (i.e., on trial 43) is the “status-quo” option that yields a natural mean of 0. In comparison, the “sequence” option yields after 42 trials a natural mean of -0.33. Including the endowment of 0.5£, deviation from maximization implies an opportunity loss of 0.33£ (0.5£-0.17£).

In Condition “Full feedback”, in each trial, participants were shown both the outcome from the option they sampled, and the outcome of the unsampled option (i.e., forgone outcome, as in Condition SampleNat.Mean in Study 1). In Condition “Partial feedback”, participants were only shown the outcome from the option they sampled in that trial. We had 99 participants in each Condition.

After reading the instructions and before starting the task, participants were required to correctly answer three multiple choice comprehension questions. These served both to measure initial understanding and attention to the instructions (we also recorded the number of wrong answers submitted by each participant), and to clarify and emphasize aspects of the instructions. Questions included which payoffs will be realized, whether forgone payoffs feedback will be provided and how a "natural mean" is calculated (see appendix D for the comprehension quiz). Participants could not start the experiment without correctly answering these three questions. Immediately after making their final choice (i.e., on trial 43), we asked participants to indicate (on a scale of 0-100) their confidence that they chose, on the very last trial, the option with the higher average outcome (i.e., higher natural mean over the 42 previous sampling trials).

Before running Study 2, we pre-registered (see https://aspredicted.org/d8yh3.pdf) the following hypotheses. First, we hypothesized that the changes implemented in Study 2 (i.e., highlighting that sampling choices are inconsequential, providing only partial feedback and including larger incentives to ignore the pattern) will not eliminate the tendency to follow patterns of outcomes during sampling and consequential choice. That is, we hypothesized participants will consistently choose according to the explicit sequence, even when there is no incentive to do so (i.e., during free sampling) and when it implies deviation from maximization (on final choice). Second, we predicted higher pattern accuracy rates in the "Full feedback" Condition compared to the "Partial feedback" Condition. This is because of an inherent information discrepancy between the full and partial feedback designs: While seeing also the forgone outcome allows easy identification of the pre-defined fixed sequence, partial feedback makes identification of one specific pattern a much more complex task (as participants choices determine the pattern they see). Third, because of this information discrepancy, we predicted an interaction between Conditions and Trial on pattern accuracy rates during sampling. Specifically, we predicted participants in the "Partial Feedback" Condition will display a slower (but significant) increase in pattern accuracy rates as a function of trial number, compared to the "Full Feedback" Condition.

Results

The number of mistakes participants made when answering the comprehension quiz did not differ significantly between the conditions (M_{Full} = 1.93, 95% CI [1.53, 2.33], $M_{Partial}$ = 1.48, 95% CI [1.15, 1.82], $t(189.5) = 1.69$, $p = .092$). We found no significant relationship between the number of mistakes and observed behavior in either of the two conditions. See Appendix E for a detailed analysis.

The left panel of Figure 4 presents the pattern accuracy rates – choice rates of the option that offered the better outcome in each trial (i.e., following the pre-defined fixed sequence of outcomes). We use a logistic mixed effects model with pattern accuracy rates as dependent variable (we set random intercepts for participants; parameter estimates use the penalized least squares method). We first focus on an analysis of behavior during the sampling trials (i.e., trials 1-42). This analysis includes fixed effects for Condition (Full Feedback/Partial Feedback), Trial number (1-42, treated as continuous variable) and their interaction. We find a significant main effect for Condition ($X^2(1) = 37.56$, $p < .001$; Bayesian *F* test, BF_{10} = 1.57e+06). As hypothesized, average pattern-accuracy rates during sampling were much higher in Condition Full Feedback, $M_{Full} = .67$, 95% CI [.63, .70], compared to Condition Partial Feedback, $M_{Partial} = .52$, 95% CI [.49, .55]. We also find a significant main effect for Trial ($X^2(1) = 161.91$, $p < .001$; Bayesian *F* test, BF_{10} = 1.574267e+33), and a significant interaction between Condition and Trial number ($X^2(1) = 58.42$, $p < .001$; Bayesian *F* test, BF_{10} = 5.31e+10). A post-hoc analysis showed that pattern accuracy rates significantly increased with trial number in both conditions, but this effect was larger in Condition Full Feedback (OR = 1.043, 95% CI [1.037, 1.050], $Z = 13.80$, $p < .001$), compared to Condition Partial Feedback (OR = 1.012, 95% CI [1.006, 1.02], $Z = 4.31$, $p < .001$). This result supports our hypothesis, that due to the inherent information discrepancy, the effect of the designated sequence would emerge later in the partial feedback condition.

Pattern-accuracy rates on final choice (i.e., on trial 43) did not differ between the conditions ($t(196) = -0.298$, $p = .766$; Bayesian t test, $BF_{10} = 0.161$; Condition Full Feedback, $M_{Full} = .67$, 95% CI [.57, .76], Conditions Partial Feedback, $M_{Partial} = .65$, 95% CI [.55, .74]). In both conditions, the pattern accuracy choice rates were significantly higher than expected under random choice ($t_{Full}(98) = 3.50$, $p < .001$; $t_{Partial}(98) = 3.03$, $p = .003$). Thus, in line with our hypothesis, participants in both conditions deviated from maximization and tended to follow the underlying pattern on their final consequential choice. Although we also hypothesised this effect will be smaller in the Partial Feedback Condition, we observed no difference between the two conditions on this measure.

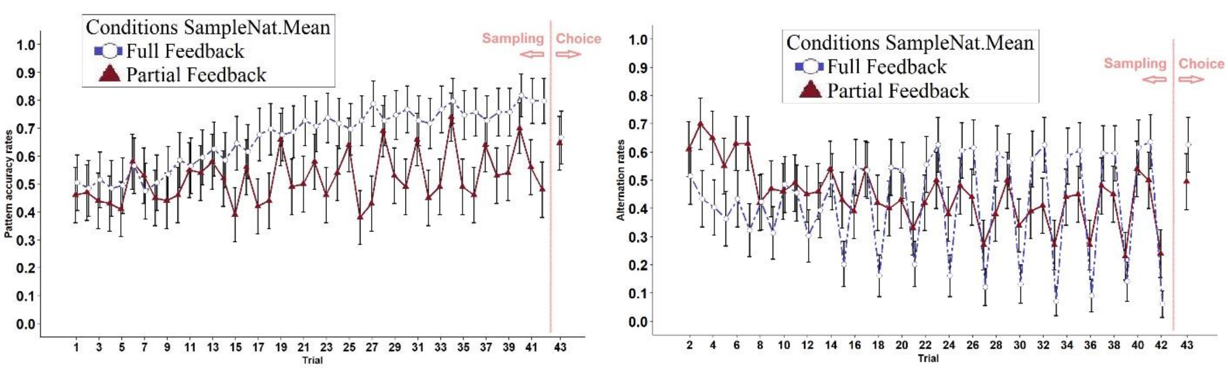

Figure 4. Left panel: Choice rates of the better-outcome alternative (“pattern-accuracy” rates) in each trial, as a function of trial number and condition. Right panel: Alternation rates between alternatives as a function of trial number and condition.

The right panel of Figure 4 presents alternation rates in the two conditions (i.e., average switching rates between the two options on each trial). A logistic mixed effect model (with the same statistical model as for pattern accuracy, with alternation rates as DV) did not find a significant main effect for Condition ($X^2(1) = 1.99$, $p = .158$; Bayesian F test, $BF_{10} = 0.03$).

Overall, average alternation rates were similar in Condition Full Feedback, $M_{Full} =$ .43, 95% CI [.38, .47] and in Condition Partial Feedback, $M_{Partial}$ = .45, 95% CI [.41, .49]. We found a significant main effect for Trial number (trials 2-43; $X^2(1) = 24.25$, $p < .001$; Bayesian F test, BF_{10} = 2021.9), and a significant interaction on alternation rates between Condition and Trial number ($X^2(1) = 53.35$, $p < .001$; Bayesian F test, BF_{10} = 4.21e+09). A post-hoc analysis showed that alternation rates increased with trial number in the Full Feedback Condition, but this increase was not statistically significant, OR_{Full} = 1.005, 95% CI [0.999, 1.010], $Z = 1.74$, $p = .081$. Conversely, the trend for alternation rates significantly decreased as a function of trial number in the Partial Feedback Condition, OR = 0.976, 95% CI [0.971, 0.982], $Z = -8.58$, $p < .001$. That is, participants in the Partial Feedback Condition alternated more during the first half of the experiment (i.e., trials 2-22) compared to the second half (i.e., trials 23-43). These results are in line with previous studies (e.g., Gonzalez & Dutt, 2011, 2016b), and can be explained with the assumption that in partial feedback settings, sampling choices reflect an exploration-exploitation tradeoff (Mehlhorn et al., 2015).

We also hypothesised that participants in the Full Feedback Condition will report higher confidence in the favourability of their final choice, compared to the Partial Feedback Condition. Although the results are qualitatively in line with our initial prediction, M_{Full} = 73.3%, 95% CI [68.5, 78.1], M_{Partial} = 67.6%, 95% CI [61.8, 73.3], this effect was not statistically significant, $t(185.5) = 1.52$, $p = .131$. In both conditions, we found no significant correlation between confidence and pattern accuracy rates either during sampling or during final choice.

Figure 5 presents the choice rates of the “Sequence” option in the two conditions of Study 2. It shows that participants in both conditions tended to exploit the positive outcomes of the fixed sequence with regularity. In Condition Full Feedback, “Sequence” choice rates on trials in which the sequence provided a positive outcome (e.g., trials 1, 4, 7 etc.) were significantly higher than random (i.e., 50%), $M_{\text{Full\_positive}}$ = .66, 95% CI [.61, .71], $t(98) = 6.67$, $p < .001$. A similar difference from random choice was observed on these trials in Condition Partial Feedback, $M_{\text{Partial\_positive}}$ = .60, 95% CI [.56, .64], $t(98) = 5.27$, $p < .001$. The difference between the two condition on "Sequence" choice rates, on these trials, was not statistically significant, $t(178.05) = 1.92$, $p = .057$.

Conversely, we found a large difference in pattern accuracy rates on trials in which the fixed sequence provided negative results. While choice rates of the "sequence" option in Condition Full Feedback were significantly lower than random $M_{\text{Full\_positive}}$ = .32, 95% CI [.27, .38], $t(98) = -6.64$, $p < .001$, these choice rates did not differ from random in Condition Partial Feedback, $M_{\text{Partial\_positive}}$ = .52, 95% CI [.48, .56], $t(98) = 1.01$, $p = .314$. This analysis is in line with the assumption that participants in the Partial Feedback Condition perceived and exploited a variety of patterns (giving the appearance of random choice when aggregated). Yet choice converged on the Sequence option when it implied positive results, leading to the predicted deviation from maximization on the final consequential choice.

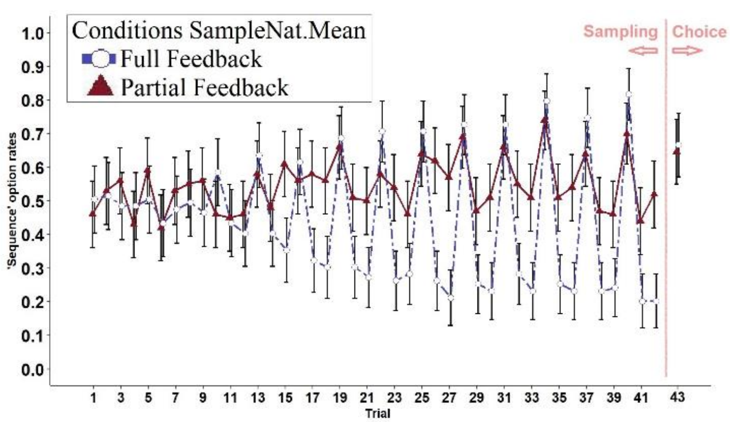

Figure 5. Choice rates of the "sequence" option as a function of trial number and condition in Study 2.

Discussion

The results of Study 2 have two main implications. First, we replicated the results of Study 1 in the Full Feedback Condition, even after assuring comprehension of instructions and with higher monetary incentives to discourage deviation from maximization. Second, Study 2 also compared the effect of the fixed sequence in settings where the forgone outcomes are not disclosed to participants (i.e., compared to Full Feedback settings). In the Partial Feedback Condition, in line with our hypothesis, the lack of information regarding the forgone outcome led to lower pattern accuracy rates during sampling. Importantly, our results suggest that the partial feedback setting did not eliminate the tendency to exploit the underlying sequence of outcomes. Our analysis of participant's choice suggests that during the sampling trials, participants in the Partial Feedback Condition tended to exploit the sequence option when the sequence implied a positive outcome. This behavior was especially evident when participants were asked to make their final, consequential choice: Choice rates that imply following the pattern (and deviation from maximization) were just as high as in the Full Feedback Condition.

One explanation for the differences observed between the two conditions relies on the fact that in the Partial Feedback Condition, participant's choices determine the pattern they perceive. This contrasts with the Full Feedback Condition, where all participants are exposed to the same pattern of outcomes. Thus, it is possible that in the Partial Feedback Condition, different participants perceived and acted on different patterns of outcomes. Another possibility is that participants in the Partial Feedback Condition recognized the underlying sequence with similar accuracy to participants in the Full Feedback Condition (albeit after more interactions with the task). Yet, because of the partial information they receive, they might put a higher premium on exploration of the task rather than exploitation of the sequence. Nevertheless, both explanations converge in predicting that consequential choice in our study would follow the underlying pattern of outcomes rather than the option with the higher natural mean.

General Discussion

In many natural circumstances, people make choices based on experiences that are sampled from the available options. While the final choice is realized, the sampling of experiences that come before it entails no financial implications (other than the information those samples convey). For example, before buying a house people typically go house-hunting, sampling many options before deciding to purchase their favorite (based on the sampled experiences). An influential experimental paradigm that was designed for the analysis of such sampling decisions is the sampling paradigm. Our aim was to examine whether people follow perceived patterns of outcomes in the sampling paradigm even when it goes against their best interest. We tested this question in an experimental setting in which following the perceived pattern implies unnecessary expenditures and contrasts with payoff maximization. We found that despite the incentives to ignore it, participants followed the perceived pattern and did not choose the option which was better overall (even when they were incentivized to do so).

Participants in our two studies faced different binary choice problems. One option always yielded a single outcome, while the other option presented a predetermined sequence of outcomes. Participants were randomly allocated to different experimental conditions that differed in the instructions and in the way monetary payoffs were determined. Specifically, participants were incentivized to find which option was better across all trials, implying they should effectively ignore the sequence to maximize payoffs. Despite this, we found that participants preferred to sample and make the consequential choice in accordance with the fixed sequence, rather than maximizing payoffs or avoiding unnecessary expenditure of effort.

In Study 1, participants were randomly allocated to one of four conditions. In two experimental conditions, participants could sample from two options before making their consequential choice. In these conditions, participant's task was to find the option with the higher average outcome. Thus, in these conditions temporal regularities should be ignored as they do not inform about the overall average of the options. In the other two (control) conditions, following the sequence coincided with the value maximizing choice strategy. We found that behavior converged to a strategy that follows the underlying pattern of outcomes in all four conditions, regardless of the instructions and incentives to ignore it. We found no differences between participant's tendency to follow the patterns during the sampling stage (i.e., when choice had no financial consequences) and when each choice was financially consequential (i.e., in the repeated choice Condition on comparable trials). Average choice on the final, consequential trial implied a strong effect for the perceived pattern in all four conditions.

In Study 2, participants were randomly allocated to either of two sampling conditions. In one condition ("Full Feedback") we replicate the results of Study 1 after ensuring participant's comprehension of the instructions and after increasing their expected monetary payoffs. Participants in that condition preferred to exploit the underlying pattern of outcomes even though this implied deviation from maximization and unnecessary investment of effort. In the second condition ("Partial Feedback"), participants received as feedback only the outcome from the option they chose in each trial (unlike the other conditions of Studies 1 and 2, where full feedback was provided). Our results suggest that, under partial feedback, although choice behavior was more variable during the sampling trials, participant's final choice followed the fixed pattern rather than exploiting the maximizing option.

Notice that in the current experimental design, the "sequence" option was also the option less frequently rewarding, and the option that provided lower outcomes in the final sampling trials. That is, our results suggest that the tendency to follow perceived patterns overshadowed not only participant's assumed effort to maximize monetary gains. Other unsupported explanations include the tendency to choose options in accordance with the (higher) frequency of rewards, with the last observed outcomes (i.e., positive recency) or with the (higher) cumulative value.

Overall, the current findings show that when the environment includes an easily identifiable pattern of outcomes, people quickly learn to match their behavior to the pattern. Importantly, our results suggest that people continue to follow perceived patterns even when it implies additional effort and at the expense of monetary gains. The current finding has several implications for decisions from experience research.

While the tendency to search for patterns is well attested by evolutionary researchers (e.g., Goldstein et al., 2010; Kolodny, Edelman, & Lotem, 2015), most existing learning models assume agents explore and exploit information about static distributions. For that reason, these classes of models are expected to fail to account for our findings. Specifically, we designed our experiments so that our results can qualitatively distinguish between the predictions implied by models with an underlying static distribution assumption and models with dynamic distribution assumptions (see Teodorescu & Usher, 2013 for a discussion of this approach). In both our studies, assuming that people react to perceived patterns (i.e., follow sequential dependencies when they observe them) leads to a qualitatively different prediction than most previously published models that do not assume this. In Study 2 for example, the sequence option consisted of a short, repeating sequence of three outcomes (i.e., +6, -3, -4) for 42 trials. Thus, when predicting choice on trial 43, any calculation that ignores sequential dependencies will inevitably include no less than two negative outcomes (i.e., a -3 and a -4) for every single positive outcome (a +6). Accordingly, models which assume maximization of the option's running average over the whole experienced sample (e.g., a natural mean heuristic, Hertwig & Pleskac, 2008) or just a subset (e.g., a recency heuristic, Erev et al., 2010) would predict participants will avoid the sequence option, in contrast to our findings.

Another popular approach to explain behavior in the sampling paradigm relies on reinforcement learning models (e.g., a value updating rule, see Frey et al., 2015). These models typically assume that valuation of each option is updated after each sample is drawn. The valuation of each option consists of a weighted average of previous outcomes from that option and the value of its most recent outcome (e.g., Hau et al., 2008). In Study 2, such models would predict that people would avoid the "sequence option" on final choice. This is because the design of the sequence ensures that both the weighted average and the last (i.e., the most recent) outcomes from the sequence option are negative (i.e., any subset of recent outcomes would include more negative outcomes and a negative mean value).

Another influential modeling approach for choice in the sampling paradigm is to base models on cumulative Prospect Theory (CPT, see e.g., Hau et al., 2008; Tversky & Kahneman, 1992) or on Instance Based Learning Theory (IBLT, Gonzalez & Dutt, 2011, 2012, 2016b; Lejarraga et al., 2012). One core assumption in CPT is differential weighting of gains and losses. In the case of Study 2, as most experienced outcomes from the "sequence" option are negative (i.e., 2 out of 3 outcomes), CPT's core assumptions predict that people will avoid it. The IBLT model in its current form assumes people rely on the most frequent and recent experiences (from each of the options) to find and exploit the choice option that maximizes expected returns. In Study 2, both frequency and recency components should have led to avoidance of the sequence option (as negative outcomes are both more recent and more frequent), again, in contrast to the current results.

Therefore, we suggest that models can increase their ecological validity by modifying the assumed object of exploration and exploitation. For example, by incorporating mechanisms that account for temporal dynamics, i.e., a process tuned to finding *when* (on which trials) an option is better rather than finding *which* option is better overall (across all experienced trials). The validity of such cognitive mechanism is also supported by studies that suggest the detection and exploitation of environmental regularities is driven by intrinsic motivation (Kolodny et al., 2014; Lotem & Halpern, 2012; Nafcha et al., 2016). This would also be in line with recent findings that show perceived patterns influence subsequent planning choices, even when this harms expected returns (Plonsky & Teodorescu ,2020).

It is also important to note that first and foremost, the sampling and repeated choice paradigms are only tools used by researchers to better understand aspects of human choice behavior. The current work clarifies the assumptions underlying the scientific usage of these tools (one such assumption is that behavior in the sampling paradigm is insensitive to perceived patterns of outcomes). Considering that these assumptions also determine the type of conclusions one can draw with regards to observed behavior (e.g., Teodorescu & Usher, 2013), clarifying them can have far-reaching implications to our understanding of human behavior. Also, note that our results can be of relevance to other tasks, such as optimal stopping tasks (e.g., the secretary problem, Seale & Rapoport, 1997). In such tasks, people might also be sensitive to sequential regularities when setting their stopping rules, even when the environment is static. This possibility should be examined in future studies.

From a practical point of view, although the tendency to follow temporal regularities in sampling decisions could be considered suboptimal in static environments, purely static environments are naturally quite rare. For example, the decision regarding which movie to watch appears to involve a static environment, i.e., the same movie does not change over time. Yet, one's utility from watching a specific movie is likely affected by mood, current affairs, and other temporal changes, some more predictable than others. Assuming reading movie reviews is akin to sampling (e.g., Wulff et al., 2015), these reviews also do not reflect a static environment as the reviewers are also influenced by context as well as by one another.

The current investigation also implies several limitations that stem from our choice of experimental paradigm and should be explicitly acknowledged. First, the possibility that higher incentives to adhere to the instructions would eliminate the tendency to follow the perceived pattern. Although Study 2's results suggest that the tendency to exploit perceived patterns is not a product of negligible incentives, we did not test the boundary conditions the incentives imply. For example, it is possible that a high enough incentive would lead participants to choose against their (perhaps more natural) tendency to follow perceived patterns of outcomes. Another limitation comes from the fact that we rely on a simplistic and relatively narrow methodology. Although we fully acknowledge this as a limitation, the challenges of shedding light on the effects of patterns in decisions from sampling constrained us to the current design. We encourage future research that can broaden the paradigm to a more general design, that will also test the robustness of the underlying phenomenon. Lastly, the current studies were conducted online, raising the possibility of too little experimental control to answer the current experimental questions. Although this is a valid concern, recent studies suggest one does not necessarily forgo experimental control when running online experiments (see Hauser & Schwarz, 2016; Jonell et al., 2020; Kees et al., 2017; Necka et al., 2016). Future research can further examine whether more controlled settings diminish the effect of perceived patterns on decisions from sampling.

In conclusion, we believe our results can be used to inform future theories and studies in DfE in general and in the sampling paradigm in particular. Specifically, these should address the impact of perceiving the world as dynamic rather than static. Assuming this tendency likely varies across individuals (e.g., Gaissmaier & Schooler, 2008), future studies should also clarify the individual differences that drive perceptions of a dynamic world. As time is a fundamental dimension, reacting to temporal patterns should be just as fundamental in decision theories, as in other aspects of human cognition.

Open Practices Statements:

Study 2 was pre-registered: https://aspredicted.org/d8yh3.pdf.

Full data and materials are available at:

https://osf.io/us9vg/?view_only=9960cc7953f446caad5adb79c8e4be40

References

Barberis, N., Shleifer, A., & Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, *49*(3), 307–343.

Barkan, R., Zohar, D., & Erev, I. (1998). Accidents and decision making under uncertainty: A comparison of four models. *Organizational Behavior and Human Decision Processes*, *74*(2), 118–144.

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*(3), 215–233.

Barron, G., & Ursino, G. (2013). Underweighting rare events in experience based decisions: Beyond sample error. *Journal of Economic Psychology*, *39*, 278–286.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, *275*(5304), 1293–1295.

Bramley, N. R., Dayan, P., Griffiths, T. L., & Lagnado, D. A. (2017). Formalizing Neurath's ship: Approximate algorithms for online causal learning. *Psychological Review*, *124*(3), 301.

Camilleri, A. R., & Newell, B. R. (2011). When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review*, *18*(2), 377–384.

Camilleri, A. R., & Newell, B. R. (2013). The long and short of it: closing the description-experience "gap" by taking the long-run view. *Cognition*, *126*(1), 54–71.

Cohen, D., Plonsky, O., & Erev, I. (2019). On the Impact of Experience on Probability Weighting in Decisions Under Risk. *Decision*. https://doi.org/10.1037/dec0000118

Dubey, R., & Griffiths, T. L. (2020). Reconciling Novelty and Complexity Through a Rational Analysis of Curiosity. *Psychological Review*, *127*(3), 455–476.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, *112*(4), 912.

Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, *124*(4), 369–409. https://doi.org/10.1037/rev0000062

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., Hertwig, R., Stewart, T., West, R., & Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*(1), 15–47.

Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, *21*(5), 575–597.

Farmer, G. D., Warren, P. A., & Hahn, U. (2017). Who "believes" in the Gambler's Fallacy and why? *Journal of Experimental Psychology: General*, *146*(1), 63.

Frey, R., Mata, R., & Hertwig, R. (2015). The role of cognitive abilities in decisions from experience: Age differences emerge as a function of choice set size. *Cognition*, *142*, 60–80.

Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, *109*(3), 416–422.

Goldstein, M. H., Waterfall, H. R., Lotem, A., Halpern, J. Y., Schwade, J. A., Onnis, L., & Edelman, S. (2010). General cognitive principles for learning structure in time and space. *Trends in Cognitive Sciences*, *14*(6), 249–258.

Gonzalez, C., & Dutt, V. (2011). Instance-based learning: Integrating sampling and repeated decisions from experience. *Psychological Review*, *118*(4), 523.

Gonzalez, C., & Dutt, V. (2012). Refuting data aggregation arguments and how the instance-based learning model stands criticism: A reply to Hills and Hertwig (2012). *Psychological Review*, *119*(4).

Gonzalez, C., & Dutt, V. (2016a). Corrigendum: Exploration and exploitation during information search and consequential choice. *Journal of Dynamic Decision Making*, *2*, 4.

Gonzalez, C., & Dutt, V. (2016b). Exploration and exploitation during information search and experimential choice. *Journal of Dynamic Decision Making*, *2*, 2.

Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description–experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, *21*(5), 493–518.

Hauser, D. J., & Schwarz, N. (2016). Attentive Turkers: MTurk participants perform better on online attention checks than do subject pool participants. *Behavior Research Methods*, *48*(1), 400–407.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534–539.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523.

Hertwig, R., & Pleskac, T. J. (2008). The game of life: How small samples render choice simpler. *The Probabilistic Mind: Prospects for Bayesian Cognitive Science*, 209–235.

Hills, T. T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science*, *21*(12), 1787–1792.

Hills, T. T., & Hertwig, R. (2012). Two Distinct Exploratory Behaviors in Decisions from Experience: Comment on Gonzalez and Dutt (2011). *Psychological Review*, *119*(4), 888-892.

Jarvik, M. E. (1951). Probability learning and a negative recency effect in the serial anticipation of alternative symbols. *Journal of Experimental Psychology*, *41*(4), 291.

Jonell, P., Kucherenko, T., Torre, I., & Beskow, J. (2020). Can we trust online crowdworkers? Comparing online and offline participants in a preference test of virtual agents. *Proceedings of the 20th ACM International Conference on Intelligent Virtual Agents*, 1–8.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk.

Econometrica, *47*(2), 263–292.

Kees, J., Berry, C., Burton, S., & Sheehan, K. (2017). An analysis of data quality: Professional panels, student subject pools, and Amazon's Mechanical Turk. *Journal of Advertising*, *46*(1), 141–155.

Kolodny, O., Edelman, S., & Lotem, A. (2014). The evolution of continuous learning of the structure of the environment. *Journal of the Royal Society Interface*, *11*(92), 20131091.

Kolodny, O., Edelman, S., & Lotem, A. (2015). Evolution of protolinguistic abilities as a by-product of learning to forage in structured environments. *Proceedings of the Royal Society B: Biological Sciences*, *282*(1811), 20150353.

Lejarraga, T., Dutt, V., & Gonzalez, C. (2012). Instance-based learning: A general model of repeated binary choice. *Journal of Behavioral Decision Making*, *25*(2), 143–153.

Lotem, A., & Halpern, J. Y. (2012). Coevolution of learning and data-acquisition mechanisms: A model for cognitive evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *367*(1603), 2686–2694.

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*(3), 191.

Nafcha, O., Higgins, E. T., & Eitam, B. (2016). Control feedback as the motivational force behind habitual behavior. *Progress in brain research*, *229*, 49-68.

Navarro, D. J., Newell, B. R., & Schulze, C. (2016). Learning and choosing in an uncertain world: An investigation of the explore–exploit dilemma in static and dynamic environments. *Cognitive Psychology*, *85*, 43–77.

Necka, E. A., Cacioppo, S., Norman, G. J., & Cacioppo, J. T. (2016). Measuring the prevalence of problematic respondent behaviors among MTurk, campus, and community participants. *PloS One*, *11*(6), e0157732.

Nevo, I., & Erev, I. (2012). On surprise, change, and the effect of recent outcomes. *Frontiers in Psychology*, *3*, 24.

Plonsky, O., & Erev, I. (2017). Learning in settings with partial feedback and the wavy recency effect of rare events. *Cognitive Psychology*, *93*, 18–43.

Plonsky, O., & Teodorescu, K. (2020). Perceived patterns in decisions from experience and their influence on choice variability and policy diversification: A response to Ashby, Konstantinidis, & Yechiam, 2017. *Acta Psychologica*, *202*, 102953.

Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, *122*(4), 621.

Schulze, C., Gaissmaier, W., & Newell, B. R. (2020). Maximizing as satisficing: On pattern matching and probability maximizing in groups and individuals. *Cognition*, *205*, 104382.

Seale, D. A., & Rapoport, A. (1997). Sequential decision making with relative ranks: An experimental investigation of the" secretary problem". *Organizational Behavior and Human Decision Processes*, *69*(3), 221–236.

Skinner, B. F. (1948). 'Superstition'in the pigeon. *Journal of Experimental Psychology*, *38*(2), 168.

Szollosi, A., Liang, G., Konstantinidis, E., Donkin, C., & Newell, B. R. (2019). Simultaneous underweighting and overestimation of rare events: Unpacking a paradox. *Journal of Experimental Psychology: General*, *148*(12), 2207.

Teodorescu, A. R., & Usher, M. (2013). Disentangling decision models: From independence to competition. *Psychological Review*, *120*(1), 1.

Teodorescu, K., & Erev, I. (2014). Learned helplessness and learned prevalence: Exploring the causal relations among perceived controllability, reward prevalence, and exploration. *Psychological Science*, *25*(10), 1861–1869.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*(4), 297–323.

Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science*, *20*(4), 473–479.

Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: risk as variance or coefficient of variation. *Psychological Review*, *111*(2), 430.

Wulff, D. U., & Hertwig, R. (2019). Uncovering the Anatomy of Search Without Technology. *A Handbook of Process Tracing Methods*, 313.

Wulff, D. U., Hills, T. T., & Hertwig, R. (2015). Online product reviews and the description–experience gap. *Journal of Behavioral Decision Making*, *28*(3), 214–223.

Wulff, D. U., Mergenthaler-Canseco, M., & Hertwig, R. (2018). A meta-analytic review of two modes of learning and the description-experience gap. *Psychological Bulletin*, *144*(2), 140.

Appendix A

Full instruction as they appeared to participants in the 4 conditions. The participants read the instructions for as long as they wished and pressed a “ready” button to indicate they are ready to start the experiment. On the experimental screen, participants were first asked to insert their Prolific ID, age and gender, and to answer the disguised attention check question.

Instructions for Condition Repeated

"This study includes two short games, each consisting of 27 trials. In each game, two unmarked keys will be presented on the screen. On each trial, your task is to select one of the keys by pressing on it. The trial's payoff can be a gain or a loss and will be presented on the selected key. At the end of the study, one trial from one of the games will be randomly selected, and your payoff in that trial will decide your bonus payment.

Your payoff from the current study will comprise of a fixed amount of 0.55£ for your participation, promised upon completion. Besides this amount, you will be paid a bonus payment, comprised of an endowment of 0.15£ + your outcome in one randomly selected trial with a conversion rate of 1 point = 0.01£.

In the next screen, you will be asked to insert your Prolific ID (important!), age and gender. In the 'Comments' text box, please only write the word 'thanks', to indicate you have read the instructions fully."

Instructions for Condition SampleLast

"This study includes two short games, each consisting of a total of 27 trials. In each game, the first 26 trials comprise the sampling stage, where you will be asked to sample the outcomes of two choice alternatives. To sample the outcomes, press on one of the unmarked keys that will appear in the next screen. After 26 trials you will get to the choice stage, in which you will be asked to choose one of the two alternatives (by pressing on it). Only the choice in the 27th trial will determine your bonus payoff.

Your payoff from the current study will comprise of a fixed amount of 0.55£ for your participation, promised upon completion. Besides this amount, you will be paid a bonus payment, comprised of an endowment of 0.15£ + the outcome from your selection in the last trial in one of the two games (randomly selected) with a conversion rate of 1 point = 0.01£.

In the next screen, you will be asked to insert your Prolific ID (important!), age and gender. In the 'Comments' text box, please only write the word 'thanks', to indicate you have read the instructions fully."

Instructions for Condition SampleEV

"This study includes two short games, each consisting of a total of 27 trials. In each game, the first 26 trials comprise the sampling stage, where you will be asked to sample the outcomes of two choice alternatives. To sample the outcomes, press on one of the unmarked keys that will appear in the next screen. After 26 trials you will get to the choice stage, in which you will be asked to choose one of the two alternatives (by pressing on it). Only the choice in the 27th trial will determine your payoff.

Your bonus payoff will be determined by the expected value of the keys you choose in the two games. The expected value represents the average outcome you would expect to get from each key, had you repeated the choice infinite number of times. For example, the expected value of a fair coin toss that gives 1 point for tails and zero points for heads is 0.5.

Your payoff from the current study will comprise of a fixed amount of 0.55£ for your participation, promised upon completion. Besides this amount, you will be paid a bonus payment, comprised of the expected value of the key you selected in the last trial of one of the two games (randomly selected) with a conversion rate of 1 point = 0.15£.

In the next screen, you will be asked to insert your Prolific ID (important!), age and gender. In the 'Comments' text box, please only write the word 'thanks', to indicate you have read the instructions fully."

Instructions for Condition SampleNat.Mean

"This study includes two short games, each consisting of a total of 27 trials. In each game, the first 26 trials comprise the sampling stage, where you will be asked to sample the outcomes of two choice alternatives. To sample the outcomes, press on one of the unmarked keys that will appear in the next screen. After 26 trials you will get to the choice stage, in which you will be asked to choose one of the two alternatives (by pressing on it). Only the choice in the 27th trial will determine your payoff.

Your bonus payoff will be determined by the average value of the keys of your choice in each of the two games. The average value is the sum of the outcomes of each key, divided by the current number of trials. This means that in each of the problems you will face, the payoff of each key in the final choice stage will be its average outcome over the previous 26 trials.

Your payoff from the current study will comprise of a fixed amount of 0.55£ for your participation, promised upon completion. Besides this amount, you will be paid a bonus payment, comprised of the average value of the key you selected in the last trial of one of the two games (randomly selected) with a conversion rate of 1 point = 0.15£.

In the next screen, you will be asked to insert your Prolific ID (important!), age and gender. In the 'Comments' text box, please only write the word 'thanks', to indicate you have read the instructions fully."

Appendix B

The right panel of Figure 2 presents alternation rates between the two alternatives (also aggregated across the two problems) as a function of trial number. The main effect of Condition on alternation rates across trials 1-27 was not statistically significant, $X^2(3) = 6.85$, $p = .077$; Bayesian F test, $BF_{10} = 0.005$. Average alternation rates were $M_{Repeated} = .48$, 95% CI [.46, .50], $M_{SampleLast} = .44$, 95% CI [.42, .47], $M_{SampleEV} = .45$, 95% CI [.42, .48] and $M_{SampleNat.Mean} = .47$, 95% CI [.45, .49].

Appendix C

To test the effect of experience with the problems on the degree with which participant's switching-behaviors aligns to the fixed sequence of outcomes, we compare alternation rates as a function of order of play, across even and odd trials. Best replying both sequences dictates switching between alternatives on every odd trial (i.e., alternate with probability = 1 on trials 3, 5, 7… 27) and exploiting the same alternative on every even trial (i.e., alternate with probability = 0 on trials 2, 4, 6… 26). To compare the effects of experience on adherence to this switching strategy, we compared the difference in alternation rates aggregating odd and even trials, across the two possible orders in which the problems were presented. When the two problems were played first (lower left panel of Figure 3), across the four conditions, the difference between odd and even trials was $M_{1st\_(odd\text{-}even)\_trials} = .30$, 95% CI [.298, .306], median = .31. This was significantly lower than the same difference when the two problems were played second (lower right panel of Figure 3), $M_{2nd\_(odd\text{-}even)\_trials} = .57$, 95% CI [.567, .575], median = .77, $t(397) = -14.72$, $p < .001$, Bayesian t test, $BF_{10} = 1.193e+36$. Thus, switching converged toward a "switch every other trial" strategy as a function of participant's experience with the task.

Appendix D

The instructional Quiz presented in Study 2. This quiz appeared immediately after participants read the experimental instructions. Participants could not proceed without answering correctly all three questions.

Quiz instructions:

"In this section we will ask you a few questions to make sure everything is clear, and the rules of the game are understood. Please answer the questions below to the best of your ability. You will not be able to proceed without answering all questions correctly.

1. *Which of your choices determine your final bonus payment?*
    A. *All my choices throughout the experiment*
    B. *One of my first 42 choices is randomly selected*
    C. *Only my last choice on trial 43 determines the bonus*
    D. *Only a correct choice will give me a bonus*

(The correct answer was <u>C</u>)

2. *What kind of feedback will you receive during the 42 sampling trials?,*
    A. *In each trial, I will only be presented with the outcome of the key I chose*
    B. *In each trial, I will be presented with both the outcome of the key I chose, and the outcome from the unchosen key*
    C. *In each trial, I will be presented with the bonus payment I can expect if I continue choosing this key*

D. *In each trial, I will be presented with the current average outcome from the key I chose*

(The correct answer was A in Condition Partial Feedback and B in Condition Full Feedback)

3. *You will be paid According to the average value of the key you choose at the final trial (average value across the 42 sampling trials). To make sure this is clear consider, the following example: Let's say you sampled three times from two keys and are now asked to make a final choice between them. You are asked to choose the key with the higher average across the sampling trials. If this is what you saw:*
   key #1: -2, 1, -2
   key #2: 3, 3, -4
   Which key has the higher average value?
   Note that in the current experiment, the average contains the outcomes from each of the 42 trials (including those you did not see).",

   A. *key #1*

   B. *key #2*

   C. *Both have equal average score*

   D. *It is impossible to know"*

(The correct answer was B)

Appendix E

Figure 1E presents pattern accuracy rates as a function of the observed number of mistakes on the comprehension quiz in Study 2 (see Methods section and Appendix D). We ran a linear mixed effects model with pattern accuracy as dependent variable, including Condition (Partial Feedback/Full Feedback) and number of Mistakes (treated as a continuous variable) as fixed effects, and their interaction. We set random intercepts for participants (we used Satterthwaite approximation to estimate degrees of freedom, restricted maximum likelihood estimations are reported). We first focus on the sampling stage (trials 1-42, see left hand side of Figure 1E). As expected (see Results section of Study 2), we find a significant main effect for condition ($X^2(1) = 38.46$, $p < .001$; Bayesian F test, $BF_{10} = 3.459e+41$). We find a non-significant main effect for Mistakes ($X^2(1) = 1.39$, $p = .238$; Bayesian F test, $BF_{10} = 0.679$), and a non-significant interaction between Condition and Mistakes ($X^2(1) = 0.02$, $p = .876$; Bayesian F test, $BF_{10} = 0.041$).

Next, we focus on choice during the final trial (trial 43, see right hand side of Figure 1E). We use a logistic mixed effects model with Condition and Mistakes as fixed effects, including their interaction (random intercepts for participants). As expected, (based on the main analyses of Study 2) we find a non-significant main effect for Condition ($X^2(1) = 0.02$, $p = .898$; Bayesian F test, $BF_{10} = 0.155$). We also find a non-significant main effect for Mistakes ($X^2(1) = 2.14$, $p = .144$; Bayesian F test, $BF_{10} = 0.401$), and a non-significant interaction between Condition and Mistakes ($X^2(1) = 0.61$, $p = .435$; Bayesian F test, $BF_{10} = 0.271$).

Overall, these results suggest that the number of mistakes participants made during the comprehension quiz cannot explain the current pattern of results. Specifically, we find that higher rates of comprehension errors do not drive the type of deviation from maximization that we observe. That is, participants tended to choose in accordance with the explicit pattern they observed, regardless of the number of mistakes they made on the comprehension quiz.

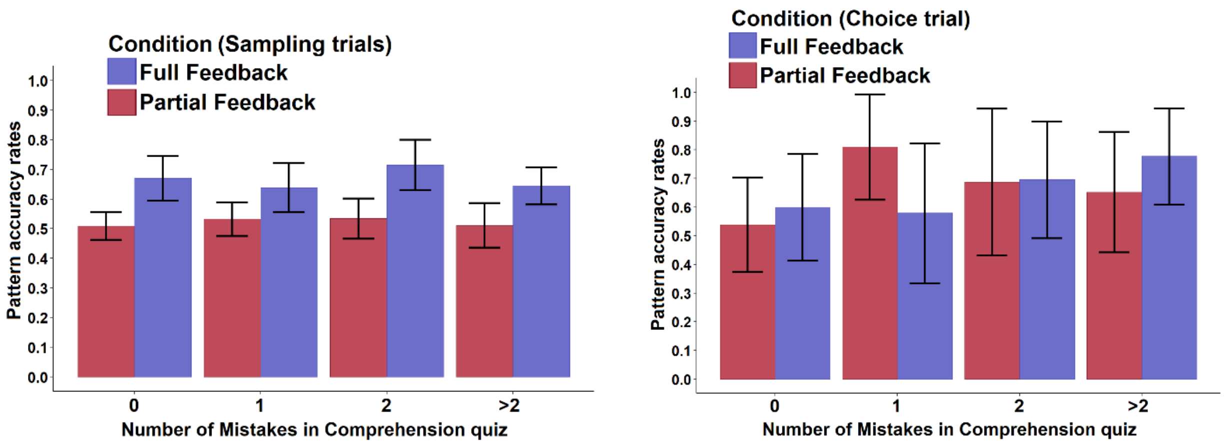

Figure 1E. Left panel: Choice rates of the better-outcome alternative ("pattern-accuracy" rates) during sampling, as a function of the number of mistakes participants made during the comprehension quiz and condition. Right panel: Pattern accuracy rates during final consequential choice, as a function of the number of mistakes participants made during the comprehension quiz and condition. In Condition Full Feedback, we had 30, 19, 23 and 27 participants that made 0, 1, 2 or >2 mistakes, respectively. In Condition Partial Feedback, we had 39, 21, 16 and 23 participants that made 0, 1, 2 or >2 mistakes, respectively.

View publication stats