UNEXPECTED UNCERTAINTY IN DECISION MAKING: HOW IT INFLUENCES SELF-GENERATED THOUGHT

Bachelor's Project Thesis
Ryanne Hauptmeijer, r.hauptmeijer@student.rug.nl
Supervisors: Dr. M.K. van Vugt & S. Huijser, MSc

Abstract: The aim of this study was to investigate whether unexpected uncertainty in decision-making can affect the degree and content of self-generated thought. Unexpected uncertainty is a sudden and fundamental change, that forces someone to adjust their beliefs, based on their experiences, to make good decisions. To address that question, we alternated a decision-making task in which uncertainty could be manipulated and used a sustained attention task in which self-generated thought could be measured. The decision-making task was a two-armed restless bandit task with jumping outcome probabilities representing unexpected uncertainty. To assess self-generated thinking, we used a metronome response task (MRT). During this task, thought content was measured by means of randomly placed questions about what the participant was thinking. We found that participants did report to think about the decision task during the MRT. However, we were unable to find a significant increase in the reports about the decision task after unexpected uncertainty was introduced, so, after a jump of outcome probability. The lack of significance could be explained by the low power of the experiment due to too little observations. In addition, not all participants appeared to understand the task.

1. Introduction

Decision-making is part of everyday life. Whether to decide what to have for breakfast or which route to take to work. If it is an important or difficult decision, you might catch yourself thinking about what to decide during routine or boring tasks that precede this decision. For example, if you want to buy a house you might think about how to finance the house while doing the dishes. Not surprisingly, uncertainty, which involves a situation which contains incomplete or even unknown information, can affect decision-making. For example, the route you take to work is under construction, however, you do not know how this affects the traffic. Do you take the normal route, or do you decide to take a detour? Bland and Schaefer (2012) suggest that uncertainty can be divided into three types; expected uncertainty, unexpected uncertainty, and volatility. Expected uncertainty arises when rules learned from past events or encounters are weak predictors of the outcomes of future actions, this unpredictability, however, is known and is also stable. An example of expected uncertainty is a decision-making environment where one knows that a certain stimulus predicts rewards 75% of the time. Expected uncertainty is quite well explored in the literature (Yu & Dayan, 2005; Mushtaq, Bland & Schaefer, 2011). Unexpected uncertainty and volatility, however, have not been investigated as much as expected uncertainty. According to Bland and Schaefer (2012), there has yet to be made an explicit distinction between unexpected uncertainty and volatility. The description they provide for unexpected uncertainty is that it can be induced by a single or infrequent fundamental change that forces one to modify their previous beliefs, based on past experiences, to make a good decision. A sudden and fundamental change, unexpected uncertainty, can increase the difficulty of decision-making.

Payzan-LeNestour used a six-arm restless bandit problem to investigate unexpected uncertainty. For the six-arm restless bandit task, participants had to choose between six locations and were given one of three possible rewards for each location. The locations had different outcome probabilities and different frequencies of reward changes. The goal was to gain as many Swiss francs as possible. They found among other things that participants had to be informed that jumps of outcome probabilities could occur, otherwise they would not be able to track the unexpected uncertainty. The participants considered the locations to be random if they were not informed. Payzan-

LeNestour investigated besides unexpected uncertainty, also risk and estimation uncertainty. Similar to these studies, Behrens, Woolrich, Walton, and Rushworth (2007) used a two-armed bandit in their research about uncertainty. The task worked as follows, participants had to choose between two platforms. One of the two platforms yielded in a reward. The reward given when the participant chose correctly, was randomly selected between 0 and 100 for the first platform and for the other platform the reward was (100 – the reward for the first platform). The two-armed bandit had a stable environment where the outcome probability stayed fixed, for example, 75% probability that the first platform yielded a reward, and an unstable environment where the outcome probabilities of the platforms jumped, hence creating unexpected uncertainty. The participants were aware of possible jumps, but they did not know when the jumps would occur or how frequent. Behrens et al. (2007) found that when subjects witnessed a new piece of information (new outcomes), activity levels in the anterior cingulate cortex reflect the salience of the new information for predicting future outcomes. The subjects were more responsive to new outcomes in the unstable environment, with the jumps. The detection of the jumps allowed the subjects to adjust their learning rate without any knowledge of the task structure. The results provided Behrens et al. with a formal explanation on how one weighs different experiences in guiding their future actions.

So, new and/or different experiences can alter your beliefs which influence future actions. Experiencing new information, caused by unexpected uncertainty, might direct your thoughts internally, which can be called self-generated thought, because it could cause you to rethink past decision or evaluate the current decision. Self-generated thought is explored well in literature. Among others, the costs and benefits of self-generated thinking have been studied widely, for example by Mooneyham and Schooler (2013). Self-generated thinking can have a negative influence on reading, mood, and sustained attention. However, Mooneyham and Schooler (2013) also found several benefits to self-generated thinking, such as future and creative thinking. “It is evident that the conditions that maximize mind-wandering can also be the most conducive to creative problem-solving." (Mooneyham and Schooler, 2013). This shows that even though a current task or action might be interrupted due to self-generated thinking, it can help create new solutions to old problems. Dijksterhuis, Bos, Nordgren, & Van Baaren. (2006) also elaborate on the benefits of self-generated thought. They provided an example that when a group of participants performed a decision task, they made a better decision after performing a distraction task compared to a group of participants who deliberately thought about the decision. The distraction task allowed the participants to wander off and think about the decision. This idea is supported by Dijksterhuis (2004) who also discovered that people that engage in self-generated thought improve their quality in complex decision making. Both pieces of research suggest that if you allow people to self-generated thinking during a decision-making task, then people would wander about the decision to be made and make a better decision. To instigate self-generated thought, you could use a boring task that interrupts the decision-making task, an interruption task so to speak. During the interruption task, self-generated thought could be measured to collect thought content.

An often-used task to measure self-generated thought is the Sustained Attention to Response Task (SART, e.g. Smallwood et al. 2004). However, according to Seli, Cheyne, and Smilek (2013), there are several issues with the SART. For example, the instruction in the task might be confusing because participants must be quick but also be accurate and participants do not know how to balance these two. One participant might focus more on being quick while another might favour being accurate. Seli, Cheyne, and Smilek argue in favour of the metronome response task (MRT) which is less susceptible to this problem. An MRT trial consists of a short tone preceded and followed by silence. The subject must press a key in sync with the tone. Each trial has the same duration which makes the MRT a rhythmic key pressing task. Seli et al. proved that there is a relation between variability in response timing around

the metronome signal and self-generated thought, hence the MRT proved to be an effective tool for measuring self-generated thought. Furthermore, unlike the SART where the key reaction is to withhold a response, the MRT always has a response.

The central aim of this research is to test whether unexpected uncertainty in decisions can affect the degree and content of self-generated thought. For this research, an adaptation of the two-armed bandit from Behrens, Woolrich, Walton, and Rushworth, is used to create unexpected uncertainty by including jumps in outcome probability, when they occur is unknown to the participants. The MRT with thought probes is used as an interruption task to measure self-generated thought. We expect that there will be an increase in self-reports of self-generated thought on the two-armed bandit task, with varying difficulty, during the MRT. We expect this increase in self-reports because the jumps make the decision to choose the optimal platform, and hence maximize the total reward, more complex. The interruption task should allow the participants to wander off and engage in self-generated thinking, possibly about the complex decision-making from the two-armed bandit task.

2. Method

2.1. Participants

A total of 25 volunteers (22 right-handed and 3 left-handed) participated in this study of whom 8 were male and 17 were female (ages 18-32; mean age 24.3). The data of 3 participants were excluded for several reasons. One participant was excluded due to a technical malfunction. Two more participants were excluded because they responded to less than 5% of the MRT trials. All participants signed an informed consent form before the start of the experiment. Participants received compensation of €8. The experiment was conducted in accordance with the Declaration of Helsinki.

2.2. Materials

The participants were tested in a small room at the University of Groningen. The experiment was presented on a 27 in. monitor (1920 x 1080 pixels). The stimulus presentation was programmed in OpenSesame 3.2.5 (Mathôt, Schreij & Theeuwes, 2012).

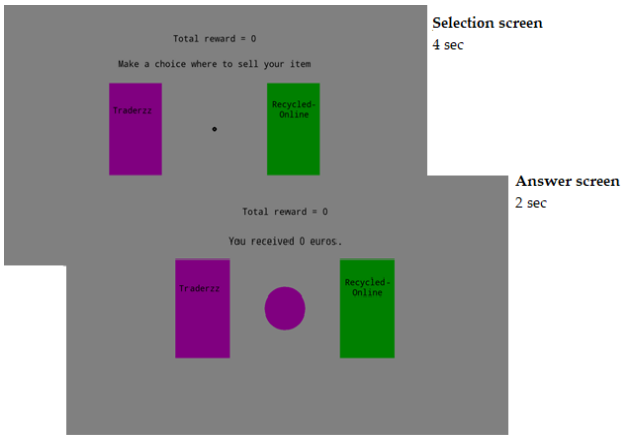

Figure 1. Overview of one trial for the bandit task where Traderzz is the correct one, whereas the participant opted for Recycled-Online.

The entire experiment used a grey background (120, 120, 120; rgb255). Stimuli for the metronome task and instructions were presented in black (0, 0, 0; rgb255). All the instructions were in English.

2.3. Design

2.3.1. Restless bandit task

Participants had to perform a two-armed restless bandit task. The layout of the task was based on the existing work of Behrens, Woolrich, Walton, and Rushworth (2007). Participants had to choose between two platforms. The cover story was that participants had to sell items on an online-selling platform similar to eBay. Uncertainty is explained by the unsteadiness of online-selling platforms. Figure 1. shows the screen participants saw during one bandit trial. A purple rectangle (125, 36, 125; rgb255) represented the online selling platform 'Traderzz'. And a green rectangle (0, 128, 0; rgb255) represented the online selling platform 'Recycled-Online'. Between the two platforms was a black fixation point. The current total reward was shown at the top of the screen (mono font, 32 pixels in height). Below that was an instruction, "Make a choice where to sell your item," (mono font, 32 pixels in height). After they responded, they were shown a screen with the total reward in the same position and the two platforms in the same

position. At the location of the fixation dot, either a green circle (0, 128, 0; rgb255) or a purple circle (125, 36, 125; rgb255) was shown, corresponding with the platform that (would have) successfully sold the item. The correct answer depended on several factors. First, there was a 75% chance that one of the two platforms successfully sold the items for several trials, the other platform had 25% of selling the item. Second, to introduce uncertainty we used jumps in the restless bandit task. After a jump, the 75% chance would change between the two platforms. For example, 'Recycled-Online' was the optimal choice, with 75% outcome probability. After a jump, 'Traderzz' would be the optimal choice with a 75% outcome probability. In total there were seven jumps. Participants were not notified when the jumps occurred. However, they were instructed with the possibility of a change over time. Third, to counteract a potential block-ordering effect, half of the participants started with 'Recycled-Online' having a 75% outcome probability and the other half started with 'Traderzz' having a 75% outcome probability.

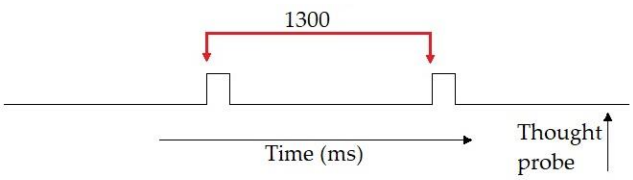

Figure 2. Overview of an MRT trial. The vertical bars represent the metronome tone (75 ms).

2.3.2. Metronome response task

The MRT is based on Seli, Cheyne, and Smilek (2012). Each trial started with 650 ms of silence after which the tone (lasting 75 ms) started, followed by 575 ms of silence. The total duration for one trial was 1300 ms. An overview of an MRT trial can be seen in figure 2. For each MRT trial, participants were shown a fixation dot in the centre of the screen.

2.3.3. Thought probe

During the MRT the participants were presented with a thought probe to collect their thought content. We accentuated that it was normal if they did not think about the MRT but were distracted instead. The thought probe questions are based on the thought probes from Unsworth and Robinson (2016). First, the participant was asked, "What were you thinking about just know?", with the following answer possibilities: (1) I was focused on the current task; (2) I was thinking about the other task; (3) I was thinking about my performance on the current task or how long it is taking; (4) I was distracted by sights/sounds/temperature or by physical sensations (hunger/thirsty); (5) My mind was wandering about things unrelated to the two tasks; (6) I was not paying attention/my mind was blank.

The first answer possibility is considered on-task, whereas all the other answer possibilities are considered off-task.

After each thought probe, there was a confirmation probe, which asked the following question: "How sure are you about your previous answer?". With the following answer possibilities: (1) Confident; (2) Fairly confident; (3) Neutral; (4) Not really confident; (5) Not confident.

2.4. Procedure

The participant was instructed beforehand to bring headphones.

When participants entered the lab, they were positioned in a testing cubicle. They were first asked to sign an informed consent form and complete a short demographic form before they started the experiment. After that, participants were handed the instructions for the experiment. The participants were instructed that they would be performing two alternating tasks. First, the bandit task was explained. They were told that they would be

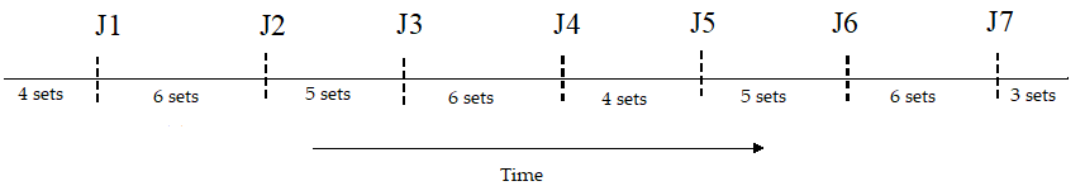

Figure 3. The jumps over time with the number of sets (bandit and MRT trials) in each block.

selling items. They could choose between two online selling platforms, similar to eBay. Their item would be sold on one of the two platforms and they would receive a reward if the item was sold. The MRT instructions told them they had to press the spacebar in sync with a rhythmic sound. After reading the instructions participants were given the time to ask questions, after which, the experiment started.

The participants were first reminded of the instructions of the bandit task after which 5 practice trials followed. Participants had 4 seconds to make a choice for the bandit task. After the choice was made, the correct answer and the received reward was shown for 2 seconds (Figure 1). Next, participants were reminded of the instructions for the MRT and 18 practice trials followed, each trial lasting 1300 ms. During the practice MRT, a thought probe was randomly shown. There was no time limit to answer the thought probe. After the thought probe, the confirmation probe was shown. Again, there was no time limit on answering the confirmation probe. The participant was informed when the practice trials finished and asked once again whether they had any questions after which the participant continued with the experiment. In total there were 195 bandit trials and 1794 MRT trials. The trials were divided into sets. Each set had 5 bandit trials and 46 MRT trials. Figure 3. shows an overview of the jumps and blocks with corresponding sets over time. In total there were 39 sets divided over 8 blocks. The first block had 4 sets, the second block had 6 sets, the third block had 5 sets, the fourth block had 6 sets, the fifth block had 4 sets, the sixth block had 5 sets, the seventh block had 6 sets and the last block had 3 sets. The blocks were unequal in sets, so the participant could not come to expect jumps after a fixed number of trials. The thought probes are randomly distributed during MRT trials in a set, 39 thought probes in total. The confirmation probes followed the thought probes directly, also 39 confirmation probes in total. After 15 minutes there was a break, in total there were three breaks. The experiment lasted approximately an hour.

2.5. Analysis methods

To perform the statistical analysis, we used the R statistical programming language (R Core Team, 2016).

To investigate the performance of the participants on the bandit task, we looked at the optimal choice. For each trial, the optimal choice, which is a binomial variable, is defined as the online selling platform associated with the 75% outcome probability. If the response of the participant corresponds with this platform optimal choice is set to 1 else it is 0.

The variance in rhythmic response times (RRTs, which is the time-based interval between the onset of the tone and the keypress) for the five trials preceding each thought probe was calculated for each participant by using the 'var' function. Variance data were positively skewed; hence they were adjusted using natural logarithm transform similar to Seli et al. (2013).

We decided not to look at all the trials but to focus on a subset of the trials, grouped into sets, for both the bandit task and the MRT. One set for the bandit consisted of five trials and one set for the MRT consisted of one thought-probe response. We are interested in one set before each jump and three sets after each jump, so we can investigate whether there is a significant effect after a jump. We used three sets after a jump because we do not know whether participants discovered the change in one set.

We statistically addressed our hypotheses using linear mixed effects models. These models are useful because they can deal with non-independence in data. Our data consists of multiple responses from the same participants. Therefore, we will use participants as a random intercept to account for the non-independence.

To see whether there is a relationship between performance on the bandit before and after a jump we used optimal choice as the dependent variable, set as a fixed factor and participants as a random intercept. We expect the optimal choice to be lower in the set(s) after a jump compared to the optimal choice of the set before a jump due to the change in outcome probability. The optimal platform has changed without the participants being informed.

We also made a model to see if the transformed RRT variance could be explained by set, the report type of the thought probe

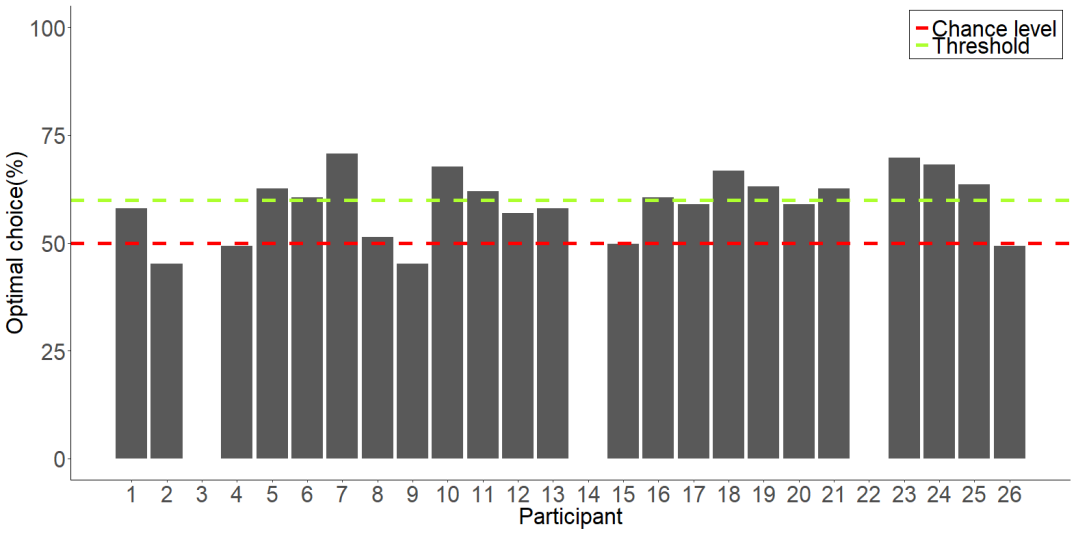

Figure 4. Performance on the restless bandit task for each participant. The red dotted line represents a 50% chance level and the dotted green line represents a 60% threshold for further analyses. Data from three participants were excluded, hence the three empty bars.

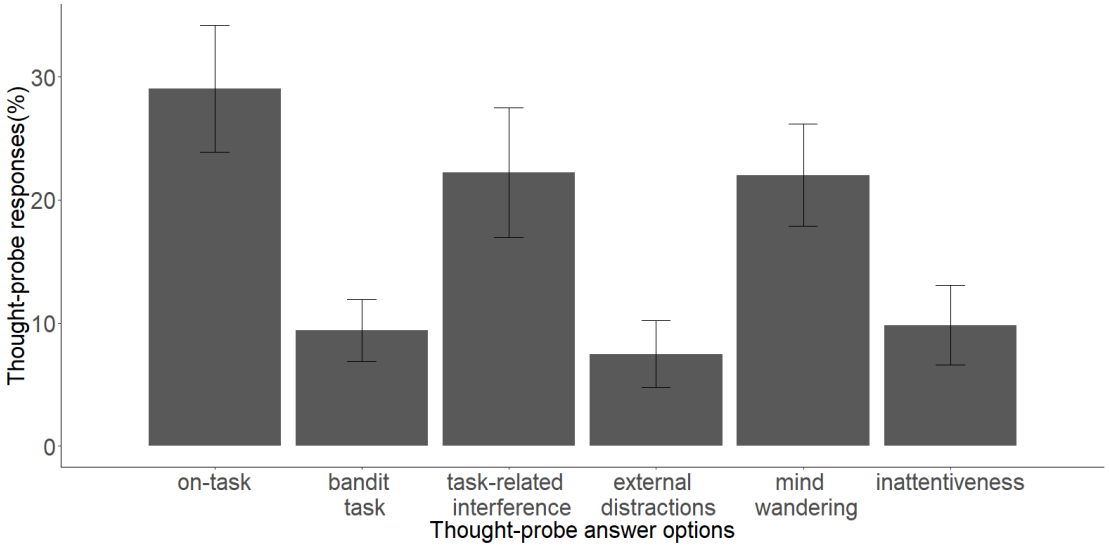

Figure 5. Overview of the different thought-probe responses, in percentages, given by participants. Error bars represent standard error of the mean.

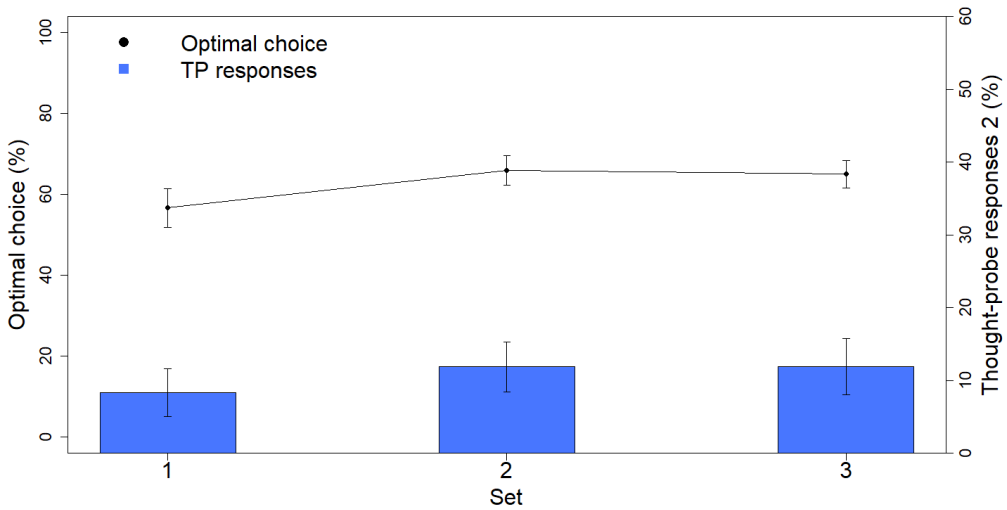

Figure 6. Mean optimal choice of first 3 sets after each jump for a subset of the participants. Mean thought-probe (TP) responses of answer option 2, thinking about the previous task, of first 3 sets after each jump for a subset of the participants. Error bars represent standard error of the mean.

(on/off task) and the response answer options. These variables are of interest because Seli et al. (2013) also looked at the effect of report type and response answer options on the transformed RRT variance. For random effect, we used participants as intercept.

Finally, we needed to look at thought probe responses with answer option 2, which was "thinking about the decision task during the MRT". For the dependent variable, we created thought probe choice which was a binomial variable. It was 1 if the thought probe response was 2 (previous task) or 4 (mind-wandering) and 0 for the other answer options. The model for only thought probe response 2 did not converge (due to too little data), therefore we also included answer option 4. Set was the fixed factor and subjects were the random intercept. We expect that the number of thought-probe responses 2 will be higher in the set(s) after a jump compared to before a jump

Hypotheses were tested using a chi-squared likelihood ratio tests comparing a model with the factor of interest against a model without that factor. Since the optimal choice and the thought probe responses 2 were binomial, we fitted them using logistic linear mixed effect models. The model with transformed RRT variance as the dependent variable was fitted using a linear mixed effects model.

For fitting the logistic linear mixed effect models and the linear mixed effect model we used the 'glmer' and 'lmer' function from the lme4 package (Bates, Maechler, Bolker & Walker, 2015) and the lmerTest package (Kuznetsova, Brockhoff & Christensen 2017).

3. Results

The main goal of the data analysis was to investigate whether participants reported thinking more about the bandit task during the MRT when uncertainty was introduced. First, we looked at the performance on the bandit task. The average performance on the bandit task was 59.4%, which is only slightly better than chance. However, there is a substantial variation in performance between individuals (Figure 4). Some participants performed below chance level on the bandit task. Since we do not know what strategy the participants, who performed poorly on the bandit task, used, participants with performance below 60% were left out for further analyses. These participants most likely did not understand the task. This leaves a subset with a total of 12 participants.

The distribution of thought probe responses for the subset of participants can be seen in figure 5. During the MRT, participants reported being off-task in 70.9% of the trials. They reported thinking about the restless bandit task while they were doing the MRT on 9.4% of the trials.

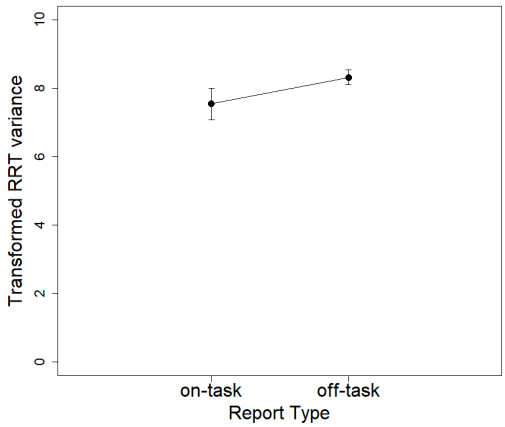

Figure 7. Transformed rhythmic response times (RRT) variance for on-task and off-task reports. Error bars are standard errors.

Mean transformed RRT variances for on- and off-task is presented in figure 7. We want to confirm whether the transformed RRT variance for on-task is lower compared to the transformed RRT variance for off-task, which was the case for Seli et al. (2013). On task are the thought-probe responses that stated, "I was focused on the current task." Off-task consists of all the other thought-probe responses. The mean transformed RRT variance for the on-task reports is lower compared to the transformed RRT variance for off-task reports similar to Seli et al. (2013). However, the linear mixed effect model showed that the report type, on or off-task, did not have a significant effect on the transformed RRT variance $(\chi^2(1) = 2.6505, p = 0.1035)$. Neither did the thought probe responses individually have a significant effect on the transformed RRT variance $(\chi^2(5) = 2.9005, p = 0.7153)$. Furthermore, the set also did not have significant effect on the transformed RRT variance $(\chi^2(3) = 2.0.5818, p = 0.9006)$.

To answer the main goal of the data analysis, we plotted the mean optimal choice for the first

3 sets after each jump with the mean thought-probe responses for answer option 2 also for the first 3 sets after each jump. Figure 6 depicts this plot. We predicted that the optimal choice over the three sets increases and that the mean thought-probe responses 2 decreases over the three sets. The graph shows that both the mean optimal choice and the mean thought-probe responses increase from the first set to the second set. However, the changes between the second set and the third set are stationary for both the mean optimal choice and the mean thought probe responses. The error bars for optimal choice get smaller over time while the error bars for thought probe responses get bigger. The generalized linear mixed effect model showed that set affected optimal choice ($\chi^2(3) = 11.62$, $p = 0.008804$). The mean optimal choice of the first set after a jump was 0.41 ± 0.15 lower compared to the mean optimal choice before a jump. This means that the performance of the participants was significantly lower after a jump, which indicates that they used a strategy to choose for the most optimal platform. It also shows that the performance of the participants did not improve after 2 sets (10 trials) after a jump. The set did not have a significant effect on thought probe choice ($\chi^2(3) = 1.1124$, $p = 0.7741$).

For the confirmation probe, only one participant mentioned being "not really confident" about the answer for the thought probe which was about external distractions, only once. This means that the participants were confident in the responses they provided for the thought-probes. The participants were aware of what they were thinking about during the MRT

4. Discussion

The central aim of this research was to test whether unexpected uncertainty in decisions can affect the degree and content of self-generated thought. We found that during the metronome response task, participants also spent some time thinking about the bandit task. However, there was no significant increase in the number of self-reports on thinking about the bandit task after a jump when uncertainty has increased. This is supported by figure 6 and the generalized linear mixed effects model. We did not find any significant effects, which is most likely due to the little number of observations since we had to exclude more than half of the participants from the analysis. Yet, figure 6. did show that while the performance increased over the sets, so did the reports on thinking about the decision task. This could mean that there is a relationship between performance on the decision task and thinking about it during the interruption task. It could also mean that participants needed more than 5 trials to notice the change in outcome probability and therefore reported to think about the bandit task more in the second set.

We also found a significant effect of set on the performance on the bandit task. The participants did perform worse on the bandit task after a jump compared to before the jump. If the participants were not using a strategy based on previous experiences, there would not have been a significant effect. However, due to the small number of observations, we cannot generalize this result.

Similar to the research by Seli, Cheyne & Smilek (2013) the transformed rhythmic response times (RRT's) variance for on-task reports was lower compared to the transformed RRT variance for off-task reports. This supports their hypothesis that during episodes of self-generated thinking, executive control resources are disconnected from the task at hand, which leads to a decreased focus on the current task, and hence, greater behavioural variability. However, the linear mixed effects model did not find any significant difference between the two report types, on and off-task, and their transformed RRT variances. Whereas Seli et al. did find a significant effect of report type on the transformed RRT variance.

The current experimental set-up was not successful in significantly increasing the number of self-generated thought reports on unexpected uncertainty in decisions. More than half the participants did not quite understand the task since they performed around chance level. Another possible interpretation is that they did not figure out that the possibility of successfully selling an item on a particular platform was independent of the reward indicated on that trial. Their behaviour could be explained by Behrens, Woolrich, Walton, and

Rushworth (2007). They found that their subjects often picked a "colour" (the equivalent of a selling platform) if it was associated with a higher reward even though it had a low outcome probability. The rewards could lead to confusion among participants. Knowing this, we still used the rewards to stay as close to the original experiment. Moreover, the varying rewards represented the variable offers one might receive for an item. If a reward was lower than that the participant expected, that could be disappointing. Just like when you receive a disappointing offer in real-life when you expected to receive more for your item. This lack of understanding could be caused by the instructions not being clear to all participants. Payzan-LeNestour & Bossaerts (2011) showed that when the instructions did not explicitly mention there would be jumps, the participants would not notice. Instructions are a crucial part in the understanding of the experiment.

An alternative issue could be that when someone mind-wanders about uncertainty in decisions it is more likely to be about one larger problem instead of small and similar problems. For example, for the study of Dijksterhuis, Bos, Nordgren, & Van Baaren. (2006) participants had to decide which car to buy out of four choices. The decisions were classified as simple or complex depending on the number of positive and negative attributes each car was given. Participants performing a distraction task made a better decision for the complex situation but not for the simple situation. Comparing picking a car with several positive and negative attributes, with choosing between two online-selling platforms, are quite different calibre problems.

Although multiple studies (Payzan-LeNestour (2010, 2011, 2013; Behrens, Woolrich, Walton, and Rushworth, 2007) used the restless bandit task to introduce unexpected uncertainty, it might be a too simple problem/task to mind-wander about. Complex decisions may need more of an emotional or personal aspect to increase self-generated thought about those decisions. Supported by Singer (1981), mind-wandering (self-generated thought) is closely related to current concerns and motivations. Jordano and Touron (2017) based their research on the "control failures × current concerns" idea by McVay and Kane (2010). Jordano and Touron demonstrated that "priming of personal, performance-related current concerns can increase mind-wandering." The decision task should relate more to current concerns from participants to increase reports of self-generated thought on the decisions.

We left out participants who scored lower than 60% on optimal choice in the bandit task in our data analyses. How could we increase the number of participants that understand the task?

The parameters of the current experimental set-up could be adjusted to improve the experiment for further research and increase the number of participants that understand the task. For example, the outcome probability of 75% might cause too much uncertainty. This could cause a larger part of the bandit task to be too uncertain and therefore too difficult. Participants could have missed the jumps and therefore could not create a good strategy for choosing the optimal platform. Increasing the outcome probability from 75% to 80% could lower the uncertainty. Another possible problem might be that the 5 bandit trials in a set are too few for the participants to figure out the optimal selling-platform. The number of bandit trials in a set could be increased to for example 7, to provide the participants with more time to figure out the optimal online-selling platform. Furthermore, thought probes are a subjective measure. The thought probes could be supported by using another measure, for example, eye-tracking.

Another possibility might be to replace the bandit task with another task but to keep the structure of two alternating tasks. A replacement task could be a navigation task. The benefit of a navigation task is that unexpected uncertainty could be manipulated in different ways. The bandit task consisted of monotonous changes. Besides, the navigation task could resemble real life more compared to the bandit task. Even though the bandit task was based on a real-life situation, selling items on unstable online-selling platforms, in real life sellers do not frequently receive immediate feedback on whether their item is successfully sold. Moreover, choosing an online selling-

platform in real life often also depends on experiences from other users. Whereas, during our experiment, they could only decide using their own experiences.

For future research, it could be interesting to investigate the relation between the performance on the decision task and the reports of self-generated thought about the decision task during the interruption task. It could be that the performance improved due to wandering about it.

In summary, the experimental set-up with the alternating restless bandit task with unexpected uncertainty and the metronome response task with thought probes did make participants report thinking about the decision task. However, it did not significantly increase reports after uncertainty was introduced. The hypothesis could neither be confirmed nor rejected. More research will be necessary to find out whether unexpected uncertainty in decision-making can affect the degree and content of self-generated thought.

References

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. Nature Neuroscience, 10(9), 1214–1221.

Bland, A. R., & Schaefer, A. (2012). Different Varieties of Uncertainty in Human Decision-Making. Frontiers in Neuroscience, 6, 85.

Christoff, K., Gordon, A., & Smith, R. (2011). The role of spontaneous thought in human cognition. In O. Vartanian & D. R. Mandel (Eds.), Contemporary topics in cognitive neuroscience series. Neuroscience of decision making (pp. 259-284). New York, NY, US: Psychology Press.

Dijksterhuis, A. (2004). Think Different: The Merits of Unconscious Thought in Preference Development and Decision Making. Journal of Personality and Social Psychology, 87(5), 586-598.

Dijksterhuis, A., Bos, M., Nordgren, L., & Van Baaren, R. (2006). On Making the Right Choice: The Deliberation-without-Attention Effect. Science, 311(5763), 1005-1007.

Bates, D., Maechler, M., Bolker, B., Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. Journal of Statistical Software, 67(1), 1-48.

Huettel, S. A., Song, A. W., & McCarthy, G. (2005). Decisions under Uncertainty: Probabilistic Context Influences Activation of Prefrontal and Parietal Cortices. Journal of Neuroscience, 25 (13) 3304-3311.

Huijser, S., van Vugt, M. K., & Taatgen, N. A. (2018). The wandering self: tracking distracting self-generated thought in a cognitively demanding context. Consciousness and cognition 58, 170-185.

Jordano, M. L, & Touron, D. R. (2017). Priming performance-related concerns induces task-related mind-wandering. Consciousness and cognition 10; 55: 126-135

Klinger, E. (2013). Goal Commitments and the content of thoughts and dreams: basic principles. Frontiers in Psychology, 4, 415.

Kuznetsova A, Brockhoff PB and Christensen RHB (2017). "lmerTest Package: Tests in Linear Mixed Effects Models." _Journal of Statistical Software_, *82*(13), pp. 1-26.

Mathôt, S., Schreij, D., & Theeuwes, J. (2012). OpenSesame: An open-source, graphical experiment builder for the social sciences. Behavior Research Methods, 44(2), 314-324.

McVay, J. C., & Kane, M. J. (2009). Conducting the Train of Thought: Working Memory Capacity, Goal Neglect, and Mind Wandering in an Executive-Control Task. Journal of Experimental Psychology. Learning, Memory, and Cognition, 35(1), 196–204.

McVay, J. C., & Kane, M. J. (2010). Does mind wandering reflect executive function or executive failure? Comment on Smallwood and Schooler (2006) and Watkins (2008). Psychological Bulletin, 136(2), 188-197.

Mooneyham, B. W., & Schooler, J. W. (2013). The costs and benefits of mind-wandering: a review. Canadian Journal of Experimental Psychology, Vol. 67, No. 1, 11–18.

Mustaq, F., Bland, A. R., and Schaefer, A. (2011). Uncertainty and cognitive control. Front. Psychol. 2:249.

Niedźwieńska, A., & Kvavilashvili, L. (2018). Reduced mind-wandering in mild cognitive impairment: Testing the spontaneous retrieval deficit hypothesis. Neuropsychology, 32(6), 711-723.

Payzan-LeNestour, E. (2010). Bayesian learning in unstable settings: experimental evidence based on the bandit problem. Swiss Finance Institute Research Paper No 10-28. pp 1-41.

Payzan-LeNestour E., & Bossaerts P. (2011). Risk, Unexpected Uncertainty, and Estimation Uncertainty: Bayesian Learning in Unstable Settings. PLoS Comput Biol 7(1): e1001048.

Payzan-LeNestour, E., Dunne, S., Bossaerts P., & O'Doherty, J., P. (2013). The neural representation of unexpected uncertainty during value-based decision making. Neuron, 79(1): 191-201.

Seli, P., Cheyne, J. A., & Smilek, D. (2013). Wandering Minds and Wavering Rhythms: Linking Mind Wandering and Behavioral Variability. Journal of Experimental Psychology: Human Perception and Performance Vol. 39, No. 1, 1–5.

Singer, J., L. (1981). Daydreaming and Fantasy. Oxford, UK: Oxford University Press.

Smallwood J., Davies J., B., Heim, D., Finnigan, F., Sudberry, M., O'Connor, R., & Obonsawin, M. (2004). Subjective experience and the attentional lapse: task engagement and

disengagement during sustained attention. Conscious Cogn.;13(4), 657-90.

Unsworth, N., & Robison, M. K. (2016). Pupillary correlates of lapses of sustained attention. Cognitive, Affective, & Behavioral Neuroscience, 16(4), 601-615.

Unsworth, N., & Robison, M. K., (2018). Tracking arousal state and mind wandering with pupillometry. Cognitive, Affective, & Behavioral Neuroscience 18, 638–664.

Vannuccia, M., Pelagattia, C., & Marchettib, I. (2017). Manipulating cues in mind wandering: Verbal cues affect the frequency and the temporal focus of mind wandering. Consciousness and Cognition 53, 61–69.

Whittle, P. (1988). Restless Bandits: Activity Allocation in a Changing World. Journal of Applied Probability, 25, 287-298.

Yu, A., J., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. Neuron 46, 681-692.