

A Within-Subject Appetitive Procedure for Investigating the Neural Mechanisms of Latent
Inhibition and Perceptual Learning

Samantha Cristallo

A Thesis
In the Department
of
Psychology

Presented in Partial Fulfillment of the Requirements
for the Degree of
Master of Arts (Psychology) at Concordia University
Montréal, Québec, Canada

August 2025

© Samantha Cristallo, 2025

CONCORDIA UNIVERSITY
SCHOOL OF GRADUATE STUDIES

This is to certify that the thesis prepared

By: Samantha Cristallo

Entitled: A Within-Subject Appetitive Procedure for Investigating the Neural
Mechanisms of Latent Inhibition and Perceptual Learning

and submitted in partial fulfillment of the requirements for the degree of

Master of Arts (Psychology)

complies with the regulations of the University and meets the accepted standards with respect to originality and quality.

Signed by the final examining committee:

Dr. Andrew Chapman Chair

Dr. Matthew Gardner Examiner

Dr. Uri Shalev Examiner

Dr. Guillem Esber Supervisor

Approved by _____

[Chair of Department or Graduate Program Director]

[Month/day/year] _____

[Dean of Faculty]

Abstract

A Within-Subject Appetitive Procedure for Investigating the Neural Mechanisms of Latent Inhibition and Perceptual Learning

Samantha Cristallo

Latent inhibition (LI) refers to the retardation of learning about a stimulus following repeated non-reinforced exposure, whereas perceptual learning (PL) reflects enhanced discriminability among preexposed stimuli. Although both phenomena are fundamental to associative learning theories, they are usually studied separately, limiting direct comparison and obscuring shared mechanisms. This thesis developed and validated a novel within-subject paradigm for examining LI and PL under common experimental conditions, with the goal of establishing a behavioral platform for future neuroscience and translational research. In Chapter 1, we manipulated preexposure duration and found that extensive preexposure (12 days) yielded robust LI, while limited preexposure (4 days) produced PL, showing that both effects can emerge within the same design. Chapter 2 applied this approach to a within-modality context. Here, LI was attenuated and variable across individuals, suggesting that perceptual similarity between preexposed and novel cues reduced discrimination. Chapter 3 tested whether LI arises from attentional decrement or associative interference by aligning preexposure and test contingencies. No net LI or facilitation was observed, indicating that attentional and associative processes may interact in complex, cancelling ways. Together, these findings highlight the utility of a within-subject paradigm for investigating LI and PL, while also revealing boundary conditions related to preexposure amount, cue modality, and theoretical mechanism. This work establishes a foundation for probing the neural circuit underlying LI, including the hippocampus, amygdala,

striatum, thalamus, prefrontal cortex, and dopaminergic midbrain regions, and offers a translational tool for studying attentional filtering deficits in disorders such as schizophrenia.

Dedication

This thesis is dedicated to my parents, Marisa and Enzo, whose love and sacrifices made this journey possible. Mom, your perseverance has been my guiding light. You have shown me what it means to push forward with grace even when the weight of the world feels heavy. Dad, your strength has been my foundation, steady and unwavering, reminding me that no challenge is too great to overcome. Together, your love has been the force that carried me here, and I am endlessly grateful for everything you have given me.

To my brother, Raffaele, who has always been by my side. My other half, my constant source of strength and companionship, who will forever walk with me through this life. You are the epitome of resilience to me. The way you face obstacles, even when they seem unbearable, with courage and determination is something I deeply admire and strive to embody. Thank you for simply being there—your presence has been a quiet but powerful comfort, and knowing you were with me made the hardest moments easier to bear. You are not only my brother but also my anchor, and I am endlessly grateful for you.

To my best friend, Alessia, who stood by me through every tear, every late night, and every moment of doubt. You reminded me to pause, to breathe, and to take care of myself when I would have otherwise kept pushing past my limits. You have been my rock and my safe place, the one who lifted me when I could not lift myself. Your love and unwavering support turned loneliness into laughter and exhaustion into hope. I could not have reached this point without you, and I will carry your friendship as one of the greatest blessings of my life.

To my family and friends, thank you for believing in me even when I doubted myself. Your unwavering support gave me the strength to keep going. This accomplishment is not mine alone; it belongs to all of you who held me up along the way.

Table of Contents

List of Figures.....	viii
Introduction	1
Latent Inhibition: What It Is and Why It Matters.....	1
Open Questions and Challenges in LI Research	2
Separating Competing Explanations	2
Interplay with Perceptual Learning.....	2
Thesis Aims and Overview	3
Establishing a Within-Subject Latent Inhibition Paradigm	3
Integrating Latent Inhibition and Perceptual Learning Measures	3
Disentangling Attentional and Interference Mechanisms.....	4
Literature Review	4
Standard Latent Inhibition Paradigms (Appetitive and Aversive Tasks).....	4
Attentional Theories of Latent Inhibition.....	5
Interference (Memory-Retrieval) Theories of Latent Inhibition	8
Hybrid Accounts and Integrative Theories	9
Neural Bases of Latent Inhibition.....	10
Relevance to the Clinical World.....	12
Chapter 1:	14
Introduction.....	14
Methods and Materials	15
<i>Experimental Animals.....</i>	<i>15</i>
<i>Apparatus</i>	<i>15</i>
<i>Procedure.....</i>	<i>15</i>
<i>Dependent Variables.....</i>	<i>16</i>
<i>Statistical Analysis.....</i>	<i>16</i>
Results.....	17
.....	17
.....	19
.....	20
Discussion	23
Chapter 2:	24

Introduction.....	24
Methods and Materials	24
<i>Experimental animals</i>	24
<i>Apparatus</i>	25
<i>Procedure.....</i>	25
Results.....	26
<i>Correlations among dependent variables.....</i>	29
<i>Examining individual differences in latent inhibition.....</i>	30
.....	31
<i>Can sex differences explain individual differences in the expression of latent inhibition?</i>	32
.....	33
<i>Can counterbalancing conditions explain individual differences in the expression of latent inhibition?</i>	33
.....	34
Discussion	35
Chapter 3:	36
Introduction.....	36
Methods and Materials	37
<i>Experimental Animals.....</i>	37
<i>Apparatus</i>	37
<i>Procedure.....</i>	37
<i>Data Analysis</i>	38
Results.....	39
<i>Examining group-level performance during Stage 2 (conditioned inhibition training)</i>	39
<i>Examining group-level performance during Stage 3 (latent inhibition test).....</i>	42
<i>Examining Individual Differences</i>	44
Factors Potentially Contributing to Individual Differences.....	47
<i>Can sex differences account for the observed individual differences?</i>	47
<i>Statistical Analysis of Sex Differences</i>	50
<i>Can Counterbalancing Conditions Account for Individual Differences in LI.....</i>	51
<i>Relationship Between Conditioned Inhibition Strength and LI.....</i>	52
Discussion	54
General Discussion	56
References	60

List of Figures

Figure 1. Mean auROC scores for two discriminations: AX ⁺ vs AY ⁻ (novel-CS based discrimination) and AX ⁺ vs BX ⁻ (preexposed-CS based discrimination) across six two-session blocks, plotted separately for Group 12 (12-day preexposure) and Group 4 (4-day preexposure). Error bars represent ± 1 standard error of the mean (SEM).	17
Figure 2. Mean percent correct scores for two discriminations: AX ⁺ vs AY ⁻ (novel-CS based discrimination) and AX ⁺ vs BX ⁻ (preexposed-CS based discrimination) across six two-session blocks, plotted separately for Group 12 (12-day preexposure) and Group 4 (4-day preexposure). Error bars represent ± 1 standard error of the mean (SEM).	19
Figure 3. Mean response latency for two discriminations: AX ⁺ vs AY ⁻ (novel-CS based discrimination) and AX ⁺ vs BX ⁻ (preexposed-CS based discrimination) across six two-session blocks, plotted separately for Group 12 (12-day preexposure) and Group 4 (4-day preexposure). Error bars represent ± 1 standard error of the mean (SEM).	20
Figure 4. Mean PC1 composite scores for two discriminations: AX ⁺ vs AY ⁻ (novel-CS based discrimination) and AX ⁺ vs BX ⁻ (preexposed-CS based discrimination) across six two-session blocks, plotted separately for Group 12 (12-day preexposure) and Group 4 (4-day preexposure). Error bars represent ± 1 standard error of the mean (SEM).	22
Figure 1. Mean auROC scores for two discriminations: AX ⁺ vs AY ⁻ (novel-CS based discrimination) and AX ⁺ vs BX ⁻ (preexposed-CS based discrimination) across six two-session blocks, plotted separately for four measures of magazine approach: (A) number of entries, (B) percent time, (C) latency, and (D) PC1 composite. Error bars represent ± 1 SEM.	26
Figure 2. Correlations among discrimination indices based on auROC distance across subjects: (A) number vs percent, (B) number vs latency, and (C) percent vs latency. Error bars represent ± 1 SEM	29
Figure 3. Distribution of individual differences in percent auROC distance scores: (A) scatterplot of individual average scores for novel-CS based (AX ⁺ vs AY ⁻) and preexposed-CS based (AX ⁺ vs BX ⁻) discriminations, (B) histogram of signed distances from the identity diagonal, (C) descriptive statistics, (D) sex-specific scatterplot of average scores, and (E) histogram of distances separated by sex. Error bars represent ± 1 SEM.	31
Figure 4. Mean auROC distance to the identity diagonal across six session blocks, shown separately for males and females. Error bars represent ± 1 SEM.	32
Figure 5. Mean percent auROC scores for two discriminations: AX ⁺ vs AY ⁻ (novel-CS based discrimination) and AX ⁺ vs BX ⁻ (preexposed-CS based discrimination) across six two-session blocks, shown separately for (A) males and (B) females. Error bars represent ± 1 SEM.	33
Figure 6. Distribution of percent auROC distance scores across the four cue-counterbalancing groups (n = 5 per group). Error bars represent ± 1 SEM.	34
Figure 1. Mean auROC scores for two conditioned-inhibition discriminations: L ⁺ vs LA ⁺ (control) and L ⁺ vs LB ⁻ (conditioned-inhibitor) across 28 session blocks of Stage 2, plotted	

separately for (A) Number of Entries, (B) Percent Time, (C) Latency, and (D) PC1 composite. Error bars represent ± 1 SEM.	39
Figure 2. Mean auROC scores for two latent-inhibition test discriminations: AX+ vs AY– (novel-CS based) and AX+ vs BX– (preexposed-CS based) across seven sessions of Stage 3, plotted separately for (A) Number of Entries, (B) Percent Time, (C) Latency, and (D) PC1 composite. Error bars represent ± 1 SEM	42
Figure 3. Individual-difference analyses of auROC scores: (A) scatterplot and (B) signed-distance histogram for Number of Entries; (C–D) the same for Percent Time; (E–F) Latency; (G–H) PC1. Distances are computed relative to the identity diagonal (AX+ vs AY– = AX+ vs BX–). Error bars represent ± 1 SEM.....	46
.....	48
Figure 4. Sex-specific individual-difference analyses for Number of Entries, Percent Time, Latency, and PC1: scatterplots (A, C, E, G) and signed-distance histograms (B, D, F, H). Error bars represent ± 1 SEM.....	49
Figure 5. Distribution of auROC distance scores across the four counterbalancing groups ($n = 5$ per group) for number of entries (A), percent time (B), latency (C), and PC1 (D).....	51
Figure 6. Correlations between Stage 2 conditioned-inhibition distance scores and Stage 3 latent-inhibition distance scores for (A) Number of Entries, (B) Percent Time, (C) Latency, and (D) PC1	53

Introduction

Latent Inhibition: What It Is and Why It Matters

Latent inhibition (LI) refers to a learning phenomenon in which prior exposure to a stimulus without any consequence impairs the subsequent acquisition of an association with that stimulus (Aranzubia-Olasolo et al., 2024). In a typical LI procedure, an organism first experiences repeated presentations of the to-be-conditioned stimulus (CS) by itself, with no outcomes, followed by a conditioning phase where the same CS is paired with an unconditioned stimulus (US) (Badiola-Lekue et al., 2025). Pre-exposure to the CS alone impairs the emergence of the conditioned response (CR) relative to a control condition where the CS is novel (Aranzubia-Olasolo et al., 2024; Badiola-Lekue et al., 2025).

The LI effect is robust and reproducible across species, sensory modalities, and conditioning procedures (Badiola-Lekue et al., 2025). It has been demonstrated in appetitive learning, such as in the conditioned magazine approach, in aversive conditioning, such as fear conditioning or conditioned taste aversion, and in more complex tasks such as and human categorization (Barad et al., 2004; Bonardi et al., 2016). Its ubiquity suggests that LI reflects a fundamental cognitive process by which organisms filter out stimuli previously deemed irrelevant (Lubow, 2010b). Its adaptive value consists in preventing stimuli that have proven to be predictively irrelevant to control behavior (Boughner & Papini, 2003). By down-regulating attention to signals that consistently occur without consequence, animals can focus cognitive resources on more predictive cues in the environment (Braunstein-Bercovitz, 2010; Buhusi et al., 1998). This idea is central to classic learning theories of attention. For instance, the Mackintosh (1975) model proposes that if a stimulus is no better than background cues at predicting outcomes, its associability or learning rate will decline (Buhusi et al., 1998). This is the case of course in the preexposure phase of an LI experiment. In contrast, the Pearce–Hall (1975) model suggests that stimuli only command high attention when their outcomes are uncertain. After repeated nonreinforced presentations during the preexposure phase, the absence of any outcome becomes fully expected, reducing attention to the cue (Kaye & Pearce, 1984).

Beyond its theoretical interpretation, LI has significant clinical relevance. In schizophrenia, the ability to filter irrelevant stimuli is often impaired (Aranzubia-Olasolo et al., 2024; Byrom et al., 2018). Early studies found that acutely psychotic patients fail to show normal LI. Unlike healthy controls or stabilized outpatients, unmedicated acute schizophrenia patients learn about a pre-exposed stimulus almost as quickly as a novel one (Caldarone et al., 2000). LI can be restored in these patients after antipsychotic treatment, which aligns with the hypothesis that a hyperdopaminergic state in acute psychosis leads to “aberrant salience” attribution. This is a breakdown in filtering so that even irrelevant stimuli capture attention (Carson, 2010).

Pharmacological models support this view, showing that dopamine-enhancing drugs, such as

low-dose amphetamine, disrupt LI, while dopamine-blocking antipsychotics can reinstate or strengthen it (Carson, 2010).

LI has also been implicated in stress and creativity. Elevated stress hormones can impair LI, producing a state of “sensory flooding” similar to information overload in anxiety (Carson et al., 2003). Conversely, unusually low LI, which is a tendency to notice and learn about inconsequential details, has been linked to high creative achievement under certain conditions (Cassaday & Moran, 2010; Carson et al., 2003). Creative individuals with high IQ may benefit from this “leaky attention” style, which facilitates novel associations among disparate information (Cassaday & Moran, 2010).

In sum, LI is a pervasive and evolutionarily conserved phenomenon showing that prior experience critically shapes what organisms learn about new events. It serves as a benchmark for theories of associative learning, bridges animal and human research, and offers insight into the attentional and relevance-detection mechanisms that underlie adaptive behavior. Disruption of these mechanisms can contribute to mental illness (Aranzubia-Olasolo et al., 2024; Byrom et al., 2018).

Open Questions and Challenges in LI Research

Despite the long history of latent inhibition research, dating back to its discovery in 1958, progress in understanding its underlying mechanisms has been hampered by several conceptual and methodological challenges (Clark et al., 1992).

Separating Competing Explanations

A central debate is whether LI reflects a failure to pay attention to the pre-exposed stimulus (an attentional decrement) or the interference of a conflicting memory (an associative interference). Both processes likely contribute, but standard LI procedures confound them by design. During pre-exposure, subjects may both learn to ignore the stimulus, supporting an attentional account, and form an association between the stimulus and the context or a “no event” representation which will interfere with expressing the subsequent stimulus → event association. When LI is reduced or absent, it is often unclear which process was affected. A subject could fail to ignore the stimulus or exhibit less interference of the stimulus → event association by the initial stimulus → no event association. Many researchers have emphasized the need for paradigms that can disentangle these factors (Clark et al., 1992). Without such disentanglement, sophisticated neural and clinical studies can be difficult to interpret.

Interplay with Perceptual Learning

A striking “opposite” effect to LI is perceptual learning (PL), where pre-exposure to stimuli facilitates rather than hinders later learning about them. For example, if two similar stimuli are presented in an intermixed fashion without outcomes, animals and humans often later

discriminate between them more quickly than if they had not been pre-exposed (Bonardi et al., 2016). In contrast, when a single stimulus is presented repeatedly in isolation, as in a typical LI arrangement, learning about that stimulus is slowed. Both LI and PL reflect experience-driven changes in learning, yet one impairs association formation while the other enhances discrimination (Bonardi et al., 2016). This apparent contradiction raises an important question: how can the same manipulation, stimulus exposure, lead to such different outcomes? One hypothesis is that the presence of alternative stimuli and the opportunity for comparison between them are critical. When an organism encounters multiple stimuli, attention is drawn to their distinguishing features, encouraging PL. In contrast, exposure to a single stimulus in isolation provides no basis for comparison, leading the organism to tune out that stimulus and produce LI. Some theories (Honey & Hall, 1989; McLaren, Kaye, & Mackintosh, 1989; McLaren & Mackintosh, 2000) attempt to unify these phenomena by suggesting that the conditions of pre-exposure, such as whether stimuli are presented together or separately, determine whether the result is learned distinctiveness, as in PL, or learned irrelevance, as in LI. Nonetheless, it remains a challenge to design procedures that capture both effects within a single framework and to determine the conditions that cause pre-exposure to promote differentiation, as in PL, versus ignoring, as in LI.

In sum, progress in LI research has been limited by the difficulty of disentangling attentional from interference mechanisms, the constraints of traditional between-subject paradigms, and the unresolved relationship between LI and perceptual learning. These challenges motivated the experiments presented in this thesis.

Thesis Aims and Overview

The aim of this thesis is to advance our understanding of latent inhibition by developing an experimental approach that overcomes the conceptual and methodological challenges outlined above. In particular, this work seeks to address three main objectives.

Establishing a Within-Subject Latent Inhibition Paradigm

This thesis presents a basic experimental design in which the same subject experiences both the preexposed and non-preexposed conditions of an LI experiment. This within-subject approach increases sensitivity to the LI effect by allowing each subject to serve as its own control, thereby reducing variability due to individual differences. It also enables the measurement of trial-level changes in behavior within the same individual. Such a design opens the possibility of integrating within-subject neural recording or imaging methods to track attention and learning signals over the course of the experiment, which is not feasible with traditional between-group designs.

Integrating Latent Inhibition and Perceptual Learning Measures

By incorporating multiple stimuli and varying the amount of pre-exposure, the behavioral design presented here can examine, under identical testing conditions, both latent inhibition, defined as the retardation of learning for a pre-exposed stimulus, and perceptual learning, defined as the facilitation of discrimination between pre-exposed stimuli. This design provides a powerful tool for directly comparing their underlying mechanisms.

Disentangling Attentional and Interference Mechanisms

A variant of the basic experimental design is presented that aims to isolating attentional decrements in LI from associative interference. This is achieved by making the associative information acquired during preexposed and that required during the subsequent LI testing phase congruent rather than incongruent as in the standard LI design (CS→no event, CS→event). impetus of this investigation is the recognition that latent inhibition is not a unitary phenomenon, but rather the aggregate outcome of attentional and interference processes that can be separated with carefully designed experiments (Hall & Rodriguez, 2003). Achieving a clearer understanding of these processes will help disentangle the role of brain regions and circuits in LI and clarify the nature of the alterations in LI exhibited by certain clinical populations.

The remainder of this introduction provides a detailed review of relevant literature to ground the empirical work. I first examine major theoretical accounts of latent inhibition, including attentional models, interference-based models, and hybrid theories. I then review neurobiological evidence identifying brain substrates and neurotransmitter systems involved in LI, and how these map onto psychological mechanisms. Finally, I consider latent inhibition's applications and observations in clinical and differential psychology, linking back to why disentangling its mechanisms is essential for a clear interpretation of the deficits.

Literature Review

Standard Latent Inhibition Paradigms (Appetitive and Aversive Tasks)

Latent inhibition (LI) has been demonstrated across a range of conditioning paradigms, which can be broadly categorized into appetitive (reward-based) and aversive (punishment-based) tasks. The fundamental design in all cases is to compare the rate of conditioning with a stimulus preexposed without consequence with conditioning to a non-pre-exposed stimulus (Byrom et al., 2018). The defining result is slower or weaker acquisition to the pre-exposed stimulus, although the exact learning metrics vary across paradigms.

In appetitive paradigms, an initially neutral stimulus (e.g., a tone) is repeatedly presented without reinforcement. Later, that same stimulus is paired with a positive reinforcer (e.g., food or water). Animals with prior stimulus exposure typically require more trials or time to reach criterion responding than non-pre-exposed controls (Lubow, 1973). This pattern has been observed across a variety of tasks, including lever pressing for food and magazine approach. The consistency of

the LI effect in appetitive learning demonstrated that it is not restricted to any specific response or reward type (Lubow, 1973).

In aversive paradigms, LI is observed when the later learning involves an unpleasant or fear-inducing outcome. For example, in a conditioned suppression paradigm, pre-exposing a tone without consequence slows the acquisition of fear (measured by suppression of ongoing lever pressing) when that tone is later paired with shock (Rescorla, 1971). Similarly, in conditioned taste aversion (CTA) tasks, pre-exposure to a flavor (e.g., saccharin) without illness delays or weakens subsequent aversion learning compared to animals tasting the flavor for the first time (Lubow & Moore, 1959; Lubow, 1973). Such findings confirm that LI is a general property of associative learning across modalities and reinforcer types.

In addition, the effect is not limited to animals. Human LI studies, beginning with Ginton et al. (1975), have shown slower learning for pre-exposed cues in signal-detection and categorization tasks. Human adaptations of animal paradigms often use target-detection or category-learning structures that mimic CS-alone pre-exposure followed by reinforced conditioning (Allan et al., 1995; Lubow et al., 1992).

Attentional Theories of Latent Inhibition

Attentional theories propose that latent inhibition arises because pre-exposure to a stimulus without any significant outcome leads organisms to decrease the amount of attention they pay to that stimulus in the future. Put another way, the organism learns to disregard the signal and stop wasting processing resources on it through repeated nonreinforced encounters. Mechanisms for learnt attentional decrements are incorporated into a number of well-known associative learning models, and these have been applied to latent inhibition. One such theory is Mackintosh's (1975) model, which states that animals preferentially focus on cues that have proven to be the best predictors of significant outcomes. In this model, each stimulus has an associability parameter (often denoted α) that is not fixed, but changes depending on the cue's predictive value relative to other cues. If a stimulus reliably predicts the unconditioned stimulus (US) better than other available cues, its α increases, meaning the animal will pay more attention to it in the future. Conversely, if a stimulus is a poor predictor of the US (especially compared to others present), its α decreases, reflecting a decline in attention to that cue. This comparative mechanism is intuitive: it allocates processing to stimuli that carry informative value and filters out those that appear irrelevant (Le Pelley et al., 2010).

Mackintosh's model explains latent inhibition as a consequence of the pre-exposed stimulus becoming irrelevant in predicting any notable outcome. During the CS-alone pre-exposure phase, the absence of any US following the stimulus is predicted just as well by the background context as by the stimulus itself. In other words, the CS does not emerge as a better predictor of "nothing happening" than the context or other cues present, so it fails to outshine any stimulus in predictiveness. According to the model's rules, this means the CS's associability will decline

over repeated exposures – essentially, the animal learns that this stimulus is not worth paying attention to. By the time the CS is later paired with a US, its salience (α) is low, yielding slow learning of the new association. Thus, Mackintosh's theory captures latent inhibition as the direct result of learned inattention: a stimulus consistently experienced without important consequences loses the competition for attention and is subsequently overshadowed by more predictive cues when reinforcement is introduced (Mackintosh, 1975).

Geoffrey Hall and John Pearce (1980) offered a different attentional principle, almost the mirror-opposite of Mackintosh's. The Pearce-Hall model argues that animals should pay attention to cues until those cues' significance is fully learned, after which continued attention is unnecessary. In their view, attention (associability) for a stimulus is governed by the uncertainty or surprise of the outcome: if the outcome following the stimulus is unexpected, the organism dedicates processing resources to that stimulus in order to learn; but if the outcome is expected (fully predicted), the stimulus no longer demands attention. Formally, Pearce and Hall proposed that a stimulus's associability on each trial is proportional to the absolute prediction error from the previous trial – the discrepancy between what occurred and what was anticipated. When this discrepancy is large (the US outcome was surprising), the stimulus gains high α on the next trial, facilitating new learning; when the discrepancy is near zero (the outcome was predicted), the stimulus's α is low, as there is nothing new to learn. In short, this model emphasizes outcome uncertainty as the driver of attention: animals learn about what they do not yet understand and stop processing cues that have become reliable or redundant (Le Pelley et al., 2010).

The Pearce-Hall model provides a straightforward account of latent inhibition by noting that during CS pre-exposure, the animal gradually comes to expect nothing to follow the CS. Early in pre-exposure, a novel CS might cause some orienting, and the first omission of any outcome could be somewhat surprising. But after repeated CS-alone presentations, the absence of a US is no longer a surprise – the animal has learned that “nothing” happens after the stimulus. Consequently, the prediction error on each trial approaches zero, driving the associability α of that CS down to a low level. The stimulus has effectively become fully predicted (albeit predicting the absence of an event), and thus the Pearce-Hall mechanism greatly reduces attention to it. When the CS is eventually paired with a US, its low associability means the animal learns about this stimulus-outcome relationship slowly, manifesting as the latent inhibition effect (Hall & Pearce, 1979; Pearce & Hall, 1980). In summary, the Pearce-Hall theory portrays latent inhibition as the result of the stimulus's outcome becoming too well-known or certain: once the animal has learned the stimulus signifies no significant outcome, it stops paying attention, impairing new learning.

Despite their differences, the Mackintosh and Pearce-Hall models endorse the notion that pre-exposure reduces stimulus “associability,” aligning with the LI effect. An alternative, behavior-centric attentional account is provided by Lubow's Conditioned Attention Theory (Lubow, 1989). This theory asserts that the organism learns that the stimulus is unimportant during the

pre-exposure phase and a result a response of inattention (disregarding the stimulus, orienting away from it) is conditioned. Later, when the stimulus becomes significant, the subject does not instantly attend to it, which explains the slower learning.

Another influential attention-related account of latent inhibition comes from Wagner's SOP (Sometimes Opponent Process) theory (Wagner, 1981). In this model, the level of processing of a stimulus depends on the activation state of its elements. When first presented, stimulus elements enter the primary activation state (A1), where they are highly effective at forming associations. Over time, they decay into a secondary state (A2), where they are less effective as associative partners, eventually decaying into an inactive state (I). Critically, stimulus representations can be placed into the secondary activation state A2 associatively, limiting their ability to form associations. During a preexposure phase, the context can form a strong association with the CS, such that presentation of that CS in that context will activate the CS representational elements into A2. When conditioning takes place, those elements in A2 will be hindered from becoming associated with the US. While this account differs substantially from those of the Mackintosh and Pearce-Hall models, it can be regarded as attentional as well because a stimulus whose representation is in the A2 state is less well processed, and thus less well attended.

Studies that demonstrate how treatments that should alter stimulus salience modify the LI effect provide empirical evidence for attentional explanations. Latent inhibition, for example, is frequently decreased if a stimulus is made more prominent or attention-grabbing during pre-exposure (for example, by intermittent presentations or increased intensity), probably because the organism is unable to completely ignore the stimulus (Lubow, 1989; Lubow & Weiner, 2010). On the other hand, LI can be strengthened by providing a distractor task during pre-exposure, which further shifts attention away from the stimulus and reinforces its irrelevance (Ginton, Urca, & Lubow, 1975; Lubow, Schnur, & Rifkin, 1976; Barak & Weiner, 2007).

Attentional modulation in latent inhibition is further supported by a wealth of experimental data. By using an attentional lens, Kaye and Pearce (1984) famously showed that rats with hippocampal lesions show attenuated latent inhibition. This suggests that the hippocampus may be required for animals to compare expected and actual events and, as a result, to decrease attention to the inconsequential stimulus. Subsequent research revealed that the LI impact can be lessened by pharmaceutical interventions that are believed to improve attention. A low amount of amphetamine administered during the pre-exposure period is one such manipulation; amphetamine is known to enhance attentional processes by increasing dopamine release. According to Weiner, Lubow, and Feldon (1988), rats administered amphetamine before exposure do not exhibit the typical latent inhibition, suggesting the drug prevents them from ignoring the stimulus. According to Weiner (2003), this finding lends credence to an attentional interpretation since the rats were able to learn about the stimuli rapidly after it became relevant

because the pharmacological attention-boosting effect maintained the stimulus processing even in the absence of reinforcement.

Interference (Memory-Retrieval) Theories of Latent Inhibition

Interference theories, as opposed to attentional explanations, credit latent inhibition to processes that take place during conditioning (or a recall test thereafter), rather than to decreased attention during pre-exposure. According to this account, during the pre-exposure phase a CS→no event association forms (Bouton, 1993; Miller & Escobar, 2001). When an outcome is later associated with the stimulus, the organism has two contradictory pieces of information: the earlier CS→no event memory and the current CS→outcome association. This conflict can cause a performance deficit, manifesting as slower acquisition of the new response because the earlier learned memory hinders the expression of the new learning (Bouton, 1993; Escobar, Arcediano, & Miller, 2002). Interference accounts, in contrast to attentional theories, do not require that the subject ignores the stimulus during conditioning; rather, the subject may process the stimulus and learn about it, but recall of the prior learning influences the subject's behavior (Weiner, 2003). In other words, latent inhibition is essentially viewed as a type of proactive interference (Lubow, 1989).

One influential interference theory is the Comparator Hypothesis proposed by Miller and Matzel (1988). According to this framework, conditioned responding is not determined solely by the strength of the CS→US association but by a comparison process at the time of testing. Specifically, the subject compares the direct CS→US association with indirect pathways (e.g., CS→context→US, or CS→other cues→US). If these comparator associations provide strong alternative activation of the US, they downmodulate behavioral expression of the CS→US association. Thus, reduced responding after CS preexposure (latent inhibition) is interpreted not as a failure to acquire the CS→US association, but as a performance deficit: the preexposed CS has formed strong comparator associations (such as CS→context), which interfere with retrieval and expression of new learning about the CS→US relationship.

Studies on context changes provide one compelling empirical example of interference theories. The latent inhibitory effect is significantly reduced or even eliminated if the context is altered between the pre-exposure and conditioning phases, according to Hall and Channell's (1986) research. If the CS–context association that was learnt in pre-exposure is context-specific, then moving to a different context means that the interfering association is no longer retrieved, allowing the animal to learn about the CS–US normally (i.e., no LI is observed). This makes sense from the perspective of interference. On the other hand, unless further assumptions are provided, attentional theory might not predict such a significant loss of LI with a context change. A key piece of evidence that suggests the effect is based on some memory of the pre-exposure (related to context) is the replication of the context-specificity of LI (e.g., Hall & Honey, 1989). Additionally, manipulating time intervals provides support: if a lengthy lag occurs between pre-

exposure and conditioning, latent inhibition tends to wane (the earlier memory becomes less significant over time), which again points to the involvement of memory processes that are susceptible to deterioration or decreased retrieval.

"Learned irrelevance" is another idea that is closely associated with interference accounts. This term was first used by Baker and Mackintosh (1977) to characterize the challenge of learning a CS-US link following uncorrelated pre-exposures to the CS and US. Animals are randomly exposed to both the CS and the US during their procedure, therefore the CS has no bearing on the US prediction. More so than the pre-exposure effects of CS alone or US alone, they discovered that learning the CS-US relationship afterward is slower than usual (Baker & Mackintosh, 1977). This implies that the animals developed a learned expectation of irrelevance that impedes subsequent acquisition—that is, that the CS has no significance in relation to the US. Latent inhibition (CS-alone pre-exposure) can be viewed as a component of learned irrelevance, and in both cases the learned “nothing there” association must be overcome for conditioning to occur (Bonardi & Hall, 1996; Mackintosh, 1973). The fact that learned irrelevance produces an even stronger retardation than CS pre-exposure alone implies that associative interference (from both CS–no outcome and context associations, and potentially from the US–no CS experiences) is a real factor in slowed learning (Bonardi, 1991; Escobar, Arcediano, & Miller, 2002). Thus, interference theorists argue that latent inhibition is fundamentally about competition between old and new learning, not necessarily a change in how much attention is paid to the stimulus (Miller & Escobar, 2001).

In summary, latent inhibition can be explained by interference theories without referring to a failure to pay attention to the stimuli. These perspectives, on the other hand, concentrate on pre-exposure learning (context connections, "no outcome" expectations), and how those memories prevent or postpone the manifestation of new learning. Findings that are hard to reconcile with a merely attentional explanation, such as the disruption of LI by context shifts and the cumulative effects of CS+US uncorrelated pre-exposures (learned irrelevance), support them.

Hybrid Accounts and Integrative Theories

Given the substantial evidence supporting both attentional and interference mechanisms, some researchers have proposed hybrid accounts that incorporate elements of both. These integrative theories suggest that latent inhibition may not stem from a single simple process; rather, multiple processes operate in parallel during the standard LI procedure, including both reduced attention to the stimulus and the formation of interference-producing associations (Lubow, 1989; Weiner, 2003). One straightforward possibility is that different stages or conditions of the experiment tap into different mechanisms. For example, during the early part of conditioning, a pre-exposed subject might show slower learning because it is still not attending fully (attention mechanism), but as conditioning continues, even once attention is regained, the subject might then experience lingering interference from the earlier context association (memory mechanism) (Bouton, 1993;

Escobar, Arcediano, & Miller, 2002). In this view, attention and memory interference effects could combine to yield the overall observed latency in learning (Miller & Escobar, 2001; Nelson, 2002).

Formal models have been developed to integrate attentional and interference accounts of latent inhibition. The most explicit example is Hall and Rodríguez's (2010) hybrid model, which combines reduced stimulus associability with interference from previously learned context–outcome associations. In this framework, pre-exposure lowers the associability of the conditioned stimulus, making it harder to form new associations, while at the same time the pre-exposed context retrieves competing “no-outcome” associations that interfere with learning. This model captures how both attentional decrement and associative interference contribute to the retardation of conditioning observed in latent inhibition.

Some empirical research has examined the contributions of memory and attention directly. In order to quantify attention, for instance, researchers have measured the orienting responses or other indices of stimulus processing throughout the pre-exposure and conditioning phases. They have also presented reminder cues or changed contexts to measure the effects of memory interference. Frequently, the findings show that neither pure theory is enough. However, a context change still disrupts LI (implying a necessary retrieval of the context-linked memory), indicating that both reduced attention and interference are real and jointly contribute (Bonardi & Hall, 1996; Escobar et al., 2002). In some cases, animals seem to stop orienting to the stimulus during pre-exposure (supporting attentional decrement). Due to these results, current evaluations have concluded that there is no one recognized mechanism for latent inhibition; instead, various preparations may highlight one element more than another. Miller, Schachtman, and Moyer (1985) pointed out that “perhaps because of the evidence, there is no consensus for a mechanism underlying LI”, several psychological processes might be at play, and their relative roles could vary depending on the procedure's particulars.

The challenge faced by hybrid theories is to empirically disentangle the contribution of attentional and interference processes, which is one aim of this thesis (Study 3).

Neural Bases of Latent Inhibition

Given LI's links to schizophrenia models, there has been a lot of attention focused on understanding the brain underpinnings of latent inhibition. Initially, scientists looked for areas of the brain whose activity would reveal whether latent inhibition is primarily mnemonic or attentional. The hippocampus has been identified as one of the important structures. According to Kaye and Pearce (1984), rats with hippocampus lesions exhibit just as rapid conditioning to pre-exposed stimuli as control rats, indicating that the lesions eliminate latent inhibition. Numerous interpretations have been offered for this data, which implies that the hippocampus is essential for the typical LI effect to occur. One theory is that the hippocampus is required for comparing stimuli in different contexts or for identifying the absence of consequences, which are

processes associated with attention or expectation. According to a different theory, the hippocampus stores the stimulus-no-outcome memory that might subsequently cause problems. Hippocampal involvement in either scenario suggests that latent inhibition depends on higher-order processing in the brain's learning and memory circuits rather than being a straightforward peripheral sensory phenomenon.

Beyond the hippocampus, extensive evidence implicates the dopamine system and associated regions (such as the nucleus accumbens) in latent inhibition (Weiner, 2003). A key target in LI research, dopamine (DA) is a neurotransmitter that plays a key role in learning and motivational salience. It is also noticeably dysregulated in schizophrenia. Animal pharmacological research has demonstrated that whereas blocking dopamine receptors tends to increase or restore latent inhibition, boosting dopamine activity tends to destabilize it. The LI effect, for example, was removed when rats were given amphetamine (which increases synaptic dopamine) before training; pre-exposed rats conditioned under amphetamine learnt at the same rate as non-pre-exposed controls, according to Solomon et al. (1975). Antipsychotic medications, many of which are dopamine D2 receptor antagonists, on the other hand, stop that disruption and, in situations where latent inhibition is lacking, can even bring it back (Weiner et al., 1988). In a seminal review, Weiner (2003) contended that antipsychotics normalize LI in a way that is similar to their clinical effects on salience attribution, while disturbed LI under elevated dopamine is a rodent analog of the attentional deficiencies in schizophrenia. A key location for these effects seems to be the nucleus accumbens (NAcc), a dopamine-rich area in the ventral striatum. After reviewing pharmacological data, Moser et al. (2000) came to the conclusion that mesolimbic dopamine acts on the NAcc to control latent inhibition: dopaminergic blockade increases LI, while hyperdopaminergic activation of the NAcc attenuates LI. This fits very nicely with both interference theories (since dopamine may influence how strongly the "no consequence" connection competes with new learning) and attentional theories (because dopamine is linked to novelty/salience processing). Furthermore, the regulation of LI has been linked to the prefrontal brain and other areas. This suggests that cognitive control networks also play a role in ignoring vs. learning about inputs. For instance, Kumari et al. (1999) employed neuroimaging in humans doing a latent inhibition task and identified abnormal activation in frontal cortical areas in schizophrenic subjects.

Theoretical debates can occasionally be clarified by neural investigations; for example, if a neural modification selectively affects latent inhibition in a way that is compatible with one mechanism, it can support that interpretation. Considering the hippocampus is frequently associated with novelty detection and attention, the results of the hippocampal lesion, for instance, may be more compatible with an attentional failure; yet some contend that they suit an encoding failure (no context association encoded). According to Kapur (2003), excessive dopamine can cause abnormal salience attribution to previously irrelevant stimuli, which explains reduced LI in psychosis. However, other researchers believe that dopamine modulates

the prediction error signals that underpin learning, which in turn affects the associative competition during conditioning. Similarly, the dopamine results could be framed as dopamine controlling a learning rate or "switch" for attention. It is evident that latent inhibition activates a distributed neural circuitry that includes prefrontal areas, striatal regions (accumbens), and limbic regions (hippocampus, amygdala). Disturbances to this circuitry can interfere with normal expression of LI and are similar to the neurochemical imbalances in schizophrenia (Weiner, 2003; Kumari et al., 1999). In addition to highlighting the difficulty of identifying a single mechanism for the impact, this convergence of pharmacological and neurological evidence further supported the use of latent inhibition as a behavioral model for psychopathology.

Relevance to the Clinical World

Latent inhibition has been most extensively studied in relation to schizophrenia. As an attentional phenomenon, LI provides a quantifiable measure of how individuals filter irrelevant stimuli, which is highly relevant to the sensory gating deficits and information overload characteristic of schizophrenia. In one of the seminal human studies, Baruch et al. (1988) found that acute, unmedicated patients showed significantly reduced LI compared to healthy controls, learning about familiar and novel cues at nearly the same rate. This suggests a breakdown in the ability to ignore previously irrelevant information. Subsequent research has replicated and extended these findings. For example, Gray et al. (2001) reported that patients off medication often show absent or even reversed LI, whereas patients on stable antipsychotic medication typically exhibit normal LI. This pattern is consistent with the idea that dopamine-blocking medications restore LI, supporting the view that a hyperdopaminergic state underlies the deficit.

LI deficits are also observed in individuals with schizotypal traits or those at high risk for psychosis, although typically to a lesser degree than in acutely ill patients (Kumari et al., 1999). Across studies, reduced LI is consistently associated with positive symptoms such as delusions and hallucinations, consistent with the idea that insufficient filtering can lead to an overload of trivial details and aberrant associations. Conversely, individuals with pronounced negative symptoms, such as apathy or cognitive impairment, sometimes display enhanced LI, suggesting an overly strong tendency to disregard stimuli (Lubow et al., 2000). This indicates that optimal cognitive functioning requires a balance in sensory gating: too little filtering may foster positive symptoms, while too much may contribute to negative symptoms.

Beyond schizophrenia, LI has been explored in other psychological and clinical contexts. Notably, Carson et al. (2003) found that eminent creative achievers often display lower LI, but only when accompanied by high IQ. In such cases, a "leaky" sensory filter may allow more irrelevant information into awareness, potentially fueling creative thinking if the cognitive capacity exists to manage it. While not strictly clinical, this finding underscores that low LI is not

universally maladaptive. In other conditions, such as severe stress, mania, or amphetamine intoxication, LI has also been shown to be reduced (Weiner, 2003).

The clinical relevance of LI has two major implications. First, it supports models of psychosis that emphasize aberrant salience attribution, such as Kapur's (2003) framework. In this view, excess dopamine causes the brain to assign undue importance to previously irrelevant stimuli, fragmenting attention and contributing to delusional thinking. LI attenuation in schizophrenia aligns with this model, as previously neutral cues capture attention and may become incorporated into psychotic ideation. Second, LI holds potential as a cognitive marker or endophenotype for psychiatric illness. While findings have been somewhat inconsistent across studies due to methodological differences and medication effects, LI could be used to assess the efficacy of pro-cognitive treatments or to help subtype patients according to their information-processing style.

In sum, latent inhibition provides a valuable bridge between basic learning theory and clinical psychology. Its modulation in schizophrenia-spectrum conditions highlights its relevance to selective attention and learning processes, and clarifying the mechanisms of LI could offer important insights into pathological cognition and inform the development of targeted interventions.

Chapter 1:

Developing a within-subject design to study the mechanisms of latent inhibition and perceptual learning within the same experimental design

Introduction

The current study employed a within-subjects experimental design to examine the effects of varying levels of stimulus preexposure on the subsequent associability and discriminability of cues relative to novel cues. Specifically, I directly compared the discriminability and conditioning performance of preexposed and novel stimuli within the same animals.

Preexposure to stimuli can have paradoxical effects on later associative learning. In some cases, limited or moderate exposure enhances the distinctiveness of cues and facilitates discrimination, a phenomenon known as perceptual learning (PL; Gibson, 1969; Hall, 2001). In other cases, extended exposure reduces a stimulus's associability, producing slower acquisition of conditioned responding when that stimulus is later paired with reinforcement. This is the latent inhibition effect (LI; Lubow & Moore, 1959; Lubow, 1989; Hall & Rodríguez, 2011).

Traditionally, LI and PL have been studied in isolation using different procedures, often in between-subjects designs, which makes direct comparisons difficult and limits the ability to examine their neural underpinnings (Hall, 2003; Honey & Hall, 1989). In the present design, I sought to overcome this limitation by creating a unified within-subject paradigm. I operationalized latent inhibition as slower learning about a preexposed cue compared to a novel cue (Lubow & Moore, 1959; Hall & Rodríguez, 2011). To capture this, I derived two discrimination indices: one based on the preexposed stimulus (AX+ vs BX-) and one based on the novel stimulus (AX+ vs AY-). Because each rat contributed both indices, the critical LI comparison could be made within the same animals rather than across groups. This within-subject approach ensured that LI and PL could be evaluated directly under the same training and testing conditions.

Two groups of rats experienced identical experimental contingencies except for the amount of preexposure they received. Group 12 received 12 days of preexposure, while Group 4 received only 4 days. This design allowed me to test two competing hypotheses about the critical factor driving LI and PL. Hypothesis 1 proposed that the mode of presentation determines the outcome, with latent inhibition expected when only one cue is preexposed and perceptual learning expected when two cues are preexposed together. Hypothesis 2 proposed that the amount of preexposure is the decisive factor, with extensive exposure producing latent inhibition and more moderate exposure producing perceptual learning. By directly contrasting these conditions within a unified design, I aimed to determine whether LI and PL reflect different mechanisms tied to presentation mode or instead represent outcomes along a continuum determined by exposure duration.

Methods and Materials

Experimental Animals. Sixteen sex-balanced Long-Evans rats (*Rattus norvegicus*), approximately 3 months of age, were used. Males weighed between 522 and 715 g, while females weighed between 264 and 522 g. All animals were sourced from Charles River Laboratories and housed individually in standard clear-plastic tubs (10.5 in × 19 in × 8 in, Charles River Laboratories) with woodchip bedding, within a colony room maintained on a 14:10 light/dark cycle schedule. Behavioral sessions were conducted between 7–10 hours after the onset of the light phase of the cycle. Throughout the experiments, water access was restricted to 1 hour per day following each session, while food was provided ad libitum. All animal care and experimental procedures were conducted in compliance with the ARRIVE guidelines and the NIH's Guide for the Care and Use of Laboratory Animals, with protocols approved by the Brooklyn College Institutional Animal Care and Use Committee.

Apparatus. Behavioral training took place in eight identical operant-conditioning chambers (Med Associates, Inc.) measuring 32 cm (length) × 25 cm (width) × 33 cm (height). Each chamber was enclosed within a ventilated, sound-attenuating cubicle (74 cm × 45 cm × 60 cm) equipped with an exhaust fan that provided a constant background noise of approximately 50 dB. The grid floor of each chamber consisted of stainless-steel rods spaced 1.1 cm apart, connected to a shock generator for other experiments but unused in the present study.

On one side wall, two white jewel lights (28 V DC, 100 mA) were mounted on the left and right panels, each positioned 9.3 cm above the floor. Directly above each light (20.6 cm from the floor) was a speaker driven by a dedicated tone generator: the left speaker emitted a 2.5 Hz, 80 dB click train, while the right speaker produced 70 dB white noise. The opposite wall contained two additional speakers positioned 24.8 cm above the floor. The left speaker produced a 12 kHz, 70 dB tone, and the right speaker produced a 1 kHz, 80 dB tone. These auditory cues served as the stimuli for the experimental tasks.

A recessed liquid delivery port was located on the right wall, 3 cm above the floor and centered between the side panels. Sucrose reinforcement consisted of 0.04 ml of a 10% sucrose solution delivered via a solenoid-operated dipper. Entries into the port were recorded via an infrared beam. Experimental events were controlled and recorded by a Med Associates interface and software system.

Procedure. A within-subjects design compared two preexposure durations (12 days vs. 4 days) to assess their effects on subsequent discrimination learning. Cues A and B were either auditory or visual, counterbalanced across subjects (12-kHz tone at 70 dB and white noise at 70 dB; or click at 80 dB and 1-kHz tone at 80 dB).

Stage 1: Preexposure. Rats received nonreinforced presentations of A and B. Each daily session contained 32 trials arranged pseudorandomly with the constraint that no more than 2 trials of the

same type occurred consecutively. Trial composition was 16 A- and 16 B-. Each trial comprised a 10-s pretrial, 10-s CS, and 10-s posttrial, separated by an intertrial interval averaging 120 s. Group 12 completed 12 daily preexposure sessions; Group 4 completed 4 daily sessions.

Stage 2: Discrimination Training (Test). Across 12 daily sessions, rats were trained on two subdiscriminations using compounds that combined preexposed and novel cues. Each session comprised 32 trials arranged pseudorandomly with the constraint that no more than 2 trials of the same type occurred consecutively and no more than 3 nonreinforced trials occurred in succession. Trial composition was 16 AX+, 8 AY-, and 8 BX-. On AX+ trials, the compound CS lasted 10 s and was followed immediately by 3 s access to 10% sucrose in the dipper; AY- and BX- trials ended without reinforcement. Each trial included a 10-s pretrial, a 10-s CS, a 10-s posttrial, and an intertrial interval averaging 120 s.

Dependent Variables: We examined conditioned magazine approach across multiple indices, including number of head entries, percent of time spent in the magazine, response latency, and the first principal component (PC1) resulting from a principal component analysis (PCA) conducted on those three measures. To determine how discrimination learning progressed over time for each cue in each group, these were monitored over the course of twelve training sessions and collapsed into two-session blocks for visualization and analysis purposes.

Statistical Analysis. Discrimination performance was quantified using the area under the receiver operating characteristic curve (auROC). The auROC compares response distributions between rewarded and non-rewarded trial types to assess how well animals distinguished between them. For each rat and each session, I computed two discrimination indices: one comparing AX+ vs BX- trials (preexposed-CS based discrimination) and one comparing AX+ vs AY- trials (novel-CS based discrimination). This procedure yielded two discrimination values per animal per session rather than three separate cue comparisons, providing a direct behavioral measure of latent inhibition within the same subjects. To calculate auROC scores, response distributions for the two trial types were sorted and evaluated across a series of thresholds. At each threshold, the false positive rate (BX- or AY- trials incorrectly classified as AX+) and true positive rate (correct classification of AX+ trials) were computed. These values were used to construct a receiver operating characteristic (ROC) curve, and the area under the curve was estimated using the trapezoidal rule. Scores ranged from -1 to 1, with 0 indicating chance-level discrimination (complete overlap of response distributions), positive values indicating discrimination in the expected direction, and negative values reflecting a reversed response bias. Mixed-design ANOVAs were conducted on these auROC discrimination scores and on other dependent measures (number of responses, percent correct, response latency, and the composite PC1 index). Where appropriate, post hoc tests with Tukey corrections were applied to evaluate specific differences between groups, discrimination types, and training sessions.

Results

Number of Responses. Group mean AUROC scores across the session blocks are shown in Figure 1, depicting performance on two separate discriminations: AX+ vs. AY- and AX+ vs. BX-. Both groups successfully learned their respective discriminations, evident by the increasing mean AUROC scores across training sessions.




Figure 1. Mean auROC scores for two discriminations: AX+ vs AY- (novel-CS based discrimination) and AX+ vs BX- (preexposed-CS based discrimination) across six two-session blocks, plotted separately for Group 12 (12-day preexposure) and Group 4 (4-day preexposure). Error bars represent ± 1 standard error of the mean (SEM).

The figure also suggests that the discriminations were solved at different rates by the two groups, and this was confirmed by a significant Group $\times$ Discrimination interaction ($F(1,14) = 9.65$, $p = .008$) and a significant Group $\times$ Discrimination $\times$ SessionBlock interaction ($F(5,70) = 4.66$, $p < .001$). There was also a significant main effect of SessionBlock ($F(5,70) = 22.26$, $p < .001$), but no main effects of Group ($F(1,14) = 0.76$, $p = .398$) or Discrimination ($F(1,14) = 0.07$, $p = .795$).

Post hoc analyses revealed significant differences in Group 12 (12-day preexposure) during sessions 4 ($p = .0125$), 5 ($p = .0116$), and 6 ($p = .0137$), with better discrimination performance on AX+ vs. AY- trials as indicated by higher auROC scores, reflecting more effective learning involving the novel cue. Conversely, in Group 4 (4-day preexposure), significant differences occurred during sessions 4 ($p = .0008$) and 5 ($p = .0417$), but with better performance on AX+ vs. BX- trials, reflecting enhanced discrimination involving the familiar cues.

The pattern supports the idea that extensive preexposure (Group 12) impaired associative learning to the preexposed cue (AX), indicative of latent inhibition. In contrast, moderate preexposure (Group 4) improved discrimination involving familiar stimuli, consistent with a perceptual learning effect.

These results are consistent with the hypothesis that extensive preexposure to AX in Group 1 impaired its conditioning, leading to relatively enhanced responding to the novel stimulus pairing (AY). In Group 2, which received less preexposure, conditioning to AX and BX proceeded more symmetrically, with a slight advantage for BX. The reversal in discrimination direction between groups supports the prediction that preexposure duration modulates associative strength, leading to measurable differences in discrimination performance.

Percent Correct. Group mean percent correct scores for discriminations across six training sessions are shown in Figure 2, separately for AX+ vs. AY- and AX+ vs. BX- discriminations. Both groups successfully improved their discrimination accuracy over training, as evidenced by increasing accuracy across session blocks.

A clear divergence in discrimination performance emerged across groups and sessions. In Group 1 (12-day preexposure), rats demonstrated stronger discrimination in favor of AX+ AY- trials as early as Session 2, with this pattern becoming more pronounced over time. From Sessions 2 through 6, AX+ AY- trials consistently yielded higher accuracy than AX+ BX-, suggesting that conditioning was more effective when the less familiar cue (AY) was present. This pattern is consistent with the prediction that extended preexposure to AX would impair learning to that cue, a behavioral signature of latent inhibition.

In contrast, Group 2 (4-day preexposure) showed an early advantage for AX+ BX- trials. Specifically, performance on AX+ BX- trials exceeded AX+ AY- in Sessions 2 and 4, suggesting that with reduced preexposure, AX retained its associative strength and was more readily conditioned. This divergence across groups illustrates the influence of preexposure duration on learning outcomes.

The discriminations were not learned equally well by the two groups, confirmed by significant interactions: Group $\times$ Discrimination ($F(1,14) = 31.02, p < .001$), Discrimination $\times$ SessionBlock ($F(5,70) = 2.84, p = .022$), and Group $\times$ Discrimination $\times$ SessionBlock ($F(5,70) = 4.61, p = .001$). There was also a highly significant main effect of SessionBlock ($F(5,70) = 23.80, p < .001$), though the main effects of Group ($F(1,14) = 0.34, p = .569$) and Discrimination ($F(1,14) = 4.41, p = .054$) were not fully significant.

Post hoc analyses clearly showed distinct patterns for each group. In Group 1 (12-day preexposure), AX+ vs. AY- discrimination accuracy was significantly higher than AX+ vs. BX- across sessions 2 through 6 (Session 2: $p = .0006$; Session 3: $p < .0001$; Session 4: $p = .0002$; Session 5: $p = .0002$; Session 6: $p < .0001$). This indicates consistently superior discrimination involving the novel cue (AY). Conversely, in Group 2 (4-day preexposure), AX+ vs. BX- discrimination accuracy was significantly higher than AX+ vs. AY- in sessions 2 ($p = .0409$) and 4 ($p = .0043$), reflecting better learning of the discrimination involving more familiar cues (BX).

Thus, these results support differential discrimination performance depending on preexposure duration: extensive preexposure (12 days) was associated with enhanced discrimination of novel cues, consistent with latent inhibition effects, while moderate preexposure (4 days) supported improved discrimination of familiar cues, aligning with perceptual learning.




Figure 2. Mean percent correct scores for two discriminations: AX⁺ vs AY⁻ (novel-CS based discrimination) and AX⁺ vs BX⁻ (preexposed-CS based discrimination) across six two-session blocks, plotted separately for Group 12 (12-day preexposure) and Group 4 (4-day preexposure). Error bars represent ± 1 standard error of the mean (SEM).

Mean percent correct discrimination scores for AX⁺ vs. AY⁻ and AX⁺ vs. BX⁻ trials across six training sessions. Group 1 (12-day preexposure) showed consistently higher accuracy on AX⁺ vs. AY⁻ discriminations, indicative of stronger discrimination involving the novel cue. Group 2 (4-day preexposure) showed higher accuracy on AX⁺ vs. BX⁻ discriminations, indicative of enhanced perceptual learning effects involving familiar cues. Error bars represent ± 1 SEM.

Latency. Group mean response latencies across the six training sessions are presented in Figure 3, with separate lines for AX⁺ AY⁻ and AX⁺ BX⁻ trials within each group. The x-axis represents session blocks (1 through 6), and the y-axis shows the mean latency to respond, measured in seconds. In general, rats in both groups showed reduced latencies over time, but the overall pattern of discrimination was less pronounced than in the previous measures.

Mean response latencies for each discrimination (AX⁺ vs. AY⁻ and AX⁺ vs. BX⁻) across the six session blocks (each block representing two of the 12 training sessions) are presented in Figure 3. Both groups showed overall improvement in response speed (shorter latencies) as training progressed, confirmed by a significant main effect of SessionBlock ($F(5,70) = 13.50, p < .001$). However, neither the main effects of Group ($F(1,14) = 2.55, p = .133$) nor Discrimination ($F(1,14) = 2.38, p = .145$) reached significance.

In Group 1 (12-day preexposure), a difference between trial types emerged early in training. Specifically, rats exhibited longer response latencies on AX+ AY-trials than on AX+ BX- trials during Sessions 2 and 3. However, this difference diminished in later sessions. In Group 2 (4-day preexposure), no consistent differences in latency were observed between AX+ AY- and AX+ BX- trials at any point during training, suggesting weaker or less differentiated cue processing.

A mixed-design ANOVA on latency data revealed a significant main effect of SessionBlock, $F(5, 70) = 13.50, p < .001$, indicating that response speeds improved with training. However, there were no significant main effects of Group ($F(1, 14) = 2.55, p = .133$) or Discrimination ($F(1, 14) = 2.38, p = .145$), and critically, the Group $\times$ Discrimination $\times$ SessionBlock interaction was not significant ($F(5, 70) = 1.63, p = .163$). Two significant two-way interactions were found: Group $\times$ SessionBlock ($F(5, 70) = 3.07, p = .015$) and Discrimination $\times$ SessionBlock ($F(5, 70) = 3.30, p = .010$), suggesting that latency patterns shifted differently across time for each group and discrimination type.

Post hoc analyses indicated that only Group 1 (12-day preexposure) showed significant latency differences between discriminations, specifically in session blocks 2 ($p = .0160$) and 3 ($p = .0023$), where response latencies were significantly higher for AX_ AY trials compared to AX_ BX trials. This suggests initial hesitation or less certainty early on when discriminating novel cue pairings (AY) compared to familiar cue pairings (BX). For Group 2 (4-day preexposure), no significant latency differences emerged in any session block.

Together, these findings indicate that latency was a less sensitive measure compared to AUROC and percent correct, capturing early-stage differences in discrimination processes primarily in the 12-day preexposed group, but failing to consistently differentiate discrimination performance in the 4-day group.




Figure 3. Mean response latency for two discriminations: AX⁺ vs AY⁻ (novel-CS based discrimination) and AX⁺ vs BX⁻ (preexposed-CS based discrimination) across six two-session

blocks, plotted separately for Group 12 (12-day preexposure) and Group 4 (4-day preexposure). Error bars represent ± 1 standard error of the mean (SEM).

Mean response latency (seconds) for discriminations involving AX+ vs. AY- and AX+ vs. BX- trials across six session blocks (each representing two training sessions). Group 1 (12-day preexposure) showed significantly higher latencies on AX+ AY- discriminations during early session blocks, indicating initial hesitation or slower responding to novel cue discriminations. Group 2 (4-day preexposure) showed no significant latency differences. Error bars represent ± 1 SEM.

PC1. PC1 composite scores, combining variance from multiple behavioral measures (number of responses, percent correct, latency), are shown in Figure 4. Higher PC1 scores indicate stronger overall discrimination. Both groups improved discrimination performance across the six session blocks (each block representing two of the twelve training sessions), confirmed by a highly significant main effect of SessionBlock ($F(5,70) = 30.69, p < .001$). Main effects of Group ($F(1,14) = 0.39, p = .545$) and Discrimination ($F(1,14) = 0.18, p = .677$) were not significant.

In Group 1 (12-day preexposure), rats showed a clear advantage for AX+ AY-trials beginning in Session 3, with the separation between AX+ AY- and AX+ BX- scores increasing across Sessions 4 through 6. This pattern mirrors the results seen in number of responses and percent correct, suggesting a consistent learning advantage when the test trial involved the less familiar cue. In contrast, Group 2 (4-day preexposure) showed significantly higher PC1 scores for AX+ BX- than AX+ AY- in Sessions 4 and 6, again reflecting the reversed pattern of discrimination seen with shorter preexposure.

However, there were significant interactions: Group $\times$ Discrimination ($F(1,14) = 16.05, p = .001$) and Group $\times$ Discrimination $\times$ SessionBlock ($F(5,70) = 4.85, p < .001$), suggesting that discriminations progressed differently depending on group and cue type. The Discrimination $\times$ SessionBlock interaction approached significance ($F(5,70) = 2.17, p = .067$).

Post hoc analyses clearly revealed distinct patterns for each group. Group 1 (12-day preexposure) showed significantly better discrimination (higher PC1) for AX+ AY- compared to AX+ BX- during sessions 3 ($p = .0028$), 4 ($p = .0019$), 5 ($p = .0038$), and 6 ($p = .0028$). This result indicates consistently stronger discrimination involving novel cue pairings (AY), suggesting effects consistent with latent inhibition.

Conversely, Group 2 (4-day preexposure) showed significantly better discrimination for AX+ BX- compared to AX+ AY- during session blocks 4 ($p = .0016$) and 6 ($p = .0346$), indicating enhanced discrimination for familiar cues (BX), supporting the presence of perceptual learning effects.

Taken together, these findings further confirm the differential discrimination outcomes associated with the duration of preexposure: longer preexposure (12 days) favors discrimination involving novel cues (latent inhibition), while shorter preexposure (4 days) favors discrimination of familiar cues (perceptual learning).




Figure 4. Mean PC1 composite scores for two discriminations: AX⁺ vs AY⁻ (novel-CS based discrimination) and AX⁺ vs BX⁻ (preexposed-CS based discrimination) across six two-session blocks, plotted separately for Group 12 (12-day preexposure) and Group 4 (4-day preexposure). Error bars represent ± 1 standard error of the mean (SEM).

Mean PC1 composite scores (based on number of responses, percent correct, and latency scores) for discriminations involving AX+ AY- and AX+ BX- trials across six session blocks (each representing two training sessions). Group 1 (12-day preexposure) demonstrated significantly stronger discrimination on AX+ AY- trials, consistent with the latent inhibition effect. Group 2 (4-day preexposure) demonstrated stronger discrimination on AX+ BX- trials, suggesting a perceptual learning effect. Error bars represent ± 1 SEM.

The findings consistently supported the hypothesis that increased stimulus preexposure hinders subsequent conditioning to that cue across all behavioral measures, including number of responses, percent correct, latency, and PC1 composite scores. Stronger discriminating favoring the novel cue (AX+AY-) was shown by Group 1, which was preexposed for 12 days. This suggests a strong latent inhibition effect. With only 4 days of preexposure, Group 2 displayed the opposite pattern, exhibiting greater discrimination on AX+ BX-trials.

These results show that, depending on the quantity of exposure, varying preexposure within the same experimental framework can produce either latent inhibition or perceptual learning effects. This study offers behavioral confirmation that perceptual learning and latent inhibition may be investigated with the same techniques and represent distinct stages of a learning continuum influenced by exposure history.

Discussion

In summary, Chapter 1 demonstrated that preexposure duration systematically influenced subsequent discrimination learning. Across all behavioral measures (auROC, percent correct, latency, and PC1), the 12-day preexposure group consistently exhibited stronger discrimination when the test discrimination involved novel cues (AX^+ vs AY^-), consistent with latent inhibition. In contrast, the 4-day preexposure group showed an advantage for the test discrimination involving the familiar cues (AX^+ vs BX^-), providing evidence of perceptual learning (Gibson, 1969; Honey & Hall, 1989).

These findings support Hypothesis 2, which proposed that the amount of preexposure is the critical factor determining whether LI or PL emerges. Extensive preexposure produced slower learning to the preexposed cue, while moderate preexposure facilitated discrimination involving familiar cues. In contrast, the results did not support Hypothesis 1, which emphasized the mode of presentation as the decisive factor (Hall, 2003; McLaren, Kaye, & Mackintosh, 1989). To my knowledge, this is the first demonstration of LI and PL under identical testing conditions, allowing a direct comparison between the mechanisms involved in these paradoxical effects.

Chapter 2:

Assessing Latent Inhibition Within-Modality Using the Same Within-Subject Design: Analysis of Individual Differences

Introduction

The magnitude of latent inhibition (LI) can vary substantially depending on factors such as stimulus properties, preexposure duration, and task design (Hall, 1991; Lubow, 1989, 2005; Weiner, 2003). One factor known to influence LI is the sensory modality relationship between preexposed and novel cues. Studies comparing cross-modality arrangements, in which preexposed and novel cues differ in modality, with same-modality arrangements, in which all cues belong to the same modality, have found that LI is generally more robust in the cross-modality case. By contrast, when preexposed and novel cues are drawn from the same modality, the LI effect is often smaller or more variable (Escobar, Arcediano, & Miller, 2002; Honey & Hall, 1989). This attenuation may arise from greater perceptual similarity between cues, which can increase generalization, or from shifts in attentional allocation when stimuli share overlapping sensory features.

In Chapter 1, I employed a within-subject design that operationalized LI as slower learning about a preexposed cue compared to a novel cue, measured using two discrimination indices: AX^+ vs BX^- (preexposed-CS based discrimination) and AX^+ vs AY^- (novel-CS based discrimination). When cues were drawn from different modalities, this design revealed a clear LI effect, with slower learning on the preexposed discrimination.

The present experiment applied the same design to a within-modality context to test whether LI remains detectable when all cues share the same sensory properties. This allowed for a direct comparison of LI expression across modality conditions while holding other procedural variables constant. Additionally, I examined potential individual differences in LI expression, including the effects of sex and cue-counterbalancing condition, to determine whether these factors might contribute to variability in performance. Based on prior findings, I expected LI to be attenuated under within-modality conditions, but predicted that systematic individual differences might still emerge and help explain variability in discrimination performance.

Methods and Materials

Experimental animals. Twenty sex-balanced adult Long-Evans rats (*Rattus norvegicus*), approximately 3 months of age, served as subjects. Males weighed between 522–715 g, and females weighed between 264–522 g. All animals were experimentally naïve and sourced from Charles River Laboratories. They were housed individually in standard clear-plastic tubs (10.5 in $\times$ 19 in $\times$ 8 in, Charles River Laboratories) with woodchip bedding, within a colony room maintained on a 14:10 light/dark cycle. Behavioral sessions were conducted between 7–10 hours after the onset of the light phase. Water access was restricted to 1 hour/day following each

session, while a 10% sucrose solution and standard laboratory chow were provided ad libitum as necessary. Rats were assigned to one of four cue–counterbalancing conditions ($n = 5$ per condition) so that each animal experienced a unique mapping of preexposed (A, B) and novel (X, Y) cues within the same modality.

Apparatus. The apparatus was identical to that used in Chapter 1.

Procedure. The procedure followed the same general structure as in Chapter 1, except that all cues (A, B, X, and Y) were drawn from the same sensory modality. The experiment consisted of three stages: preexposure, conditioning, and latent inhibition testing. Stimulus identities for A, B, X, and Y were counterbalanced across animals to form four cue–counterbalancing conditions ($n = 5$ per condition).

Stage 1: Preexposure. Rats received daily sessions in which cues A and B were presented individually without reinforcement. Each session comprised 16 A trials and 16 B trials (32 total), intermixed in a pseudorandom order with no more than two consecutive presentations of the same cue. Each cue was presented for 10 s, followed by no sucrose delivery. Intertrial intervals (ITIs) averaged 50 s (range: 30–70 s).

Stage 2: Conditioning. Across 12 daily sessions, rats were trained on two separate discriminations: AX+ vs AY- (novel-CS based discrimination) and AX+ vs BX- (preexposed-CS based discrimination). Each session comprised 16 AX+ trials, 8 AY- trials, and 8 BX- trials. On reinforced trials (AX+), the compound stimulus ended with 3-s access to a sucrose-filled dipper; on nonreinforced trials (AY-, BX-), the trial ended without reward. Stimulus presentations were 10 s in duration, and ITIs matched those in Stage 1.

Stage 3: Latent Inhibition Test. Discrimination training continued under the same parameters, allowing for direct assessment of learning rates for novel versus preexposed discriminations. Latent inhibition was defined as slower acquisition of the preexposed discrimination (AX+ vs BX-) relative to the novel discrimination (AX+ vs AY-).

Results

The results are presented for each quantitative measure of conditioned magazine approach. Analyses first addressed the presence of latent inhibition across the various quantitative measures of magazine activity.




Figure 1. Mean auROC scores for two discriminations: AX+ vs AY- (novel-CS based discrimination) and AX+ vs BX- (preexposed-CS based discrimination) across six two-session blocks, plotted separately for four measures of magazine approach: (A) number of entries, (B) percent time, (C) latency, and (D) PC1 composite. Error bars represent ± 1 SEM.

Number of Entries. Group mean AUROC scores across the six two-session blocks are shown in Figure 1A, depicting performance on the two within-subject discriminations: AX+ vs AY- and AX+ vs BX-. Both discriminations improved steadily over training, indicating that rats learned to distinguish rewarded from non-rewarded cues in terms of their number of magazine entries.

The figure also shows that both discriminations were acquired at similar rates, as indicated by a significant main effect of SessionBlock ($F(2.75, 52.26) = 4.66, p = .007$, Greenhouse–Geisser corrected), but no main effect of Discrimination ($F(1, 19) = 0.29, p = .598$) and no Discrimination $\times$ Block interaction ($F(3.63, 68.93) = 0.83, p = .502$). These results indicate that

rats' number of magazine entries increased across blocks, reflecting overall learning, but the rate and strength of learning did not differ between AX+ vs AY- and AX+ vs BX- discriminations.

Post hoc analyses confirmed that AUROC scores for AX+ vs AY- and AX+ vs BX- were statistically indistinguishable within every session block (all $p > .12$). This pattern suggests that, for the number-of-entries measure, there was no evidence of a latent inhibition effect under within-modality conditions. Both discriminations followed parallel trajectories, with no selective impairment or enhancement of learning for either cue.

Percent Time. Group mean AUROC scores for percent time spent in the magazine are shown in Figure 1B, separately for AX+ vs AY- and AX+ vs BX- discriminations. Both conditions showed steady improvement in discrimination over the six two-session blocks, with AX+ vs AY- consistently yielding higher AUROC scores than AX+ vs BX-, indicating stronger latent inhibition effects in percent time.

The discriminations were not learned equally well across conditions, as confirmed by a significant main effect of Discrimination ($F(1, 19) = 12.95, p = .002$) and a significant main effect of SessionBlock ($F(2.03, 38.65) = 4.37, p = .019$, Greenhouse–Geisser corrected). There was no Discrimination $\times$ Block interaction ($F(3.67, 69.63) = 1.15, p = .34$), indicating that while both discriminations improved with training, the performance advantage for AX+ vs AY- was stable across blocks.

Post hoc analyses further revealed that AX+ vs AY- discrimination produced significantly higher AUROC scores than AX+ vs BX- in session blocks 2 ($p = .019$), 3 ($p = .004$), and 5 ($p = .014$), with a marginal trend in block 4 ($p = .076$), and no significant differences in blocks 1 or 6 ($p > .25$). This pattern demonstrates a robust latent inhibition effect within modality for percent time, with the effect most pronounced in the middle training blocks and maintained throughout the majority of the sessions. The results confirm that percent time is the most sensitive measure for detecting LI under these conditions, consistently favoring superior discrimination involving the novel cue.

Thus, these results indicate that latent inhibition can be detected within a single sensory modality when measured by percent time in the magazine. The consistently higher AUROC scores for AX+ vs AY- reflect stronger discrimination for the novel cue, supporting the presence of latent inhibition. This effect was most pronounced in the middle session blocks, highlighting that within-modality LI is both present and most easily detected with sensitive behavioral measures like percent time.

Latency. Group mean AUROC scores based on latency to enter the magazine are shown in Figure 1C, separately for AX+ vs AY- and AX+ vs BX- discriminations. Both discriminations improved over the six two-session blocks, with overall increases in AUROC scores indicating learning. Although the two conditions followed a similar trajectory, AX+ vs AY- tended to

outperform AX+ vs BX-, particularly in the middle blocks, suggesting a possible but limited latent inhibition effect.

The discriminations were not learned equally well across blocks, as indicated by a highly significant main effect of SessionBlock ($F(3.43, 65.09) = 11.69, p < .001$, Greenhouse–Geisser corrected), but only a marginal main effect of Discrimination ($F(1, 19) = 2.88, p = .1058$), and no Discrimination $\times$ Block interaction ($F(3.64, 69.11) = 1.51, p = .3838$). This pattern indicates that rats' latency-based discrimination improved over training, but the difference between AX+ vs AY- and AX+ vs BX- was less consistent compared to the percent time measure.

Post hoc analyses revealed that AX+ vs AY- discrimination produced significantly higher AUROC scores than AX+ vs BX- only in block 3 ($p = .024$), with no significant differences observed in any other block (all $p > .13$). This suggests a limited and transient latent inhibition effect in latency, emerging only at the midpoint of training. Both discriminations showed improvement across blocks, but the discrimination effect was more restricted and less robust compared to percent time.

Thus, these findings suggest that, while rats improved their latency-based discrimination with training, evidence for latent inhibition in this measure was limited to a brief window in the middle of the experiment. The LI effect observed in latency was transient and less robust than that seen with percent time, indicating that the sensitivity of latency as an index for latent inhibition within modality is relatively weak and may only emerge under specific training conditions.

PC1. Group mean AUROC scores based on the first principal component (PC1) are shown in Figure 1D, separately for AX+ vs AY- and AX+ vs BX- discriminations. Both discriminations showed steady improvement over the six two-session blocks, with AUROC scores increasing for both conditions as training progressed. Although AX+ vs AY- discriminations tended to yield slightly higher AUROC scores, especially around the middle training blocks. The difference between the two conditions diminished by the end of training.

A mixed-design ANOVA on the PC1 dataset revealed a highly significant main effect of Block ($F(2.82, 53.66) = 18.02, p < .001$, Greenhouse–Geisser corrected), reflecting strong improvement in discrimination over training. The main effect of Discrimination was marginally significant ($F(1, 19) = 3.05, p = .097$), with no Discrimination $\times$ Block interaction ($F(3.72, 70.73) = 1.70, p = .16$). This pattern suggests overall learning for both discriminations, with a modest advantage for AX+ vs AY-.

Post hoc analyses showed that AUROC scores for AX+ vs AY- were significantly higher than AX vs BX- only in block 3 ($p = .035$), with all other blocks showing no significant difference between discriminations (all $p > .26$). This block-specific effect indicates a brief, mid-training

period where latent inhibition may have influenced PC1-based discrimination, but the effect was not sustained across blocks.

These results suggest that PC1, as a composite measure, captures robust learning for both discriminations but only weak evidence of within-modality latent inhibition. Any LI effect in PC1 was brief, emerging only mid-training before both discriminations converged by the end of the experiment. This highlights that, like latency, PC1 is less sensitive than percent time for detecting sustained LI effects under these conditions.

To sum up, across all measures, discrimination performance improved with training, but LI effects were modest, inconsistent across measures, and marked by large individual variability. Only percent time in the magazine consistently revealed a significant LI effect. Given this effect, I next examined how well the measures correlated.

Correlations among dependent variables




Figure 2. Correlations among discrimination indices based on auROC distance across subjects: (A) number vs percent, (B) number vs latency, and (C) percent vs latency. Error bars represent ± 1 SEM.

Number vs. Percent. The correlation between the number of magazine entries and the percent of time spent in the magazine during discrimination trials was weak and not statistically significant. As shown in Figure 2A, the Pearson correlation coefficient was $r = 0.236$ with a p-value of .317 and a 95% confidence interval ranging from -0.231 to 0.614 . Although the scatterplot indicates a slight positive relationship, there was considerable variability in the data, and the association did not reach significance. This finding suggests that number and percent measures of magazine approach do not consistently reflect the same underlying pattern of discrimination performance and may be capturing partially independent aspects of behavior.

Number vs. Latency. A moderate positive correlation was observed between the number of magazine entries and the latency auROC distance, with Pearson's $r = .574$ and a p-value of 0.008. The 95% confidence interval ranged from 0.176 to 0.811, and the relationship was statistically significant. As shown in Figure 2B, this result indicates that animals exhibiting greater discrimination by number of entries also tended to show greater discrimination in latency, suggesting that these two measures capture related aspects of learning.

Percent vs. Latency. A moderate positive correlation was observed between percent auROC distance and latency auROC distance across subjects, with a Pearson correlation coefficient of $r = 0.544$ and a p-value of 0.013. The 95% confidence interval ranged from 0.134 to 0.795, indicating that the relationship was statistically significant. This result, as shown in Figure 2C, shows that animals exhibiting greater discrimination in percent time tended to also show greater discrimination in latency.

To conclude, the dependent variables were not particularly well correlated. Given that percent time spent in the magazine revealed a reliable latent inhibition effect with cues belonging to the same sensory modality, the subsequent analyses focus on this measure.

Examining individual differences in latent inhibition

The inclusion of 20 subjects in this study allowed for a preliminary investigation of individual differences in the expression of latent inhibition. To this end, I calculated individual auROC scores based on each of the subdiscriminations presented in Stage 2 (AX+ vs. AY- and AX+ vs. BX-), focusing on the percent time measure. The degree of latent inhibition expressed by each rat is captured by the distance of each circle mark to the identity (diagonal) line. Thus, marks falling along this line represent animals that solved the two subdiscriminations equally well, marks that fall below this line depict animals that showed latent inhibition, and those above the line represent animals showing the opposite pattern of behavior.




Figure 3. Distribution of individual differences in percent auROC distance scores: (A) scatterplot of individual average scores for novel-CS based (AX+ vs AY-) and preexposed-CS based (AX+ vs BX-) discriminations, (B) histogram of signed distances from the identity diagonal, (C) descriptive statistics, (D) sex-specific scatterplot of average scores, and (E) histogram of distances separated by sex. Error bars represent ± 1 SEM.

Inspection of the scatterplot reveals notable individual differences in the expression of latent inhibition (Figure 3A). To further examine individual differences in latent inhibition, I calculated for each animal the perpendicular distance between each marking and the identity line and plotted the distribution of these distance values (Figure 3B). Descriptive statistics are provided in Figure 3C, and the second scatterplot separated by sex is shown in Figure 3D. The mean percent auROC distance was 0.10 (SD = 0.12), with values ranging from -0.02 to 0.41 and a median of 0.05. The distribution was positively skewed (skewness = 1.46), indicating that most animals clustered at lower scores, while a smaller number showed much higher levels of latent inhibition. The Shapiro-Wilk test confirmed significant deviation from normality ($p = 0.001$). This pattern is visualized in the colored histogram (Figure 3E), which highlights the predominance of lower scores and the presence of a long right tail. Taken together, these results indicate that while most subjects showed modest or no latent inhibition effects under the current within-modality conditions, a subset displayed a strong effect as measured by percent time in the magazine.

Can sex differences explain individual differences in the expression of latent inhibition?

Given that the sample of rats in this study consisted of both males ($n = 12$) and females ($n = 8$), I examined whether sex might account for the variability observed in latent inhibition. Visual inspection of the histogram (Figure 3D) and scatterplot (Figure 3E) indicated substantial overlap between the sexes, with no obvious separation in percent auROC distance scores. To further investigate whether sex might contribute to the variability observed in latent inhibition, I examined percent auROC distance scores separately for males and females. Initial visual inspection of the data revealed no obvious large-scale separation between sexes, as both males and females displayed a wide range of scores with substantial overlap. However, given the potential for temporal patterns to emerge across training, a mixed-design Sex $\times$ Session Block ANOVA on distance scores was conducted to test for differences in how latent inhibition developed over time.




Figure 4. Mean auROC distance to the identity diagonal across six session blocks, shown separately for males and females. Error bars represent ± 1 SEM.

The mixed-design Sex $\times$ Session Block ANOVA on distance scores, using the Huynh–Feldt correction, revealed a significant interaction between sex and session block, $F(5,90) = 2.49$, $p = 0.037$. Post hoc comparisons indicated that this effect was driven primarily by a difference in Session Block 2, where females showed higher percent auROC distance scores than males, reflecting a stronger latent inhibition effect early in training. Females maintained relatively high levels through Blocks 3 and 4, whereas males displayed little latent inhibition overall, with any effects being restricted to isolated session blocks. Although the difference was not large, these results suggest that sex may modulate the expression of latent inhibition at specific points in training, with females sustaining higher discrimination performance across multiple blocks and males showing more limited and sporadic effects.




Figure 5. Mean percent auROC scores for two discriminations: AX+ vs AY- (novel-CS based discrimination) and AX+ vs BX- (preexposed-CS based discrimination) across six two-session blocks, shown separately for (A) males and (B) females. Error bars represent ± 1 SEM.

To further assess how latent inhibition developed in each sex, mean percent auROC scores for AX+ vs AY- and AX+ vs BX- discriminations were plotted across session blocks separately for males and females. As shown in Figure 5A, males exhibited little or no evidence of latent inhibition, with discrimination performance remaining relatively flat across most session blocks. In contrast, Figure 5B shows that females demonstrated a pronounced latent inhibition effect early in training, with higher auROC scores in Session Block 2 and sustained elevated scores through Blocks 3 and 4. Although this sex difference was not dramatic, it suggests that females tended to express stronger and more persistent latent inhibition during the early and middle phases of training compared to males.

Can counterbalancing conditions explain individual differences in the expression of latent inhibition?

While the individual differences observed are pronounced, it is possible that they do not reflect differences in latent inhibition per se, but rather differences in the ease with which the subdiscriminations as instantiated in the various counterbalancing conditions could be solved, or interaction between the two (latent inhibition manifesting for some cues, but not others). To examine this possibility, I performed a between-subjects ANOVA comparing latent inhibition performance as captured by distance-to-identity-line scores (Figure 5).




Figure 6. Distribution of percent auROC distance scores across the four cue-counterbalancing groups ($n = 5$ per group). Error bars represent ± 1 SEM.

This analysis revealed a significant effect of counterbalancing group on Percent auROC distance, $F(3,16) = 6.57$, $p = 0.004$. Post hoc Tukey tests indicated that Group 4 performed significantly higher than all other groups on this measure (Group 4 mean = 0.253), with adjusted p -values of 0.010, 0.006, and 0.028 for comparisons with Groups 1, 2, and 3, respectively. In contrast, Groups 1, 2, and 3 did not significantly differ from one another (all $p > 0.9$). As shown in Figure 6, animals in Group 4 consistently showed higher percent auROC distance scores compared to the other groups, who tended to cluster at lower values. This suggests that Group 4 demonstrated stronger discrimination performance overall. These results point to a strong counterbalancing effect in the percent time measure, indicating that the presence or strength of latent inhibition in Group 4 depended on which cues were assigned to each discrimination. In other words, how the cues were paired seemed to influence the extent to which LI was expressed.

Thus, both sex and counterbalancing condition appeared to have had an influence on individual differences in the expression of latent inhibition.

Compared to Chapter 1, the latent inhibition effect in the within-modality design appeared weaker and more variable, with greater overlap in discrimination performance between novel and preexposed conditions. This pattern is consistent with previous findings that modality separation enhances LI by reducing generalization between cues (Escobar, Arcediano, & Miller, 2003; Killcross & Dwyer, 1999).

Discussion

In summary, Chapter 2 revealed a modest latent inhibition effect when all cues were drawn from the same sensory modality. Percent time in the magazine showed reliably slower acquisition for the preexposed discrimination (AX+ vs BX-) compared to the novel discrimination (AX+ vs AY-). However, this effect was absent in number of entries, latency, and PC1, suggesting that LI was not robust across measures. This pattern aligns with previous findings that within-modality LI is weaker and less reliable than cross-modality LI (Escobar, Arcediano, & Miller, 2003; Killcross & Dwyer, 1999).

Compared to the robust cross-modality effect observed in Chapter 1, the LI effect here was smaller in magnitude and less consistent. One likely explanation is that greater perceptual similarity among cues belonging to the same modality increased generalization and reduced the distinctiveness of the preexposed stimulus, thereby attenuating the impact of preexposure on subsequent discrimination learning (Lubow, 1989; Hall, 2001).

Individual differences were also more pronounced under within-modality conditions. Substantial variability emerged across animals, counterbalancing conditions, and sexes, though these factors did not provide consistent predictors of LI magnitude. These findings suggest that the expression of LI under within-modality conditions is not only weaker overall, but also more sensitive to idiosyncratic and procedural factors.

Taken together, these results extend the findings of Chapter 1 by demonstrating that while LI can be observed under both cross-modality and within-modality arrangements, its magnitude and reliability depend strongly on cue distinctiveness. By showing that LI is attenuated when cues share overlapping sensory properties, the present experiment highlights the role of generalization and attentional competition in constraining the robustness of latent inhibition.

Chapter 3:

Isolating Attentional from Interfering Processes in Latent Inhibition

Introduction

As discussed in the General Introduction, latent inhibition (LI) can be explained by two competing accounts. The attentional account proposes that repeated preexposure to a conditioned stimulus (CS) reduces its associability, leading to slower learning when that CS is later paired with reinforcement (Lubow & Moore, 1959; Mackintosh, 1975; Pearce & Hall, 1980). By contrast, the interference account suggests that preexposure results in a learned CS → no outcome association that interferes with the acquisition of a new CS → US association at test (Bouton, 1993; Miller & Escobar, 2001). This theoretical ambiguity complicates interpretation of LI effects: when a brain region or process is shown to be critical for LI, it is often unclear whether it supports attentional control or resolves associative interference.

LI is clinically and cognitively relevant, making it important to understand why it occurs. Both accounts have empirical support. For example, a decline in orienting to preexposed cues parallels the behavioral LI effect, consistent with an attentional decrement explanation (Lubow, Weiner, & Schnur, 1981). On the other hand, LI can be abolished by context shifts introduced between preexposure and conditioning, supporting an interference-based account (Bouton, 1993). Thus, the available evidence leaves open whether LI reflects a loss of attention to the cue, interference from a competing CS → “nothing” association, or a combination of both.

The present experiment was designed to address this question by adapting the within-subject design used in Chapters 1 and 2. As in those studies, LI was measured as the relative delay in solving a discrimination based on preexposed cues (AX+ vs BX-) compared with one based on novel cues (AX+ vs AY-). However, unlike the earlier experiments in which A and B were presented without consequence during preexposure, here pretraining was introduced to minimize interference. Specifically, cue A was trained as a redundant cue (L+ and LA+), and cue B was trained to signal the omission of reward (L+ and LB-). If animals acquired a B → no food association, this learning should facilitate acquisition of the AX+ vs BX- discrimination at test, thereby removing the possibility that interference could account for any observed LI effect.

Under these conditions, if discrimination involving the preexposed cues (AX+ vs BX-) is still retarded relative to the novel-cue discrimination (AX+ vs AY-), this would provide strong evidence for a loss of associability caused by preexposure, consistent with an attentional mechanism. This logic parallels the negative transfer effect described by Hall and Pearce (1979). In their seminal experiments, when a CS was first paired with a weak unconditioned stimulus (US) and later paired with a stronger US, learning of the new association was retarded relative to a control CS. Within the Pearce–Hall model, this effect was interpreted as reduced associability:

once a stimulus predicts an outcome, the need to attend to it diminishes, slowing subsequent learning even if the outcome changes in magnitude.

The design used here extends this reasoning. By equating the associative demands of preexposure and test while holding the outcome constant (food omission), the experiment rules out an interference-based explanation. Thus, any observed delay in learning about the preexposed cues would isolate an attentional decrement as the underlying mechanism of latent inhibition.

Methods and Materials

Experimental Animals. Thirty-two sex-balanced adult Long–Evans rats (*Rattus norvegicus*), approximately 3 months of age, served as subjects. Males weighed 452–721 g, while females weighed 277–340 g at the start of the study. All animals were experimentally naïve, sourced from Charles River Laboratories, and housed individually in standard clear-plastic tubs (10.5 in × 19 in × 8 in; Charles River Laboratories) with woodchip bedding, within a colony room maintained on a 14:10 light/dark cycle. Behavioral sessions were conducted 7–10 hr after light onset. Water access was restricted to 1 hr per day following each session, while food was provided ad libitum. Rats were assigned to one of four cue–counterbalancing conditions ($n = 8$ per condition) such that each animal experienced a unique mapping of preexposed (A, B) and novel (X, Y) cues within the same modality. Auditory cues consisted of a click (80 dB), white noise (70 dB), low tone (1 kHz, 80 dB), and high tone (12 kHz, 70 dB), with assignments counterbalanced across subjects. All animal care and experimental procedures complied with the ARRIVE guidelines and the NIH Guide for the Care and Use of Laboratory Animals, with protocols approved by the Brooklyn College Institutional Animal Care and Use Committee.

Apparatus. The apparatus was identical to that described in Chapters 1 and 2.

Procedure. The design followed the within-subject latent inhibition paradigm used in Chapters 1 and 2, with the addition of a pretraining phase to establish differential associative histories for cues A and B prior to LI testing. The experiment consisted of three phases: conditioning, pretraining, and LI testing. All sessions were conducted in Med Associates operant-conditioning chambers under the conditions described in the Apparatus section. All cues (A, B, X, Y) were auditory, and cue identities were fully counterbalanced across animals.

Phase 1: Conditioning. Rats first received one daily session of 60 trials in which a panel light (L) was presented alone and reinforced with sucrose delivery (L+). Each trial consisted of a 10-s light presentation followed by 3-s access to the sucrose dipper. Intertrial intervals (ITIs) averaged 50 s (range: 30–70 s). This phase established the light as an excitatory cue.

Phase 2: Pretraining. Over four daily sessions, the light continued to be presented alone (L+), in compound with cue A and reinforced (LA+), and in compound with cue B without reinforcement (LB-). Cue A was also presented alone without reinforcement (A-), as was cue B (B-). Each session comprised 20 L+, 10 LA+, 10 LB-, 10 A-, and 10 B- trials, intermixed pseudorandomly with no more than two identical trial types in succession. All cues were 10 s in duration, and ITIs matched those in Phase 1. This arrangement was designed to render cue B a conditioned inhibitor (predictor of reward omission) and cue A a nonpredictive stimulus.

Phase 3: Latent Inhibition Test. LI was assessed by training two concurrent subdiscriminations: a novel discrimination (AX+ vs AY-) and a preexposed discrimination (AX+ vs BX-). Each daily session consisted of 20 AX+, 10 AY-, and 10 BX- trials, intermixed pseudorandomly with no more than two consecutive presentations of the same type. On reinforced trials (AX+), the compound stimulus terminated with sucrose delivery; on nonreinforced trials (AY-, BX-), no reward was delivered. Each trial consisted of a 10-s pretrial, 10-s CS period, and a 3-s US period on reinforced trials, followed by a 10-s posttrial. ITIs matched those used in earlier phases. LI was defined as slower acquisition of the preexposed discrimination (AX+ vs BX-) relative to the novel discrimination (AX+ vs AY-).

Data Analysis. Magazine activity was quantified using the same measures as in Chapters 1 and 2: number of head entries, percent time spent in the magazine, and latency to enter during the 10-s CS period. A composite index (PC1) was derived from a principal component analysis (PCA) of these measures. Discrimination performance was assessed using the area under the receiver operating characteristic curve (auROC) for each subdiscrimination (AX+ vs AY-; AX+ vs BX-) in each session. auROC scores were centered and rescaled so that 0.5 indicated chance performance and 1.0 indicated perfect discrimination.

Mixed-design ANOVAs were conducted separately for each measure, with Subdiscrimination (AX+ vs AY-; AX+ vs BX-) and Session Block as within-subjects factors, and Counterbalancing Group and Sex as between-subjects factors. Significant main effects and interactions were followed by planned comparisons and Tukey-adjusted post hoc tests. Although the omnibus ANOVAs did not consistently yield significant effects, planned within-discrimination contrasts were conducted at each session block because the hypotheses specified a directional comparison between preexposed-CS and novel-CS discriminations. These targeted comparisons were scientifically justified given the theoretical interest in distinguishing attentional from interference accounts of latent inhibition. All p-values were adjusted using Tukey's HSD procedure to correct for multiple testing.

In addition, a unique LI score for each animal was calculated by taking the perpendicular distance between each auROC score pair (AX+ vs BX- and AX+ vs AY-) and the identity line. These LI scores were used to analyze individual differences, including correlations across behavioral measures and with other indices of performance.

Results

Examining group-level performance during Stage 2 (conditioned inhibition training)

During the pretraining phase, rats successfully learned the discrimination required for cue B to function as a conditioned inhibitor. Across sessions, responding was reliably lower on LB- trials relative to LA+ trials, indicating that B acquired inhibitory properties by signaling the omission of reward when compounded with the light cue. This ensured that, at the start of the latent inhibition (LI) test, cue B served as an established signal for reward omission, while cue A remained a nonpredictive stimulus.




Figure 1. Mean auROC scores for two conditioned-inhibition discriminations: L+ vs LA+ (control) and L+ vs LB- (conditioned-inhibitor) across 28 session blocks of Stage 2, plotted separately for (A) Number of Entries, (B) Percent Time, (C) Latency, and (D) PC1 composite. Error bars represent ± 1 SEM.

Number of Entries. Group mean auROC scores across the 28 conditioned inhibition sessions are shown in Figure 1A for the two discriminations: L+LA+ (control) and L+LB- (conditioned inhibitor). Performance on the L+LB- discrimination increased rapidly during the initial sessions, rising from approximately 0.30 auROC in Session 1 to around 0.70–0.75 by the end of training, whereas scores for the L+LA+ discrimination remained near zero throughout. This pattern

indicates that rats reliably learned to distinguish rewarded from non-rewarded compounds, with cue B serving as a consistent signal for reward omission. The trajectories for the two discriminations diverged early and remained widely separated for the duration of training, reflecting a robust conditioned inhibition effect.

A repeated-measures ANOVA on auROC scores revealed a significant main effect of Session, $F(12.77, 396) = 49.84$, $p < .001$, Greenhouse–Geisser corrected, $\eta^2 = 0.116$, reflecting overall improvement in discrimination performance across sessions. There was also a significant main effect of Discrimination, $F(1, 31) = 258.58$, $p < .001$, $\eta^2 = 0.642$, with higher auROC scores for the L+LB- discrimination. The Session $\times$ Discrimination interaction was significant, $F(12.77, 396) = 36.78$, $p < .001$, $\eta^2 = 0.126$, indicating that the rate of learning differed between the two discriminations. Post hoc comparisons confirmed that auROC scores for L+LB- were significantly higher than for L+LA+ from Session 3 onward (all $p < .001$). Thus, animals successfully acquired the conditioned inhibition contingency in terms of number of magazine entries, with rapid and stable suppression of responding to the inhibitory cue.

Percent. Group mean auROC scores across the 28 conditioned inhibition sessions are shown in Figure 1B for the two discriminations: L+LA+ (control) and L+LB- (conditioned inhibitor). Performance on the L+LB- discrimination increased sharply during the first few sessions, rising from approximately 0.35–0.40 auROC at the start to a stable high level of around 0.75 by the end of training. In contrast, scores for the L+LA+ discrimination remained near zero throughout. This large and consistent separation between the two discriminations suggests that percent time in magazine captured the conditioned inhibition effect robustly.

A repeated-measures ANOVA on auROC scores revealed a significant main effect of Session, $F(6, 186) = 17.54$, $p < .001$, Greenhouse–Geisser corrected ($p_{\text{HF}} < .001$), $\eta^2 = 0.1197$, indicating reliable improvement in discrimination performance across training. There was also a significant main effect of Discrimination, $F(1, 31) = 2.28$, $p = .141$, $\eta^2 = 0.0029$, as well as a significant Session $\times$ Discrimination interaction, $F(6, 186) = 4.00$, $p < .001$, $\eta^2 = 0.0019$, indicating that the rate of improvement differed between the two cue types. Post hoc comparisons confirmed that auROC scores for L+LB- were significantly higher than for L+LA+ in multiple sessions ($p < .05$). Thus, animals successfully acquired the conditioned inhibition contingency in terms of percent time in magazine, with the inhibitory cue reliably suppressing responding relative to the control cue.

Latency. Group mean auROC scores across the 28 conditioned inhibition sessions are shown in Figure 1C for the two discriminations: L+LA+ (control) and L+LB- (conditioned inhibitor). Scores for the L+LB- discrimination rose modestly during the early sessions, reaching a moderate plateau of approximately 0.35–0.45 auROC, while scores for the L+LA+ discrimination remained near zero throughout training. The separation between the two curves

was smaller and more variable than for the Number and Percent measures, suggesting a weaker conditioned inhibition effect for latency to first magazine entry.

A repeated-measures ANOVA on auROC scores revealed a significant main effect of Session, $F(27, 837) = 60.16$, $p < .001$, Greenhouse–Geisser corrected ($p_{\text{HF}} < .001$), $\eta^2 = 0.0850$, indicating overall improvement in performance across sessions. There was also a significant main effect of Discrimination, $F(1, 31) = 139.76$, $p < .001$, $\eta^2 = 0.3644$, as well as a significant Session $\times$ Discrimination interaction, $F(27, 837) = 38.93$, $p < .001$, $\eta^2 = 0.0610$, showing that the rate of improvement differed between the two cue types. Post hoc comparisons confirmed that L+LB- scores were significantly higher than L+LA+ scores in most sessions ($p < .05$), consistent with effective acquisition of the conditioned inhibition contingency for this measure.

PC1. Group mean auROC scores across the 28 conditioned inhibition sessions are shown in Figure 1D for the two discriminations: L+LA+ (control) and L+LB- (conditioned inhibitor). Scores for the L+LB- discrimination rose sharply in the early sessions, reaching high levels and maintaining a consistent advantage over the L+LA+ discrimination for the remainder of training. In contrast, L+LA+ scores stayed near zero across all sessions. This pattern suggests that the multivariate PC1 measure captured a conditioned inhibition effect comparable to that observed in the Number and Percent measures.

A repeated-measures ANOVA on auROC scores revealed a significant main effect of Session, $F(13, 403.96) = 55.28$, $p < .001$, Greenhouse–Geisser corrected ($p_{\text{HF}} < .001$), $\eta^2 = 0.0499$, indicating reliable improvement in performance across training. There was also a significant main effect of Discrimination, $F(1, 31) = 10.17$, $p = .003$, $\eta^2 = 0.1731$, and a significant Session $\times$ Discrimination interaction, $F(27, 837) = 46.56$, $p < .001$, $\eta^2 = 0.0127$, reflecting greater and faster gains for the L+LB- discrimination. Post hoc comparisons confirmed that L+LB- scores were significantly higher than L+LA+ scores in nearly every session from Session 3 onward (all $p < .05$). These results indicate that prior cue pairing with reward omission (B) produced a clear and persistent conditioned inhibition effect when measured with the composite PC1 index.

In summary, across all four quantitative measures of magazine activity (Number of Entries, Percent Time in Magazine, Latency to First Magazine Entry, and the composite PC1 index), rats reliably acquired the conditioned inhibition discrimination (L+LB-) while showing little to no responding to the control discrimination (L+LA+). Performance for L+LB- rose rapidly in the early sessions and remained well above L+LA+ levels throughout the remainder of training, indicating that cue B was effectively learned as a signal for reward omission. This pattern was evident in each measure, with the strongest and most consistent separation between discriminations observed for Number, Percent, and PC1, and a somewhat weaker but still reliable

effect for Latency. These results confirm that Stage 2 successfully established B as a conditioned inhibitor, providing the necessary preexposure for the subsequent latent inhibition test in Stage 3.

Examining group-level performance during Stage 3 (latent inhibition test)

In Stage 3, group mean performance was examined for the two within-subject discriminations: AX+ vs AY- (novel) and AX+ vs BX- (preexposed). Figure 2A–D shows that both discriminations improved steadily across the seven test sessions, increasing from near-chance levels in early sessions to approximately 0.65–0.70 auROC by the end of training.




Figure 2. Mean auROC scores for two latent-inhibition test discriminations: AX+ vs AY– (novel-CS based) and AX+ vs BX– (preexposed-CS based) across seven sessions of Stage 3, plotted separately for (A) Number of Entries, (B) Percent Time, (C) Latency, and (D) PC1 composite. Error bars represent ± 1 SEM.

Number of Entries. Group mean auROC scores across the seven LI test sessions are shown in Figure 2A for the two within-subject discriminations: AX+ vs AY- (novel) and AX+ vs BX- (preexposed). Both discriminations improved steadily over training, rising from approximately 0.30 auROC in Session 1 to around 0.65–0.70 by Session 7. The trajectories were largely parallel, with the novel discrimination tending to score slightly higher from mid-training onward, although this difference was not statistically reliable.

A repeated-measures ANOVA revealed a significant main effect of Session, $F(6, 186) = 14.83$, $p < .001$, Greenhouse–Geisser corrected ($p_{\text{HF}} < .001$), $\epsilon = 0.0847$, indicating overall improvement across sessions. There was no significant main effect of Discrimination, $F(1, 31) = 0.39$, $p = .532$, and no Session $\times$ Discrimination interaction, $F(6, 186) = 0.09$, $p = .997$, indicating similar acquisition rates and endpoints. Post hoc comparisons confirmed no significant differences between AX+ vs AY- and AX+ vs BX- in any session (all $p > .12$). Thus, prior cue exposure did not significantly affect learning for this measure.

Percent. Group mean auROC scores across the seven LI test sessions are shown in Figure 2B. Both discriminations improved steadily, increasing from approximately 0.30–0.35 in early sessions to about 0.65–0.70 by the final session. The novel discrimination tended to score higher throughout, with a more visible separation than for Number of Entries, but this difference was not statistically reliable.

A repeated-measures ANOVA revealed a significant main effect of Session, $F(3.74, 115.89) = 11.43$, $p < .001$, Greenhouse–Geisser corrected, indicating reliable improvement over time. There was no significant main effect of Discrimination, $F(1, 31) = 0.14$, $p = .714$, and no Session $\times$ Discrimination interaction, $F(5.08, 157.37) = 1.33$, $p = .247$. Post hoc comparisons confirmed no significant differences between the two discriminations at any session (all $p > .10$). These results indicate that both discriminations were acquired at similar rates and final performance levels.

Latency. Group mean auROC scores for latency are shown in Figure 2C. Both discriminations improved gradually from near-chance levels in Session 1 to approximately 0.60–0.65 by Session 7. Differences between the novel and preexposed discriminations were small and inconsistent, suggesting little evidence of a latent inhibition effect for this measure.

A repeated-measures ANOVA revealed a significant main effect of Session, $F(3.94, 122.59) = 6.12$, $p = .012$, Greenhouse–Geisser corrected, indicating that response latencies differentiated rewarded from non-rewarded cues increasingly over time. There was no significant main effect of Discrimination, $F(1, 31) = 0.08$, $p = .777$, and no Session $\times$ Discrimination interaction, $F(5.28, 163.75) = 1.51$, $p = .177$. Post hoc comparisons confirmed no significant differences in any session (all $p > .17$). Thus, latency-based performance improved similarly for both discriminations.

PCI. Group mean auROC scores for PC1, combining Number, Percent, and Latency, are shown in Figure 2D. Performance improved steadily from near-chance in early sessions to approximately 0.65–0.70 by Session 7, mirroring the Percent measure. The novel discrimination tended to score slightly higher in some sessions, but differences were inconsistent and nonsignificant.

A repeated-measures ANOVA revealed a significant main effect of Session, $F(4.26, 132.15) = 17.54, p < .001$, Greenhouse–Geisser corrected, indicating reliable improvement over time. There was no significant main effect of Discrimination, $F(1, 31) = 2.28, p = .141$, and no Session $\times$ Discrimination interaction, $F(5.12, 158.72) = 0.61, p = .719$. Post hoc tests confirmed no significant differences at any session (all $p > .08$). Thus, PC1 performance improved comparably for novel and preexposed discriminations. A repeated-measures ANOVA revealed a significant main effect of Session, $F(4.26, 132.15) = 17.54, p < .001$, Greenhouse–Geisser corrected, indicating reliable improvement in the combined measure of discrimination performance over time. There was no significant main effect of Discrimination, $F(1, 31) = 2.28, p = .141$, and no Session $\times$ Discrimination interaction, $F(5.12, 158.72) = 0.61, p = .719$, suggesting that the novel and preexposed discriminations reached similar performance levels and improved at comparable rates. Post hoc comparisons confirmed that AX+ vs AY- and AX+ vs BX- did not differ significantly at any session (all $p > .08$). Overall, the PC1 results indicate that both discriminations were learned to a similar degree, with no clear statistical evidence of a latent inhibition effect under within-modality conditions.

To sum up, across all four measures of magazine activity (Number of Entries, Percent Time in Magazine, Latency to First Magazine Entry, and PC1), rats successfully learned both within-subject discriminations (AX+ vs AY- and AX+ vs BX-) over the course of the seven LI test sessions. Performance improved steadily for both novel and preexposed cue conditions, with no statistically reliable differences in acquisition rate or final accuracy between them. Although the novel discrimination tended to score slightly higher in some measures, these trends were not supported by significant effects of Discrimination or by significant interactions with Session. Overall, the results indicate that prior cue exposure did not produce a robust latent inhibition effect under within-modality conditions in Stage 3.

Examining Individual Differences

Although group means showed no LI or facilitation effect, individual animals displayed considerable variability. Scatterplots of mean auROC scores for the two discriminations (AX+ vs AY- and AX+ vs BX-) revealed that some subjects performed better on the novel discrimination (consistent with LI), others on the preexposed discrimination (consistent with facilitation), and many near the identity diagonal (balanced performance). Histograms of signed distances from the diagonal were centered near zero but showed a widespread in both positive and negative directions, indicating heterogeneity in individual performance patterns (Figure 3A–H).











Figure 3. Individual-difference analyses of auROC scores: (A) scatterplot and (B) signed-distance histogram for Number of Entries; (C–D) the same for Percent Time; (E–F) Latency; (G–H) PC1. Distances are computed relative to the identity diagonal (AX+ vs AY– = AX+ vs BX–). Error bars represent ± 1 SEM.

Number of Entries. Inspection of the scatterplot of individual average auROC scores for AX+ vs AY– and AX+ vs BX– (Figure 3A) revealed notable variability across animals. Several subjects clustered near the identity diagonal, indicating comparable discrimination performance for novel and preexposed cues, while others fell further away, reflecting a stronger bias toward one discrimination. To quantify these differences, we calculated each animal’s perpendicular signed distance from the diagonal and plotted the distribution (Figure 3B). The mean signed distance was -0.05 ($SD = 0.11$), with values ranging from -0.35 to 0.27 and a median of -0.06 . The distribution was slightly negatively skewed (skewness = -0.25), suggesting that more animals scored below zero, consistent with a modest tendency toward facilitation rather than latent inhibition. The Shapiro–Wilk test indicated no significant deviation from normality ($p = .21$). Overall, these results indicate that, on the Number of Entries measure, most animals showed little or no latent inhibition, and in some cases, performance was modestly biased toward facilitation.

Percent. The scatterplot of individual average auROC scores for AX+ vs. AY– and AX+ vs. BX– (Figure 3C) showed substantial variability in performance across animals. Some clustered close to the identity diagonal, indicating similar discrimination for novel and preexposed cues, while others deviated more strongly, suggesting a bias toward one discrimination. To quantify these patterns, each animal’s perpendicular distance from the diagonal was calculated, and the distribution is shown in Figure 3D. The mean signed distance was 0.04 ($SD = 0.14$), with scores ranging from -0.37 to 0.39 and a median of 0.05 . The distribution was slightly positively skewed (skewness = 0.16), suggesting a mild tendency toward latent inhibition. A Shapiro–Wilk test found no significant deviation from normality ($p = 0.17$). These findings indicate that while the majority of animals showed relatively balanced discrimination, a subset exhibited notable individual differences in both the direction and magnitude of bias.

Latency. The scatterplot of individual average auROC scores for AX+ vs. AY– and AX+ vs. BX– (Figure 3E) revealed wide variability in performance across subjects. Several animals fell close to the identity diagonal, suggesting comparable discrimination for novel and preexposed cues, while others deviated substantially, indicating a bias toward one discrimination over the other. To quantify this variability, each animal’s perpendicular distance from the diagonal was calculated, and the resulting distribution is shown in Figure 3F. The mean signed distance was -0.01 ($SD = 0.12$), with scores ranging from -0.29 to 0.28 and a median of -0.02 . The distribution was slightly negatively skewed (skewness = -0.19), indicating a mild tendency toward stronger discrimination for the preexposed cue. A Shapiro–Wilk test found no significant deviation from

normality ($p = 0.29$). Overall, most animals demonstrated relatively balanced discrimination, though a subset showed marked individual differences in both magnitude and direction of bias.

PCI. The scatterplot of individual average auROC scores for AX+ vs. AY- and AX+ vs. BX- (Figure 3G) showed moderate clustering along the identity diagonal, indicating generally similar discrimination performance for novel and preexposed cues in most animals. A smaller subset deviated more markedly, suggesting individual biases toward one discrimination. Perpendicular distances from the diagonal were calculated for each subject, and their distribution is presented in Figure 3H. The mean signed distance was -0.03 ($SD = 0.11$), with values ranging from -0.29 to 0.19 and a median of -0.01 . The distribution was slightly negatively skewed (skewness = -0.42), indicating a mild tendency for stronger performance on the preexposed discrimination. A Shapiro–Wilk test indicated no significant deviation from normality ($p = 0.23$). These results suggest that while discrimination performance was balanced for most animals, a subset showed consistent biases in the opposite direction to that expected under latent inhibition.

Factors Potentially Contributing to Individual Differences

Can sex differences account for the observed individual differences?

Sex-specific scatterplots and histograms (Figure 4A–H) showed overlapping distributions for males and females in all four measures. Mixed-design ANOVAs revealed no significant main effects of Sex, no significant main effects of Session, and no Sex $\times$ Session interactions (all $ps > .09$), indicating that sex was not a reliable source of the observed variability in LI/facilitation magnitude.











Figure 4. Sex-specific individual-difference analyses for Number of Entries, Percent Time, Latency, and PC1: scatterplots (A, C, E, G) and signed-distance histograms (B, D, F, H). Error bars represent ± 1 SEM.

Number of Entries. Sex-specific analyses of the Number of Entries measure are shown in Figure 6A–B. In the sex-specific scatterplot (Figure 4A), males (blue) and females (red) displayed broadly overlapping distributions, with both sexes exhibiting a range of discrimination scores from near the identity diagonal to more pronounced deviations. Animals clustered along the diagonal showed similar discrimination performance for novel and preexposed cues, while those further from the diagonal demonstrated a stronger bias toward one discrimination. The histograms of signed distances (Figure 4B) further illustrate this pattern, revealing that both males and females exhibited variability in both positive and negative directions without a clear skew toward either discrimination. Together, these results suggest that discrimination performance patterns were not systematically driven by sex, and that individual differences were present within each group.

Percent. Sex-specific analyses of the Percent measure are shown in Figure 4C–D. In the sex-specific scatterplot (Figure 4C), males (blue) and females (red) displayed largely overlapping distributions, with both sexes showing discrimination scores ranging from near the identity diagonal to more pronounced deviations. Animals positioned along the diagonal demonstrated comparable discrimination performance for novel and preexposed cues, whereas those farther from the diagonal exhibited a stronger bias toward one discrimination. The histograms of signed distances (Figure 4D) support this pattern, indicating that both sexes displayed variability in both positive and negative directions without a consistent skew toward either discrimination. Overall, these findings suggest that sex did not systematically influence discrimination performance for the Percent measure, and that individual differences were evident within each sex.

Latency. Sex-specific analyses of the Latency measure are shown in Figure 6E–F. In the sex-specific scatterplot (Figure 4E), males (blue) and females (red) displayed overlapping distributions, with most animals clustering close to the identity diagonal. This suggests that, for many subjects, discrimination performance was similar for novel and preexposed cues. A subset of animals showed greater deviation from the diagonal, indicating stronger bias toward one discrimination, but this pattern was present in both sexes. The histograms of signed distances (Figure 4F) reinforce this observation, revealing that males and females exhibited variability in both positive and negative directions without a consistent skew toward either discrimination. These findings suggest that sex did not exert a systematic influence on discrimination performance in the Latency measure, and that individual differences were present within each sex.

PC1. Sex-specific analyses of the PC1 measure are shown in Figure 4G–H. In the sex-specific scatterplot (Figure 4G), males (blue) and females (red) exhibited partially overlapping but

distinct distributions, with some separation between the sexes in their discrimination scores. Most animals scored near or above the identity diagonal, indicating similar or slightly stronger discrimination for novel cues, while a smaller subset fell below the diagonal, suggesting a bias toward preexposed cue discrimination. The histograms of signed distances (Figure 4H) show that both sexes displayed variability, though males appeared slightly more concentrated around positive distances and females showed a wider spread toward negative values. These patterns suggest that while both sexes engaged in similar discrimination strategies overall, there may be subtle sex-linked differences in the distribution of bias magnitude.

Statistical Analysis of Sex Differences

Number of Entries. A mixed-design ANOVA was conducted on the auROC distance scores with Sex (male, female) as the between-subjects factor and Session (1–7) as the within-subjects factor. The main effect of Sex was not significant, $F(1, 30) = 2.99$, $p = .094$, indicating no reliable overall performance difference between males and females. The main effect of Session was non-significant, $F(6, 180) = 0.09$, $p = .997$, and the Sex $\times$ Session interaction was also non-significant, $F(6, 180) = 1.72$, $p = .119$. Mauchly's test of sphericity was non-significant for both main effects and the interaction.

Percent. A mixed-design ANOVA revealed no significant main effect of Sex, $F(1, 30) = 1.06$, $p = .313$, $ges = .012$, no significant main effect of Session, $F(6, 180) = 1.34$, $p = .243$, $ges = .028$, and no Sex $\times$ Session interaction, $F(6, 180) = 1.23$, $p = .294$, $ges = .026$. Mauchly's test of sphericity was non-significant for both Session and the interaction.

Latency. A mixed-design ANOVA revealed no significant main effect of Sex, $F(1, 30) = 0.32$, $p = .578$, $ges = .0026$, no significant main effect of Session, $F(6, 180) = 1.54$, $p = .168$, $ges = .037$, and no Sex $\times$ Session interaction, $F(6, 180) = 1.57$, $p = .160$, $ges = .038$. Mauchly's test was non-significant for both Session and the interaction.

PCI. A mixed-design ANOVA revealed no significant main effect of Sex, $F(1, 30) = 0.32$, $p = .578$, $ges = .0026$, no significant main effect of Session, $F(6, 180) = 1.54$, $p = .168$, $ges = .037$, and no Sex $\times$ Session interaction, $F(6, 180) = 1.57$, $p = .160$, $ges = .038$. Mauchly's test was non-significant for both Session and the interaction.

Can Counterbalancing Conditions Account for Individual Differences in LI




Figure 5. Distribution of auROC distance scores across the four counterbalancing groups ($n = 5$ per group) for number of entries (A), percent time (B), latency (C), and PC1 (D).

Number of Entries. Group means and scatterplots (Figure 5A) indicated variability across the four counterbalancing conditions, but there was no consistent pattern suggesting a systematic advantage for any one group. Although some groups contained individual high values, overall means were similar, and there was substantial overlap in distributions. A one-way ANOVA revealed no significant effect of counterbalancing group, $F(3, 28) = 2.05, p = .130$. Post hoc Tukey tests confirmed that no pairwise comparisons were significant (all $ps > .22$). These results suggest that cue assignment did not meaningfully influence LI/facilitation magnitude for the Number of Entries measure.

Percent. For Percent time in magazine (Figure 5B), Groups 2 and 3 had higher average scores compared to Groups 1 and 4. A one-way ANOVA revealed a significant effect of counterbalancing, $F(3, 16) = 6.57, p = .004$. Post hoc Tukey tests indicated that both Group 2 and Group 3 scored significantly higher than Group 1 ($ps < .05$) and Group 4 ($ps < .05$), while Groups 2 and 3 did not differ significantly from each other ($p > .90$). These results suggest that cue identity influenced Percent scores, with the assignments in Groups 2 and 3 producing generally stronger discrimination.

Latency. For Latency (Figure 5C), Group 4 showed the highest scores compared to the other groups. A one-way ANOVA revealed a significant effect of counterbalancing, $F(3, 16) = 3.79, p = .031$. Post hoc Tukey tests indicated that Group 4 was significantly higher than Group 2 ($p < .05$) and marginally higher than Group 3 ($p \approx .067$), while differences with Group 1 were not significant. These results suggest that cue identity influenced Latency-based performance, with the Group 4 assignment producing the strongest discrimination.

PC1. For the composite PC1 measure (Figure 5D), Group 3 had the highest mean scores. A one-way ANOVA revealed a significant effect of counterbalancing ($p < .05$). Post hoc Tukey tests showed that Group 3 scored significantly higher than Group 1 ($p < .05$) and Group 4 ($p < .05$), while other comparisons were not significant. This indicates that cue assignment influenced the combined discrimination measure, with Group 3 producing the strongest overall performance when all behavioral indices were integrated.

Relationship Between Conditioned Inhibition Strength and LI

Finally, we tested whether the strength of conditioned inhibition to B in Stage 2 predicted Stage 3 LI/facilitation scores (Figure 6A–D). Correlations were generally weak and non-significant across measures (Number: $r = -.324, p = .070$; Percent: $r = -.093, p = .614$; Latency: $r = -.287, p = .111$; PC1: $r = .014, p = .938$). Although Number of Entries showed a trend toward a negative correlation, suggesting that stronger inhibition might be associated with weaker LI, this did not reach significance. Overall, inhibitory strength in Stage 2 was not a consistent predictor of LI or facilitation in Stage 3.




Figure 6. Correlations between Stage 2 conditioned-inhibition distance scores and Stage 3 latent-inhibition distance scores for (A) Number of Entries, (B) Percent Time, (C) Latency, and (D) PC1.

Number. The correlation between Stage 2 (conditioned inhibition) and Stage 3 (latent inhibition) distance scores for the number of magazine entries was moderate and approached statistical significance. As shown in Figure 6A, the Pearson correlation coefficient was $r = -0.324$ with a p-value of .070 and a 95% confidence interval ranging from -0.604 to 0.028 . Although the scatterplot indicates a negative relationship; suggesting that animals showing stronger inhibitory learning during Stage 2 tended to display weaker latent inhibition during Stage 3. The association did not reach significance. This pattern suggests that individual differences in conditioned inhibition may contribute to variability in latent inhibition for this measure, but the evidence remains inconclusive.

Percent. The correlation between Stage 2 (conditioned inhibition) and Stage 3 (latent inhibition) difference scores for the percent of time spent in the magazine was very weak and not statistically significant. As shown in Figure 6B, the Pearson correlation coefficient was $r = -0.093$ with a p-value of .614 and a 95% confidence interval ranging from -0.428 to 0.264 . The scatterplot shows no discernible relationship between the two measures, with data points distributed broadly around the zero line. This result suggests that performance in conditioned inhibition during Stage 2 was unrelated to the magnitude of latent inhibition expressed in Stage 3 for this measure, indicating that percent time and number of entries may capture different aspects of behavioral variation.

Latency. The correlation between Stage 2 (conditioned inhibition) and Stage 3 (latent inhibition) difference scores for magazine entry latency was weak and not statistically significant. As shown in Figure 6C, the Pearson correlation coefficient was $r = -0.287$ with a p-value of .111 and a 95% confidence interval ranging from -0.578 to 0.069 . The scatterplot indicates a slight negative trend, with higher conditioned inhibition scores tending to be associated with lower latent inhibition scores; however, this relationship did not reach statistical significance. These results suggest that variation in inhibitory performance during Stage 2 was not strongly predictive of latency-based measures of latent inhibition in Stage 3.

PC1. The correlation between Stage 2 (conditioned inhibition) and Stage 3 (latent inhibition) difference scores for the first principal component (PC1) was essentially zero and not statistically significant. As shown in Figure 6D, the Pearson correlation coefficient was $r = 0.014$ with a p-value of .938 and a 95% confidence interval ranging from -0.336 to 0.361 . The scatterplot indicates no discernible relationship between the two variables, with points distributed evenly around the regression line. These results indicate that PC1-based performance in Stage 2 was not predictive of PC1-based measures of latent inhibition in Stage 3.

Discussion

The most salient finding of this study was that, despite extensive preexposure to cues A and B, no reliable latent inhibition (LI) effect emerged at the group level. This null result is informative given that the design was specifically intended to isolate attentional decrements from interference processes by establishing B as a conditioned inhibitor prior to test. One possibility is that the facilitation produced by making B an explicit inhibitor counteracted any losses in associability typically associated with LI (Lubow & Moore, 1959; Hall & Rodríguez, 2011). Another possibility is that the mode of preexposure used in Stage 2, embedded within a complex discrimination task rather than through simple nonreinforced presentations, was less effective in producing the kind of attentional decrement emphasized by associative models of LI (Pearce & Hall, 1980; Mackintosh, 1975).

If no associability losses had occurred, the expected outcome would have been the opposite of LI: net facilitation of the preexposed-based subdiscrimination. This expectation rests on two considerations. First, A and B were explicitly discriminated during Stage 2, which should have promoted perceptual learning (Gibson, 1969). Second, B was trained as a conditioned inhibitor, which should have transferred to facilitate the acquisition of AX+ vs. BX– relative to AX+ vs. AY– at test (Rescorla, 1969). The absence of either an LI or facilitation effect suggests that these competing processes may have cancelled each other out, leaving no clear group-level bias in discrimination learning.

One limitation of this study is that cue B was not independently validated as a conditioned inhibitor. Although its training history strongly implied inhibitory properties, the absence of a direct summation or retardation test means that this interpretation rests on assumption rather than direct confirmation.

At the individual level, performance patterns varied considerably. Some animals acquired the novel-cue subdiscrimination more rapidly than the preexposed-cue subdiscrimination, consistent with LI, while others showed little difference or the reverse pattern. This heterogeneity could reflect genuine individual differences in susceptibility to LI, but it may also have arisen from cue-identity effects, counterbalancing variability, or stochastic noise (Le Pelley, 2004). Given that counterbalancing group influenced performance in other chapters, it is plausible that cue assignment contributed to these discrepancies.

Although the omnibus ANOVAs did not yield significant discrimination effects, I conducted planned post hoc contrasts because the hypotheses specified a directional comparison between preexposed- and novel-based discriminations. From a theoretical perspective, testing these contrasts was important for evaluating whether attentional decrement or interference best accounted for the observed performance. This ensured that even in the absence of significant omnibus results, the analyses remained aligned with the scientific questions driving the design.

In sum, Chapter 3 provided no evidence of a robust group-level LI effect under conditions where both perceptual learning and inhibitory transfer might also have been expected to enhance discrimination. While some individuals displayed patterns consistent with LI, these were neither consistent across measures nor strong enough to establish attentional decrement as the sole mechanism.

General Discussion

The primary purpose of this thesis was to develop and validate a novel within-subject latent inhibition (LI) paradigm that establishes a behavioral foundation for future neuroscience and translational research. The experimental design tested here can be applied to address fundamental questions about the neural circuitry underlying both LI and perceptual learning (PL), as well as to examine how these mechanisms may be altered in clinical populations. Given that attentional filtering and associative learning deficits are core features of disorders such as schizophrenia (Lubow & Gewirtz, 1995; Weiner, 2003), and that these deficits may manifest differently depending on preexposure conditions, this paradigm provides a promising link between basic learning theory, neuroscience, and translational research.

In Chapter 1, we manipulated the amount of preexposure to a conditioned stimulus and found that extensive preexposure (12 days) produced a robust latent inhibition (LI) effect, whereas limited preexposure (4 days) under intermixed conditions facilitated discrimination, consistent with perceptual learning (PL). This finding is significant because it demonstrates that LI and PL can emerge under identical testing conditions, allowing direct comparisons of neural and behavioral outcomes. The primary limitations of Chapter 1 were methodological. The relatively small sample size likely reduced statistical power (Lakens, 2022), and the absence of a no-exposure control group prevents us from definitively concluding that, without preexposure, the two test discriminations would have been learned at the same rate. This limits the interpretation of the PL effect. Future behavioral work should replicate these findings with a larger sample and include an explicit no-exposure control group. Additionally, parametrically varying preexposure duration would help clarify whether the transition between PL and LI occurs gradually or categorically. Clinically, this procedure could reveal whether individuals with schizophrenia differ in their sensitivity to preexposure amount. For instance, they may show stronger PL but little or no LI, fail to show either effect, or display a shifted balance between the two. Such patterns have been observed in both behavioral and neural studies of LI in schizophrenia (Lubow & Gewirtz, 1995; Weiner, 2003; Gal, Schiller, & Weiner, 2009).

The experimental design employed in Chapter 1 allows testing of theoretical accounts of latent inhibition (LI) and perceptual learning (PL) that have previously been inaccessible due to the absence of within-subject assays. One such account is McLaren and Mackintosh's (2000, 2002) model. According to this account, when two stimuli, A and B, are preexposed, both their unique elements ("a" and "b") and their common element ("c") are preexposed. Thus, an A trial consists of "ac" elements, while a B trial consists of "bc" elements. Because "c" is preexposed on both A and B trials, it is presented twice as often as "a" or "b." As a result, the salience of "c" diminishes more rapidly, making "a" and "b" more conspicuous after repeated exposure. Since "a" and "b" are unique, this enhanced distinctiveness facilitates the solution of the subsequent discrimination relative to novel stimuli, consistent with PL. With continued preexposure, however, the salience of all elements ("a," "b," and "c") is expected to decline to such low levels

that the discrimination becomes more difficult to solve than one involving novel stimuli, consistent with LI. In addition, the model proposes that PL is further supported by mutual inhibition between the unique elements “a” and “b.” This inhibition develops because the common element “c” activates a memory of “b” on “ac” trials (stimulus A) and of “a” on “bc” trials (stimulus B), making “a” a conditioned inhibitor for “b” and vice versa.

While the results of this thesis do not directly adjudicate between theoretical mechanisms, they provide a foundation for future investigations at the neural level. For example, in vivo electrophysiological recording or calcium imaging could be used to track changes in the unique and common patterns of activity elicited by A and B presentations during the preexposure phase, and to examine how those changes relate to the degree of LI observed in the test phase. During the test stage itself, it would also be possible to compare the neural correlates of preexposed- versus novel-based subdiscriminations. Previous research implicates a distributed neural circuit in LI, including the hippocampus, amygdala, nucleus accumbens, striatum, thalamus, and prefrontal cortex, with dopaminergic modulation from midbrain structures such as the ventral tegmental area and substantia nigra (Weiner, 2003). These regions represent particularly promising candidates for future studies using the present paradigm, offering the potential to clarify how distinct neural systems contribute to LI and PL.

In Chapter 2, we examined whether LI could be obtained within the same sensory modality by using both preexposed and novel cues from that modality. The LI effect observed here was relatively weak, reaching significance on only one behavioral measure (percent time), and showed substantial variability across individuals. This finding suggests that perceptual similarity between cues may attenuate LI by reducing the distinctiveness of the preexposed cue relative to the novel cue (Hall, 2001). Increased generalization between the cues would have blurred any differences between the test discriminations. Future behavioral work could systematically manipulate cue similarity to clarify how perceptual similarity moderates LI. Neural investigations could further test whether reduced LI in high-similarity conditions is associated with weaker neural differentiation in sensory cortical areas or with reduced hippocampal pattern separation (Yassa & Stark, 2011). Clinically, this line of work may help determine whether patients with schizophrenia, or other groups with perceptual discrimination deficits (Silverstein et al., 2013), are particularly prone to showing little or no LI when cues are drawn from the same modality. Alternatively, they may display an exaggerated PL effect under these conditions, since generalization across cues from the same modality may facilitate PL rather than impair it. More broadly, future studies should compare the effects of same- versus cross-modality training on LI and PL. It is interesting to speculate that the size of the within-subjects LI effect might benefit from using cues from different modalities (by reducing generalization), whereas PL might be enhanced when cues belong to the same modality (by providing a greater generalization backdrop against which PL can increase stimulus discriminability).

Chapter 3 addressed a key theoretical question: whether LI is primarily driven by attentional decrement or by interference from previously learned associations (Bouton, 1993; Pearce & Mackintosh, 2010). To test this, the experimental design encouraged the transfer of congruent associative properties from training to testing, effectively transforming the typical potential for interference in LI designs into an opportunity for facilitation. If LI had still been observed under these conditions, it would have provided strong evidence for associability losses in the preexposed cues. Although this approach was theoretically compelling, the results showed no net LI or facilitation effect, suggesting that LI mechanisms and associative transfer may have cancelled one another out. This outcome is consistent with the view that attentional and associative retrieval processes are intertwined, making them difficult to experimentally separate without additional attentional probes (Bouton, 1993).

Two further limitations of Chapter 3 are worth noting. First, cue B was not independently validated as an inhibitor. Its inhibitory status was assumed from training history but was not confirmed through summation or retardation tests, leaving uncertainty about the strength of its inhibitory properties. Second, although none of the omnibus ANOVAs yielded significant effects, post hoc comparisons were still conducted. These were retained because the hypotheses specified a directional contrast between novel- and preexposed-based discriminations, and testing them was scientifically justified despite the null omnibus outcomes. Including these contrasts ensured that the analyses directly addressed the theoretical questions motivating the experiment.

Chapter 3 also revealed pronounced individual differences, which appeared to be influenced by counterbalancing condition but not by sex. Notably, responding during the conditioned inhibition training stage did not predict the magnitude of LI in the subsequent test stage, indicating that individual variability in LI expression was not due to differences in acquiring inhibitory learning. Future studies will need to clarify whether these individual differences reflect random variability around the absence of an average effect between preexposed- and novel-based subdiscriminations, or whether they represent meaningful positions along an LI–facilitation continuum. One approach would be to incorporate attentional measures, such as orienting responses during preexposure, to directly link observed LI effects to changes in attentional allocation. Distinguishing attentional from interference processes in this way would provide a stronger foundation for mechanistic neuroscience and translational research into attentional decrements and their role in individual differences in LI.

Taken together, the three chapters of this thesis establish and refine a novel within-subject paradigm for measuring latent inhibition and perceptual learning. Chapter 1 demonstrated that both LI and PL can emerge under identical testing conditions, with preexposure amount determining the balance between them. Chapter 2 showed that LI is weaker and more variable when cues are drawn from the same sensory modality, underscoring the importance of cue similarity and generalization. Chapter 3 attempted to isolate attentional from interference

processes but revealed the difficulty of disentangling these mechanisms without additional probes, as LI and facilitation appeared to counterbalance one another. Across studies, individual variability emerged as a consistent theme, highlighting the need to integrate attentional, associative, and perceptual processes to explain differences in LI expression. By uniting LI and PL within a common framework, this paradigm provides a foundation for future neural and translational work aimed at understanding how basic learning mechanisms map onto brain function and dysfunction in clinical populations.

References

- Aranzubia-Olasolo, M., Nelson, J. B., & Sanjuán Artegain, M. D. C. (2024). Latent inhibition in humans from simple stimulus exposure. *Journal of Experimental Psychology: Animal Learning and Cognition*, 50(2), 118–130. <https://doi.org/10.1037/xan0000374>
- Badiola-Lekue, I., Arriola, N., & Rodríguez, G. (2024). The relationship between latent inhibition, divergent thinking, and eyewitness memory: A study on attention to irrelevant stimuli. *bioRxiv*. <https://doi.org/10.1101/2024.11.23.625026>
- Baker, A. G., & Mackintosh, N. J. (1977). Excitatory and inhibitory conditioning following uncorrelated presentations of CS and UCS. *Animal Learning & Behavior*, 5(3), 315–319. <https://doi.org/10.3758/BF03209246>
- Barad, M., Blouin, A. M., & Cain, C. K. (2004). Like extinction, latent inhibition of conditioned fear in mice is blocked by systemic inhibition of L-type voltage-gated calcium channels. *Learning & Memory*, 11(5), 536–539. <https://doi.org/10.1101/lm.78304>
- Barak, S., & Weiner, I. (2007). Scopolamine induces disruptions in latent inhibition and extinction: Differential reversal by antipsychotic drugs. *Neuropsychopharmacology*, 32(9), 2092–2102. <https://doi.org/10.1038/sj.npp.1301323>
- Baruch, I., Hemsley, D. R., & Gray, J. A. (1988). Latent inhibition and “psychotic proneness” in normal subjects. *Personality and Individual Differences*, 9(4), 777–783. [https://doi.org/10.1016/0191-8869\(88\)90067-0](https://doi.org/10.1016/0191-8869(88)90067-0)
- Bonardi, C., Brilot, B., & Jennings, D. J. (2016). Learning about the CS during latent inhibition: Preexposure enhances temporal control. *Journal of Experimental Psychology: Animal Learning and Cognition*, 42(2), 187–199. <https://doi.org/10.1037/xan0000096>
- Boughner, R. L., & Papini, M. R. (2003). Appetitive latent inhibition in rats: Now you see it (sign tracking), now you don’t (goal tracking). *Learning & Behavior*, 31(4), 387–392. <https://doi.org/10.3758/BF03195999>
- Bouton, M. E. (1993). Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychological Bulletin*, 114(1), 80–99. <https://doi.org/10.1037/0033-2909.114.1.80>
- Braunstein-Bercovitz, H. (2010). Latent inhibition as a function of anxiety and stress: Implications for schizophrenia. In R. Lubow & I. Weiner (Eds.), *Latent inhibition* (pp. 457–476). Cambridge University Press. <https://doi.org/10.1017/CBO9780511730184.020>
- Buhusi, C. V., Gray, J. A., & Schmajuk, N. A. (1998). Perplexing effects of hippocampal lesions on latent inhibition: A neural network solution. *Behavioral Neuroscience*, 112(2), 316–351. <https://doi.org/10.1037/0735-7044.112.2.316>
- Byrom, N. C., Msetfi, R. M., & Murphy, R. A. (2018). Human latent inhibition: Problems with the stimulus exposure effect. *Psychonomic Bulletin & Review*, 25(6), 2102–2118. <https://doi.org/10.3758/s13423-018-1455-4>

- Caldarone, B. J., Duman, C. H., & Picciotto, M. R. (2000). Fear conditioning and latent inhibition in mice lacking the high affinity subclass of nicotinic acetylcholine receptors in the brain. *Neuropharmacology*, 39(13), 2779–2784. [https://doi.org/10.1016/S0028-3908\(00\)00137-4](https://doi.org/10.1016/S0028-3908(00)00137-4)
- Carson, S. (2010). Latent inhibition and creativity. In R. Lubow & I. Weiner (Eds.), *Latent inhibition* (pp. 183–198). Cambridge University Press. <https://doi.org/10.1017/CBO9780511730184.010>
- Carson, S. H., Peterson, J. B., & Higgins, D. M. (2003). Decreased latent inhibition is associated with increased creative achievement in high-functioning individuals. *Journal of Personality and Social Psychology*, 85(3), 499–506. <https://doi.org/10.1037/0022-3514.85.3.499>
- Cassaday, H. J., & Moran, P. M. (2010). Latent inhibition and other salience modulation effects: Same neural substrates? In R. Lubow & I. Weiner (Eds.), *Latent inhibition* (pp. 342–371). Cambridge University Press. <https://doi.org/10.1017/CBO9780511730184.016>
- Clark, A. J. M., Feldon, J., & Rawlins, J. N. P. (1992). Aspiration lesions of rat ventral hippocampus disinhibit responding in conditioned suppression or extinction, but spare latent inhibition and the partial reinforcement extinction effect. *Neuroscience*, 48(4), 821–829. [https://doi.org/10.1016/0306-4522\(92\)90270-C](https://doi.org/10.1016/0306-4522(92)90270-C)
- Collins, A. C., Wehner, J. M., & Gould, T. J. (2001). Nicotine enhances latent inhibition and ameliorates ethanol-induced deficits in latent inhibition. *Nicotine & Tobacco Research*, 3(1), 17–24. <https://doi.org/10.1080/14622200125450>
- Coutureau, E., Galani, R., Gosselin, O., Majchrzak, M., & Di Scala, G. (1999). Entorhinal but not hippocampal or subicular lesions disrupt latent inhibition in rats. *Neurobiology of Learning and Memory*, 72(3), 143–157. <https://doi.org/10.1006/nlme.1998.3895>
- Coutureau, E., Léna, I., Dugé, V., & Di Scala, G. (2002). The entorhinal cortex-nucleus accumbens pathway and latent inhibition: A behavioral and neurochemical study in rats. *Behavioral Neuroscience*, 116(1), 95–104. <https://doi.org/10.1037/0735-7044.116.1.95>
- Crowell, C. R., & Anderson, D. C. (1972). Variations in intensity, interstimulus interval, and interval between preconditioning CS exposures and conditioning with rats. *Journal of Comparative and Physiological Psychology*, 79(2), 291–298. <https://doi.org/10.1037/h0032550>
- Davis, J. A., & Gould, T. J. (2005). Rolipram attenuates MK-801-induced deficits in latent inhibition. *Behavioral Neuroscience*, 119(2), 595–602. <https://doi.org/10.1037/0735-7044.119.2.595>
- De la Casa Rivas, L. G., Traverso Arcos, L. M., & Márquez Zamora, R. (2010). Effect of a retention interval between pre-exposure and conditioning on latent inhibition in humans using a blink conditioning procedure. *Psicothema*, 22(4), 708–714.
- De la Casa, G., & Lubow, R. E. (1995). Latent inhibition in conditioned taste aversion: The roles of stimulus frequency and duration and the amount of fluid ingested during preexposure. *Neurobiology of Learning and Memory*, 64(2), 125–132. <https://doi.org/10.1006/nlme.1995.1051>

- De la Casa, L. G., & Lubow, R. E. (2001). Latent inhibition with a response time measure from a within-subject design: Effects of number of preexposures, masking task, context change, and delay. *Neuropsychology*, 15(2), 244–253. <https://doi.org/10.1037/0894-4105.15.2.244>
- De la Casa, L. G., & Pineño, O. (2010). Inter-stage context and time as determinants of latent inhibition. In R. Lubow & I. Weiner (Eds.), *Latent inhibition* (pp. 40–61). Cambridge University Press. <https://doi.org/10.1017/CBO9780511730184.004>
- Escobar, M., & Miller, R. R. (2010). Latent inhibition: Acquisition or performance deficit? In R. Lubow & I. Weiner (Eds.), *Latent inhibition* (pp. 62–93). Cambridge University Press. <https://doi.org/10.1017/CBO9780511730184.005>
- Escobar, M., Arcediano, F., & Miller, R. R. (2002). Latent inhibition and contextual associations. *Journal of Experimental Psychology: Animal Behavior Processes*, 28(2), 123–136. <https://doi.org/10.1037/0097-7403.28.2.123>
- Escobar, M., Arcediano, F., & Miller, R. R. (2005). Disruption of latent inhibition by interpolation of task-irrelevant stimulation between preexposure and conditioning. *Learning & Behavior*, 33(3), 371–385. <https://doi.org/10.3758/BF03192865>
- Esnal, A., Sánchez-González, A., Río-Álamos, C., Oliveras, I., Cañete, T., Blázquez, G., Tobeña, A., & Fernández-Teruel, A. (2016). Prepulse inhibition and latent inhibition deficits in Roman high-avoidance vs. Roman low-avoidance rats: Modeling schizophrenia-related features. *Physiology & Behavior*, 163, 267–273. <https://doi.org/10.1016/j.physbeh.2016.05.020>
- Evans, L. H., Gray, N. S., & Snowden, R. J. (2007). A new continuous within-participants latent inhibition task: Examining associations with schizotypy dimensions, smoking status, and gender. *Biological Psychology*, 74(3), 365–373. <https://doi.org/10.1016/j.biopsycho.2006.09.007>
- Forrest, D. R. L., Mather, M., & Harris, J. A. (2018). Unmasking latent inhibition in humans. *Quarterly Journal of Experimental Psychology*, 71(2), 380–395. <https://doi.org/10.1080/17470218.2016.1249894>
- Friedman, B. X., Blaisdell, A. P., Escobar, M., & Miller, R. R. (1998). Comparator mechanisms and conditioned inhibition: Conditioned stimulus preexposure disrupts Pavlovian conditioned inhibition but not explicitly unpaired inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*, 24(4), 453–466. <https://doi.org/10.1037/0097-7403.24.4.453>
- Gaisler-Salomon, I., & Weiner, I. (2003). Systemic administration of MK-801 produces an abnormally persistent latent inhibition which is reversed by clozapine but not haloperidol. *Psychopharmacology*, 166(4), 333–342. <https://doi.org/10.1007/s00213-002-1311-z>
- Gal, G., Schiller, D., & Weiner, I. (2005). Latent inhibition is disrupted by nucleus accumbens shell lesion but is abnormally persistent following entire nucleus accumbens lesion: The neural site controlling the expression and disruption of the stimulus preexposure effect. *Behavioural Brain Research*, 162(2), 246–255. <https://doi.org/10.1016/j.bbr.2005.03.019>
- Gershman, S. J., & Niv, Y. (2012). Exploring a latent cause theory of classical conditioning. *Learning & Behavior*, 40(3), 255–268. <https://doi.org/10.3758/s13420-012-0080-8>

- Gibson, E. J. (1969). *Principles of perceptual learning and development*. Appleton-Century-Crofts.
- Ginton, A., Urca, G., & Lubow, R. E. (1975). Effects of a competing task on stimulus preexposure phenomena in selective attention. *Journal of Experimental Psychology: Human Perception and Performance*, 1(2), 161–171. <https://doi.org/10.1037/0096-1523.1.2.161>
- Gould, T. J. (2010). A comparison of mechanisms underlying the CS–US association and the CS–nothing association. In R. Lubow & I. Weiner (Eds.), *Latent inhibition* (pp. 252–275). Cambridge University Press. <https://doi.org/10.1017/CBO9780511730184.013>
- Gould, T. J., & Lewis, M. C. (2005). Coantagonism of glutamate receptors and nicotinic acetylcholinergic receptors disrupts fear conditioning and latent inhibition of fear conditioning. *Learning & Memory*, 12(4), 389–398. <https://doi.org/10.1101/lm.89105>
- Gray, J. A., Joseph, M. H., Hemsley, D. R., Young, A. M. J., Warburton, E. C., Boulenguez, P., Grigoryan, G. A., Peters, S. L., Rawlins, J. N. P., Taib, C.-T., Yee, B. K., Cassaday, H., Weiner, I., Gal, G., Gusak, O., Joel, D., Shadach, E., Shalev, U., Tarrasch, R., & Feldon, J. (1995). The role of mesolimbic dopaminergic and retrohippocampal afferents to the nucleus accumbens in latent inhibition: Implications for schizophrenia. *Behavioural Brain Research*, 71(1–2), 19–31. [https://doi.org/10.1016/0166-4328\(95\)00154-9](https://doi.org/10.1016/0166-4328(95)00154-9)
- Gray, J. A., Moran, P. M., Grigoryan, G., Peters, S. L., Young, A. M. J., & Joseph, M. H. (1997). Latent inhibition: The nucleus accumbens connection revisited. *Behavioural Brain Research*, 88(1), 27–34. [https://doi.org/10.1016/S0166-4328\(97\)02313-9](https://doi.org/10.1016/S0166-4328(97)02313-9)
- Gray, N. S., Fernandez, M., Williams, J., Ruddle, R. A., & Snowden, R. J. (2002). Which schizotypal dimensions abolish latent inhibition? *British Journal of Clinical Psychology*, 41(3), 271–284. <https://doi.org/10.1348/014466502760379136>
- Gray, N. S., Williams, J., Fernandez, M., Ruddle, R. A., Good, M. A., & Snowden, R. J. (2001). Context dependent latent inhibition in adult humans. *Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*, 54(3), 233–245. <https://doi.org/10.1080/02724990143000027>
- Hall, G. (2001). *Perceptual and associative learning*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198524211.001.0001>
- Hall, G. (2003). Learned changes in the processing of stimuli. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 435–462). Erlbaum.
- Hall, G., & Channell, S. (1986). Context specificity of latent inhibition in taste aversion learning. *Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*, 38(2), 121–139.
- Hall, G., & Pearce, J. M. (1979). Latent inhibition of a CS during CS–US pairings. *Journal of Experimental Psychology: Animal Behavior Processes*, 5(1), 31–42. <https://doi.org/10.1037/0097-7403.5.1.31>

- Hall, G., & Rodriguez, G. (2010). Associative and nonassociative processes in latent inhibition: An elaboration of the Pearce-Hall model. In R. Lubow & I. Weiner (Eds.), *Latent inhibition: Cognition, neuroscience and applications to schizophrenia* (pp. 114–136). Cambridge University Press. <https://doi.org/10.1017/CBO9780511730184.007>
- Hall, G., & Rodriguez, G. (2011). Blocking of potentiation of latent inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*, 37(1), 127–131. <https://doi.org/10.1037/a0020716>
- Hall, G., & Rodríguez, G. (2011). Latent inhibition: A review and integration of findings. *Neuroscience & Biobehavioral Reviews*, 35(9), 2065–2078. <https://doi.org/10.1016/j.neubiorev.2011.04.004>
- Hall, G., & Rodríguez, G. (2019). Attention to perceive, to learn and to respond. *Quarterly Journal of Experimental Psychology*, 72(2), 335–345. <https://doi.org/10.1080/17470218.2017.1339719>
- Hall, G., & Rodríguez, G. (2020). When the stimulus is predicted and what the stimulus predicts: Alternative accounts of habituation. *Journal of Experimental Psychology: Animal Learning and Cognition*, 46(3), 327–340. <https://doi.org/10.1037/xan0000237>
- Hall, G., Stockhorst, U., Enck, P., & Klosterhalfen, S. (2016). Overshadowing and latent inhibition in nausea-based context conditioning in humans: Theoretical and practical implications. *Quarterly Journal of Experimental Psychology*, 69(6), 1227–1238. <https://doi.org/10.1080/17470218.2015.1080739>
- Hall, G., Symonds, M., & Rodriguez, M. (2009). Enhanced latent inhibition in context aversion conditioning. *Learning and Motivation*, 40(1), 62–73. <https://doi.org/10.1016/j.lmot.2008.05.001>
- Han, J. S., Gallagher, M., & Holland, P. (1995). Hippocampal lesions disrupt decrements but not increments in conditioned stimulus processing. *Journal of Neuroscience*, 15(11), 7323–7329. <https://doi.org/10.1523/JNEUROSCI.15-11-07323.1995>
- Harris, J. A., Jones, M. L., Bailey, G. K., & Westbrook, R. F. (2000). Contextual control over conditioned responding in an extinction paradigm. *Journal of Experimental Psychology: Animal Behavior Processes*, 26(2), 174–185. <https://doi.org/10.1037/0097-7403.26.2.174>
- Haselgrove, M., Lagator, S., Mah, S. L., & Gray, E. K. (2025). Novelty mismatch as a determinant of latent inhibition. *Journal of Experimental Psychology: Animal Learning and Cognition*, 51(1), 13–34. <https://doi.org/10.1037/xan0000388>
- Holt, W., & Maren, S. (1999). Muscimol inactivation of the dorsal hippocampus impairs contextual retrieval of fear memory. *Journal of Neuroscience*, 19(20), 9054–9062. <https://doi.org/10.1523/JNEUROSCI.19-20-09054.1999>
- Holtzman, O., Siette, J., Holmes, N. M., & Westbrook, R. F. (2010). Additional exposures reverse the latent inhibitory effects of recent and remote exposures. *Journal of Experimental Psychology: Animal Behavior Processes*, 36(3), 368–380. <https://doi.org/10.1037/a0017674>

- Honey, R. C., & Good, M. (1993). Selective hippocampal lesions abolish the contextual specificity of latent inhibition and conditioning. *Behavioral Neuroscience*, 107(1), 23–33. <https://doi.org/10.1037/0735-7044.107.1.23>
- Honey, R. C., & Hall, G. (1989). Acquired equivalence and distinctiveness of cues. *Journal of Experimental Psychology: Animal Behavior Processes*, 15(4), 338–346. <https://doi.org/10.1037/0097-7403.15.4.338>
- Honey, R. C., & Hall, G. (1989). Enhanced discrimination following prior exposure to the stimuli. *Quarterly Journal of Experimental Psychology*, 41B(1), 33–45. <https://doi.org/10.1080/14640748908401180>
- Honey, R. C., Iordanova, M. D., & Good, M. (2010). Latent inhibition and habituation: Evaluation of an associative analysis. In R. Lubow & I. Weiner (Eds.), *Latent inhibition* (pp. 163–182). Cambridge University Press. <https://doi.org/10.1017/CBO9780511730184.009>
- Ishii, K. (1999). Attenuation of latent inhibition after compound conditioning. *Japanese Psychological Research*, 41(2), 102–111. <https://doi.org/10.1111/1468-5884.00109>
- Kapur, S. (2003). Psychosis as a state of aberrant salience: A framework linking biology, phenomenology, and pharmacology in schizophrenia. *American Journal of Psychiatry*, 160(1), 13–23. <https://doi.org/10.1176/appi.ajp.160.1.13>
- Lakens, D. (2022). Sample size justification. *Collabra: Psychology*, 8(1), 33267. <https://doi.org/10.1525/collabra.33267>
- Le Pelley, M. E. (2004). The role of associative history in models of associative learning: A selective review and a hybrid model. *Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*, 57(3), 193–243. <https://doi.org/10.1080/02724990344000141>
- Le Pelley, M. E., & Schmidt-Hansen, M. (2010). Latent inhibition and learned irrelevance in human contingency learning. In R. Lubow & I. Weiner (Eds.), *Latent inhibition* (pp. 94–113). Cambridge University Press. <https://doi.org/10.1017/CBO9780511730184.006>
- Lubow, R. E. (1973). Latent inhibition. *Psychological Bulletin*, 79(6), 398–407. <https://doi.org/10.1037/h0034425>
- Lubow, R. E. (1989). *Latent inhibition and conditioned attention theory*. Cambridge University Press.
- Lubow, R. E., & Gewirtz, J. C. (1995). Latent inhibition in humans: Data, theory, and implications for schizophrenia. *Psychological Bulletin*, 117(1), 87–103. <https://doi.org/10.1037/0033-2909.117.1.87>
- Lubow, R. E., & Moore, A. U. (1959). Latent inhibition: The effect of nonreinforced pre-exposure to the conditional stimulus. *Journal of Comparative and Physiological Psychology*, 52(4), 415–419. <https://doi.org/10.1037/h0046700>

- Lubow, R. E., & Moore, A. U. (1959). Latent inhibition: The effect of nonreinforced pre-exposure to the conditional stimulus. *Journal of Comparative and Physiological Psychology*, 52(4), 415–419. <https://doi.org/10.1037/h0046700>
- Lubow, R. E., Schnur, P., & Rifkin, B. (1976). Latent inhibition and conditioned attention theory. *Journal of Experimental Psychology: Animal Behavior Processes*, 2(2), 163–174. <https://doi.org/10.1037/0097-7403.2.2.163>
- Mackintosh, N. J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, 82(4), 276–298. <https://doi.org/10.1037/h0076778>
- McLaren, I. P. L., & Mackintosh, N. J. (2000). An elemental model of associative learning: I. Latent inhibition and perceptual learning. *Animal Learning & Behavior*, 28(3), 211–246. <https://doi.org/10.3758/BF03200258>
- McLaren, I. P. L., & Mackintosh, N. J. (2002). Associative learning and elemental representations: II. Generalization and discrimination. *Animal Learning & Behavior*, 30(3), 177–200. <https://doi.org/10.3758/BF03192828>
- McLaren, I. P. L., Kaye, H., & Mackintosh, N. J. (1989). An associative theory of the representation of stimuli: Applications to perceptual learning and latent inhibition. In R. G. M. Morris (Ed.), *Parallel distributed processing: Implications for psychology and neurobiology* (pp. 102–130). Oxford University Press.
- Miller, R. R., & Escobar, M. (2001). Contrasting acquisition-focused and performance-focused models of acquired behavior. *Current Directions in Psychological Science*, 10(4), 141–145. <https://doi.org/10.1111/1467-8721.00135>
- Miller, R. R., & Matzel, L. D. (1988). The comparator hypothesis: A response rule for the expression of associations. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 22, pp. 51–92). Academic Press. [https://doi.org/10.1016/S0079-7421\(08\)60038-9](https://doi.org/10.1016/S0079-7421(08)60038-9)
- Miller, R. R., & Schachtman, T. R. (1985). The several roles of context at the time of retrieval. In P. D. Balsam & A. Tomie (Eds.), *Context and learning* (pp. 167–194). Erlbaum.
- Moser, H. W., Raymond, G. V., Dubey, P., Moser, A. B., Xu, J., & Bezman, L. (2000). X-linked adrenoleukodystrophy: Overview and prognosis as a function of age and brain magnetic resonance imaging abnormality. A study involving 372 patients. *Neuropediatrics*, 31(5), 227–239. <https://doi.org/10.1055/s-2000-9236>
- Pearce, J. M., & Hall, G. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, 87(6), 532–552. <https://doi.org/10.1037/0033-295X.87.6.532>
- Pearce, J. M., & Mackintosh, N. J. (2010). Two theories of attention: A review and a possible integration. In C. J. Mitchell & M. E. Le Pelley (Eds.), *Attention and associative learning: From brain to behaviour* (pp. 11–39). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199230142.003.0002>

Rescorla, R. A. (1971). Summation and retardation tests of latent inhibition. *Journal of Comparative and Physiological Psychology*, 75(1), 77–81. <https://doi.org/10.1037/h0030694>

Rodríguez, G., & Hall, G. (2017). Human latent inhibition and the density of predictive relationships in the context in which the target stimulus occurs. *Quarterly Journal of Experimental Psychology*, 70(4), 610–618. <https://doi.org/10.1080/17470218.2016.1143957>

Silverstein, S. M., Keane, B. P., Wang, Y., Mikkilineni, D., Paterno, D., Papathomas, T. V., & Feigenson, K. (2013). Effects of short-term inpatient treatment on sensitivity to a visual context in schizophrenia. *Psychiatry Research*, 210(2), 145–153. <https://doi.org/10.1016/j.psychres.2013.05.001>

Solomon, P. R., & Moore, J. W. (1975). Latent inhibition and stimulus generalization of the classically conditioned nictitating membrane response in rabbits (*Oryctolagus cuniculus*) following hippocampal ablation. *Journal of Comparative and Physiological Psychology*, 89(10), 1192–1203. <https://doi.org/10.1037/h0077183>

Weiner, I. (2003). The “two-headed” latent inhibition model of schizophrenia: Modeling positive and negative symptoms and their treatment. *Psychopharmacology*, 169(3–4), 257–297. <https://doi.org/10.1007/s00213-002-1313-x>

Yassa, M. A., & Stark, C. E. L. (2011). Pattern separation in the hippocampus. *Trends in Neurosciences*, 34(10), 515–525. <https://doi.org/10.1016/j.tins.2011.06.006>