Learning Without Feedback: Detection, Quantification and Implications of Implicit Learning

by

Stephen J.C. Luehr
Bachelor of Science, University of British Columbia, 2016

A Thesis Submitted in Partial Fulfillment
of the Requirements for the Degree of

MASTER OF SCIENCE

in Interdisciplinary Studies

Supervisory Committee

Learning Without Feedback: Detection, Quantification and Implications of Implicit Learning

by

Stephen J.C. Luehr
Bachelor of Science, University of British Columbia, 2016

Supervisory Committee

Olav E. Krigolson, School of Exercise Science, Physical & Health Education
Supervisor

Adam Krawitz, Psychology
Co-Supervisor

Abstract

Mounting evidence has suggested that structures such as the anterior cingulate cortex (ACC) and other areas within the medial-frontal cortex are part of a reinforcement learning system responsible for the optimization of behaviour (Holroyd & Coles, 2002). However, we also learn without reinforcement and it has been less clear what neural structures are recruited in these instances. The P300 component of the human event-related brain potential (ERP) has been intensely researched in regards to context updating and the processing of novel stimuli (Spencer, Dien, & Donchin, 2001). Here, I sought to elaborate on the role of the P300 ERP component in implicit learning of stimulus frequencies – learning driven by the stimulus itself and not reward feedback. I propose over the course of three experiments that I have provided evidence indicating that the P300 and its neural sources play a role in feedback-free learning mechanisms. Specifically, in a feedback-free paradigm participants are shown to learn stimulus frequencies. While this occurs, P300 amplitude scales in line with participant behaviour and stimulus frequency. A common trend is revealed in how quickly this amplitude scaling occurs, suggesting further mechanisms are at play. Trial-by-trial analysis ultimately shows that behavioural prediction error formula and neural correlate prediction errors utilize a nearly identical function. These trends hold even in a passive auditory task in which the participant is fully distracted.

Supervisory Committee
Olav E. Krigolson, School of Exercise Science, Physical & Health Education
Supervisor
Adam Krawitz, Psychology
Co-Supervisor

Table of Contents

List of Tables

List of Figures

Dedication

None of this would be possible without the special people I'm surrounded by in my life. Most of all my wife, Teesha Luehr. You always encourage me to strive to be better than I was yesterday. Together we can achieve all our dreams. I truly gained my will to do this from you, my brilliant and inspiring support. Next I would like to thank all of my lab mates in the Krigolson Lab. We are a strange group at times, but whether it was bouncing around ideas or exploring at conferences we always bolstered each other. Finally, I want to thank my parents. You've fostered curiosity, provided a home full of love and pride, and raised a man who can revel in learning even insignificant factoids.

Chapter One: Introduction and Review

1.1 Overview

Feedback, in whatever form it takes, has been shown to be crucial for human learning (Chansky, 1960; Gill & Martens, 1975). We often instinctively think of feedback in explicit terms. For example, "You burnt my toast!" instructs the other person of the outcome of their action - leaving the toast in too long. However, in order to give such feedback, the person eating the toast has already received feedback in the form of sensory information. When they have eaten the toast the see the burnt edges and taste the charcoaled finish of the bread.

Reinforcement learning theory posits that we learn via prediction errors – discrepancies between outcomes and expectations that drive changes in behaviour (Glimcher, 2011). If we again consider the toast example the person eating the toast has an expectation of the taste and texture before they eat the toast, once they take the first bite there is an outcome, and if the expectation differs from the outcome a prediction error occurs – the toast was left in the toaster to long, next time take it out sooner. By utilizing predictions in this manner, we as humans can shape our behaviour and optimize our decision-making processes (Cohen & Ranganath, 2007; Glimcher, 2011; O'Doherty, Cockburn, & Pauli, 2017; Trepel, Fox, & Poldrack, 2005).

Here, I intend to investigate whether these prediction errors take place without explicit feedback and potential methods of assessing these prediction errors utilizing neural correlates. First, I will review the literature of how we might learn at the most basic on a neural level. Next, I will elaborate on implicit learning, the specific type of learning that I wish the investigate in relation to prediction errors and its importance to human learning overall. Finally, I will review some recent methods of assessing learning before outlining in three experiments how I intend to investigate these mechanisms.

1.2 Learning and the Brain

Seminal work done by Schultz and colleagues (Mirenowicz & Schultz, 1996; Schultz, Dayan, & Montague, 1997; Schultz, 1998; Schultz, Apicella, & Ljungberg, 1993) found that the firing rates of midbrain dopamine system neurons in monkeys aligned with some important theoretical predictions of dopamine neuron responses to motivationally significant stimuli. Primarily that as the monkeys learned the task at hand, the dopamine neuron firing rates began to fire to the cues that predicted rewards rather than the reward itself. This was tested by placing monkeys in an environment where they would be made to predict appetitive events. First, fruit juice would be given to the monkey without warning or requirement, producing a positive-error of predicting reward. This lead to a large burst in dopamine firing rate after receiving the reward. The monkey was then conditioned by having a visual or auditory cue followed by the same reward. After conditioning, the dopamine neurons would fire to this cue, and fail to fire to the reward itself. By now providing a cue and *no* reward, it was observed that the dopamine response remained for the cue and produced depression for the expected reward. Schultz and colleagues proposed that throughout the experiment, the monkey was making predictions of the outcome and adjusting expectations on a trial-by-trial basis based on the error of that prediction.

These findings have since then been further reflected in a large number of human brain imaging studies utilizing both electroencephalography (EEG) and (fMRI) functional-magnetic resonance imaging (Bray & O'Doherty, 2007; Brown & Braver, 2005; Cohen & Ranganath, 2007; Hajcak, Moser, Holroyd, & Simons, 2007; Haruno & Kawato, 2006; Hassall, MacLean, & Krigolson, 2014; Holroyd & Krigolson, 2007; Jessup, Busemeyer, & Brown, 2010; Krigolson, Hassall, & Handy, 2014; Morris, Heerey, Gold, & Holroyd, 2008; Nieuwenhuis, Heslenfeld, et al., 2005; O'Doherty et al., 2004; Tanaka et al., 2004). Of note here is Holroyd & Coles (2002)

EEG study on human dopamine reward systems. They have posited that negative reinforcement learning signals are conveyed by the mesencephalic dopamine system to the anterior cingulate cortex to modify participant behaviour on the task being learned. Further work (Holroyd & Krigolson, 2007) elaborated on this detail explicitly: the neural response is largest for the difference between unexpected rewards and punishments as compared to the expected rewards and punishments.

The feedback related negativity (FRN) (Miltner, Braun, & Coles, 1997) was the primary brain activity being observed in this study. It had been proposed (Holroyd & Coles, 2002; Nieuwenhuis, Holroyd, Mol, & Coles, 2004) that the FRN indexes reward prediction error signals. The amplitude of the FRN should then be affected by the unexpectedness of feedback stimuli, producing larger differences between unexpected errors and correct feedback as opposed to expected errors and correct feedback. The study discussed here (Holroyd & Krigolson, 2007) tested this using a time estimation task with varying levels of difficulty. It was ultimately shown that FRN amplitude was much more negative for the unexpected forms of feedback (mistakes in the easy condition, correct answers in the hard condition) as opposed to expected feedback (unexpected: -10.8μV, expected -6.7μV). In addition, behavioural adjustments were larger following these unexpected pieces of feedback. The FRN amplitude also correlated well both across participants and with performance. All facets of this study combined provide strong evidence that the FRN is reflecting a prediction error signal.

These findings have a similar conclusion as Schultz's work on monkey dopamine systems. Prediction errors play a role in sculpting behaviour through learning systems, specifically utilizing midbrain dopamine systems (Fiorillo, Tobler, & Schultz, 2003; Holroyd & Coles, 2002; Holroyd & Krigolson, 2007; Schultz, 1998).

1.3 Prediction Errors in Implicit Learning

A key question remains arises from these studies however: is a prediction error computation utilized when we learn without explicit feedback? Both Schultz and Holroyd's experiments discussed involve providing the participant with explicit feedback (i.e., a reward such as a treat, money, or check mark). The participant then is able to make the adjustments based on feedback information of each trial. One would expect that predictions would steadily change in such a scenario. There is however the case of error-related negativity (ERN) response tasks. These tasks elicit ERN responses to errors as they are being made. The process thus starts with the *response* to a stimulus, not the interpretation of the stimulus itself. When discussing prediction errors and expectancies, we are more interested in the response to the stimulus itself, or specifically it's violation of our expectancies in appearing in the first place. Explicit reinforcement learning is distinct from how we learn in many situations in the real world such as an infant babbling their first words, or a long jumper repeating jumps until attaining their goal. A babbling infant rarely receives a direct input for each word spoken. Instead, they will continue to babble and constantly adjust the noises they are making to produce new, novel combinations. As these combinations approach real words, the infant may hear this or begin to be reinforced more by the parents. This still leaves a large portion of exploration between "*mah*" and "*Dada*" with only their own experience as feedback. Our long jumper is learning in much the same process: incrementally through trial and error. Not every jump will have a coach reviewing footage and tweaking small features of stance, gait, and other techniques. Rather this jumper will reattempt many times with small adjustments happening on a more intuitive level. As they learn which small adjustments are improving their distance, they accumulate a more and more optimized set of movements towards their goal. Their only feedback has been their final jump distance.

Implicit learning as a whole has long been investigated from a variety of fields and applications such as artificial grammar learning (Cleeremans, Destrebecqz, & Boyer, 1998), sequence learning (Dienes & Berry, 1997), language acquisition (Michas & Berry, 1994), and even what will be used here: probability learning (Yellott, 1969).

Yellott's study made some foundational insights using only reaction time and accuracy measures in a speeded choice-making task. Much like the common oddball paradigm (Donchin, 1981; Picton, 1992; Polich, 1989), only two stimuli were presented to participants. On each trial, participants were to predict whether "*X*" or "*Y*" would appear on the display screen. After a "ready" signal a guess would be inputted, and the letter would immediately appear. This was repeated 450 times for each participant. Over these trials there would be six different "reinforcement schedules" which varied in how often the stimulus appearing after a guess would be "*Y*", ranging from 50% to 100% probability. The probabilities were selected to observe both noncontingent success (NCS) and noncontingent event (NCE) reinforcement schedules. An NCS schedule would mean that the participants' odds of correctly guessing were the same regardless of what response they had given. An NCE schedule means that the event probabilities are not dependant on the participants action. For example, using an 80% probability condition, in an NCS schedule if the participant responds with "*X*" as their prediction, their chance of success (i.e., the next stimulus being an "*X*") is precisely 80%. Meanwhile in an NCE schedule, the odds of the next stimulus being *X* would not be 80%, but rather *Y* is predetermined to show up precisely 80% of the time. This makes NCE schedules learnable over time, while NCS schedules cannot be improved upon.

The core findings of this experiment were that on the fastest reaction time trials, it seemed that participants used some proportion of simple guesses. Meanwhile longer reaction

time trials were using a more complete process to analyse for a learnable pattern. When the proportion of successes was fixed to 100% “superstitious solutions” (Yellott, 1969) became more prevalent. It was noted that participants fixate on more or less complex, idiosyncratic response patterns much as they do in operant conditioning (Skinner, 1948).

Most importantly, predictions made by Yellot for the 80% success condition aligned with participant behaviour. Specifically, it was found that participants behaviour in an NCS environment matched those in an NCE environment, and better suited a *N* element model that would suggest a common learning mechanism to all prediction experiments with contingent schedules. Yellot discusses little of prediction error mechanisms at this point, only alluding that prediction errors may be disruptive on memory to explain why his participants appeared unable to detect longer run patterns in stimuli.

While a small selection of studies have looked into the effect of outright restricting explicit feedback (Erhel & Jamet, 2013; Ishikura, 2008), there is a dearth of studies that seek to investigate whether prediction error mechanisms are still carried out on trials without feedback.

Ishikura (2008) for example composed of a golf-putting task with restricted knowledge of results (KR). Over each participants’ 60 trials, they were shown the path and results of their putt either 100% of the time or 33% of the time. The 33% condition were shown the results of every third putt that was attempted. The participants were then retested both 10 minutes and 24 hours after the experiment. The 33% KR group was found to have a lesser constant error than the 100% KR group during the retest phase. Ishikura goes on to suggest that reduced relative frequency of feedback aids in the learning process of an accurate golf putt. This is suggested to be more a result of the guidance hypothesis (Salmoni, Schmidt, & Walter, 1984), where the 100% KR group was relying on external feedback whereas the 33% KR group could not see the path of

their putt the majority of the time. This meant that the 100% KR group would be hindered in their memory retention according to this hypothesis. While quite interesting, the authors do not delve further into the specific mechanisms behind why this may be occurring, or which mechanisms take place during feedback-free trials.

1.4 The Relevance of Implicit Learning

It is important to interject at this point with an operationalization of what is meant by implicit learning. In this case I am using the description most supported by literature which requires that the task at hand be learnable, in order for a prediction error to accrue to a correct choice being made later on (Bray & O'Doherty, 2007; Glimcher, 2011; Holroyd, Krigolson, Baker, Lee, & Gibson, 2009; Sambrook & Goslin, 2015). As such, a task of complete random chance will invoke prediction errors of some degree, but they are not inherently a guarantee of implicit learning. For example, a fair coin flip can have a predicted outcome of "heads". When the coin lands on tails, this is indeed a violation of a prediction and therefore could be considered a prediction error. However, because the coin flip is 50/50% odds on every subsequent flip, this prediction error correction will not lead the observer any closer to learning a more likely outcome. This means that while the outcome distribution can be learned, an optimal decision-making process cannot be refined. Here, only tasks in which an optimal response pattern can be learned are of interest. Prediction errors themselves may be spontaneous and meaningless in a vacuum, there must be some aspect of non-random information to glean knowledge of. As a result of this definition, we are tied to the midbrain dopamine system as playing at least some role in what we observe, as its importance in learning cannot be overstated (Bromberg-Martin, Matsumoto, & Hikosaka, 2010; Cohen & Ranganath, 2007; Holroyd & Coles, 2002; Montague,

Dayan, & Sejnowski, 1996; Schultz, 1998). As established, midbrain dopamine systems have been found to be a central part of regular reinforcement learning processes. As I will discuss shortly, this systems involvement in novelty processing may also make it key to implicit learning mechanisms as well (Nieuwenhuis, Aston-Jones, & Cohen, 2005).

The work of William Estes (Estes, 1950, 1974) into understanding implicit learning processes behind stimulus probabilities (Atkinson & Estes, 1962) highlights why learnability is a requirement and the potential impact of research into these mechanisms. This early work brought into light the concept that although we may not learn in single instances, we can learn through a cumulative and gradual movement to optimal behaviour. In addition, his work highlights that variability is inherent in the process and environment, and so not every instance of learning will be as valuable as the next.

Estes work was primarily in developing a mathematical model of learning much like Yellott (1969) would later spend a considerable amount of time expanding upon. Estes noted that in T-maze experiments utilizing rats and food rewards, rats would rarely choose a direction with chance. Almost immediately they would begin progressively favouring one direction over another, the one with the food reward. He reasoned that the rats must be learning on a trial-by-trial basis the optimal path to travel. He concluded that our responses were summations of all stimuli around us. The rats decided a direction according to a formula by his theory:

$$Probability\ of\ R_k = x/S$$

Equation 1. William Estes Stimulus Sampling Theory

Where the response depended on *S*, the number of stimuli in a situation, compared to *x*, the number of stimuli conditioned for the response. This definition leaned the field of learning away from staunchly defined pathways, and towards a more probabilistic approach (Estes, 1974).

Random chance permeated this model by way of his definition of a stimulus. Estes considered all subtle changes in environment to be a part of the "noise" that could influence a response outcome (Atkinson & Estes, 1962; Estes, 1950, 1974). This noise was why an individual trial may not be useful for learning, as variability from moment to moment would affect both the learner and the response being made. The final cornerstone of his theory was that all of the stimuli leading to this response would become at least partially conditioned to the response as well (Estes, 1950).

These early studies of implicit learning laid a groundwork for how research developed into a variety of related topics including memory (Artigas & Prados, 2017; Inman & Pearce, 2018; Murphy, Byrom, & Msetfi, 2017; Pacchiarini, Fox, & Honey, 2017; Rodríguez, Blair, & Hall, 2008). Pacchiarini, Fox, & Honey (2017) even utilized the basis of this theory to begin investigating perceptual learning with tactile stimuli in rodents. The principal of pre-exposing the rodents to two tactile stimuli facilitates later discrimination between them is extremely simple yet had not been investigated until recently. The formula used to analyse the rodents' behaviour uses Estes work as a basis.

The operational definition stated previously for implicit learning may leave some to be desired. This is especially true when considering that "implicit" and "explicit" have differing definitions depending on the field framework used to study them. For example, learning and memory researchers will often distinguish these systems based on conscious awareness or attention involved in response (Squire & Zola, 1996). Alternatively, these terms may be couched within reinforcement learning terminology as two separate pathways rather than simply differing levels of awareness (Barch et al., 2017). To clarify this, we must define the feedback, rewards, and stimuli being studied. In this case, I will be investigating learning where the reward is not

explicit, and also provides intrinsic as opposed to extrinsic reward. That is to say, feedback is not provided in the form of rewards but in the form of confirmatory evidence. As an environment is learned, the participant has intrinsic reward to confirming their beliefs and improving performance based on their own observations (Singh, Lewis, Barto, & Sorg, 2010). Here-after, "implicit learning" refers to this specific case of implicit, intrinsically rewarding reinforcement learning as opposed to a more traditional explicit and extrinsic reinforcement learning environment.

With a bevy of work having been done on both reinforcement and implicit learning and their processes such as prediction errors, it is only sensible to now begin to probe for similarities in structures or processes between the two. We know the use of ERN from previous studies used to assess reinforcement learning (Hajcak et al., 2007; Holroyd & Krigolson, 2007; Morris et al., 2008), and so I propose to also use neural correlates to explore implicit learning mechanisms.

1.5 Assessing Implicit Learning

The P300 has been heavily investigated (Donchin, 1981; Duncan-Johnson, 1981; Käthner, Wriessnegger, Müller-Putz, Kübler, & Halder, 2014; Patel & Azzam, 2005; Picton, 1992; Polich, 1989, 2007; Qiu, Tang, Chan, Sun, & He, 2014; Turetsky et al., 2015; Uetake & Murata, 2000; Ullsperger & Baldeweg, 1991) since its first discovery in 1965 (Sutton, Braren, Zubin, & John, 1965). Despite the constantly growing literature around this event-related potential (ERP) component, it's role in implicit learning mechanisms has been relatively underexposed.

1.5.1 Utilizing Electroencephalography

EEG Signals Defined

As in previous notable work (Holroyd & Coles, 2002) I utilized electroencephalography (EEG) in combination with behavioural data measures to assess learning processes. EEG

provides considerable temporal resolution as to ongoing brain activity (Luck, 2014), which is a benefit to the paradigms I have used that rely on immediate sequential stimulus response. The scalp level electrical data can also be provided relatively non-intrusively, allowing participants to focus on the learning task at hand while being unrestrained. Combining sub-millisecond temporal resolution, ease of application and affordability, EEG provides many benefits over more cumbersome and expensive methods such as MRI at the expensive of source accuracy (DellaBadia Jr, Bell, Keyes Jr, Mathews, & Glazier, 2002; Dien, Spencer, & Donchin, 2003; Luck, 2014; Song et al., 2015; Urbach & Kutas, 2002).

EEG signals are produced directly from neural activity, in contrast to the blood-oxygen level-dependant (BOLD) method of fMRI. To begin, the neural signal of interest is generated as a result of neurotransmitters binding to receptors and producing postsynaptic potentials (PSPs) (Buzsáki, Anastassiou, & Koch, 2012; Jackson & Bolger, 2014). This small electrical dipole is thus the direct result of neural activity, whereas MRI must rely on factors such as blood flow which can be affected by multiple factors other than neuronal activity (Desjardins, Kiehl, & Liddle, 2001; Kalisch, Elbel, Gössl, Czisch, & Auer, 2001; Van Dijk, Sabuncu, & Buckner, 2012). Of course, the small size of this dipole means that numerous dipoles must be clustered together in similar orientations to be measured at the scalp even with amplifiers. This results in most neural activity measured being the generated by pyramidal cells of the cerebral cortex, which are typically oriented perpendicular to the cortical surface (Buzsáki et al., 2012; Coenen, 1995; Contreras & Steriade, 1995). After these clusters of pyramidal cells propagate electric current, the charge is conducted through the various layers of fluid, bone, and tissue to the scalp. This volume conduction action means that by the time neural signals are present at the scalp where we may measure them with EEG they are often much more broadly distributed, and may

have even summated or been negated by other neural signals before reaching the surface (Buzsáki et al., 2012; Jackson & Bolger, 2014; Spencer et al., 2001; Urbach & Kutas, 2002). With sufficient electrode arrays, a reasonable accuracy in source localization can be proposed or even supported by functional magnetic resonance imaging (Dien et al., 2003; Mosher & Leahy, 1998; Qin et al., 2003) and mathematical modelling (Delorme & Makeig, 2004; Mosher & Leahy, 1998; Song et al., 2015).

While these advances have been achieved, there is still the issue of variance that comes with measuring summed activity from such a relative distance. To combat this, EEG researchers adopted the event-related potential (ERP) technique first used to study cognitive components by Walter, Cooper, Aldridge, McCallum, & Winter (1964).

Event-Related Potentials

ERP experiments follow a simple premise based on averaging numerous trials of EEG data time locked to a stimulus presentation or participant response. Averaging many of these events together produces a waveform representative of the brain activity resulting from the target event, and many trials will improve the signal to noise of this final waveform. In these waveforms, strong characteristic patterns will arise which are often named after their deflection, order, latency and/or processes (i.e., P1, N1, P200, N200, P300, RewP) which correlate to neuronal aggregates (Kok, 1997).

These waveforms are quite useful for the ability to examine neural activity within a specific task, with high time resolution due to the aforementioned association of EEG signals to direct neural activity (Buzsáki et al., 2012). Experiments will often focus on specific, repeatable characteristics of the waveform responses labelled as “components” (Coenen, 1995; Foti, Hajcak, & Dien, 2009; Sambrook & Goslin, 2015). With many trials averaged across conditions,

an ERP component can reveal small within-subject changes in well controlled experimental designs (Luck, 2014). Examples of ERP components include the aforementioned FRN (Nieuwenhuis et al., 2004), the N200 component (Holroyd, Pakzad-Vaezi, & Krigolson, 2008; Patel & Azzam, 2005), and the P300 component (Donchin, 1981; Polich & Kok, 1995; Soltani & Knight, 2000).

1.5.2 The P300 Event-Related Potential Component

P300 Generation and Detection

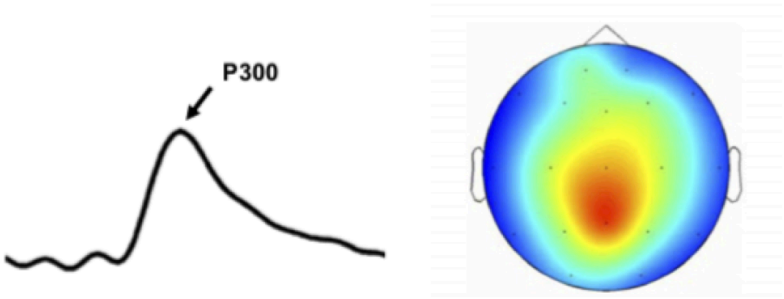

Figure 1. The P300 Component. (Left) The P300 average waveform produced in response to many averaged stimuli. (Right) Visual topography of the P300 response to target stimuli, generally centered around electrode Pz. Figure adapted from Polich, 2007.

The P300 (*Figure 1*) is a large positive deflection that occurs around 300ms after stimulus onset. Its amplitude is quantified as the voltage difference between a pre-stimulus baseline and the mean peak (+/-25ms) of the components peak latency window (Polich, 2007). The initial discovery of this component is tied to Sutton (1965) and his colleagues work involving presenting participants with stimuli on an uncertain or certain basis. They of course found that this large positive deflection around 300ms after stimulus onset was of much greater amplitude for unexpected stimuli. The production of this ERP component was achieved using uncertain stimuli: either flashes of light or auditory clicks. Participants would be cued with a

stimulus which was always followed by a sound or a light, or they could be cued such that they could not predict whether it would be a light or sound following 3 to 5 seconds later. The uncertain condition presented either 66% light and 33% clicks or the inverse pairing. The more uncertain stimuli – or rather lower frequency stimuli produced a greater amplitude P300 response. This early work required an extremely modest setup: only a single electrode placed 1/3 of the way down from the vertex of the scalp and the auditory meatus (ear), and two reference electrodes clipped to the earlobes.

The P300 was studied ahead of many other components due to its ease of elicitation and detection. The most common task for eliciting the component is the "oddball" paradigm, a signal-detection task first used by Ritter, Vaughan, & Costa (1968). This is quite similar to the paradigm described previously (Sutton et al., 1965) however there are no cueing stimuli and typically no certain versus uncertain conditions. Participants are presented with a series of two stimuli in which the frequency of one differs from the other. The less frequent "oddball" is considered a target stimuli which the participant must attend to either by pressing a button or mentally counting its occurrences (Polich, 2004; W. Ritter & Vaughan, 1969). The stimuli can be either auditory as in Sutton or Ritter's work, or visually based such as in later works (Picton, 1992; N. K. Squires, Squires, & Hillyard, 1975). The P300 observed is typically on the difference wave of this oddball stimuli average minus the frequent control stimuli average.

The Many Roles of the P300

It is worth noting that the P300 is often broken down into two subcomponents; the P3a and the P3b. The P3a is typically slightly earlier, has a more frontal topography and is more associated with stimulus novelty itself (Patel & Azzam, 2005; Snyder & Hillyard, 1976; N. K. Squires et al., 1975). It was labelled as the P3a to distinguish itself from the P3b which is

associated with target stimulus processing (Polich, 2007; Snyder & Hillyard, 1976; N. K. Squires et al., 1975). Further investigation by Polich (Conroy & Polich, 2007) utilizing distracter stimuli in an oddball paradigm to elicit both P3a and P3b components found that these two components did indeed differ in both scalp topography and latency.

The P3b meanwhile has been tied to context updating, being produced in scenarios in which the mental representation of incoming stimuli needs to be changed (Donchin, 1981). If there is no need to update stimulus representations the P3b component is not evoked, and sensory ERPs are produced instead (Polich, 1989, 2007) as detailed in *Figure 2* (Polich, 2003). Despite the distinction between P3a and P3b, both components are often produced in standard oddball paradigms by both the target and non-target stimuli (Spencer et al., 2001) as they are both producing either context updating response or a novelty response on each given trial.

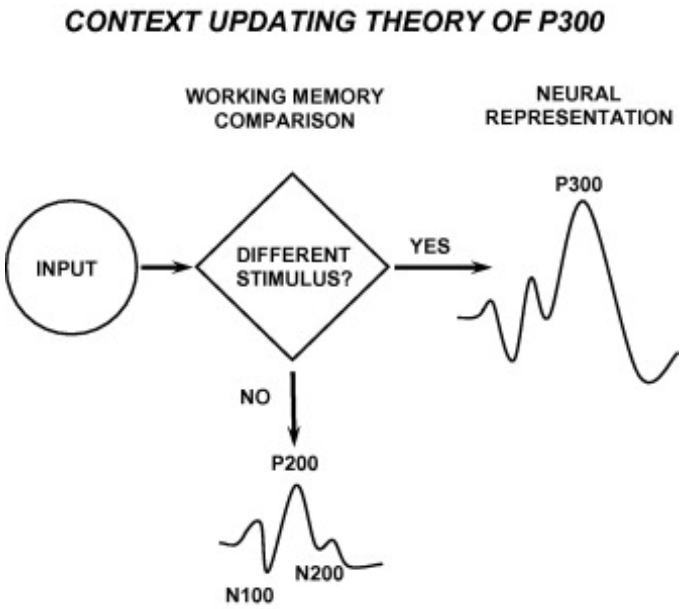

Figure 2. Context updating theory of P300. The context updating model of P300 suggests that stimuli requiring memory or processing are compared to previous representations of that stimuli. If it is the same, sensory ERPs are produced without the P300 as there is no context updating. An

unexpected or novel stimulus such as in an oddball paradigm produces a P3b response. Both waveforms presented here are averaged from many trials. Figure from Polich, 2003.

Manipulation of stimulus probability has also lead to the P300 being tied to resource allocation. Difficult discrimination of a target from a non-target stimulus leads to a larger amplitude P300, and this amplitude increases as the probability of the target decreases (Duncan-Johnson & Donchin, 1977; Duncan-Johnson, 1981; Snyder & Hillyard, 1976). Attentional resource allegation is further supported by the behaviour of the P300 in dual-task paradigms. In these paradigms the participant must complete a task that requires a varying level of cognitive load, while also mentally tracking target oddball stimuli. P300 amplitude in response to oddball stimuli is found to increase as the cognitive load task is made easier and vice versa (Foerde, Knowlton, & Poldrack, 2006; Isreal, Chesney, Wickens, & Donchin, 1980; Kok, 1997; Polich & Kok, 1995; Wickens, Kramer, Vanasse, & Donchin, 1983).

While it may be tempting to focus on stimulus probability alone as a manipulator for P300 amplitude, stimulus probability does not always affect P300 amplitude. The timing between stimuli in a typical oddball task has been shown to greatly affect the P300 (Croft, Gonsalvez, Gabriel, & Barry, 2003; Gonsalvez & Polich, 2002). These findings showed that the greater the number of non-target stimuli between the two target stimuli, the larger the P300 amplitude. Short intervals thus produced smaller P300 components. Importantly, long intervals of 6-8s between targets removed probability effects on the P300 amplitude (Polich, 1990; Polich & Bondurant, 1997).This emphasizes that stimulus sequence is a key part of the effect of stimulus frequency on P300 amplitude proposed by Squires (Squires, Petuchowski, Wickens, & Donchin, 1977).

The P300 component has also been used as an indicator of memory processing. Based on the premise of P300 involvement in attentional allocation and context updating, researchers assessed P300 amplitude correlation with recall performance (Foerde et al., 2006; Karis, Fabiani, & Donchin, 1984; Polich & Kok, 1995; Spencer, Vila Abad, & Donchin, 2000). In Karis' work (1984) words were presented sequentially with occasional words being presented in smaller or larger font size to facilitate memory. Those distinct words which could later be recalled elicited larger P300 components during their initial presentation than those that could not be recalled. Later experiments instructed participants to use rote rehearsal or elaborative memory strategies, with elaborative strategies removing the relationship between P300 amplitude and recall (Fabiani, Karis, & Donchin, 1986, 1990). From these results it follows that tasks involving memory processing affect P300 amplitude, and indeed delayed retrieval tasks elicit larger P300 amplitudes supporting memory engagement (Azizian & Polich, 2007; Donchin, 1981).

It may be clear by now that over the course of the past half-decade, a variety of explanations have arisen for the role of the P300 in underlying cognitive processes. These numerous explanations include but are not limited to: attention and memory (Kok, 1997; O'Doherty et al., 2017; Polich, 1989; Rushby, Barry, & Doherty, 2005), probability effects (Mecklinger & Ullsperger, 1995; Nieuwenhuis, Aston-Jones, et al., 2005; Ullsperger & Baldeweg, 1991), context-updating (Donchin, 1981), attentional resource allocation (Isreal et al., 1980; Wickens et al., 1983), fatigue (Aidman, Chadunow, Johnson, & Reece, 2015; Lal & Craig, 2005; Zhao, Zhao, Liu, & Zheng, 2012), memory processing through habituation and dishabituation (Kok, 1997; Polich, 1989; Rushby et al., 2005), and even norepinephrine activity in the locus coeruleus (Nieuwenhuis, Aston-Jones, et al., 2005).

Neural Basis of P300

The neural sources and neurotransmitters responsible for P300 production have remained unclear, however numerous theories have been put forward to shed light on this issue (Knight, Grabowecky, & Scabini, 1995; Nieuwenhuis, Aston-Jones, et al., 2005). The P3a subcomponent has been shown to decrease in amplitude in response to non-target stimuli in patients with focal hippocampal lesions, while the P3b was unaffected for targets (Knight et al., 1995). The temporal-parietal junction has been shown to severely reduce P300 amplitude when compromised, especially over the parietal area (Verleger, Heide, Butt, & Kömpf, 1994). With P3a being affected by frontal lesions and P3b being affected by the temporal-parietal junction, there is an implication that these two components are part of a circuit pathway between frontal and parietal areas of the brain (Polich, 2003, 2007; Soltani & Knight, 2000).

In the frontal lobe, P3a activity in regards to working memory and attentional resource allocation have been suggested to be the result of dopaminergic activity while the P3b in the parietal area is heavily encouraged by dense norepinephrine inputs found in this region of the brain (Braver & Cohen, 2000; Nieuwenhuis, Aston-Jones, et al., 2005). The work of Nieuwenhuis et al. (2005) highlighted the role of the locus coeruleus-norepinephrine (LC-NE) system in P3b production for target detection tasks. This work was consistent with views of P300 characteristics in terms of scalp distribution and component latency and amplitude (Aston-Jones & Cohen, 2005).

Importantly, the conclusions found by Nieuwenhuis (2005) showed the empirical findings that the LC-NE phasic response is a result of internal decision-making processes, and key to information processing as a whole. This phasic response is theorized to improve processing in areas engaged with the task at hand, and the P3 is presumed to reflect this action.

This may explain the differing topography of P3a and P3b, as pre-frontal structures contribute to novelty processing and experience LC-NE engagement.

Each of these hypotheses implicate the P300 in a variety of cognitive processes that still has not converged in over 50 years. It seems more likely at this point that the P300 is a multifaceted component, being an indicator of many processes culminating at the scalp given the plethora of roles it has been found to play, and the factors that may influence it (Duncan-Johnson & Donchin, 1982; Gonsalvez & Polich, 2002; Patel & Azzam, 2005; Picton, 1992; Polich, 2004; Polich & Kok, 1995; Wu & Zhou, 2009). While the P300 might vary due to many conditions such as arousal, acute fatigue, exercise and more (Polich & Kok, 1995), it is still reproduced with remarkable consistency between recordings or differing experiments (Williams et al., 2005), with numerous studies recreating the P300 with altered paradigms (Fabiani et al., 1986; Jeon & Polich, 2003; Katayama & Polich, 1996; Krigolson, Williams, Norton, Hassall, & Colino, 2017; Picton, 1992; Polich, 1989; Rushby et al., 2005), or eliciting it in sensitive patient populations such as schizophrenics (Jeon & Polich, 2003; Qiu et al., 2014; Turetsky et al., 2015; Winterer et al., 2003).

1.6 The Proposed Study

Over the course of three experiments I aim to show that the P300 component is both a reliable indicator of implicit learning processes and is robust enough to be elicited in complex environments. The experiments laid out are incremental in nature and rely on tightly controlled environments to isolate any changes that should occur in neural or behavioural activity.

There are several questions I wish to investigate based on the literature, requiring several incremental experimental designs. The first question will be whether the P300 is indicative of implicit learning processes taking place. To assess this an extremely simplified learning task will

need to be performed, to allow little confounding factors to influence the paradigm. As such, I propose the use of a modified oddball paradigm. In this task, participants will only be asked to classify whether a dot on the screen is “frequent” or “infrequent”. When analysing the data, I will be observing the P300 amplitude on a trial-by-trial basis. By doing this, I can assess whether the P300 amplitude is immediately tied to a stimulus (Coenen, 1995; N. K. Squires et al., 1975), or whether there is some time where behavioural learning and neural correlates are influenced in tandem as part of a learning process. Participants will never be given feedback on whether their classifications are correct, only the appearance of the next stimuli in the sequence. This should mean that they are learning the accuracy of their estimates purely by prediction errors. I predict that P300 amplitude will correlate strongly with these behavioural prediction errors and should respond in a similar manner to behavioural measures such as accuracy percentage.

Should there be signs of learning processes taking place, Experiment Two will validate this process by disrupting the learning process. If the P300 is indeed an index of prediction error activity, it should stand to reason that P300 amplitude will take time to readjust to new stimulus frequencies when they are changed without participant awareness. I hypothesize that after stimulus frequencies are changed, P300 amplitude will change over several trials to respond to the new relative frequency in step with behavioural adjustments concordant to prediction errors taking place. This paradigm will be identical to the first aside from a switch in stimulus frequencies half-way through. Thus I will be both replicating the first experimental conditions, while further testing the conclusions made in a more complex environment.

Experiment three will be yet another capstone on this test. In theory, the P300 (and especially P3b) requires task engagement to be elicited (Coenen, 1995; Donchin, 1981; Polich & Kok, 1995; Soltani & Knight, 2000). Here, I seek to investigate the implicit learning process on a

task that is not engaging the participant. Participants will be distracted with a modified time-estimation task while and auditory oddball paradigm is overlaid on the task beyond their upfront attention. If the P300 is elicited, I predict that it should demonstrate similar scaling properties to the first two experiments by increasing in amplitude over the first few trials of the experiment relative to the stimulus frequencies.

By targeting one facet of the of this process at a time, I hope to have reduced the limitations that could be implied by overly complex tasks. These experiments will develop from a measurement of P300 amplitude in a sterile and simple task to an assessment made in a convoluted environment with multiple tasks being learned and switched between. I ultimately propose that P300 amplitude will correlate strongly to prediction error behaviour in an implicit learning environment. I further predict that P300 amplitude will predict implicit learning processes in non-engaging tasks.

Chapter Two: Experiment One – Learning Without Feedback: Does the P300 Encode an Implicit Prediction Error?

2.1 Introduction

Reinforcement learning theory posits that we learn via discrepancies between outcomes and expectations that drive changes in behaviour, aka. prediction errors (Glimcher, 2011). By utilizing predictions in this manner, we as humans can shape our behaviour and optimize our decision-making processes (Cohen & Ranganath, 2007; Glimcher, 2011; O'Doherty, Cockburn, & Pauli, 2017; Trepel, Fox, & Poldrack, 2005). In a foundation work, Schultz and colleagues (Schultz et al., 1993) found that as monkeys learned the task at hand, their dopamine neurons would fire to a cue that predicted rewards, and fail to fire to the reward itself. By now providing a cue and *no* reward, it was observed that the dopamine response remained for the cue and produced depression for the expected reward. Schultz and colleagues proposed that throughout the experiment, the monkey was making prediction errors. Therefore, dopamine systems may be a key place to start investigating new learning processes.

However, this work leads to a reasonable question: how do these prediction errors allow us to learn when we are not given feedback? Feedback seems critical to allow us to make adjustments to our behaviour on subsequent attempts (Glimcher, 2011). There are many scenarios in which we learn without immediate explicit extrinsic reinforcement such as exploratory speech production (Cleeremans et al., 1998; Michas & Berry, 1994). These feedback-free environments however still have been found to have easily observed learning behaviour, often driven by prediction errors (Bray & O'Doherty, 2007; Ishikura, 2008; Montague et al., 1996; Walsh & Anderson, 2012).

These feedback-free conditions can also be called conditions of implicit learning. These are tasks that are learnable, but lack immediate reward or punishment feedback to the participant

(Bray & O'Doherty, 2007; Glimcher, 2011; Sambrook & Goslin, 2015). As these environments are learnable, they are often studied using similar methods to typical explicit reinforcement learning paradigms. This includes the utilization of EEG to examine direct neural responses to novel stimuli, such as the P300 ERP component (Coenen, 1995; Donchin, 1981; Krigolson & Holroyd, 2007; Patel & Azzam, 2005).

The P300 is a large positive deflection that occurs around 300ms after stimulus onset. Its amplitude is quantified as the voltage difference between a pre-stimulus baseline and the mean peak (+/-25ms) of the components peak latency window (Polich, 2007). The P300 was studied ahead of many other components due to its ease of elicitation and detection utilizing an "oddball" paradigm (Ritter et al., 1968; Sutton et al., 1965). Participants are presented with a series of two stimuli in which the frequency of one differs from the other. The less frequent "oddball" is considered a target stimuli which the participant must attend to either by pressing a button or mentally counting its occurrences (Polich, 2004; W. Ritter & Vaughan, 1969). The stimuli can be either auditory as in Sutton (1965) or Ritter's work (1968), or visually based such as in later works (Picton, 1992; Squires et al., 1975). The P300 observed is typically on the difference wave of this oddball stimuli average response minus the frequent control stimuli average response.

The P300 is often broken down into a more frontal and lower latency P3a component and a later, more parietal P3b component. In the frontal lobe, P3a activity in regards to working memory and attentional resource allocation have been suggested to be the result of dopaminergic activity while the P3b in the parietal area is heavily encouraged by dense norepinephrine inputs found in this region of the brain (Braver & Cohen, 2000; Nieuwenhuis, Aston-Jones, et al., 2005). Utilizing this knowledge, and the sensitivity of the P300 component to stimulus

probability, an oddball paradigm is a prime candidate to begin investigation of the P300 components role in implicit learning (Duncan-Johnson, 1981; Katayama & Polich, 1996; Polich & Kok, 1995).

In the present study we sought to examine whether or not the processing of stimuli would elicit a prediction error like response in the absence of explicit feedback. First, a standard oddball task will be performed by participants in which they are required to respond to the stimuli based on perceived frequency. This data will be analyzed by the production of grand averages to confirm the presence of the P300. Secondly, a trial-by-trial analysis will investigate the process of P300 scaling to stimulus frequency, with the aim to detect patterns involved in the process and determine if the P300 scaling process is dependent on stimulus probability as the participant learns. Finally, the trial-by-trial data will be compared to the behavioural performance of the participants. By comparing the behavioural evidence of learning with the neural evidence of learning, we will draw our conclusions about the role of the P300 as an indicator of implicit learning processes. Given the nature of the task, I predict participants to have gradually increasing accuracy rates as the block progresses and they determine the relative frequencies of the coloured squares. Accuracy would also be expected to increase more quickly in a block where the frequency discrepancy is larger, as this discrepancy is more readily assessed in a short period of time. I ultimately hypothesize that the P300 amplitude will gradually increase from a common start point before scaling based on the stimulus frequency (Katayama & Polich, 1996; Polich, 1990), in a manner that correlates with behavioural learning and prediction error mechanisms (Montague et al., 1996; Schultz, 2016; Walsh & Anderson, 2012; Zarr & Brown, 2016).

2.2 Method

2.2.1 Participants

Participants (n = 18: age range 18-25) from the University of Victoria participated in the experiment. All participants had normal or corrected-to-normal vision, no known neurological impairments, and were recruited through voluntary extra course credit in a psychology course. Informed consent, approved by the Human Research Ethics Board at the University of Victoria was obtained. The study followed ethical standards as prescribed by the 1964 version of the Declaration of Helsinki policy statement and all following revisions.

2.2.2 Procedure

Participants were seated in a sound dampened room, in front of a 19" LCD computer monitor. Using a standard USB mouse, participants completed a modified oddball task while EEG data were recorded (ActiCAP, Brainproducts GmbH, Munich, Germany). The experimental task was coded in MATLAB programming environment (Version 8.6, Mathworks, Natick, U.S.A.) using the Psychophysics Toolbox extension (Brainard, 1997). Participants completed a variant of the standard visual oddball paradigm (Picton, 1992; N. K. Squires et al., 1975). In an oddball task participants are presented with a series of stimuli which are then occasionally interrupted with a infrequent deviant stimulus, aka the "oddball". The neural response to this oddball is then recorded. Here, we did not inform participants of the frequencies of the stimuli they would be presented. Not knowing which was the oddball, they were to respond by keypress whether they believed the presented stimulus was the frequent or infrequent in the given block. Note that while both stimuli were pre-labelled as frequent or infrequent in the experiment code, the participants beliefs would be based on observed frequency, not nominal frequency. For example, if the colour series for a block of trials was blue and green squares and a blue square appeared on a given trial participants would classify the blue square as frequent if they believed

they had seen more blue that green squares or they would classify the blue square as infrequent if they believed they had seen more green than blue squares. An example of four trials is seen in *Figure 3*.

2.2.3 Experimental Task

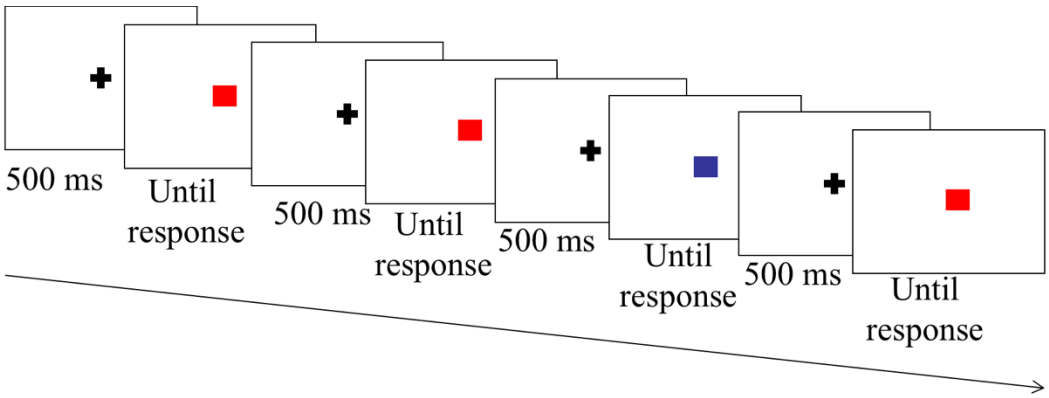

Figure 3. Example of Experimental Procedure. Participants would view the stimulus until it was classified as either "frequent" or "infrequent". The presentation times of the stimulus and the fixation cross were jittered randomly by up to 200ms to prevent frequency effects of the presentation.

Each trial of our task began with a black fixation cross being presented for 300 to 500ms on a dark grey background. This was followed by the presentation of a randomly coloured square from a possible pair of colours (two colours were used for each block of trials, colours were changed randomly between blocks). Squares were presented until the participant responded by depressing a button to indicate the presented square was either "frequent" or "infrequent" in appearance. As noted above, participants were asked to classify the stimuli based on frequency although they were not informed about the underlying frequency distribution of square appearance. Within the course of the experiment, participants encountered one of two different

block conditions at a time - the frequencies of the squares were either 40% or 60% (Condition 1) or 10% and 90% (Condition 2). Following the classification of a presented square as either frequent or infrequent, the black fixation cross reappeared initiating the next experimental trial. If a participant did not respond, the next trial was initiated after 2.5 seconds. Participants completed 60 blocks of 30 trials and unique square colours were used for each block. The distribution of the two frequency conditions within the 60 blocks was random as was the order of square appearance within each block of trials.

2.2.4 Data Acquisition

Participant's responses to the stimuli were recorded using a standard USB computer mouse in the MATLAB (Version 7.1, Mathworks, Natick, U.S.A.) programming environment utilizing Psychophysics Toolbox extension (Brainard, 1997). Accuracy ratings were calculated by trial as a grand average across all participants. For example, on the first trial if a participant responded correctly they would gain one point or zero if responding incorrectly. This participant's 1st trial would be taken from all 60 blocks completed and averaged. This is then repeated for every trial position from 1 to 30. This average from each participant can then be averaged across all participants to find the grand average accuracy percentage. This accuracy was presented as the percentage of participants who correctly selected whether a stimulus was frequent or infrequent on each trial position for both conditions across all blocks of trials.

EEG data were recorded using Brain Vision Recorder software (Version 1.3, Brain Products, GmbH, Munich, Germany) via 41 electrodes that were attached to a fitted cap, according to the standard 10-20 layout. Once fitted on the cap, electrodes were initially referenced to the whole head average. On average, electrode impedances were kept below 10 kΩ. EEG data were sampled at 250 Hz, amplified (Quick Amp, Brain Products GmbH, Munich, Germany), and filtered through a passband of 0.017Hz – 67.5Hz (90 dB octave roll off).

2.2.5 Data Processing

Data were processed offline with Brain Vision Analyzer 2 software (Version 2.1.1, Brainproducts, GmbH, Munich, Germany) using methods we have previously employed (see http://www.neuroeconlab.com/data-analysis.html). First, excessively noisy or faulty electrodes were removed. Then data was down sampled to 250Hz. The EEG data were then re-referenced to an average mastoid and then filtered using a dual pass 4th order roll off Butterworth filter with a passband of 0.1 Hz to 30 Hz in addition to a 60 Hz notch filter. Next, segments encompassing the onset of each event of interest (1000 ms before to 2000 ms after) were extracted from the continuous EEG. Following segmentation, independent component analysis was used to correct ocular artifacts (Delorme & Makeig, 2004; Luck, 2014). Data were reconstructed after the independent component analysis and any channels that were removed initially were interpolated using the method of spherical splines. Following this, all segments were baseline corrected using a 200 ms window preceding stimulus onset. New, shorter epochs were then constructed – from 200 ms before to 600 ms after the onset of each event of interest (presentation of coloured square stimulus), separated into each trial position between 1 and 30. All segments within each trial were then submitted to an artifact rejection algorithm that marked and removed segments that had gradients of greater than 10 μV/ms and/or a 100 μV absolute within segment difference (rejection rates; 10% condition: 5.6%, 40% condition: 4.5%, 60% condition: 4.1%, 90% condition: 4.0%). Finally, an average of the remaining segments for each stimulus was created for all given trials in that position.

For each trial position and stimulus, grand average ERP waveforms were created by averaging the data obtained for each participant. For example, the 10% condition segments were averaged for the first trial of each block. This was then repeated for every remaining trial position, for all stimulus frequencies, producing 30 trial position grand average waveforms for

each of the four stimulus frequencies. The P300 ERP component of interest was quantified as the maximal positive difference from baseline between 250 and 400ms of stimulus onset (Duncan-Johnson, 1981; Krigolson & Holroyd, 2007; Patel & Azzam, 2005; Picton, 1992; Polich, 2003). The electrode Pz was used with reference from previous literature (Dien et al., 2003) and based on visual inspection of scalp topographies as seen in *Figure 16*.

2.2.6 Data Analysis

Statistical analyses were performed using R Statistical Software (R. C. Team, 2016; Rs. Team, 2015) and data were plotted for inspection using Brain Vision Analyzer 2 software (Version 2.1.1, Brainproducts, GmbH, Munich, Germany) and R Studio (Rs. Team, 2015).

Participant accuracy was subjected to a simple linear regression analysis for both block types to determine if accuracy increased over subsequent trials, i.e. participants were learning the stimulus frequencies. Reaction time and accuracy was compared between stimulus frequencies using Repeated Measures ANOVA, followed by Tukey post-hoc testing. Correlations were also made between P300 amplitude and reaction time and accuracy.

The trial-by-trial data was subjected to an analysis similar to traditional inflection point analysis (Bronshtein, Semendyayev, Musiol, & Mühlig, 2015). Traditional inflection point analysis observes points in which the second derivative of a function is equal to zero. This find points in which the line changes sign but is quite overzealous for noisy trial by trial data such as what is being analyzed here. Here I use a different method meant to conservatively note points of change using similar principles. Rather than looking for the trial in which the second derivative passes through zero, this method will observe where the function has essentially consistently passed from a non-zero to a near zero second derivative. This process will highlight areas of distinct change in curvature as a result. First, the difference is calculated between each trial position and the position before it. Then, starting at the 6^{th} trial position, we take the absolute

value of the mean difference for each of the trials +/-5 trial positions from this trial. The absolute value is used as we are looking for changes in curvature, and an equal change on both sides of the trial we are analysing may be masked if not using absolute values. Inflection points are now chosen as any point greater than or equal to a predefined cut-off value. A cut-off value must be chosen for changes in slope to have granular enough sensitivity to not be triggered by small changes expected of variance, but not so insensitive as to say that slopes were not changing on curved lines. In this case, a cut-off value of 10 was chosen for detection. This value was shown on test data to not mark linear functions, while marking a sample sine wave with an amplitude comparable to the grand average P300 amplitude reliably and only in peak areas. A sample of this function created in R (R. C. Team, 2016) is shown in *Supplementary Figure 31.* Inflection Point R Script.

In other analyses Repeated Measures ANOVA comparisons were made between trial position and both frequency and amplitude. These ANOVA results were followed by paired t-tests. Further analyses were conducted used Bayes Factors, comparing once again trial positions effect on amplitude, and the effect of stimulus frequency. Bayes Factor analyses compares the likelihood ratio of multiple hypothesis based on Bayes' Theorem (Kass & Raftery, 1995). Given a set of data, each model has a reported K value that represents the likelihood that each the given model arises from that data set. K values over 20 are generally considered to be strong evidence for a given hypothesis (Kass & Raftery, 1995).

2.3 Results

Behavioural Analysis

Behavioural data showed expected results in terms of participant performance. Simple linear regression showed a regression equation was found ($F(1,28)= 77.87$, $p < 0.001$), with an

R^2 of 0.74 for the more difficult 60/40% condition. Participant accuracy increased 1% for each trial that passed. Similarly for the 90/10% condition a regression equation was found (F(1,28)= 11.63, p = 0.002), with an R^2 of 0.29. Participant accuracy increased 0.2% for each trial that passed. Mean accuracy by trial was lower in the more difficult 60/40% condition (M=79.1%, 95% CI [75.5, 82.8]) compared to the easier 90/10% condition (M=93.9%, 95% CI [92.9, 95.0]), $t(29) = 9.45$, $p < 0.001$. Mean reaction time by trial was also slower in the more difficult 60/40% condition (M=351.97ms, 95% CI [339.32, 364.62]) compared to the easier 90/10% condition (M=263.56ms, 95% CI [248.82, 278.29]), $t(29) = 25.37$, $p < 0.001$. This is shown in detail in *Figure 4* and *Figure 7*.

Inflection point analyses showed that the mean accuracy inflection point was the 11th trial [8.89, 13.94], while the mean reaction time inflection point was the 10th trial [8.66, 11.50]. These inflection points are shaded in *Figure 4* and *Figure 7*.

However, of note is the strong correlations between P300 amplitude and the measures of selection accuracy (r = 0.76) and reaction time (r = -0.85) shown separately and as mean comparisons in *Figure 8* and *Figure 9* respectively.

Finally, the observed stimulus frequency was plotted for comparison to P300 scaling effects (*Figure 11*). In addition, participant responses to the observed stimuli (*Figure 5*) were used to calculate d prime, a common index of learning progress in signal detection theory (Azzopardi & Cowey, 1998; McFall & Treat, 1999; Stanislaw & Todorov, 1999).

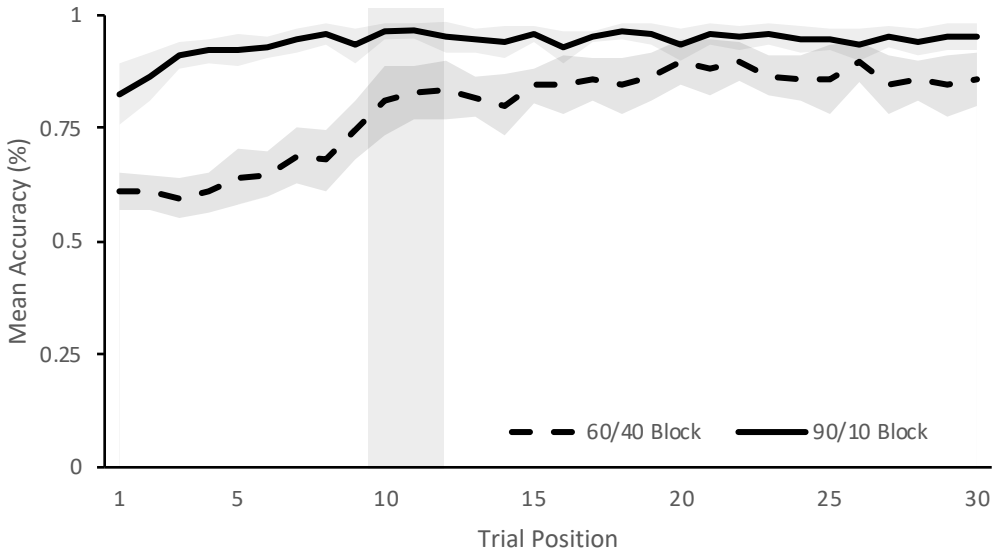

Figure 4. Accuracy by Trial. Behavioural data shows the running accuracy average of all participants over the course of a block. Line type indicates the block frequency type. Shaded areas represent 95% confidence interval. Participants learned the 90/10 condition much more quickly, and a large inflection is seen around the 9th trial of the 60/40 condition where participants appear to have reduced exploration behaviour.

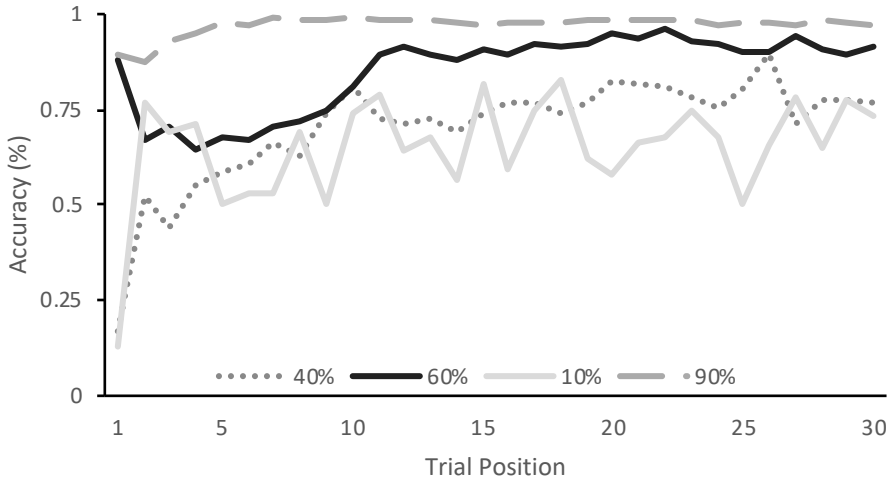

Figure 5. Accuracy by Trial by Stimulus. Behavioural data showing the running accuracy average of all participants for each stimulus over the course of a block. Line type indicates

stimulus frequency. This shows the rapid adjustment in behaviour within a few trials to the 90/10% stimulus condition.

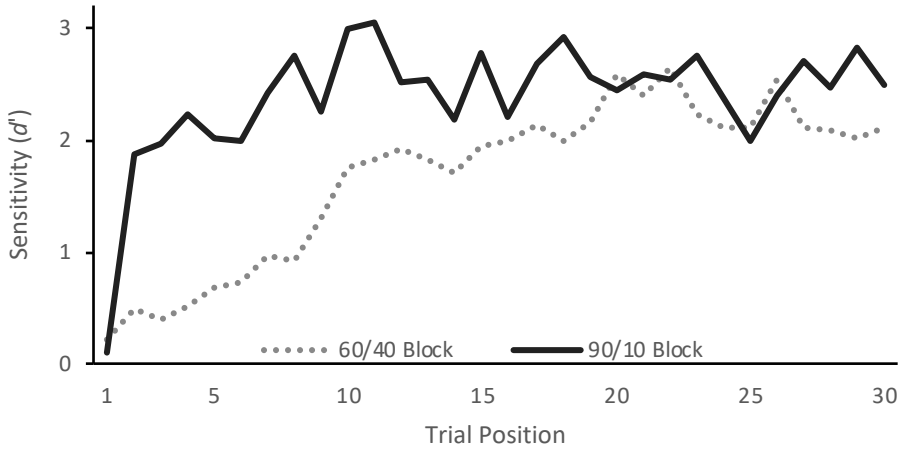

Figure 6. Sensitivity by Trial. The *d*' calculated on a trial by trial basis shows steadily increasing *d*'. This is in line with accuracy measures for each block type, and the increasing value of *d*' is suggestive that learning is taking place according to Signal Detection Theory.

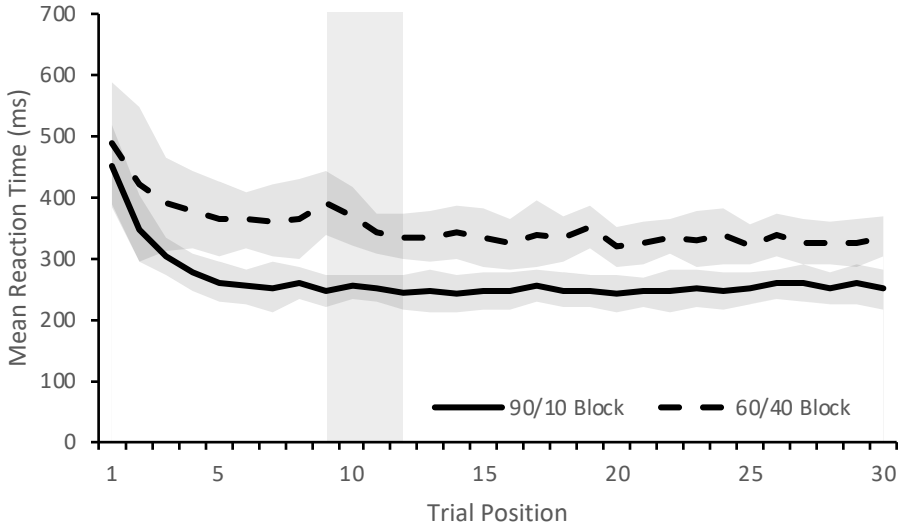

Figure 7. Reaction Time by Trial. Shaded areas represent 95% confidence interval. Reaction time steadily quickened for participants as the task progressed. The faster acceleration and

overall RT difference of the 10/90 condition is likely due to the much greater ease of classifying for this block condition.

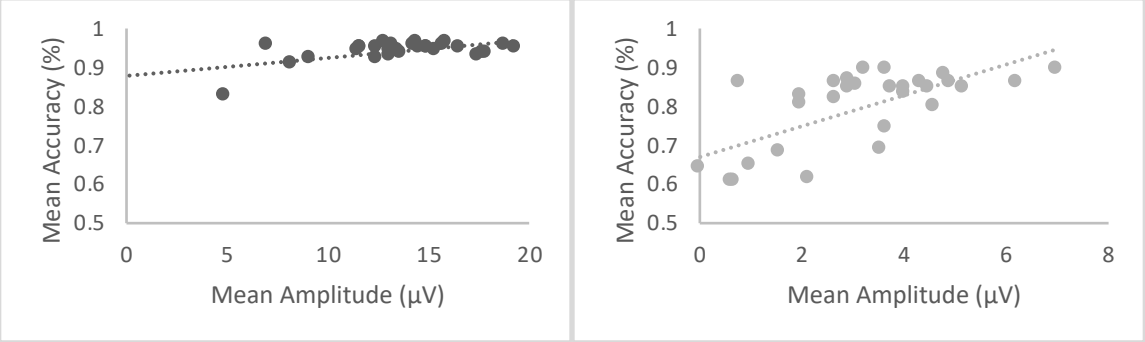

Figure 8. Correlation of P300 amplitude to accuracy. Here, data is shown in both the 90/10% block (left) and 60/40% block (right). Mean participant accuracy is compared to mean P300 amplitude for each of the 30 possible trial positions. The overall correlation between mean accuracy and mean amplitude is large and positive (r = 0.83).

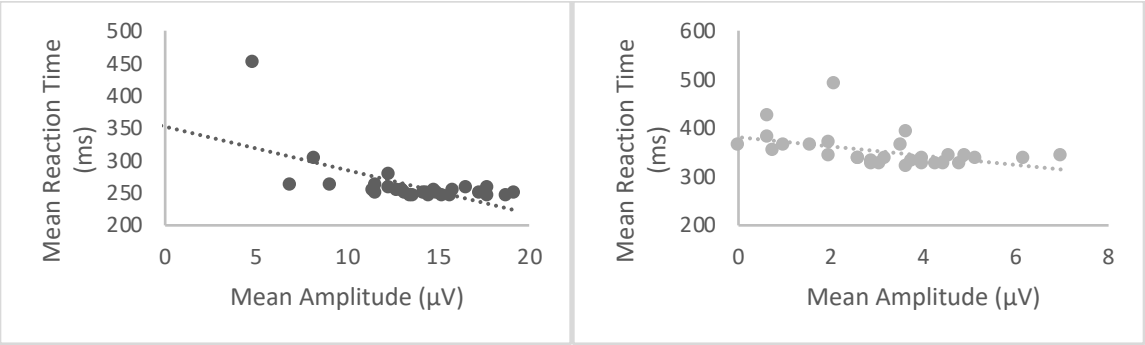

Figure 9. Correlation of P300 amplitude to reaction time. Here, data is shown in both the 90/10% block (left) and 60/40% block (right). Mean participant reaction time is compared to mean P300 amplitude for each of the 30 possible trial positions. The correlation between mean RT and mean amplitude is large and negative (r = -0.71).

Trial-by-Trial Analysis

When focusing on the trial-by-trial analysis, the inflection point between the more linear and steeply scaling sections of the data reveals some commonalities (*Figure 10*). Participants appear to have a shift from a rapidly scaling amplitude to a more linear trend, based on inflection point, at the 9th trial [8.59 – 9.77]. This inflection occurs for all of the stimuli presented. A Repeated Measures ANOVA showed no difference in inflection point position between stimulus types $F(3,40) = 0.59$, $p = 0.63$. In addition, the linear behaviour of P300 amplitude beyond this inflection point produces distinctly different mean amplitudes for each frequency, consistent with the above findings (elaborated in *Figure 12*). Bayes Factor analysis also supports the contribution of trial position to the changes that occur by frequency throughout the block ($2\log_e B_{10} = 166.5$). These changes by frequency are detailed below. P300 amplitude additionally moderately correlated with observed stimulus frequency ($r = -0.45$) on a trial-by-trial basis.

An identical analysis pathway was applied to P300 latency on a trial-by-trial basis, shown in *Figure 13*. However, an inflection point could not be found for these trends.

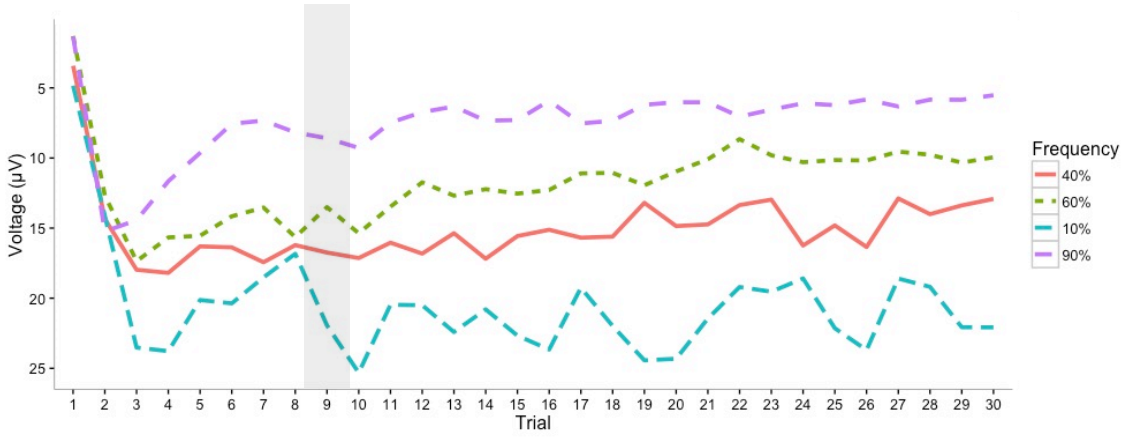

Figure 10. Peak Amplitude by Trial. Peak P300 amplitude by trial separated by stimulus frequency. Each stimulus begins with a comparable neural response, however by trial 9 [8.59-9.77] appearances of a given stimulus the P300 amplitude scales to stimulus frequency.

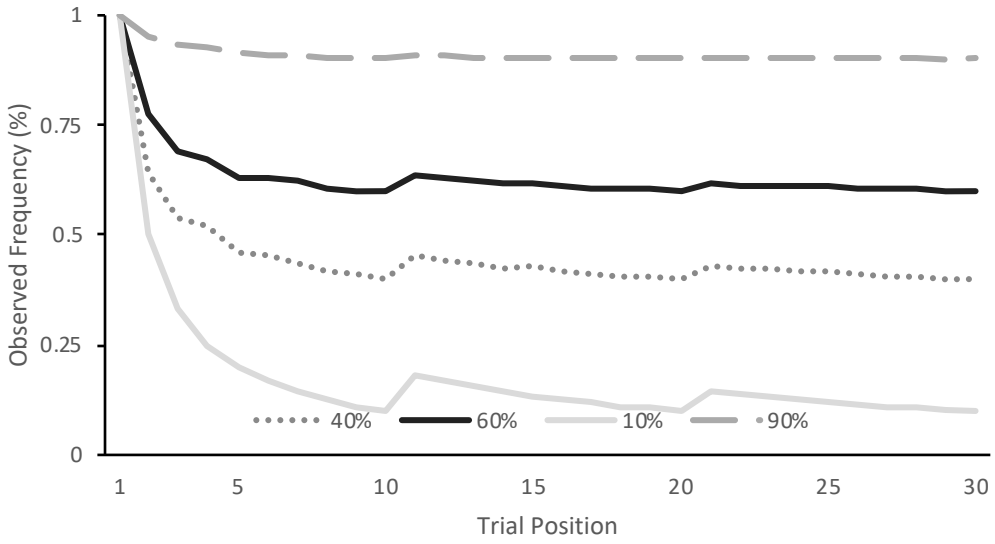

Figure 11. Observed Frequency by Trial. The observed average frequency of a stimulus by trial for each nominal frequency. Here we see the steady scaling of observed stimulus probability to actual stimulus probability over the course of the task block.

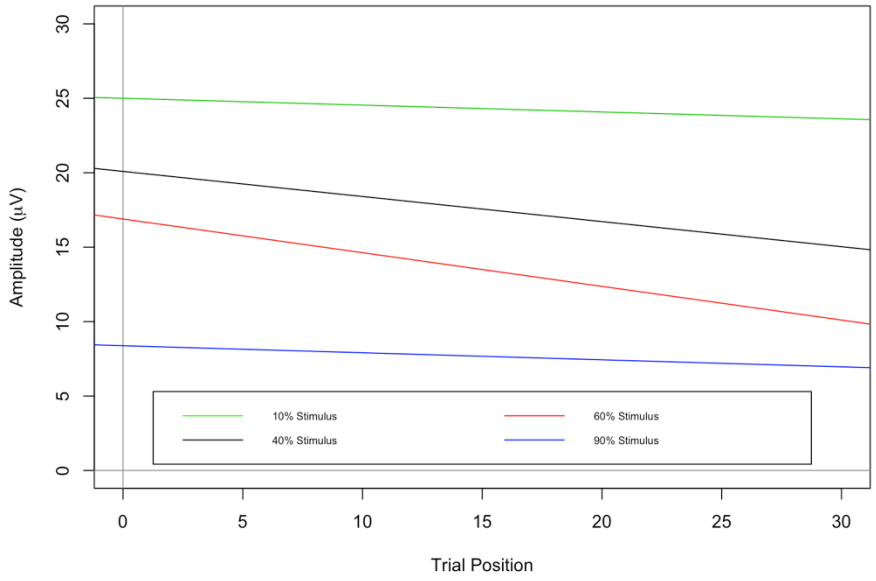

Figure 12. Amplitude by Trial with Linear Regression. The 40/60 condition appears to take an overall negative slope, though this may be a function of the tasks difficulty extending beyond 30 trials. Legend: Blue – 90%, Green – 10%, Red – 60%, Black – 40%.

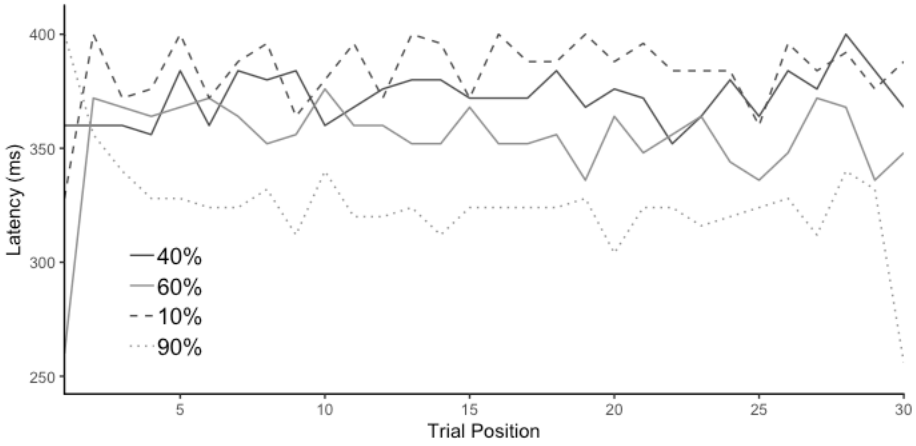

Figure 13. Peak Latency by Trial. Peak P300 latency by trial separated by stimulus frequency. Each stimulus begins with a varied neural response. There is a general visual trend to scale to stimulus frequency, however no inflection point is produced upon analysis, with latency scaling not easily discernable.

Grand Average Data

Repeated measures analysis of variance revealed a difference in both latency and amplitude of the P300 based on stimulus frequency in line with previous research (Mecklinger & Ullsperger, 1995; Nieuwenhuis, Aston-Jones, et al., 2005; Ullsperger & Baldeweg, 1991). To elaborate, repeated measures analysis of variance showed an effect of stimulus frequency on P300 latency, $F(3,51) = 31.48$, $p < 0.001$. Post-hoc with paired t-tests support this finding, with differences between all frequency's effects on latency, and no effect of subject. A repeated measures analysis of variance showed a similar effect of stimulus frequency on P300 amplitude, $F(3,51) = 88.68$, $p < 0.001$. Post-hoc with paired t-tests supported this finding, with differences between all frequencies effects on amplitude, and no effect of subject. These findings were further supported by Bayes Factor analyses, with $2\log_e B_{10}$ for these tests being 81.5 for stimulus frequency and subject effecting amplitude, and 40.3 for the same effect on latency. Latency and amplitude by frequency are shown in *Figure 14* and *Figure 15* and summarized in Table **1** with 95% CI. It is worth nothing that these grand averages are what is found when stimuli are collapsed across all trials (i.e., all of the data in the trial-by-trial analysis for a given stimulus is averaged). This is why the amplitudes found here are similar to the linear portion amplitude from the trial-by-trial seen in *Figure 10.*

Table 1

Mean Latency and Amplitude by Stimulus Frequency

Stimulus Frequency	Mean Latency (ms)	Mean Amplitude (μV)
10%	397 [381.6, 412.4]	23.34 [19.35, 27.33]
40%	380 [362.4, 397.6]	17.26 [14.43, 20.09]
60%	357 [347.9, 366.1]	13.81 [11.33, 16.29]
90%	332 [320.2, 343.8]	9.26 [7.17, 11.35]

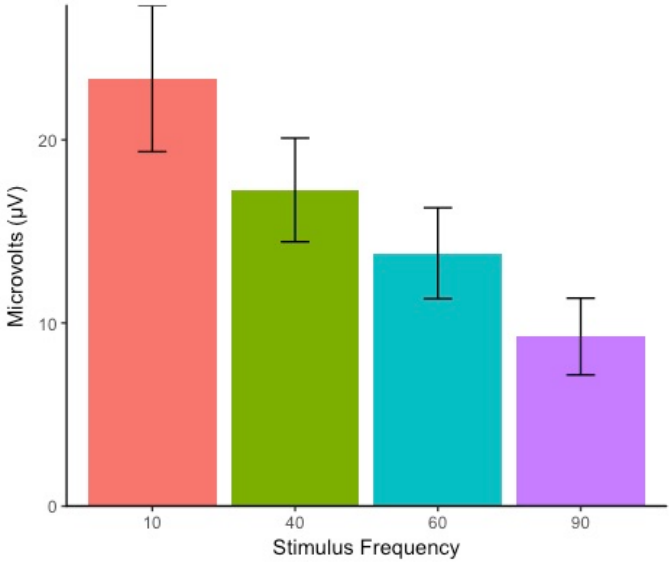

Figure 14. P300 amplitude by stimulus frequency. Amplitude is seen to decrease as stimuli become more frequent. Amplitude was calculated as mean voltage +/- 25ms around the peak value between 250-400ms. Stimuli were paired 40/60% frequency and 10/90% frequency in their respective blocks.

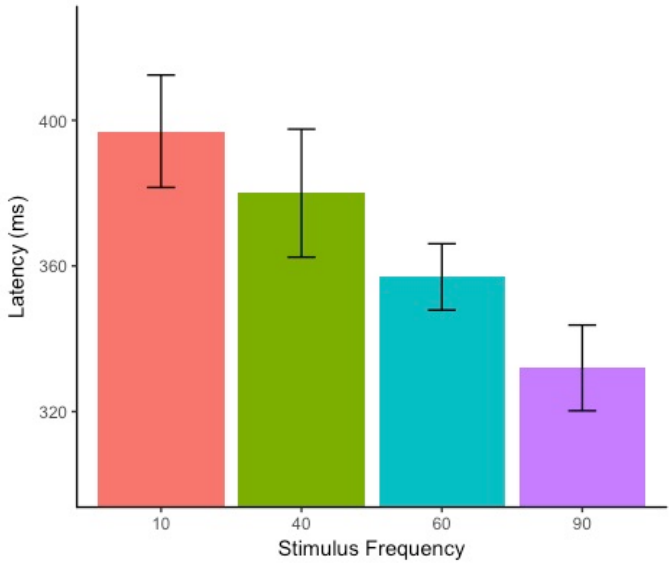

Figure 15. P300 latency by stimulus frequency. Latency is seen to decrease as stimuli become more frequent. Stimuli were paired 40/60% frequency and 10/90% frequency in their respective blocks.

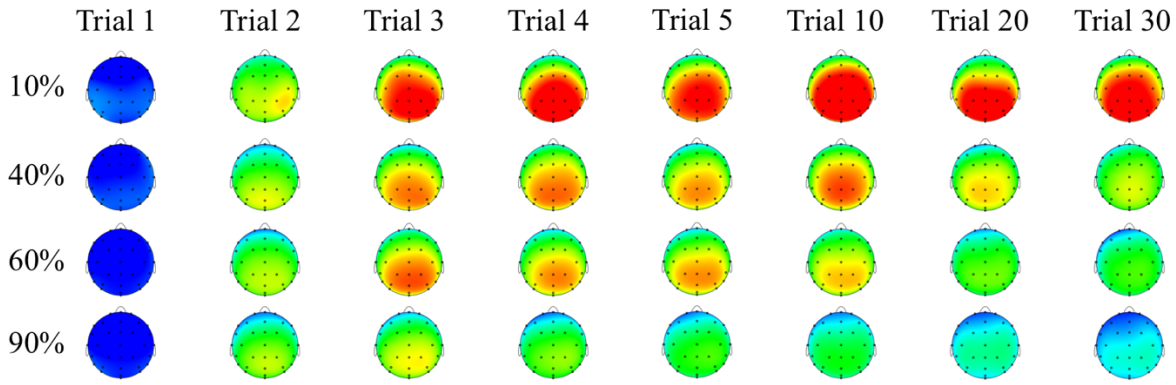

Figure 16. Topographic plot of P300 amplitude from various selected trials. Intensity can be seen to differentiate over time as participants acclimate to the frequencies of the stimuli.

A final result of note is the topographies presented in *Figure 16*. These topographical images show the above-mentioned increases in P300 amplitude over each trial, and the

subsequent scaling of amplitude to stimulus frequency. They also show a shift in topography from the midfrontal cortex region associated with the P3a to the parietal region of the brain which is more typically associated with the P3b (Polich & Kok, 1995).

2.4 Discussion

Not surprisingly, the grand average P300 amplitude was dependent on stimulus probability, as found in previous studies (Duncan-Johnson, 1981; Katayama & Polich, 1996; Nieuwenhuis, Aston-Jones, et al., 2005; Picton, 1992; Polich, 1990; Shijian Lu, Cuntai Guan, & Haihong Zhang, 2009; Trepel et al., 2005). Also in line with the previous research, there was a latency effect affecting a 65ms delay between the most frequent and infrequent stimuli (Donchin, 1981; Duncan-Johnson, 1981; Luck, 2014; Picton, 1992; Snyder & Hillyard, 1976; N. K. Squires et al., 1975).

Reassessing the data using a trial-by-trial analysis revealed that the first trial in each block for a given stimulus produced extremely similar P300 amplitudes for both frequent and infrequent stimuli. However, over the next few trials P300 amplitude began to scale to become representative of the stimulus probabilities (demonstrated in *Figure 10*). Following this adjustment period, P300 amplitude diverged and the classical P300 amplitude scaling to stimulus probability became clear by the 9th experimental trial. The P300 therefore is suggested to be related to a participant's trial-by-trial expectancies of a particular stimulus (Ishikura, 2008; Montague et al., 1996; N. K. Squires et al., 1975). This scaling behaviour was not found for P300 latency, although grand average analysis shows this behaviour. This suggests that while the P300 amplitude may scale by these expectancies, latency is not a reliable factor in this system.

Combining this information with the behavioural data, we can begin to see strong correlates between the behavioural and physiological data. An analysis of d' reveals a positive

going slope showing support that learning progress is being made (Stanislaw & Todorov, 1999). The high correlation shown between amplitude and both reaction time and accuracy suggest the P300 amplitude plays a key role as an indicator of the learning process. As the 9th trial inflection point in amplitude correlates with the approximately 10th trial inflection of the behavioural data, it seems that the leap in performance occurs nearly simultaneously with the neurophysiological changes as the participant learns. This suggests that the P300 at the very least is indexing exploratory behaviour processes, or as indicated by the trial-by-trial data reflecting a prediction of the upcoming stimulus. It is only as the data shows that on average participants have learnt the probabilities of the stimuli that the rate of change of the P300 amplitude by trial begins to approach zero. These prediction errors are informed not by explicit feedback, but by the presentation of the next stimulus. If the prediction is for an additional green dot to appear, the appearance of one would confirm the hypothesis made internally. A violation of this hypothesis would therefore elicit a larger P300 response (Nieuwenhuis, Aston-Jones, et al., 2005).

It is interesting to note that the inflection point is the same regardless of stimulus frequency. This suggests that the stimulus frequency effect is not the deciding factor for the rate of adaptation itself, instead it only affects the final P300 amplitude. Past research indicates that the P300 should scale to stimulus probability as some function of the processing required to update internal representations of that stimulus (Nieuwenhuis, Aston-Jones, et al., 2005; Patel & Azzam, 2005; Polich, 2007). It appears as though the *rate* of P300 adaptation, and therefore engagement of the LC-NE also scales to the stimulus in some way from very early on in implicit learning processes. In order for this prediction error compensation to consistently conclude around the same number of stimulus presentations in two different task difficulties with widely ranging stimulus frequencies, there may be a further system at play that enhances LC-NE

activity. The early trials frontal P3a topography shifting to posterior P3b supports that the LC-NE is facilitating the connection between the initial novel stimuli and the later context updating and signal processing (Nieuwenhuis, Aston-Jones, et al., 2005). I posit that the LC-NE phasic activity decreases after a "stimulus saturation" point, relegating further presentations of the stimulus to the posterior lobe for simple signal processing. Therefore, while the initial trials represent a learning and discovery phase the P300 is indicating an accruing of prediction errors, the later trials represent the more classical view of the P300 as an indication of processing. These later trials are therefore purely a function of stimulus probability effects that have been established (Katayama & Polich, 1996; Polich, 1990; Squires et al., 1975), while early trials are a function of both stimulus frequency and novelty processing arising from prediction errors. The novel and surprising stimuli generate larger prediction errors over subsequent trials, leading to the trend of the rarer stimuli producing the larger P300 as in previous literature. This also aligns with the initial trials having a near-zero P300 response, as predictions have not yet been made about stimulus presentation.

This view may be supported by comparing the observed frequencies in *Figure 11* with the P300 amplitudes shown in *Figure 10*. Observed frequency scales in a similar manner to the P300 amplitude, though it lacks the distinct sharp rise and following turn to a linear function. Instead, observed frequency slowly accumulates in a gradual curve. This suggests that some mechanism may be amplifying the effect of P300 amplitude in response to the earlier trials, i.e., a learning mechanism. The otherwise similar appearance of observed frequency and P300 amplitude supports the prediction error theory as a result of constantly diminishing corrections being made.

Conclusions

Our results suggest that the amplitude of the P300 reflects an expectancy driven prediction error. Specifically, a reinforcement learning process is taking place based on implicit violations of frequency generating intrinsic reward as opposed to explicit reward. Topography of the P300 supports the idea that the posterior parietal cortex may play a role in prediction error driven implicit learning mechanisms when the medial-frontal cortex cannot be used because of the absence of feedback (Glimcher, 2011; Ribas-Fernandes et al., 2011). Additionally, the data suggests that some further mechanism is at play during the initial trials that moves neural systems beyond novelty processing and initial stimulus categorization (Nieuwenhuis, Aston-Jones, et al., 2005; Polich, 2007).

Chapter Three: Experiment Two – Implicit Learning Response to Unknown Stimulus Changes

3.1 Introduction and Proposal

Implicit learning requires of course that a task be able to be learned and an optimal behaviour be discoverable (Bray & O'Doherty, 2007; Glimcher, 2011; Sambrook & Goslin, 2015). However, in paradigms that require a lack of feedback to the participant, we must somehow detect whether the participant is able respond to sudden changes in their learning environment. If a participant is forced away from an environment in which they can accumulate evidence for optimal behaviour, will their behaviour and subsequently their neural responses accommodate when the environment is now changed to be unclear? This area is ideal to be studied by methods such as EEG, where we can probe on a trial-by-trial basis whether implicit learning behaviours are responding to such changes. This includes changes in response based on prediction errors (Bray & O'Doherty, 2007; Ishikura, 2008; Montague et al., 1996; Walsh & Anderson, 2012). A prime candidate for this work is the P300 ERP component (Coenen, 1995; Donchin, 1981; Krigolson & Holroyd, 2007; Patel & Azzam, 2005), due to its sensitivity to disparities in stimulus frequencies (Coenen, 1995; Picton, 1992; Soltani & Knight, 2000; N. K. Squires et al., 1975; Wronka, Kaiser, & Coenen, 2008).

The P300 is a large positive deflection that occurs around 300ms after stimulus onset (Conroy & Polich, 2007). The P300 is most often elicited utilizing an "oddball" paradigm (Ritter et al., 1968; Sutton et al., 1965). Participants are presented with a series of two stimuli in which the frequency of one differs from the other. The less frequent "oddball" is considered a target stimuli which the participant must attend to either by pressing a button or mentally counting its occurrences (Polich, 2004; W. Ritter & Vaughan, 1969). Utilizing the sensitivity of the P300 component to stimulus probability, an oddball paradigm is a prime candidate to investigate the

P300 components role in implicit learning (Duncan-Johnson, 1981; Katayama & Polich, 1996; Polich & Kok, 1995).

This experiment aims to clarify the role of the P300 as an implicit learning indicator. Importantly we must limit the possibility that we are merely detecting the P300 resetting or scaling to block changes themselves and not to the stimulus frequencies each time. To do so, participants must be placed in an environment where they will both implicitly learn stimulus probabilities, as in an oddball-learning paradigm, and then be unknowingly made to learn new probabilities without any indication of change. This will be accomplished utilizing a mid-block change in stimulus frequencies, incorporating a mix of distinguishable and indistinguishable frequency pairings. This would elucidate if the implicit learning process will eventually scale to recognize that the task is not discernable, but still show similar behaviour leading up to this point. This would be in line with the theory that P300 scaling is tied to implicit learning processes.

I predict that P300 amplitude will indeed respond to, and immediately begin scaling to new stimulus probabilities. Specifically, I predict that after changing stimulus probabilities P300 amplitude will begin adjusting immediately and over several trials reach amplitudes similar to blocks where the newly presented stimulus frequency was constant.

3.2 Method

3.2.1 Participants

Participants (n = 15: age range 18-25) from the University of Victoria participated in the experiment. All participants had normal or corrected-to-normal vision, no known neurological impairments, and were recruited through voluntary extra course credit in a psychology course. Informed consent, approved by the Human Research Ethics Board at the University of Victoria

was obtained. The study followed ethical standards as prescribed by the 1964 version of the Declaration of Helsinki policy statement and all following revisions.

3.2.2 Procedure

Participants were seated in a sound dampened room, in front of a 19" LCD computer monitor. Using a standard USB mouse, participants completed a modified oddball task while EEG data were recorded (ActiCAP, Brainproducts GmbH, Munich, Germany). The experimental task was coded in MATLAB programming environment (Version 8.6, Mathworks, Natick, U.S.A.) using the Psychophysics Toolbox extension (Brainard, 1997). Participants completed a variant of the standard visual oddball paradigm (Picton, 1992; N. K. Squires et al., 1975). In an oddball task, participants are presented with a series of stimuli which are then occasionally interrupted with an infrequent deviant stimulus, aka the "oddball". The neural response to this oddball is then recorded. Here, we did not inform participants of the frequencies of the stimuli they would be presented. They were to respond by keypress whether they believed the presented stimulus was the frequent or infrequent in the given block.

3.2.3 Experimental Task

Each trial of our task began with a black fixation cross being presented for 300 to 500ms on a dark grey background. This was followed by the presentation of a randomly coloured square from a possible pair of colours (two colours were used for each block of trials, colours were changed randomly between blocks). Squares were presented until the participant responded by depressing a button to indicate the presented square was either "frequent" or "infrequent" in appearance. As noted above, participants were asked to classify the stimuli based on frequency although they were not informed about the underlying frequency distribution of square appearance. As such, the participants responses would be based on observed frequency, not nominal frequency. Atypical to a standard oddball task, there are four possible block conditions.

Each block was presented in randomized and counterbalanced order. Two of the four block conditions were a standard visual oddball paradigm with a 25/75% or 50/50% stimulus frequency pairing. The remaining two block types begin with a 25/75% or 50/50% stimulus pairing before switching to the other possible frequency pairing without informing the participant. This occurred on the 41st of 80 total trials that were within each block. Which stimuli of the pair that was made infrequent was randomized at the time of the switch. Following the classification of a presented square as either frequent or infrequent, the black fixation cross reappeared initiating the next experimental trial. If a participant did not respond, the next trial was initiated after 2.5 seconds. Participants completed 20 blocks (5 of each condition) of 80 trials and unique square colours were used for each block.

3.2.4 Data Acquisition

Participant's responses to the stimuli were recorded using a standard USB computer mouse in the MATLAB (Version 7.1, Mathworks, Natick, U.S.A.) programming environment utilizing Psychophysics Toolbox extension (Brainard, 1997). Accuracy ratings were calculated by trial as a grand average across all participants. For example, on the first trial if a participant responded correctly they would gain one point or zero if responding incorrectly. This participant's 1st trial would be taken from all 60 blocks completed and averaged. This is then repeated for every trial position from 1 to 30. This average from each participant can then be averaged across all participants to find the grand average accuracy percentage. This accuracy was presented as the percentage of participants who correctly selected whether a stimulus was frequent or infrequent on each trial position for both conditions across all blocks of trials.

EEG data were recorded using Brain Vision Recorder software (Version 1.3, Brain Products, GmbH, Munich, Germany) via 41 electrodes that were attached to a fitted cap, according to the standard 10-20 layout. Once fitted on the cap, electrodes were initially

referenced to the whole head average. On average, electrode impedances were kept below 10 kΩ. EEG data were sampled at 250 Hz, amplified (Quick Amp, Brain Products GmbH, Munich, Germany), and filtered through a passband of 0.017Hz – 67.5Hz (90 dB octave roll off).

3.2.5 Data Processing

Data were processed offline with Brain Vision Analyzer 2 software (Version 2.1.1, Brainproducts, GmbH, Munich, Germany) using methods we have previously employed (see http://www.neuroeconlab.com/data-analysis.html). First, excessively noisy or faulty electrodes were removed. Then data was down sampled to 250Hz. The EEG data were then re-referenced to an average mastoid and then filtered using a dual pass 4th order Butterworth filter with a passband of 0.1 Hz to 30 Hz in addition to a 60 Hz notch filter. Next, segments encompassing the onset of each event of interest (1000 ms before to 2000 ms after) were extracted from the continuous EEG. Following segmentation, independent component analysis was used to correct ocular artifacts (Delorme & Makeig, 2004; Luck, 2014). Data were reconstructed after the independent component analysis and any channels that were removed initially were interpolated using the method of spherical splines. Following this, all segments were baseline corrected using a 200 ms window preceding stimulus onset. New, shorter epochs were then constructed – from 200 ms before to 600 ms after the onset of each event of interest (presentation of coloured square stimulus), separated into each trial position between 1 and 80. All segments within each trial were then submitted to an artifact rejection algorithm that marked and removed segments that had gradients of greater than 10 μV/ms and/or a 100 μV absolute within segment difference. Finally, an average of the remaining segments for each stimulus was created for all given trials in that position.

For each trial position and stimulus, grand average ERP waveforms were created by averaging the data obtained for each participant. The P300 ERP component of interest was

quantified as the maximal positive difference from baseline between 250 and 400ms of stimulus onset (Duncan-Johnson, 1981; Krigolson & Holroyd, 2007; Patel & Azzam, 2005; Picton, 1992; Polich, 2003). The electrode Pz was used with reference from previous literature (Dien et al., 2003) and based on visual inspection of scalp topographies.

3.2.6 Data Analysis

Statistical analyses were performed using R Statistical Software (R. C. Team, 2016; Rs. Team, 2015) and data were plotted for inspection using Brain Vision Analyzer 2 software (Version 2.1.1, Brainproducts, GmbH, Munich, Germany) and R Studio (Rs. Team, 2015).

Participant accuracy was subjected to a simple linear regression analysis for both block types to determine if accuracy increased over subsequent trials, i.e. participants were learning the stimulus frequencies. Reaction time and accuracy was compared between stimulus frequencies using Repeated Measures ANOVA, followed by Tukey post-hoc testing. Correlations were also made between P300 amplitude and reaction time and accuracy.

The trial-by-trial data was subjected to an inflection point analysis (Bronshtein et al., 2015). Traditional inflection point analysis observes points in which the second derivative of a function is equal to zero. This find points in which the line changes sign but is quite overzealous for noisy trial by trial data such as what is being analyzed here. Here I use a different method meant to conservatively note points of change using similar principles. First, the difference is calculated between each trial position and the position before it. Then, starting at the 6^{th} trial position, we take the absolute value of the mean difference for each of the trials +/-5 trial positions from this trial. The absolute value is used as we are looking for changes in curvature, and an equal change on both sides of the trial we are analysing may be masked if not using absolute values. Inflection points are now chosen as any point greater than or equal to a predefined cut-off value. A cut-off value must be chosen for changes in slope to have granular

enough sensitivity to not be triggered by small changes expected of variance, but not so insensitive as to say that slopes were not changing on curved lines. In this case, a cut-off value of 10 was chosen for detection. This value was shown on test data to not mark linear functions, while marking a sample sine wave with an amplitude comparable to the grand average P300 amplitude reliably and only in peak areas. A sample of this function created in R (R. C. Team, 2016) is shown in *Supplementary Figure 31.* Inflection Point R Script.

Here there is a focus on the comparison of pre and post flip in each block by 3x2 mixed-factorial ANOVA. Here, pre-post is compared for all stimulus P300 response amplitudes. The ANOVA were followed by post-hoc paired t-tests. Further analyses were conducted used Bayes Factors, comparing once again trial positions effect on amplitude, and the effect of stimulus frequency. Bayes Factor analyses compares the likelihood ratio of multiple hypothesis based on Bayes' Theorem (Kass & Raftery, 1995). Given a set of data, each model has a reported K value that represents the likelihood that each the given model arises from that data set. K values over 20 are generally considered to be strong evidence for a given hypothesis (Kass & Raftery, 1995).

3.3 Results

Behavioural data shows expected exploratory behaviour. Simple linear regression was run on the mean accuracy for each block condition. The results are shown in Table 2.

Table 2

Summary of Regression Analysis for Participant Accuracy per Trial

Condition	R^2	*df*	*F*	Intercept	Slope
B1 50/50%	0.92	1, 38	435.1***	72.91	-0.45
B1 25/75%	0.99	1, 38	6046.7***	41.76	0.38
B2 25/75%	0.89	1, 38	303.8***	54.52	0.94

B2 50/50%	0.71	1, 38	92.9***	84.88	0.04
B3 25/75%	0.68	1, 78	162.6***	77.66	0.23
B4 50/50%	0.79	1, 78	289.4***	50.01	-0.04

Note: *$p < .05$. **$p < .01$. ***$p < .001$.

Participant accuracy increased over trials for each condition that was learnable (non-equal stimulus probability). Similarly the 50/50% conditions had negative or near zero accuracy slopes in their regressions (Table 2). Correlation between P300 amplitude and accuracy was low in this task (*Figure 18*).

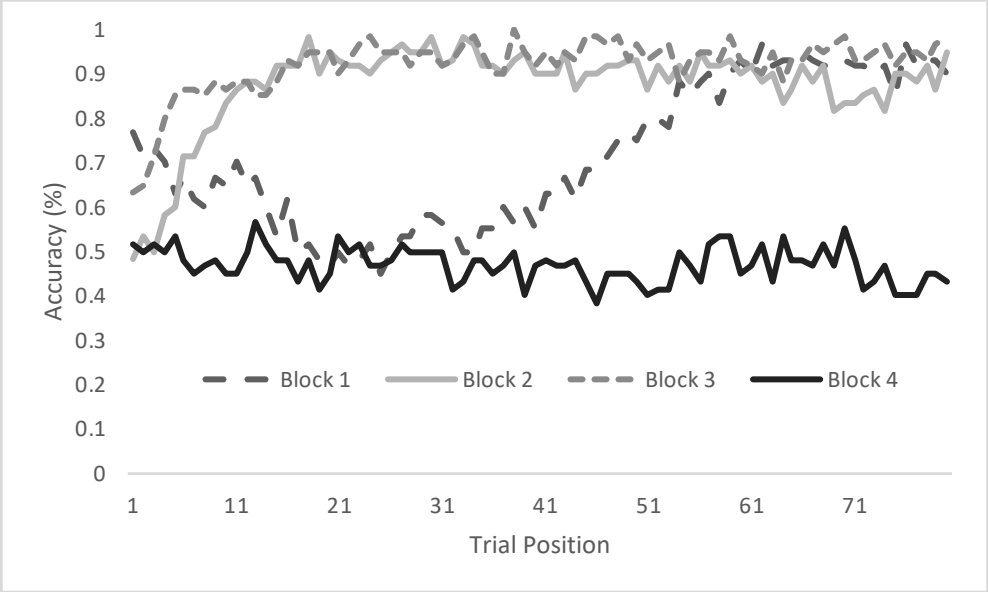

Figure 17. Participant Accuracy by Trial. The line plot shows the instantaneous accuracy by trial for each block type. The frequency pairings before and after trial 40 are as follows: (A) Block 1 Pre: 50/50, Block 1 Post: 25/75. (B) Block 2 Pre: 25/75, Block 2 Post: 50/50. (C) Block 3: 25/75. (D) Block 4: 50/50.

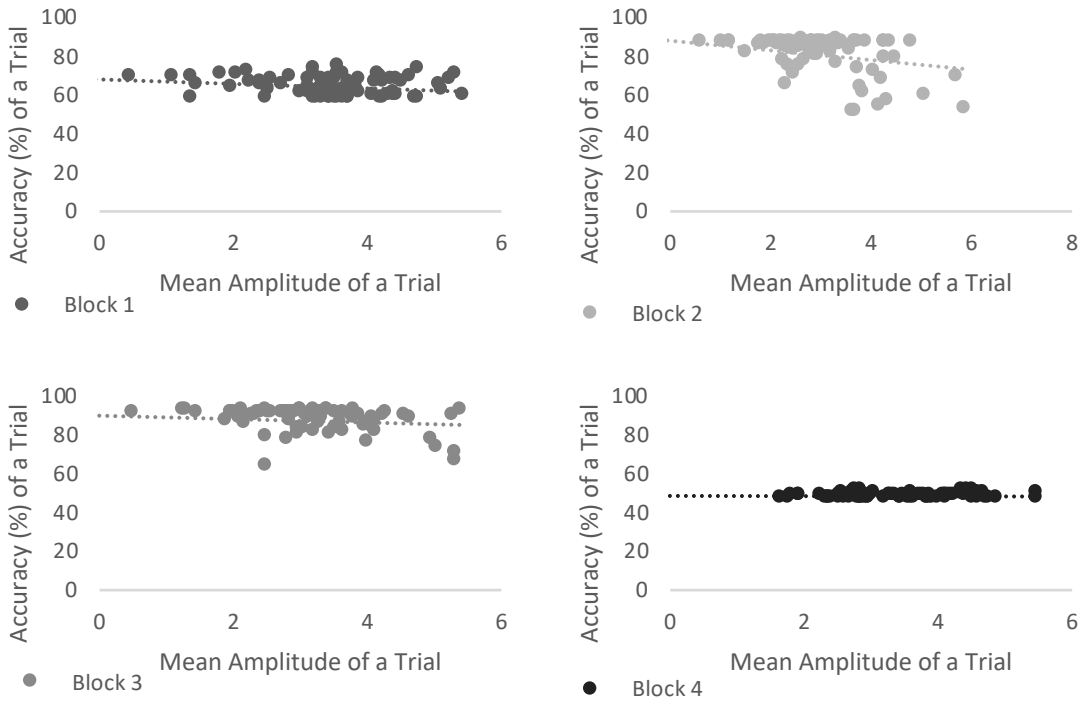

Figure 18. Mean Amplitude by Accuracy Correlation. The correlation between amplitude and accuracy is quite small despite visual impact (r = -0.19). The scaled P300 amplitude by block condition are easily visible.

A plot of amplitude by trial can be seen in *Figure 19* with Local Regression smoothing (Shyu, Grosse, & Cleveland, 2017) applied purely for ease of visualization, all statistics were performed on the raw data. Blocks which changed stimulus frequencies (Block 1 and 2) showed differences in P300 amplitude after the switch occurs ($F(1,7) = 14.60$, $p < 0.001$). This was supported by the paired t-tests, detailed in Table 3.

A plot of observed frequency by trial can be seen in *Figure **20***, which correlates to the trial-by-trial amplitude with a moderate correlation (r = -0.37).

Table 3

Significance report for ERP comparison.

Stimulus	Block 1	Block 2	Block 3	Block 4
Stim 1 (Post – Pre)	1.31ΔμV*	-1.55ΔμV**	-1.51ΔμV**	0.85ΔμV
Stim 2 (Post – Pre)	-1.32 ΔμV*	0.58ΔμV	-0.28ΔμV	0.40ΔμV
Pre Oddball	-1.53μV**	4.04μV***	2.97μV***	-0.12μV
Post Oddball	1.1μV	1.91μV***	1.74μV***	0.33μV

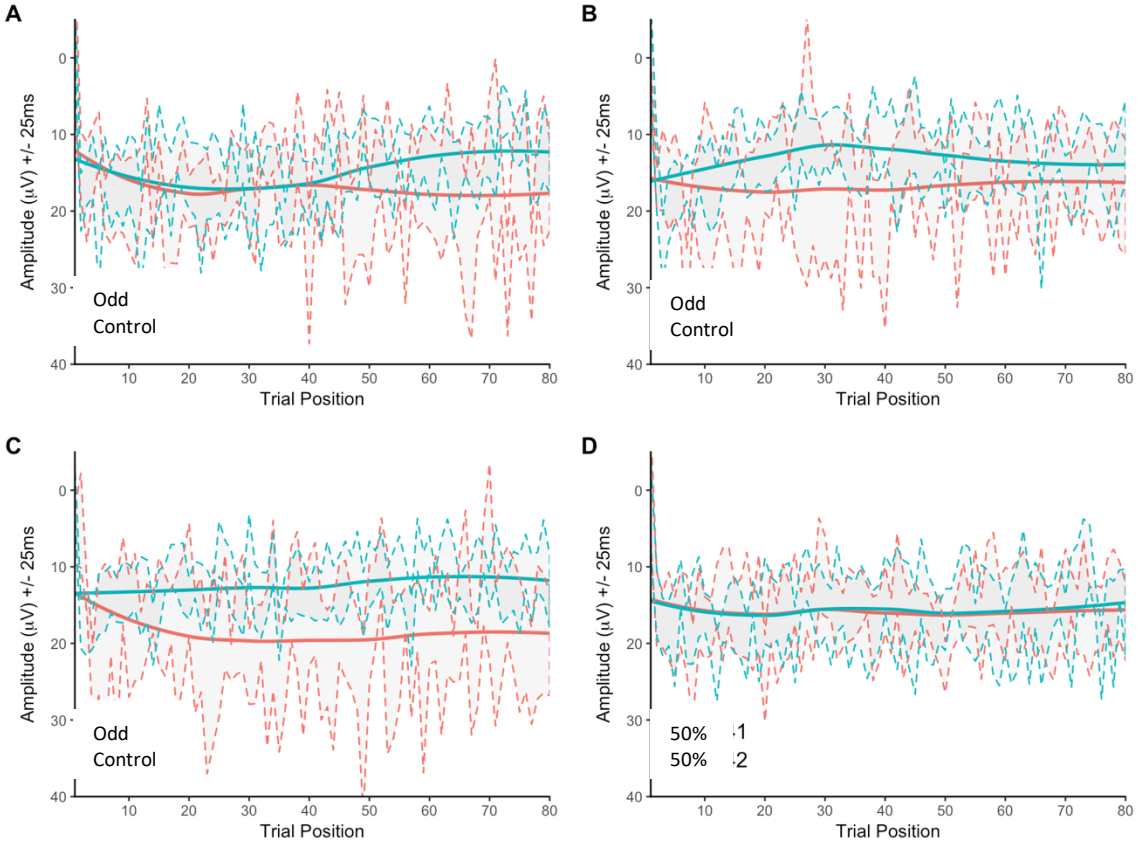

Figure 19. Smoothed Peak P300 Amplitude by Trial and Block. Peak P300 amplitude was calculated as +/-25ms about the mean peak value in a 250-400ms time window for each trial. Here local regression smoothing is applied for ease of visualization. Shaded areas represent 95% confidence intervals for the waveform. The frequency pairings before and after trial 40 are as

follows: (A) Block 1 Pre: 50/50, Block 1 Post: 25/75. (B) Block 2 Pre: 25/75, Block 2 Post: 50/50. (C) Block 3: 25/75. (D) Block 4: 50/50.

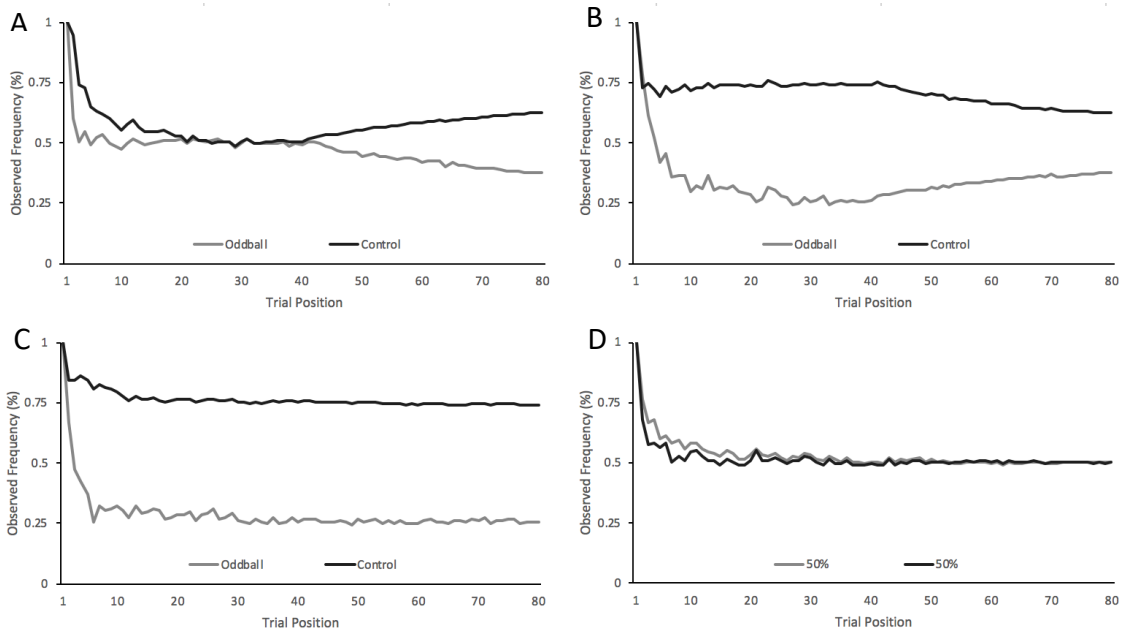

Figure 20. Observed Frequency by Trial. The observed average frequency of a stimulus by trial for each nominal frequency. Observed frequency scales in a similar manner to what is observed in the P300 amplitude shown in *Figure 19.* Smoothed Peak P300 Amplitude by Trial and Block. Peak P300 amplitude was calculated as +/-25ms about the mean peak value in a 250-400ms time window for each trial. Here local regression smoothing is applied for ease of visualization. Shaded areas represent 95% confidence intervals for the waveform. The frequency pairings before and after trial 40 are as follows: (A) Block 1 Pre: 50/50, Block 1 Post: 25/75. (B) Block 2 Pre: 25/75, Block 2 Post: 50/50. (C) Block 3: 25/75. (D) Block 4: 50/50. (A) Block 1 Pre: 50/50, Block 1 Post: 25/75. (B) Block 2 Pre: 25/75, Block 2 Post: 50/50. (C) Block 3: 25/75. (D) Block 4: 50/50.

3.4 Summary

This experiment aimed to validate the role of the P300 as an indicator of implicit learning by disrupting the learning process. One would expect the P300 to rescale to new stimulus probabilities in a timely manner if the participant is still being mindful of the stimulus probabilities and updating their internal representation of that stimulus (Polich & Kok, 1995; Soltani & Knight, 2000). Indeed, this stimulus probability switch was a necessary inclusion to the oddball-learning paradigm in order to exclude the possibility that P300 amplitude was not simply resetting and scaling due to the start of a new block, rather than starting from zero due to the stimulus itself.

This experiment showed results as expected of a prediction error driven system (Bray & O'Doherty, 2007; Schultz, 2016; Schultz et al., 1997). After switching stimulus probabilities mid-block, P300 amplitude scaled to the new stimulus probabilities rapidly. Further, amplitude changes after the switch were similar to amplitudes when the block had always contained those frequencies. This is to say, P300 amplitudes for 25/75% stimuli after the switch were similar to those found in the block that was entirely 25/75%. It is worth noting as well that similar scaling from near-zero microvolt P300 amplitude up to an amount relative to stimulus frequency is once more visually distinct in this data set. Amplitudes did not reset to zero after a frequency switch, but rather shifted over several trials from their current value to one representative as the new frequencies. This is suggestive that prediction errors are guiding a change in amplitude until new frequencies are learned by the participant.

This is supported by behavioural data showing participant accuracy steadily climbing in learnable environments, for example in Block 1 accuracy climbs when moved from a 50/50 condition to a 25/75 condition without informing the participant (second half of Block 1, and all

of Block 3 in *Figure 17*). Note that *Figure 17* shows no decline in accuracy in block 2, which contained a 25/75% stimulus frequency pairing before switching to 50/50%. This block does not properly visualize expected participant response due to an error in the experiment coding. As the frequent/infrequent label does not change with the block switch, participants were able to carry on with their choices in frequency labelling and maintain the same level of accuracy as they had before the block switch. This may also explain the lower correlation between P300 amplitude and accuracy shown in *Figure 18* as the P300 was indexing the sudden change in frequencies while our inadequate accuracy measure was not. However, it is important to note that the observed frequencies in *Figure 20* align with P300 amplitude trends. This lends further credence to the P300 amplitude tracking prediction errors on a trial-by-trial basis.

A comparison of observed stimulus frequency to P300 shows a more expected result by comparing the data shown in *Figure 19* and *Figure **20***. Here, the moderately correlation observed frequency supports the prediction error based explanation of P300 amplitude. Most interestingly, it evens lends support to the slower rate of P300 scaling in the Block B condition, where the stimuli switch from 72/25 to 50/50 frequency pairings. The P300 is responding at a rate congruent to the adjustment of expectations as more of one stimuli begin to appear.

Much of this data is supported by Walsh & Anderson (2012). This study focuses on neural reward correlates, specifically the feedback related negativity (FRN). However, they note the established difficulty in separating the N2, P300, and FRN (Donkers & van Boxtel, 2005). The P300 is quite sensitive to reward magnitude, but not valence (Walsh & Anderson, 2012; Wu & Zhou, 2009). Gambling tasks have previously supported the P300 (Hajcak et al., 2007; Mühlberger, Angus, Jonas, Harmon-Jones, & Harmon-Jones, 2017) as being sensitive to a combination of stimulus effects and overlap with feedback related negativity. The P300 as an

encoder of not just stimulus probability, but also stimulus value may fit the data being observed in this study.

Ultimately, I concluded that the P300 is indeed likely indicative of a prediction error driven learning process taking place due to the familiar scaling of P300 amplitude to stimulus frequencies. P300 amplitude scales in step with observed stimulus frequency as opposed to nominal stimulus frequency, further lending support to this mechanism. Behavioural data also appears to align with these data, although with less support then was seen previously.

Chapter Four: Experiment Three – Passive Implicit Learning in a Complex Task Environment

4.1 Introduction and Proposal

This experiment aims to answer whether the implicit learning P300 I have measured is present in non-attended stimuli. Unattended tasks such as an auditory oddball have a variety of potential uses for non-intrusive assessment. This includes scenarios such as fatigue assessment in medical and transport environments (Kaseda, Jiang, Kurokawa, Mimori, & Nakamura, 1998; Käthner et al., 2014; Murata, Uetake, & Takasawa, 2005; Polich & Kok, 1995; Uetake & Murata, 2000; Zhao et al., 2012). If an auditory oddball overlaid over a complex task can still assess implicit learning processes, it may be applicable to use this as an index of current learning capability in subject engaged in more complicated tasks.

Importantly, a passive task eliciting this detectable process would run counter to some expectations on P300 behaviour (Coenen, 1995; Donchin, 1981; Polich & Kok, 1995; Soltani & Knight, 2000) which indicate that P300 production requires task engagement. Despite this, implicit learning may be running through systems that are atypical of the type of P300 generated in these studies. Specifically, running under the hypothesis that the P300 I am detecting is produced by the LC-NE system adjusting neural processes associated with novelty and prediction errors (Nieuwenhuis, Aston-Jones, et al., 2005), I expect a P300 to still be produced. Indeed, I hypothesize that the same scaling system would be present as in previous experiments.

There is support for this P300 being producible in the paradigm I will be using specifically: auditory oddball paradigms are known to elicit remarkably similar P300 components in both attended and unattended stimuli (Polich, 1989; Ritter et al., 1968; Snyder & Hillyard, 1976; K. C. Squires et al., 1977; N. K. Squires et al., 1975). This has been a mainstay for P300 elicitation in participants who cannot perform discriminatory tasks. By showing that the

passive auditory P300 is produced in a not only passive, but distracted task and ultimately if there is evidence of prediction errors taking place, I hope to elaborate a method by which we can assess prediction error systems in nearly any situation or population.

4.2 Method

4.2.1 Participants

Undergraduate students (n = 12: age range 18-25) from the University of Victoria participated in the experiment. All participants had normal or corrected-to-normal vision, no known neurological impairments, and were recruited through voluntary extra course credit in a psychology course. Each participant provided written informed consent approved by the Human Research Ethics Board at the University of Victoria. The study followed ethical standards as prescribed by the 1964 Declaration of Helsinki policy statement and all following revisions.

4.2.2 Procedure

Participants were seated in a sound dampened room, in front of a 19" LCD computer monitor. Using a standard USB mouse and keyboard, participants completed both a time estimation task and a visual oddball task with an additional passive embedded audio oddball task while EEG data were recorded (ActiCAP, Brainproducts GmbH, Munich, Germany). The experimental task was coded in MATLAB programming environment (Version 8.6, Mathworks, Natick, U.S.A.) using the Psychophysics Toolbox extension (Brainard, 1997). Each participant would complete 10 blocks of another visually based task. Overlaid on these tasks was an auditory oddball task, in which two differing tones of beep were played at random intervals simultaneously over the entire experiment shown in Figure 21.

4.2.3 Experimental Task

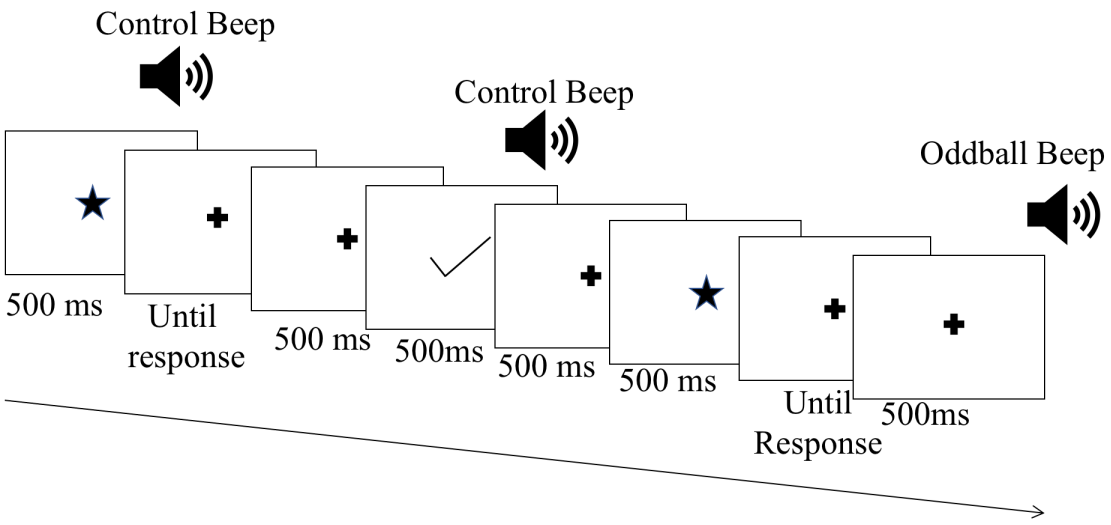

Figure 21. Example of Experimental Procedure. Participants completed a visually based task on the computer monitor. Throughout the experiment, auditory stimuli were presented 3 seconds apart +/- 1200ms. The oddball stimuli was jittered 200ms from visual cues to prevent frequency effects of the presentation.

Time Estimation Task. Participants were required to estimate the duration of 1 second (Miltner, Braun, & Coles, 1997). After the presentation of a fixation cross, each trial began with a cue that lasted for 50ms. Participants responded by mouse click when they estimated 1 second had passed.

Auditory Oddball Task. Overlaid over all blocks was an auditory oddball probe. Over intervals averaging 2.5s +/- 1.5 seconds, either an oddball tone (1600Hz, 65dB) would play for 50ms or a control tone would be played (800Hz, 65dB) for 50ms. The oddball tone played at a fixed 25% ratio of the auditory trials. Participants were not asked to attend to these stimuli and were told to focus on the stimuli directly related to the time estimation task at hand. Auditory oddball trials were intentionally jittered by 200ms to avoid overlap with time estimation cue.

4.2.4 Data Acquisition

Response time (ms) and accuracy (percentage on time) were recorded on each trial of the time estimation task using a standard USB mouse and keyboard in MATLAB ("MATLAB," 2016).

EEG data were recorded using Brain Vision Recorder software (Version 1.21, Brainproducts, GmbH, Munich, Germany) via 64 electrodes that were attached to a fitted cap, according to the standard 10-20 layout (ActiCAP, Brainproducts GmbH, Munich, Germany), in addition to two mastoid electrodes for offline reference. Once fitted on the cap, electrodes were initially referenced to a common ground. On average, electrode impedances were kept below 20 kΩ. EEG data were sampled at 500 Hz, amplified (ActiCHamp, Revision 2, Brainproducts GmbH, Munich, Germany), and filtered through an antialiasing low-pass filter of 8 kHz. To ensure temporal coincidence of event-markers with experimental stimuli a DATAPixx stimulus unit was used (VPixx, Vision Science Solutions, Quebec, Canada).

4.2.5 Data Processing

Data were processed offline with Brain Vision Analyzer 2 software (Version 2.1.1, Brainproducts, GmbH, Munich, Germany). First, excessively noisy or faulty electrodes were removed. The data was not down-sampled further. The EEG data were then re-referenced to an average mastoid and then filtered using a dual pass 4^{th} order Butterworth filter with a passband of 0.1 Hz to 30 Hz in addition to a 60 Hz notch filter. Next, segments encompassing the onset of each event of interest (1000 ms before to 2000 ms after) were extracted from the continuous EEG. Following segmentation, independent component analysis was used to correct ocular artifacts (Delorme & Makeig, 2004; Luck, 2014). Data were reconstructed after the independent component analysis and any channels that were removed initially were interpolated using the method of spherical splines. Following this, all segments were baseline corrected using a 200 ms

window preceding stimulus onset. New, shorter epochs were then constructed – from 200 ms before to 600 ms after the onset of each event of interest, i.e., the auditory beeps of the oddball task. All segments were then submitted to an artifact rejection algorithm that marked and removed segments that had gradients of greater than 15 μV/ms and/or a 100 μV absolute within segment difference. Approximately 1% of trials were removed for both auditory oddball and auditory control stimuli. Finally, an average of the remaining segments for each stimulus was created for all given trials of interest. This led to an overall average auditory oddball and auditory control waveform for each participant.

From these averages, grand average ERP waveforms were created by averaging the data obtained for each participant. The P300 ERP component of interest was quantified as the maximal positive difference from baseline between 250 and 400ms of stimulus onset (Duncan-Johnson, 1981; Patel & Azzam, 2005; Picton, 1992; Polich, 2003). The electrode Pz was used with reference from previous literature (Dien et al., 2003) and based on visual inspection of scalp topographies.

4.2.6 Data Analysis

During this analysis we will assess if the auditory oddball task did in fact elicit a distinct oddball response from our two stimuli by creating a difference wave of the subtracted grand average waveforms (odd – control). By comparing this difference wave to zero in the time range of the P300 using a paired t-test we will detect if the oddball was elicited.

Repeated Measures ANOVA comparisons were made between trial position and both frequency and amplitude. These ANOVA results were followed by paired t-tests.

The trial-by-trial data was subjected to an inflection point analysis (Bronshtein et al., 2015). Traditional inflection point analysis observes points in which the second derivative of a function is equal to zero. This find points in which the line changes sign but is quite overzealous

for noisy trial by trial data such as what is being analyzed here. Here I use a different method meant to conservatively note points of change using similar principles. First, the difference is calculated between each trial position and the position before it. Then, starting at the 3rd trial position, we take the absolute value of the mean difference for each of the trials +/-1 trial positions from this trial. The absolute value is used as we are looking for changes in curvature, and an equal change on both sides of the trial we are analysing may be masked if not using absolute values. Inflection points are now chosen as any point greater than or equal to a predefined cut-off value. A cut-off value must be chosen for changes in slope to have granular enough sensitivity to not be triggered by small changes expected of variance, but not so insensitive as to say that slopes were not changing on curved lines. In this case, a cut-off value of 1 was chosen for detection. This value was shown on test data to not mark linear functions, while marking a sample sine wave with an amplitude comparable to the grand average P300 amplitude reliably and only in peak areas. A sample of this function created in R (R. C. Team, 2016) is shown in *Supplementary Figure 31.* Inflection Point R Script.

4.3 Results

Grand Average Data

An examination of the oddball by way of a difference between the odd stimulus minus the control (*Figure 22*) showed a typical oddball response ($t(48) = 7.32$, $p < 0.001$, $d = 0.70$, $1\text{-}\beta = 0.69$), with a peak latency of 250ms, however with a very low mean voltage of 0.3μV [0.22, 0.38] (*Figure 23*).

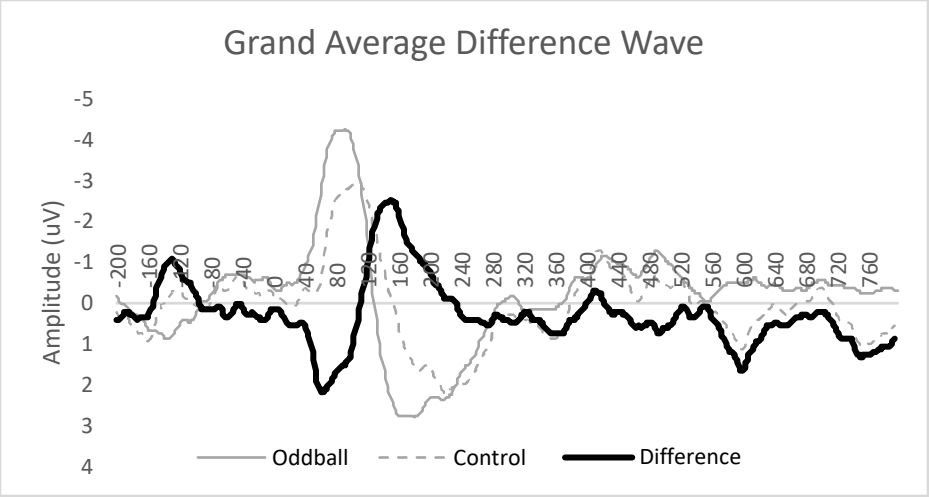

Figure 22. Auditory Oddball Grand Average. Shown is the grand average waveforms of both the oddball and control stimuli. A difference wave was constructed by subtracting the control waveform from the oddball. Analysis was performed +/-25ms about the mean peak of 250ms.

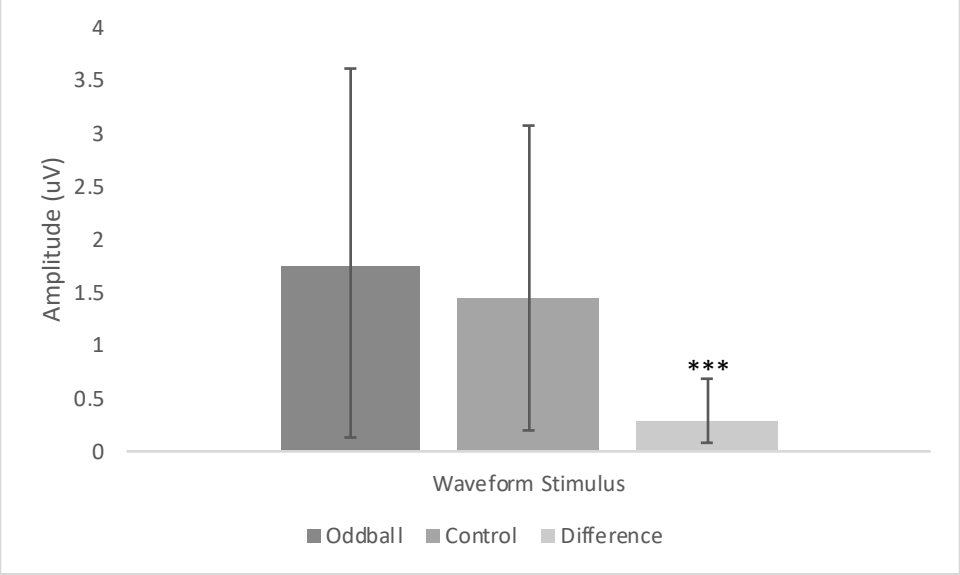

Figure 23. Auditory Oddball Peak Mean Comparison. Whiskers represent 95% CI. Peak mean was calculated from 250ms +/-25ms as calculated from the grand average waveforms. A difference was found between the oddball and control P300 stimulus.

Trial-By-Trial Data

Figure 24. Peak Amplitude by Trial. Peak P300 amplitude by trial separated by stimulus type in the auditory oddball task. Each stimulus begins with a comparable neural response, however by the 3rd trial the P300 amplitude scales to stimulus frequency. Shaded areas are 95% CI, vertical bar is inflection point.

4.4 Discussion

The auditory task overlaid with this experiment produced a small but significant oddball effect, which runs counter to previous literature which states that the P300 should be elicited only on tasks which are attended to by the participant (Donchin, 1981; Patel & Azzam, 2005; Picton, 1992). This is surprising, given that visual inspection of the waveforms shows the potential that a P300 ERP component is only minimally produced by both stimuli (*Figure 22*). However, as had been noted previously auditory oddball paradigms are more reliable than visually based oddballs in unattended tasks (Polich, 1989; Ritter et al., 1968; Squires et al., 1977).

Given the relatively low power of this comparison in regards to its moderate effect size, I am inclined to conclude that this result is likely due to low signal-to-noise ratio and sample size issues (Cohen, 1988; Ioannidis, 2005; Smith, 2004). These results are promising for the potential

use of unattended, non-intrusive stimulus to probe various aspects of brain function; however, this must be tested further. It would not be unheard of for a completely unattended stimulus to produce a reliable P300 effect. While the P3b subcomponent typically requires that a stimulus be task-relevant, the earlier P3a subcomponent is more easily produced by auditory stimuli as they are still unexpected in the participants environment (Delplanque, Silvert, Hot, & Sequeira, 2005; Donchin, 1981; Polich, 1989, 2003, 2007). Polich (1989) used a passive paradigm in that participants were not required to respond to the target tone being heard. These participants were still left in such a scenario that they were still undistracted, so they may have been attending to the stimulus more than participants in the current experiment. The distracted nature of participants in this paradigm may explain the non-typical nature of the P300 shown here as compared to other auditory paradigms.

The most exciting aspect of this study is the trial-by-trial analysis. While there was little data to work with as the trials were not broken into many smaller blocks, there is still some interesting insights in the first twelve exposures of the experiment. As before, a distinct trend can be seen as the first few exposures to a stimulus lead to a steadily increasing P300 before plateauing after the third trial of the experiment. This could indicate that the implicit learning processes shown previously are indeed taking place even in an unattended task.

There is no behavioural data to compare in this task paradigm, but it is interesting to note that the trial-by-trial amplitude in this design reached an inflection point quite early. This would hamper earlier statements of the LC-NE system potentially relegating stimuli to other cortical areas based on novelty after a saturation point. Or perhaps more simply there is a differing response to auditory unattended stimuli than to visual and attended stimuli, as supported by previous literature (Finoia et al., 2015; Squires et al., 1977). The low amplitude of the P300

produced in this paradigm may be indicative that there is little LC-NE engagement as well (Nieuwenhuis, Aston-Jones, et al., 2005), as the task demands are very low when the participant is not tasked with learning the probabilities outright.

Regardless, it is interesting that the prediction error driven patterns are visible in such an environment where the participant is not tasked with attending the stimuli. This suggests that the implicit learning process at hand is truly not a consciously effortful task. The systems involved in learning such as the medial-frontal, or more likely the parietal areas I have spoken of, may be participating in task engagement on some high level beyond participant awareness.

It would be fruitful to repeat this task paradigm, but to alter the auditory oddball frequencies over time to assess if this system is truly able to index and provide predictions of the auditory cues when unattended. It may be simply that we have captured the "growing pains" of hearing the first few auditory cues interrupting the participants focus, though the presence of a P300 in the grand average suggests that the scaling effect measured remains even further into the experiment.

Chapter Five: Limitations and Discussion

5.1 Summary

Experiment one produced a scenario in which small predictions had to be made very frequently. The intriguing results showed that the P300 ERP component was in fact slowly scaling its amplitude in line with participants' learning processes. The eventual scaling of P300 amplitude to stimulus frequency was expected, however the process leading up to it was of great interest. Further analysis showed that this scaling was highly correlated to the learning behaviours taking place. Indeed, evidence points to the P300 being a potential indicator of prediction errors taking place regardless of the presence of feedback.

While experiment one suggested that the P300 amplitude was scaling as learning processes moved forward, it did not definitively prove that implicit learning was the sole affecter of the P300. It may very well have been that the P300 amplitude was scaling to some other factor in the experiment such as stimulus novelty, or a process related to the block reset and break. To counter this limitation, Experiment 2 incorporates a frequency switch midway through half of the blocks. By having a mid-task change in frequency, it was expected that the P300 should reset in amplitude and again have to scale to stimulus frequency. This would be in line with the theory that P300 scaling is tied to implicit learning processes. In addition, frequency pairings are altered such that a 50/50 frequency pair is sometimes shown. This would elucidate if the implicit learning process will eventually scale to recognize that the task is not providing clear evidence accumulation, but still show similar behaviour leading up to this point.

Experiment two further solidified the findings of experiment one by showing that when stimulus frequencies were changed without the participants knowledge the P300 amplitude would again scale to the new stimulus frequency. Experiment Two showed results as expected of

a prediction error driven system. After switching stimulus probabilities mid-block, P300 amplitude scaled to the new stimulus probabilities rapidly. Further, amplitude changes after the switch were similar to amplitudes when the block had always contained those frequencies. This is to say, P300 amplitudes for 25/75% stimuli after the switch were similar to those found in the block that was entirely 25/75%. It is worth noting as well that similar scaling from near-zero microvolt P300 amplitude up to an amount relative to stimulus frequency is once more present in this data set. Amplitudes did not reset to zero after a block switch, but rather shifted over several trials from their current value to one representative as the new frequencies. This is suggestive that prediction errors are guiding a change in amplitude until new frequencies are learned by the participant.

This is supported by behavioural data showing participant accuracy steadily climbing in learnable environments, for example in Block 1 accuracy climbs when moved from a 50/50 condition to a 25/75 condition without informing the participant, and Block 3 shows steadily increasing accuracy throughout its 25/75 condition.

A positive result was found in the detection of the auditory oddball laid over a complex task in Experiment 3. Despite the participants full engagement in the task at hand and the cognitive load it must require to perform accurately, a reliable P300 was produced. Trial-by-trial analysis again provided the most novel insights in this task. A distinct trend can be seen as the first few exposures to a stimulus lead to a steadily increasing P300 before plateauing after the third trial of the experiment. It appears that the implicit learning mechanism at play in Experiments 1 and 2 are at play again here in a fully passive task.

Inflection Points

The inflection point of the P300 amplitude is a difficult issue in these experiments, and perhaps the area most deserving of follow-up experimentation. In each experiment, the P300 amplitude would scale in unison for all stimulus frequencies. This initially lead me to believe that some sort of novelty system was at hand aiding in P300 amplitude scaling. Most likely, I had concluded that this would be part of the LC-NE system as proposed by Nieuwenhuis' LC-P3 hypothesis (Nieuwenhuis, Aston-Jones, et al., 2005). However, at the conclusion of these experiments we can note another pattern: the inflection point is not the same in these paradigms. Indeed, they are not reliant on stimulus frequency (hence the inflection point is typically for all stimulus frequencies at once in a given experiment) and seem to have little relation between experiments even when stimulus probabilities are comparable. If a single system involved in assessing only stimulus frequency over trials were responsible for P300 amplitude scaling, we would expect that this effect should be far more consistent. The large gap in inflection point between Experiment 1's trial 9 and Experiment 3's trial 3 may be explained away by the differing stimulus presentation (auditory vs. visual) or the differing engagement (passive vs. active), but it still begs the question as to why there exists such a difference. If the LC-P3 hypothesis is correct, I would state that it is quite reasonable for these differences to exist. By the LC-P3 hypothesis, the LC-NE system acts to enhancement the processing of stimuli. In a passive task, with non-complex stimuli, it would be reasonable for this system to be engaged far less. As a result, learning may be slowed considerably.

Indeed, another study (Wronka et al., 2008) has added evidence for the P300 to behave in this way by directly comparing passive and active auditory P300 ERP components in a variety of stimulus pairings. They found that the passive oddball paradigm produced far lower, but still

significant amplitude responses. Counter to this experiment, they found that only the active P3 condition produced a parietal response, whereas here we find it more parietally at electrode Pz.

Prediction Errors

I have noted the presence of evidence for prediction error mechanisms, the crux of this inflection point issue is critical. First, a refresher on the typical prediction error formula (Glimcher, 2011; Schultz, 2016), as it was first described in relation to Pavlov's dogs:

$$A_{next-trial} = A_{last-trial} + \alpha(R_{current-trial} - A_{last-trial})$$

Equation 2. Prediction Error Formula

This basic concept forms the basis of prediction errors in learning: you have an expectation of an outcome based on the previous trial, that is updated by the current trial's outcome. This is the piece shown in brackets, typically this is the "error" that is used to learn from. Alpha is shown as a learning rate between 0 and 1. In the experiments presented here, we see P300 amplitudes that correlate to behavioural processes. Note that the prediction error formula shown above utilized scalar rewards, and here this framework is used to estimate stimulus frequency rather than reward. These P300 amplitudes are following similar patterns to what is produced with the scalar reward prediction error formula. The catch comes in the fact that each stimulus probability is either scaling with a unique alpha value, or alternatively is experiencing much larger prediction errors. Given how all stimuli in an experiment reach inflection at the same trial position, I am inclined to say the latter. If there are differing alpha values at play, they would all be similarly high as demonstrated in *Figure 25*, which

demonstrates standard prediction error curves which vary only in alpha assuming that the P300 scales in step with observed arbitrary stimulus frequency that is matched for all four alpha values.

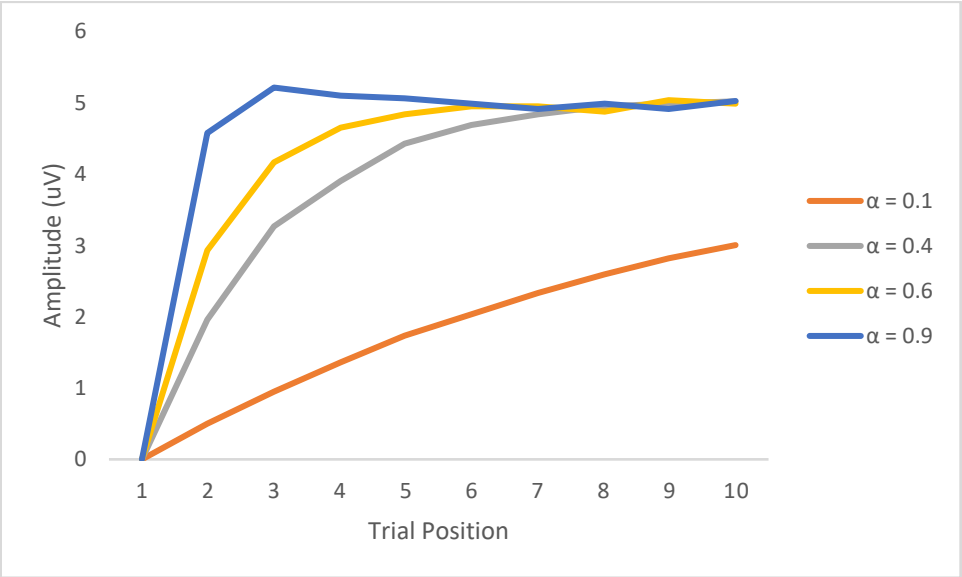

Figure 25. Theoretical Prediction Error Curves. Here, hypothetical prediction error paths are presented with varying alpha values. The ten trials are simulated from 500 "participants" and averaged.

The ideal fit shown in *Figure 25* is promising that the inflection point scaling being measured is likely in line with other prediction error work (Bray & O'Doherty, 2007; Glimcher, 2011; Schultz, 2016). The fact that the ideal prediction error curve here does not use a "two-system" model as I have previously proposed proves interesting. It may be that the two-system explanation is not necessary to provide a neural basis for the P300 scaling effect. P300 scaling is in line with other work such that the inflection point is merely a mathematical artifact of a standard learning process (Glimcher, 2011; Schultz, 2016). However, I would still posit that the migration of neural generators from the prefrontal cortex to a more parietal source is indicative

that even this one system model involves different neural processes at differing stages, and that the lack of consistency of the P300 scaling between experiments and stimulus frequencies warrants investigation.

Neural Generators

There is little evidence in these proposed experiments to disrupt current thinking on neural generators of the P300. In fact, current thinking has guided the decisions behind this work. It is still very much posterior parietal areas tied to P3b generation, with early trials being more tied to medial-frontal activity (Dien et al., 2003; Nieuwenhuis, Aston-Jones, et al., 2005; Soltani & Knight, 2000). As there has been no source analysis in this process, I cannot make any strong claims as to the true sources of the P300 component being measured. Building on previous work however, it is reasonable to claim that the LC-NE system is playing a role in initially aiding the interpretation of the stimuli being presented (Nieuwenhuis, Aston-Jones, et al., 2005). As the P3 topography shifts over trials, it is likely a mediated act by this system as it begins to play a lesser role in the processing of each stimulus. As the stimuli is merely updated in context, the P3b plays a more central role in low level processing as trials continue (Dien et al., 2003; Knight et al., 1995; Mecklinger & Ullsperger, 1995; Soltani & Knight, 2000).

5.2 Limitations and Future Directions

One of the key limitations across this study is the issue of signal-to-noise ratio. Experiment 1 and 2 rely heavily on trial by trial analysis of an ERP component. A typical study will employ the basic principle of ERP research; collect many trials of each stimulus in each condition of interest and average them (Luck, 2014). By taking a close look at each individual trial and the trends between them, it is effectively parsing data down to a handful of segments per

trial, even in a large scale repetitive experiment. For example, Experiment 2 contains 20 blocks of 80 trials for 2 stimulus types. With a typical oddball paradigm this would provide a total of 1,600 trials from which to perform artifact rejection and segment down to 1,200 control and 400 oddball stimuli *per participant* in a typical oddball paradigm. Even when broken down into the four block conditions (300 control, 100 oddball each), there is a vast amount of data available which is then multiplied by every participant available. In this experiment, the large number of trials and blocks was chosen to intentionally try to counteract this issue, as typically only approximately 30 trials over 10 blocks (300 trials total) would produce an oddball such as in Experiment 3's visual oddball paradigm, some experiments using even half that (Coenen, 1995; Krigolson et al., 2017; Picton, 1992; Polich & Kok, 1995). This however vastly underestimated the need for more data as even with this increase a perfect participant and EEG data set would only produce 20 segments from which to average for each trial position. This is before subdividing by stimulus type and block condition. From this, we can conclude that these 80 trial studies would require nearly 100,000[1] trials to produce comparable data quality to a standard 300 trial oddball paradigm in each trial position.

While it may not be necessary to have model data for each and every trial position for the purpose of these comparisons, this does highlight the gulf in either trial or participant count needed when performing trial-by-trial analysis which is inherently noisy (Blankertz, Lemm, Treder, Haufe, & Müller, 2011; Jung et al., 1999). Since it is not reasonably possible to increase participant count by a factor of 320, and participants are difficult to engage for the time necessary for a 1,200-block task, it is reasonable that those wishing to perform trial-by-trial

[1] A typical oddball paradigm produces 75 odd stimulus trials and 225 controls. By multiplying this out across 80 trials, and 4 block conditions as in Experiment 2, we find ourselves needing 96,000 trials!

analysis studies be expected to either manage the SNR ratio in analysis as best as possible or to simplify their experimental designs. In the case of this study simplification of design would have introduced more limitations in the claims being made so I believe recruiting more participants to bring the noise and power to an acceptable level would be a productive future direction.

The studies at hand also contain limitations very typical to the field of psychology in that all participants were university aged participants. This is a limitation endemic to the field, and is especially problematic in this case due to the fact that explicit reinforcement learning, and so potentially the implicit mechanisms at play here, have been shown to decline with age in both performance and neural correlates (Campagne, Pebayle, & Muzet, 2004; Picton, 1992; Sojitra, Lerner, Petok, & Gluck, 2018). This limits the generalizability of these findings until more data is collected on a wider demographic group. However, if these findings follow the patterns of other studies into learning mechanisms and aging, useful predictions may still be made from this base data.

One of the key future directions necessary for these studies is to expand on experiment 3. It would be fruitful to perform a similar probability switching design as in experiment 2, to assess the effect of engagement on the ability for context to be updated and the stimuli attended to once more.

5.3 Conclusions

Our results suggest that the amplitude of the P300 reflects an expectancy driven prediction error. Specifically, a reinforcement learning process is taking place based on implicit violations generating intrinsic reward as opposed to explicit reward. Additionally, the data suggests that some further mechanism is at play during the initial trials that moves neural

systems beyond novelty processing and initial stimulus categorization (Nieuwenhuis, Aston-Jones, et al., 2005; Polich, 2007).

Further experimentation supported the presence of implicit learning by requiring new stimuli to be learned without warning, and further strengthened the relationship between P300 amplitude and behaviour. The final experiment verified that this mechanism is in action and measurable during a passive auditory task. This could prove to be a useful cognitive function test design in later studies.

References

Aidman, E., Chadunow, C., Johnson, K., & Reece, J. (2015). Real-time driver drowsiness feedback improves driver alertness and self-reported driving performance. *Accident Analysis & Prevention*, *81*, 8–13. doi: 10.1016/j.aap.2015.03.041

Artigas, A. A., & Prados, J. (2017). Perceptual learning transfer in an appetitive Pavlovian task. *Learning & Behavior*, *45*(2), 115–123. doi: 10.3758/s13420-016-0245-y

Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annual Review of Neuroscience*, *28*, 403–50. doi: 10.1146/annurev.neuro.28.061604.135709

Atkinson, R. C., & Estes, W. K. (1962). *Stimulus sampling theory (No. 48)*. Applied Mathematics and Statistics Laboratories, Stanford University: Institute for Mathematical Studies in the Social Science.

Azizian, A., & Polich, J. (2007). Evidence for attentional gradient in the serial position memory curve from event-related potentials. *Journal of Cognitive Neuroscience*, *19*(12), 2071–81. doi: 10.1162/jocn.2007.19.12.2071

Azzopardi, P., & Cowey, A. (1998). Blindsight and Visual Awareness. *Consciousness and Cognition*, *7*(3), 292–311. doi: 10.1006/ccog.1998.0358

Barch, D. M., Carter, C. S., Gold, J. M., Johnson, S. L., Kring, A. M., MacDonald, A. W., … Strauss, M. E. (2017). Explicit and implicit reinforcement learning across the psychosis spectrum. *Journal of Abnormal Psychology*, *126*(5), 694–711. doi: 10.1037/abn0000259

Blankertz, B., Lemm, S., Treder, M., Haufe, S., & Müller, K.-R. (2011). Single-trial analysis and classification of ERP components — A tutorial. *NeuroImage*, *56*(2), 814–825. doi: 10.1016/j.neuroimage.2010.06.048

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436. doi: 10.1163/156856897X00357

Braver, T. S., & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In *Control of cognitive processes Attention and performance XVIII* (Vol. XVIII, pp. 713–737). doi: 10.1016/S0165-0173(03)00143-7

Bray, S., & O'Doherty, J. P. (2007). Neural Coding of Reward-Prediction Error Signals During Classical Conditioning With Attractive Faces. *Journal of Neurophysiology*, *97*(4), 3036–3045. doi: 10.1152/jn.01211.2006

Bromberg-Martin, E. S., Matsumoto, M., & Hikosaka, O. (2010). Dopamine in Motivational Control: Rewarding, Aversive, and Alerting. *Neuron*, *68*(5), 815–834. doi: 10.1016/j.neuron.2010.11.022

Bronshtein, I. N., Semendyayev, K. A., Musiol, G., & Mühlig, H. (2015). *Handbook of mathematics, sixth edition*. *Handbook of Mathematics, Sixth Edition*. doi: 10.1007/978-3-662-46221-8

Brown, J. W., & Braver, T. S. (2005). Learned predictions of error likelihood in the anterior cingulate cortex. *Science (New York, N.Y.)*, *307*(5712), 1118–21. doi: 10.1126/science.1105783

Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*, *13*(6), 407–420. doi: 10.1038/nrn3241

Campagne, A., Pebayle, T., & Muzet, A. (2004). Correlation between driving errors and vigilance level: Influence of the driver's age. *Physiology and Behavior*, *80*(4), 515–524.

doi: 10.1016/j.physbeh.2003.10.004

Chansky, N. M. (1960). Learning : a Function of Schedule and Type of Feedback, 1960.

Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: news from the front. *Trends in Cognitive Sciences*, *2*(10), 406–16. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/21227256

Coenen, A. M. L. (1995). Neuronal activities underlying the electroencephalogram and evoked potentials of sleeping and waking: Implications for information processing. *Neuroscience and Biobehavioral Reviews*, *19*(3), 447–463. doi: 10.1016/0149-7634(95)00010-C

Cohen, J. D. (1988). *Statistical Power Analysis for the Behavioral Sciences*. (L. Erlbaum, Ed.) (2nd ed.). Hillsdale, N.J. doi: 10.1016/C2013-0-10517-X

Cohen, M. X., & Ranganath, C. (2007). Reinforcement Learning Signals Predict Future Decisions. *Journal of Neuroscience*, *27*(2), 371–378. doi: 10.1523/JNEUROSCI.4421-06.2007

Conroy, M. A., & Polich, J. (2007). Normative Variation of P3a and P3b from a Large Sample. *Journal of Psychophysiology*, *21*(1), 22–32. doi: 10.1027/0269-8803.21.1.22

Contreras, D., & Steriade, M. (1995). Cellular basis of EEG slow rhythms: a study of dynamic corticothalamic relationships. *The Journal of Neuroscience*, *15*(1), 604–622. doi: 10.1523/JNEUROSCI.15-01-00604.1995

Croft, R. J., Gonsalvez, C. J., Gabriel, C., & Barry, R. J. (2003). Target-to-target interval versus probability effects on P300 in one- and two-tone tasks. *Psychophysiology*, *40*(3), 322–328. doi: 10.1111/1469-8986.00036

DellaBadia Jr, J., Bell, W. L., Keyes Jr, J. W., Mathews, V. P., & Glazier, S. S. (2002). Assessment and cost comparison of sleep-deprived EEG, MRI and PET in the prediction of

surgical treatment for epilepsy. *Seizure*, *11*(5), 303–309. doi: 10.1053/seiz.2001.0648

Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21. doi: 10.1016/j.jneumeth.2003.10.009

Delplanque, S., Silvert, L., Hot, P., & Sequeira, H. (2005). Event-related P3a and P3b in response to unpredictable emotional stimuli. *Biological Psychology*, *68*(2), 107–120. doi: 10.1016/j.biopsycho.2004.04.006

Desjardins, A. E., Kiehl, K. A., & Liddle, P. F. (2001). Removal of Confounding Effects of Global Signal in Functional MRI Analyses. *NeuroImage*, *13*(4), 751–758. doi: 10.1006/nimg.2000.0719

Dien, J., Spencer, K. M., & Donchin, E. (2003). Localization of the event-related potential novelty response as defined by principal components analysis. *Cognitive Brain Research*, *17*(3), 637–650. doi: 10.1016/S0926-6410(03)00188-5

Dienes, Z., & Berry, D. C. (1997). Implicit learning: Below the subjective threshold. *Psychonomic Bulletin & Review*, *4*(1), 3–23. doi: 10.3758/BF03210769

Donchin, E. (1981). Presidential address, 1980. Surprise!...Surprise? *Psychophysiology*. doi: 10.1111/j.1469-8986.1981.tb01815.x

Donkers, F. C. L., & van Boxtel, G. J. M. (2005). Mediofrontal Negativities to Averted Gains and Losses in the Slot-Machine Task. *Journal of Psychophysiology*, *19*(4), 256–262. doi: 10.1027/0269-8803.19.4.256

Duncan-Johnson, C. C. (1981). P300 latency: A new metric for information processing. *Psychophysiology*, *18*, 207–215.

Duncan-Johnson, C. C., & Donchin, E. (1977). On quantifying surprise: the variation of event-

related potentials with subjective probability. *Psychophysiology*, *14*(5), 456–67. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/905483

Duncan-Johnson, C. C., & Donchin, E. (1982). The P300 component of the event-related brain potential as an index of information processing. *Biological Psychology*, *14*(1–2), 1–52. doi: 10.1016/0301-0511(82)90016-3

Erhel, S., & Jamet, E. (2013). Digital game-based learning: Impact of instructions and feedback on motivation and learning effectiveness. *Computers and Education*, *67*, 156–167. doi: 10.1016/j.compedu.2013.02.019

Estes, W. K. (1950). Toward a statistical theory of learning. *Psychological Review*, *57*(2), 94–107. doi: 10.1037/h0058559

Estes, W. K. (1974). Learning theory and intelligence. *American Psychologist*, *29*(10), 740–749. doi: 10.1037/h0037458

Fabiani, M., Karis, D., & Donchin, E. (1986). P300 and recall in an incidental memory paradigm. *Psychophysiology*, *23*(3), 298–308. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/3749410

Fabiani, M., Karis, D., & Donchin, E. (1990). Effects of mnemonic strategy manipulation in a Von Restorff paradigm. *Electroencephalography and Clinical Neurophysiology*, *75*(2), 22–35. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1688770

Finoia, P., Mitchell, D. J., Hauk, O., Beste, C., Pizzella, V., & Duncan, J. (2015). Concurrent brain responses to separate auditory and visual targets. *Journal of Neurophysiology*, jn.01050.2014. doi: 10.1152/jn.01050.2014

Fiorillo, C. D., Tobler, P. N., & Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science*, *299*(5614), 1898–1902. doi:

10.1126/science.1077349

Foerde, K., Knowlton, B. J., & Poldrack, R. A. (2006). Modulation of competing memory systems by distraction. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(31), 11778–11783. doi: 0602659103 [pii]

Foti, D., Hajcak, G., & Dien, J. (2009). Differentiating neural responses to emotional pictures: Evidence from temporal-spatial PCA. *Psychophysiology*, *46*(3), 521–530. doi: 10.1111/j.1469-8986.2009.00796.x

Gill, D. L., & Martens, R. (1975). The informational and motivational influence of social reinforcement on motor performance. *Journal of Motor Behavior*, *7*(3), 171–82. doi: 10.1080/00222895.1975.10735031

Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences*, *108*(Supplement_3), 15647–15654. doi: 10.1073/pnas.1014269108

Gonsalvez, C. J., & Polich, J. (2002). P300 amplitude is determined by target-to-target interval. *Psychophysiology*, *39*(3), S0048577201393137. doi: 10.1017/S0048577201393137

Hajcak, G., Moser, J. S., Holroyd, C. B., & Simons, R. F. (2007). It's worse than you thought: The feedback negativity and violations of reward prediction in gambling tasks. *Psychophysiology*, *44*(6), 905–912. doi: 10.1111/j.1469-8986.2007.00567.x

Haruno, M., & Kawato, M. (2006). Heterarchical reinforcement-learning model for integration of multiple cortico-striatal loops: fMRI examination in stimulus-action-reward association learning. *Neural Networks*, *19*(8), 1242–1254. doi: 10.1016/j.neunet.2006.06.007

Hassall, C. D., MacLean, S., & Krigolson, O. E. (2014). Hierarchical Error Evaluation: The Role of Medial-Frontal Cortex in Postural Control. *Journal of Motor Behavior*, *46*(6), 381–387.

doi: 10.1080/00222895.2014.918021

Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis. of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review-New York-*, *109*(4), 679–709. doi: 10.1037//0033-295X.109.4.679

Holroyd, C. B., & Krigolson, O. E. (2007). Reward prediction error signals associated with a modified time estimation task. *Psychophysiology*, *44*(6), 913–917. doi: 10.1111/j.1469-8986.2007.00561.x

Holroyd, C. B., Krigolson, O. E., Baker, R., Lee, S., & Gibson, J. (2009). When is an error not a prediction error? An electrophysiological investigation. *Cognitive, Affective, & Behavioral Neuroscience*, *9*(1), 59–70. doi: 10.3758/CABN.9.1.59

Holroyd, C. B., Pakzad-Vaezi, K. L., & Krigolson, O. E. (2008). The feedback correct-related positivity: Sensitivity of the event-related brain potential to unexpected positive feedback. *Psychophysiology*, *45*(5), 688–697. doi: 10.1111/j.1469-8986.2008.00668.x

Inman, R. A., & Pearce, J. M. (2018). The discrimination of magnitude: A review and theoretical analysis. *Neurobiology of Learning and Memory*. doi: 10.1016/j.nlm.2018.03.020

Ioannidis, J. P. A. . (2005). Why most published research findings are false. *PLoS Medicine*, *2*(8), e124. doi: 10.1371/journal.pmed.0020124

Ishikura, T. (2008). Reduced Relative Frequency of Knowledge of Results without Visual Feedback in Learning a Golf-Putting Task. *Perceptual and Motor Skills*, *106*(1), 225–233. doi: 10.2466/pms.106.1.225-233

Isreal, J. B., Chesney, G. L., Wickens, C. D., & Donchin, E. (1980). P300 and Tracking Difficulty: Evidence For Multiple Resources in Dual-Task Performance. *Psychophysiology*, *17*(3), 259–273. doi: 10.1111/j.1469-8986.1980.tb00146.x

Jackson, A. F., & Bolger, D. J. (2014). The neurophysiological bases of EEG and EEG measurement: A review for the rest of us. *Psychophysiology*, *51*(11), 1061–1071. doi: 10.1111/psyp.12283

Jeon, Y.-W., & Polich, J. (2003). Meta-analysis of P300 and schizophrenia: patients, paradigms, and practical implications. *Psychophysiology*, *40*(5), 684–701. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/14696723

Jessup, R. K., Busemeyer, J. R., & Brown, J. W. (2010). Error Effects in Anterior Cingulate Cortex Reverse when Error Likelihood Is High. *Journal of Neuroscience*, *30*(9), 3467–3472. doi: 10.1523/JNEUROSCI.4130-09.2010

Jung, T. P., Makeig, S., Westerfield, M., Townsend, J., Courchesne, E., & Sejnowski, T. J. (1999). *Analyzing and visualizing single-trial event-related potentials. Advances in neural information processing systems* (11th ed.).

Kalisch, R., Elbel, G.-K., Gössl, C., Czisch, M., & Auer, D. P. (2001). Blood Pressure Changes Induced by Arterial Blood Withdrawal Influence Bold Signal in Anesthesized Rats at 7 Tesla: Implications for Pharmacologic MRI. *NeuroImage*, *14*(4), 891–898. doi: 10.1006/nimg.2001.0890

Karis, D., Fabiani, M., & Donchin, E. (1984). "P300" and memory: Individual differences in the von Restorff effect. *Cognitive Psychology*, *16*(2), 177–216. doi: 10.1016/0010-0285(84)90007-0

Kaseda, Y., Jiang, C., Kurokawa, K., Mimori, Y., & Nakamura, S. (1998). Objective evaluation of fatigue by event-related potentials. *Journal of the Neurological Sciences*, *158*(1), 96–100.

Kass, R. E., & Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical Association*, *90*(430), 773. doi: 10.2307/2291091

Katayama, J., & Polich, J. (1996). P300, probability, and the three-tone paradigm. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, *100*(6), 555–562. doi: 10.1016/S0168-5597(96)95171-0

Käthner, I., Wriessnegger, S. C., Müller-Putz, G. R., Kübler, A., & Halder, S. (2014). Effects of mental workload and fatigue on the P300, alpha and theta band power during operation of an ERP (P300) brain–computer interface. *Biological Psychology*, *102*(1), 118–129. doi: 10.1016/j.biopsycho.2014.07.014

Knight, R. T., Grabowecky, M. F., & Scabini, D. (1995). Role of human prefrontal cortex in attention control. *Advances in Neurology*, *66*, 21-34; discussion 34–6. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/7771302

Kok, A. (1997). Event-related-potential (ERP) reflections of mental resources: A review and synthesis. In *Biological Psychology* (Vol. 45, pp. 19–56). doi: 10.1016/S0301-0511(96)05221-0

Krigolson, O. E., Hassall, C. D., & Handy, T. C. (2014). How We Learn to Make Decisions: Rapid Propagation of Reinforcement Learning Prediction Errors in Humans. *Journal of Cognitive Neuroscience*, *26*(3), 635–644. doi: 10.1162/jocn_a_00509

Krigolson, O. E., & Holroyd, C. B. (2007). Hierarchical error processing: Different errors, different systems. *Brain Research*, *1155*(1), 70–80. doi: 10.1016/j.brainres.2007.04.024

Krigolson, O. E., Williams, C. C., Norton, A., Hassall, C. D., & Colino, F. L. (2017). Choosing MUSE: Validation of a low-cost, portable EEG system for ERP research. *Frontiers in Neuroscience*, *11*(MAR), 1–10. doi: 10.3389/fnins.2017.00109

Lal, S. K. L., & Craig, A. (2005). Reproducibility of the spectral components of the electroencephalogram during driver fatigue. *International Journal of Psychophysiology*,

55(2), 137–143. doi: 10.1016/j.ijpsycho.2004.07.001

Luck, S. J. (2014). A Closer Look at Averaging: Convolution, Latency Variability, and Overlap. *An Introduction to the Event-Related Potential Technique*, Online Chapter 11.

MATLAB. (2016). Natick, Massachusetts: The MathWorks Inc. doi: 10.1007/s10766-008-0082-5

McFall, R. M., & Treat, T. A. (1999). Quantifying the Information Value of Clinical Assessments with Signal Detection Theory. *Annual Review of Psychology*, *50*(1), 215–241. doi: 10.1146/annurev.psych.50.1.215

Mecklinger, A., & Ullsperger, P. (1995). The P300 to novel and target events: a spatio- temporal dipole model analysis. *Neuroreport*.

Michas, I. C., & Berry, D. C. (1994). Implicit and explicit processes in a second-language learning task. *European Journal of Cognitive Psychology*, *6*(4), 357–381. doi: 10.1080/09541449408406520

Miltner, W. H. R., Braun, C. H., & Coles, M. G. H. (1997). Event-Related Brain Potentials Following Incorrect Feedback in a Time-Estimation Task: Evidence for a "Generic" Neural System for Error Detection. *Journal of Cognitive Neuroscience*, *9*(6), 788–798. doi: 10.1162/jocn.1997.9.6.788

Mirenowicz, J., & Schultz, W. (1996). Preferential activation of midbrain dopamine neurons by appetitive rather than aversive stimuli. *Nature*, *379*(6564), 449–451. doi: 10.1038/379449a0

Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience*, *16*(5), 1936–1947. doi: 10.1.1.156.635

Morris, S. E., Heerey, E. A., Gold, J. M., & Holroyd, C. B. (2008). Learning-related changes in

brain activity following errors and performance feedback in schizophrenia. *Schizophrenia Research*, *99*(1–3), 274–285. doi: 10.1016/j.schres.2007.08.027

Mosher, J. C., & Leahy, R. M. (1998). Recursive MUSIC: A framework for EEG and MEG source localization. *IEEE Transactions on Biomedical Engineering*, *45*(11), 1342–1354. doi: 10.1109/10.725331

Mühlberger, C., Angus, D. J., Jonas, E., Harmon-Jones, C., & Harmon-Jones, E. (2017). Perceived control increases the reward positivity and stimulus preceding negativity. *Psychophysiology*, *54*(2), 310–322. doi: 10.1111/psyp.12786

Murata, A., Uetake, A., & Takasawa, Y. (2005). Evaluation of mental fatigue using feature parameter extracted from event-related potential. *International Journal of Industrial Ergonomics*, *35*(8), 761–770. doi: 10.1016/j.ergon.2004.12.003

Murphy, R., Byrom, N., & Msetfi, R. M. (2017). The problem with explaining symptoms: The origin of biases in causal processing. *European Journal for Person Centered Healthcare*, *5*(3), 344. doi: 10.5750/ejpch.v5i3.1318

Nieuwenhuis, S., Aston-Jones, G., & Cohen, J. D. (2005). Decision making, the P3, and the locus coeruleus--norepinephrine system. *Psychological Bulletin*, *131*(4), 510–532. doi: 10.1037/0033-2909.131.4.510

Nieuwenhuis, S., Heslenfeld, D. J., Alting von Geusau, N. J., Mars, R. B., Holroyd, C. B., & Yeung, N. (2005). Activity in human reward-sensitive brain areas is strongly context dependent. *NeuroImage*, *25*(4), 1302–1309. doi: 10.1016/j.neuroimage.2004.12.043

Nieuwenhuis, S., Holroyd, C. B., Mol, N., & Coles, M. G. H. (2004). Reinforcement-related brain potentials from medial frontal cortex: origins and functional significance. *Neuroscience and Biobehavioral Reviews*, *28*(4), 441–8. doi:

10.1016/j.neubiorev.2004.05.003

O'Doherty, J. P., Cockburn, J., & Pauli, W. M. (2017). Learning, Reward, and Decision Making. *Annual Review of Psychology*, *68*(1), 73–100. doi: 10.1146/annurev-psych-010416-044216

O'Doherty, J. P., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R. J. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science (New York, N.Y.)*, *304*(5669), 452–4. doi: 10.1126/science.1094285

Pacchiarini, N., Fox, K., & Honey, R. C. (2017). Perceptual learning with tactile stimuli in rodents: Shaping the somatosensory system. *Learning & Behavior*, *45*(2), 107–114. doi: 10.3758/s13420-017-0269-y

Patel, S. H., & Azzam, P. N. (2005). Characterization of N200 and P300: selected studies of the Event-Related Potential. *International Journal of Medical Sciences*, *2*(4), 147–54. doi: 10.7150/ijms.2.147

Picton, T. W. (1992). The P300 wave of the human event-related potential. *Journal of Clinical Neurophysiology : Official Publication of the American Electroencephalographic Society*, *9*(4), 456–79. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1464675

Polich, J. (1989). P300 from a passive auditory paradigm. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, *74*(4), 312–320. doi: 10.1016/0168-5597(89)90061-0

Polich, J. (1990). P300, Probability, and Interstimulus Interval. *Psychophysiology*, *27*(4), 396–403. doi: 10.1111/j.1469-8986.1990.tb02333.x

Polich, J. (2003). Overview of P3a and P3b. In *Detection of Change: Event-Related Potential and fMRI Findings.* (pp. 83–98). Boston, MA: Kluwer Academic Press.

Polich, J. (2004). Clinical application of the P300 event-related brain potential. *Physical*

Medicine and Rehabilitation Clinics of North America. doi: 10.1016/S1047-9651(03)00109-8

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

Polich, J., & Bondurant, T. (1997). P300 Sequence Effects, Probability, and Interstimulus Interval. *Physiology & Behavior*, *61*(6), 843–849. doi: 10.1016/S0031-9384(96)00564-1

Polich, J., & Kok, A. (1995). Cognitive and biological determinants of P300: an integrative review. *Biological Psychology*, *41*(2), 103–146. doi: 10.1016/0301-0511(95)05130-9

Qin, Y., Sohn, M.-H., Anderson, J. R., Stenger, V. A., Fissell, K., Goode, A., & Carter, C. S. (2003). Predicting the practice effects on the blood oxygenation level-dependent (BOLD) function of fMRI in a symbolic manipulation task. *Proceedings of the National Academy of Sciences of the United States of America*, *100*(8), 4951–4956. doi: 10.1073/pnas.0431053100

Qiu, Y., Tang, Y., Chan, R. C. K., Sun, X., & He, J. (2014). P300 Aberration in First-Episode Schizophrenia Patients: A Meta-Analysis. *PLoS ONE*, *9*(6), e97794. doi: 10.1371/journal.pone.0097794

Ribas-Fernandes, J. J. F., Solway, A., Diuk, C., McGuire, J. T., Barto, A., Niv, Y., & Botvinick, M. M. (2011). A Neural Signature of Hierarchical Reinforcement Learning. *Neuron*, *71*(2), 370–379. doi: 10.1016/j.neuron.2011.05.042

Ritter, W., & Vaughan, H. G. (1969). Averaged Evoked Responses in Vigilance and Discrimination: A Reassessment. *Science*, *164*(3877), 326–328. doi: 10.1126/science.164.3877.326

Ritter, W., Vaughan, H. G., & Costa, L. D. (1968). Orienting and habituation to auditory stimuli:

A study of short terms changes in average evoked responses. *Electroencephalography and Clinical Neurophysiology*, *25*(6), 550–556. doi: 10.1016/0013-4694(68)90234-4

Rodríguez, G., Blair, C. A. J., & Hall, G. (2008). The role of comparison in perceptual learning: effects of concurrent exposure to similar stimuli on the perceptual effectiveness of their unique features. *Learning & Behavior*, *36*(2), 75–81. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/18543708

Rushby, J. A., Barry, R. J., & Doherty, R. J. (2005). Separation of the components of the late positive complex in an ERP dishabituation paradigm. *Clinical Neurophysiology*, *116*(10), 2363–2380. doi: 10.1016/j.clinph.2005.06.008

Salmoni, A. W., Schmidt, R. A., & Walter, C. B. (1984). Knowledge of results and motor learning: a review and critical reappraisal. *Psychological Bulletin*, *95*(3), 355–86. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/6399752

Sambrook, T. D., & Goslin, J. (2015). A neural reward prediction error revealed by a meta-analysis of ERPs using great grand averages. *Psychological Bulletin*, *141*(1), 213–235. doi: 10.1037/bul0000006

Schultz, W. (1998). Predictive Reward Signal of Dopamine Neurons. *Journal of Neurophysiology*, *80*(1), 1–27. doi: 10.1152/jn.1998.80.1.1

Schultz, W. (2016). Dopamine reward prediction error coding. *Dialogues in Clinical Neuroscience*, *18*(1), 23–32. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/27069377

Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of Monkey Dopamine Neurons to Reward and Conditioned Stimuli during Successive Steps of Learning a Delayed Response Task. *The Journal of Neuroscience*, *13*(March), 900–913.

Schultz, W., Dayan, P., & Montague, P. R. (1997). A Neural Substrate of Prediction and Reward. *Science*, *275*(5306), 1593–1599. doi: 10.1126/science.275.5306.1593

Shijian Lu, Cuntai Guan, & Haihong Zhang. (2009). Unsupervised Brain Computer Interface Based on Intersubject Information and Online Adaptation. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, *17*(2), 135–145. doi: 10.1109/TNSRE.2009.2015197

Shyu, W. M., Grosse, E., & Cleveland, W. S. (2017). Local regression models. In *Statistical models in S* (pp. 309–376). Routledge.

Singh, S., Lewis, R. L., Barto, A. G., & Sorg, J. (2010). Intrinsically Motivated Reinforcement Learning: An Evolutionary Perspective. *IEEE Transactions on Autonomous Mental Development*, *2*(2), 70–82. doi: 10.1109/TAMD.2010.2051031

Skinner, B. F. (1948). "Superstition" in the pigeon. *Journal of Experimental Psychology*, *38*(2), 168–172. doi: 10.1037/h0055873

Smith, M. H. (2004). A Sample/Population Size Activity: Is it the sample size of the sample as a fraction of the population that matters? *Journal of Statistics Education*, *12*(2). doi: 10.1080/10691898.2004.11910735

Snyder, E., & Hillyard, S. A. (1976). Long-latency evoked potentials to irrelevant, deviant stimuli. *Behavioral Biology*, *16*(3), 319–331. doi: 10.1016/S0091-6773(76)91447-4

Sojitra, R. B., Lerner, I., Petok, J. R., & Gluck, M. A. (2018). Age affects reinforcement learning through dopamine-based learning imbalance and high decision noise-not through Parkinsonian mechanisms. *Neurobiology of Aging*, *68*, 102–113. doi: 10.1016/j.neurobiolaging.2018.04.006

Soltani, M., & Knight, R. T. (2000). Neural origins of the P300. *Critical Reviews in*

Neurobiology, *14*(3–4), 199–224. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12645958

Song, J., Davey, C., Poulsen, C., Luu, P., Turovets, S., Anderson, E., … Tucker, D. (2015). EEG source localization: Sensor density and head surface coverage. *Journal of Neuroscience Methods*, *256*, 9–21. doi: 10.1016/j.jneumeth.2015.08.015

Spencer, K. M., Dien, J., & Donchin, E. (2001). Spatiotemporal analysis of the late ERP responses to deviant stimuli. *Psychophysiology*, *38*(2), 343–358. doi: 10.1017/S0048577201000324

Spencer, K. M., Vila Abad, E., & Donchin, E. (2000). On the search for the neurophysiological manifestation of recollective experience. *Psychophysiology*, *37*(4), 494–506. doi: 10.1017/S0048577200982210

Squire, L. R., & Zola, S. M. (1996). Structure and function of declarative and nondeclarative memory systems. *Proceedings of the National Academy of Sciences*, *93*(24), 13515–13522. doi: 10.1073/pnas.93.24.13515

Squires, K. C., Petuchowski, S., Wickens, C., & Donchin, E. (1977). The effects of stimulus sequence on event related potentials: A comparison of visual and auditory sequences. *Perception & Psychophysics*, *22*(1), 31–40. doi: 10.3758/BF03206077

Squires, N. K., Squires, K. C., & Hillyard, S. A. (1975). Two varieties of long-latency positive waves evoked by unpredictable auditory stimuli in man. *Electroencephalography and Clinical Neurophysiology*, *38*(4), 387–401. doi: 10.1016/0013-4694(75)90263-1

Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, and Computers*, *31*(1), 137–149. doi: 10.3758/BF03207704

Sutton, S., Braren, M., Zubin, J., & John, E. R. (1965). Evoked-Potential Correlates of Stimulus

Uncertainty. *Science*, *150*(3700), 1187–1188. doi: 10.1126/science.150.3700.1187

Tanaka, S. C., Doya, K., Okada, G., Ueda, K., Okamoto, Y., & Yamawaki, S. (2004). Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nature Neuroscience*, *7*(8), 887–893. doi: 10.1038/nn1279

Team, R. C. (2016). R: A Language and Environment for Statistical Computing. Vienna, Austria: R Foundation for Statistical Computing.

Team, Rs. (2015). RStudio: Integrated Development for R. Boston, MA: RStudio Inc.

Trepel, C., Fox, C. R., & Poldrack, R. A. (2005). Prospect theory on the brain? Toward a cognitive neuroscience of decision under risk. *Cognitive Brain Research*, *23*(1), 34–50. doi: 10.1016/j.cogbrainres.2005.01.016

Turetsky, B. I., Dress, E. M., Braff, D. L., Calkins, M. E., Green, M. F., Greenwood, T. A., … Light, G. (2015). The utility of P300 as a schizophrenia endophenotype and predictive biomarker: Clinical and socio-demographic modulators in COGS-2. *Schizophrenia Research*, *163*(1–3), 53–62. doi: 10.1016/j.schres.2014.09.024

Uetake, a., & Murata, a. (2000). Assessment of mental fatigue during VDT task using event-related potential (P300). *Proceedings of the 2000 IEEE International Workshop on Robot and Human Interactive Communication*, 235–240. doi: 10.1109/ROMAN.2000.892501

Ullsperger, P., & Baldeweg, T. (1991). Interpreting P300 amplitude changes with adaptation level theory. *Behavioral and Brain Sciences*, *14*(4), 733–734. doi: 10.1017/S0140525X00072216

Urbach, T. P., & Kutas, M. (2002). The intractability of scaling scalp distributions to infer neuroelectric sources. *Psychophysiology*, *39*, 791–808. doi: 10.1017/S0048577202010648

Van Dijk, K. R. A., Sabuncu, M. R., & Buckner, R. L. (2012). The influence of head motion on

intrinsic functional connectivity MRI. *NeuroImage*, *59*(1), 431–438. doi: 10.1016/j.neuroimage.2011.07.044

Verleger, R., Heide, W., Butt, C., & Kömpf, D. (1994). Reduction of P3b in patients with temporo-parietal lesions. *Brain Research. Cognitive Brain Research*, *2*(2), 103–16. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/7833690

Walsh, M. M., & Anderson, J. R. (2012). Learning from experience: Event-related potential correlates of reward processing, neural adaptation, and behavioral choice. *Neuroscience & Biobehavioral Reviews*, *36*(8), 1870–1884. doi: 10.1016/j.neubiorev.2012.05.008

Walter, W. G., Cooper, R., Aldridge, V. J., McCallum, W. C., & Winter, A. L. (1964). Contingent Negative Variation: An Electric Sign of Sensorimotor Association and Expectancy in the Human Brain. *Nature*, *203*, 380–4. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/14197376

Wickens, C., Kramer, A., Vanasse, L., & Donchin, E. (1983). Performance of concurrent tasks: a psychophysiological analysis of the reciprocity of information-processing resources. *Science*, *221*(4615), 1080–1082. doi: 10.1126/science.6879207

Williams, L. M., Simms, E., Clark, C. R., Paul, R. H., Rowe, D., & Gordon, E. (2005). The Test-Retest Reliability of a Standardized Neurocognitive and Neurophysiological Test Battery: "NeuroMarker." *International Journal of Neuroscience*, *115*(12), 1605–1630. doi: 10.1080/00207450590958475

Winterer, G., Egan, M. F., Raedler, T., Sanchez, C., Jones, D. W., Coppola, R., & Weinberger, D. R. (2003). P300 and genetic risk for schizophrenia. *Archives of General Psychiatry*, *60*(11), 1158–67. doi: 10.1001/archpsyc.60.11.1158

Wronka, E., Kaiser, J., & Coenen, A. M. L. (2008). The auditory P3 from passive and active

three-stimulus oddball paradigm. *Acta Neurobiologiae Experimentalis*, *68*(3), 362–72. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/18668159

Wu, Y., & Zhou, X. (2009). The P300 and reward valence, magnitude, and expectancy in outcome evaluation. *Brain Research*, *1286*, 114–122. doi: 10.1016/j.brainres.2009.06.032

Yellott, J. I. (1969). Probability learning with noncontingent success. *Journal of Mathematical Psychology*, *6*(3), 541–575. doi: 10.1016/0022-2496(69)90023-6

Zarr, N., & Brown, J. W. (2016). Hierarchical error representation in medial prefrontal cortex. *NeuroImage*, *124*, 238–247. doi: 10.1016/j.neuroimage.2015.08.063

Zhao, C., Zhao, M., Liu, J., & Zheng, C. (2012). Electroencephalogram and electrocardiograph assessment of mental fatigue in a driving simulator. *Accident Analysis and Prevention*, *45*, 83–90. doi: 10.1016/j.aap.2011.11.019

Appendix A – Additional Figures

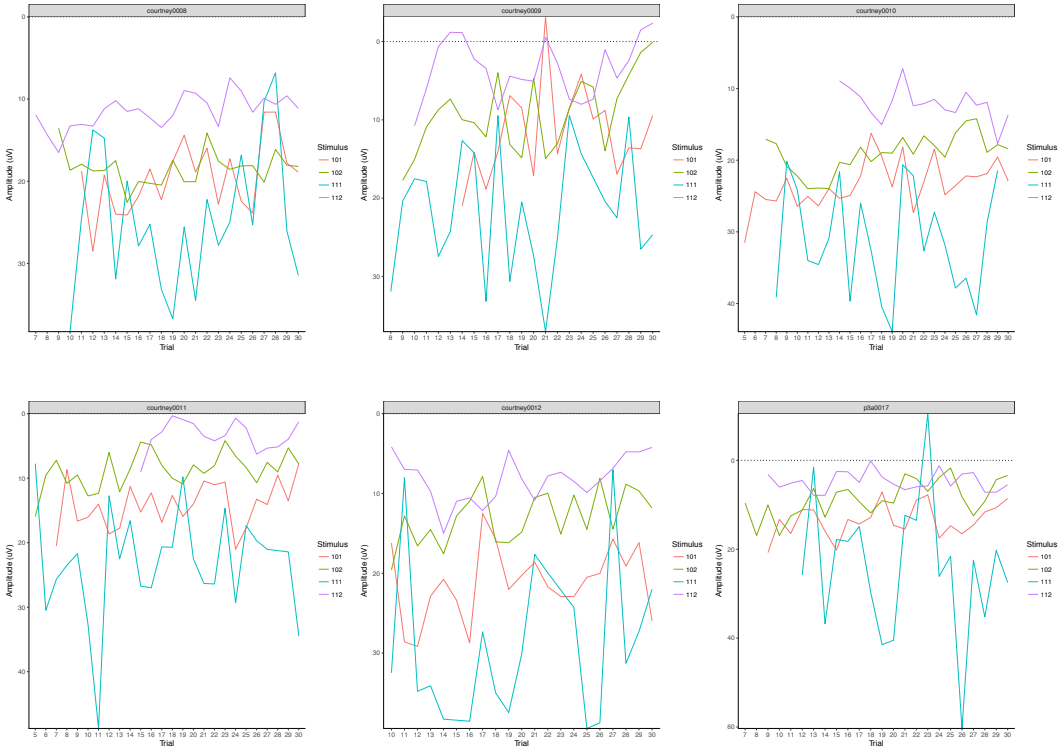

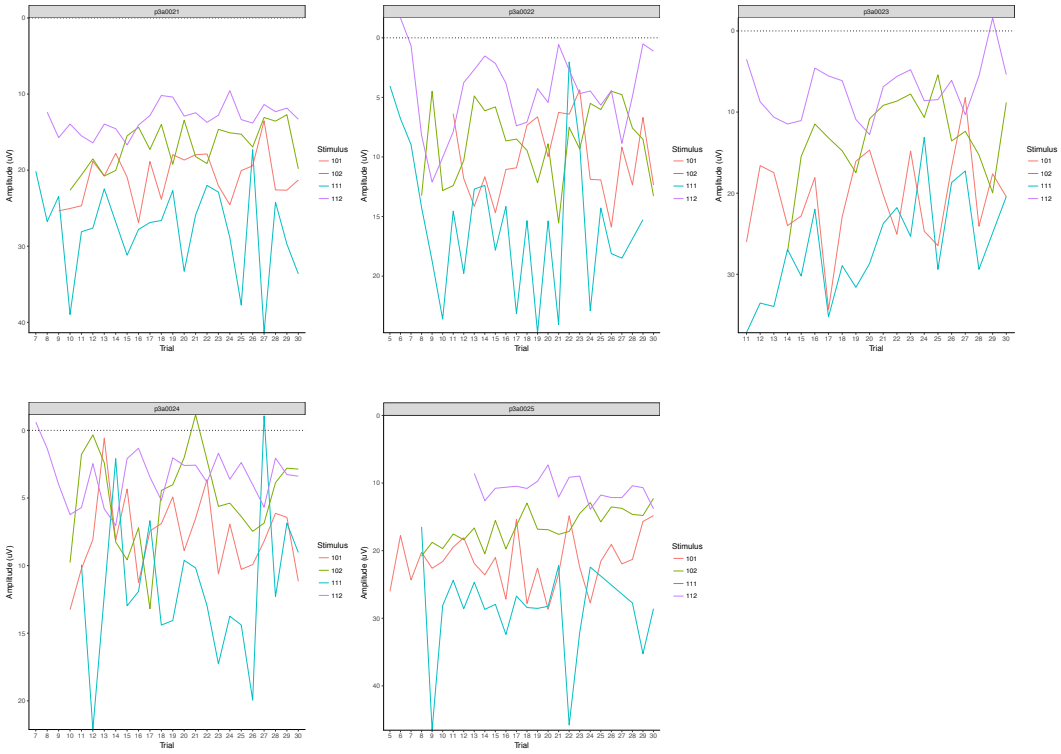

Supplementary Figure 26. Experiment 1 trial by trial P300 peak data per participant. Trials before the participants calculated inflection point are removed in this visualization.

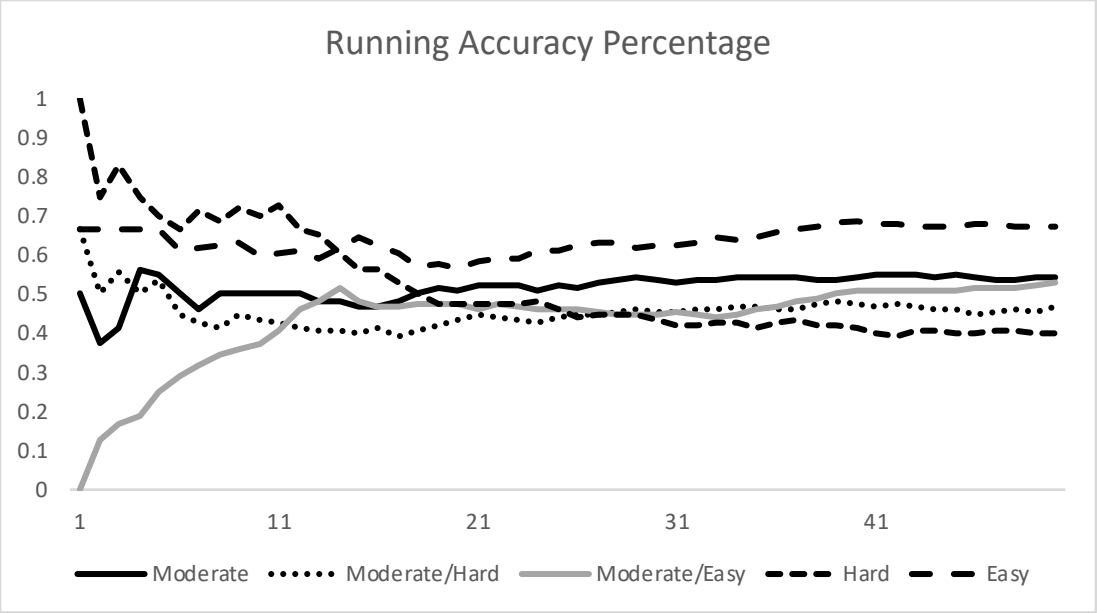

Supplementary Figure 27. Experiment 3 Participant Running Accuracy. Participant running accuracy by trial expressed over the 50 trials of a given block and separated by block difficulty. As expected, accuracy ran progressively to distinct ordered values for each condition.

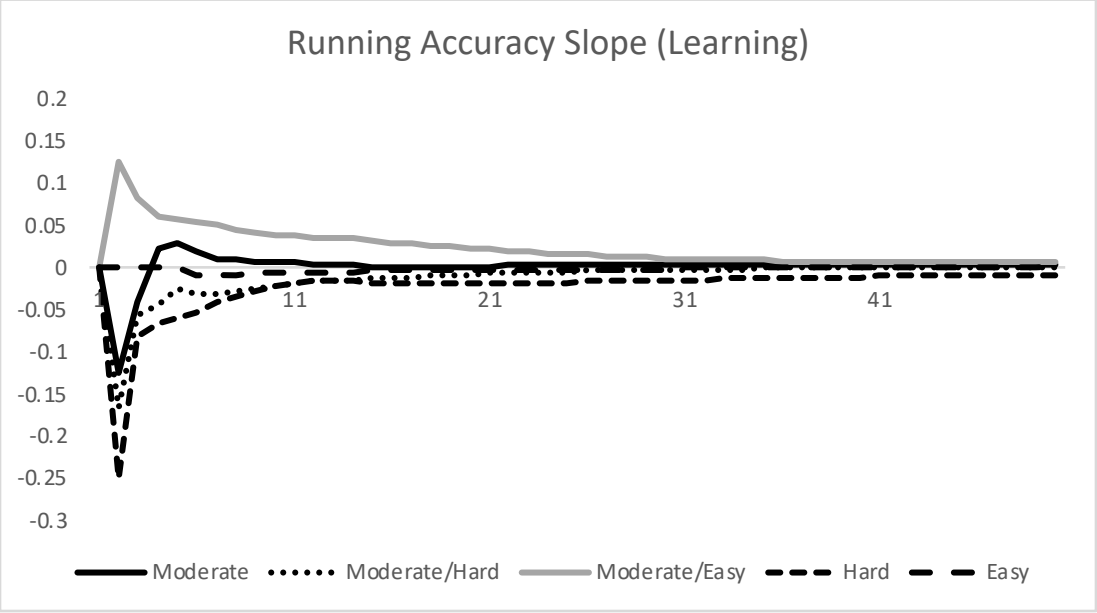

Supplementary Figure 28. Experiment 3 Learning as a Function of Accuracy Slope. Slope of running accuracy is used here to express the rate of learning taking place by the participant. It is

clear that participants accuracy is no longer changing near the final 10 trials of a block as expected.

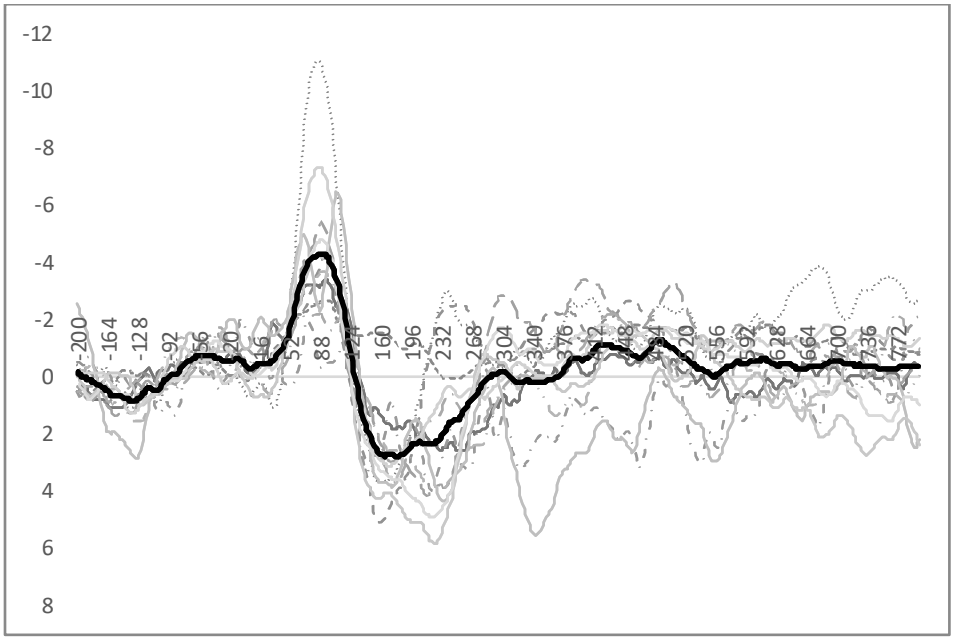

Supplementary Figure 29. Experiment 3 Auditory Control Stimulus by Participant. Participants are shown individually as greyed lines, black line indicates grand average.

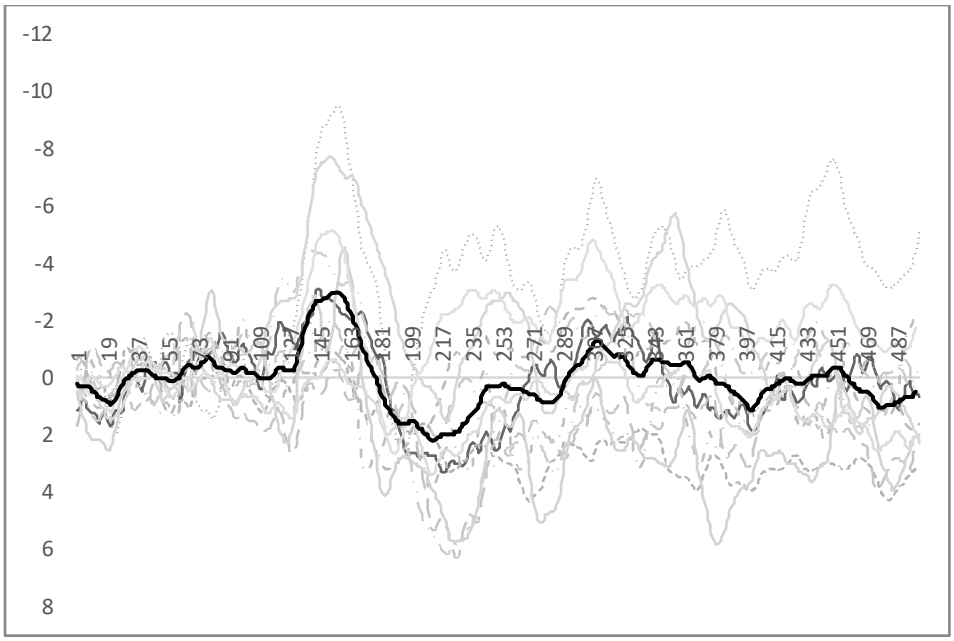

Supplementary Figure 30. Experiment 3 Auditory Oddball Stimulus by Participant. Participants are shown individually as greyed lines, black line indicates grand average.

Appendix B – Code Samples

```
#Inflection Point Function

#Note that data must be in long format, as a data.table type variable “dataa”

#Participant variable, Stimulus type variable.
InflectionPoint <- function(Participant, Stimulus){
v2 <- subset(dataa[(dataa$File == Participant) & (dataa$Stimulus == Stimulus),])
v2$variable <- as.numeric(v2$variable)
lines(v2$value, lwd=1, col="red")
d2 <- diff(v2$value)
d2 <- d2>0
d2 <- d2*2 -1
width <- 5
cutoff <- 10
scores <- sapply(width:(length(d2)- width), FUN=function(i){
  score <- abs(mean(-d2[ i-1: width ], na.rm=T) + mean(d2[ i+0: width ], na.rm=T))
})


scores <- sapply(width:(length(v2$value)- width), FUN=function(i){
  left <- (v2$value[sapply(i-1: width, max, 1) ]<v2$value[i])*2-1
  right <- (v2$value[sapply(i+1: width, min, length(v2$value)) ]<v2$value[i])*2-1

  score <- abs(sum(left) + sum(right))
})

inflections <- (width:(length(v2$value)- width))[scores>=cutoff]

plot(v2$variable, v2$value, type="l")
abline(v=inflections, col="red", lwd=3)
print(inflections)
```

Supplementary Figure 31. Inflection Point R Script.