9-2023

Brainwaves, Memory, and Reward

Rebecca McCune
The Graduate Center, City University of New York

BRAINWAVES, MEMORY, AND REWARD

by

REBECCA D. MCCUNE

A master's thesis submitted to the Graduate Faculty in Cognitive Neuroscience in partial fulfillment of the requirements for the degree of Master of Science, The City University of New York

2023

APPROVAL

Brainwaves, Memory and Reward
by

Rebecca D. McCune

This manuscript has been read and accepted for the Graduate Faculty in Cognitive Neuroscience in satisfaction of the thesis requirement for the degree of Master of Science.

Approved: September 2023

Jennifer Mangels, Advisor

Tony Ro, Executive Officer

THE CITY UNIVERSITY OF NEW YORK

ABSTRACT

Brainwaves, Memory, and Reward
by

Rebecca D. McCune

The development of effective educational curricula for enhancing learning involves the crucial consideration of effort and rewards. In the realm of education, teachers commonly employ rewards as motivational tools. Traditionally, these rewards are given to students as a recognition of their successful performance. However, a thought-provoking idea emerges: What if we were to extend rewards to students not solely based on accurate answers, but also on the effort they invest, even in cases where their actual response might be incorrect? Our study explores the potential impact of this approach on the way information is absorbed and subsequently retained, specifically focusing on corrective information.

Using Electroencephalography (EEG) we examined both behavioral indicators of error correction subsequent to the presence or absence of reward-based feedback, alongside extensively studied event-related potential (ERP) indicators associated with negative feedback and reward processing.

We found no statistically significant effect of reward on error correction in this small sample. However, participants were more likely to correct items they had initially answered incorrectly if the answer they gave had been “eligible” for a reward or was given with high confidence, regardless of whether it was actually rewarded or not, corresponding to previous findings. Interestingly, a trend for an interaction emerged suggesting that to the extent that rewarding effort had an influence on learning, it did so by motivating greater attention to the task overall, resulting in greater error correction following lower quality (i.e., reward ineligible) responses rather than to those items that were specifically rewarded. Although these results did not

yet reach statistical significance, they do support the value of continued research to explore the complex interrelationships existing between effortful rewards and task engagement in learning.

ACKNOWLEDGMENTS

I would like to thank my advisor Dr. Jennifer A. Mangels for allowing me the privilege of working with her on this project. It was truly an honor to learn from a leading expert in the field. Thank you for giving me the opportunity to work in the Dynamic Learning Lab, whose members' support and feedback has been invaluable.

Many thanks to Caesar Ekya for his EEG and coding expertise and Daisy Reyes for lending her exceptional statistical skills. I would like to acknowledge their unwavering support and abundance of patience throughout many late nights in the lab.

To my research assistants, Alexandra Bunde and Jennifer Lai, this project would not have been possible without the hours of work you devoted to the team. I cannot express how much your dedication means to me.

Thank you to my parents, brother and friends for their unconditional love and support throughout the entire program.

I am grateful to Dr. Gregory E. Tasian and my research team at the Children's Hospital of Philadelphia for providing me with the fundamental skills necessary to conduct research.

Lastly, I would like to acknowledge my neuroscience cohort for whose shared struggles turned into amazing friendships and research. We did it!

Contents

FIGURES

Chapter 1: Introduction

Effort and reward are key elements to consider in developing educational curricula for efficient learning and memory. Rewards can be an essential tool in facilitating optimal learning both inside and outside the classroom. Teachers often use rewards to incentivize students with everything from a gold star sticker to extra credit points. Traditionally, however, rewards are only given when students perform well, but that may discourage students from attempting to answer questions for which they are less sure about the answer.

Research indicates that putting in effortful attempts to engage with a problem, even if those attempts aren't entirely accurate, results in more significant learning benefits for corrective information compared to simply being exposed to that information passively (Metcalfe, 2011). To encourage this type of engaged problem solving, what if we reward students, not only for correct responses but also for effort, even if their actual answer is incorrect? We want to investigate how this approach potentially alters the encoding and subsequent memory of corrective information. We will test the effects of rewards both at the task engagement level (Does the prospect of a reward for effort improve the quality of answers given and learning regardless of whether the specific item is rewarded?) and at the specific item level (Does reward boost learning more for those items that are rewarded than not rewarded?). Utilizing Electroencephalography (EEG) we will measure event related potentials (ERPs) to gain insight into the mechanisms underlying how the presence of effortful rewards may influence the neural response to feedback and encoding of corrective information.

Specifically, we aim to understand the cognitive and neural mechanisms by which rewards given for putting effort into retrieving answers to general knowledge questions (i.e., a common, educationally relevant type of retrieval task) influence the response to and ability to rebound from (i.e., correct) errors. We will be answering our research questions by using a semantic memory retrieval and learning trivia task which asks questions providing accuracy feedback followed by a subsequent surprise retest. In this study we operationalize effort as a semantically plausible answer to a question (referenced to a semantic knowledge database www.mangelslab.org/bknorms), regardless of response accuracy.

Effort

Effort can be difficult to measure because there are various ways of quantifying it, from measuring physical exertion with stress gauges to measuring subjective experience of effortful engagement (Epel et al., 2018). Here, we are interested in measures of mental effort during task performance in learning environments. According to Smith and Walker (1993), one way to characterize mental effort toward a task is the quality of a response based on its proximity to the accurate response. In a study with general knowledge, there will be some more difficult questions where the correct answer will not come immediately to mind. In that case, it would take more "memory search/retrieval effort" to find a plausible answer than to just move on to the next question.

Previous research such as Huelser & Metcalfe (2011) has observed a positive impact on learning when initial retrieval attempts were semantically linked to the target answer, as opposed to attempts with no relation. In their experiment using related and unrelated word pairs they found when the cues and targets in their task were related, error-generation led to the highest correct retention while no benefit was derived from generating an error when the cue and target were unrelated. However, they could not conclude producing an error was due to mental effort or engagement during generation itself, as there was no benefit from incorrect guessing for the unrelated materials. Similarly, Butterfield & Mangels (2003) found that when individuals corrected their errors, they were more likely to already possess familiarity with the correct answer. This suggests that both the familiarity of the correct answer and the semantic relatedness of the incorrect answer play a role in mediating error correction. In other words, even when the student gets the initial answer incorrect, they have a higher chance of remembering it after giving a semantically related "effortful" attempt rather than if they give an unrelated answer or only see the answer without making any attempt (i.e., passive viewing (Metcalfe, 2017). Giving effort toward retrieving higher-quality, semantically related answer requires attentional engagement, which refers to "the process of intentional, sustained allocation of cognitive resources to guide problem solving, planning, sensemaking, and decision making" (Ocasio, 2011). Thus, motivating attentional engagement forward effortful answers should ultimately yield greater benefits from corrective feedback. Reward may be a potentially powerful method of increasing this type of motivation.

Reward

Rewards come in various forms and are contingent upon the specific context of the individual's surroundings. Rewards are reinforcers which can be physical and satisfying of a physical need like hunger or thirst, social (praise from important other), emotional (thrills), mental (positive accuracy feedback) or indirect (i.e., getting a raffle ticket that is a route to reward such as obtaining candy) (Davis et al., 2006; Schultz, 1998). Cognitive Evaluation Theory (CET) highlights the decline in intrinsic motivation often resulting from the use of over-rewarding to extrinsically induce behavior, diminishing the inherent drive of individuals to engage in the assigned task (Deci et al., 1999). CET specifically addresses how internal motivation suffers when rewards are offered solely for participation, decreasing the feeling of rewards as related to agency and competence. Thus, while rewarding for participation is commonly used in young children to promote self-esteem and excitement, it can interfere with longer-term motivation to put forth effort on a task (Bühren, 2013).

In contrast, intermittently rewarding individuals for their competency in task engagement, regardless of the final outcome — essentially acknowledging their effort to generate a high-quality response even in the face of an unsuccessful attempt — holds particular promise for a demanding declarative memory task that requires prolonged attention and may be particularly challenging for individuals (Abraham et al., 2019). Granting rewards as recognition for successful learning holds the potential to enhance overall task involvement and enhance attentive focus on relevant information to earn rewards (Wittmann et al., 2005). Although research on the impact of rewards has been well studied, much of the research has been on rewarding for learning specific things. However, there is a paradox because reward could reinforce the wrong answer (Eckert et al., 2023). Our study hopes to go beyond this and add to the field of memory by providing information on how non-traditional rewarding (rewarding effort, not just accuracy) leads to improvements in learning.

Rewards seem to be more influential on supporting learning for less interesting or less confident information. Murayama and Kuhbandner (2011) demonstrated that memory improvement linked to rewards occurred for information that participants regarded as "boring," whereas there was no such enhancement for information they deemed interesting. Due to the overall lower intrinsic interest of the participant, this left them more malleable to the influence of extrinsic rewards (Kang et al., 2009; Gruber et al., 2014). In line with this finding, Abraham et al. (2019) found that the general effect of reward context was most apparent for trials endorsed with low confidence.

Although it might appear logical that errors made with a strong sense of certainty would pose a greater challenge to rectify compared to errors made with less ingrained, low confidence, evidence presents a contradictory outcome: errors initially made with high confidence are actually more likely to be corrected successfully than those made with low confidence. This unexpected discovery is termed the hypercorrection effect (Butterfield & Metcalfe, 2001). Originally, it was thought that high confidence errors would be harder to correct due to there being a firmly established memory. Previous memory models suggested that making an error would only strengthen the incorrect response, leading to overall low recall of accurate information and high frequency of remembering the incorrect response given (Mozer, Howe, & Pashler, 2004). In order to correct this error going forward, you would need to update the already well-established memory, which was thought to be more difficult than correcting a weaker memory (with a low confidence) as it would be easier to override a weak connection.

However, Butterfield & Metcalfe (2001) found evidence of the opposite. One reason for the hypercorrection effect is the element of surprise. Due to the high confidence the person would be surprised that they are receiving unexpected feedback (negative feedback that they were wrong) leading to better correction as this information was "flagged" due to the surprise result. This is similar to Kulhavy's (1976) findings that proposed a large discrepancy between the participant's initial beliefs and the correct answer leads the participant to expend more effort to correct the misunderstanding. Based this previous research we predict the effects of rewards in our study will be greatest for low confidence incorrect but eligible answers at the trial level, and for items that may not even be eligible for reward if reward is helping at the task level.

Event-related Potentials

Electroencephalography (EEG) helps investigate the neural population under an electrode which is part of activity across a larger network. While this method is often used clinically to diagnose epilepsy and other neurological disorders, it can be used in research to record cognitive processes. EEG provides a unique temporal insight down to the millisecond of brain functioning. For example, this level of precision can detect neurological changes occurring during the encoding and retrieval of information.

Event-related potentials (ERPs) arise by extracting and subsequently averaging numerous temporal segments from the continuous EEG signal, delineating the targeted event. The amplitude, latency, and distribution of the resulting positive and negative deflections are employed as indicators for the fundamental cognitive processes occurring (Luck, 2005). The three main ERPs of interest for this study are the P3 or P300, Feedback related negativity (FRN) and P2 or P200 which are all well studied and validated measures of neural processing (Ibanez, 2012). ERPs related to feedback and attention are both key elements of interest in our study due to their relationship with initial encoding, and retrieval of memories. Certain ERPs associated with memory and their presence may indicate subsequent memory effects (Ibanez, 2012).

P3 (P300)

The P300 (or P3) component is one of the most studied ERPs and is described as engaging higher-order cognitive operations related to selective attention (Polich, 2007). There exist two subcomponents known as P3a and P3b. The P3a, which displays a more frontal distribution, emerges after an unexpected event at around 250-350 ms, irrespective of the stimulus type and is commonly linked with the automatic modulation of attention. P3b occurs more broadly from 300-500ms and is connected to attention, working memory, and higher-order cognitive functions. It is maximal over centro-parietal locations (Koivisto & Revonsuo, 2010) and is also linked to motivation, sustained attention, novelty, and other psychological processes associated with social cognition tasks (Friedman, Cycowicz, & Gaeta, 2001). In general, the frontal and parietal P3 components may represent anterior automatic and posterior controlled attentional processes, respectively. In the current study we expect the P3 to the accuracy feedback (regardless of accuracy) in the reward framing condition, will be enhanced because of overall greater engagement and task effort, compared to the control framing. Effects on the P3 are expected to be greatest for trials that are effortful/reward eligible as opposed to clearly omitted type responses (e.g., “idk” responses).

Feedback Related Negativity (FRN)

Feedback related negativity (FRN) is another well-studied ERP component (Ibanez, 2012). The FRN is a negative-going waveform that is particularly sensitive to feedback valence (positive vs. negative; REF). It is measured 200-300 ms after the onset of the accuracy feedback over fronto-central electrodes and likely generated in part by the anterior cingulate cortex (ACC) (Miltneretal.,1997; Gehring and Willoughby, 2002). As the name suggests, it is related to receiving negative feedback (such as monetary losses or incorrectness). Research has supported the idea that a greater FRN is associated with better encoding because the FRN indexes attention to the stimulus as well as registration of valence (de Bruijn et al., 2020). However, in Butterfield and Mangels (2003) and all subsequent studies in our lab, we found that the FRN is not predictive of encoding in declarative memory, even though it is still sensitive to feedback valence. To the extent that the FRN is particularly sensitive to valence, we predict that reward may reduce the amplitude of the FRN if the prospect of reward reduces the negative valence of negative feedback.

P2 (P200)

The previous two ERPs have a performance focus on expectation and effort. The P2 is also important for this study as previous research has shown that it is most impacted by the presence of a reward, thus giving it the name the "reward" ERP (RewP). While ERPs can't directly measure the midbrain and basal ganglia structures associated with dopamine-mediated reward processes, the P2/RewP *does* seem to demonstrate association with reward (i.e., reward sensitivity). It is typically measured 200-400 ms after the onset of a reward stimulus (Ibanez, 2012) maximally over centro-frontal and the parieto-occipital areas. The RewP is likely related to the P2 that has been researched using oddball paradigms and visual search tasks and is thought to be a neurological response to visual stimuli exhibiting selective attention (Cao, 2021). In the present study we predict the P2 to the reward-relevant stimuli will be bigger to the designated target stimulus (yellow disk/coin) vs. the non-target stimulus (blue disk) in both the reward and control blocks because both targets should be associated with greater visual attention regardless of condition. However, when the target represents a discrete monetary reward, the difference should be larger between the target and non-target.

Previous Research

Abraham et al. (2019) used a test-retest general knowledge task paradigm to identify the potential influence of rewards based on effort on incidental learning. Their trial-level reward was a raffle of a ticket which increased the probability of receiving a larger reward later. There were four different groups used in that study design; the effort group (reward for effortful responses), luck group (rewards as related to participant-chosen lottery numbers), random award group (computer generated reward, no control, not competence-based and a control group (no reward). Framing the task as giving rewards for effortful responses resulted in greater error correction overall compared to the control condition (no differences between the other conditions), but only in women. This improvement seemed to be related to the prospect of getting rewards rather than trial-level effects of receiving a reward because there were no differences in error correction for rewarded and non-rewarded trials overall. However, as previously mentioned, the reward framing generally benefited the lower-confidence errors more than higher-confidence errors. These results suggest that both reward context and gender are important factors contributing to the effectiveness of rewards as tools to enhance learning from errors.

Recently, Farber and Mangels [in prep] conducted two studies (survey and experimental) the results of which piqued a continuing interest in adding to this literature. Their survey study showed significant differences between the motivational potential of monetary and social rewards, and a trend for males' preference of money and females' preference of private social praise. The experimental study looked at the effects on memory and showed that there was no effect of reward on item-specific memory, only an effect of reward eligibility. Subsequent memory was best for the reward eligible items that had shown "effort" compared to the ineligible items, in line with previous research. However, this study has not yet included a control condition (for receiving no rewards) and that addition will be essential in understanding whether both monetary and social rewards improve memory compared to the control, as seen in Abraham et al. (2019).

The Present Study

We will use both behavioral measures of learning and neurophysiological measures of brain activity to understand how incentivizing effort, rather than accuracy, changes how people attend to negative feedback and learning opportunities in a general knowledge task. We will look at both behavioral measures of error correction after receiving or not receiving reward feedback and well-researched event-related potential (ERP) markers of negative feedback and reward processing. We are also interested in reward presence impact on task engagement, hypothesizing reward presence will enhance overall task engagement. This would be evident in the number of reward-eligible responses or the overall difficulty of the questions needed to maintain a similar level of performance across the reward and control conditions, as well as in self-report measures of concentration and motivation.

We hope to use the information we learn from the ERPs to better understand the cognitive and physiological mechanisms of how the presence or absence of a reward following an error impacts the response to and ability to rebound from (i.e., correct) negative feedback. We will measure the amplitude of the feedback related negativity (FRN), P2, and P3 waveforms. Although past research (including research in this lab) has shown that giving intermittent rewards for effortful responses, even when the actual answer itself is incorrect, can improve task engagement and error correction (Abraham et al, 2020; Farber & Mangels, in prep) we hope to expand upon these findings by adding ERP metrics of underlying cognitive mechanisms to help understand behavioral outcomes.

The current research study is set up similarly to Abraham et al. (2020) and Farber & Mangels (in prep) and aims to replicate and extend their findings. We will examine both behavioral indicators related to error correction in response to the presence or absence of reward feedback, as well as extensively studied event-related potential (ERP) indicators associated with negative feedback and the processing of rewards. Specifically, we hope to better understand how effortful rewards may impact incidental learning by utilizing the same test-feedback-retest paradigm. However, unlike either of those previous studies, this study uses a within-subjects study design to minimize noise in potential participant differences as everyone will act as their own control. Similar to Farber & Mangels (in prep), we will use more discrete rewards after each trial, using the monetary (25¢) option, but also include a control condition similar to the counting control used by Abraham et al. (2020). Additionally, by looking at the specific aspects of the EEG (i.e., event-related potentials [ERPs]) that are tied to the onset of accuracy feedback and reward events, we can examine whether the opportunity for reward changes the immediate response to that negative (and positive) feedback. Additionally, we can examine the neural response at the time of both the feedback stimulus and subsequent reward stimulus, and whether these are related to the successful encoding of the correct answer.

As the previous reward learning studies in the lab, participants will be asked to answer a series of general knowledge questions across a range of common academic topics (literature, geography, science, arts, etc.) during an initial test in which feedback regarding the correct answer is provided, and then will be given a surprise retest on the same questions. We can determine the measure of incidental learning occurred by looking at their retest performance. We are interested in incidental learning which is similar to typical everyday learning most commonly occurring in real life (Do Carmo-Blanco, 2019).

The influence of reward will be manipulated by contrasting two conditions: reward vs. counting/control conditions. These conditions will be counterbalanced with differences in task framing directions and the stimuli relevant disks shown. After participants enter a free response answer to the question, they will be given "accuracy feedback" which provides them the correct answer and thus also signals whether their provided answer was incorrect. Following the accuracy feedback, a colored disk will be shown that, in the reward condition, signals whether they have received a reward of 25¢ for their effort on that trial or not. In the counting control condition, it will simply signal the need to count. Although the yellow disk does not signal reward in the counting condition, participants are instructed that the closer their estimate of the number of disks to the accurate total number of yellow disks shown in each block, the more money they will earn. The frequency of receiving a reward or counting stimulus will be the same and about 40-50% of trials; the no reward or ignore blue disk stimulus occurring during the remaining 50-60% of the trials.

Our Research Questions

The three main research questions of this study will be 1) How does reward influence error correction in males and females and does it provide any specific trial-level benefit for correction or is that benefit only the task level?; 2 Does reward stimulus framing change how individuals process and attend to negative feedback as measured by the amplitude of the feedback related negativity (FRN) and P3?; and 3) How do neural responses P2 and P3 ERPs, which are associated with the presentation of a reward and should also be modulated by the reward manipulation, relate to successful error remediation?

To address the first research question, we will compare rewarded, non-rewarded/eligible, non-rewarded/ineligible items. While we expect to replicate prior work and find that reward-eligible answers (i.e., effortful answers that are semantically related to the question/correct answer) to have a greater error correction than non-eligible items, if reward influences error correction at the trial-level, we would expect the rewarded and eligible items would be remembered better than the non-rewarded but eligible items. If reward effects error correction overall (i.e., at the task level), we expect the reward condition to have a greater effect on error correction than the control condition overall.

We will also look at error correction as a function of confidence level in each condition. While we expect that we will see evidence of the hypercorrection effect (higher confidence error correction > lower confidence error correction), we expect that the opportunity for reward might be particularly beneficial to error correction of lower confidence error correction by generally increasing motivational relevance of all items.

To address the second question, we will look at the feedback-related negativity (FRN) and P3 event-related potentials (ERPs) that are elicited by the negative and positive feedback. We predict that in both the reward and control conditions, negative feedback will elicit a larger FRN than positive feedback, but it is possible that in the reward condition the FRN would have a smaller FRN due to the remaining possibility of reward. However, the potential for reward may correspondingly enhance the amplitude of the subsequent anterior P3 to incorrect answers. The prospect of reward, even after error, may increase attention to the accuracy feedback, even when an error is made with lower confidence.

To address the third research question, we will look at the P2 and P3 ERPs associated with the presentation of a reward and see if they will predict items that are later corrected and not corrected on a surprise retest. We hypothesize reward presence will enhance P2 and P3 ERPs and therefore encoding of the correct response leading to increased error correction on the surprise retest compared to the control trials.

Chapter 2: Methods

Participants

Seven subjects (6 females) were recruited from a larger cohort of CUNY Baruch College and Graduate Center researchers. Six of the subjects volunteered as lab members for the pilot phase of the study and did not receive compensation (pre-IRB approval). One participant volunteered through the SONA Psychology and Management voluntary subject pool (post-IRB approval) and received monetary compensation at the rate of $15.00/hour and a bonus of $30.00.

Design and Procedure

Overview

The study took place over two days separated by 24-48 hours. On the first day subjects provided consent, were set up for EEG monitoring and answered 200 general knowledge questions divided into two sets of 100 questions each, which were further broken up into two blocks of 50 trials. Each set of 100 questions was tested under either the reward condition or counting control condition, in a counterbalanced order. Following each set of 50 questions, participants filled out a questionnaire about their experiences during that block.

A titration algorithm was used to maintain a stable accuracy rate of 30% correct (i.e., 70% failure) for all subjects across all four blocks. When participants returned to the lab 24–48 hours later, they completed a second set of general knowledge questions with feedback but no reward or counting component. The questions were a repeat of the 200 questions they had received during the first test, although during the first day of testing they had just been told that there would be a second set of questions when they returned but not given any further information. No EEG measures were captured on the second day of testing. After completing the questions, they filled out a questionnaire regarding their experiences in the task as a whole.

First Day of Testing

Instructions

After providing informed consent (pre-IRB pilot participants did not provide consent), general instructions for the task were presented visually with a simultaneous voice-over read by the principal investigator who was also a professor and chair of the Psychology Department at Baruch College. The principal investigator framed the task overall as emphasizing how the Psychology department is invested in the success of their students and wants to learn more about what motivates students to succeed in their classes, particularly their tougher classes where students have to invest a fair amount of effort because the class isn't an easy "A".

Then participants were instructed about the task itself and that both correct answers and "effortful" but wrong answers were valued. The instructions further clarified that an "effortful" incorrect response that were which could be considered a "good try" as opposed to simply writing "idk" and not attempting to answer or an unrelated response (Abraham et al., 2020).This was operationalized as a response that had a short semantic distance between the answer they gave and the correct answer (e.g., "pine" has a shorter semantic distance to spruce than "shoe"). Semantic similarity was coded in a large database of general knowledge questions and answers created by our lab (B-KNorms; www.mangelslab.org/bknorms).

The general introduction video was followed by specific instructions for either the monetary reward condition or the counting control condition detailed below. In the reward condition participants read and simultaneously listened via headphones to the professor read aloud the following instructions:

Introduction to both conditions:

> “In this research study, you’ll be asked to answer some general knowledge questions — basically “trivia” questions — some of which are relatively easy, and you’ll know the answer to right away, but some of which are pretty challenging. For these, you might find yourself only able to take an educated guess...
>
> Each set will take about 30-40 minutes to complete, so it will be somewhat of an effort to stay on task. Since it’s a long task, we’ve added some things to hopefully give you the motivation to keep giving your best effort on every question — even the harder ones and even throughout the whole task.
>
> Before we get to the motivation part, what do we mean by “best effort.”
>
> We just mean trying to provide at least a “plausible” response rather than just skipping the question or putting in a clearly wrong response like “IDK” just to “get on with it.”
>
> How will we know if you’re making an effort?
>
> We will be matching your answer to a database of responses to determine whether your answer is correct. If it’s not correct, we can use the database to look at whether it was a “good, educated guess.” This database was created from over 85,000 responses to these questions by Baruch students just like you.
>
> How are we going to try to motivate you to keep giving this best effort throughout the task?
>
> In each set of 100 questions, you are going to be able to earn some cash rewards based on your performance, in addition to the basic compensation you receive for your time in the study.

Next, I'll tell you a bit about what you'll be rewarded for on this first set of 100 questions. For the second set, the way you'll earn rewards is different, but I'll tell you about those instructions right before that set so not to confuse things.

Each trial will begin with a question.

SAMPLE QUESTION: What California city did the last Pony Express ride end in?

Below each question, there will be a space where you will type in your best 1-word response to the question. When you are happy with your response, press the ENTER/RETURN key. You will have a limit of 3 minutes to answer each question. You should always do your best to provide a reasonable answer in that time. Many times, an educated guess turns out to be correct.

SAMPLE: What California city did the last Pony Express ride end in?

In the example above, if an answer does not immediately come to mind, you should narrow your choices to cities you know within California. Next, you also know that the answer must be only 1 word, so you should eliminate cities such as "San Francisco" or "Los Angeles". Remember, if you write either of these 2-word answers as a compressed 1-word answer (e.g., SANFRANCISCO), it will automatically be scored as incorrect. Also, note that educated guesses to this question would NOT include answers such as "TIBET", which is not a city and not in North America; "EXPRESS" which just repeats a word in the question - remember, correct answers will never be words taken from the question unless otherwise specified; or "CLUELESS", which is a "cop-out" answer.

If you misspell an answer, don't worry. The program has a built-in spell checker to help you. The spell checker will not necessarily provide you with the correct answer, and it has limited alternatives. So, if you don't want any of the provided options, you may return to the question screen by selecting the "Go Back" option at the bottom.

Once you submit your response, you will use the number pad to rate your confidence in the accuracy of your response on a 1-7 scale.

Let's review the sample shown below:

How confident are you in the answer you provided?

1 = you are ABSOLUTELY SURE that your response is WRONG

2 and 3 = you are PRETTY SURE that your response is WRONG

4 = you are UNSURE whether your response is WRONG or RIGHT

5 and 6 = you are PRETTY SURE that your response is RIGHT

7 = you are ABSOLUTELY SURE that your response is RIGHT

We encourage you to USE THE FULL CONFIDENCE SCALE. Please use '1' only when you are absolutely sure your response is wrong and '7' only when you are absolutely sure your response is right. are absolutely sure your response is right.

After you have entered your response and confidence for a given question, you will be given feedback about the accuracy of your answer. First an orientation cross (+) will appear in the center of the screen indicating where to direct your gaze as you receive the feedback. If your response to the question is correct, you will see the correct answer in GREEN and hear a higher-pitched tone. However, if your response is incorrect, you will see the correct answer in RED and hear a lower-pitched tone."

The following portions were shown based on the condition being explained (counterbalanced):

Reward Condition

"In this set of questions, you will be eligible to receive a 25¢ reward on trials where you give an answer that's correct or, if it's wrong, is at least a good guess that is close" to the correct answer (i.e., an effortful response).

You won't get a reward on every trial that is eligible, but aiming for correct responses gives you a chance at a reward. You definitely won't get a reward if you don't even try to give a good answer.

When you get a reward, you'll see a yellow disk with 25¢ in the center and hear a "cha-ching" sound. That 25¢ is real money you've earned!

If there's no reward, you'll see a blue disk and hear a short ”boop” sound instead.

While you won't know exactly which trials you'll be rewarded on, you can increase the chances that you'll get a reward by putting in your best effort, even when you aren't 100% sure of the correct response.

If you are putting in consistent effort, you should see a yellow disk with the 25¢ reward on close to half the questions. This can add up to an extra $15 BONUS.

You receive this bonus at the end of Day 2 after we've had a chance to verify your responses and add up the rewards.

After you find out if you are right or wrong and what the correct answer is, you'll see if you got a reward or not.

Remember that if you get a reward, you'll hear that "cha-ching" sound and see a yellow disk with a 25¢ in the center.

If you see blue disk and hear a "boop" noise you don't get a reward.

You'll receive your bonus for this set of questions on Day 2, after we've had a chance to review your answers and count up your reward total.

Remember that giving a "good" answer to a question makes you eligible for a reward, although it doesn't guarantee one. Still, giving your best effort throughout the set increases the chances of a larger reward total at the end.

Finally, remember these hints about how to give a "good" answer.

1. All of the answers in this study are only <u>one word</u> (if you type in 2 words, it's wrong).

2. You don’t have to have perfect spelling. But if you are prompted with the spell-checker, use it to help increase your chances of getting the answer correct. But while the spell-checker is pretty good, it isn’t perfect either and may sometimes offer a word that doesn’t make sense.

3. Even if you are not sure of your answer, try to give a response that is at least in the same category as the answer the question is looking for – you might just end up correct!”

or

Counting Condition

“In this set of questions, we are interested in how people answer questions when they have to keep other information in mind at the same time. You will be rewarded based on your ability to do both tasks well.

Specifically, we will be asking you to keep a running count of the number of yellow disks you see during a block of 50 questions. You won’t see these disks at the end of every trial, just some trials.

When you need to count the trial, you’ll see a yellow disk and hear a “tweet-tweet” sound. When you don’t need to count the trial, you’ll see a blue disk and hear a short ”boop” sound.

While you won’t know exactly which trials you’ll be asked to count, try to keep an accurate running tally because you’ll be asked for the total number at the end of the block and you will receive a bonus based on the accuracy of that count.

If you get a perfect count on both blocks you will receive a $15 bonus. If you are “off” from the correct count (either above or below), we will subtract 25¢ from that bonus for each one you are off. For example, if the accurate count is 20 and you say 18 or 22, you will receive $14.50.

You only have to remember the number of yellow disks in a block of 50 items, not over the entire 100 items. The size of your bonus on this set of questions is related to the accuracy of your count on both blocks. Whether you see a yellow disk or not isn't tied to your performance on the question before it. You'll see yellow disks you need to count after both correct and incorrect responses.

However, even though your bonus amount in this set is tied to the accuracy of your count, you should still try to give a "good" answer to each question. If your effort on the general knowledge questions is too low, we will conclude that you were not trying hard enough on that part of the task, and you will forfeit the bonus for the counting task, no matter how accurate your count was.

You'll receive your bonus for this set of questions on Day 2, after we've had a chance to verify your count and review your answers to the questions."

After you find out if you are right or wrong and what the correct answer is, you'll see if you have to count this trial or not.

If you hear a "tweet-tweet" sound and see a yellow disk, you'll need to add that trial to your total count for that block.

If you see blue disk and hear a "boop" noise you don't have to count it.

You'll receive your bonus for this set of questions on Day 2, after we've had a chance to verify your count and review your answers to the questions.

Remember that even though the color of the disk has nothing to do with your accuracy on that trial, you should still try to give "good" answers to each question because that effort still needs to meet a certain threshold for you to be able to get the bonus for the counting task.

Finally, remember these hints about how to give a "good" answer.

1. All of the answers in this study are only one word (if you type in 2 words, it's wrong).

2. You don't have to have perfect spelling. But if you are prompted with the spell-checker, use it to help increase your chances of getting the answer correct. But while the spell-checker is pretty good, it isn't perfect either and may sometimes offer a word that doesn't make sense.

3. Even if you are not sure of your answer, try to give a response that is at least in the same category as the answer the question is looking for – you might just end up correct!"

Participants then answered 100 questions in two 50-trial blocks under one condition, after which they took a 15-minute break, then received instructions for the other condition and answered a second set of 100 questions, also broken into two 50-trial blocks. Condition order was counterbalanced across participants.

The 200 questions asked of each participant were randomly selected from a larger set of 434 normed questions tapping a variety of academic domains, including science, history, music, art history, literature, geography, and religion. The current version of this question set is available as the B-KNorms; http://www.mangelslab.org/bknorms. The B-KNorms is a database collected from 498 City University of New York (CUNY) students. These students provided over 85,000 responses to general knowledge questions included in the database. The average accuracy of each question by dividing the raw frequency with which the correct answer was given by the number of times the question was sampled. Thus, a question with a higher accuracy score is relatively easy because the correct answer was given more frequently.

A titration algorithm was used to maintain each participant at a stable, but low, accuracy rate of 30% throughout the task (for algorithm details see Butterfield & Mangels, 2003). Titration was done to ensure consistency in intra and inter-individual levels of negative feedback regardless of participants' pre-existing knowledge.

Trial Structure

Figure 1. shows a sample trial

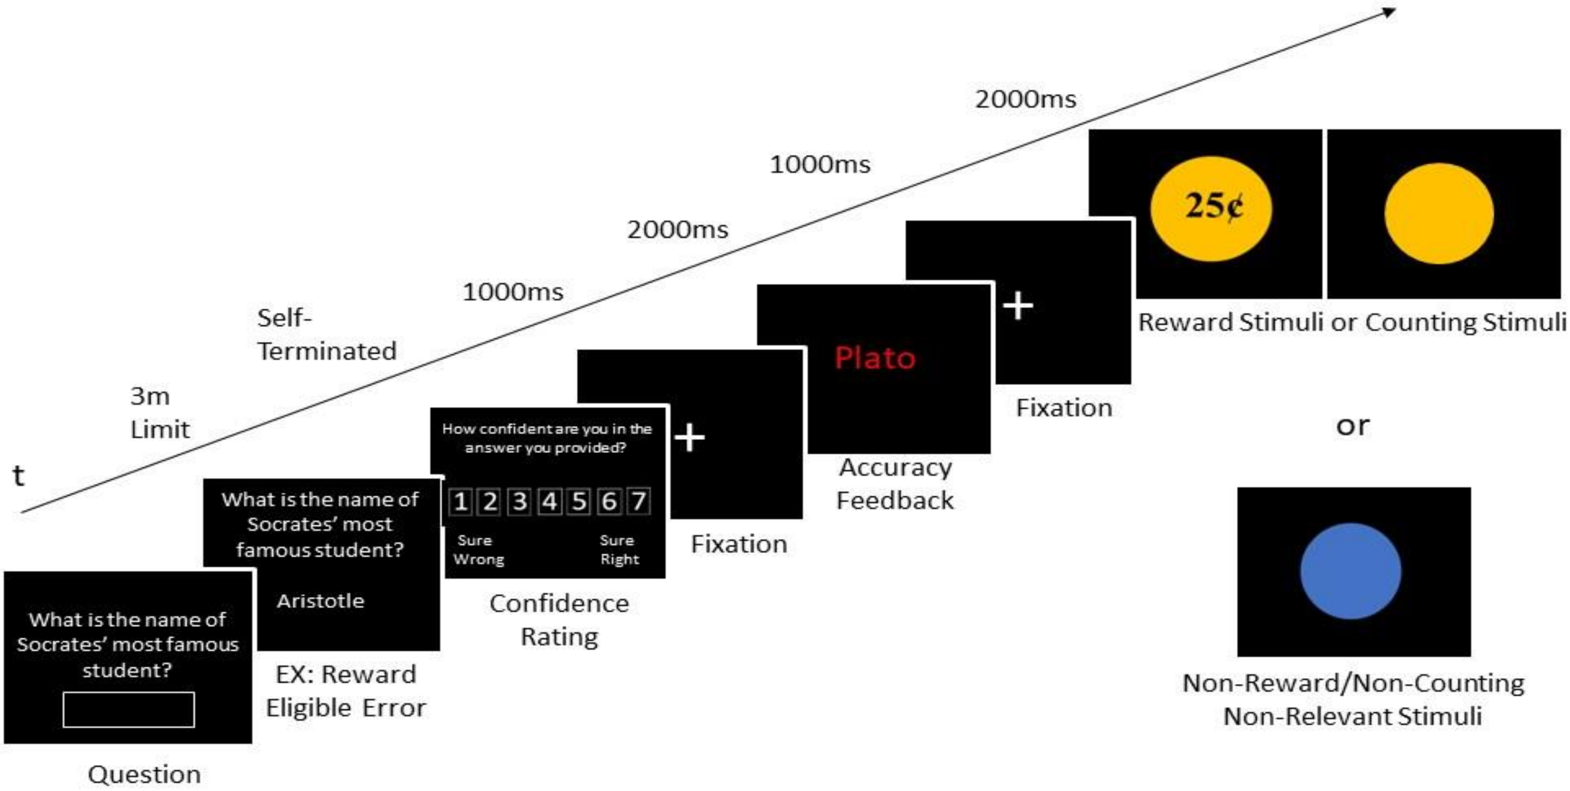

FIGURE 1 | First-test trial structure. At the start of each trial, the participant is shown a question and types their answer in a blank space below. In this example, the answer entered is ARISTOTLE, which is incorrect, but reward-eligible because it is a plausible, semantically-relevant response. After the answer is entered, the participant would be prompted to rate their confidence in their answer's accuracy on a 7-point scale (1 = sure wrong, 7 = sure right). Then, after a 1 s fixation point to orient attention to the center of the screen, the correct answer would appear for 2 s, either in red if their initial answer was incorrect, as in this example, or in green if it was correct. This accuracy/learning feedback is then followed by another 1 s fixation. Then, because this example shows a reward-eligible answer, the final stimulus would be either reward or no-reward feedback, assuming that this trial occurred in any of the three rewarded groups (in the Control group these stimuli function as target or non-target symbols to be counted). Both of these reward options would also have been possible if the participant had initially provided the correct answer (i.e., PLATO). However, if the participant had given an incorrect response that showed little effort and thus, was reward ineligible (e.g., SMITH or IDK), only non-reward feedback would have been possible as the final trial stimulus.

Questions were presented one at a time, and participants had a 3-minute window to provide their best response. If they did not know an answer, they were prompted to type their best guess. If they waited out the 3-minute window without responding they would be marked incorrect and automatically move to the next question. After submitting their response, participants rated their confidence in its accuracy (1 = “sure wrong,” 4 = “unsure,” 7 = “sure right”). The feedback sequence that followed was a blank screen for 500 ms, followed by a fixation crosshair for 1000ms.

Next, accuracy feedback (2000ms) was given both visually and aurally in the pairing of the correct answer (green-colored if correct with a high tone and red-colored with a low tone for incorrect responses). After another 1000ms crosshair, reward/counting relevant stimuli was presented for 2000ms. This consisted of either a yellow twenty-five cent coin accompanied by a cha-ching money sound in the experimental group (this target stimuli indicated whether a “reward” was earned or not for the given response) or a plain yellow disk with a high-pitched sound in the counting control group. Both conditions saw a plain blue disk with a neutral (sound if the non-target stimuli was shown. Following the reward/counting relevant stimuli was a 500ms blank pause and then the next question.

Post-Block Questionnaire

After each block, participants were asked to complete a post block questionnaire. At this point, the participants were asked to indicate the number of yellow disks they had been shown. We then asked questions (utilizing a Likert scale format) about participants' subjective experience of receiving feedback and rewards, as well as the difficulty of the task. The responses provide self-reported information reflecting how participants feel towards the negative and positive feedback and to the reward-relevant stimulus when it shows up or does not show up. We also asked about overall concentration and motivation during that block.

Second Day of Testing

The retest trial sequence differed from that of first-test in that no reward/counting relevant stimuli were shown or post block questionnaires. Day 2 was identical to Day 1 as correct answers appeared in green text and were paired with the same high tone, and incorrect answers appeared in red text and were paired with a low tone. Similar to Day 1, this answer feedback was presented immediately following subjects' confidence ratings and lasted for 2000 ms. Question order was randomized with the exception that, to decrease variability in study-test delay and preserve some aspects of test context, all questions from first-test blocks 1 and 2 (first condition) were randomized within the same condition/group initially shown separately from the questions from first-test blocks 3 and 4 (second condition).

Only at the outset of this second day of testing were subjects made aware that they were being retested on a subset of questions they had answered on the first day. After the final block, participants were shown an incidental memory probe and reported on overall task effort and task difficulty via an online Qualtrics survey. The EEG cap was then removed, and participants were debriefed and compensated (pre-IRB participants were not compensated).

EEG Recording

EEG was recorded continuously using a sintered Ag/AgCl 64-electrode Quick-Cap. The analog signal was amplified using Neuroscan Synamps 2 and converted to digital at a rate of 500 Hz with a bandpass of DC-100 Hz. Impedance was kept below 11 kW. EEG was initially referenced to Cz during recording and afterwards converted to an average reference off-line.

We plan to use 4–6 PCA-derived ocular components to compensate for blinks and other eye movement artifacts. Prior to averaging, both low-pass (35 Hz) and high-pass (0.12 Hz) filters will be applied to the EEG data. The continuous EEG data then will be cut into epochs separately for performance- and reward/counting-relevant feedback and then time-locked to the onset for each type (100 ms to 1000 ms). After conducting baseline correction to the 100 ms window preceding the feedback stimulus, we will reject any epochs containing excessive noise (such as eyeblinks, movements or other artifacts not representing neurological activities of interest) and averaged all remaining epochs for event-related potential (ERP) analysis. Single-subject averages will be generated at the overall and block levels.

Chapter 3: Results

Data Analytic Strategy

Statistical Tests

We utilized a variety of statistical tests to interpret our behavioral data. Across each of these analyses, we set the criterion for significance as the conventional alpha level of $p = 0.05$. Main effects or interactions with an alpha level between 0.05 and 0.1 were considered marginal but explored and reported as trends. Significant and marginal effects were further investigated by carrying out post hoc tests.

Firstly, we conducted a manipulation check to ensure we limited any potential confounding variables that may have occurred between the conditions during the initial day 1 test. We defined this as overall test performance and question difficulty because differences in number of errors to correct and/or difficulty of the questions to correct could influence error correction over and above the manipulation of the reward condition. We also predicted that if participants were following directions and attending to disk count in the counting condition, that accuracy in that condition should be better than in the reward condition. We conducted paired samples t-tests for each condition (counting vs. reward) for the measures listed above.

Next, we measured the overall retest performance and overall proportion corrected during retest by using a paired samples t-tests for each condition. Using repeated measures ANOVAs, we then took a closer look at error correction as a function of answer “eligibility” and reward (yellow disk (rewarded/counted) items, blue disk (non-rewarded/ignored) but eligible items and blue disk but ineligible items.

Past research would predict that ineligible items would be corrected less than eligible items, regardless of whether they were accompanied by a reward or not. If there is an effect of additionally receiving a reward, however, the yellow disk eligible items should be corrected more than the blue disk (but eligible) items. We also used a repeated measure ANOVA to analyze error correction as a function of high or low confidence ratings. We are not including analysis of the post-block questionnaires as they were not completed by all the participants.

Manipulation Check

First, we conducted a paired samples t-test on overall first test performance for each condition (counting & reward). The overall first test performance for the counting condition (*M* = 0.45, *SD* = 0.14) did not significantly differ from the reward condition (*M* = 0.46, *SD* = 0.07), $p = 0.92$ (see figure 2), indicating that the titration algorithm had successfully matched performance in these two conditions. Participants would have experienced equal levels of failure in each condition.

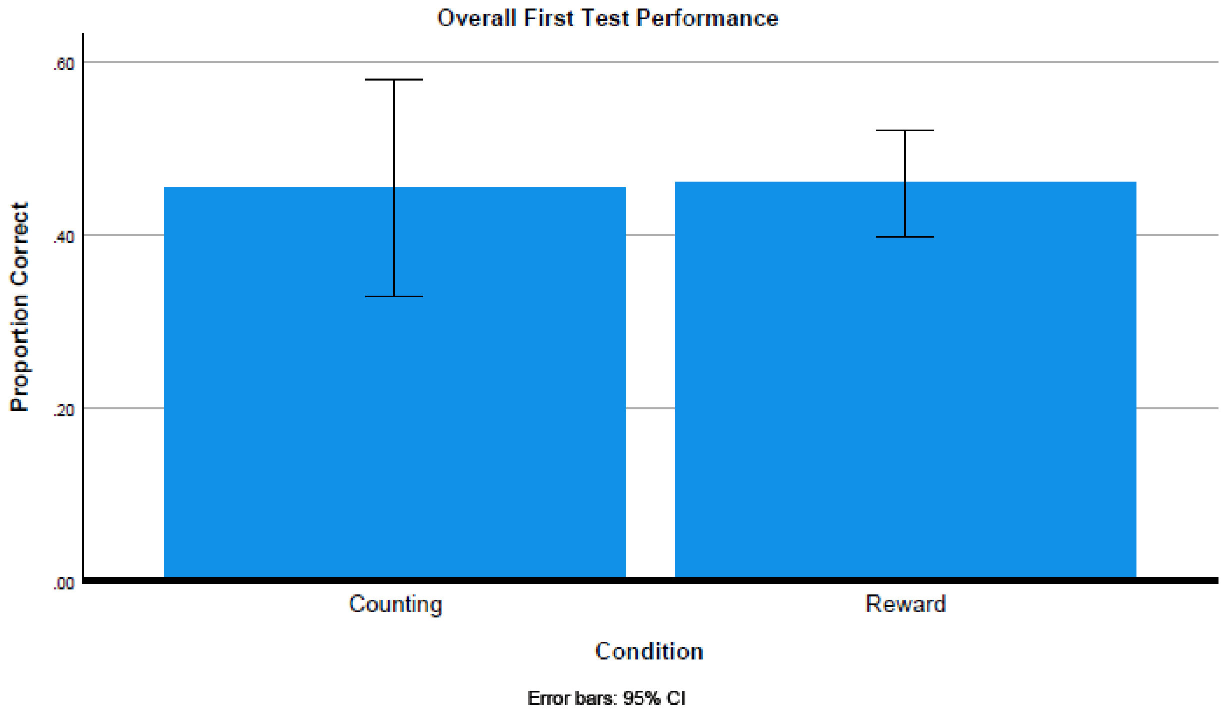

Figure 2

FIGURE 2 | The overall first test performance (mean accuracy) for the counting condition and the reward condition.

Next, we conducted a paired samples t-test first test on item difficulty (ranging from 0 to 1; based on B-KNorms) for each condition with no responses correct being 0, and all correct being 1. The first test item difficulty for the counting condition (M = 0.40, SD = 0.13) did not significantly differ from the reward condition (M = 0.42, SD = 0.17), p = 0.65 (see figure 2), indicating that the titration algorithm did not have to use more difficult questions in one of the conditions to achieve matched performance. Participants would have experienced equal levels of difficulty in both conditions.

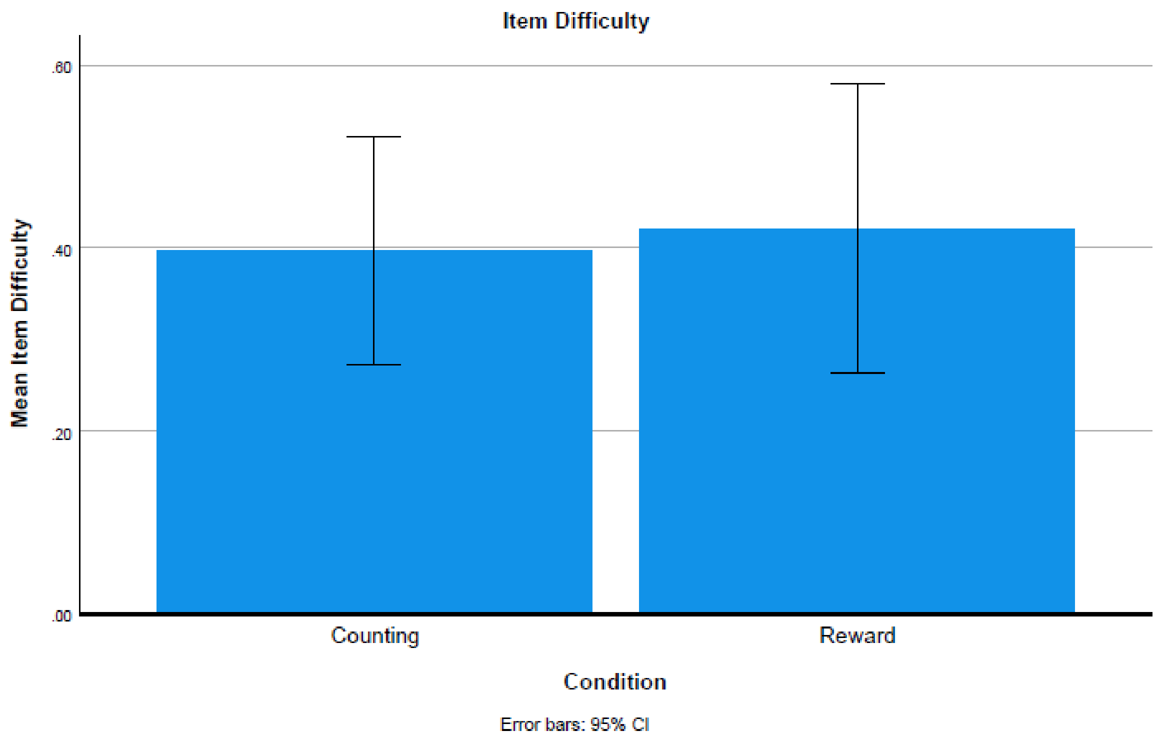

Figure 3

FIGURE 3 | First test mean difficulty by condition

Next, we ran a paired samples t-test disk count accuracy for each condition in the first test. Because the number of disks that participants were shown varied, the disk count accuracy was measured as the proportion of the subject's reported number of disks seen compared to the total number that actually appeared. A disk count accuracy of 1.0 would indicate that they were 100 % accurate in their report of how many disks they were shown, whereas a disk count accuracy of 0.0 would indicate that they were 0% accurate in their report of how many disks were shown.

The test disk count accuracy for the counting condition (M = 0.66, SD = 0.20) did significantly differ from the reward condition (M = 0.33, SD = 0.23), p = 0.03 (see figure 3), indicating that participants were paying more attention to the count of the yellow disks during the counting condition and thus, followed instructions. Participants were more accurate when asked to explicitly count the disks in the counting condition than in the reward condition where no explicit count was requested.

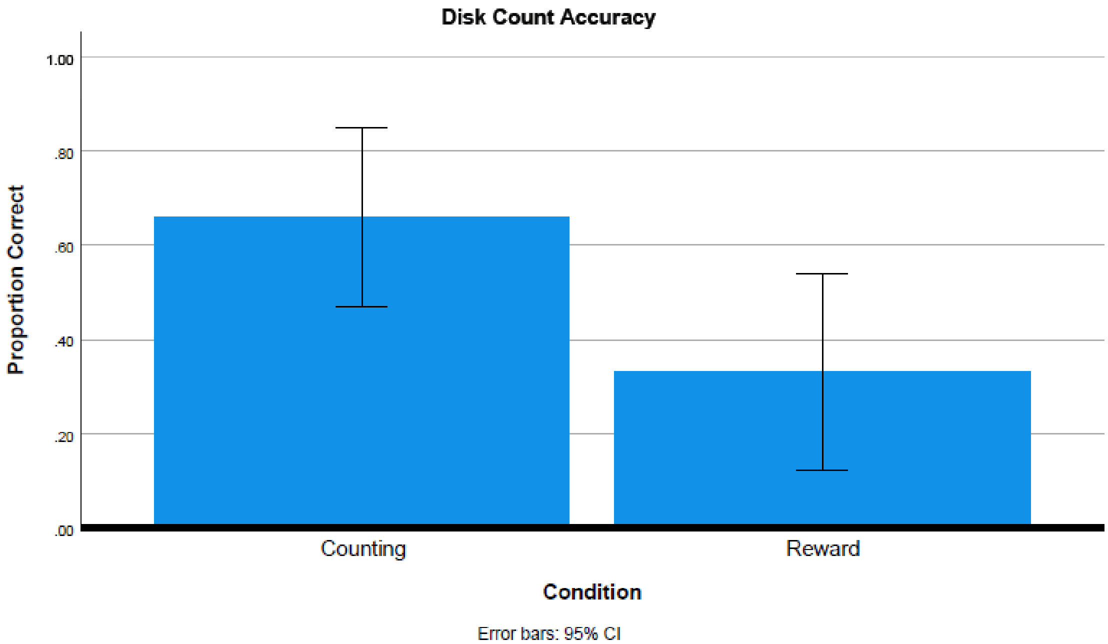

Figure 4

FIGURE 4 | First test disk count accuracy by condition

Overall Performance

Looking at overall retest performance by condition, the counting condition (M = 0.51, SD = 0.08) did not significantly differ from the reward condition (M = 0.58, SD = 0.09), p = 0.12 (see figure 4). Looking at overall error correction by condition, the counting condition (M = 0.52, SD = 0.12) did not significantly differ from the reward condition (M = 0.58, SD = 0.11), p = 0.18 (see figure 5). Although there were numerical differences favoring the reward condition, with this small sample size these differences did not reach significance.

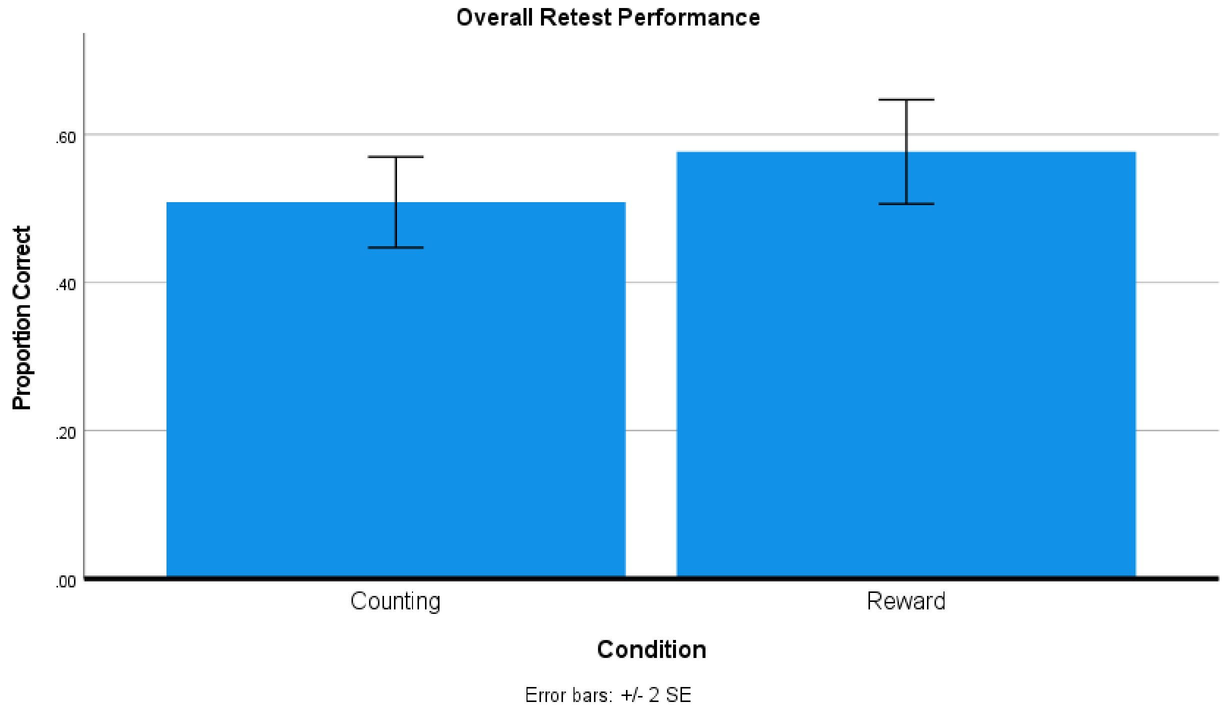

Figure 5

FIGURE 5 | Overall retest performance (accuracy) by condition

Next, we considered the proportion of errors corrected by each item type. For the reward condition, this would be defined as rewarded items, non-rewarded but eligible items and non-rewarded ineligible items. For the control condition, the corresponding types would be counted items, non-counted but eligible items, and non-counted ineligible items (yellow disks in the counting condition only appeared for semantically eligible items, as in the reward condition). We grouped rewarded and counted items together as they have the same frequency (40%) and contingency (eligible only) and are supposed to be attended to differently than the non-rewarded/non-counted items.

There was a statistically significant main effect of item type on error correction, $F(1.82, 10.93) = 9.63$, $p = 0.004$ (see figure 6). These significant results prompted a post hoc test to identify which item was causing the main effect. Post hoc tests revealed that participants were more likely to correct errors in the rewarded and eligible item type ($M = 0.65$, $SD = 0.10$) than non-rewarded and noneligible items ($M = 0.46$, $SD = 0.02$), $p = 0.007$. Moreover, participants were more likely to correct errors in the non-rewarded and eligible items ($M = 0.64$, $SD = 0.13$) than non-rewarded and noneligible items ($M = 0.46$, $SD = 0.02$), $p = 0.02$. Although here was no statistically significant main effect of condition on error correction, F(1,W 6) = 1.92, $p = 0.22$, there was a trend toward a significant interaction between item type and condition, $F(1.54, 9.2) = 3.20$, $p = 0.10$ (see figure 6). *Post hoc* tests revealed that there was a statistically significant difference in error correction between reward ($M = 0.55$, $SD = 0.17$) and counting ($M = 0.37$, $SD = 0.13$) condition for nonrewarded and noneligible items, $p = 0.01$. No other significant differences were found.

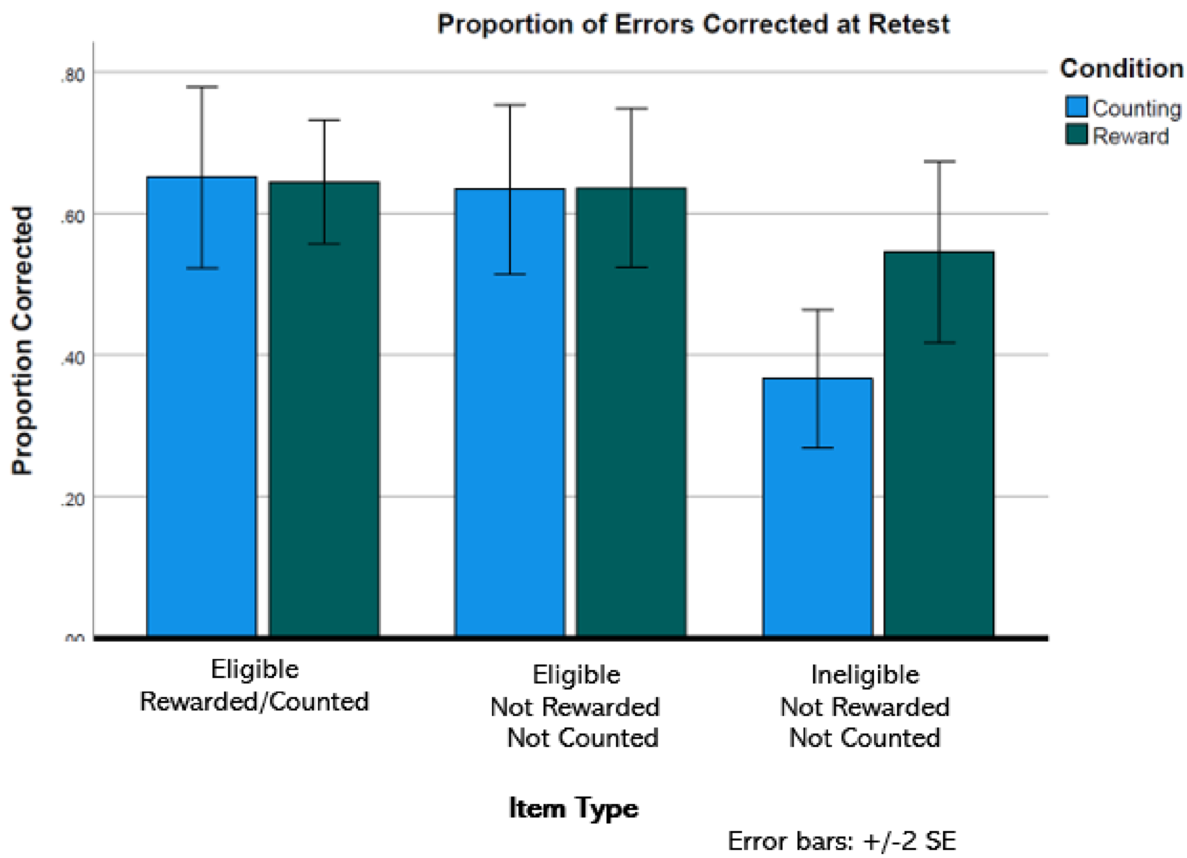

Figure 6

Figure 6 | Proportion of Errors Corrected at Retest (Error Correction of 3 item types)

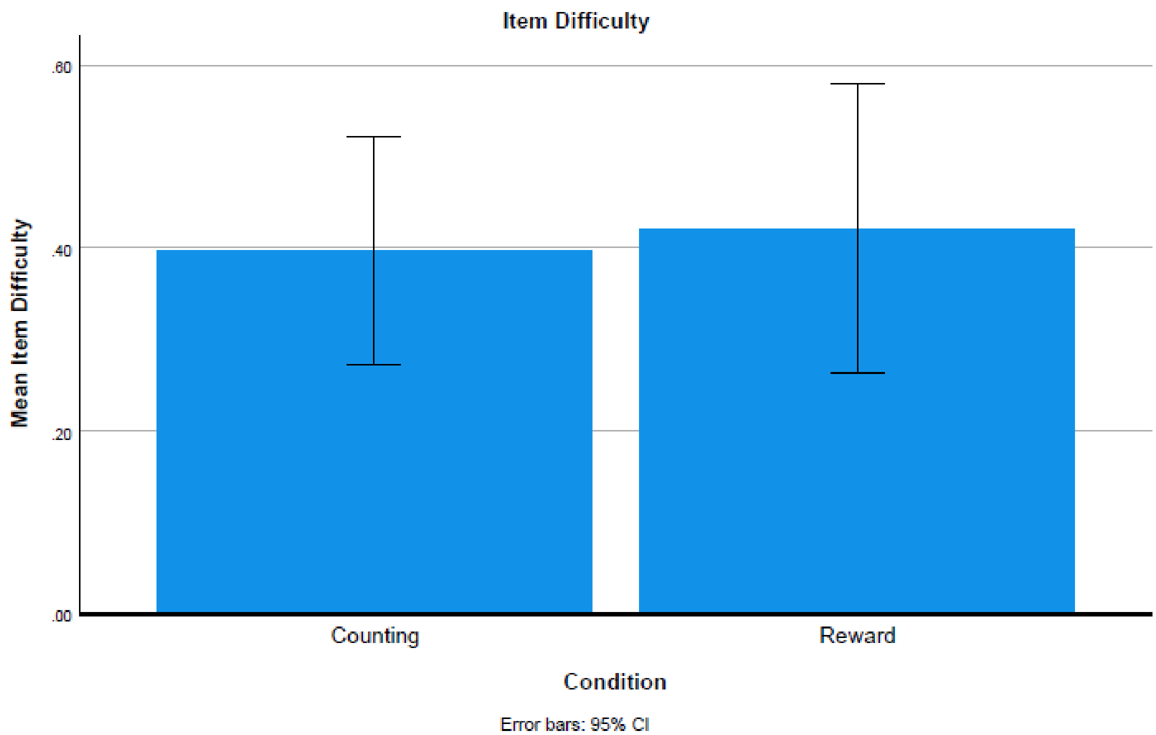

Figure 7

Figure 7 | The first test item difficulty by condition

Error Correction: High- vs. Low-Confidence Errors

When looking at the proportion of errors corrected at retest, we measured items considered high confidence vs. low confidence. Because participants may calibrate their confidence differently along the scale, we defined higher and lower confidence errors for each participant individually, based on a median split of their confidence ratings for incorrect responses. There was a statistically significant main effect of confidence on error correction, $F(1, 4) = 9.88$, $p = 0.04$. Participants corrected more high confidence items ($M = 0.63$, SD = 1.27) than low confidence items ($M = 0.51$, $SD = 0.10$). There was no statistically significant main effect of condition on error correction, $F(1, 4) = 1.10$, $p = 0.35$. There was no statistically significant interaction between confidence and condition, $F(1, 4) = 2.94$, $p = 0.16$ (see figure 7).

Post hoc tests revealed that although there was no significant difference in error correction between reward ($M = 0.65$, $SD = 0.14$) and control condition ($M = 0.61$, $SD = 0.19$) for high confidence items ($p = 0.62$), there was a marginally significant difference in error correction between reward ($M = 0.58$, $SD = 0.14$) and control condition ($M = 0.45$, $SD = 0.10$) for low confidence items, $p = 0.13$.

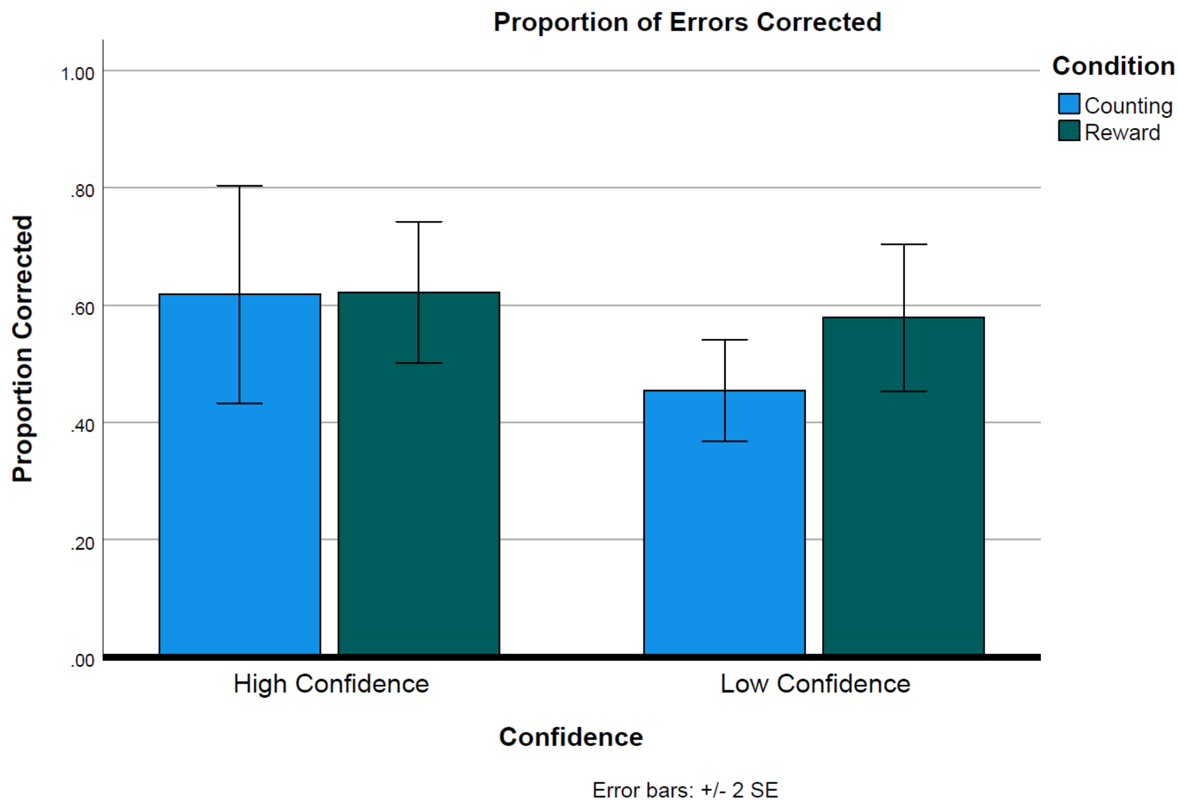

Figure 8

FIGURE 8 | Proportion of Errors Corrected at Retest for Low Confidence and High Confidence

Chapter 4: Discussion

Our statistical analyses to measure manipulation check confirmed that performance during the initial test was the same for all participants, regardless of framing by condition (control/counting vs. reward) in the overall first test performance and item difficulty level. This confirmation was essential to ensure our baseline first test conditions were as similar as possible and any variances in the post test would not be due to first test differences. We also found differences in the disk count accuracy between the reward condition and counting condition, confirming our manipulation utilizing different framing instructions worked.

We hypothesized that compared to the control framing, the reward framing would enhance encoding and consolidation of the corrective response leading to more corrected answers at the retest. These reward-enhancement effects may be specific to the individually-rewarded items or be more general to the block of reward framing as a whole.

Overall retest performance and overall proportion corrected (error correction during retest from first test) did not significantly vary by condition as predicted. However, in each analysis, the reward condition did show slightly greater proportion corrected suggesting this result may become significant in a greater sample size, similar to what was found in Abraham et al. (2019). We also plan to run an equal number of women and men and include gender as a variable to determine if effects are gender-specific as they were in Abraham et al. (2019) or not, as they were in Farber and Mangels (in prep).

There was a statistically significant main effect of item type on error correction. Participants were more likely to correct eligible items, regardless of whether they had been rewarded or not, than ineligible items. This finding supports Huelser & Metcalfe (2011) research showing the greater the degree of semantic similarity between items, the more probable it is that there will be a correction. This is consistent with the view that items which are eligible (even if not rewarded) are likely semantically closer to the correct answer than something completely different and therefore ineligible. However, rewarded items did not enjoy any particular advantage for error correction above the eligible but non-rewarded items, indicating no trial-level influence of reward.

There was no statistically significant main effect of condition on error correction and there was no statistically significant interaction between item type and condition. However, there was a trend toward an interaction which preliminary post-hoc analyses indicated was due to differences in the effect of reward on the ineligible items rather than the eligible items. Specifically, items were more likely to be corrected in the reward condition than the control condition.

Finally, we looked at the proportion of errors corrected at retest for low confidence and high confidence. There was a statistically significant main effect of confidence on error correction which supports our hypothesis of confidence levels impacting error correction. This is in line with previous research on the hypercorrection effect which refers to the finding that high-confidence errors are more likely to be corrected after feedback than low-confidence errors (Butterfield & Metcalfe, 2001). While there was no interaction with condition, there was a similar pattern as in the analysis by item type. However, inspection of the means indicate there was more of a difference between reward and control in the low than high confidence error correction, similar to what Abraham et al. (2019) found, and this difference may become significant with additional participants. Analysis of confidence and eligibility produced similar results, likely because ineligible items are more likely to be associated with low confidence due to the subject being aware their response is not semantically close and therefore have low confidence in their unrelated guess.

ERP Predictions

Based on past research we predict the ERPs to occur as follows; In both the reward and control framing conditions, negative accuracy feedback will elicit a larger feedback-related negativity (FRN) compared to positive accuracy feedback, but in the reward framing, the FRN may be smaller than in the control framing because when a reward is still possible, negative feedback is? less negatively-valanced and arousing.

In the reward framing condition, the P3 to the accuracy feedback, regardless of accuracy, will be enhanced because of overall greater engagement and task effort, compared to the control framing. Effects on the P3 are expected to be greatest for trials that are effortful/reward eligible as opposed to clearly omitted type responses (e.g., “idk” responses). We expect the P3 to be greatest for trials with high confidence/eligible items as opposed to low confidence/ineligible items.

The P2 to the reward-relevant stimuli will be bigger to the designated target stimulus vs. the non-target stimulus in both the reward and control blocks but when the target represents a discrete monetary reward the difference will be larger between the target and non-target. The behavioral and neural effects will be related to each other such that participants who show a larger difference in error correction between the reward and control conditions will also show a larger difference in ERPs between these conditions.

Unfortunately, our samples were unmatched along the dimension of gender, as women outnumbered men. We had an extremely small sample size which may be why very few comparisons reached statistical significance was found due to the limited power. Going forward we plan to continue this protocol to increase n $\geq$ 35. We expect the presence of a larger sample size will better highlight the trends we hypothesized.

Chater 5: Conclusion

Extending Abraham and Mangels’ study by utilizing EEG to observe differently implemented reward framing, the findings from the current study support the view of how extrinsic rewards that are effort-contingent, rather than performance-contingent, may benefit long-term memory for declarative information. We address the intersection of reward-based learning systems, effortful engagement and test-enhanced (i.e., feedback-based) learning of declarative knowledge. To accomplish this, we asked what extent does reward impact error correction in both males and females, and does it confer any specific advantages at the trial level for correction, or are these benefits primarily observed at the task level? Our study also researched the question of if framing of reward stimuli alters how individuals perceive and engage with negative feedback, as indicated by changes in the amplitude of the feedback-related negativity (FRN) and P3? Lastly, how are neural responses in the form of P2 and P3 event-related potentials (ERPs), which are associated with reward presentation and should also be influenced by reward manipulation, related to the successful error correction.

While our behavioral results were not as strong as expected with respect to predicted differences between conditions our findings suggest there may be considerable opportunity for future research studies to delve into the intricate connections among effortful rewards, task engagement and incidental learning.

REFERENCES

Abraham, D., McRae, K., & Mangels, J. A. (2019a). “A” for Effort: Rewarding Effortful Retrieval Attempts Improves Learning From General Knowledge Errors in Women. *Frontiers in Psychology*, *10*, 1179. https://doi.org/10.3389/fpsyg.2019.01179

Abraham, D., McRae, K., & Mangels, J. A. (2019b). “A” for Effort: Rewarding Effortful Retrieval Attempts Improves Learning From General Knowledge Errors in Women. *Frontiers in Psychology*, *10*. https://www.frontiersin.org/articles/10.3389/fpsyg.2019.01179

Bjork, R. A., Dunlosky, J., & Kornell, N. (2013). Self-regulated learning: Beliefs, techniques, and illusions. *Annual Review of Psychology*, *64*, 417–444. https://doi.org/10.1146/annurev-psych-113011-143823

Bühren, C., & Kundt, T. (2013). *Imagine Being a Nice Guy: A Note on Hypothetical vs. Incentivized Social Preferences* (SSRN Scholarly Paper 2487072). https://doi.org/10.2139/ssrn.2487072

Butterfield, B., & Mangels, J. A. (2003). Neural correlates of error detection and correction in a semantic retrieval task. *Cognitive Brain Research*, *17*(3), 793–817. https://doi.org/10.1016/S0926-6410(03)00203-9

Butterfield, B., & Metcalfe, J. (2001). Errors Committed with High Confidence Are Hypercorrected. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *27*, 1491–1494. https://doi.org/10.1037/0278-7393.27.6.1491

Cao, C., Huang, Y., Chen, A., Xu, G., & Song, J. (2021). Improvement in Attention Processing After Surgical Treatment in Functional Pituitary Adenomas: Evidence From ERP Study. *Frontiers in Neurology*, *12*. https://www.frontiersin.org/articles/10.3389/fneur.2021.656255

Davis, K. D., Winsler, A., & Middleton, M. (2006). Students’ Perceptions of Rewards for Academic Performance by Parents and Teachers: Relations With Achievement and Motivation in College. *The Journal of Genetic Psychology*, *167*(2), 211–220. https://doi.org/10.3200/GNTP.167.2.211-220 de Bruijn, E. R. A., Mars, R. B., & Hester, R. (2020). Processing of performance errors predicts memory formation: Enhanced feedback-related negativities for corrected versus repeated errors in an associative learning paradigm. *European Journal of Neuroscience, 51*(3), 881–890. https://doi.org/10.1111/ejn.14566

Deci, E. L., Koestner, R., & Ryan, R. M. (1999). A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin*, *125*(6), 627–668; discussion 692-700. https://doi.org/10.1037/0033-2909.125.6.627

Do Carmo-Blanco, N., & Allen, J. J. B. (2019). Neural correlates of cue predictiveness during intentional and incidental associative learning: A time-frequency study. *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, *143*, 80–87. https://doi.org/10.1016/j.ijpsycho.2019.06.010

Eckert, M., Scherenberg, V., & Klinke, C. (2023). How a token-based game may elicit the reward prediction error and increase engagement of students in elementary school. A pilot study. *Frontiers in Psychology*, *14*. https://www.frontiersin.org/articles/10.3389/fpsyg.2023.1077406

Epel, E. S., Crosswell, A. D., Mayer, S. E., Prather, A. A., Slavich, G. M., Puterman, E., & Mendes, W. B. (2018). More than a feeling: A unified view of stress measurement for population science. *Frontiers in Neuroendocrinology*, *49*, 146–169. https://doi.org/10.1016/j.yfrne.2018.03.001

Friedman, D., Cycowicz, Y. M., & Gaeta, H. (2001a). The novelty P3: An event-related brain potential (ERP) sign of the brain's evaluation of novelty. *Neuroscience and Biobehavioral Reviews*, *25*(4), 355–373. https://doi.org/10.1016/S0149-7634(01)00019-7

Friedman, D., Cycowicz, Y. M., & Gaeta, H. (2001b). The novelty P3: An event-related brain potential (ERP) sign of the brain's evaluation of novelty. *Neuroscience and Biobehavioral Reviews*, *25*(4), 355–373. https://doi.org/10.1016/S0149-7634(01)00019-7

Gehring, W. J., & Willoughby, A. R. (2002). The medial frontal cortex and the rapid processing of monetary gains and losses. *Science (New York, N.Y.)*, *295*(5563), 2279–2282. https://doi.org/10.1126/science.1066893

Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014a). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*, *84*(2), 486–496. https://doi.org/10.1016/j.neuron.2014.08.060

Gruber, M. J., Gelman, B. D., & Ranganath, C. (2014b). States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron*, *84*(2), 486–496. https://doi.org/10.1016/j.neuron.2014.08.060

Huelser, B. J., & Metcalfe, J. (2012). Making related errors facilitates learning, but learners do not know it. *Memory & Cognition*, *40*(4), 514–527. https://doi.org/10.3758/s13421-011-0167-z

Ibanez, A., Melloni, M., Huepe, D., Helgiu, E., Rivera-Rei, A., Canales-Johnson, A., Baker, P., & Moya, A. (2012). What event-related potentials (ERPs) bring to social neuroscience? *Social Neuroscience*, *7*(6), 632–649. https://doi.org/10.1080/17470919.2012.691078

Koivisto, M., & Revonsuo, A. (2010). Event-related brain potential correlates of visual awareness. *Neuroscience and Biobehavioral Reviews*, *34*(6), 922–934. https://doi.org/10.1016/j.neubiorev.2009.12.002

Kornell, N., Hays, M. J., & Bjork, R. A. (2009). Unsuccessful retrieval attempts enhance subsequent learning. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *35*(4), 989–998. https://doi.org/10.1037/a0015729

Krigolson, O. E. (2018). Event-related brain potentials and the study of reward processing: Methodological considerations. *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, *132*(Pt B), 175–183. https://doi.org/10.1016/j.ijpsycho.2017.11.007

Mangels, J. A., Butterfield, B., Lamb, J., Good, C., & Dweck, C. S. (2006). Why do beliefs about intelligence influence learning success? A social cognitive neuroscience model. *Social Cognitive and Affective Neuroscience*, *1*(2), 75–86. https://doi.org/10.1093/scan/nsl013

Metcalfe, J. (2017). Learning from errors. *Annual Review of Psychology*, *68*, 465–489. https://doi.org/10.1146/annurev-psych-010416-044022

Metcalfe, J., & Finn, B. (2011). People's hypercorrection of high-confidence errors: Did they know it all along? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*(2), 437–448. https://doi.org/10.1037/a0021962

Murayama, K., & Kuhbandner, C. (2011). Money enhances memory consolidation—But only for boring material. *Cognition*, *119*(1), 120–124. https://doi.org/10.1016/j.cognition.2011.01.001

Ocasio, W. (2011). Attention to attention. *Organization Science*, *22*(5), 1286–1296. https://doi.org/10.1287/orsc.1100.0602

Pashler, H., Cepeda, N. J., Wixted, J. T., & Rohrer, D. (2005). When Does Feedback Facilitate Learning of Words? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(1), 3–8. https://doi.org/10.1037/0278-7393.31.1.3

Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, *118*(10), 2128–2148. https://doi.org/10.1016/j.clinph.2007.04.019

Retrieval-based learning: A perspective for enhancing meaningful learning. (n.d.). Retrieved August 30, 2023, from https://psycnet.apa.org/record/2012-23428-006

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, *80*(1), 1–27. https://doi.org/10.1152/jn.1998.80.1.1

Smith, V. L., & Walker, J. M. (1993). Monetary Rewards and Decision Cost in Experimental Economics. *Economic Inquiry*, *31*(2), 245–261. https://doi.org/10.1111/j.1465-7295.1993.tb00881.x

Test-potentiated learning: Distinguishing between direct and indirect effects of tests. (n.d.). Retrieved August 30, 2023, from https://psycnet.apa.org/record/2012-18091-001

Tirri, K., & Kujala, T. (2016). Students' Mindsets for Learning and Their Neural Underpinnings. *Psychology*, *7*(9), Article 9. https://doi.org/10.4236/psych.2016.79125

When Does Feedback Facilitate Learning of Words? (n.d.). Retrieved August 30, 2023, from https://psycnet.apa.org/record/2004-22496-001

Whiteman, R. C., & Mangels, J. A. (2020). State and Trait Rumination Effects on Overt Attention to Reminders of Errors in a Challenging General Knowledge Retrieval Task. *Frontiers in Psychology*, *11*. https://www.frontiersin.org/articles/10.3389/fpsyg.2020.02094

Wittmann, B. C., Schott, B. H., Guderian, S., Frey, J. U., Heinze, H.-J., & Düzel, E. (2005). Reward-related FMRI activation of dopaminergic midbrain is associated with enhanced hippocampus-dependent long-term memory formation. *Neuron*, *45*(3), 459–467. https://doi.org/10.1016/j.neuron.2005.01.010