First Impressions and updating of beliefs: Biases in smoking cessation treatment exploitation

Ritu Mishra

Faculty of Social and Behavioural Sciences, Utrecht University

Social, Health and Organizational Psychology

Supervisor: Harris, C.A.

Second assessor: Dr. Wieby Altink-Van den Berg

Word count: 7590

Date: 8th July 2021

Make publicly accessible

Abstract

First impressions can be biased and influence future decision-making behavior where people tend to exploit one option rather than engaging in exploration (Mehlhorn et al., 2015). The present study investigates whether, in a reward-rich environment, the individuals will form pseudocontingency biases towards a treatment option and maintain it in a free sampling task. The research is a conceptual replication of the study by Harris et al. (2020), investigating the underlined concept in a health context. In a two-armed bandit task, the participants sample between two smoking cessation treatments to minimize their nicotine symptoms. The results supported our first hypothesis as the participants sampled the frequent treatment option based on their pseudocontingency bias. However, our second hypothesis stating that the initial biases will persist in a free sampling phase was not supported. The participants sampled both the treatments and engaged in exploration. Finally, the third hypothesis tested was also supported. We observed that the intervention group in the free sampling phase maintained their initial biases in later trials. Hence, even though people form biases in reward-rich environments, these biases are not necessarily always influential in their subsequent decision making. Moreover, the initial biases were upheld in a reward-rich environment. Therefore, we could not draw firm conclusions regarding the maintenance of initial biases because we observe two groups with contrasting sampling behaviour. Thus, future research should investigate further by exploring different health contexts.

Keywords: Exploration-exploitation trade-off, pseudocontingency, treatment, initial bias, reward-rich environments

Introduction

The order in which humans perceive information significantly affects their interpretation and judgment about subsequent possibilities (Dennis & Ahn, 2001). An important aspect of deciding which option is favourable is the number of times an option is presented. People observe the recurrence of an option and draw inferences about its favourability (Vogel & Kutzner, 2017). However, when choosing between options, humans tend to be anxious; they fear making the wrong decision (Kusev et al., 2017). Therefore, they commonly tend to persist with a choice they made in the past or the choice which guarantees minimum uncertainty and a favourable outcome. Whereas exploring options brings uncertainty, fear of unpleasant outcomes, and no guarantee of the consequence. This decision-making dilemma can be explained through risk aversion; it states that people tend to opt for the immediate favourable outcome, which has low uncertainty, rather than gambling their outcome with uncertainty (Holt & Laury, 2002). Thus, although humans are good at making everyday decisions, their learning and decision-making can be biased over time with repetitive instances.

Consider the following instances: (1) You have decided to take up an online course to improve your writing skills. When searching for online courses, you find many websites offering great courses, but how do you decide which website to choose? Do you take several courses simultaneously or choose one and start learning? After choosing one, will you keep going back to the same website for other courses because you know that one writing course was good, or will you explore other websites? (2) You go to a restaurant and have to decide between two dishes. One dish is what you always eat, whereas the other dish is new. Do you eat the same thing again or try to go for the new dish? (3) You come across a machine at an arcade that shows a reward of €20 and a mystery box. Which one do you choose? If you chose €20, will you choose that option for each turn or take a chance at the mystery box, which may have a higher monetary reward? (4) When going to a pharmacy to buy medicine for the common cold, there are many medicines to choose from; how do you choose which medicine is best for you? Do you buy the one primarily used by everyone or gather more information to choose the one best for your body? When you have decided on the medicine, will you keep taking the same medicine when you are sick again or try another?

People are often confronted with such instances in daily life where they need to find the balance between two or more seemingly good options and predict which decision will lead them to a favourable outcome. Nevertheless, sometimes, people stay with suboptimal

options when making such a decision due to false beliefs, for example, in the previous case of buying medicine for a common cold. If you decide to consume the medicine you always choose when you get sick, you rely on your initial bias. This initial bias is guided by the first impression of the event and its outcome. Hence, the memory of recovering (favourable outcome) by taking medicine each time influences you to exploit that medicine option repeatedly.

Similarly, in choosing between €20 and a mystery box, we may see people repeatedly choosing €20 in each turn rather than choosing the mystery box. This behaviour of repeatedly choosing the same action over varied instances rather than learning or choosing an alternative option is called exploitation. Over time, we expect people to learn the best alternatives by taking risks (exploration). Yet, the instances described earlier show that people tend to persist with their initial biases, driven by uncertainty and anxiety, in making wrong decisions. For example, in recent research, Krypotos et al. (2021) explore how individuals solve this exploration-exploitation dilemma in a task associated with receiving pain or reward. They found that individuals relied more on exploitation, and this behaviour persisted when they could maximize rewards. The current study investigates how the exploration-exploitation trade-off plays a role in the persisting bias resulting in exploitation. Specifically, we explore this behaviour in a health context to study how people behave when making health decisions, do they learn more and explore or do they know enough to exploit it.

Exploration- exploitation trade-off

Exploitation involves choosing the option that may be relatively safe, already learnt and less risky to execute. Examples of exploitation behavior are using the same phone for years, eating the same food at your favourite restaurant and taking the same route to the park. In contrast, exploration involves opting for different alternatives, learning new information and taking risks. Examples of exploration behaviors are changing your phone when the new model is launched, trying new restaurants, and exploring new paths in the city. However, when confronted with a situation requiring decision-making between two or more options, individuals rely on previously experienced outcomes of exploration and exploitation similar to that situation (Humphreys et al., 2015). Hence, people often rely on their initial biases and stick with the option, which is safe or already learned, thus exploiting that option. The tendency to exploit the initial choice rather than exploring potentially more optimizing and beneficial options persisting due to initial biases is called the exploration-exploitation trade-off (Mehlhorn et al., 2015). For example, a study by Morris and colleagues (2016) explored

the compulsive behaviour of alcohol use or food in people with alcohol use disorder and binge-eating disorder by using exploration and exploitation tasks and found that participants engaged more in exploitation and reward maximizing to avoid uncertainty or loss. The trade-off between the exploitation of known sources of reward versus exploring new alternatives is a common dilemma faced in various domains such as business (Kim et al., 2012), management (Hardy et al., 2019), psychology (Cohen et al., 2007), neuroscience (Humphries et al., 2012) and other disciplines.

People continue to exploit sources to maximize rewards even when new information is received; this is steered from their strong first impression (bias) and the need to maximize rewards (Harris et al., 2020). Although there is nothing wrong with exploiting an option as long as the first impression is accurate, we are not always correct about these impressions. We know that initial biases can easily influence these first impressions. Hence when making decisions, choosing a less optimizing option can be decision making can be disadvantageous.

Biases and heuristics

The process of decision making is more complex than it seems because of limited resources, time, and knowledge (Campo et al., 2016). Hence, humans tend to rely on mental shortcuts to make everyday decisions quickly. Heuristics are mental shortcuts that help us make frugal and less time-consuming decisions, allowing us to overlook part of the information while making effective decisions (Gigerenzer & Gaissmaier, 2011). For example, Tversky & Kahneman (1973) have demonstrated that people tend to overestimate the probability of an event's availability by the ease with which the event's occurrence comes to mind. This bias of relying on immediately recalled information is called the availability heuristic. It can lead to systematic biases in judgement and contingency tasks, leading to misinterpreting the probability of an event and inferring false beliefs from it. For example, when an unlikely health risk happens, people tend to get insurance. Thus, although it is easy to use heuristics in everyday decision-making, it is also how cognitive biases are formed (Todd, 2000).

One such important cognitive bias is the Primacy effect. The primacy effect in impression formation explains that humans tend to have a strong memory of information presented at the beginning of a series task compared to the middle of the task (Sullivan, 2019). This is because the working memory is less loaded at the start of a task, which allows the brain to retain and encode the first information much more concretely than the information presented later (Waugh & Norman, 1965). This would explain why some people

exploit their initial preference instead of exploring or gathering more information on other options.

Choice-supportive bias is another cognitive bias; it is the tendency to recall and attach more positive features to a choice once made, for example, recalling choosing a specific medicine when you were sick and over-generalizing its positive effect of treating ailments (Mather et al., 2000). Recall the example discussed earlier, when choosing between common cold medicines, one may assign higher positive features to medicine taken earlier than the other option. Another cognitive bias is the confirmation bias. It is the tendency to search, recollect and process the received information in a way that favours one's existing beliefs (Nickerson, 1998). This would explain why some people may develop unwarranted beliefs.

Further, the type of environment also influences the decision and beliefs people make in certain situations. For example, in a reward-rich environment where the outcomes are usually favourable, the exploitation behaviour is easily adapted (Harris et al., 2020). Due to the initial inference of a positive outcome, the desire to maximize rewards causes this exploitation behavior. On the other hand, in a reward-impoverished environment, the outcome is more frequently negative than positive, beginning an exploratory behavior. However, a positive bias towards an option when exploring different options can create an initial bias causing people to engage in exploitation because then the option is believed to be better than the alternative. Thus, in a reward-rich environment, the need to explore is not experienced as the positive rewards reinforce further exploitative behavior.

Pseudocontingency

These biases formed from heuristically inclined thinking is developed because the environment is often skewed. For example, we may encounter that one option is more prevalent than another in the environment, causing us to perceive that one option is better. This belief formation is investigated through a field of research called 'Pseudocontingency'. Pseudocontingencies are logically unwarranted beliefs that create an illusion of likelihood between two variables deduced from information about the observed variables in other environmental contexts (Fiedler et al., 2009). Therefore, when people see one option occurring more frequently than the other and find one outcome more prevalent, they perceive an association between the option and the outcome. Pseudocontingency interpretations may be convenient and useful, but they can also override genuine contingencies (Fiedler, 2010). People may perceive contingencies between option and outcome that are unrelated and form such biases in environments where both options are equally rewarding (Fiedler, 2000; Meiser

& Hewstone, 2010). Most literature on pseudocontingency and its role has mainly explored decision making (Fiedler & Wanke, 2009; Vogel & Kutzner, 2017) and stereotype formation (Kutzner & Fiedler, 2017; Meiser & Hewstone, 2010). Expansion of the field has led researchers to explore different questions. Recent studies have explored the role of pseudocontingencies in actual sequential contingency tasks. In the study by Meiser and colleagues (2018), they presented the participants with options and outcomes (positive or negative) over several trials in the learning phase. After this, they made choices based on the already learnt contingencies. The results show that participants assumed pseudocontingency during the learning phase, which influenced their choice behaviour. It was also found that pseudocontingency was principally more influential in the case of positive outcomes.

Similarly, Harris et al. (2020) tested how the initial biases persisted in reward-rich and reward-impoverished environments in a sequential contingency task. They specifically focused on how these initial beliefs persist in the contingency task and its influence on future choice behaviour. The results revealed that in a reward rich condition, the participants perceived pseudocontingency and engaged in exploitation. Whereas in a reward-impoverished condition, the biases weakened.

This suggests that people see a favourable option A and exploit that choice if it is in a reward-rich environment, forming an initial bias for option A and considering option B as inferior. The research findings by Harris et al. (2020) further help establish a basis for understanding and testing why humans can have such strong yet false beliefs upheld for a long time. Conclusions by Harris et al. (2020) can be helpful to investigate actual situations in a health context. For example, in the context of treatment choice, one treatment may be preferred more than the other due to initial bias favouring one treatment more than the other.

Similarly, in the current study, based on the research of Harris et al. (2020), we attempt to explore how initial biases will be formed and sustain in future decision behavior in a reward-rich environment, which is hard to attenuate even over extended trials. It is also expected that this initial belief will be upheld and will also influence the decision making in similar instances. We test the claims of Harris et al. (2020) in a hypothetical health scenario where participants will be asked to choose between two treatments. We test this in a health context because if biases persist in reward-rich environments, it may be damaging for individuals in a real-life health scenario where they might fail to optimize the best option.

Present study

In the current research, we investigated how the initial belief regarding two smoking cessation treatments are maintained and how this first impression influences the decision-making behaviour. We conceptually replicate and adopt the same method as used by Harris et al. (2020). They claim that biased sampling through pseudocontingencies in reward-rich environments could explain the initial biases in the health context. Hence, we test if these claims are valid. We conceptually replicate the study to help further understand and establish the validation and generalizability of their claims.

In this research, the participants will repeatedly choose between two smoking cessation treatments throughout the experiment in a reward-rich environment. We test three hypotheses. Participants will perceive one treatment as more favourable than the other, thus forming an initial bias towards the frequent option by means of a pseudocontingency bias. Hence, our first hypothesis is "In a reward-rich environment, participants will show their preference towards the frequent option as they will perceive it as favourable". Second, when a pseudocontingency bias is achieved, we expect that the initial bias towards the frequent treatment will persist in the free sampling phase, influencing the choice behavior in the remaining trials. Thus, our second hypothesis is "The initial bias will be maintained in a free sampling trial, and the participants will exploit the choice perceived as more favourable".

Additionally, we test the assumption made by Harris et al. (2020) that the initial biases will be upheld in reward-rich environments. We test an intervention to investigate the accuracy of this assumption by randomly dividing participants in the experiment into two groups when they begin the free-sampling phase. The intervention group will receive additional instruction to sample equally at the beginning. This will help us understand that if an initial bias is induced, and when participants engage in exploration for a while, do they continue exploring or are eventually guided by their initial biases when sampling repeatedly. Thus, our third hypothesis is "Initial biases will attenuate when participants engage in exploration and will continue exploring in this sequential sampling". We expect that participants will not sample based on their initial biases. Instead, initial exploration will help them overcome their initial biases. Hence, they will continue to explore throughout the trial, overcoming the initial bias as they sample both treatment options.

Method

The experiment is set out to validate the hypothesis based on the claims by Harris et al. (2020). Specifically, we test if, in a reward-rich environment, participants form a

pseudocontingency bias towards a treatment option, if these biases sustain over repeated trials and if these biases can be updated. The sample size of our study was estimated to be around 100; we estimated this using power analysis G*power (Faul, Erdfelder, Lang, & Buchner, 2007). These calculations were based on a 5% alpha-level, 80% statistical power, and effect sizes between ηp2 = .081 and ηp2 = .270 as used by Harris et al. (2020) and as reported by Meiser, Rummel, & Fleig (2018). Thus, a larger sample is expected to give meaningful results and increase the power of the study to draw significant inferences.

Participants for this study were recruited via the online crowd-sourcing platform, Prolific Academy (https://www.prolific.co/) and the study was run on SoSci survey, a professional tool for academic online questionnaires. One hundred twenty participants (Nfemale = 70) with an average age of 22 years (SD = 5.68) participated in the study, with 64.2% participants who had or are enrolled in a psychology program. This sample was acquired through a convenience sample, using the university's SONA SYSTEM with a .25 PPU as a reward for 10 minutes of participation, and it was distributed via email and other social media platforms to recruit participants. The experiment was run in English. The research was conducted according to the guidelines and approval of the Ethics Review Board of the Faculty of Social and Behavioral Sciences at Utrecht University.

Design

The experimental design was constructed to match the design used by Harris et al. (2020), i.e., a two-armed bandit task where the participants sampled between two treatment options. They were instructed to imagine that they were smokers who had decided to refrain from engaging further in this bad health habit. Hence, they enrolled in a 'smoking cessation programme'. They were then presented with a choice-based task to choose between two (hypothetical) smoking cessation medicines; namely, Cefaparcin and Nexadotin, which either helped in alleviating nicotine symptoms or the symptoms returned with no change. We counterbalanced which treatment option would be presented frequently in the induction phase and which treatment option would be asked for when giving estimates of preference and conditional estimates. The first hypothesis was tested in the induction phase.

Further, the second hypothesis was tested in the free sampling phase. The participants were free to sample either cessation treatments with the goal to minimize their symptoms. Lastly, the third hypothesis was tested in the free sampling phase, with a slight difference in instructions. Some participants were randomly assigned to receive instructions to sample

between the two treatment options equally in the initial trials to minimize their symptoms. Therefore, we had set the study to a between-subject design.

Procedure

The experiment was divided into four phases; first, in the induction phase, where the participants were forced to sample only one treatment option at a time. Second, the estimate phase, where participants indicated their estimation regarding the treatment. Third, a free sampling phase, where the participants could sample between the two treatments freely. Lastly, a final estimate phase, where the participants gave their estimation regarding the treatments again.

First, in the induction phase, participants were presented with only one treatment option at a time. After selecting the treatment option presented, the following feedback was shown 'you took Cefaparcin/Nexadotin, and your nicotine cravings went down(positive)/up (negative)' with a smiling or frowning face. This feedback disappeared after a delay of one second before going onto the subsequent trial. We used the same distribution for the initial evidence as Harris et al. (2020). Hence, the participants sampled for 16 consecutive trials. Specifically, they forced sampling of one option 12 times (9 positives and three negative outcomes) and the other option 4 times (3 positives and one negative outcome).

Second, after the induction phase, participants were asked to give estimates. In this estimate phase, they were first asked to give their preference estimate. Specifically, they were asked which medicine was more likely to decrease their nicotine cravings. They could respond by moving the slider freely towards the image of the treatment option, which was anchored at either end of the slider (see Figure 1). Participants also answered the question regarding the conditional probability for both medicines. They could answer the percentage likelihood (0% - 100%) of it occurring through their observation during the trial by moving a slider. Specifically, they were asked to indicate how likely was it that if they selected a medicine, the outcome would be a decrease in nicotine cravings. They indicated this separately for both the treatment options. In addition, participants were asked to give their confidence estimate (how confidently they can say) that the estimates made were reasonable. Again, this was answered by adjusting a slider anchored at ‘*not confident at all*’ and ‘*very confident*’.

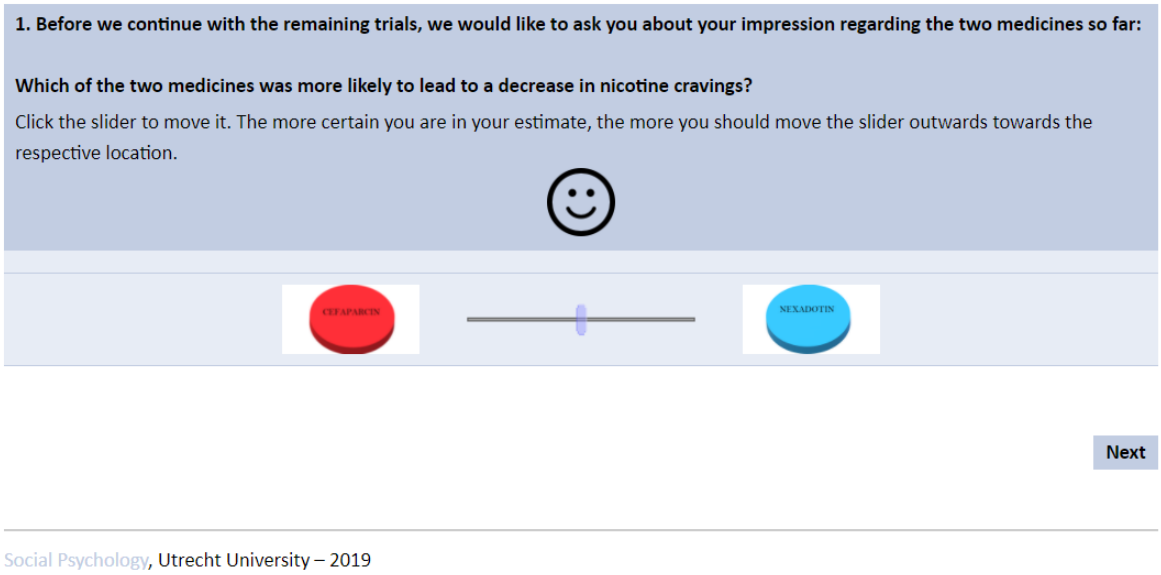

Figure 1. Relative preference estimate

Before the third phase began, the participants were presented with information about how they would receive additional visual information regarding their health progress for the remaining trials. This additional feedback measure was an image of the lungs. This graphic image of the lungs indicates the health progress. Namely, the pink lungs indicated healthy lungs, purple indicated unhealthy lungs, and grey indicated smoker's lung. In addition, they were instructed to personalize the treatment to reduce nicotine cravings. We randomly divided participants into two groups: the control and equal group. While both the groups could personalize the treatment and sample freely between the two treatments, the equal group received an additional instruction asking them to sample both the treatments equally often in the following ten trials. After these instructions, the free sampling phase began. The participants would choose between the two treatments, and feedback appeared after the selection, same as in the induction phase. The lung image changed according to the feedback. Participants were free to sample between both the treatment options for the remaining 84 trials. Throughout the experiment, the participants were able to see the current trial number from the total number of trials ("Trial: x/100") at the top right corner of the screen.

In the final phase, the participants were asked to give their estimates. Again, the questions asked were the same as the second phase. Lastly, the participants' demographics were collected, namely, age, gender, education, and if they are/were psychology students. The survey ended by thanking the participants for their participation and informed about what the study attempts to explore. Finally, the collected participant data was counterbalanced. We

did this for the frequent treatment option, the direction of the question (increase vs decrease), primacy (left vs right), the estimates given and intervention (control vs equal).

Data preparation

All the dependent variables, i.e., estimates, were recoded. Recoding was done so that higher positive values would indicate a preference for the initially frequently presented option. For example, if Cefaparcin was presented more frequently in the induction phase, and if participants indicated higher positive values in their estimates, this would indicate a bias towards Cefaparcin. Since the values recorded for the conditional probability estimate included two estimates (one for each treatment option), we calculated the difference between the two to determine the contingency of the treatments (Allan, 1980), which they indicated as the one to decrease their symptoms.

Data preparation and analysis was undertaken using Statistical Program for Social Sciences (IBM SPSS Statistics for Windows, Version 26). We expected a bias towards the frequent treatment option across the estimate measures, preference estimate, conditional estimate, confidence estimate, behavioural measure, and choice index. Therefore, one-sample *t*-tests were performed to analyze these measures. In addition, we used repeated-measures ANOVA for investigating our research questions.

Results

Pre-Relative preference measure

Following the initial trials in the induction phase, the participants made a relative preference estimate after the initial trials. We expected that they would prefer one treatment more over the other. Hence, to investigate if the bias was successfully induced through pseudocontingency and if the participants perceived one treatment as better than the other, we performed a one-sample *t*-test. The results show that the bias effect was significant ($M = 6.61$, $SD = 29.71$), $t(119) = 2.22$, $p = .026$, $d = 0.22$, 95% C.I [0.74, 11.48]. Therefore, the bias induction was successful, supporting our first hypothesis as the participants perceived an initial bias towards the frequent option as we had expected.

Pre-Conditional probability measure

The second measure determined the conditional probability estimates where the participants were asked the likelihood of choosing a treatment and that it led to a decrease in symptoms. As expected, the participants estimated that the frequently presented treatment led to more of a decrease in nicotine symptoms than the infrequent treatment option. The *t*-test

was statistically significant as can be seen by a positive ΔP value (ΔP = 0.05, *SD* = 0.29), *t*(119) = -16.47, *p* < .001, *d* = 0.19, 95% C.I [-0.49, -0.38].

Pre-Confidence probability measure

Lastly, participants were asked to give their confidence in judgement regarding the estimate they made about the treatment options. The result was significant, *M* = -.00, *SD* = 0.26, *t*(119) = -0.10, *p* < .000, *d* = -0.00, 95% C.I [-0.05, 0.04]. Therefore, the participants were more confident in their estimate regarding the frequent treatment option than the infrequent treatment option.

Free sampling phase and Intervention

After the preference estimates, the participants were free to sample for the remaining 84 trials. We wanted to test if participants continue to favour one treatment over the other in the remaining trials based on their initial bias.

We conducted a one-sample *t*-test to see if there was a difference in their choice behavior after the induction phase. The result showed a non-significant effect, *M* = 0.51, *SD* = 0.23, *t*(119) = 0.91, *p* = .365, *d* = 2.20, 95% C.I = [-0.02, 0.06]. This indicates that the participants did not sample the frequent treatment option as expected throughout the free sampling phase. Therefore, our second hypothesis, which explored if the initial bias would persist in the free sampling phase, was unsuccessful. In contrast, we observed the opposite result, where the participants sampled both the frequent and infrequent options.

We further investigated what the sampling behaviour of the equal group (intervention) was in our study. A repeated-measures ANOVA was performed to measure the group differences between the control and equal group in our study. The within-subject design consisted of 84 levels corresponding to the number of trials in the free sampling phase (dependent variables). The between-subject factor (independent variable) comprised of equal group (intervention) data. We compared the group differences between the equal group and the representative group. We first began the analysis by checking for outliers. The data showed no outliers as assessed by examination of studentized residuals for values greater than ±3. The test of between-subject effects shows that there was a non-significant main effect of intervention, *F*(1,115) = .000, *p* = .990. Hence, the third hypothesis under investigation to explore if the initial bias can be overcome was rejected. The participants in the equal group exploited the frequent treatment persisting due to the initial bias instead of exploring the other treatment option.

Figure 2 shows the percentage of participant sampling for the frequent treatment throughout the experiment, i.e., for each trial. Notice the yellow line indicating participant sampling in the 'control group' for all the trials. Whereas the blue line indicates the equal group. The control group initially sampled the frequent treatment, but the participants explored both the treatment options as the trials progressed. Eventually, we can observe that their preference shifts towards the infrequent option as they sample that option more often. This can be observed as they cross the 0.5 chance level. Whereas the equal group showed the same sampling behaviour as the control group, halfway through the experiment, we observe that they change their sampling behavior and begin sampling the frequent option more often than the infrequent, thus maintaining their initial bias.

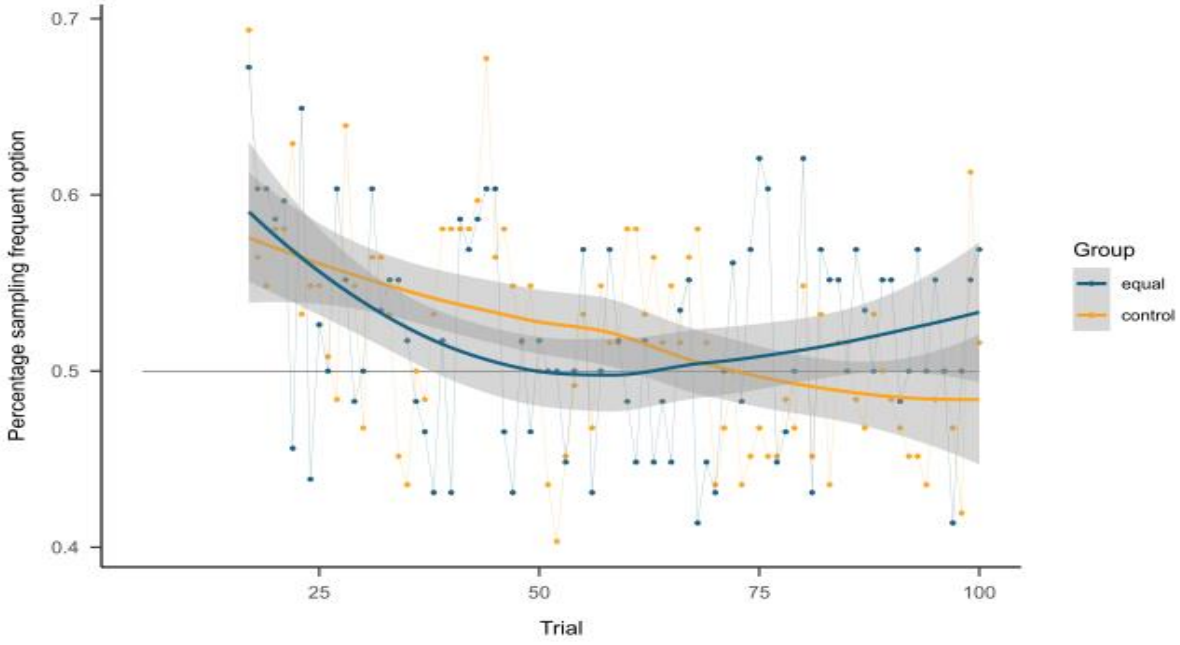

Fig.2 Percentage of participants sampling frequent treatment per trail (Chance level: 0.5)

Post-Relative preference measure

After the free sampling phase, we asked the participants about their relative preference estimate again, expecting the initial bias to persist. To measure their preference, we performed a one-sample *t*-test again as we did in the pre-estimate measures. We expected that participants would maintain their initial bias towards the frequent treatment when they sampled freely for 84 trials, however the results indicated no significant bias effect, $M = 4.51$, $SD = 31.31$, $t(119) = 1.40$, $p = .163$, $d = 0.14$, 95% C.I. [-1.64, 9.67]. Thus, participants did not show a biased preference towards the frequent option; instead, we observed they explored the treatments. Hence, our hypothesis testing if the initial bias will persist in free sampling

phase was not successful. Although, as we discussed above, the equal group (intervention), on the other hand, showed a bias towards the frequent option.

Post-Conditional probability measure

The participants gave the estimates of how likely they were to receive a decrease in nicotine symptoms given they has selected either treatment options (Cefaparcin and Nexadotin). The results show a significant conditional estimate, as can be seen by a positive ΔP value of (M = 0.06, SD = 0.37), $t(119) = -12.90$, $p < .001$, d = 0.16, 95% C.I. [-0.50, -0.37].

Post-Confidence probability measure

Lastly, another t-test was performed to record how confident were the participants in the estimates they made regarding the two treatment options. The result showed a significant confidence estimate, M = .00, SD = 0.25, $t(119) = -21.50$, $p < .001$, d = .01, 95% C.I. [-0.54, -0.44]. Thus, indicating that they were confident in their judgement of the estimates they made reading their treatment choice.

Discussion

In this conceptual replication of Harris et al. (2020) study, we tested their claims to increase the generalizability and test the validity of the research. The present study tested if the pseudocontingency in a reward-rich environment persists due to initial biases in a health context. We asked the participants to sample between the two treatment alternatives and give their estimates regarding their choice in a sequential contingency task. The first hypothesis was accepted and also supported the claim by Harris et al. (2020). In the forced-choice task (induction phase), the participants perceived the frequent treatment as a more favourable option than the other treatment option. Thus, forming an initial bias and supporting our hypothesis. The second hypothesis tested if the initial bias is maintained or attenuates in the free sampling phase. We expected participants would sample the perceived favourable option more often, thus maintaining the bias. However, this was not the case as participants attenuated their initial biases. Lastly, we tested if these initial biases can be updated. The results showed that the initial biases persisted in the equal group. Thus, our hypothesis was not supported as their initial biases guided participants even after exploring.

The results in the free sampling phase contradict the findings by Harris et al. (2020). We found that the participants initially sampled the frequent treatment more in the free sampling phase; however, they did not carry this initial bias throughout the experiment. Instead, we observe that the participants sample both the treatment options more equally

throughout the trials, thus engaging in exploration. Eventually, their initial biases attenuate towards chance level as their preference shifts towards the infrequent treatment option. Hence, as hypothesized, that in a reward rich environment, we expected the initial biases to maintain while free sampling was found to be false in our study. This is a good strategy in a health context because exploration gives new information. This way, participants will not be biased in their decision making and stuck with a suboptimal treatment choice for themselves. The end goal of the participants was to minimize their nicotine symptoms; therefore, we believe that participants engaged in exploration as they attempted to reduce the uncertainty and learn more about the infrequent option (Mehlhorn et al., 2015). Another explanation for this contrasting result could be because the initial bias may not have been strong enough to influence their free sampling preference. Since 64 per cent of the participants in our study had been and were current students of psychology, they may have recognized or counteract the study's underlying mechanism, thus shifting their preference.

Further, we assume that introducing the lung image in the free sampling phase, which indicated lung health after each trial as feedback, gave the participants more clarity of their sampling outcome. In the initial phase, the participants did not see the lung image as feedback. However, in the free sampling phase, participants may have attempted to sample the two treatment options while focusing on the changing colors of the lung as a deciding factor to their goal of minimizing nicotine symptoms. Thus, causing a shift in preference and attenuating towards chance level. Another explanation could be that the participants perceived the free sampling phase as a new task. Although this was emphasized in the instructions, we assume that the participants could have formed a new belief regarding the treatment options at the beginning of the free sampling phase. This could explain why they engaged in exploration behaviour in sampling as the trial's progress. Another explanation for the opposite sampling behaviour to the expected exploitation behaviour could be the recency effect. The recency effect is the tendency for a person to retain the most recent piece of information, i.e., the ending of a list or task (Troyer, 2011). In our case, the most recent information would be the last trials of the induction phase before beginning the free sampling trials. Therefore, if the participants saw the infrequent option more often in the last few trials of the induction phase, they might have weighed more heavily on that option; thus, their initial biases attenuated, and they sampled both options in the free sampling phase instead of relying on one treatment option. Furthermore, we found that the post preference estimate was

not significant in our results, thus confirming that the initial bias did not persist in the free sampling phase.

Finally, we tested the intervention to draw inferences regarding the updating of beliefs. We assumed that when we instruct the participants in the equal group to try sampling both the treatment options equally often, their initial biases will attenuate as they sample further in the experiment. However, in our sample, we observed a contradicting result to our hypothesis and the assumption made by Harris et al. (2020) was supported. We observe that initially, the participants sampled both the treatment options equally often and engaged in exploration behaviour as we expected; however, halfway in free sampling, their preference shifts, and they begin deviating towards their initial bias in the induction phase. Hence, they start sampling the frequent option more often and engage in exploitation behaviour. The curve (refer to figure 1) indicating a change in treatment preference is very interesting as we expected to obtain the opposite results. Therefore, we believe that participants in the equal group returned to their initial bias because they perceived the frequent option as favourable, and exploration may not have yielded the feedback that matched their goal of minimizing nicotine symptoms. This could be negative in situations because first impressions are not always correct. Thus, when choosing between treatments, people may be stuck with a suboptimal option. In a health context, decision making based on falsely upheld beliefs and first impressions can keep people from making the optimal choices.

Limitations

The results we obtained did not fully support the hypothesis of our study. Here, we discuss some limitations of the study that may have contributed to a different result. Firstly, we did not see an initial bias persisting in the free sampling phase; this could be due to boredom. Repeatedly choosing between treatments for 100 trials could be a tedious task. This may have led to a change in preference, causing participants to explore rather than exploit in the pursuit to ease boredom (Mehlhorn et al., 2015). It could also be possible that the participants deviated from the bias and explored because the participants were able to recognize the contingency design of the study (Harris et al. 2020). Although we do not have any evidence for this, we only assume that they could have recognized the task and thus engaged in exploring rather than exploiting the treatment. Secondly, the study design was hypothetical and asked participants to imagine that they were smokers; however, not everyone smokes. Hence, the experimental scenario may not have been applicable enough for the participants to relate to. Lastly, the participants in the study may not have been motivated

enough to engage with the experiment fully. In our experiment, we did not provide a monetary incentive for participation and performance, whereas participants received a financial reward in the study by Harris et al. (2020). This could have influenced the motivation for participation in the study and affected the performance in this repeated sampling task.

Future research

Although there were some limitations in the study and not all the results were as desired, the study holds importance in increasing the generalizability of the subject matter under investigation. Future research should improve and further explore pseudocontingency and its role in future decision making in different health contexts. The first recommendation for future research is to use a different design to explore contingency behavior in a health context. This could be done by employing participants who have an underlying health issue, such as real smokers. This will give more credibility to the research, and the results could help understand their choice behavior better. It could also be informative and helpful to devise real interventions to counter this unhealthy health habit and choice behavior based on false beliefs. Secondly, we recommend testing the study in a reduced number of trials but with similar validity and quality of research. Reducing the number of trials will help gain insight into sampling behavior as it will eliminate the boredom factor in the study. Thirdly, in our intervention test, we asked the participants to sample only the first ten trials equally often; this may not have impacted the sampling behaviour as we expected. Hence, future research should ask the participants to sample a minimum of twenty samples equally often. Lastly, we could recommend that the study's statistical analysis be done in line with the calculation used by Harris et al. (2020), namely including the base rate calculations. Other relevant statistical tests which can provide in-depth detail into the participant sampling performance and behavior is worth exploring for future research.

Conclusion

The current research explores the influence of first impressions and biases in smoking cessation treatment options. We found that an initial bias was induced. However, the pseudocontingency bias did not persist in the free sampling phase, and the biases attenuated. Further, as claimed by Harris et al. (2020), participants in our study upheld their belief in a reward rich environment when sampling the two treatment options. Although the initial biases were upheld, we believe that these results do not essentially mean that the biases and beliefs cannot be updated in the future. Interventions to understand updating false beliefs

need to be explored in other health contexts to gain a deeper insight into the influence of pseudocontingency bias in similar frameworks. Although not all the results were as expected, we cannot draw precise conclusions about the pseudocontingency and sampling behavior in a reward-rich environment.

References

Allan, L. G. (1980). A note on measurement of contingency between two binary variables in judgment tasks. Bulletin of the Psychonomic Society, 15, 147–149. http://dx.doi.org/10.3758/BF03334492

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. Philosophical Transactions of the Royal Society B: Biological Sciences, 362(1481), 933–942. https://doi.org/10.1098/rstb.2007.2098

Del Campo, C., Pauser, S., Steiner, E., & Vetschera, R. (2016). Decision making styles and the use of heuristics in decision making. *Journal of Business Economics*, *86*(4), 389-412. https://doi.org/10.1007/s11573-016-0811-y

Dennis, M. J., & Ahn, W. K. (2001). Primacy in causal strength judgments: The effect of initial evidence for generative versus inhibitory relationships. Memory and Cognition, 29(1), 152–164. https://doi.org/10.3758/BF03195749

Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. Behavior research methods, 39(2), 175-191. https://doi.org/10.3758/BF03193146

Fiedler, K. (2000). Beware of samples! A cognitive-ecological sampling approach to judgment biases. Psychological Review, 107(4), 659–676. https://doi.org/10.1037/0033-295X.107.4.659

Fiedler, K. (2010). Pseudocontingencies can override genuine contingencies between multiple cues. Psychonomic Bulletin & Review, 17(4), 504–509. https://doi.org/10.3758/PBR.17.4.504

Fiedler, K., & Wänke, M. (2009). The cognitive-ecological approach to rationality in social psychology. *Social Cognition*, *27*(5), 699-732.10.1521/soco.2009.27.5.699

Fiedler, K., Freytag, P., & Meiser, T. (2009). Pseudocontingencies : An Integrative Account of an Intriguing Cognitive Illusion. Psychological Review, 116(1), 187–206. https://doi.org/10.1037/a0014480

Geana, A., Wilson, R., Daw, N. D., & Cohen, J. (2016). Boredom, Information-Seeking and Exploration. In *CogSci*.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, *62*, 451-482. https://doi.org/10.1146/annurev-psych-120709-145346

Hardy, J. H., Day, E. A., & Arthur, W. (2019). Exploration-exploitation trade-offs and information-knowledge gaps in self-regulated learning: Implications for learner-controlled training and development. Human Resource Management Review, 29(2), 196–217. https://doi.org/10.1016/j.hrmr.2018.07.004

Harris, C., Fiedler, K., Marien, H., & Custers, R. (2020). Biased preferences through exploitation: How initial biases are consolidated in reward-rich environments. *Journal Of Experimental Psychology: General*, *149*(10), 1855-1877. https://doi.org/10.1037/xge0000754

Holt, C. A., & Laury, S. K. (2002). Risk Aversion and Incentive Effects. In The American Economic Review (Vol. 92, Issue 5, pp. 1644–1655). https://doi.org/10.1257/000282802762024700

Humphreys, K. L., Lee, S. S., Telzer, E. H., Gabard-Durnam, L. J., Goff, B., Flannery, J., & Tottenham, N. (2015). Exploration-exploitation strategy is dependent on early experience. Developmental Psychobiology, 57(3), 313–321. https://doi.org/10.1002/dev.21293

Humphries, M. D., Khamassi, M., & Gurney, K. (2012). Dopaminergic Control of the exploration-exploitation trade-off via the basal ganglia. Frontiers in Neuroscience. https://doi.org/10.3389/fnins.2012.00009

Kim, C., Song, J., & Nerkar, A. (2012). Learning and innovation: Exploitation and exploration trade-offs. Journal of Business Research, 65(8), 1189–1194. https://doi.org/10.1016/j.jbusres.2011.07.006

Krypotos, A.-M., Crombez, G., Alves, M., Claes, N., & Vlaeyen, J. W. S. (2021). The exploration-exploitation dilemma in pain. In Pain: Vol. Publish Ah (Issue May). https://doi.org/10.1097/j.pain.0000000000002352

Kusev, P., Purser, H., Heilman, R., Cooke, A. J., Van Schaik, P., Baranova, V., ... & Ayton, P. (2017). Understanding risky behavior: the influence of cognitive, emotional and hormonal factors on decision-making under risk. Frontiers in psychology, 8, 102. https://doi.org/10.3389/fpsyg.2017.00102

Kutzner, F., & Fiedler, K. (2017). Stereotypes as pseudocontingencies. European Review of Social Psychology, 28(1), 1–49. https://doi.org/10.1080/10463283.2016.1260238

Mather, M., Shafir, E., & Johnson, M. K. (2000). Misremembrance of options past: Source Monitoring and Choice. Psychological Science, 11(2), 132–138. https://doi.org/10.1111/1467-9280.00228

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the Exploration-Exploitation Trade-off: A Synthesis of Human and Animal Literatures. Decision, 2, 191–215. https://doi.org/http://dx.doi.org/10.1037/dec0000033

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M., Morgan, K., Braithwaite, V. A., Hausmann, D., Fiedler, K., & Gonzalez, C. (2015). Unpacking the Exploration-Exploitation Trade-off: A Synthesis of Human and Animal Literatures. Decision, 2, 191–215. https://doi.org/http://dx.doi.org/10.1037/dec0000033

Meiser, T., & Hewstone, M. (2010). Contingency learning and stereotype formation: Illusory and spurious correlations revisited. European Review of Social Psychology, 21(1), 285–331. https://doi.org/10.1080/10463283.2010.543308

Meiser, T., Rummel, J., & Fleig, H. (2018). Pseudocontingencies and choice behavior in probabilistic environments with context-dependent outcomes. Journal of Experimental Psychology: Learning Memory and Cognition, 44(1), 50–67. https://doi.org/10.1037/xlm0000432

Morris, L. S., Baek, K., Kundu, P., Harrison, N. A., Frank, M. J., & Voon, V. (2016). Biases in the Explore-Exploit Trade-off in Addictions: The Role of Avoidance of Uncertainty. Neuropsychopharmacology, 41(4), 940–948. https://doi.org/10.1038/npp.2015.208

Raymond S. Nickerson. (1998). Confirmation bias: a ubiquitous phenomenon in many guises. Review of General Psychology, 2(2), 175–220.

Sullivan, J. (2019). The Primacy Effect in Impression Formation: Some Replications and Extensions. Social Psychological and Personality Science, 10(4), 432–439. https://doi.org/10.1177/1948550618771003

Todd, P. M. (2000). Todd-Gigerenzer-Précis of Simple heuristics that make us smart-2000. 727–780. https://doi.org/10.1017/S0140525X00003447

Troyer A.K. (2011) Recency Effect. In: Kreutzer J.S., DeLuca J., Caplan B. (eds) Encyclopedia of Clinical Neuropsychology. Springer, New York, NY. https://doi.org/10.1007/978-0-387-79948-3_1146

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. Cognitive Psychology, 5(2), 207–232. https://doi.org/10.1016/0010-0285(73)90033-9

Vogel, T., & Kutzner, F. (2017). Pseudocontingencies in Consumer Choice: Preference for Prevalent Product Categories Decreases with Decreasing Set Quality. Journal of Behavioral Decision Making, 30(5), 1193–1205. https://doi.org/10.1002/bdm.2034

Waugh, N. C., & Norman, D. A. (1965). Primary memory. Psychological Review, 72(2), 89–104. https://doi.org/10.1037/h0021797