First Impressions in the Health Context: The Generalizability of Initial Biases on Sequential Drug Choices

Florine M.V. Kist
Department of Social Sciences, Utrecht University
Social, Health and Organisational Psychology
Supervisor: Harris, C.A.

Word count: 8968

June 22, 2020

The paper may be made publicly accessible

Abstract

Previous research has shown that false first impressions can guide subsequent decision making in the form of biased sampling in which people exploit seemingly better alternatives (Harris et al., 2020). The present research tests the hypothesis that pseudocontingency biases towards a certain medicine will be maintained and generalized onto other similar, yet different medicines in a free sampling context. The current research demonstrates that in reward-rich environments, pseudocontingencies between medicine options and health outcomes guide subsequent decision making so that the frequently presented medicine is preferred more than the infrequently presented option. In line with our expectations, people appear to use a strategy of exploitation, however, this was only evident at the beginning of the sampling phase. When providing participants with a virtually identical alternative of the preferred medicine, the previous bias spills over into their sampling of this new medicine. However, explicit preference measures reveal that participants do not prefer this new medicine more than the alternative. The present findings provide evidence that pseudocontingency biases may transfer onto similar alternatives, however, it is still unclear whether this transfer demonstrates generalization. Finally, it cannot be concluded whether the bias towards the new alternative shows a strategy of exploitation or exploration.

Keywords: pseudocontingency, exploration-exploitation trade-off, drug choice, initial evidence, generalization, categorization

Introduction

Imagine that you are a patient suffering from hay fever that has to choose a medicine to relieve your complaints. One strategy is to choose an option that you have tried before and experienced to yield a positive effect. This process is known as exploitation. The benefit is that you know from experience that this medicine will relieve your complaints. The pitfall here is that you might overexploit this medicine without investigating other medicines that may yield even more positive outcomes, such as a quicker recovery. As an alternative strategy, you could choose to switch options in order to investigate whether other medicines could provide even more rewarding effects. This strategy is called exploration. The advantage of exploration is that by actively exploring which medicine works best, your chances of finding the best medicine may increase. The downside is that through this process you might come across medicines that have uncertain rewards and variable values, which is not just costly in terms of time, but also in terms of health.

The scenario described above demonstrates an interesting tradeoff that is at core in every decision made in environments in which there are multiple options to choose from. This tradeoff is known as the exploration-exploitation tradeoff, in which people strive to find the right balance between exploiting options and exploring alternatives (Mehlhorn et al., 2015). In this tradeoff, reward maximization on the one hand and information search on the other hand are competing strategies to optimize the decision-making process.

In the current study, the trade-off between exploration and exploitation of medicine choices will be investigated. Specifically, the research will focus on the maintenance of initial beliefs with regards to two types of medicines, and whether people's first impression of a certain medicine will guide their subsequent decisions. More precisely, the study aims to explore to what extent false initial beliefs are maintained and whether this type of biased exploitation will be transferred onto different, yet similar types of medicines. This research will rely on the foundations of pseudocontingency literature (Fiedler et al., 2009; Kutzner et al., 2008; Harris et al., 2020) as a framework for investigating the effect of initial biases on future decision-making processes.

Exploitation, Heuristics and the Rise of Biases

In trading off exploration and exploitation, the key is to find the right balance (Cohen et al., 2007). However, one can imagine how exploiting a seemingly superior alternative overrules the intent to explore other potentially better alternatives. Although exploitation can be the right strategy to maximize outcomes in simple environments in which the relationship between actions and outcomes is known (Sims et al., 2013), most of human decision making takes place in much more complex environments in which humans have limited time, knowledge and resources to make decisions (Simon, 2000; Gigerenzer & Selten, 2002). In such complex environments, the relationship between actions and future outcomes is uncertain (Sims et al., 2013). Our decisions oftentimes have delayed effects and indirect consequences, and therefore require a different type of approach.

In complex choice environments like the world today, the relationship between choice and outcome must be learned from experience (Sims et al., 2013). Contingency learning, the human ability to detect contingencies between events (De Houwer & Beckers, 2002), plays an essential role in enabling people to explain the past, control the present and predict the future (Crocker, 1981). Contingency learning allows us to predict and control events in the environment by teaching us which events predict or cause the presence or absence of other events (De Houwer & Beckers, 2002). However, when people have experienced an option to be beneficial, their reliance on the positive contingency between that option and the positive outcome may tempt them to exploit this option prematurely (Mehlhorn et al., 2015). This exploitation becomes specifically problematic in cases when these initial judgements of experiences are biased, which may happen when people rely on heuristics.

Heuristics are inevitable mental shortcuts that people use to guide their behavior (Shah & Oppenheimer, 2008). In environments that do not allow us to always make fully informed judgments (Gigerenzer & Gaissmaier, 2011), people oftentimes make use of heuristics instead (Gigerenzer & Selten, 2002; Kahneman, 2011; Thaler & Sunstein, 2008). By ignoring part of the information, heuristics allow decision making to be more quickly, frugally and or accurately than more complex methods, such as statistics (Gigerenzer & Gassmaier, 2011). Despite their everyday usefulness, heuristics also have a downside, namely the rise of biases (Tversky &

Kahneman, 1974). A well-known example of this is when people rely on the representativeness heuristic and estimate the likelihood of a person A belonging to a category B by drawing on how similar they consider A to be to their stereotype of B (Tversky & Kahneman, 1974). This heuristic also explains misperceptions in decision making (Kahneman & Tversky, 1972), such as the gambler's fallacy. When one option yields positive outcomes upon several subsequent runs, people do not perceive this series of good outcomes as representative of the statistical probability, leading them to switch options. Both examples show that despite the potential usefulness of heuristics on some occasions, it can also lead to misperceptions of patterns resulting in biased judgments and biased choices.

Pseudocontingencies

One field of research that investigates this phenomenon of biased judgements that are heuristically inferred, is pseudocontingency research. Pseudocontingencies are logically unwarranted inferences of contingencies between two variables from information other than observations of these variables in a specific ecology (Fiedler et al., 2009). More precisely, pseudocontingencies arise in the presence of skewed samples of both the different options and the different outcomes. When there are more samples of one option than of another option and one type of outcome is more frequent than another type of outcome, people may falsely perceive contingencies between options and outcomes by aligning their base rates. Even when a contingency is nonexistent and two options are on average equally rewarding (Fiedler, 2000), people may perceive one option to be more favorable than another, merely because it is presented more frequently. As has been described by Kutzner, Freytag, Vogel and Fiedler (2008), the judgement heuristic that people rely on is "if the frequent (rare) level is observed in the one variable, then the frequent (rare) level is likely to be observed in the other variable as well" (p. 3). Although pseudocontingencies are not wrong per definition (Fiedler et al., 2013; Klauer, 2015; Kutzner et al., 2011), they can be misleading when the pseudocontingency is not in concordance with the actual contingency, as the prefix 'pseudo' discloses. As useful as this mental shortcut may seem, it can lead to biased judgments and decisions.

To date, pseudocontingency research has predominantly explored the mechanisms behind pseudocontingencies and the conditions under which they are likely to arise (Fiedler et al., 2009).

Pseudocontingency literature hardly says anything, though, about how we interact with the world around us based on these initial beliefs. How are pseudocontingency inferences held onto or updated as we continue interacting?

Only recently have these questions been investigated by researchers interested in the effects of pseudcontingency inferences on actual, sequential choice behavior (Meiser et al., 2018; Bott & Meiser, 2020; Harris et al., 2020). As Meiser et al. (2018) demonstrated, after a learning phase in which participants were presented with skewed samples, pseudocontingencies between options and outcomes were inferred and participants were influenced by these false contingencies in their subsequent decision behavior. Likewise, the paper by Bott and Meiser (2020) provides additional support for the maintenance of pseudcontingency biases on following decision making. An interesting addition by Meiser and Bott (2020) is that they compare the extent to which these false preferences are maintained in either predominantly positive or predominantly negative contexts. Similarly, Harris et al. (2020) make this distinction by referring to reward-rich and reward-impoverished environments. In reward-rich environments, the absolute base rates of positive outcomes are high for both alternatives. Initial erroneous contingencies are less likely to be corrected, as people are generally likely to receive a positive outcome regardless of their choice. Simultaneously, the temptation to exploit seemingly better alternatives is reinforced over and over again as they continue sampling. In other words, "the reward structure fosters exploitation" (Harris et al., p. 2, 2020). In reward-impoverished environments, where the absolute base rates of positive outcomes are low, exploration is rendered since earlier inferences are more likely to be corrected upon sampling.

Biases in Experience Sampling

Previous pseudocontingency research conducted in the field of social psychology has shown that the process of experience sampling can be inherently biased. A study examining pseudocontingency effects of positive versus negative interactions showed that despite basing one's impressions on all available, all accurately interpreted and remembered observations, the setup of experience sampling in social interactions allows for impressions to nevertheless be biased (Denrell, 2005). This was demonstrated by the finding that in case of a positive interaction, people exploited this option by continuing to interact with the other person, whereas

a negative interaction stopped someone from interacting. Because of this systematic bias in experience sampling, false negative initial impressions are unlikely to be corrected as negative first impressions are not being put to the test. Positive experiences, on the contrary, will encourage future interactions, which will lead to repeated sampling. This statement supports the claim by Harris et al. (2020) that the reward structure of reward-rich environments fosters exploitation.

Rather than focusing on social interactions, the research by Harris, Fiedler, Marien and Custers (2020) involves an experience sampling task in which participants repeatedly sample between two object options. Their paper will form the basis for the current thesis because it explores not just the rise, but also the maintenance of biases on successive decision making. As briefly discussed before, the study by Harris et al. (2020) tested the maintenance of initial pseudocontingency biases over multiple trials and found that false initial biases are maintained in reward-rich environments, leading to the exploitation of the favored option. They explain this as a dysfunctional interaction between a primacy effect of the first evidence and positive outcomes encountered in reward-rich environments and claim that any type of initial bias can lead to exploitation in reward-rich environments.

The findings by Harris et al. (2020) may explain many real-life instances in the health context. With regards to drug choice, for instance, patients may falsely believe one drug to be superior in alleviating one's symptoms, based on a positive first impression. Instead of exploring other medicines and testing which drug works best for them at which instance, they may exploit the medicine that they deem most rewarding. Exploitation in the context of drug choice is alarming. Not only does exploitation hinder the potential search for other, better alternatives, it may also lead to serious health hazards when people rely on the wrong information. When someone continuously and over a longer period of time uses a medicine that does not optimally relieve their symptoms, they may not only suffer for longer, or face a stagnation in terms of their health, but they may also find their health worsening when this medicine simply is not good enough.

Generalization of Biases

One of the most interesting and novel claims made by Harris et al. (2020) is that any initial bias in a reward-rich environment may be maintained over a period of time. Since their paper is one of the first to investigate the maintenance of initial pseudocontingency biases, it is relevant to replicate these findings to see to what extent their claim is generalizable. However, this research aims to go beyond mere replication, and sets out to explore under which conditions these biases are likely to be maintained. Specifically, the current study will investigate whether a bias towards a medicine will be maintained once people are presented with another similar, yet novel medicine from the same category. The hypothesis that a bias towards a medicine will be generalized onto another similar medicine draws on the claim by Kutzner et al. (2008) that humans tend to base their predictions on perceived similarity. If it is the similarity of the base rates that drives predictions, then it is a plausible hypothesis that the similarity of variable properties can also drive predictions.

In order to substantiate this new avenue of research, we will briefly address a chapter that summarizes a large body of attitude generalization literature (Fazio et al., 2015). One of the core claims of this paper is that attitude generalization of novel stimuli depends heavily on valence weighting, on weighting the positive versus negative features of the stimulus in question. In the process of valence weighting, people assess how much the new stimulus resembles past instances that proved to be positive versus past occurrences that yielded negative attitudes. This finding seems to be in concert with pseudocontingency findings related to the maintenance of biases in reward-rich versus reward-impoverished environments. According to this argument, people will generalize their positive attitudes towards novel objects in reward-rich environments, as long as they deem the options to be sufficiently similar. Therefore, it is plausible to hypothesize that biased positive attitudes towards one medicine will be transferred onto a very similar, novel medicine in a reward-rich environment.

The Current Paper

The present paper will both rely and elaborate on the paper by Harris et al. (2020) in three principal ways. Firstly, this research aims to replicate their findings with regards to the instigation of initial biases through pseudocontingencies in the health context. Transferring their

exact methods to this new context will expand the generalizability of their claims, such as their statement that any type of initial bias in reward-rich environments will lead to exploitation. The first research question addressed in this paper therefore is: “To what extent do pseudocontingency biases arise in the health context?”.

Secondly, the current paper will also further explore the boundaries of their finding that initial biases are maintained in reward-rich environments. Presupposing that biases will arise in the health context; the next research question is: “To what extent do first impressions of a certain medicine influence later choice behavior?”. As Harris et al. (2020) found that biases guide subsequent decisions, it is expected that biased first impressions of medicines will also guide subsequent drug choices.

Thirdly, this research will give a novel insight into the transferability of initial biases onto other similar, yet distinct alternatives. This additional avenue of research aims to answer the following research question: “If biased first impressions of medicines are maintained; will they also generalize to other similar yet different medicines within the same category?”. As the experiment takes place in a reward-rich environment, positive pseudocontingencies are expected to arise and the process of valence weighting is expected to be positive (Fazio et al., 2015), therefore making attitude categorization likely to happen.

Methods & Design

Methods

This experiment set out to validate the three main hypotheses underlying this paper; that pseudocontingencies induce a bias towards a certain medicine, that this bias will be maintained over several trials, and that this bias will generalize to other, similar medicines within the same category. The sample size was estimated to be approximately 100, based on power analyses using G*Power (Faul et al., 2007). These calculations rely on a 5% alpha-level, 80% statistical power, and effect sizes between $\eta_p^2 = .081$ and $\eta p2 = .270$ as used by Harris et al. (2020) and reported by Meiser et al. (2018).

Participants for this experiment were recruited via an online database, Prolific Academy (https://www.prolific.co). As the experiment was run in English on Soscisurvey (https://www.soscisurvey.de), all participants were required to be fluent in English. Additionally, they were pre-screened on the basis of their age (18-55), approval rates (95-100%) and previous studies that they participated in, to exclude those individuals that had already taken part in previous experiments by Harris and colleagues. 110 participants[1] (N_{female} = 56) with an average age of 28 years (*SD* = 8.52) took part in the experiment for a financial reward of £1 per 10 minutes. Over 83% of participants had an educational degree of College, A levels or higher. This research was conducted according to the guidelines of the Ethics Review Board of the Faculty of Social and Behavioral Sciences at Utrecht University.

Design

The current research relies on the same design as used by Harris et al. (2020), namely a two-armed bandit task in which participants repeatedly choose between two medicines in a reward-rich environment. They were instructed to imagine themselves being patients suffering from hay fever having to choose between two types of medicines in an online pharmacy. Choosing either medicine Hydrox or Sofrix resulted in a positive health-improving outcome or a negative outcome in which symptoms were not relieved. In addition to this behavioral data of participants' choice behavior, attitudinal data of their medicine preferences was collected.

Procedure

A repeated measures design was used for this experiment. It was divided into six phases: an induction phase, a first preference estimate phase, a first free sampling phase, a second preference estimate phase, a second free sampling phase and a third preference estimate phase.

In the induction phase, participants were told that the program would randomly determine which medicine was to be chosen in order to get participants familiar with the task. Participants thus were unable to choose freely, and they only saw a picture of the one available option. After clicking on this picture, participants received feedback about their decision both verbally and visually. A positive outcome resulted in a reinforcing message including the text "you chose [medicine name] and it helped" and a smiley face, whereas a negative outcome resulted in a

[1] Due to issues with the recruitment platform Prolific the last 33 participants were recruited about a week later. There is no reason to assume this to have affected the outcomes.

discouraging message reading “you chose [medicine name] but it had no effect” and a frowning face. After a delay of a second, this feedback vanished and the next choice was presented. To allow participants to know the remaining number of trials they still had to either exploit or explore, the scope of the sampling phase was made salient by presenting the current trial number (“Trial: x/100”) on the display throughout the entire experiment. Since the experiment took place in a reward-rich environment, participants were presented with 9 wins and 3 losses for one option, making this the frequent option, and 3 wins and 1 loss for the other, infrequent option. Therefore, there was a 75% chance to receive a positive reward for both options over the course of 16 forced trials.

In the second phase of the experiment, participants were asked to indicate their relative preference estimates, base rate estimates and conditional estimates including confidence estimates. From this point onwards, this phase will be referred to as the first preference measures, P.M.1. To start with the relative preference estimates, participants were asked which medicine was more likely to relieve their symptoms. This question was answered by moving a slider which was anchored with the image of Hydrox at one end, and Sofrix at the other. For the base rate estimates, participants moved sliders to indicate how likely it was that they selected each medicine and how likely it was that their symptoms were (not) relieved. In addition, participants also gave conditional estimates for both medicines regarding how likely it was that their symptoms were relieved when they chose either one of the medicines. Finally, participants had to indicate how confident they were in making a reasonable judgment regarding their conditional estimates of each medicine. To this end, the anchor was adjusted at *not confident at all* and *very confident*. For all estimates except the relative preference measure, each scale indicated the slider marker’s current position in percentages, which were updated as the slider was moved.

At the beginning of the third phase, we introduced participants to a stick figure that visually represented their health progress by changing colors from red to orange to green depending on the outcomes. This figure was added to personalize the hypothetical scenario and to make people more motivated to make good decisions. To encourage this even more, we programmed the figure so that one positive outcome would instantly make the figure green

again, even if it was red before, thus emphasizing the effect of positive outcomes over negative ones.

Following this information, participants were free to choose either one of the medicines on any of the remaining 42 trials. Important to note is that both medicines were equal in terms of their probability of having positive outcomes, namely 75% of the time. In contrast to the induction phase in which the distribution of the outcomes was manipulated, the outcomes in the free sampling phase were randomly determined on each trial. Thus, across all trials the two options were completely equal. Upon completion of this free sampling phase, in the fourth phase (second preference measures, P.M.2), participants answered the exact same questions related to their preference measures as during P.M.1.

At the start of the fifth phase, participants were told that Hydrox (Sofrix) was no longer available due to high concentrations of pollen this year. However, they were also told that a similar medicine from the same manufacturer, Hydrax (Sofrex), was available. They were informed that this medicine contained the same active ingredients as the previous medicine. After reading this information, participants could now choose between medicines Hydrax (Sofrex) and Sofrix (Hydrox) over the course of 42 free sampling trials. In the sixth phase (third preference measures, P.M.3), the same preference questions were asked as before only this time with the new replacement medicine.

Finally, participants had to answer three questions related to the generalizability of the new medicine. Specifically, they moved sliders to indicate the extent to which they perceived Hydrax (Sofrex) to be a good alternative, to be effective in relieving their symptoms and to be similar to Hydrox (Sofrix).

Data Preparation

The current experiment made use of counterbalancing in order to avoid systematic biases for both the type of medicine that was presented and the side this medicine was displayed on. Therefore, in preparation of the analyses, all slider values related to questions about either the side or the medicine were modified and recoded. Specifically, they were transformed in such a way that zero always presents neutrality between both options and all positive values represent the option that was shown more frequently in the induction phase. To illustrate, when Hydrox

was presented more frequently in the induction phase, any positive value on the relative preference estimate indicates a bias towards Hydrox and any positive value on the conditional estimates demonstrates a positive contingency between Hydrox and the likelihood of symptoms being relieved.

Data preparation and analyses were undertaken using SPSS (IBM SPSS Statistics for Apple, Version 26). Across all three measures of preference - the sampling decisions, the relative preference estimates and the conditional estimate measures - biases are expected, and therefore one-sided t-tests were performed.

In order to analyze the behavioral measures, a mean sample score of all trials was created and recoded in order to create a choice index of participants' overall preference. This score ranged from 0 to 1, in which a score closer to 1 described a preference towards the frequent option, and a score closer to 0 represented no bias. Confidence intervals for the effect sizes were added.

Results

Sampling

Over the 42 trials of the first free sampling phase, participants chose the frequent option on average 58% (*SD* = .27) of the time. In the second free sampling phase, during the last 42 trials, they chose the frequent option 55% (*SD* = .25) of the time. Figure 1 shows separate trial analyses of the average sampling behavior per trial, indicating the percentage of participants sampling the frequent option during both sampling phases. Across all trials this is significantly above chance level for both the first free sampling phase ($t(109) = 2.96$, $p = .002$, $d = .28$, 95% CI_d [0.09, 0.47]), as well as the second free sampling phase ($t(109) = 2.32$, $p = .01$, $d = .22$, 95% CI_d [0.01, 0.10]).

Figure 1

Average frequent option choices per trial

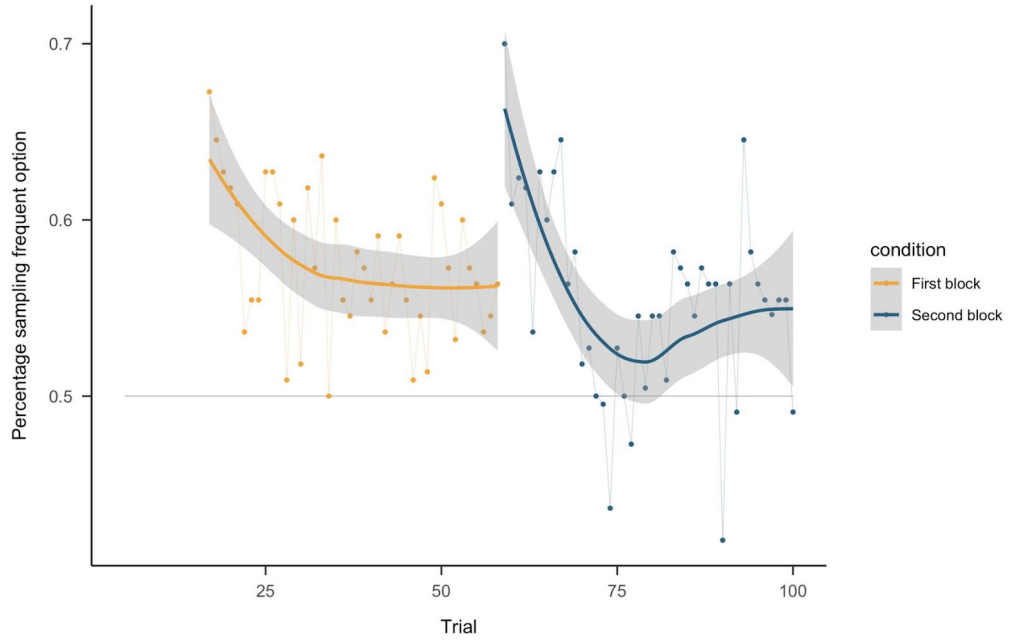

Note. The yellow curve represents the trials during the first free sampling phase, the blue curve represents the trials during the second free sampling phase in which the new medicine was sampled.

Relative preference estimates

In addition to the behavioral sampling measures, participants' preferences were also measured by relative preference estimates that people made before and after finishing sampling. Before the first sampling phase, at P.M.1, participants had a preference towards the frequent medicine 61.45% (*SD* = 28.96) of the time. This preference was 58.29% (*SD* = 33.56) at P.M2, and only 54.00% (*SD* = 29.57) at P.M.3. Preference scores indicated a significant effect at both P.M.1 ($t(109) = 3.78$, $p < .001$, $d = .36$, 95% CI_d [0.16, 0.55]) and P.M.2 ($t(109) = 2.28$, $p = .012$, $d = .22$, 95% CI_d [0.03, 0.41]). A bias towards the frequent option thus was found right after the initial evidence as well as after the first free sampling phase. The preference scores at P.M.3 were not significant, $t(109) = 1.06$, $p = .15$, $d = .10$, 95% CI_d [-0.09, 0.29].

Conditional estimates

Before and after the two sampling phases, participants' estimates for the likelihood of receiving a positive outcome upon choosing either the frequent or infrequent medicine were recorded. From these conditional estimates difference scores were calculated (ΔP; Allan, 1980) that indicate the perceived contingency between medicine options and outcomes in participants' estimates. At P.M.1 (frequent medicine: $M = 54.58$, $SD = 24.26$; infrequent medicine: $M = 45.22$, $SD = 25.09$) this resulted in a mean Δ*P*-score of $\Delta P = .09$ ($SD = 0.27$). At P.M.2 (frequent medicine: $M = 55.06$, $SD = 27.58$; infrequent medicine: $M = 47.67$, $SD = 26.08$), this resulted in a mean Δ*P*-score of $\Delta P = .07$ ($SD = 0.41$). At P.M.3 (frequent medicine: $M = 51.55$, $SD = 27.56$; infrequent medicine: $M = 47.78$, $SD = 26.62$), this resulted in a mean Δ*P*-score of $\Delta P = .04$ ($SD = 0.38$).

A pattern similar to the results observed for the relative preference estimates was found, with estimates yielding significant results at both P.M.1 ($t(109) = 3.65$, $p < .001$, $d = .35$, 95% CI_d [0,15, 0,54]) and P.M.2, $t(109) = 1.91$, $p = .03$, $d = .18$, 95% CI_d [-0,01, 0,37]. Conditional estimates at P.M.3 did not reach significance, $t(109) = 1.04$, $p = .15$, $d = .10$, 95% CI_d [-0,09, 0,29].

Confidence estimates

Overall, the mean of confidence estimates ranged from 44.98% ($SD = 24.46$) for the infrequent medicine at P.M.1 to 55.10% ($SD = 28.57$) for the frequent medicine at P.M.2. The only confidence score that yielded a significant effect was the estimate for the infrequent medicine at P.M.1, $t(109) = -2.58$, $p = .01$, $d = -.25$, 95% CI_d [-0,06, -0,44]. Using ΔP-values of the confidence estimates, bias was found in people's confidence estimates at both P.M.1 ($t(109) = 4.44$, $p < .001$, $d = .42$, 95% CI_d [0,23, 0,62]) and P.M.2 ($t(109) = 2.07$, $p = .02$, $d = .20$, 95% CI_d [0,01, 0,39]). Confidence estimates were not significantly skewed at P.M.3, $t(109) = .32$, $p = .38$, $d = .00$, 95% CI_d [-0,16, 0,22]).

Base rate estimates

Base rate estimates at P.M.1 disclose that participants perceived the initial evidence as rather skewed. On average, they estimated to have encountered the frequent medicine 67.95% ($SD = 14.43$) of the time but the infrequent medicine only 35.30% ($SD = 14.11$) of the time.

Likewise, they estimated to have encountered the frequent outcome more often (*M* = 74.47, *SD* = 14.47) than the infrequent outcome (*M* = 27.43, *SD* = 13.49). At P.M.2, participants perceived the evidence to be less skewed, as the frequent medicine was chosen 59.77% of the time (*SD* = 27.72) and the infrequent medicine was picked only 43,47% of the trials (*SD* = 28,12). The perceived win-loss ratio was slightly more skewed at P.M.2 than at P.M.1, as they estimated their likelihood of winning 78.14% (*SD* = 10.67) and their likelihood of losing 24.52% (*SD* = 11.82). At P.M.3, participants chose the frequent medicine 54.50% (*SD* = 27.53) of the time and the infrequent medicine 48.45% (*SD* = 28.10) of the trials. They estimated their chances quite similarly to those in the previous sampling phase, namely as having won 75.72% (*SD* = 12.47) of the time and their chances of having lost 27.02% (*SD* = 13.95). Participants' perceptions thus are near perfect for the outcomes, but regressive for the options.

Additionally, log-transformed base rate ratio scores were created from the original base rate estimates in order to quantify the perceived skewness of participants' estimates. Strongly perceived skewness of base rates in the same direction were demonstrated by larger log scores, whereas scores around zero signify no skew on either variable and negative log scores indicate skews in the opposite direction. The log-transformations were created using the following formula (Kutzner, 2009) that was used by Harris et al. (2020):

$$log_{BR} = log_{10}\left(\frac{ab}{cd}\right) \times log_{10}\left(\frac{ac}{bd}\right)$$

In this formula, ab, cd, ac and bd are the base rates of a regular four-cell contingency table, ab = a + b.

In line with the perceived skewness of the initial evidence as demonstrated by the mean scores at P.M.1, the transformed log score (*log* = 0.16, *SD* = .18) also yielded a significant effect, $t(109) = 9.26$, $p < .001$, $d = .88$, 95% CI_d [0,66, 1,10]. At P.M.2, this effect for the log score (*log* = .14, *SD* = .60) was significant as well, $t(109) = 2.37$, $p = .01$, $d = .23$, 95% CI_d [0,04, 0,41]. At P.M.3, the perceived skewness of evidence (*log* =.05, *SD* = .52) was not significant, $t(109) = 1.12$, $p = .13$, $d = .11$, 95% CI_d [-0,08, 0,29].

A significant, positive correlation between choices in the first sampling phase and the perceived skewness of the base rates at P.M.2 was found, $r(110) = 0.80$, $p < 0.001$. Additionally, regression analyses were run to further investigate the relationship between the sampling behavior and the perceived skewness of the evidence. The log scores at P.M.2 significantly predicted sampling and explained 64% of the variance by itself, $F(1,108) = 190.25$, $p < .001$, $R^2 = 0.64$. When log scores of perceived skewness of the base rates at P.M.1 were added, a significant regression equation was found as well, $F(2,107) = 95.51$, $p < .001$, $R^2 = 0.64$. However, the log score at P.M.1 was not significant, $p = .34$, only the log score at P.M.2 was a significant predictor, $p < .001$.

For the second sampling phase, a positive correlation between sampling and the perceived skewness of base rates at P.M.3 was found as well, $r(110) = 0.82$, $p < 0.001$. A regression analysis demonstrated that sampling could significantly be predicted by the perceived skewness at P.M.3, $F(1,108) = 225.92$, $p < .001$), $R^2 = .677$. Additionally, a regression equation for the relationship between the last sampling phase and both the perceived skewness at P.M.2 and P.M.3 yielded significant results, $F(2,107) = 116.04$, $p < .00$, $R^2 = .684$. Accordingly, a positive correlation was found between sampling and base rate log scores at P.M.2, $r(110) = 0.37$, $p <0.001$. The base rate log scores at P.M.2 were found to significantly predict sampling, $F(1,108) = 16.60$, $p < .001$), however, only 13,3% of the variance was explained ($R^2 =.133$). Similarly to the results in the first sampling phase, most of the variance in sampling behavior thus was explained by the perceived skewness of base rates after sampling.

Generalizability

To get more insight into the potential generalization of biases upon encountering a new medicine, additional variables related to perceived similarity, effectiveness and alternativeness were investigated. On average, participants indicated the new medicine to be around 65% as effective as the old medicine ($M = 65.85$, $SD = 25.97$). With regards to similarity, they estimated the new medicine to be just over 70% similar to the old medicine ($M = 70,59$, $SD = 25.97$). Moreover, they indicated the new medicine to be a good alternative of the old medicine ($M = 70.88$, $SD = 25.97$). A reliability analysis between these three variables was run that yielded a Cronbach α of 0.85, which indicates a high level of internal consistency. All items appeared to be

worthy of retention, demonstrating a decrease in α if deleted. Therefore, a new variable 'generalizability' (M = 69.11, SD = 24.11) was computed that combined the mean scores on all three variables. Generalizability significantly predicts sampling of the new medicine ($F(1,108)$ = 6.41, p = .013), however only 5,6% of the variance in sampling can be explained by how similar participants perceived the medicine to be to the previous medicine (R^2 = 0,56).

Discussion

The current research tested the novel hypothesis that an initial pseudocontingency bias would generalize to similar alternatives and would lead to biased sampling of this alternative. In doing so, the study simultaneously replicated and validated previous findings by testing the hypotheses that pseudocontingency inferences induce an initial bias and that this bias will subsequently lead to biased sampling in the health context. For the purpose of exploring these hypotheses, the current study induced pseudocontingency biases by presenting participants with skewed base rates of options and outcomes, followed by two free sampling phases in which the actual choice behavior was measured.

Support was found for the hypothesis that pseudocontingencies induce a preference towards the frequent option in a reward-rich environment. After presenting participants with skewed base rates of frequencies of options and outcomes, they indicated a clear preference towards the frequently presented medicine as opposed to the infrequent medicine. Accordingly, the other preference measures also revealed a clear bias towards the frequent option. Since bias towards the preferred option was detected in participants' subsequent decisions during the first free sampling phase, the second hypothesis was supported as well. Congruently, this bias was reflected in the self-reported measures. Altogether, the current findings provide clear evidence that pseudocontingencies have the potential to induce biases that guide successive decision making in the health context.

The last hypothesis, which proposed that an initial bias would generalize onto a similar, novel alternative, was partially supported. As participants' behavior revealed, the new medicine representing the alternative to the frequent option was chosen significantly more often than the

infrequent option. However, the self-reported preference measures related to participants' perceptions of this new medicine were not significantly more preferred than the alternative on any of the preference estimates. Still, a positive correlation between sampling choices and the perceived skewness of base rates shows that, although people did not perceive the evidence to be significantly skewed, there is a relationship between the perceptions of the evidence and the sampling decisions.

Alternative Explanations

Altogether, the results demonstrate a clear general trend towards the frequently presented medicine that is in line with the induced positive pseudocontingency. However, there is some incongruence in the findings that needs to be explained. For instance, in discord with both the second and third hypothesis, there is no evidence suggesting that people's sampling decisions were based on how skewed they perceived the evidence to be. Although positive correlations were found between sampling and perceived skewness of evidence for both phases, when investigating the direction of this relationship, there was only weak evidence suggesting that sampling could be predicted on the basis of perceived skewness before the respective sampling phase. Rather, it became evident that perceived skewness after sampling predominantly predicted the variance in sampling. Do these results say no more than that people base their base rate predictions on how they sampled, which is precisely what they were instructed to do? Or could this also mean that we do not yet have a complete understanding of how pseudocontingency effects work?

Separate trial analyses

One way to interpret these incongruent results is by redirecting the focus to each sampling trial separately (see Figure 1), instead of interpreting the mean sampling score as was done for the analyses. A striking general trend for both sampling phases can be observed in which the frequent option is chosen remarkably more often upon the first few trials but is chosen less and less with every trial until it stagnates above chance level. Looking at the first couple of trials only, it is plausible to expect that people's preferences towards the frequent option have been influenced by their perceptions of base rates before sampling, as this preference is in the expected direction of the frequent option. However, the graph also demonstrates that the frequent

medicine is chosen observably less after the first 25 trials, which may explain why there is no overall predictive relationship between sampling and base rates, as people appear not to base their decisions as much on the initial evidence as they did during the earlier trials.

Biased sampling the frequent medicine. Why does this graph show a huge spike of frequent option choices during the first few trials and why does this curve stagnate above chance level? Examining the first free sampling phase, the parsimonious explanation is that this distribution is exactly how one would expect a bias and subsequent learning to look like. As participants have a 75% chance to win on each trial and were made to believe this to be the case during the forced trials, the spike presents a clear pseudocontingency effect as the option was chosen around 60% of the time during the first couple of trials. This seems to suggest that pseudocontingency effects may be especially strong directly after having observed the initial evidence but decrease over time.

With regards to the peak in sampling of the new medicine, an alternative explanation could be that novelty effects (Tulving & Kroll, 1995) may have played an important role in determining their sampling behavior. Novelty effects induce preferences towards novel stimuli over old stimuli (Berlyne, 1970). Whether people's sampling behavior has been influenced by novelty has an important implication, since people's motivation to choose the new medicine may come from a source of exploration of the new, interesting option. When choices are influenced by novelty effects this could imply that people may not have relied on a strategy of exploitation, as the present paper aims to show, but on an exploration strategy. This would be an interesting and relevant claim for future research to validate.

Stagnation of the curve. A very plausible interpretation of the stagnation of the curve could be that people have successfully learned their chances of winning, regardless of the option. That is, in line with current findings, they estimated their likelihood of winning, regardless of the medicine, quite accurately as being around 70%. Having noticed that the experiment took place in a reward-rich environment, people may have changed their decision-making strategies accordingly, leading to a strategy that resembles exploration in which people were somewhat indifferent about which option to choose, as winning was likely to happen anyhow. If this interpretation is correct, then it may imply that as pseudocontingency effects wear off after a

certain number of trials, people may simultaneously switch from exploitation strategies to exploration strategies.

Another explanation is related to the alleviation of boredom. As introduced by Cohen et al. (2007) and used in the context of the exploitation-exploration paradigm by Mehlhorn et al. (2015), a reason to switch to a strategy of exploration after exploitation is not just to increase information, but also to reduce boredom. Especially in the context of the current experiment, it is plausible that participants experienced boredom and intended to alleviate this by switching options, since having to press the same button over the course of a hundred trials is hardly exciting. Interestingly, by occasionally choosing the option that supposedly is less rewarding in the context of the experiment, one may argue that participants actually received a personally more rewarding effect, namely reducing boredom and speeding up the experiment. Perhaps participants' goal of the experiment, improving their hypothetical health, was not personally relevant enough and the benefits of reducing boredom were greater than the benefits of improving their hypothetical health. Despite adding the stick figure as a personalized visual reminder of their health progress, it could be that people were still not personally motivated enough to treat this hypothetical health as their real health.

Alternatively, sampling could have also stagnated due to negative recency effects. Research on probability learning has shown that negative recency effects, in which a tendency towards the less likely option is observed (Jarvik, 1951), are likely to happen for longer runs of trials of random events. This relates to the previously mentioned gambling fallacy (Kahneman & Tversky, 1972), in which people expect runs of the same outcome to be less likely to happen than they actually are. As people noticed that they continued receiving positive outcomes upon choosing the frequent medicine, they may have estimated their chances of being rewarded to become less likely with every trial. The experience that they are winning so often is not representative of the statistical probability, making them switch options out of disbelief. Yet, as opposed to the randomness of sequences of events and outcomes that Kahneman and Tversky (1972) discussed, the experimental instructions of the present study actually intended to make participants believe that options and outcomes were linked. However, the game-like setup and the experimental setting in general may have encouraged participants to rely on this heuristic.

Finally, it is important to note that although the curve stagnates, it does so well above chance level. This relates back to the main findings of this study, namely that people guide their behavior on the basis of the induced bias. Apparently, despite the potential knowledge of the reward structure of the experiment, boredom alleviation attempts, and negative recency effects, people were still more inclined to choose the frequent option over the infrequent option. In the end, a clear tendency towards the frequent option was shown, which - statistically speaking - should not be the case if they were unbiased.

Generalization

According to attitude generalization literature, generalization depends on people's assessments of how similar a stimulus is to past instances that either proved to be positive or negative (Fazio et al., 2015). Therefore, the current study analyzed how similar people perceived the new medicine to be compared to the original one and showed that the more similar people perceived the medicine to be predicted whether people would subsequently choose this new medicine. The extent to which a bias is generalized onto another alternative thus appears to depend on how similar people perceive this alternative to be. However, this claim should be interpreted with caution, since only a very small percentage of the variance in sampling was predicted by generalizability.

Although this new medicine was introduced as an alternative that is virtually identical to the old medicine, having the same active ingredients and belonging to the same manufacturer, the average generalizability score was only 70%. Although this may seem relatively high at first glance, when interpreting this on the basis of the questions that it relates to, it actually demonstrates that despite there being no real difference other than the name, people appeared not to perceive this medicine as effectively identical. This interpretation makes sense in light of the current finding that people did not explicitly prefer the new medicine over the alternative, which may reveal that people have thus treated the new alternative as if it were a totally new option.

In terms of their behavior, however, they do seem biased in their decisions, as they choose the new medicine significantly more. Whereas people seem to treat the newly introduced medicine as novel in their explicit self-reported measures, it could be that in less explicit forms of preference, namely their sampling decisions, they somehow still perceive the medicine to be

similar enough to the previous medicine to treat it as such. Could it be that pseudcontingency effects are predominantly implicit? Are pseudocontingency effects indeed strong enough to spillover onto similar stimuli, but is this only shown in implicit measures? These types of questions about implicit and explicit measures of pseudocontingency biases introduces a new topic that goes beyond the scope of this paper but would be interesting to revisit in future research.

Conclusion

The present inquiry into the rise, maintenance and generalization of pseudocontingency biases in the health context has given some valuable insights into pseudocontingency effects. First of all, the current findings add further evidence for pseudocontgingency effects in the health context. Moreover, it adds to the existing body of research by providing support for the recent claim that pseudocontingency biases are maintained in reward-rich environments. Finally, building on research by Harris et al. (2020), the present study is the first research to investigate the potential generalization of pseudocontingency biases onto novel stimuli. Although no firm conclusions can be drawn, it appears that pseudocontingency biases at the very least have the potential to be generalized onto similar stimuli.

However, the current experiment also provides evidence that adds the necessary bit of nuance to the observed general trend. Although pseudocontingency effects evidently occurred and were maintained, these results should be interpreted with caution, as it remains unclear which type of strategy people have relied on for their decisions. Specifically relating to the topic of the generalization of biases, the present study voices some serious doubts about whether the demonstrated spillover effect was influenced by a strategy of exploitation. Moreover, no firm conclusions can be drawn about the relationship between perceived similarity, attitude generalization and the observed transfer of biases. In sum, this paper provides interesting new insights and has paved the way for new avenues of research to be conducted in the field of pseudocontingencies.

Practical importance

A final note is related to the practical importance of pseudocontingency effects in the health context. A recent study suggests that many Dutch people buy over-the-counter drugs that

do not require prescriptions by medical professionals (CBS, 2013). Therefore, the current findings could be considered a warning. False inferences between pieces of health information may arise in health contexts. When people rely on wrong information or biased judgments, they might choose medicines that may not improve their health, or even worsen it. Even more, this type of bias may guide further decision making, potentially leading to the repeated choosing of suboptimal medicines for a longer period of time. Finally, it is uncertain, yet plausible, that a bias towards a certain medicine may spill over onto other, similar medicines.

Nonetheless, the current findings should not just be considered a warning. They may also serve as a helpful reminder of the potential threats to be found in the perception of health information. Perhaps this reminder may actually raise awareness in individuals, suppliers and professionals involved in providing and receiving health information and activate them to prevent false perceptions from happening.

References

Berlyne, D. E. (1970). Novelty, complexity, and hedonic value. *Perception & psychophysics*, *8*(5), 279-286. https://doi.org/10.3758/BF03212593

Bott, F. M., & Meiser, T. (2020). Pseudocontingency inference and choice: The role of information sampling. *Journal of Experimental Psychology: Learning, Memory, and Cognition***.** Advance online publication. http://dx.doi.org/10.1037/xlm0000840

Centraal Bureau voor de Statistiek. (2014, 30 juni). *Medische contacten, opnamen, medicijnen; leeftijd, geslacht 2010-2013.* StatLine. https://opendata.cbs.nl/statline/#/CBS/nl/dataset/81178NED/table?fromstatweb.

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

Crocker, J. (1981). Judgment of covariation by social perceivers. *Psychological Bulletin*, *90*(2), 272-292. https://doi.org/10.1037/0033-2909.90.2.272

Denrell, J. (2005). Why most people disapprove of me: experience sampling in impression formation. *Psychological review*, *112*(4), 951-978. https://doi.org/10.1037/0033-295X.112.4.951

De Houwer, J.D., & Beckers, T. (2002). A review of recent developments in research and theories on human contingency learning. *The Quarterly Journal of Experimental Psychology: Section B*, *55*(4), 289-310. https://doi.org/10.1080/02724990244000034

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

Fazio, R. H., Pietri, E. S., Rocklage, M. D., & Shook, N. J. (2015). Positive versus negative valence: Asymmetries in attitude formation and generalization as fundamental individual differences. *Advances in experimental social psychology, 51*(1), 97-146. https://doi.org/10.1016/bs.aesp.2014.09.002

Fiedler, K. (2000). Illusory correlations: A simple associative algorithm provides a convergent

account of seemingly divergent paradigms. *Review of General Psychology*, *4*(1), 25-58. https://doi.org/10.1037/1089-2680.4.1.25

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

Fiedler, K., Kutzner, F., & Vogel, T. (2013). Pseudocontingencies: Logically unwarranted but smart inferences. *Current Directions in Psychological Science, 22*(4)*,* 324 –329. https://doi.org/10.1177/0963721413480171

Gigerenzer, G., & Selten, R. (2002). The Adaptive Toolbox. In G. Gigerenzer & R. Stelten (Eds.), *Bounded Rationality, the Adaptive Toolbox* (pp. 36-50). MIT. Press. http://www.alice.id.tue.nl/references/gigerenzer-2001.pdf

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, *62*(1), 451-482. https://doi.org/10.1146/annurev-psych-120709-145346

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,* 1-23. Publish ahead of print. https://doi.org/10.1037/xge0000754

Jarvik, M. E. (1951). Probability learning and a negative recency effect in the serial anticipation of alternative symbols. *Journal of experimental psychology*, *41*(4), 291-297. https://doi.org/10.1037/h0056878

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive psychology*, *3*(3), 430-454. https://doi.org/10.1016/0010-0285(72)90016-3

Kahneman, D. (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.

Klauer, K. C. (2015). Mathematical modeling. In B. Gawronski & G. V. Bodenhausen (Eds.), *Theory and explanation in social psychology* (pp. 371–389). New York, NY: Guilford Press.

Kutzner, F., Freytag, P., Vogel, T., & Fiedler, K. (2008). Base‐rate neglect as a function of base rates in probabilistic contingency learning. *Journal of the Experimental Analysis of Behavior*, *90*(1), 23-32. https://doi.org/10.1901/jeab.2008.90-23

Kutzner, F. (2009). Pseudocontingencies - rule based and associative. (Doctoral thesis). University of Heidelberg, Heidelberg, Germany. http://archiv.ub.uni-heidelberg.de/volltextserver/9912/1/Dissertation_Kutzner_UB.pdf

Kutzner, F., Vogel, T., Freytag, P., & Fiedler, K. (2011). Contingency inferences driven by base rates: Valid by sampling. *Judgment and Decision Making, 6*(3), 211–221. http://journal.sjdm.org/11/9727/jdm9727.pdf

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., ... & Gonzalez, C. (2015). Unpacking the exploration–exploitation tradeoff: A synthesis of human and animal literatures. *Decision*, *2*(3), 191-215. https://doi.org/10.1037/dec0000033

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. http://dx.doi.org/10.1037/xlm0000432

Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological bulletin*, *134*(2), 207-222. https://doi.org/10.1037/0033-2909.134.2.207

Simon, H. A. (2000). Bounded rationality in social science: Today and tomorrow. *Mind & Society*, *1*(1), 25-39. https://doi.org/10.1007/BF02512227

Sims, C. R., Neth, H., Jacobs, R. A., & Gray, W. D. (2013). Melioration as rational choice: Sequential decision making in uncertain environments. *Psychological Review*, *120*(1), 139-154. https://doi.org/10.1037/a0030850

Thaler, R. H., & Sunstein, C. R. (2008). Nudge: improving decisions about health. *Wealth, and Happiness*. New Haven: Yale University Press.

Tulving, E., & Kroll, N. (1995). Novelty assessment in the brain and long-term memory encoding. *Psychonomic Bulletin & Review*, *2*(3), 387-390. https://doi.org/10.3758/BF03210977

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *science*, *185*(4157), 1124-1131. https://doi.org/10.1126/science.185.4157.1124