



Universiteit Utrecht

Does a First Bias that goes against the Correct Normative Answer Persist?

Michelle Mojet (5724023)

Department of Social Sciences, Utrecht University

Master thesis Social Organization and Health Psychology

Academic Supervisor: Chris Harris

Amount of words: 4336

June 23rd, 2020

Abstract

The current study tested whether a first bias that goes against the normative correct answer persists in a free sampling phase. The current study presented participants with two options and aimed to induce a bias towards a supposedly better option using a pseudocontingency. According to pseudocontingency theory, participants are expected to form inferences by associating the frequent option and the frequent outcome, and the infrequent option and infrequent outcome. If a bias was successfully induced, we hypothesized that participants would not overcome their initial bias even though they were faced with a clear normative correct answer. The results showed that the current study failed to induce an initial bias towards the supposedly better option. Because of this, there was no maintenance towards this supposedly better but actually worse option during the free sampling phase. Instead participants did not show a preference for either options. Based on the current study's results, future research is encouraged to continue exploring the boundary conditions in which a first bias is upheld even during repeated interactions.

People are confronted with situations in their daily lives that require them to make choices. Ideally these choices lead to the most optimal outcome. However, there is a large variety of examples where people do not always make the most optimal decision. One of those examples is people choosing homeopathic medicines over evidence-based medicines. Another example is people choosing crash diets over evidence-based weight loss programs. The similarity in these examples is that people seem to have unwarranted beliefs because they tend to stick to their initial choice, even though they repeatedly encounter evidence of the contrary. This is, for example beliefs about healthy food. Some people chose to eat meat every day since they believe it is important for their health. Whereas it has been shown by research (e.g. De Boer, Schösler, & Aiking, 2014), presented on the news (Nu.nl, 2017), and advised by national food boards (Voedingscentrum, n.d.) that it is unhealthy to eat meat every day and not sustainable either. That person's choice is not only non-beneficial for them, but it also causes negative consequences for society as a whole. Nevertheless, if someone believes that eating meat every day is healthy, it is understandable that they stick with this way of maintaining their health. On a day to day basis eating meat might feel good to that person and indeed continued overall health would suggest that things are going well. However, it also results in that we do not feel the necessity to explore alternatives and test whether perhaps without meat we would feel even healthier. People tend to stick to their choice and exploit this allegedly best option, all based on a first bias towards that option. However, sticking to this first bias could be non-adaptive for a person, since they prevent themselves from optimizing their choices, which is picking the choice that is the most beneficial for them. Therefore, it is interesting to explore whether a first bias persists because it could lead to people choosing to exploit inferior options.

In the process of making repeated choices two phases can be distinguished; information sampling, which is the gathering of information about all the options, and reward maximization, which is the exploitation of one, considered most rewarding, option (Mehlhorn et al., 2015). During the process of making a choice, people could prematurely start to exploit the seemingly most rewarding option. For example when someone moves to a new town and they want to explore possible restaurants for their Friday evening pizza night they have every week. The person tries a couple of different restaurants of which they thought that some were positive and some were negative. While that person was still in their search of trying out different restaurants, they hear from their co-workers that one of the restaurants they visited is, according to the co-

workers the best in town. The person decides to pick that restaurant as their weekly Friday night pizza place, because they in general experienced the most frequent positive feedback on it, even though the person had not explored all pizza places across town. In this example of looking for the best pizza place, if the recommended pizza place is in fact the best in town, they should start to exploit it prematurely and not waste a lot of time exploring other options when they already know the correct answer. If someone happens to come across the correct answer during the phase of exploration it is not problematic. However, in the case of someone not knowing the correct answer, premature exploitation is not beneficial because of two reasons. First of all they unknowingly exploit the seemingly most rewarding option. Second of all they prevent themselves from trying other options and finding a better one. This could be disadvantageous because people are easily influenced by biases such as primacy effect or availability in the environment.

People are prone to the primacy effect (Asch, 1946; Anderson & Barrios, 1961) according to which people tend to weigh the initial evidence more heavily than the evidence that follows, even if it is contradicting evidence proving that the first trails were wrong (Rosnow, 1966). Initial evidence that conforms with a communicated belief has been found to influence subsequent biased choice behaviour and facilitate maintenance of unsupported beliefs (Pildtich & Custers, 2018; Staudinger & Büchel, 2013). The primacy effect could illustrate why people willingly exploit their initial preference instead of looking for and accessing more evidence. In addition to that, one's surroundings could also influence the availability of a certain option. People could be confronted with situations that could already have a skewed base rate of certain options being presented. One option could be more available than the other. A person acting in this environment, will interact more with the more frequent option, which causes them to have more opportunities to receive a positive outcome.

An interesting example of people coming to wrong conclusions because of a formed bias towards one option, are pseudocontingencies (Fielder & Freytag, 2004; Fiedler, Freytag, & Meiser, 2009). According to pseudocontingency theory, people form inferences by associating the frequent cue and the frequent outcome, and the infrequent cue and infrequent outcome. In the above mentioned example of the search for a new perfect pizza place, one pizza restaurant was frequently experienced more positively through the person's own encounters and positive feedback from co-workers, and the other pizza places were less frequently experienced and less

frequently experienced in a positive way. This example illustrates that a pseudocontingency inference is formed when people associate the frequent cue with the frequent outcome, and the infrequent cue with the infrequent outcome. One could argue that choosing one's pizza based on a first bias, due to pseudocontingencies, is not problematic per se. However, it is a representation of how forming a belief in a new unknown situation could go. The effect of a pseudocontingency causes a person to form an initial bias towards a certain option, making them want to exploit that option but meanwhile preventing them from exploring and exploiting possible better options.

Environments can also influence this trade-off between exploration and exploitation. In a reward-rich environment, options at baseline yield predominantly positive outcomes whereas in a reward-impooverished environment options at baseline yield predominantly negative outcomes. It could be argued that a reward-rich environment would stimulate a person to prematurely exploit their initial preference because the positive outcome reinforces this behaviour. They have no reason to explore alternative options since their current seemingly best option is already rewarding. However in a reward-impooverished environment, options with a negative outcome do not reinforce people to stick with one option which in turn motivates them to explore alternatives for possible better outcomes. In other words, the options are not rewarding enough for a person stimulating them to continue exploring and not maintain their first bias. A proposed mechanism is that as long as people sample in a reward-rich environment, they do not feel the need to start exploring other options if they believe one option to be superior. As a consequence, people have the tendency to stick to their initial bias regardless of multiple opportunities where they can encounter evidence of the contrary. Risk aversion explains that people tend to avoid uncertain outcomes even if they could be better or similar to the outcome they currently have (Holt & Laury, 2002; Kahneman & Tversky, 1984). In other words, they like to stick to what they know, by exploiting the option based on their initial bias, because it is a safe bet. Harris, Fiedler, Marien and Custers (2020) tested whether initial biases would be sustained in a reward-rich and reward-impooverished environment. Compared to other experiments in the field of pseudocontingencies (Fiedler & Freytag, 2004; Fiedler, Freytag, & Meiser, 2009), Harris et al. (2020) not only presented an induction phase in which one option was shown more frequently than the alternative and one outcome occurred more frequently, but added a phase in which participants could sample freely between both options. They found that participants in the

reward-rich environment continued to exploit the option based on their initial bias. This suggests that in reward-rich environments people are prone to exploiting a possible inferior option due to their first bias. However, there was no normatively correct answer in the free sampling phase because both options were equally rewarding. An alternative explanation for their findings could therefore be that because the surrounding in both phases were the same, participants did not start the process of exploration and exploitation from the beginning but stuck with their first bias. In addition to that, because options A and B were equally rewarding, participants were not actually exploiting inferior options. Following this reasoning, there is no clear evidence yet that people, in reward-rich environments, stick to their initial bias and exploit a clear inferior option. The question remains whether people, after forming a first biased impression, notice that they are exploiting an inferior option when faced with clear evidence. Therefore, the current study aims to address this open question in the literature, by investigating whether a first bias in a reward-rich environment persists even though there is a correct normative strategy.

The current study will test whether a first bias that goes against the normative correct answer persists in a free sampling phase, the following two hypotheses are composed. Firstly, in the initial evidence phase participants will perceive the more frequent option as more rewarding and the less frequent option as less rewarding if a pseudocontingency is successfully achieved. If so, then secondly it is hypothesized that this pseudocontingency persists in a free sampling phase, regardless of the correct normative answer. That is, participants continue favouring one option even when the alternative option is objectively the better choice.

Methods

Participants

One hundred participants participated in the current study with an average age of 30.92 years old ($SD = 9.25$), and of which 58 were female. The participants were approached through the online crowdsourcing platform Prolific Academic (<https://prolific.co/>) and the study was run in English on Soscisurvey. This means that they completed the experiment in their environment. They were financially rewarded with, on average, £1.00 with earnings dependent on their performance (min £0.85 and max £1.15). According to a power analysis using G*power (Faul,

Erdfelder, Lang, & Buchner, 2007) the study will need $N = 67$ participants using a 5% alpha-level, 80% statistical power, and a small to medium effect size (between $d=.35$ and $d=.40$) as reported by Harris et al (2020). The experiment was conducted according to the guidelines of the Ethics Review Board of the Faculty of Social & Behavioural Sciences at Utrecht University.

Design

The experiment had one reward-rich experimental group. The current study was divided into four phases which will be described below: an induction phase; an evaluation phase; a free sampling phase; and an evaluation phase.

The first phase was the induction phase of which the goal was to induce a bias. Participants were forced to sample the distribution presented in Table 1. Which bag was shown more frequently and which colour was shown more frequently was counterbalanced.





The second phase was an evaluation phase to measure whether a pseudocontingency was successfully achieved. Participants were asked to give three estimates regarding the task; participants' relative preference, their expectation of outcome, and their confidence in making a reasonable estimate.

The third phase was the free sampling phase, in which participants were able to choose freely between both bags. In this phase the previously frequent option in the initial evidence phase had a probability for positive outcome 66% of the time (8/12) and the previously infrequent option in the initial evidence phase had a probability for a positive outcome 80% of the time (4/5). The probability of winning in the free sampling phase, therefore, remained identical to what participants perceived in the induction phase.

The fourth phase was an evaluation phase again. Participants were asked to give the above mentioned three estimates again regarding the task.

Table 1

Distributions of initial evidence

Induction phase		
		
	8	4
Frequently shown bag		
	4	1
Infrequently shown bag		

Procedure

Participants completed the experiment online using a computer in their settings. The first instructions were to silence their phone, e-mail, music and anything else distracting to minimize external distractions that could influence their concentration. In addition to that, the instructions included an attention check.

In the induction phase participants were presented with two bags, bag A and bag B. In this phase they could not choose freely between both options, the computer chose which bag they

could click on and either a yellow or blue ball would appear (based on the ratio presented in Table 1).

The induction phase was followed with an evaluation phase in which participants were asked to give three estimates regarding the task. The first estimate was the relative preference measure. Participants were asked to rate which of the two bags they were more likely to grab a blue or yellow ball on a scale with images of both bags at the end as an anchor. The second estimate was the conditional estimate. Participants were asked how likely it was to grab a blue or yellow ball if they chose bag A (and bag B) on a scale anchored at 0% and 100%. The third estimate was participants' confidence in making a reasonable estimate. They were asked how confident they were in making a reasonable estimate regarding bag A (and bag B) on a scale from 0% to 100%.

In the free sampling phase, participants were instructed that nothing would change about the two bags and the amount of blue and yellow balls in the two bags, but that they would win 10 points for every yellow ball and lose 10 points for every blue ball. They were also told that their score would be displayed in the top right corner. As well as that their score would be converted in a monetary reward at the end of the study.

After the free sampling phase, the participants were asked again to give the three estimates regarding the task. They were asked the same questions as mentioned above in the first evaluation phase.

The experiment finished with a final series of demographic questions. Participants were asked what age and gender they had and what their highest level of education was. They were also asked if they currently or in the past were a psychology student. Finally, a series of 10 questions were asked to assess participants' level of curiosity. However, these questions were not used for testing the hypotheses in the current experiment (for more information see: <https://aspredicted.org/j8r5a.pdf>).

Data analysis

To test whether participants continue favouring one option even when the alternative option is objectively the better choice in the free sampling phase, a t-test was used to compare the difference between the group and chance level. The t-test was done with four different dependent variables. The first was the behavioural outcome which was the choice of bag across

trials in the free sampling phase. The second was the relative preference measure towards the choice options. The third and fourth were the conditional estimates and confidence estimates regarding the conditional estimates.

To analyse the four dependent variables the data was recoded. For the behavioural outcome, the choice of the frequent bag was +1, and the choice of the infrequent bag as -1. This created a value that represented participants' overall preference for either two options. For the relative preference, conditional estimate, and confidence estimate all the slider values were recoded so that a value of 0 would represent a neutral preference for both options and a positive value would represent the option that was encountered more frequently in the induction phase. In other words, if bag A was shown more positively eight times, all measurements for these participants were recoded in a way that a positive value on these estimates would indicate a positive contingency between bag A and the yellow balls.

Results

To test whether a first bias that goes against the normative correct answer persists in a free sampling phase, it is necessary to analyse whether a bias for the more frequent option was achieved in the initial evidence. A t-test showed a non-significant effect in preference $t(99) = -.204, p = .838, 95\% \text{ CI} [-6.8512, -5.5712]$, as well as a non-significant effect in how likely participants thought they would get a positive outcome for the more frequent option $t(99) = -.410, p = .683, 95\% \text{ CI} [-.0507, .0771]$, and a non-significant effect in their confidence in making reasonable estimate $t(99) = -.700, p = .487, 95\% \text{ CI} [-.0228, .0476]$. These results indicate that a bias in the initial evidence phase was not successfully achieved.

To test whether there was a difference in choice behaviour in the free sampling phase compared to the initial evidence phase a t-test was done which showed a negative significant effect, $t(99) = -3.556, p = .001, 95\% \text{ CI} [-.1492, -.0423]$. This indicates that participants chose the infrequent option, the normative correct answer, more often compared to the frequent option

over time. See figure 1 for participants' aggregated choices over time and Figure 2 for the mean percentage of frequent choices per participant.




Figure 1. Percentage of participants sampling the frequent option per trail.




Figure 2. Proportion of the frequent choice per participant.

After the free sampling phase, the participants' opinion on the two options was tested. A t-test showed a significant effect in preference $t(99) = -3.667, p = .000, 95\% \text{ CI } [-18.8317, -5.6083]$, and a significant effect in how likely participants thought they would get a positive outcome for the more frequent option $t(99) = -3.060, p = .003, 95\% \text{ CI } [-.1866, .0398]$. These results show that participants showed a preference for the normative correct answer and that they would get a positive outcome for this option. However a t-test showed a non-significant effect in participants' confidence in making a reasonable estimate $t(99) = -1.157, p = .250, 95\% \text{ CI } [-.0798, .0210]$. This indicates that they were not sure about making a good estimate regarding the two choices. These results would not be expected when participants prefer the normative correct answer.

Discussion

The current study explored whether a first bias that goes against the normative correct answer persists in a free sampling phase. To test this, the current study presented participants with two options of which one option was more frequent but less rewarding and the other option was less frequent but more rewarding. In other words, the more frequent option was the supposedly better option and the less frequent option was the actual better option. The hypothesis was that if a pseudocontingency was successfully achieved, participants would have a bias towards the frequent option. That is, they would perceive the more frequent option as more rewarding and the less frequent option as less rewarding. It was hypothesized that participants would maintain this bias towards the supposedly better option in the free sampling phase, even though the other option was more rewarding. However, the results showed that the current study failed to induce an initial bias towards the supposedly better option. Because of this, there was no bias towards this supposedly better but actually worse option during the free sampling phase. Instead, participants did not show a preference for either option.

These results are in contrast with the findings of Harris et al. (2020) who successfully induced a bias using a pseudocontingency. The length of the initial evidence phase of the current study is similar to Harris et al. (2020). However, the distribution that was used is different. Specifically, the distribution that Harris et al. (2020) used had a contingency of 0. Accordingly, there was no difference between the two options in how likely each one would result in a win (for more information see Harris et al., 2020). Whereas in the current study the contingency was negative, such that the less frequent option was more likely to result in a positive outcome. It could be that because of this difference a bias was not successfully induced.

Another explanation for unsuccessfully inducing a bias could be due to the task difficulty. Task difficulty can moderate the effect of pseudocontingencies. Research has shown that with simple tasks, people are less susceptible to pseudocontingencies as they are more able to rely on cell frequencies instead of base rates (Fleig, Meiser, Ettl, & Rummel, 2017). One could argue that the task in the current study was too simple and that participants could fully rely on base rates.

Despite of unsuccessfully inducing a bias, the participants showed interesting choice behaviour over time in the free sampling phase. Since participants were not biased in the initial evidence phase, it could be that they were still information sampling and did not start reward

maximization. In other words, they were still trying both options and did not start exploiting an option that they considered most rewarding. Following this reasoning, they would have entered the free sampling phase while still being in the information sampling phase. It could be expected that over time they learned that the correct normative answer is the most rewarding option and once they did so, they would have started to exploit it. The results of the conditional estimates show that participants had a preference for the normative correct option and considered it more likely to get a positive outcome when choosing the normative correct option. When looking at Figure 1, the percentage of participants sampling the frequent option per trail decreases until halfway of the free sampling phase. This confirms the aforementioned reasoning that participants learned through information sampling that the previously infrequent option in the initial evidence phase is actually the correct normative answer. However, half way through the free sampling phase the percentage of participants sampling the frequent option remains around 38% and only drops to about 35% in the last few trials. Whereas it would be expected that the percentage of participants sampling the frequent option would drop even lower than roughly one-third of the sample. This finding goes against the expectation that participants would exploit the more rewarding option. The results support this observation since participants' confidence regarding the conditional estimates was non-significant. This could indicate that participants were unsure about making a good estimate regarding the two choices.

On the other hand, the finding that participants did not exploit the more rewarding option could be a version of probability matching (Vulkan, 2000). Probability matching can occur when people are faced with the task of choosing between two possible outcomes. They then tend to match their choice to the corresponding probability of the outcome. Instead of maximizing the choice that has the highest probability of a positive outcome, which would be more beneficial for them (Koehler & Jamers, 2010). However, the current experiment is different from the task in which probability matching usually occurs. Namely, the choices do not have dependent probabilities. The outcomes of the choices are independent from each other, which makes it unclear how exactly a probability matching rule would look like here.

Despite the unsuccessful induction of a first bias in the current study, it does give more insight into what future research could do to test whether a first bias persists when a clear correct answer is given. Future research should continue to test the composed hypotheses but they should change the operationalization. Based on the current experiment there are two suggestions

that future research should consider while exploring this question. First of all, to ensure that a bias is successfully achieved, a different distribution should be used. Future research could consider distributions similar to the ones used in Harris et al. (2020) or Fleig et al. (2017). Second of all, the current study used a conservative testing method by comparing the behaviour to chance level. Statistically chance level is an appropriate value that represents no preference for either option. Nevertheless, people rarely maximize completely (Vulkan, 2000), which could mean that no preference for either option could be higher or lower than chance level. As well as, that no preference would probably also result in a varying value over trials. Instead, future research could use a control group to have a more accurate comparison.

Looking at these results from an overarching perspective, people make a lot of repeated decisions in their daily lives. These repeated decisions could range from something unimportant such as picking the best pizza place, to something important such as choosing your diet that ultimately impacts your health and society. While decision making usually goes well, it does not always result in people choosing the best option for themselves (Harris et al., 2020). But people do seem to be capable of figuring out the task that was used in the current experiment. When we are confronted with a clear distribution we do notice it. Which speaks to how good people are in general in these tasks. But it remains an open question in the literature under what boundary conditions we do or do not figure out to overcome initial biases and when we remain stuck on our first bias even during repeated interactions.

Literature

- Anderson, N. H. & Barrios, A. A. (1961). Primacy effects in personality impression formation. *Journal of abnormal and social psychology, 63*(2), 346-350.
- Ash, S. E. (1946). Forming impressions of personality. *Journal of abnormal and social psychology, 41*, 258-290.
- Button, K. S., Loannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature reviews Neuroscience, 14*, 365-376.
- De Boer, J., Schösler, H., & Aiking, H. (2014). “Meatless days” or “less but better”? Exploring strategies to adapt Western meat consumption to health and sustainability challenges. *Appetite, 76*(1), 120-128. doi: 10.1016/j.appet.2014.02.002
- 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. doi:10.3758/BF03193146
- Fiedler, K., & Freytag, P. (2004). Pseudocontingencies. *Journal of Personality and Social Psychology, 87*(4), 453-467. doi:10.1037/0022-3514.87.4.453
- Fiedler, K., Freytag, P., & Meiser, T. (2009). Pseudocontingencies: an integrative account of an intriguing cognitive illusion. *Psychological Review, 116*(1), 187-206. doi:10.1037/a0014480
- Fiedler, K. (2011). Voodoo correlations are everywhere—not only in neuroscience. *Perspectives on psychological science, 6*(2), 163-171. doi: 10.1177/1745691611400237
- Fleig, H. Meiser, T. Ettl, F. & Rummel, J. (2017). Statistical numeracy as a moderator of (pseudo)contingency effects on decision behavior. *Acta psychologica, 174*, 68-79, doi: 10.1016/j.actpsy.2017.01.002

- 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*. Advance online publication. <http://dx.doi.org/10.1037/xge0000754>
- Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American psychologist*, 39(4): 341-350. doi: 10.1037/0003-066X.39.4.341
- Koehler, D. J. & James, G. (2010). Probability matching and strategy availability. *Memory & Cognition*, 38(6), 667-676. doi: 10.3758/MC.38.6.667
- 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. doi:10.1037/dec0000033
- Nu.nl. (2017, December 5). Europeanen moeten hun voedingspatroon drastisch aanpassen. Retrieved January 16, 2020, from <https://www.nu.nl/>.
- Pilditch, T. D., & Custers, R. (2018). Communicated beliefs about action-outcomes: The role of initial confirmation in the adoption and maintenance of unsupported beliefs. *Acta Psychologica*, 184, 46-63. doi: 10.1016/j.actpsy.2017.04.006
- Rosnow, R. L. (1966). Whatever happened to the “law of primacy”? *Journal of communication* 16(1), 10-31. doi: 10.1111/j.1460-2466.1966.tb00013.x
- Staudinger, M. R., & Büchel, C. (2013). How initial confirmatory experience potentiates the detrimental influence of bad advice. *Neuroimage*, 76, 125-133. doi: 10.1016/j.neuroimage.2013.02.074
- Voedingscentrum. (n.d.). Vlees. Retrieved January 16, 2020, from <https://www.voedingscentrum.nl/nl.aspx>.
- Vulkan, N. (2000). An economist’s perspective on probability matching. *Journal of Economic Surveys*, 14, 101-118.