SOCIAL LEARNING FOLLOWING SOCIAL INTERACTIONS

by

Megan Bayne
A Thesis
Submitted to the
Graduate Faculty
of
George Mason University
in Partial Fulfillment of
The Requirements for the Degree
of
Master of Arts
Psychology

Committee:

____ Director

____

____

____ Department Chairperson

____ Dean, College of Humanities and Social Sciences

Date: ____________________________________ Fall Semester 2017
George Mason University
Fairfax, VA

Social Learning Following Social Interactions

A Thesis submitted in partial fulfillment of the requirements for the degree of Master of Arts at George Mason University

by

Megan Bayne
Bachelor of Science
George Mason University, 2016

Director: James Thompson, Principal Investigator
George Mason University

Fall Semester 2017
George Mason University
Fairfax, VA

DEDICATION

This is dedicated to my parents, who support, love, and believe in me unconditionally. Thank you, Mom and Dad.

ACKNOWLEDGEMENTS

I would like to greatly thank the many friends, relatives, and supporters who have made this happen. My supportive parents gave me this wonderful opportunity and supported me the whole way, I could not be more grateful. My loving boyfriend, Pat, offered unconditional support and encouragement throughout this long process, which I could not have done without. My sisters generously assisted me in my research with enthusiasm. My good friend, Thanh, offered valuable research advice and creative perspectives for my project. Drs. Thompson, Fischer, and McDonald of my committee provided plenty of thoughtful help and guidance. Finally, thanks to Mike Hock and other graduate faculty members for helping me with every detail of this process.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF EQUATIONS

LIST OF ABBREVIATIONS

Reward Invigoration .... RewInv
Punishment Suppression .... PunSup
Pavlovian Performance Bias .... PPB
Go-to-Win .... GTW
Go-to-Avoid .... GTA
NoGo-to-Avoid .... NGTA
NoGo-to-Win .... NGTW

ABSTRACT

SOCIAL LEARNING FOLLOWING SOCIAL INTERACTIONS

Megan Bayne, M.A.

George Mason University, 2017

Thesis Director: Dr. James Thompson

This research explored the effect of social interactions on social learning. Previous research has indicated that our tendency to approach rewarding social cues or avoid negative social cues can interfere with our ability to learn from social feedback (Cavanaugh, 2013). In this study we examined if these tendencies are dependent on our social motivational state. The study consisted of two parts: part one of this study used an online ball-tossing game, Cyberball, as a social exclusion task, part two of this study used a social rejection picture task as a negative social interaction. The control for both parts involved a neutral social interaction. Both tasks were followed immediately by a reinforcement learning task which utilizes Pavlovian biases. Part one and two of the study were analyzed separately using a mixed ANOVA. The results of this test were hypothesized to show a significant increase in accuracy on the withdrawal trials of the Pavlovian influence task following a negative social interaction. There was no effect found from the social motivational state on Pavlovian reinforcement, however, results did point to an approach bias within individuals. The findings from this study contribute to

our understanding of motivational influences on learning from the feedback of other people.

INTRODUCTION

Social feedback from other individuals is one way in which we learn which behaviors to continue and which behaviors to stop. Humans are equipped with a Pavlovian system that has the ability to predict outcomes based on cues that signal the coming of an appetitive or aversive consequence (Thompson & Westwater, 2013). This predictive system allows for the increase in acceptable behaviors. While social feedback has been shown to have the potential to produce a variety of social emotions, little is known about the lasting effect of social interactions, particularly in relation to subsequent social learning. The goal of this study was to determine if negative social interactions affect reinforcement learning tendencies under the Pavlovian influence.

Within psychology, social interactions are often studied through the use of experimenter-controlled situations. Experimenter-controlled situations allow for the simulation of real-life experiences, and the controlled study of social situations. A program called Cyberball was developed by Williams, Cheung, and Choi in 2000 for the use of studying ostracism over the internet. Cyberball provides an online ball-tossing game between a real-life participant and other fake “participants.” The real participant plays the game online with what they believe to be other participants playing from a remote location. The experimenter controls how many times the participant receives the ball from other “participants.” Social exclusion can be induced by having the participant receive the ball only once or twice, and then not again throughout the game, as they watch the other participants play with each other (Williams & Jarvis, 2006). Cyberball has been used in many situations to study ostracism, discrimination, conformity, and prejudice.

A study published in 2013 by Kawamoto, Nittono, and Ura confirmed the efficacy of Cyberball using facial electromyogram (EMG) and electroencephalogram (EEG) to view the biological and neurological effects during the task. Facial EMG in this study was used to investigate activity of the corrugator supercilii, which is the muscle just above the eyebrow. Activity of these muscles is said to be negatively related to the valence of a subjective experience, therefore showing more activity during a negative experience (Kawamoto et al., 2013). EEG was used to view electrical activity of the brain and measure ERPs (event related potentials) occurring during exclusion through Cyberball. The results showed that facial EMG increased throughout the task for all three conditions- fair play, observation, and exclusion, but that the recordings increased the most dramatically in the exclusion condition (Kawamoto et al., 2013). EEG results showed that frontal activation was stronger on the left (associated with approach motivation) in the first half of the task, but was stronger on the right (associated with withdrawal motivation) during the second half of the task. This is evidence of the change in motivation that occurs during exclusion, in order for a human to protect their fundamental need to belong. These results demonstrate how exclusion from Cyberball is similar to that experienced with real humans, and that the neural correlates indicate a withdrawal motivation occurring during exclusion (Kawamoto et al., 2013). A study by

Gray (2009) also found that negative affect can induce withdrawal states, but also that this state affects different tasks in a different manner. This points to the idea that withdrawal states can separately influence distinct components of cognitive control.

`In addition to the experience of exclusion during the task, Cyberball has been shown to cause lasting changes in behavior and decision making. A study by Williams, Cheung, and Choi (2000) reported that participants who were ostracized through Cyberball were more likely to conform on a subsequent task. This research provides the foundation for the use of Cyberball to examine the effects of social exclusion on social reinforcement learning utilizing Pavlovian conflict.

In comparison to exclusion, rejection is also a common negative social interaction that people experience. Although similar in their appearance, exclusion and rejection differ in terms of their explicitness, with exclusion being more implicit and rejection being more explicit. Kipling Williams (2007) differentiates between the two in the *Annual Review of Psychology,* and states that, "Being ignored and excluded … often occurs without excessive explanation or explicit negative attention," whereas rejection is described as, "A declaration by an individual or group that they do not (or no longer) want to interact or be in the company of the individual."

Previous studies have evaluated the effects of rejection primarily using picture tasks, in which an individual is specifically told they were disliked by the person in the picture they are viewing. An fMRI study by Somerville, Heatherton, and Kelley (2006) developed a rejection task in which participants have their picture taken ~1 week prior to coming in for the experiment, under the guise that their picture will be sent out to other universities to be evaluated for a study on first impressions. Upon coming in for the experiment, participants were given fictitious feedback as to whether they were evaluated by other people in terms of “like” or “do not like.” In addition to being either positive or negative, feedback was also labeled as either consistent or inconsistent with the participant’s expectations. The study showed that while the dorsal ACC responds to expectancy violations, the ventral ACC is responsive to social feedback (Somerville et al., 2006).

A study by Gunther, Crone, and van der Molen (2010) used a similar paradigm to evaluate the cardiac response to social rejection. The study found that rejection feedback caused a slowing of the heart rate that was significantly greater when acceptance was expected (Gunther et al., 2010). These results show how the parasympathetic nervous system responds to unexpected rejection. This research provides evidence that rejection is a specific negative social situation that can be replicated, and leads to neurological changes that may alter thinking and learning.

Social reinforcement learning involves an increase or decrease in the probability of an action being repeated due to perceived social feedback. This can be an important aspect in studying how social stimuli affect an individual’s decision making. A significant feature of reinforcement learning, known as the Pavlovian bias, explains how people tend to couple reward seeking with action invigoration, while linking punishment avoidance with action suppression (Cavanagh, Eisenberg, Guitart-Masip, Huys, & Frank, 2013). This results in a tendency to strengthen an action in response to a reward, and suppress action in response to punishment. This bias can hinder learning in conflicting situations when action invigoration is required in response to a punishment, or when action suppression is required in response to a reward. However, it is suggested that humans have systems that can help to overcome Pavlovian conflict (Cavanagh et al., 2013).

The situations and systems involved in the Pavlovian bias can be studied using tasks that pair conflicting, as well as non-conflicting, action and valence requirements. A study by Cavanagh et al. (2013) adapted a learning task from Guitart-Masip et al. (2012) which utilizes the Pavlovian bias. The task involved various cues which lead to either a reward or punishment outcome based on the participants' decision to click on a target or not, following the cue. Clicking the target was considered a "Go" response, while not clicking was a "NoGo" response. The task used four different shapes as separate cues, which were 70% predictive of an outcome based on the participant's decision to "Go" or "NoGo." The outcome for each trial of this task was either a reward "+$", a punishment "-$", or a neutral "-". This provides four combinations of action and outcome: Go-to-Win, where a participant provides a "Go" response in order to gain the reward, NoGo-To-Win, in which the participant provides a NoGo response in order to gain the reward, Go-To-Avoid, where a participant provides a "Go" response in order to avoid the punishment, and NoGo-to-Avoid, in which a participant provides a NoGo response in order to avoid the punishment. In the win trials, the neutral outcome switches to the reward outcome once the participant begins to provide the appropriate response. In the avoid trials, the punishment outcome switches to the neutral outcome once the participant begins to provide the appropriate response. The study found that on average, participants performed better on Pavlovian congruent trials (Go-to-Win, NoGo-to-Avoid) than the conflict trials (Go-to-Avoid, NoGo-to-Win), with NoGo-to-Win resulting in the poorest performance (Cavanagh et al., 2013). The study also revealed that participants with greater frontal theta power were better at overriding the Pavlovian bias and displayed improved performance on conflicting Pavlovian trials (Cavanagh et al., 2013).

The culmination of this research provides evidence that exclusion and rejection can lead to neurological and behavioral changes, and that these may alter the reinforcement of subsequent social stimuli. This previous literature led to the hypothesis that a negative social interaction involving exclusion or rejection would cause an increased accuracy in NoGo (withdrawal) trials during a subsequent social Pavlovian reinforcement task.

METHOD

Participants

This study included 75 participants- 39 in Part 1 and 36 in Part 2. The participants were drawn from George Mason University's population using SONA systems, in which undergraduate Psychology students can sign up to participate in the study for class credit, as well as flyer and email recruitment in which participants were paid. Participants were eligible for the experiment if they are over 18 years of age. Participants were involved in deception, as they were told that they are interacting with real people through the computer in both Parts 1 and 2, in order to fabricate social interactions. In reality, all social interactions were controlled by the experimenter for specific results. Participants were fully debriefed immediately following completion of participation.

Materials

The materials for this study included a brief demographics survey, social interaction tasks, social Pavlovian influence task, and debrief script. The demographics survey asked for gender, age, and race. Participants were told they do not have to respond to any questions they do not feel comfortable answering. The social interaction task in Part 1 was an online ball-tossing game, Cyberball. This game allows the experimenter to control how many players the participant interacts with and how many times they are tossed the ball (Williams & Jarvis, 2006). In this experiment, participants played online with what they believed to be two other participants. In the experimental condition, participants were thrown the ball only twice to induce social exclusion (Williams & Jarvis, 2006). In the control condition, participants received the ball roughly the same amount as the other players.

The social interaction task in Part 2 was a rejection task involving pictures of what appear to be other participants. Approximately one week prior to coming in for the experiment, participants sent in a photo of themselves that they are deceptively told will be used in the experiment. During the task, participants viewed photos of the other people one at a time and selected whether they liked or disliked the person in the photo. Participants were then shown whether that individual liked or disliked them from their photo. There were multiple areas of deception in this task: the people in the photos during the task were not actually participants, but random pictures of people taken by the experimenter; the pictures participants sent in are actually immediately deleted and never seen by anyone; and positive or negative feedback was controlled by the experimenter. Participants were not actually being evaluated, this is all controlled by the experimenter to induce rejection in the experimental condition, which involved negative feedback (dislikes from the fake participants) 2/3 of the time. In the control condition, participants received positive feedback 2/3 of the time. Participants were fully debriefed and had all deception explained following completion of the experiment.

The social Pavlovian influence task is a computer based program adapted from Cavanaugh et al. (2013). This task is an implicit learning task associated with motivational (Pavlovian) bias. In the task, there are specific shapes used as cues that are associated with a positive (happy face) or negative (angry face) social outcome. The outcome is based on whether the participant chooses to click the mouse or not click the mouse following the specific cues. This is referred to as a "Go" or "NoGo" response, respectively. The outcomes in this situation are positive and negative facial expressions shown on the screen, or a neutral outcome that is a yellow equal's sign. Figure 1, adapted from Thompson and Westwater (2017), demonstrates the design of this task. The goal of the task is to learn how to react in order to elicit as many positive outcomes as possible. The task produces four different types of events: Go-to-win (GTW), Go-to-avoid (GTA), NoGo-to-avoid (NGTA), and NoGo-to-win (NGTW). The Go events involve the clicking of the mouse in order to receive a reward or avoid a punishment, while the NoGo events involve the not clicking of the mouse to receive a reward or avoid a punishment. The events are not isolated within the task and it is goal of the participant to figure out when to perform either action.

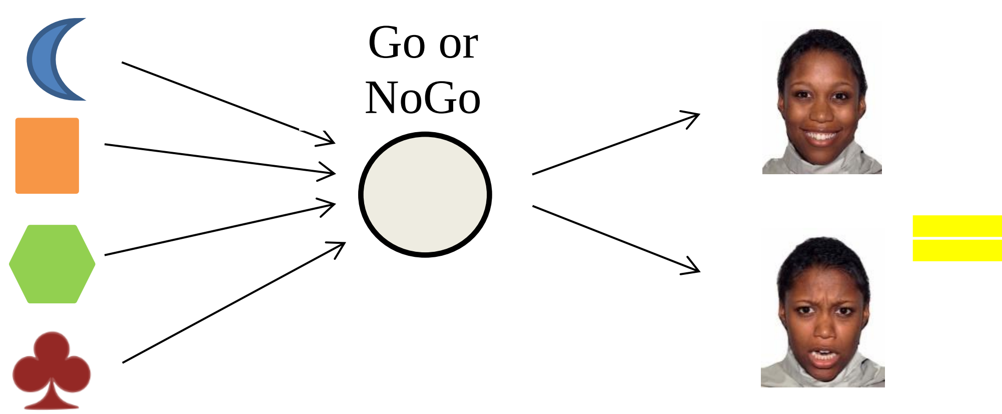

Figure 1. Task Design. Cues lead to reward or punishment outcome, sometimes neutral outcome.

Procedure

This two-part study implemented a between-subjects experimental design in both parts. It was conducted to determine if a prior negative social interaction would alter the learning patterns on the subsequent Pavlovian influence learning task. The study used two conditions, a control (or neutral) condition and a negative social interaction condition. Participants were recruited from George Mason University and randomly assigned to either condition.

Part 1 of the experiment involved the participant entering the experiment analysis room and signing an informed consent in the presence of a researcher. The participant was then read instructions to complete the Cyberball task. The participants were told they were completing this task as part of a mental visualization experiment. This deception was to protect the goal of eliciting feelings of exclusion within the experimental group. The game then began and appeared as though two other players were playing the game with the participant, although these were really fake characters with predetermined actions controlled by the experimenter. The control condition received the ball an equal amount of times as the other players within the task, while the experimental condition received the ball only twice, thus creating the illusion of being excluded. The task lasted approximately 5-10 minutes. Once finished the participant was read instructions for the Pavlovian influence reinforcement task. A paper demonstration was given for the participant to have a better idea of what they would be seeing on the screen. The participant then completed the task, which lasted approximately 30 minutes, and was debriefed immediately after. The debrief provided the participant with complete knowledge on the nature of the experiment, including the information that the other characters in Cyberball were created by the researcher and were controlled to either provide a neutral situation (in the control condition) or an exclusionary social situation (in the experimental condition). The debrief also included contact information for George Mason University's CAPS program located on campus should the participant feel any distress following the experiment. The participant was compensated with cash or SONA credits immediately following completion of the experiment.

Part 2 of the experiment followed a similar protocol with a different social task. The participant first sent a selfie to the experimenter approximately one week prior to coming in to participate. The participant was told that this study looks at how social evaluations are made, and that their picture would be seen by other participants for evaluation. The selfie was only part of the deception and was immediately deleted and not shown to anyone. On the day of the experiment, the participant entered the analysis room and signed the informed consent in the presence of a researcher. The participant was then read instructions to complete the picture task. The participant was told they were completing the task as part of a social evaluation experiment, and that the pictures of the people they would be seeing were previous participants who saw their picture in the same way as the participant was seeing these pictures now. In reality, the pictures in the task were of random people, and the participant's picture was never shown to anyone. This deception was to protect the goal of eliciting feelings of rejection within the experimental group. The participant was instructed to select like or dislike for each picture they see, and then they would be shown what that "participant" selected for them. This appeared to the participant as, "This participant said they liked you," or "This participant said they disliked you." The control condition received 66% "like" feedback and 33% "dislike" feedback, to provide a realistic neutral social situation. The experimental condition received 66% "dislike" feedback and 33% "like" feedback, to create the illusion of being rejected. The task duration was approximately 5-10 minutes. Once finished the participant was read instructions for the Pavlovian influence reinforcement task. A paper demonstration was given for the participant to have a better idea of what they would be seeing on the screen. The participant then completed the task, which lasted approximately 30 minutes, and was debriefed immediately after. The debrief provided the participant with complete knowledge on the nature of the experiment. This included the information that the pictures the participant saw in the picture task were random people who never saw their photo and never evaluated them, and that the feedback was controlled by the researcher to either provide a neutral situation (in the control condition) or a rejection social situation (in the experimental condition). Participants were also informed that the selfie was immediately deleted and not used in the experiment. The debrief also included contact information for George Mason University's CAPS program located on campus should the participant feel any distress following the experiment. Lastly participants were asked to rate their experience of the rejection task on a scale of 0-5, 0 being "No feelings of rejection", and 5 being "Worst feeling of rejection possible." The participant was then compensated with cash or SONA credits immediately following completion of the experiment.

RESULTS

The hypothesis of this study was that negative social interactions would improve the accuracy of withdrawal trials in the Pavlovian influence task. This means that participants who experienced a negative social interaction prior to the task would be better at learning to suppress an action (provide a “NoGo” response) in order to gain a reward (happy face) or avoid a punishment (angry face). This was based on the finding that right frontal brain activation increases during continued exposure to exclusion and is associated with withdrawal motivation (Kawamoto et al., 2013). Individuals may therefore be primed to better withdrawal themselves in order to protect from further ostracism or rejection, and thus would be better at learning to provide a NoGo response.

The Pavlovian reinforcement learning task results in seven score categories giving a percentage. The first of these categories is reward invigoration, which refers to how often a participant was driven to perform an action in response to a reward (happy face) (Cavanaugh et al., 2013). This is measured by the equation:

Equation 1 Reward Invigoration

$$((\text{Go on Go} - \text{to} - \text{Win} \; + \; \text{NoGo} - \text{to} - \text{Win}) / \text{ Total Go})$$

Notice from the equation that the Reward Invigoration score does not refer to accuracy, because providing a “Go” response on a NoGo-to-win trial would not be accurate. It simply refers to the degree to which a participant is driven to “Go” in an attempt to get a reward. The second category, Punishment Suppression, refers to how often a participant was driven to suppress an action to avoid punishment (angry face). This is measured by the equation:

Equation 2 Punishment Suppression

$$((\text{NoGo on Go} - \text{to} - \text{Avoid} \ + \ \text{NoGo} - \text{to} - \text{Avoid}) / \ \text{Total NoGo})$$

This similarly is not a measure of accuracy, but rather the degree to which a participant is driven to "NoGo" in an attempt to avoid punishment. The average of reward invigoration and punishment suppression produces the third category, Pavlovian Performance Bias (PPB). This refers to how heavily a participant relies on the Pavlovian bias when making decisions (Cavanaugh et al., 2013). A participant with perfect learning throughout the task would result in a score of 50% on Pavlovian Performance Bias, since half of the trials are action-valence congruent (GTW and NGTA), and the other half are action-valence incongruent, therefore contradicting Pavlovian responses (GTA and NGTW) (Cavanaugh et al., 2013). Higher scores indicate a higher dependency on Pavlovian biases. The remaining four categories are percentages reflecting accuracy on each of GTW, GTA, NGTA, and NGTW trials. These categories are described in the mixed ANOVA in terms of valence, response, and condition. Valence and response are both within-subjects variables, while condition is a between-subjects variable. Valence has two levels and refers to whether the participant expected a reward or a punishment. Response also has two levels and refers to whether the participant decided to "Go" or "NoGo" throughout reinforcement. The combination of these two variables gives all four trials within the task:

Valence(1) Response(1) = GTW Valence (2) Response(1) = GTA

Valence(1) Response(2) = NGTW Valence(2) Response (2) = NGTA

Lastly, the between-subjects variable condition has two levels and produces the two groups, accepted or rejected. The interactions between these three variables provide further analysis of the data. This method of analysis allows for the examination of both within- and between-subjects factors.

Part 1

A mixed ANOVA was conducted separately for each part to compare the Pavlovian task scores in exclusion/rejection and control conditions. Part one of the study included 19 participants in the control and 20 participants in the experimental condition. The participants ranged in age from 18 to 43 and included 14 males and 25 females. Analysis revealed no significant results in terms of the experimental manipulation, but did result in significance pertaining to the efficacy of the Pavlovian influence task. There was no significant difference in accuracy due to valence within subjects. This means that there was no difference in performance accuracy between the combined win trials (GTW and NGTW) and combined avoidance trials (GTA and NGTA) within subjects; $F(1,34)=2.374$, $p = .132$. Looking at Table 1 it can be seen that the combined means for win trials is very similar to the combined means for the avoidance trials. There was also no significant difference in accuracy due to valence between the conditions. This means that both conditions received statistically similar accuracy on the combined win and combined avoidance trials; $F(1,34)= .351$, $p = .557$. There was a significant difference in accuracy due to response within subjects. This means that there was a difference in accuracy when participants were required to "Go" and when they were required to "NoGo"; $F(1,34)= 22.393$, $p = .000$. Table 1 shows that the combined means for "Go" trials (GTW and GTA) is greater than the combined means for "NoGo" trials (NGTA and NGTW). This finding is reflective of a "Go" bias in terms of response within subjects. Participants performed better when they were required to "Go" than when they were required to "NoGo". This is a central finding in other papers, including Cavanaugh et al. (2013). The difference in accuracy due to response was the same between the two conditions, resulting in no significance in the response*condition interaction; $F(1,34)= .098$, $p = .756$. The valence*response interaction was significant; $F(1,34)= .19.676$, $p = .000$. This interaction reflects the presence of the Pavlovian bias and is similar to the Pavlovian Performance Bias measure. More specifically, it means that participants are driven to "Go" when expecting a reward, and are driven to "NoGo" when expecting a punishment; these are similar to the Reward Invigoration measure and the Punishment Suppression measure, respectively. The level of Pavlovian bias was not different between the two conditions, exclusion and control, resulting in no significance in the valence* response *condition interaction; $F(1,34)= .002$, $p = .996$. This indicates that the level of reliance on the Pavlovian bias was very similar between the two groups. These results reject the hypothesis and suggest that the mental state following online exclusion does not have an effect on subsequent Pavlovian influence learning. However, the results do show the strong efficacy of the Pavlovian influence task, as it is an appropriate measure of Pavlovian influence.

Table 1
Descriptive Statistics of Part 1

Pavlovian Reinforcement Learning	Group					
	Control			Exclusion		
	M	SD	n	M	SD	n
RewInv	.5964	.1636	19	.5924	.1351	20
PunSup	.6475	.1998	19	.6624	.1914	20
PPB	.6220	.1763	19	.6274	.1582	20
GTW	.7627	.2385	19	.8063	.1903	20
GTA	.5802	.2494	19	.6032	.2452	20
NGTA	.5911	.2709	19	.5789	.2463	20
NGTW	.3184	.2488	19	.3357	.2970	20

Table 2
Mixed ANOVA Results Part 1

Tests of Within-Subjects Effects	df	F	Sig
Valence	1	2.374	.132
Valence*Condition	1	.351	.557
Response	1	22.393	.000
Response*Condition	1	.098	.756
Valence*Response	1	19.676	.000
Valence*Response *Condition	1	.002	.966

Part 2

Part two of the study included 18 participants in the control and 18 participants in the experimental condition. The participants ranged in age from 18 to 27 and included 8 males (22.2%) and 28 females (77.8%). Statistics of subjective experience of the rejection task on a scale of 0-5 (0 being no rejection and 5 being worst possible rejection) revealed a significant difference in the experience of the two conditions within the task. This was included to evaluate how effective the task was at eliciting rejection. Results showed that within the rejection condition (n = 18), 27.8% of participants rated the experience a 2, 61.1% rated the experience a 3, 5.6% rated the experience a 4, and 5.6% rated the experience a 5, (M = 2.8889, SD = .758). Within the control condition, 66.7% of participants rated the experience a 0 and 33.3% rated the experience a 1, (M = .3333, SD = .485). Results of a t-test between means of the subjective rejection ratings showed a significant difference between the rejection experience of the two conditions, shown in Table 3; $t(34) = -12.044$, $p = .000$. This provides evidence that the rejection manipulation was significant.

Table 3
Descriptive Statistics of Rejection Rating Between Groups

	Group							
	Control			Rejection				
	M	SD	n	M	SD	n	t	Sig (2-tailed)
Rejection Rating	.333	.485	18	2.889	.758	18	-12.04	.000

Mixed ANOVA analysis of main results were similar to Part 1 and revealed no significant results in terms of the experimental manipulation, but further confirmed the effectiveness of the Pavlovian influence task. There was no significant difference in accuracy due to valence within subjects. This means that there was no difference in performance accuracy between the combined win trials (GTW and NGTW) and combined avoidance trials (GTA and NGTA) within subjects; $F(1,34)=.109$, $p = .743$. Looking at Table 4 it can be seen that the combined means for win trials is very similar to the combined means for the avoidance trials. The valence*condition interaction was not significant, referring to no significant difference in accuracy due to valence between the conditions. This means that the difference in accuracy (or in this case, lack thereof) between the combined win and combined avoidance trials was the same between groups; $F(1,34)= 1.177$, $p = .286$. There was a significant difference in accuracy due to response within subjects. This refers to a difference in accuracy when participants were required to "Go" and when they were required to "NoGo"; $F(1,34)= 41.648$, $p = .000$. Table 4 shows that the combined means for "Go" trials (GTW and GTA) is greater than the combined means for "NoGo" trials (NGTA and NGTW). Similar to Part 1, this finding is reflective of a "Go" bias in terms of response within subjects, in which participants have higher accuracy when required to "Go" than when required to "NoGo". The difference in accuracy due to response was the same between the two conditions, resulting in no significance in the response*condition interaction; $F(1,34)= .220$, $p = .642$. The valence*response interaction was significant; $F(1,34)= 27.012$, $p = .000$. This interaction reflects the presence of the Pavlovian bias and more specifically, means that participants are driven to "Go" when expecting a reward (Reward Invigoration), and are driven to "NoGo" when expecting a punishment (Punishment Suppression). The level of Pavlovian bias was not different between the two conditions, rejection and control, resulting in no significance in the valence*response*condition interaction; $F(1,34) = .757$, $p = .390$. This indicates that the level of reliance on the Pavlovian bias was very similar between the two groups. These results reject the hypothesis and suggest that the mental state following online rejection does not have an effect on subsequent Pavlovian influence learning. However, the results do show the strong efficacy of the Pavlovian influence task.

Table 4
Descriptive Statistics of Part 2

Pavlovian Reinforcement Learning	Group					
	Control			Exclusion		
	M	SD	n	M	SD	n
RewInv	.6072	.1575	18	.6605	.1696	18
PunSup	.6463	.1874	18	.7000	.1878	18
PPB	.6267	.1674	18	.6802	.1681	18
GTW	.8450	.1335	18	.8563	.1085	18
GTA	.5894	.2767	18	.5644	.2912	18
NGTA	.6171	.2028	18	.6542	.2613	18
NGTW	.4021	.3239	18	.2862	.2411	18

Table 5
Mixed ANOVA Results Part 1

Tests of Within-Subjects Effects	df	F	Sig
Valence	1	.109	.743
Valence*Condition	1	1.177	.286
Response	1	41.648	.000
Response*Condition	1	.220	.642
Valence*Response	1	27.012	.000
Valence*Response *Condition	1	.757	.390

DISCUSSION

It is valuable to understand the natural responses to rewarding and punishing social interactions and how cues that predict these situations can alter learning and behavior. The automatically elicited responses that occur during social learning strongly influence the repetition or cessation of behaviors (Thompson & Westwater, 2017). These responses are often under the control of the Pavlovian bias. Examination of the contribution of the Pavlovian bias to behavior and learning within the context of social interactions is therefore worthwhile. The transient effect of social interactions has been well studied; however, it is not well known how the nature of these social interactions affects subsequent social learning. The goal of this study was to determine if negative social interactions affect reinforcement learning under the influence of Pavlovian biases.

The results of this experiment revealed no effect of the feeling of rejection or exclusion on Pavlovian reinforcement learning, rejecting the hypothesis that a negative social interaction would increase accuracy on withdrawal trials within the task (NGTW and NGTA). Part 1 did not include a measure of subjective affect, so it is difficult to know the level of effectiveness of the exclusion manipulation through Cyberball. Part 2, however, did include a subjective experience measure that showed the rejection task produced significantly greater feelings of rejection in the experimental condition than in the control. It is possible that these Pavlovian tendencies are largely innate within each individual and are not affected by mental state or affect. They are simply too engrained to be swayed by emotion.

An alternative explanation as to why the experimental manipulation in Part 2 did not affect Pavlovian reinforcement is that the level of cognitive load within the specific Pavlovian influence task may distract from any emotional state. The Pavlovian influence task is a complex task that involves a very high level of attention. It is possible that the amount of cognitive effort required to complete the task may cause an "overwriting" of any emotional mental state. Although evidence of the Pavlovian bias certainly exists within this task, it is nevertheless presented in a stimulus heavy manner, requiring constant attention for 30 minutes. Previous studies have shown that active distraction through playing an online game significantly increases pain tolerance, in comparison to baseline, or even passive distraction such as television (Jameson, Trevena, & Swain, 2011). Although the study was looking at physical pain, other studies have shown that brain activation during social exclusion is very similar to physical pain (Eisenberg, Lieberman, & Williams, 2003). This pressure of the task may distract from an emotional state, causing it to be forgotten about. It is possible that in a real setting, with more meaningful associations and less of a distraction from the pain, that negative social situations would affect the influence that the Pavlovian bias has over reinforcement learning.

Despite a lack of effect from experimental manipulation, this study did show how this task is extremely effective at eliciting and measuring Pavlovian bias. This is evident through the significant effect on performance accuracy due the response*valence interaction, reflecting the presence of Pavlovian bias. This interaction points to the Pavlovian tendencies to invigorate action in anticipation of reward, and suppress action in anticipation of punishment. The task also shows a significant "Go" bias in which participants are better at learning when to "Go" than learning when to "NoGo". This is also consistent with the Pavlovian bias. Similar Pavlovian reinforcement tasks have been used throughout a variety of experiments all resulting in similar accuracy patterns, with combined Pavlovian congruent trials having higher accuracy than Pavlovian conflict trials, in line with the Pavlovian bias (Guitart-Masip et al., 2012; Cavanaugh, et al. 2013; Thompson and Westwater, 2017). These findings continue to represent high external validity. The stability of the trends that appear in this task due to the influence of the Pavlovian bias is impressive.

Limitations

The limitations of this study include minor shortcomings in the sample. A larger sample size would have allowed for more in-depth review of the manipulation involved in this experiment by including more participants in each condition. The sample was also uneven in the ratio of males to females. Both parts included significantly more females than males, making it difficult to analyze the task based on gender. This would have been an interesting component to add as it is thought that males and females process emotions differently, with women being more likely to internalize negative emotions (Chaplin, 2015).

The rejection task involved in Part 2 was proven to be effective through subjective reports of the level of rejection (on a scale of 0-5) for both conditions. The reports, however, were mostly bimodal, with the majority of the control condition rating a 0 and the majority of the rejection condition rating a 3. This resulted in no correlation between the level of rejection and the Pavlovian influence measures. A finer grained scale for subjective reporting may have provided a greater range of answers, which would have benefited the analysis of correlations between these measures.

Lastly, the lack of a subjective rating scale in Part 1 was largely an oversight. This would certainly want to be included in future studies to allow for the evaluation of task effectiveness.

REFERENCES

Cavanagh, J. F., Eisenberg, I., Guitart-Masip, M., Huys, Q., & Frank, M. J. (2013). Frontal Theta Overrides Pavlovian Learning Biases. *The Journal of Neuroscience, 33*(19), 8541– 8548. DOI: 10.1523/JNEUROSCI.5754-12.2013

Chaplin, T. M. (2015). Gender and Emotion Expression: A Developmental Contextual Perspective. *Emotion Review : Journal of the International Society for Research on Emotion, 7*(1), 14–21. http://doi.org/10.1177/1754073914544408

Eisenberg, N. I., Lieberman, N. D., Williams, K. D. (2003). Does rejection hurt? An fMRI study of social exclusion. *Science, 302*(5643), 290-292. DOI:10.1126/science.1089134

Gray, J. R. (2001). Emotional modulation of cognitive control: Approach–withdrawal states double-dissociate spatial from verbal two-back task performance. *Journal of Experimental Psychology: General*, *130*(3), 436-452. http://dx.doi.org/10.1037/0096-3445.130.3.436

Guitart-Masip, M., Huys, Q. J. M., Fuentemilla, L., Dayan, P., Duzel, E., Dolan, R. J. (2012). Go and nogo learning in reward and punishment: interactions between affect and effect. *Neuroimage*, *62*(1), 154-66. https://doi.org/10.1016/j.neuroimage.2012.04.024

Gunther, M. B., Crone, E. A., & van der Molen, M. W. (2010). The heartbrake of social rejection: heart rate deceleration in response to unexpected peer rejection. *Psychological Science, 21*(9), 1326-33. doi: 10.1177/0956797610379236

Jameson, E., Trevena, J., Swain, N. (2011). Electronic Gaming as Pain Distraction. *Pain Research & Management: The Journal of the Canadian Pain Society*, *16*(1), 27-32.

Kawamoto, T., Nittono, H., & Ura, M. (2013). Cognitive, Affective, and Motivational Changes during Ostracism: An ERP, EMG, and EEG Study Using a Computerized Cyberball Task. *Neuroscience Journal 2013*. http://dx.doi.org/10.1155/2013/304674

Somerville, L. H., Heatherton, T. F., & Kelley, W. M. (2006). Anterior cingulate cortex responds differentially to expectancy violation and social rejection. *Nature Neuroscience, 9*(8), 1007-1008. doi:10.1038/nn1728

Thompson, J. C. & Westwater, M. L. (2017). Alpha EEG power reflects the suppression of Pavlovian bias during social reinforcement learning. *bioRxiv*. doi: http://dx.doi.org/10.1101/153668.

van der Veen, F. M., van der Molen, M. W., Sahibdin, P. P., & Franken, I. H. (2013). The Heart- Break of Social Rejection Versus the Brain Wave of Social Acceptance. *Social, Cognitive, Affect Neuroscience*. doi: 10.1093/scan/nst120

Williams, K. D., (2007). Ostracism. *Annual Review of Psychology*. *58*, 425-452.

Williams, K. D., Cheung, C. K. T., & Choi, W. (2000). CyberOstracism: Effects of being ignored over the Internet. *Journal of Personality and Social Psychology*, *79*(5), 748-762. doi:10.1037/0022-3514.79.5.748

Williams, K. D., & Jarvis, B. (2006). Cyberball: A program for use in research on interpersonal ostracism and acceptance. *Behavior Research Methods, 38*(1), 174-180. DOI:10.3758/BF03192765

BIOGRAPHY

Megan Bayne graduated from Fairfax High School, Fairfax, Virginia, in 2012. She received her Bachelor of Science in Psychology from George Mason University in 2016. This thesis is partial completion of her degree, Accelerated Master of Arts in Psychology with a concentration in Cognitive and Behavioral Neuroscience, from George Mason University in 2017.