Bachelor's thesis - Comparing the predictiveness of reward learning and mind-wandering in depression

Bachelor's Project Thesis

Boris Benjamin Winter, b.b.winter@student.rug.nl,
Supervisor: Dr. M.K. van Vugt

Abstract: Depression is usually measured through either a self-report questionnaire or through a structured clinical interview with an expert, with both options having their own advantages and limitations. Different behavioral measures of depression have been proposed, such as reward learning and mind-wandering, but these have never been compared. This study compared the correlations between self-report depression questionnaire scores and data from tasks that quantified reward learning and mind-wandering. Reward learning was quantified through a task based on signal detection theory, and mind-wandering data came from a Sustained Attention to Response Task (SART). Both tasks contained aspects that correlated significantly with depression questionnaire scores, and could therefore be used in the prediction of these scores. Neither of the tasks could be used to predict the scores from all questionnaires and both cognitive functions seem to have their own strengths when trying to predict depression.

1 Introduction

Major Depressive Disorder (MDD), commonly known as depression, is one of the most common health issues in the world, with an estimated 4.4% of the world population (approximately 322 million people) suffering from it in 2015 (World Health Organization, 2017). Depressive disorders like MDD can and often do cause severe cognitive dysfunctions during episodes of depression (see Marazziti et al. (2010) for a summary), and recently it has been suggested that they may even leave certain cognitive functions impaired for much longer than the duration of a depressive episode (Hammar & Årdal, 2009). Examples include a study that found a remaining difficulty in visual search with participants after 6 months, even though their depression severity lowered (Hammar et al., 2003), and another study that found a long lasting impairment in sustained attention with patients in remission (Majer et al., 2004). While the co-occurrence of cognitive dysfunctions and depression is unfortunate, the link between them might also be a helping factor in recognizing and predicting depression. Rumination, the act of continuously thinking about the same (mostly negative) things, has for example been suggested to predict the presence and new onsets of major depressive disorders (Nolen-Hoeksema, 2000).

The huge prevalence of MDD, combined with the significant influences that depressive disorders can have on someone's life, make it critical to try to understand the disease, find cures for it, and most importantly: To find methods that can predict the presence of or predisposition to MDD, resulting in the ability to respond more quickly to an onset or relapse of depression.

It has been suggested that someone's ability to learn through reward learning is correlated with the severity of their depression. People that score low on depression severity assessments seem to do much better at learning through rewards than people that score high on depression severity assessments. In two studies that tested this phenomenon, participants were asked to respond to the size of a smileys mouth by pressing one button for a short mouth and another for a long mouth. When one mouth size was rewarded three times more often than the other mouth size, participants without an

MDD diagnosis developed a tendency to choose the more rewarded mouth size, while this phenomenon did not occur in participants with MDD (Pizzagalli et al., 2005; Vrieze et al., 2013).

Mind-wandering, the act of thinking about something else than what one is doing at the moment, also seems to correlate with depression severity. Mind-wandering in general appears to correlate with unhappiness (Killingsworth & Gilbert, 2010; Smallwood et al., 2009), and a person's tendency to engage in mind-wandering has been suggested to correlate with the presence of depressive thoughts (Smallwood et al., 2007). While mind-wandering is generally seen as something negative, it does have its benefits, and the effects of mind-wandering seem to depend largely on both context and content (Smallwood & Andrews-Hanna, 2013). Recently, it has been suggested that mind-wandering correlates with worse moods and depression specifically when it consists of ruminative and worrying thoughts (Ottaviani et al., 2015, 2013).

Given the correlations between depression and the cognitive functions reward learning and mind-wandering, the ability to learn through rewards and the tendency to engage in mind-wandering are both potential predictors for depression. Working towards the ultimate goal of finding practical methods for diagnosing depression, this study will focus on comparing the predictive ability of both cognitive functions by asking the question: Which cognitive function, reward learning or mind-wandering, can more accurately predict depression? It is expected that both functions will contain aspects that correlate with depression, but there is no expectation as to which function is of better use when predicting depression.

2 Methods and Materials

2.1 Participants

Participants were initially recruited through a so-called HIT (Human Intelligence Task) on Amazon Mechanical Turk (also known as MTurk) (https://www.mturk.com), which entailed filling out a survey. To qualify for the HIT, MTurk Workers had to have been granted the Master's qualification and had to be located in the United States (for more information about MTurk and an overview of the lessons learned about using it for this study, please refer to appendix A). In total, 139 people filled out the initial survey. From this group, 107 people (77.0%) were approved based on criteria like answering catch-questions correctly and finishing the survey in a reasonable amount of time. Because the end goal was to compare how well reward learning and mind-wandering could predict depression separately, two follow-up tasks (one measuring reward learning and one measuring mind-wandering) were created for the approved participants. Eventually, 41 participants (38.3%) took part in both tasks. As a comparison could only be made for participants who took part in both tasks, these 41 participants are seen as the actual participants from here on. The sample consisted of 12 men and 29 women, aged 43.29 ± 11.15 (between 29 and 69). Nine participants had been treated for depression at one point in their life and five had been treated for some other mental illness. Participants were paid \$3.00 for doing the mind-wandering task, which had an average completion time of 27 minutes and 56 seconds. For the reward learning task, participants were given a base reward of \$1.00 and earned a bonus reward of \$0.02 for each correct trial that had a reward connected to it, which averaged to a bonus of \$1.92. The reward learning task was completed in an average time of 13 minutes and 37 seconds. Participants who finished both tasks were given an additional reward of \$2.00. During recruitment, special attention was paid to self-report questionnaire scores in order to ensure that the sample consisted of participants with a wide range of scores, which came about organically.

2.2 Tasks and Procedure

The general procedure for each participant was set up as follows: Participants would find the survey HIT on MTurk if they qualified for it and would choose to take part in the HIT, which meant filling out the survey. The survey started with a general information and consent form. After this, participants were asked to list five important and recent accomplishments/happy moments and five important and recent concerns in their life, which were (unbeknownst to the participants at this time) later used in the mind-wandering task to stimulate mind-wandering. To get an indication of the participant's depression severity that was more general than what one self-report questionnaire would provide, each participant was asked to fill out three different questionnaires (questionnaire names were not used in the survey in order to avoid influencing participants' answers): The Ruminative Response Scale (Nolen-Hoeksema & Morrow, 1991), the Perseverative Thinking Questionnaire (Ehring et al., 2011), and the Beck Depression Inventory 2 (Beck et al., 1996) respectively.

Within a few days of filling out the survey, participants were invited through MTurk to make both the reward learning and the mind-wandering follow-up task. Both tasks ran on a server from the University of Groningen with JATOS software (Lange et al., 2015), which allowed the tasks to be accessed through a web browser. Both tasks could be accessed through a link posted on MTurk. In order to account for sequence effects, the task with which participants started was counterbalanced across all invited participants (but this was not enforced, meaning that participants were able to make the other task first without the researcher's knowledge). Creation of these counterbalanced groups was done while trying to keep an approximately equal distribution of questionnaire scores.

2.2.1 Reward learning task

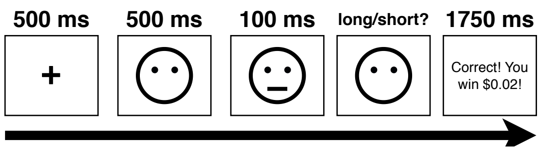

Figure 2.1: Schematic representation of the reward learning task. Participants had to press the 'z' key or the 'm' key (counterbalanced across participants) based on if they saw a short or a long mouth. The last image ("Correct! You win $0.02!") was only shown if the participant answered correctly and the trial had a reward connected to it.

In order to measure the participants' reward learning performance, a well-established signal detection task by Pizzagalli et al. (2005) was used (see also Pizzagalli et al. (2008); Vrieze et al. (2013)). This task is an adapted version of the task in Tripp and Alsop (1999) and the workings of the task are explained in the current section. To avoid on-screen distractions, the task was automatically started in full screen mode. On the screen, the participants were first presented with written instructions, which can be found in appendix B. After clicking on "I understand", the instructions were followed by a new screen that started with a demonstration of a trial. The demonstration was given through still images (appendix C) and participants could in their own time go to the next image by pressing any key. Two practice trials were performed after the demonstration, one for each stimulus. After these practice trials, a message was displayed on screen stating that the experiment was about to begin, and asking the participant to press any key once they were ready to start. Each trial went as follows: A fixation point was presented in the middle of the screen for 500 ms. Then, a mouthless emoticon would appear for another 500 ms. On this emoticon, one out of two possible stimuli would be shown. This stimulus was a mouth of either 11.5 mm (denoted as a small mouth) or 13 mm (denoted as a large mouth). To make sure that participants kept their attention on the task, the mouth was only shown for 100 ms each trial. The participant then had to press either the 'z' key or the 'm' key on the keyboard based on the mouth that they think they saw, which was counterbalanced across subjects. After the participant had pressed one of the two keys, there were three possible scenarios: Firstly, they could have chosen the wrong mouth. In this case, no feedback was given and the next trial was started. If they had chosen the correct answer, two things could happen: They could either get no feedback and start with the next trial (just as if they had chosen wrong), or feedback was given in the form of a message on the screen, stating: "Correct! You win $0.02!". It is important to emphasize that participants would therefore not always get feedback and be rewarded when they answered correctly. The awarding of rewards was done through a predetermined semi-random schedule in which reward trials were selected beforehand. Each participant performed a total of 300 trials, divided into three blocks of 100 trials. A 30 second break occurred between each block with a warning to get ready for the next block after 25 seconds. After completing the last trial, a screen was shown notifying the participant of the end of the experiment. For each block, rewards were set to be given for 40 out of the 100 trials, and were unevenly divided between the two mouths in order to create a response bias. One mouth (the *rich* stimulus) would have 30 reward trials per block while the other mouth (the *lean* stimulus) would only have 10 reward trials per block. For counterbalancing purposes, both the short and the long mouth were chosen as the rich stimulus for half of the participants. If a trial was scheduled to yield a reward but was not answered correctly, the reward would be postponed to the next trial with the same stimulus type.

2.2.2 Mind-wandering task

Participants' tendency to engage in mind-wandering was measured through a Sustained Attention to Response Task (SART) (Robertson et al., 1997) as this is a common method of measuring mind-wandering (e.g. McVay & Kane, 2009; Smallwood et al., 2004). The task was automatically initialized in full screen mode to avoid on-screen distractions. During the task, participants would be shown a generic word from a randomly selected set of English words, and were to press the space bar if the word was lower case, or do nothing if the word was upper case. At the start of each word trial, participants would see a fixation point in the middle of the screen for 1000 ms. Then, a word stimulus appeared on the screen for 500 ms. In order to prevent recall of the word (and thereby require the participant to pay attention if they wanted to answer correctly), this was followed by a mask in the form of a horizontal line of X's, which was presented for 500 ms. Lastly, the mask disappeared and an empty screen would be shown for 1000 ms. The goal for participants during each trial was to respond as quickly and accurately as possible to the word stimulus. Intermittently throughout the word trials, question blocks would come up, asking the following three multiple choice questions: "What were you thinking about just now?", "If you were not thinking about the task itself, what was the content of your thought?", and "How difficult was it to disengage from the thought?". The current paper refers to these questions as the on-task-, valence-, and sticky question respectively. The questions and their answer choices can be found in appendix D. To always allow for some time in which the participant could engage in mind-wandering, all question blocks would be separated by at least 8 word trials. All question blocks asked the same three questions in the same sequence.

The sequence of the task was as follows: Participants were first reminded of the fact that they named five achievements/happy moments and five concerns in the initial survey. Participants were asked to try to remember these words and told that they would be asked to name these as quickly as possible later in the task (which was not the case). It was desired that participants had these ten items in their short-term memory, as this would stimulate mind-wandering for all participants during the task. A short summary of why and how this would be the case can be found in Smallwood and Schooler (2006). The task continued with more instructions (appendix D), followed by practice trials. The practice trials contained two lower case words, two upper case words, and ended with a question block. After the practice trials were over, the participant could start the main experiment by pressing any key on the keyboard. The main experiment contained 540 word trials, 30 question blocks, and no breaks.

2.3 Data analysis

2.3.1 Reward learning analysis

In order to make the results comparable with previous research (Pizzagalli et al., 2005) and to remove trials where the response time was unusually fast or slow, trials with a response time lower than 150 ms and higher than 2500 ms were excluded from the data. Additionally, participants with more than 30 outlier trials (10% of total) were completely excluded from the reward learning and mind-wandering experiments (Pizzagalli et al., 2005). This was the case for eight participants. Like in the original experiment, the response bias (RB) was identified as the main variable of interest. The response bias variable signifies the participants' tendency to choose the stimulus that is rewarded more frequently. Response bias therefore measures reward learning, with a higher response bias indicating better learning through rewards (Pizzagalli et al., 2005). Calculation of the response bias was done in the same way as in the original experiment by Pizzagalli et al. (2005), who derived the formula from the behavioral model of signal detection (e.g. McCarthy & Davison, 1979; Tripp

& Alsop, 1999). To allow for the calculation of the response bias when one of its factors equals zero, .5 was added to each factor (Pizzagalli et al., 2008).

Response bias:

$$log\ b = \frac{1}{2} log \left(\frac{(R_{correct} + .5)(L_{incorrect} + .5)}{(R_{incorrect} + .5)(L_{correct} + .5)} \right) \quad (2.1)$$

In equation 2.1, R represents the rich stimulus (the stimulus with a higher reward probability) and L represents the lean stimulus (the stimulus with a lower reward probability).

2.3.2 Mind-wandering analysis

In order to exclude data from participants that did not appear to take the mind-wandering task seriously, participants with an accuracy below 65% for the word trials were excluded from the reward learning and mind-wandering experiments, which was the case for two participants. The main variable of interest for the mind-wandering task was the proportion of choices for each answer to the three questions in the question block. For each participant, the number of choices for each answer were recorded and Pearson correlation coefficients were calculated between the proportion of choices for an answer and questionnaire scores. The answers that correlated the highest with questionnaire scores were chosen as the variables that were to be used in the comparison.

2.3.3 Comparison

With the goal of comparing the cognitive functions based on how well they predict depression, the most highly correlating variables from both tasks were used in (multiple) linear regression in order to predict the three depression questionnaire scores (RRS, PTQ, and BDI-2). Linear regression was chosen to be used for the prediction as this simple form of regression can act as a basis to work from and compare to in the future. To have one score as a more general measurement of depression, a combination score (COMB) was created in addition to the questionnaire scores by summing up all three scores, and was also to be predicted. For each of the four depression scores, a linear model was made using data from each task, and these models were thereafter compared by how well they fit the data, based on adjusted R^2 and AIC values.

3 Results

3.1 Depression scores

To examine to what extent the different questionnaires tap into related concepts (i.e. if people who score high on one questionnaire also score high on the other questionnaires), pairwise Pearson correlation coefficients were calculated, an overview of which can be found in table 3.1. All scores were found to have significant positive correlations with each other, but the correlations between the BDI-2 score and other scores are much less extreme than the other correlations.

	PTQ	**BDI-2**	**COMB**
1. RRS	$.77^a$	$.43^b$	$.89^a$
2. PTQ		$.40^b$	$.87^a$
3. BDI-2			$.72^a$
4. COMB			

Table 3.1: Pairwise Pearson correlations between depression scores of the (1) Ruminative Response Scale, (2) Perseverative Thinking Questionnaire, (3) Beck Depression Inventory 2, and the (4) combination score ($^a p < .0005$, $^b p < .05$).

3.2 Reward learning

3.2.1 Change in response bias

To examine the learning through rewards and to disregard any preexisting biases, each subject's change in response bias ($\triangle RB$) was calculated (rather than their final response bias) for three different intervals: Block1 to block2, block2 to block3, and block1 to block3 (Pizzagalli et al., 2005). For each interval, change in response bias was plotted against all three questionnaire scores and the combination score. Pearson correlation coefficients were calculated to determine which interval had the strongest correlation with the depression scores. For the first interval (block1-block2), Pearson correlation tests revealed significant negative correlations between $\triangle RB_{1-2}$ and RRS score ($r(29)$ = -.39, p = .032), BDI score ($r(29)$ = -.37, p = .043), and COMB score ($r(29)$ = -.43, p = .015). PTQ score also had a negative correlation with $\triangle RB_{1-2}$, but failed to be significant ($r(29)$ = -.33 p = .073).

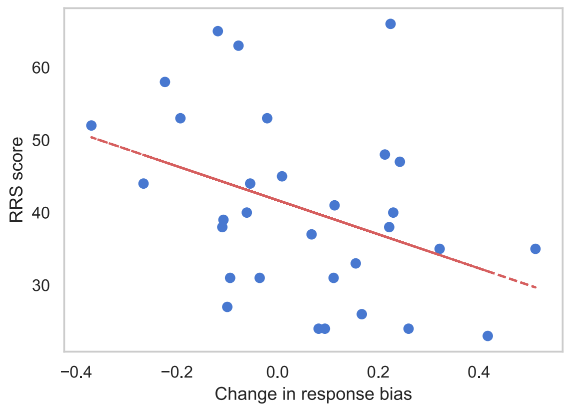

(a)

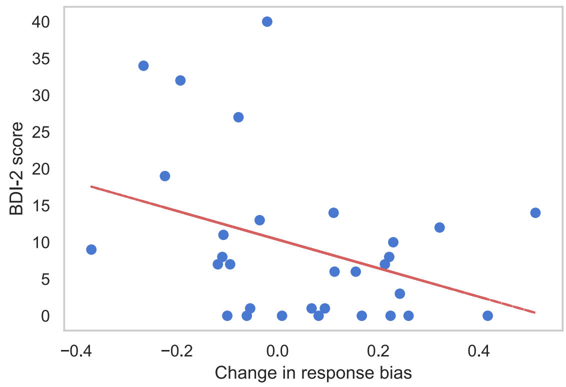

(b)

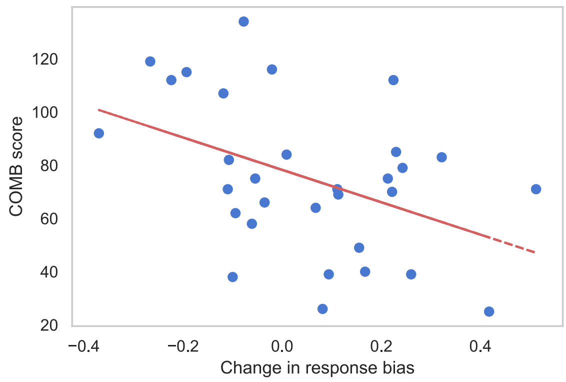

(c)

Figure 3.1: The relation between the (a) RRS, (b) BDI-2, and (c) COMB score, and the change in response bias early in the task (between block 1 and block 2).

Scatterplots of the significant results can be found in figure 3.1. In the second interval (block2-block3), change in response bias was actually found to correlate positively with depression scores, though only the correlation between $\triangle RB_{2-3}$ and RRS score showed significance ($r(29) = .37$, $p = .039$). This suggests that participants with higher depression scores did develop a response bias, but developed it later than participants with lower scores. No strong correlations between $\triangle RB_{1-3}$ and depression scores were found in the third interval. These results indicate that of all intervals, the correlation between $\triangle RB_{1-2}$ and the depression scores is the strongest.

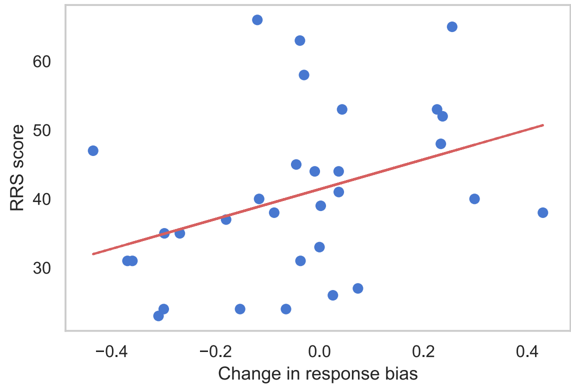

Figure 3.2: The relation between RRS score and the change in response bias late in the task (between block 2 and block 3).

3.2.2 Score distribution by interval of highest response bias development

To determine whether participants with higher depression scores in fact developed a response bias later than participants with lower scores, subjects were divided into two groups based on the interval in which they experienced their highest response bias development (positive change), which could be either block1-block2 or block2-block3. The score distributions for each questionnaire were plotted for both groups and are shown in figure 3.3. To determine whether the difference in depression scores between the two groups was significant, independent t-tests were performed. This revealed a significant difference in RRS score between the subjects that

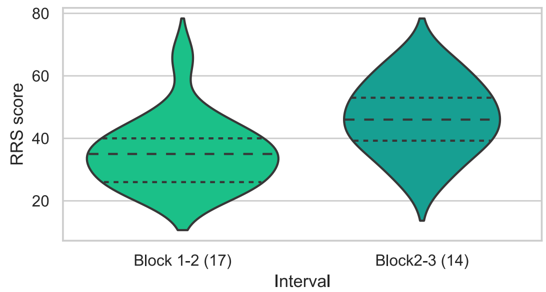

(a)

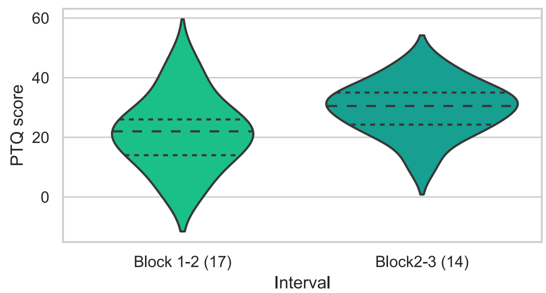

(b)

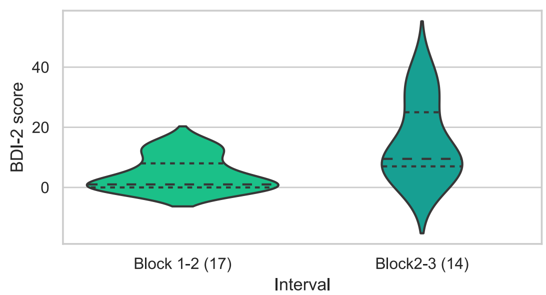

(c)

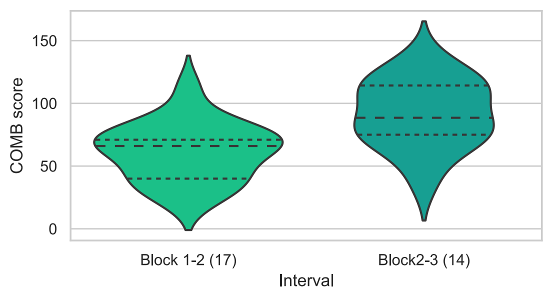

(d)

Figure 3.3: Differences in (a) RRS, (b) PTQ, (c) BDI-2, and (d) COMB score, between participants who developed their response bias in earlier and later blocks of the task (number of subjects).

developed their response bias mostly early in the task (block1-block2) and subjects that developed their response bias mostly late in the task (block2-block3) ($t(29) = 2.87$, $p = .008$). Similar results were found for the same comparison with regards to PTQ score ($t(29) = 2.13$, $p = .042$), BDI-2 score ($t(29) = 3.05$, $p = .005$), and COMB score ($t(29) = 3.39$, $p = .002$).

3.3 Mind wandering

3.3.1 SART performance

Neither a subject's average response time nor accuracy over the whole task was found to have a significant correlation with any of the depression scores. To examine the relation between response time, accuracy, and question answers, the average response time and accuracy for each five word trials preceding a question set were calculated. A one-way ANOVA revealed no significant correlations between answers given to the on-task/valence/sticky question and the accuracy or response time preceding a question set.

3.3.2 Correlations between SART answers and depression scores

Pearson correlation coefficients were calculated to assess the relationships between the proportion of answers to each on-task, valence, and sticky question, and the depression scores. Significant results found were the following: For the on-task question, the proportion of answers indicating that the participant was daydreaming (option 5) correlated positively with RRS score ($r(29) = .43$, $p = .015$). For the valence question, the proportion of answers indicating negative self-related thought (option 2) had a positive correlation with PTQ score ($r(29) = .40$, $p = .025$). The proportion of self-related thoughts in general (positive and negative, options 1 and 2) correlated with PTQ score as well ($r(29) = .41$, $p = .021$). Answer proportions to the sticky question showed no significant correlation with any of the depression scores.

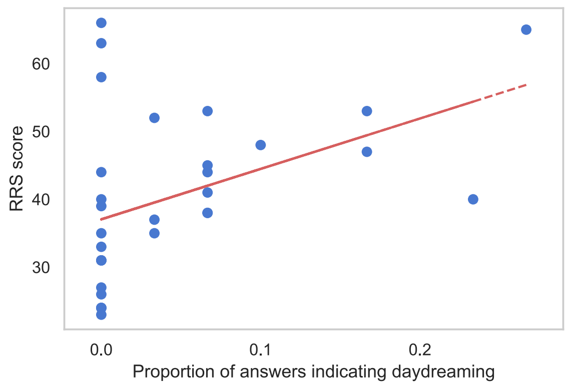

Figure 3.4: The relation between RRS score and the proportion of answers to the on-task question that indicate daydreaming.

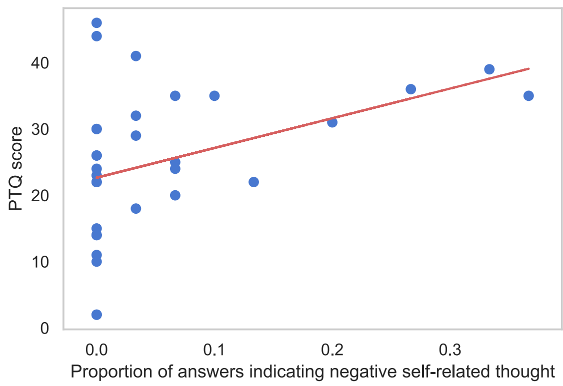

(a)

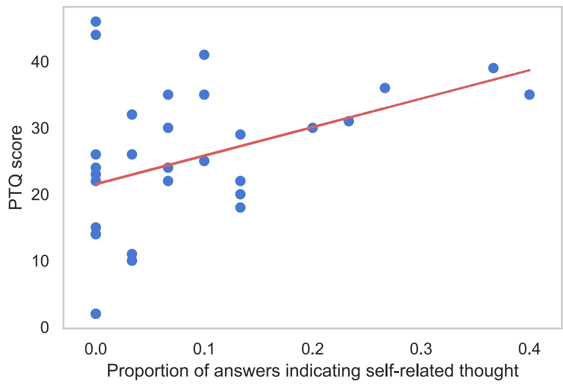

(b)

Figure 3.5: The relation between PTQ score and (a) the proportion of answers to the valence question indicating negative self-related thought, and (b) self-related thought in general.

3.4 Comparing score prediction between reward learning and mind-wandering

To answer the question of which cognitive function, reward learning or mind-wandering, can more accurately predict depression, this section compares both tasks based on how well their most highly correlating variable(s) can predict the questionnaire scores and combination score through (multiple) linear regression. For the reward learning task, the change in response bias early in the task ($\triangle RB_{1-2}$) was used as the independent variable, as this was shown to correlate the strongest with all four scores. For the mind-wandering task, the proportion of answers to the on-task question indicating daydreaming (hereafter referred to as DAYDREAMING) and the proportion of answers to the valence question indicating self-related thought (SELF-RELATED) were compared through a stepwise regression for each score independently, in order to determine the best combination (either one of them or both) for predicting each depression score. The best combination for each depression score was then used in the comparison with $\triangle RB_{1-2}$. The linear models were compared based on their adjusted R^2 (R^2_{adj}) and AIC values. R^2_{adj} was used instead of R^2 because it allows for a fairer comparison when the models use a different number of independent variables. To compare the models on more than one aspect, AIC values were compared as well, as AIC focuses more on parsimony while R^2_{adj} focuses more on predictive power.

3.4.1 Ruminative Response Scale (RRS)

In order to determine whether reward learning or mind-wandering is better at predicting RRS scores, linear models based on each task's data were compared. With the mind-wandering task, RRS scores were best predicted through a linear model using both DAYDREAMING and SELF-RELATED ($R^2_{adj}(2{,}28) = .21$, $p = .014$, AIC $= 151.35$). A linear model based on $\triangle RB_{1-2}$ explained less variance than the mind-wandering model ($R^2_{adj}(1{,}29) = .12$, $p = .032$, AIC $= 153.86$). Based on the higher R^2_{adj} and lower AIC value, mind-wandering seems to be more useful for predicting RRS scores than reward learning. Interestingly, a linear model combining reward learning and mind-wandering showed the best fit ($R^2_{adj}(3,27) = .27$, $p = .009$, AIC = 149.88).

3.4.2 Perseverative Thinking Questionnaire (PTQ)

When predicting PTQ scores, a linear model using only SELF-RELATED showed the best results with the mind-wandering data ($R^2_{adj}(1,29) = .14$, $p = .021$, AIC = 146.89). The linear model based on $\triangle RB_{1-2}$ explained less variance ($R^2_{adj}(1,29) = .08$, $p = .073$, AIC = 149.19), and was not significant. Again, a combination between reward learning and mind-wandering resulted in the best fit ($R^2_{adj}(2,28) = .17$, $p = .027$, AIC = 146.73). R^2_{adj} and AIC values from the separate models were compared in order to determine whether mind-wandering or reward learning better predicts PTQ scores. A higher R^2_{adj} and a lower AIC value from the mind-wandering model suggested that the mind-wandering model resulted in a better fit and is therefore better at predicting PTQ scores than reward learning.

3.4.3 Beck Depression Inventory 2 (BDI-2)

No combination of DAYDREAMING and SELF-RELATED resulted in a model that could predict BDI-2 scores (all $R^2_{adj} < 0$). A linear model based on $\triangle RB_{1-2}$ did find significant results ($R^2_{adj}(1,29) = .10$, $p = .043$, AIC = 146.28). As mind-wandering could not be used to predict BDI-2 scores, reward learning is clearly the better predictor for BDI-2 scores.

3.4.4 Combination of scores (COMB)

Using the mind-wandering data, the best fit based on R^2_{adj} was found through a linear model using both DAYDREAMING and SELF-RELATED, but this result did not reach significance ($R^2_{adj}(2,28) = .07$, $p = .138$, AIC = 207.85). The linear model based on $\triangle RB_{1-2}$ did show significance $R^2_{adj}(1,29) = .16$, $p = .015$, AIC = 204.53). A linear regression using both mind-wandering and reward learning data resulted in a slightly higher R^2_{adj} than the model with only $\triangle RB_{1-2}$, but showed a worse p- and AIC value ($R^2_{adj}(3,27) = .17$, $p = .043$, AIC = 205.76). To determine whether reward learning or mind-wandering is the better predictor of COMB scores, R^2_{adj} and AIC values from the separate models were again compared. It was found that (in addition to actually being significant) the reward learning model resulted in a better fit than the mind-wandering model. This suggests that reward learning is of better use than mind-wandering when predicting the combination score.

4 Discussion

This study wanted to compare two cognitive functions, reward learning and mind-wandering, by how accurately tasks that measure these functions can predict the presence of depression. Based on previous research (e.g. Pizzagalli et al., 2005; Smallwood et al., 2007), it was hypothesized that tasks measuring these functions would generate data that would correlate with depression scores from self-report questionnaires. The most highly correlating variables from each task were to be used in linear models in order to predict the questionnaire scores. As no comparison had been made before, no expectations were set for which of the two cognitive functions would be the better predictor, but it was expected that both functions would have a predictive ability.

As hypothesized, both tasks contained aspects that correlated with depression scores. Data from the mind-wandering task revealed positive correlations between Ruminative Response Scale (RRS) and Perseverative Thinking Questionnaire (PTQ) scores, and the proportion of certain answers to the questions in the task. The reward learning task showed results that were only partly similar to previous findings. As expected, change in response bias early on in the task correlated negatively with most depression scores. However, previous research also found that participants' response bias changes much less in the second half of the task (Pizzagalli et al., 2008, 2005). Instead of reproducing these results from the second half of the task, this study found an exact opposite (but weaker) correlation to the correlation that occurred early on in the task. This indicated that participants with higher depression scores developed their response bias later than participants with lower scores, instead of developing it less. One reason for this might be that people with higher scores are completely capable of developing the same response bias, but need more reinforcement than people with lower scores in order to do so.

It was found that both reward learning and mind-wandering could predict aspects of depressive thinking, but neither one of the functions could be used to predict all four depression scores through linear regression. Mind-wandering appeared to be a better predictor for depression in participants with more rumination-type symptoms, as the mind-wandering models explained more variance than reward learning models when predicting the RRS and PTQ scores. At the same time, mind-wandering could not be used to predict Beck Depression Inventory 2 (BDI-2) scores at all and a model for predicting the combination score was not significant. Reward learning, on the other hand, was able to predict RRS, BDI-2, and combination scores. It also found a predictive model for PTQ scores but that model was insignificant. Therefore, while reward learning was a worse predictor than mind-wandering for the RRS and PTQ scores, results indicated that it can be used to predict a wider variety of questionnaire scores. Because of this, neither one of the cognitive functions was found to be the better predictor of depression, but each one seems to have its own speciality. One of the reasons why this might be the case is that the depression questionnaires partly focus on different aspects of depression. For example, the Ruminative Response Scale focuses more on rumination, making it almost expected that a closely related cognitive function like mind-wandering is better at predicting RRS scores than a less closely related function. These findings also suggest that mind-wandering might be better at predicting depression with patients who tend to exhibit more ruminative type symptoms during depressive episodes.

This study had several limitations. For starters, the use of Amazon Mechanical Turk (MTurk) for the recruitment of participants should be questioned. While MTurk provided a good and very quick way to recruit participants through the initial survey, the use of MTurk also meant that participants made both follow-up tasks at home, resulting in all tasks being made in an uncontrolled environment. As both tasks were cognitive tasks that required the participant to react in a timely manner, distractions coming from an uncontrolled environment could have influenced the task results. In the reward learning task, distractions might cause participants to pick the incorrect stimulus more often, thereby influencing the measured response bias. Distractions during the mind-wandering task will have an influence on the answers, as the questions are directly related to distraction. Secondly, it should be noted that the Bayes factor analyses, which were performed as a supplement to every calculation and can be found in appendix E, mostly resulted in a Bayes factor that was not extreme enough to make definite conclusions (between 0.3 and 3). This indicated that more data is needed before the results found in this study can be trusted completely.

Further research should focus on performing this experiment on a larger scale in order to determine whether the results found in the current study hold up. As this study only included participants from the United States, the experiment should also be repeated with participants from outside the United States in order to verify generalisability. Additionally, the influences of using an online platform for this study should be investigated by also performing the experiment in a controlled environment. Moreover, future studies should use and compare more advanced machine learning techniques in order to find better methods of prediction, and recompare the usefulness of both tasks' data based on those findings.

In sum, the current study reaffirmed that significant correlations exist between depression and the cognitive functions reward learning and mind-wandering. It also found that these correlations can be used in order to predict depression. Neither one of the functions is an obviously better predictor, and both functions should be investigated further to find out more about their usefulness in predicting depression.

References

Beck, A. T., Steer, R. A., & Brown, G. K. (1996). Beck depression inventory-ii. *San Antonio*, *78*(2), 490–498.

Casler, K., Bickel, L., & Hackett, E. (2013). Separate but equal? a comparison of participants and data gathered via amazon's mturk, social media, and face-to-face behavioral testing. *Computers in Human Behavior*, *29*(6), 2156-2160. doi: https://doi.org/10.1016/j.chb.2013.05.009

Ehring, T., Zetsche, U., Weidacker, K., Wahl, K., Schönfeld, S., & Ehlers, A. (2011). The perseverative thinking questionnaire (ptq): Validation of a content-independent measure of repetitive negative thinking. *Journal of Behavior Therapy and Experimental Psychiatry*, *42*(2), 225-232. doi: https://doi.org/10.1016/j.jbtep.2010.12.003

Getting great survey results from MTurk and Qualtrics. (2017). Retrieved January 15th, 2020, from `https://blog.mturk.com/getting-great-survey-results-from-mturk-and-qualtrics-be1704ff9786`

Hammar, A., Lund, A., & Hugdahl, K. (2003). Long-lasting cognitive impairment in unipolar major depression: a 6-month follow-up study. *Psychiatry Research*, *118*(2), 189-196. doi: https://doi.org/10.1016/S0165-1781(03)00075-1

Hammar, A., & Årdal, G. (2009). Cognitive functioning in major depression-a summary. *Frontiers in Human Neuroscience*, *3*, 26. doi: https://doi.org/10.3389/neuro.09.026.2009

Killingsworth, M. A., & Gilbert, D. T. (2010). A wandering mind is an unhappy mind. *Science*, *330*(6006), 932. doi: https://doi.org/10.1126/science.1192439

Lange, K., Kühn, S., & Filevich, E. (2015). "just another tool for online studies"(jatos): An easy solution for setup and management of web servers supporting online studies. *PloS one*, *10*(6). doi: https://doi.org/10.1371/journal.pone.0130834

Majer, M., Ising, M., Künzel, H., Binder, E. B., Holsboer, F., Modell, S., & Zihl, J. (2004). Impaired divided attention predicts delayed response and risk to relapse in subjects with depressive disorders. *Psychological Medicine*, *34*(8), 1453-1463. doi: https://doi.org/10.1017/S0033291704002697

Marazziti, D., Consoli, G., Picchetti, M., Carlini, M., & Faravelli, L. (2010). Cognitive impairment in major depression. *European Journal of Pharmacology*, *626*(1), 83-86. doi: https://doi.org/10.1016/j.ejphar.2009.08.046

McCarthy, D., & Davison, M. (1979). Signal probability, reinforcement and signal detection. *Journal of the Experimental Analysis of Behavior*, *32*(3), 373–386. doi: https://doi.org/10.1901/jeab.1979.32-373

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. doi: https://doi.org/10.1037/a0014104

Nolen-Hoeksema, S. (2000). The role of rumination in depressive disorders and mixed anxiety/depressive symptoms. *Journal of Abnormal Psychology*, *109*(3), 504-511. doi: https://doi.org/10.1037/0021-843X.109.3.504

Nolen-Hoeksema, S., & Morrow, J. (1991). A prospective study of depression and posttraumatic stress symptoms after a natural disaster: the 1989 loma prieta earthquake. *Journal of Personality and Social Psychology*, *61*(1), 115. doi: https://doi.org/10.1037/0022-3514.61.1.115

Ottaviani, C., Shahabi, L., Tarvainen, M., Cook, I., Abrams, M., & Shapiro, D. (2015). Cognitive, behavioral, and autonomic correlates of mind wandering and perseverative cognition in major depression. *Frontiers in Neuroscience*, *8*, 433.

Ottaviani, C., Shapiro, D., & Couyoumdjian, A. (2013). Flexibility as the key for somatic health: From mind wandering to perseverative cognition. *Biological Psychology*, *94*(1), 38-43. doi: https://doi.org/10.3389/fnins.2014.00433

Pizzagalli, D. A., Iosifescu, D., Hallett, L. A., Ratner, K. G., & Fava, M. (2008). Reduced hedonic capacity in major depressive disorder: evidence from a probabilistic reward task. *Journal of Psychiatric Research*, *43*(1), 76-87. doi: https://doi.org/10.1016/j.jpsychires.2008.03.001

Pizzagalli, D. A., Jahn, A. L., & O'Shea, J. P. (2005). Toward an objective characterization of an anhedonic phenotype: a signal-detection approach. *Biological Psychiatry*, *57*(4), 319-327. doi: https://doi.org/10.1016/j.biopsych.2004.11.026

Pricing. (n.d.). Retrieved January 14th, 2020, from `https://www.mturk.com/pricing`

Robertson, I. H., Manly, T., Andrade, J., Baddeley, B. T., & Yiend, J. (1997). Oops!': performance correlates of everyday attentional failures in traumatic brain injured and normal subjects. *Neuropsychologia*, *35*(6), 747–758.

Smallwood, J., & Andrews-Hanna, J. (2013). Not all minds that wander are lost: the importance of a balanced perspective on the mind-wandering state. *Frontiers in Psychology*, *4*, 441. doi: https://doi.org/10.3389/fpsyg.2013.00441

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. *Consciousness and Cognition*, *13*(4), 657–690. doi: https://doi.org/10.1016/j.concog.2004.06.003

Smallwood, J., Fitzgerald, A., Miles, L. K., & Phillips, L. H. (2009). Shifting moods, wandering minds: negative moods lead the mind to wander. *Emotion*, *9*(2), 271. doi: https://doi.org/10.1037/a0014855

Smallwood, J., O'Connor, R. C., Sudbery, M. V., & Obonsawin, M. (2007). Mind-wandering and dysphoria. *Cognition and Emotion*, *21*(4), 816-842. doi: https://doi.org/10.1080/02699930600911531

Smallwood, J., & Schooler, J. W. (2006). The restless mind. *Psychological Bulletin*, *132*(6), 946. doi: https://doi.org/10.1037/0033-2909.132.6.946

Tripp, G., & Alsop, B. (1999). Sensitivity to reward frequency in boys with attention deficit hyperactivity disorder. *Journal of Clinical Child Psychology*.

Tutorial: Best practices for managing Workers in follow-up surveys or longitudinal studies. (2017). Retrieved January 14th, 2020, from `https://blog.mturk.com/tutorial-best-practices-for-managing-workers-in-follow-up-surveys-or-longitudinal-studies-4d0732a7319b`

Use MTurk. (2019). Retrieved January 15th, 2020, from `https://www.jatos.org/Connect-to-Mechanical-Turk.html`

Vrieze, E., Pizzagalli, D. A., Demyttenaere, K., Hompes, T., Sienaert, P., de Boer, P., ... Claes, S. (2013). Reduced reward learning predicts outcome in major depressive disorder. *Biological Psychiatry*, *73*(7), 639-645. doi: https://doi.org/10.1016/j.biopsych.2012.10.014

World Health Organization. (2017). *Depression and other common mental disorders: global health estimates* (Tech. Rep.). World Health Organization.

A Performing a (follow-up) study through Amazon MTurk: Lessons learned

A.1 Introduction

This guide will take a more in-depth look at Amazon Mechanical Turk and its use in the current study and follow-up studies in general. In order to make it more useful for people who have little experience with MTurk, the guide will start by covering the main MTurk components that were used during this study one by one. The guide will end with a summary of the lessons that were learned while conducting the current study on MTurk. Experienced MTurk Requesters can most likely skip to section A.8 immediately.

A.2 What is MTurk?

Amazon MTurk is an online platform that can be used for all sorts of data gathering. Lately, it is also being used more by researchers to perform scientific studies, and with good reason. Data quality does not seem to suffer when conducting studies through MTurk instead of in-lab, and using MTurk even seems to be beneficial to the diversity of a sample (Casler, Bickel, & Hackett, 2013). Mturk is also a powerful tool in that it is really helpful for gathering a lot of data quickly and from specific and possibly hard to reach audiences. The survey from this study exemplifies this: Within two hours of publishing it from a personal computer in Groningen, 100 people from the United States had filled out the survey.

A.3 How does it work?

MTurk knows two types of users: Workers and Requesters. The MTurk Workers are users who take part in HIT's (Human Intelligence Tasks) in exchange for money. Requesters are the users who create HIT's and request Workers to take part in their HIT's. Each Worker has earned their own share of Qualifications, and only when they have all the required Qualifications for a HIT, they are able to take part in the task. A Requester will specify all required Qualifications when creating the HIT. Examples of Qualifications are a Worker's age, job, location, or percentage of approved HIT's. Once a Worker has successfully made a HIT, they will let the Requester know by submitting an "assignment". For studies not completely made within MTurk (see A.7 for some options), a confirmation code is usually generated at the end of the task for the Worker to submit.

Once the Requester receives assignments for a HIT, they are to review the assignments, and approve or reject the assignment based on if the task was performed reasonably well. After a certain time period (three days by default, but this can be changed upon HIT creation), assignments will be approved automatically. Once an assignment has been approved, it cannot be rejected anymore. However, if an assignment has been incorrectly rejected, it can be approved up until 30 days after being rejected.

The HIT will stay open for new Workers to accept as long as the number of Workers who accepted the HIT is lower than the requested amount of assignments. The HIT will be completed once the number of submitted (not accepted!) assignments equals the number of requested assignments. If a Worker accepts a task but does not complete it within a set amount of time, the assignment will be republished for other Workers. When rejecting an assignment, the requester can choose to republish the assignment, so to make sure that enough useful data is gathered.

A.4 Quality assurance

The main challenge in working with MTurk is in assuring that the quality of the participants and data is up to par, and getting good quality workers to make your HIT's is the main part of that. MTurk has some good ways to weed out the poor Workers, although not all are free of charge.

A.4.1 Masters

MTurk Masters are Workers that have been granted the Masters qualification. The Masters qualification is a special qualification given out by Amazon, for Workers that have done at least 1000 HIT's and have maintained an approval rate $> 99.0\%$. Amazon Masters are therefore generally high-quality Workers and using this qualification is highly advised for scientific studies.

A.4.2 Qualifications

Besides using the Masters qualification, other qualifications can be used to ensure the quality of the Workers. While most qualifications are premium and require an additional fee, the "Number of HIT's approved" and "HIT approval rate" qualifications are free to use and are helpful in only selecting well performing Workers, especially when used in combination with a Masters qualification.

A.4.3 Data quality tests

Depending on the task, seeing if the data that you receive from Workers is good enough to approve can be difficult to do in a short period of time. It is therefore helpful to use smaller parts of the task to check if the Worker has been paying attention during the task. One straightforward example of this is using catch-questions in surveys, but also the time it took them to complete the task can be very useful in judging task engagement (do be aware that MTurk Workers make a lot of surveys, and are therefore generally a lot faster in answering them than participants with less survey experience).

A.4.4 Rejecting assignments

Once it has been determined that the quality of an assignment is too low, for example because all catch-questions were answered incorrectly, the Requester can choose to reject the assignment. While this leads to the Requester not having to pay the Worker, it is important to note that rejections also influence a Worker's account. Especially Masters and people with high approval rates (this study encountered a Worker with 133,895 approved and 35 rejected HIT's), will not take it lightly if their assignment is rejected. Rejecting assignments should be done with care and it is advised to approve all cases of doubt and to only reject assignments for which multiple rejection reasons can be given. Rejecting all cases of doubt can result in a lot of e-mails and severe customer support-like headaches, and is probably not worth the money.

A.5 Costs and payment

A.5.1 Fees

While MTurk Workers can be quite cheap participants for one's study, an important thing to take into account when calculating costs are the (hefty) fees that Amazon charges ("Pricing", n.d.). The base fee that Amazon charges is 20% of the amount that workers get paid, with a minimum of \$0.01. Additionally, Amazon charges an extra 20% over the participants' reward for HIT's with more than 9 assignments. Most qualifications (so called "premium qualifications") come with a fee as well, this time not as a percentage but simply a dollar amount, which is usually around \$0.50 per qualification. Lastly, a fee of 5% is charged for the usage of the Masters qualification. It can easily be seen how the costs of doing a study through MTurk could come out much higher than previously anticipated because of these fees. For scientific studies, where needing more than 9 participants is highly likely and a Masters qualification is very useful, fees are 45% almost by default. Luckily there are a few ways in which one can lessen the fee amounts drastically, which will be discussed next.

A.5.2 Circumventing fees

One significant way in which the fee amount can be lowered is by only creating batches of 9 or less assignments, and thereby circumventing the 20% fee for batches with more assignments. This may seem like a straightforward and easy solution, but there is more to it than meets the eye. When wanting to split up a study into multiple batches, the qualifications for doing the HIT will almost definitely stay the same between all batches. This means that the same Worker would be able to do the HIT from every batch once, resulting in multiple entries from the same participant in the data. Solving this issue in order to circumvent the fee can be done using custom qualifications (which are free), but takes a lot of effort and time, as each batch would have to be published only after the previous batch has finished. Custom qualifications and their use in follow-up studies are discussed more in depth in A.6.
The Masters qualification fee can be circumvented in follow-up studies, also using custom qualifications. As long as not too much time has passed, it is safe to assume that Workers who had the Masters

qualification previously, still have this qualification. Simply selecting only people who performed the first HIT will therefore be enough, as long as that HIT required Workers to have the Masters qualification.

A.5.3 Bonuses

Besides the standard reward that Requesters can give Workers, they are also able to give specific Workers one or more bonuses for every assignment that was approved. This can be done to reward Workers that performed exceptionally well, or to partly pay Workers based on their task performance (as in the current paper's reward learning experiment). Bonuses can also be used to communicate with workers, which is useful as MTurk does not offer this ability for privacy reasons. For each bonus that is sent, a message can be attached. Sending bonuses of $0.01 with a message to Workers is therefore a common and cheap way to contact them after they have performed one or more HIT's, and it does not seem to bother Workers.
Workers are, rightfully, wary of HIT's that have a small base reward but that promise high bonuses. When using bonuses as part of the payment, make sure to advertise clearly what the bonus depends on. Try to mention an average or minimum bonus amount in the title of the HIT. A disadvantage of using bonuses as part of the payment is that there exists no overview of paid bonuses, and keeping track of a comprehensive spreadsheet of the Requester's dealings with Workers is advised.

A.6 Follow-up studies

It is possible to do follow-up/longitudinal studies on MTurk, but managing all your Workers can be quite a hassle. The current study consisted of one survey followed by two different cognitive tasks that were to be performed if the participant had made the survey reasonably well (answered catch-questions correctly, etc.). A lot of lessons were learned in trying to achieve this on MTurk, and this section will focus on how the follow-up process should be handled for a study that looks like the current study.

A.6.1 Survey

Firstly, the survey should be made by a lot of participants from our target population, > 100 in this case. One could choose to use batches of 9 sequentially in order to skip the extra 20% fee, but once more than 50 participants are needed this is probably not worth the trouble and it would be advised to pay the extra 20% to be able to gather everything at once.
However, when trying to avoid the extra fee on a survey like this, the following can be done: Start a batch with at most 9 assignments, and wait for it to complete. Meanwhile, create a custom qualification "Made survey already" and once the first batch is completed, assign the qualification to all Workers that submitted to that batch before publishing the next batch. Keep doing this for all batches until enough data has been gathered ("Tutorial: Best practices", 2017).

A.6.2 Follow-up tasks

Once the Workers that are to make the follow-up tasks have been selected, qualifications should be made for the follow up tasks. Now that the Requester already possesses the Worker ID's of the targeted set of Workers, circumventing the extra 20% fee is actually easier than before. Instead of having to wait for one batch to complete before publishing the next one, the Requester can now simply make a multiple of batches of 9 assignments with a specific (free) custom qualification for each batch. The set of Workers that has been chosen can now be divided into these qualifications, and all batches can be published at once.
The current study faced the main challenge of having the participants make both of the follow-up assignments, as only the data from Workers that made both assignments could be used. In order to maximize the percentage of Workers that make both tasks instead of only one (usually the most profitable), the following approach seems to work best: Instead of inviting the selected Workers for both tasks using one batch qualification, create a specific custom qualification for each batch/-task combination (so Task1Batch1, Task2Batch1, instead of Batch1 for both tasks). Give the Workers a qualification and invite them for the least attractive (usually the longest) task. Only once they

have made that task successfully, give them a qualification and invite them for the second task. If you assign the qualifications of both tasks simultaneously, Workers will be able to find the published batch of a second task without you inviting them for it, and thereby potentially making the second task before the first.
Inviting workers for follow-up tasks can most easily be done through a message attached to a bonus of $0.01 (see A.5.3 for more information about bonuses). Success has also been found in using bonus messages to send reminders about unmade tasks.
With regards to the payment for two follow-up tasks, the following is suggested: From the start of the study (in this case the survey), clearly communicate what the payment looks like over the entire study. Make approximately 25% of the payments for the tasks a bonus that is only paid once a Worker submits both HIT's, in order to stimulate the completion of both tasks. MTurk Workers are more inclined to do shorter HIT's, so consider paying disproportionally much for a task if it is much longer than the other (which can be subtracted from the shorter task's reward).
Lastly, remember that MTurk Workers are (usually) not using MTurk full-time, and being quick in responding to e-mails, approving assignments and inviting Workers for new tasks has a significant effect on the amount of fully completed experiments.

A.7 Software

While MTurk is a good platform for connecting Workers and Requesters, it can fall short in other aspects. Luckily a lot of software can be integrated with MTurk, giving a lot more options for what tasks can look like. Two examples of useful software, which were both also used in the current study, are given below.

A.7.1 Qualtrics

Qualtrics (https://www.qualtrics.com) is very easy to use survey software that can be integrated seamlessly with MTurk. Instructions on how to combine Qualtrics and MTurk can be found on "Getting great survey results" (2017).

A.7.2 JATOS

Besides surveys, it is also possible to run full cognitive tasks on MTurk. One way of doing this is by using JATOS (Lange et al., 2015). JATOS can be used to run cognitive tasks that were made using a tool like OpenSesame/OSWeb, jsPsych, or PsyToolkit, and is also easily integrated with MTurk ("Use MTurk", 2019). It should be noted, however, that this study had some issues with JATOS' generation of the confirmation codes that Workers need to enter in order to complete the HIT. This can be solved easily though by instructing the Workers to submit something specific that is not the confirmation code (a part of the task URL in this case), if something goes wrong.

A.8 Lessons learned

In summary, important lessons learned about the use of MTurk from this study are the following:

1. MTurk workers, especially Masters, care a lot about being rejected. Think very carefully about rejecting Workers if their work is subpar. Simply approving the HIT and then disregarding the data might be the cheaper option both time wise and financially.

2. To assure data quality in surveys, use a combination of catch-questions and time data. But keep in mind that MTurk Workers are faster than average in making surveys, and that their Time To Complete cannot fully be compared to non-MTurk participants.

3. In order to circumvent the additional 20% fee for batches of more than 9 assignments, one can try to create multiple batches with the same task. However, this process takes longer as only one of the batches can be published at the same time, because Workers that made the HIT from a previous batch should be excluded from all further batches.

4. For follow-up tasks, the 20% fee for larger batches can be omitted more easily as you already possess the Worker ID's. In this case, simply create multiple batches, each one for nine specific Workers. As each Worker only qualifies for one of the batches, all batches can be published simultaneously.

5. Also for follow-up tasks, the extra Masters qualification cost can be omitted as one can assume that all Workers still have this qualification if it was also needed to do the initial HIT.
6. Bonuses can be awarded to Workers, and are useful in encouraging certain behaviours. For example, awarding a bonus for making multiple HIT's could be used to encourage Workers to make two different tasks if data from both tasks is needed before the data is useful. When doing so, it is also wise to let Workers make the longest/least paying task first.
7. When using bonuses, make sure to communicate very clearly how they work and advertise them in the title of the HIT. Workers will generally avoid HIT's with a low base payout and you will therefore have to convince them that the awarding of a bonus is either certain or very likely.
8. Messages can be attached to a bonus. Use a bonus of $0.01 if you need to contact workers, for example to invite them for a follow-up task.
9. Be responsive, especially in follow-up studies. Responding quickly to questions and problems has a significant effect on how many Workers complete the full study.

B Reward learning instructions

Screen 1

In this experiment, you will have to determine whether an emoticon has a long or a short mouth.

The experiment will consist of three blocks of 100 trials. After each block, you will have a 30 second break.

Each trial will go as follows:

- A fixation point appears
- A mouthless emoticon appears at the point
- A mouth appears on the emoticon for 1/10 of a second
- If you think the mouth you saw was
  - short, you press the 'z' key.
  - long, you press the 'm' key.

If you answered correctly, you may be rewarded with $0.02. But this is not necessarily the case for every correct answer!

You will first get a demonstration of a trial. Then there will be 2 practice trials and then the experiment will start.

Click on 'I understand' once you're ready to start.

C Reward learning task demonstration

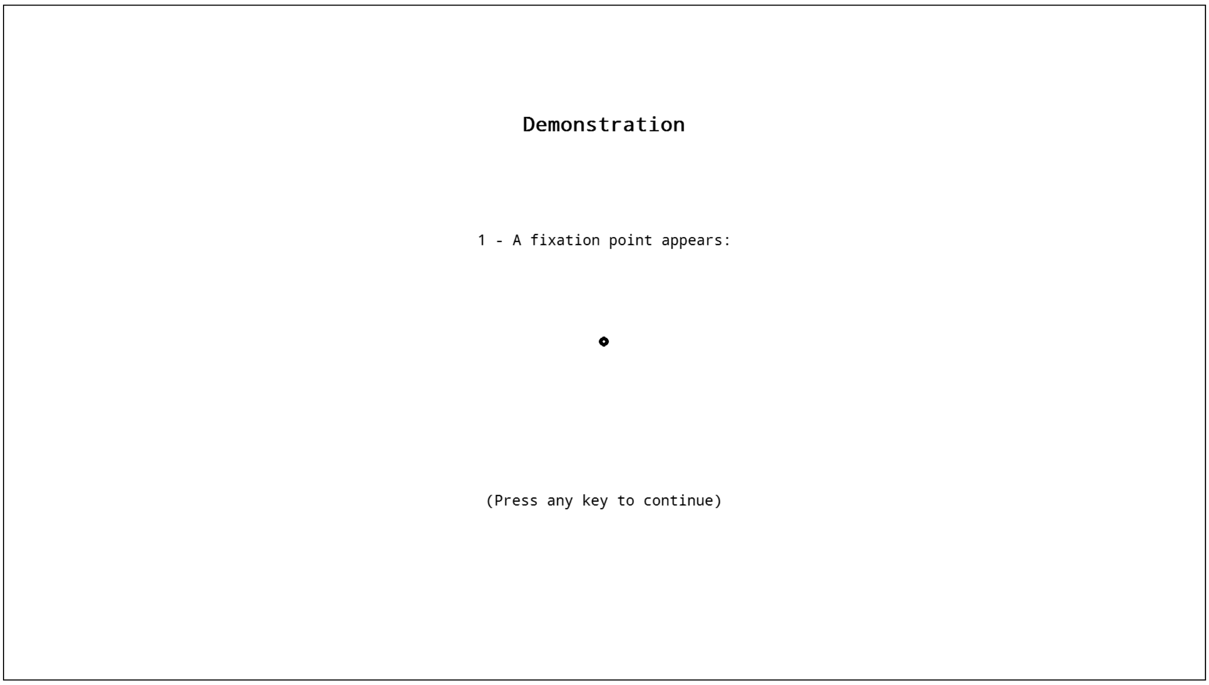

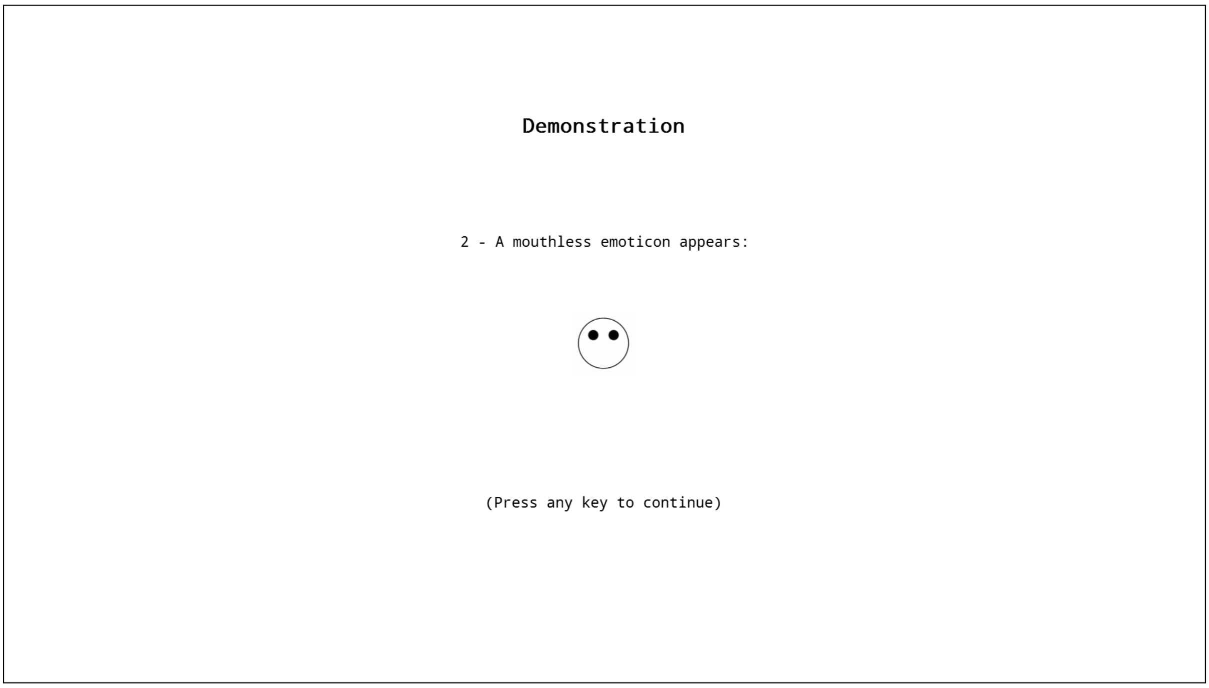

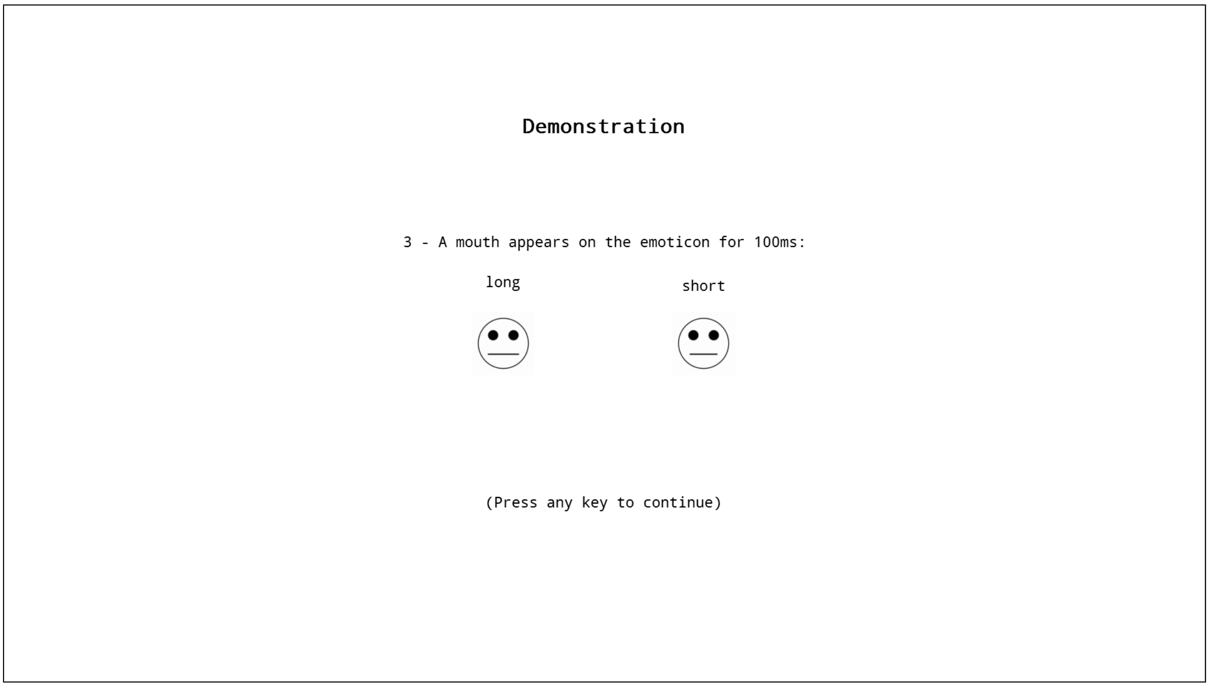

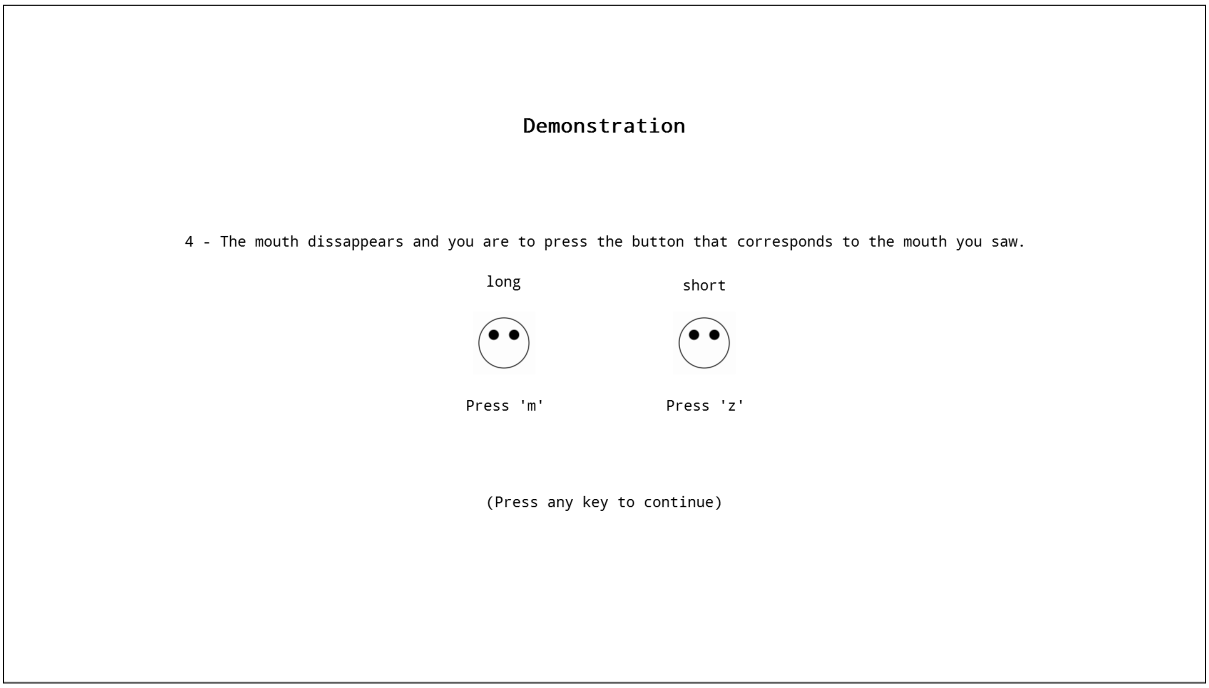

Figure C.1: Example of the sequence of still images that was used in the demonstration of the reward learning task. Key press instructions differed for each participant due to counterbalancing.

D Mind-wandering task instructions

Screen 1

A few days ago, you filled out our survey. In this survey, you were asked to name 5 achievements/happy moments and 5 concerns in your life.

Please take a minute to try to remember what you answered to these questions, as later in this experiment you will be asked to repeat them as quickly as possible.

Screen 2

In this experiment you will view words written in upper- or lower case. Press the space bar when the word is written in lower case (e.g., dog). Do NOT press any button when the word is written in upper case (e.g., TREE).

It is important that you respond as accurately as possible.

Press any key to continue.

Screen 3

Every now and then, the task will be interrupted with the question of what you were thinking about just now.
You can answer this question by pressing the number key corresponding to your answer.

Time is not an issue here, so please take your time to think about the questions and answer them.

Press any key to see the example questions.

Screen 4

What were you thinking about just now?

1) The task itself
2) An aspect of the task (e.g., how I was doing or how long the task was taking)
3) Personal things
4) I was distracted by my environment/myself
5) I was daydreaming / thinking about something else
6) I was not paying attention, and did not think about anything in particular

Press the key that corresponds to your answer to continue

Screen 5

If you were not thinking about the task itself,what was the content of your thought?

1) Positive, self-related
2) Negative, self-related
3) Positive, other-related
4) Negative, other-related
5) I was thinking about the task itself

Press the key that corresponds to your answer to continue

Screen 6

How difficult was it to disengage from the thought?

1) Very easy
2) Easy
3) Neither difficult nor easy
4) Difficult
5) Very difficult

Press the key that corresponds to your answer to continue

Screen 7

We will now start with a few practice trials. Remember: Press the space bar for lower-case words. Press any key to continue.

E Bayes Factors

On-task				
Answer	**RRS**	**PTQ**	**BDI-2**	**COMB**
1.	.41	.46	.39	.40
2.	.78	.45	.41	.54
3.	.71	.43	.42	.44
4.	.39	.49	.80	.40
5.	4.82^{a}	.43	.40	.77
6.	.40	.40	.53	.43
1&2.	1.88	.40	.41	.53
Valence				
Answer	**RRS**	**PTQ**	**BDI-2**	**COMB**
1.	.53	.44	.40	.44
2.	1.16	3.34^{a}	.40	1.14
3.	.99	.43	1.10	.62
4.	.88	.66	1.54	1.37
5.	.53	.40	.44	.40
1&2.	1.68	3.81^{a}	.39	1.34
3&4.	.43	.73	.50	.58
1&3.	1.18	.40	.72	.65
2&4.	.39	.43	.73	.41
Sticky				
Answer	**RRS**	**PTQ**	**BDI-2**	**COMB**
1.	1.44	1.11	.46	1.15
2.	.80	.45	.42	.56
3.	.41	.41	.63	.42
4.	.79	.96	.58	.50
5.	49	.42	.46	.49
1&2.	.41	.49	.40	.43
4&5.	.44	.52	.59	.40

Table E.1: Bayes factors of the correlations between the proportion of choices for one answer or a combination of answers to the on-task/valence/sticky question, and each depression score. Answers are indicated by their choice number and the content of the answers can be found in appendix D. ^{a}Enough evidence that the correlation exists (BF > 3.0).

RRS	PTQ	BDI-2	COMB
1.97	2.39	N/A	$.18^{a}$

Table E.2: Bayes factors comparing the linear models of reward learning and mind-wandering for each depression score. ^{a}Enough evidence that reward learning is the better predictor (BF < .30).

Score	PTQ	BDI-2	COMB
1. RRS	$>1000.00^{a}$	4.66^{a}	$>1000.00^{a}$
2. PTQ		3.40^{a}	$>1000.00^{a}$
3. BDI-2			$>1000.00^{a}$

Table E.3: Bayes factors of the pairwise depression score correlations. ^{a}Enough evidence that the correlation exists (BF > 3.0).

Score	Δ RB$_{1-2}$	Δ RB$_{2-3}$	Δ RB$_{1-3}$
RRS	2.80	2.42	0.39
PTQ	1.55	0.72	0.44
BDI-2	2.25	0.63	0.53
COMB	4.87^{a}	1.47	0.44

Table E.4: Bayes factors of the correlation between change in response bias per interval and depression scores. ^{a}Enough evidence that the correlation exists (BF > 3.0).

RRS	PTQ	BDI-2	COMB
6.26^{a}	1.79	8.80^{a}	17.59^{a}

Table E.5: Bayes factors of the t-tests comparing score distributions by when participants learn their response bias, divided by depression score. ^{a}Enough evidence that the correlation exists (BF > 3.0).

RT/ACC	RRS	PTQ	BDI-2	COMB
RT	.40	.40	.42	.40
ACC	.45	.43	.47	.40

Table E.6: Bayes factors for the relation between response time (RT) and accuracy (ACC), and survey scores.

RT/ACC	On-task	Valence	Sticky
RT	$.04^{a}$	.36	$<.01^{a}$
ACC	.58	$.06^{a}$	$.02^{a}$

Table E.7: Bayes factors for the ANOVA that assessed the relation between response time (RT) and accuracy (ACC) preceding a question block, and the answers to that question block. ^{a}Enough evidence that the correlation does not exist (BF < .30).