RUNNING HEAD: PRIOR EXPERIENCE AND ACQUISITION

Prior Experience Modifies Acquisition Trajectories Via Response-Strategy Sampling

Carter W. Daniels^{1,2} and Peter D. Balsam^{1,2,3}

¹Department of Psychiatry, Columbia University

²New York State Psychiatric Institute

³Department of Psychology, Barnard College

Author Note

We would like to thank Charles Gallistel, David Freestone, Başak Akdoğan, and Jorge Mallea for helpful conversations that clarified analysis, interpretation, and presentation of data.

We would also like to thank Kimberly Bowman for help with data collection. This work was conducted while Carter Daniels was a post-doctoral fellow at Columbia University.

Correspondence should be addressed to Carter W. Daniels. Email: carter.wa.daniels@gmail.com

Abstract

Few studies have considered how signal detection parameters evolve during acquisition periods. We addressed this gap by training mice with differential prior experience in a conditional discrimination, auditory signal detection task. Naïve mice, mice given separate experience with each of the later correct choice options (Correct Choice Response Transfer, CCRT), and mice experienced in conditional discriminations (Conditional Discrimination Transfer, CDT) were trained to detect the presence or absence of a tone in white-noise. We analyzed data assuming a two-period model of acquisition: a pre-solution and solution period (Heinemann, 1983). The pre-solution period was characterized by a selective sampling of biased response strategies until adoption of a conditional responding strategy in the solution period. Correspondingly, discriminability remained low until the solution period; criterion took excursions reflecting response-strategy sampling. Prior experience affected the length and composition of the pre-solution period. Whereas CCRT and CDT mice had shorter pre-solution periods than naïve mice, CDT and Naïve mice developed substantial criterion biases and acquired asymptotic discriminability faster than CCRT mice. To explain these data, we propose a learning model in which mice selectively sample and test different response-strategies and corresponding task structures until they exit the pre-solution period. Upon exit, mice adopt the conditional responding strategy and task structure, with action values updated via inference and generalization from the other task structures. Simulations of representative mouse data illustrate the viability of this model.

Keywords: Acquisition, signal detection, conditional discrimination, response strategies, prior experience, reinforcement learning

Many psychological processes can be studied within the theoretical framework of signal detection (Wickens, 2002; Wixted, 2019), including time perception (e.g., Brown et al., 1995), reinforcement learning (e.g., Dayan & Daw, 2008), recognition memory (e.g., Wixted, 2007), and perceptual learning (e.g., Gold et al., 2004). Importantly, signal detection analyses have yielded novel and interesting predictions in the aforementioned fields (Alves-Pinto et al., 2012; Andrzejewski & Holder, 2016; Balci et al., 2011; Dayan & Daw, 2008; Gallun & Kampel, 2016; Horga & Abi-Dargham, 2019; MacMillan, 2002; McNamara & Trimmer, 2019; Ratcliff et al., 2017; Ritchie & Carlson, 2016; Wickens, 2002; Wixted, 2007, 2019). Numerous studies have demonstrated that prior experience affects perceptual discriminations (Karim et al., 2012; Kim et al., 2015; Meier & McGee, 1959; Odic et al., 2014), implying that the earlier experiences might affect the learning that underlies perceptual discriminations. This makes it surprising that signal detection models have largely been restricted to characterization of asymptotic, steady-state performance. Indeed, most but not all (Mill et al., 2014) signal detection models do not make contact with the acquisition of steady-state performance or transitions between adjacent steady-states (i.e., new steady-states induced by experimental manipulations); signal detection models are largely silent on how it is individuals *learn* to discriminate the signal from the noise (but see Turner et al., 2011). As such, they are also silent on the sensitivity of component processes (e.g., response-sampling, motivation, attention, memory, etc.; Krakauer et al., 2017) of learning to manipulation. Here we address this gap, isolating acquisition dynamics and their sensitivity to prior experience.

Briefly, the classic signal detection framework assumes that individuals are tasked with discriminating at least two states—a noise state and a noise plus signal state (Killeen, 2016;

Wickens, 2002). Discrimination of such states relies on noisy sensory input communicating whether individuals are in the noise state or the noise plus signal state. Such discriminations may be described by either the likelihood functions corresponding to sensory inputs or the convolution of these likelihood functions and prior beliefs about likelihood function parameters (Horga & Abi-Dargham, 2019; Turner et al., 2011). As shown in Figure 1, it is typically assumed that these likelihood functions can be described by a sensitivity or discriminability (d') parameter and that decisions are based on a bias or criterion (c) parameter (Stanislaw & Todorov, 1999; Wickens, 2002). Whereas discriminability describes the distance between the central tendency of the noise state and noise plus signal state likelihood functions, criterion describes a threshold invoked for labeling some noisy sensory input as originating from the noise likelihood or the noise plus signal likelihood function. In well trained individuals, criterion is typically set at the midpoint of normalized discriminability (i.e., zero).

Figure 1. Schematic of Signal Detection Theory. Shows the probability of a stimulus as a function of stimulus dB level under the Noise State (dashed line) and the Noise + Signal State (solid line) and the corresponding responses (Hit, False Alarm, Correct Rejection, and Miss) given discriminability (blue line) and criterion (orange line) parameters.

The canonical task in which to study signal detection is a two-alternative forced-choice (2AFC) or conditional discrimination task in which a tone is sometimes embedded in white-noise (Green, 1960), having been studied in humans (Zeng et al., 2000), pigeons (Heinemann, 1983), cats (Benovitski et al., 2014), ferrets (Mill et al., 2014), rats (Xiong et al., 2015), and mice (Klink et al., 2006). For example, a mouse might be trained to press a left lever following a tone in white-noise and a right lever following just white-noise, with correct responses reinforced and incorrect responses non-reinforced. While the comparative psychology of auditory signal detection is interesting per se (Fay, 1988; Heffner & Heffner, 2007), it is worth noting that there are substantial commonalities across species, justifying studies in animal models (Carandini & Churchland, 2013; Hanks & Summerfield, 2017; Jaramillo & Zador, 2014). Studies in mice have provided insight into the role and tuning of the auditory cortex and surrounding cortical (Chen et al., 2019) and striatal (Chen et al., 2019; Xiong et al., 2015; Znamenskiy & Zador, 2013) structures, as well as the basal forebrain (Hangya et al., 2015) and thalamus (Chen et al., 2019). Such research has yielded useful neurobehavioral computational models largely targeted at steady-state behavior (Hangya et al., 2015; Mill et al., 2014).

To our knowledge, only a few acquisition studies have been published on auditory signal detection. For example, Heinemann (1983) showed that the rate of acquisition is directly related to the signal-to-noise ratio. Additionally, Kurt and Ehret (2010) showed that although amplitude discriminations are more difficult to acquire than frequency discriminations, amplitude discrimination training facilitates acquisition of frequency discriminations. More recently, Kuchibhotla et al. (2019) showed that evidence of acquisition appears sooner in a mass of

63 signaled non-reinforced probe trials than in training trials alone, suggesting that competition
64 between the reinforcer and task-relevant stimuli for behavioral control interferes with acquisition
65 expression. More recently, while Xiong et al (2015) revealed a role for cortico-striatal plasticity
66 in acquisition, Hangya et al (2015) demonstrated a role for basal forebrain cholinergic neurons
67 projecting to the cortex in signaling and communicating both reward and punishment.
68 Nonetheless, there has yet been any research on the component processes (Krakauer et al., 2017)
69 comprising auditory signal detection acquisition. This makes it unclear which processes affect
70 discriminability and criterion.

71 In principle, performance in 2AFC such as signal detection tasks can be described by
72 popular associative or reinforcement learning theories (Bathellier et al., 2013; Daniels &
73 Sanabria, 2018; Dayan & Daw, 2008; Kuchibhotla et al., 2019; Mill et al., 2014; Sewell et al.,
74 2019). These models assume that in any 2AFC signal detection task, individuals seek to
75 maximize outcomes by making choices with the greatest value and do so by learning the
76 mapping and values between stimuli (tone plus white-noise, white-noise) responses (left lever,
77 right lever), and outcomes (reward, no-reward). To learn this mapping, it is assumed that
78 individuals' sample available responses according to their expected value given the sensory
79 input, where the value of all responses are equal at the beginning of training such that responding
80 is initially random. At the end of each trial the expected value of the response is compared to the
81 value of the obtained outcome, a difference is calculated and the value of the response given the
82 sensory input is updated by some proportion of that difference. Such prediction-error algorithms
83 increment and decrement the value of responses in a cumulative and stochastic manner until
84 conditional responding emerges. An implication of these theories is that acquisition consists of a
85 single solution period in which responding comes under progressive control of sensory input.

The success of reinforcement learning models is a testament to the success of generalizing the insights of associative learning models (Gershman, 2015). However, common to both model classes is the tendency to build and thus assume task structure into the model rather than let the model learn the task structure (for discussions see p. 2 of Jones & Canas, 2010; and p. 54 of Sutton & Barto, 2018). In other words, these models assume individuals are already aware of the stimulus-response-outcome map. This is evident by the fact that these models tend to represent task structure as a two-by-two matrix with cells corresponding to stimulus dependent response-outcome values. All reinforcement learning models need to do is update the value of each cell as responses are randomly sampled. In contrast, Heinemann (1983) suggested individuals learn the stimulus-response-outcome map during a pre-solution period. In the pre-solution period, it is assumed that individuals learn in parallel the stimulus-response-outcome map and associated values (St. Claire-Smith & MacLaren, 1983) by testing different response strategies, including random responding, “Yes” biased responding, “No” biased responding, and conditional responding guided by sensory input. By testing these different response strategies, individuals can learn both the associations among stimuli, responses, and outcomes as well as the response strategy that reliably nets different outcomes based on sensory input. After individuals become aware of the stimulus-response-outcome map, they transition into the solution period in which learning unfolds following reinforcement learning theories.

In the pre-solution period, it is typically assumed that response strategies are randomly sampled with equal probability, meaning that performance in the pre-solution period should be indistinguishable from random responding when aggregated over the whole period. Thus, following the logic of Heinemann’s two-period framework of acquisition, each period may be uniquely characterized by specific patterns of responding: the pre-solution period, random

responding; and the solution period, conditional responding. This description, however, ignores the possibility that during the pre-solution period individuals do not always sample strategies randomly. Within and across sessions, individuals appear to enter phases of biased responding characterized by persistent emission of “Yes” and “No” responses (Ashwood et al., 2020; Roy et al., 2021). Such biased responding is not typically attributed a causal role in acquisition beyond that of hypothesis testing. Other times, biased responding is viewed as a nuisance rather than an important and informative response strategy to facilitate acquisition. It is possible that during the pre-solution period individuals may become biased for prolonged periods of time and that such bias actually plays a causal role above and beyond hypothesis testing as suggested by Heinemann (1983).

Of course, this raises the following questions: What is the composition of the pre-solution period? How does the composition of the pre-solution period relate to acquisition as expressed in terms of signal detection? Additionally, are the effects of prior experience due to changes in the pre-solution period? We focus on the effect of prior experience for both theoretical and practical reasons. Theoretically, despite prior experience clearly facilitating acquisition (Eck & Thomas, 1970; Lany et al., 2007; Rabin, 1988; Saksida, 1999), it is yet unclear how specific kinds of experiences facilitate acquisition and how such facilitation manifests when data are analyzed within the context of the two-period model. Prior experience could affect acquisition by shortening or altering the composition of the pre-solution period; alternatively, it could affect acquisition by inducing higher learning rates across the pre-solution and solution period. Practically, prior experience induced facilitation of acquisition should cut down on the time it takes to train subjects in a signal detection task, which can take upwards of 6000 plus trials (e.g., Heinemann & Chase, 1970; Kuchibhotla et al., 2019).

To answer these questions, we trained mice that differed in their prior experience. One group of mice (Naïve) did not receive any prior training beyond learning to press a bar to earn a reward. A second group, Correct Choice Response Transfer (CCRT) learned the correct stimulus response mappings prior to the discrimination training and a third group, Conditional Discrimination Transfer (CDT) mice had prior training in a conditional discrimination that did not involve auditory cues. We expected that both learning the stimulus-response mapping (CCRT) and having general experience with conditional discriminations (CDT) would facilitate acquisition of auditory signal detection relative to complete naivete. However, because there is little theory on acquisition of auditory signal detection, we made no predictions about how such facilitation might manifest as a function of the quality and quantity of prior experience. Surprisingly, in the current experiment, we found that prior experience in general shortens the pre-solution period length; and that the content of prior experience affects the composition of the pre-solution period, which in turn affects the rate at which subjects reach asymptotic performance. These effects were mirrored in both discriminability and criterion but only discriminability changed systematically during the transition from the pre-solution to the solution period. We discuss the implications of these data for both signal detection and learning theories; offering a proof-of-concept theory that builds off the theorizing initiated by Heinemann (1983).

Methods

Subjects

Sixteen naive and 8 experienced ($N = 24$) adult male C57/BL6 mice (Jackson Labs, Bar Harbor, ME, USA) served as subjects in the present experiments. The naïve mice arrived from Jackson labs and were kept on free-feeding for a week while mice acclimated to the colony room and prior to implementation of food-restriction. The 8 experienced mice were previously trained

in an externally paced (i.e., trials started automatically after an inter-trial interval elapsed) visual working-memory task wherein subjects were required to choose the lever over which a yellow light was illuminated; choice of the lever over which there was no light was not reinforced. To challenge working memory, the duration of the yellow light was varied across conditions (the associated data are not yet published). These mice were off-task, under free-feeding conditions for roughly a month prior to onset of food restriction for this experiment. Mice were housed four per cage in a temperature- and humidity-controlled vivarium on a 12 h light/12 h dark light cycle; with lights on at 7 am and lights off at 7 pm. Free access to water was provided in home cages. Mice were food restricted such that before every experimental session, mice were approximately 85-90% of their free feeding weight. All animal handling procedures in the present study followed National Institutes for Health guidelines and were approved by both the Columbia University and New York State Psychiatric Institute Institutional Animal Care and Use Committee.

Apparatus

Experiments were conducted in eight Med Associates (St. Albans, VT, USA) modular test chambers (ENV-307W), enclosed in a sound- and light-attenuating cubicle (MED-OFA-022). Each cubicle was equipped with an exhaust fan, which provided approximately 60 dB of masking noise inside the chamber. Chambers measured 240 mm long $\times$ 200 mm wide $\times$ 165 mm tall. Each chamber was equipped with a floor of metal rods spaced 0.87 cm apart (ENW-307W-GFW), a stainless-steel waste pan, back and front aluminum walls and two plexiglass side walls, one of which was hinged and served as the door to the chamber. A trough furnished with an infrared head-entry detector (ENW-302HD), ceiling mounted yellow light (trough-light; ENV-302RL-1), and liquid dipper (ENW-302W-S) providing 0.01 cc of reinforcement (evaporated

milk; Red & White; Federated Group, Buffalo, NY, USA) was centered on the front wall. Flanking the trough were two retractable levers (ENV-312-3W), one on the left and one on the right. Above each retractable lever was a 3-stimulus LED panel (ENV-322W). On the back-aluminum wall, a house-light (ENV-315W) was mounted top-center, flanked by two speakers (ENV-324W), one on the left and one on the right; the left speaker was connected to a tone-generator (ANL-926) and a right speaker was connected to a white-noise generator (ENV-225SM). Experimental events were programmed and controlled by a Dell PC with Windows 10 and MED-PC® IV Software. Sessions were one hour long, allowing mice to complete approximately between 100-120 trials/session and conducted 7 days a week.

Procedure

Dipper & Lever Training. All mice were first trained to retrieve reinforcement (0.01 cc evaporated milk) from the liquid dipper on a variable time 45-s schedule of reinforcement. After three sessions of dipper training, mice were trained to press the left lever, right lever, or make a head-entry into the trough for reinforcement on a fixed-ratio 1 schedule of reinforcement with a 10-s inter-trial interval. This training continued for 6-10 sessions.

Response-Initiation Training. Mice were then trained to self-pace trials. At the beginning of each trial, the house-light was turned on, signaling that a trial could be initiated. A single head-entry into the trough turned off the house-light, turned on the yellow LEDs above both levers and initiated a 1-s delay. At the end of this delay the yellow LEDs were turned off and red LEDs turned on along with either the left or right lever. A single response on the activated lever turned off the red LEDs and retracted the lever while activating the dipper and turning on the trough-light. A single subsequent head-entry into the trough turned off the trough-light and 2.5 s later deactivated the dipper. Deactivation of the dipper was followed by a dark 10-s inter-trial

interval (ITI) during which no stimuli were present. At the end of the ITI, the house-light was again turned on. The 1-s delay between trial initiation and lever presentation was increased to 5-s after two sessions; after an additional three sessions, the 5-s delay was converted into a variable-time delay that averaged 10-s (VT10s), with a minimum delay of 2.5 s. Note there was no requirement that mice initiate trials within sometime window; mice could initiate trials quickly or slowly. Training continued for another three sessions. The LEDs were discontinued after this phase.

Auditory Signal Detection Training. After learning to self-pace trials, all 8 experienced mice (CDT) and 8 randomly chosen naïve mice (Naïve mice) were trained with all choice trials and the remaining 8 naïve mice (CCRT) were initially trained with forced correct trials with choice trials faded in (see below).

Figure 2. Schematic of Auditory Signal Detection Task. Trials start with illumination of the house-light, signaling a trial is available. Following a head-entry into the trough, the house light turns off, white-noise turns on, and a variable 10-s delay is started. At the end of the delay, with probability $q = 0.5$ either a 2-s 80 dB 3.5 kHz tone is added to the white-noise or a 2-s delay is initiated. At the end of the 2-s tone or 2-s delay all auditory stimuli are terminated and the levers extended. Following a correct choice, the dipper is activated until a head-entry occurs confirming receipt of milk, followed by deactivation of the dipper 2.5-s later and a 10-s ITI. Incorrect choices are followed by a 2.5-s black out and a 20-s ITI.

1
2 A schematic of the auditory signal detection task is shown in Figure 2. For Naïve and
3 CDT mice trial-initiation was followed by the onset of 60 dB white-noise. A trial-type was
4 determined pseudo-randomly such that Tone and No-Tone trials were equally likely but a single
5 trial type could appear no more than 3 times in any block of 6 consecutive trials. In Tone trials,
6 the VT 10-s terminated with a 2-s 3.5-kHz 80-dB tone at the end of which both the tone and
7 white-noise were turned off and both levers extended; in No-Tone trials, the VT 10-s terminated
8 with an additional 2-s of white-noise after which the white-noise was turned off and both levers
9 extended. Following lever extension, a single response on either lever retracted both levers and if
10 correct was reinforced and followed by an ITI as described in *Response-Initiation Training*.
11 Incorrect trials were followed by a 20-s rather than a 10-s ITI. For half of the mice in each group,
12 the correct lever was randomly determined to be the left lever in Tone trials and the right lever in
13 No-Tone trials; for the other mice, the opposite was true. There was no time limit for mice to
14 make their choices, they could take a short time or a long time to make a choice; thus, mice
15 could not make errors of omission. Sessions terminated after one hour. Most mice completed
16 between 100 and 120 trials within a session.

17 CCRT mice had their programs modified as above but were initially trained with either
18 all Tone and then No-Tone or all No-Tone and then Tone trials terminating with extension of
19 only the correct lever for four sessions (2 sessions per trial type). This was followed by
20 intermixing these two single-lever trial types for two sessions. Tone and No-Tone trial types
21 with extension of both levers (i.e., choice trials) were then introduced in the following sequence:
22 five sessions of 50%, five sessions of 75%; and for the remainder of the experiment all trials

were 100% choice trials. The correct responses in the choice trials were identical to the responses made during the single-lever pretraining.

All mice were trained for a minimum of 25 sessions from onset of choice and until the probability of a correct choice in both Tone and No-Tone trials was greater than 0.75 for at least three consecutive sessions.

Data Analysis

In the present paper, we focus on acquisition dynamics as expressed in choice behavior because we largely expected the effect of prior experience to be expressed on choice rather than initiation and choice latencies. Nonetheless, we present a latency analysis in the Supplementary Materials.

Acquisition periods and in turn response strategies were demarcated via a two-stage maximum likelihood algorithm that first determined whether the number of obtained reinforcers indicated whether mice were in the pre-solution or the solution period as a function of session. If the number of reinforcers indicated the pre-solution period, we determined whether the distribution of “Yes” tone responses were most likely under the random responding or biased responding strategies, the latter of which was further subdivided into biased “Yes” and biased “No” responding (see Supplemental Materials for more details). From this analysis we calculated three measures: the length of the pre-solution period as the session immediately followed by at least three consecutive sessions of conditional responding; the probability of switching response strategies as the number of response strategy transitions divided by the pre-solution length; and the proportion of time spent in each response strategy as the number of sessions in a response strategy divided by the pre-solution length.

Signal detection parameters discriminability (d' ; a.k.a. sensitivity) and criterion (c ; a.k.a. bias) were calculated from hits (reporting a tone when there was a tone) and false alarms (reporting a tone when there was not a tone) as previously described (Stanislaw & Todorov, 1999),

$$\text{Discriminability} = d' = \theta^{-1}(\text{Hits}) - \theta^{-1}(\text{False Alarms}) \quad (1)$$

$$\text{Criterion} = c = \frac{\theta^{-1}(\text{Hits}) + \theta^{-1}(\text{False Alarms})}{2} \quad (2)$$

where θ^{-1} denotes the Z-score of the hit and false alarm rate, and thus discriminability is the difference in the Z-scores and criterion is the average of the Z-scores. Note that $d' < 0$ is possible, thus to facilitate description of d' , we set $d' = 0$ if $d' < 0$.

All data were log or log-odds transformed if visual inspection suggested deviation from a normal distribution and/or the domain of the data was between 0 and 1, or 0 and infinity. Such a transformation is consistent with the observation that many underlying neurobiological processes are consistent with a log normal distribution (Limpert et al., 2001). Null hypothesis testing and their Bayesian variants were conducted in JASP (Daniels & Sanabria, 2017b; JASP Team, 2019; Love et al., 2019; Wagenmakers, Love, et al., 2018; Wagenmakers, Marsman, et al., 2018). For Bayesian variants, the natural log transformed Bayes Factor is reported comparing the model of interest to the null model; whereas a $\ln BF_{10} \geq 1.098$ indicates substantial evidence for the model of interest, the $\ln BF_{10} \leq -1.098$ indicates substantial evidence for the null model (Jarosz & Wiley, 2014). Few significant statistical tests were not also supported by Bayesian tests; statistical tests that were statistically significant but not supported by substantial evidence are noted. For all analyses and claims reported below, corresponding data for individual mice if not visualized, including visualization of model fits, can be found in both the Supplemental Materials or OSF

materials (DOI 10.17605/OSF.IO/FJZ87). Note that due to operant chamber malfunctions, data on session 22 was omitted for Naïve and CCRT mice.

Results

Defining the Pre-Solution and Solution Period

Hit rates increased and false alarm rates decreased as a function of session for all three groups. As shown in Figure 3A, CCRT and CDT mice increased their hit rates sooner than Naïve mice. CCRT mice decreased their false alarm rates sooner than Naïve and CDT mice. These differences were also reflected in signal-detection parameters as shown in Figure 3B & 3C, with discriminability increasing sooner for CCRT and CDT than Naïve mice; and with criterion taking excursions substantially above zero, corresponding to “Yes” biases, or below zero, corresponding to “No” biases, for CDT and Naïve mice, respectively. This average data suggests mice differ in (1) how long they spend in the pre-solution period and (2) how quickly they reach asymptotic performance. CCRT and CDT mice appear to exit the pre-solution period sooner than Naïve mice. Once they have left the pre-solution phase, Naïve and CDT mice reach asymptotic performance more abruptly than CCRT mice.

Figure 3. Acquisition of Tone vs. No-Tone Discrimination. Mean ($\pm$ SEM: transparent colored patches) P(Hit = Blue) and P(False Alarm = Orange), Discriminability (d' = blue; Panel B), and Criterion (c = orange; Panel C) as a function of session for each training group (Naïve= darkest color, solid line; CCRT= middle color, dashed line; CDT = lightest color, dotted line). CDT mice were each trained for 28 sessions; Naïve and CCRT mice required variable amounts of training, with modes of 31 and 27 sessions, respectively; as such, data are only shown up to the modes, but all data comprise analyses. Data in Panels A and B are back transformed from the log scale.

Inspection of the performance of individual representative mice (see Supplemental Figures 3 and 4 for individual data of all mice) confirms the above observations, but with one caveat: performance prior to asymptote is idiosyncratic. As shown in Figure 4A-C, Naïve and CCRT mice start by responding randomly but then the Naïve mouse develops a substantial “No” bias before abruptly showing conditional responding. Rather than develop a substantial bias, the CCRT mouse begins to progressively show conditional responding soon after random responding. In contrast, the CDT mouse starts with a substantial “Yes” bias and then comes under conditional responding. Such dynamics are also reflected in signal detection parameters. As shown in Figure 4D-F, discriminability starts at zero for Naïve and CCRT mice but near 1 for CDT mice. As acquisition proceeds, discriminability either abruptly or progressively increases to asymptotic levels between 2.5 and 3.5. Interestingly, transitions to asymptote appear to be immediately preceded or followed by criterion developing substantial “Yes” or “No” biases. This is particularly interesting given that otherwise the criterion appears to make random and sudden changes around the idealized mean of zero.

Figure 4. Acquisition of Tone vs. No-Tone Discrimination in Representative Mice. $P(\text{Hit} = \text{Blue})$ and $P(\text{False Alarm} = \text{Orange})$ in Panels A-C, Discriminability ($d' = \text{blue}$), and Criterion ($c = \text{orange}$) in Panels D-F as a function of session for each representative mouse from each training group (Naïve= darkest color, solid line, Panels A and D; CCRT= middle color, dashed line, Panels B and E; CDT = lightest color, dotted line, Panels C and F). Representative mice were chosen as the mice closest to the median of parameter estimates of Eq. 3 for each group.

1
2
3
4
5
6
7
8
9
10
11
12
13
14

To relate these observations to the pre-solution and solution period, we first passed our data through our two-stage maximum likelihood classification algorithm. (see Data Analysis and Supplemental Materials). As shown in Figure 5A-C, there is no clear order in how response strategies are deployed during the pre-solution period. Indeed, there doesn't appear to be a distinct pattern in which subjects progress from a random responding strategy to a biased responding strategy prior to adopting the conditional responding strategy. Most mice start with a biased response strategy and then unpredictably switch into a random responding strategy and back to a biased response strategy, sometimes with a different direction of bias. There is, however, a clear transition from the pre-solution period to the solution period. Despite this clear transition, it appears that sometimes conditional responding emerges in the pre-solution period but is not always adopted; likewise, sometimes biased responding or random responding emerges during the solution period but is quickly abandoned.

Figure 5. Session-by-session Classification of Response Strategies during Acquisition. Random responding (white), biased “Yes” responding (blue), biased “No” responding (orange), and conditional responding (black) as determined by our two-stage algorithm (see Data Analysis and Supplemental Materials) as a function of session for each mouse in each group (Panel A = Naïve, Panel B = CCRT, and Panel C = CDT). Although not marked in this figure, the pre-solution period length was defined as the session followed by three consecutive sessions of conditional responding.

Table 1 (for individual data see Supplemental Table 1) shows the relative composition of the pre-solution period for each group of mice. Analysis of the pre-solution period revealed that the length of the pre-solution period differed with prior experience [$F(2,21) = 5.265, p = 0.014, \ln BF_{10} = 1.481$]. Post-hoc tests indicated the pre-solution period was longer for Naïve than CCRT [$t(14) = 3.025, p = 0.019, \ln BF_{10} = 1.290$] and CDT mice [$t(14) = -2.530, p = 0.039, \ln BF_{10} = 3.617$] but similar for CCRT and CDT mice [$t(14) = 0.495, p = 0.626, \ln BF_{10} = -0.790$]. The composition of the pre-solution period was also affected by prior experience. Prior experience substantially affected the proportion of time spent in “Yes” [$F(2,21) = 12.780, p < 0.001, \ln BF_{10} = 4.723$] and conditional [$F(2,21) = 5.246, p = 0.014, \ln BF_{10} = 1.471$] responding; prior experience significantly but did not substantially affect the proportion of time spent in “No” [$F(2,21) = 4.098, p = 0.031, \ln BF_{10} = 0.892$] and did not affect random [$F(2,21) = 0.270, p = 0.766, \ln BF_{10} = -1.252$] responding. Specifically, CDT mice spent proportionally more time responding “Yes” than Naïve [$t(14) = 3.105, p = 0.014, \ln BF_{10} = 1.312$] or CCRT [$t(14) = 5.008, p < 0.001, \ln BF_{10} = 4.238$] mice, indicating CDT mice were “Yes” biased. Naïve mice spent proportionally more time responding “No” than CDT [$t(14) = 2.565, p = 0.046, \ln BF_{10} = 1.278$] but not CCRT [$t(14) = 2.384, p = 0.066, \ln BF_{10} = -0.838$] mice indicating that Naïve and CCRT mice were both “No” biased. Interestingly, CCRT mice spent proportionally more time responding conditionally than Naïve [$t(14) = 2.862, p = 0.024, \ln BF_{10} = 1.111$] or CDT [$t(14) = -2.744, p = 0.031, \ln BF_{10} = 0.841$] mice indicating that CCRT mice had already discovered and were testing the conditional responding strategy during the pre-solution period. Prior experience did not affect the probability of switching response strategies [$F(2,21) = 1.790, p = 0.192, \ln BF_{10}$

= -0.358]. Thus, it appears that prior experience or lack thereof can result in particular response biases dominating the pre-solution period without altering the rate of response strategy sampling.

Table 1. Mean ($\pm$ SEM) metrics of the pre-solution period

	Naïve	CCRT	CDT
Pre-Solution Length	21.87 (19.91-21.87)	8.37 (5.95-11.77)	9.79 (8.33-11.52)
Proportion Random	0.34 (0.24-0.45)	0.25 (0.16-0.36)	0.26 (0.21-0.33)
Proportion Yes	0.07 (0.05-0.11)	0.02 (0.01-0.02)	0.47 (0.65-0.29)
Proportion No	0.51 (0.37-0.64)	0.47 (0.31-0.64)	0.09 (0.05-0.16)
Proportion Conditional	0.01 (0.01-0.02)	0.07 (0.04-0.12)	0.01 (0.01-0.02)
Probability Switch	0.23 (0.21-0.26)	0.33 (0.28-0.38)	0.28 (0.25-0.32)

Note. Bolded numbers indicate the group with the largest estimate as indicated by statistical analysis. Data are back transformed from the log scale.

Discriminability

To characterize the relationship between acquisition periods and signal detection parameters, we first conducted a time series analysis to determine the degree of hysteresis within each parameter and thus whether they could be easily described by tractable closed-form functions typically invoked for acquisition (e.g., logistic growth or cumulative Weibull). The time series analysis of discriminability indicated that it showed a substantial degree of hysteresis, indicating it could be easily modeled (see Supplemental Material; including Supplemental Figures 3-7). Visual inspection suggested that discriminability could be described as a function of session for each mouse by fitting a cumulative Weibull equation (Gallistel et al., 2004). Importantly, the cumulative Weibull has two important features: it is relatively flexible allowing it to capture many different shapes of acquisition; and it also offers an estimate of the pre-solution period that can be compared to the our two-stage algorithm. Specifically,

$$\text{Discriminability}_i = A(1 - 2^{-(i/L)^s}) \quad (3)$$

where i is the i th session, A is asymptotic discriminability, L is the midpoint latency at which half of A is obtained and may be considered a proxy measure of the pre-solution length, and s is the abruptness of transition or learning rate; as s increases, acquisition approaches a step-function like process. Note that s is normalized to L , making comparisons across groups of mice with substantially different estimates of L meaningful. Eq. 3 was fit via a mixed-effects approach as described in the supplemental materials.

Eq. 3 provided a good description of discriminability as a function of session, approximating a sigmoid-like function for almost all mice (for individual fits and parameter estimates see Supplemental Figure 8 and Supplemental Table 2). Interestingly, as shown in Figure 6A-C, prior experience affected the number of sessions to reach the midpoint (L) and the abruptness of the transition to asymptotic performance (s) but not asymptotic performance (A). All mice obtained similar levels of asymptotic discriminability [$F(2,21) = 1.954, p = 0.166$; $\ln\text{BF}_{10} = -0.264$] indicating that previous experience does not affect the solution period but rather the route taken through the pre-solution to reach the solution period.

Figure 6. Parameter Estimates of Discriminability Dynamics (fits of Eq. 3). Boxplots of parameter estimates (Asymptotic Discriminability, A , Panel A; Midpoint Latency, L , Panel B; Sensitivity, s , Panel C) of each training group (Naïve= darkest color; CCRT= middle color; CDT = lightest color) overlaid with the parameter estimates of each mouse. Each boxplot shows the inter-quartile range (patches), 95% confidence interval (vertical lines), mean (horizontal solid line) and median (horizontal dashed line). *Indicates a statistically significant comparison ($p < .05$) corroborated by substantial evidence for a difference between groups ($\ln BF_{10} \geq 1.098$). Data were analyzed on the log scale but are presented back transformed on the linear scale.

Previous experience affected estimates of L [$F(2,21) = 11.331, p < 0.001; \ln BF_{10} = 4.163$; see Figure 6B] just as it did pre-solution period: prior experience resulted in shorter estimates of L . Whereas Naïve mice had longer estimates of L than CCRT [$t(14) = 3.414, p < 0.004; \ln BF_{10} = 2.283$] and CDT [$t(14) = 6.417, p < 0.001; \ln BF_{10} = 6.790$], CCRT and CDT mice had similar estimates of L [$t(14) = 0.667, p = 0.515; \ln BF_{10} = -0.699$]. Because of the similar effect of prior experience on both estimates of L and pre-solution period, we tested whether these estimates were related. We found a strong positive correlation ($r = 0.942, p < .001; \ln BF_{10} = 20.703$), indicating that L may serve as a proxy measure for the length of the pre-solution period (see Supplemental Figure 9). These data corroborate the finding that prior experience reduces the length of the pre-solution period.

Previous experience also affected estimates of s [$F(2,21) = 18.295, p < 0.001; \ln BF_{10} = 6.624$; see Figure 6C]. Naïve [$t(14) = 7.162, p < 0.001; \ln BF_{10} = 7.828$] and CDT mice [$t(14) = -3.352, p = 0.005; \ln BF_{10} = 2.192$] had larger estimates of s compared to CCRT mice, Naïve and CDT mice had similar estimates of s [$t(14) = 2.106, p = 0.053; \ln BF_{10} = 0.517$]. Such a finding is curious, because it is unclear why CCRT mice, who have knowledge of correct choices, would acquire discriminability slower than both Naïve mice who have no knowledge of the task or CDT mice who only have knowledge of a task-relevant response rule.

Criterion

Unlike discriminability, the time series analysis (see Supplemental Materials; including Supplemental Figures 10-14) suggested less criterion hysteresis and that if anything, the criterion may be best thought of as a mean reverting, shock sensitive variable randomly sampled from a

normal distribution. As such, when the time-series analysis was taken together with visual inspection it was clear that criterion could not be easily described by a closed form equation. Thus, we decided to empirically analyze the criterion distribution in the pre-solution and solution period.

We aligned the criterion around the pre-solution period as determined by our two-stage classification algorithm. Figure 7 shows the aligned criterion as a function of the median length of the pre-solution period and solution period for each group of mice. Visual inspection of these data confirmed the biases observed for each group of mice and suggested slight changes in criterion as they exited the pre-solution period and into the solution period. Whereas for Naïve mice criterion biases appear to resolved as they exit the pre-solution period, for both CCRT and CDT it appears criterion biases are resolved after exiting the pre-solution period. Inspection of individual data largely confirmed these observations (see Supplemental Figure 4), but also revealed that for most mice most large excursions in the criterion preceded the pre-solution period or immediately followed the transition into the solution period. It is worth reiterating that there are strong individual differences in criterion trajectories. Whereas some mice show decreasing criterions, others increasing criterions, and others flat criterions; most mice show strong excursions into biased criterions, but eventually recover to the ideal criterion with a mean of zero.

Figure 7. Criterion aligned to the last session of the pre-solution period. Mean ($\pm$ SEM: transparent colored patches) criterion as a function of session relative to the last session of the pre-solution period (as determined by our two-stage algorithm) for each training group (Naïve= darkest color, solid line; CCRT= middle color, dashed line; CDT = lightest color, dotted line). The solid vertical line denotes the last session of the pre-solution period. The number of sessions represented in the pre-solution and the solution periods are the median of each period for each group of mice.

To characterize the average difference between the pre-solution and solution period as a function of group as depicted in Figure 7, we calculated the following parameters of criterion distributions during both the pre-solution and solution period: minimum, maximum, mean, and root mean square. Figure 8A-D shows these distribution parameters when the pre-solution and solution period are defined by the two-stage algorithm (for individual data see Supplemental Table 3). Although prior experience significantly but not substantially affected the minimum [$F(2,21) = 3.961, p = 0.035, \ln BF_{10} = 0.473$; see Figure 8A], prior experience significantly and substantially affected both the mean [$F(2,21) = 4.906, p = 0.018, \ln BF_{10} = 1.253$; Figure 8B] and the maximum [$F(2,21) = 4.154, p = 0.030, \ln BF_{10} = 1.299$; Figure 8C]. Post-hoc tests revealed that these group effects were driven by CDT mice having a higher mean [$t(14) = 3.119, p = 0.016, \ln BF_{10} = 1.253$] and maximum [$t(14) = 2.687, p = 0.041, \ln BF_{10} = 2.198$] criterion than

60 Naïve mice. Although the group effect was not substantial, post-hoc tests revealed that Naïve
 61 mice had smaller criterion minimums than CDT mice [$t(14) = 2.683, p = 0.042, \ln BF_{i0} = 2.100$].
 62 In contrast, the root mean square was insensitive to prior experience [$F(2,21) = 0.199, p = 0.821,$
 63 $\ln BF_{i0} = -1.521$] but sensitive to the acquisition period [$F(1,21) = 5.356, p = 0.031, \ln BF_{i0} =$
 64 1.706 ; see Figure 8D]. The root mean square decreased from the pre-solution to the solution
 65 period. All other comparisons were neither significant nor substantial [all $ps > 0.05$ and all $\ln BF_{i0}s$
 66 < 1.098]. Taken together, these data indicate that as mice transition from the pre-solution to the
 67 solution period, criterion variance decreases and prior experience biases the criterion to the
 68 greatest degree for mice with either no knowledge or knowledge of a task-relevant response rule.

69

1

2 **Figure 8. Empirical Estimates of Criterion Distributions.** Boxplots of criterion distribution parameters (Minimum, Panel A; Mean,
3 Panel B; Maximum, Panel C; Root Mean Square, Panel D) of each training group (Naïve= darkest color; CCRT= middle color; CDT =
4 lightest color) as a function of acquisition period overlaid with estimates of each mouse with lines connecting the estimates of each
5 mouse between acquisition periods. Each boxplot shows the inter-quartile range (patches), 95% confidence interval (vertical lines),
6 mean (horizontal solid line) and median (horizontal dashed line). *Indicates a statistically significant comparison ($p < .05$)
7 corroborated by substantial evidence for a difference between groups ($\ln BF_{10} \geq 1.098$). All data except for the root mean square was
8 analyzed on the linear scale, the root mean square was analyzed on the log scale and back transformed to the linear scale.

Discussion

The purpose of the present study was to characterize the effect of prior experience on learning of signal detection manifest in pre-solution period composition and corresponding changes in signal detection parameters. Generally, we found that prior experience facilitates discovery of the correct stimulus-response-outcome mapping and thus exiting of the pre-solution period (Figure 3, 5, and Figure 6B; Supplemental Figures 3 and 4). However, the content of prior experience determines the composition of response strategies that dominate the pre-solution period and in turn the rate of learning once subjects have exited the pre-solution period.

Specifically, we found that the pre-solution period is comprised of three response strategies: “Yes” biased, “No” biased and random responding; for mice experienced with correct choices (CCRT) conditional responding sometimes emerged transiently during the pre-solution period (Table 1, Figure 4 and 5; Supplemental Table 1). Rather than sampling each strategy equally, mice tended to select one strategy over another; frequently the “Yes” biased or “No” biased rather than random response strategy. Initially, sampling of these response strategies was associated with discriminability close to zero and noisy, somewhat biased criterion. Around the end of the pre-solution period, criterions became transiently and substantially biased (Figure 3C, Supplemental Figure 4). After such bias, mice adopted the conditional responding strategy, with criterion abruptly returning to zero with less noise (Figure 7 and Figure 8D; Supplemental Table 3) and discriminability either abruptly or progressively increasing from zero to asymptote (Figure 6C; Supplemental Table 2). Whereas discriminability was well described by a closed form mathematical function (a cumulative Weibull; see Supplemental Figure 8), the criterion is likely best thought of as a shock sensitive, mean reverting random sample from a normal

distribution with a mean of zero. In contrast to canonical signal detection theories, these data indicate that discriminability but not criterion are most obviously sensitive to feedback.

Interestingly, prior experience affected both the composition of the pre-solution period and signal detection parameters. Although prior experience shortened the pre-solution period (Table 1, Figure 6 Panel B), the content of prior experience caused mice to select certain response strategies during the pre-solution period (Figure 5). Naïve mice selected the “No” biased, CCRT mice selected either “No” biased or random with conditional responding emerging for a single transient session, and mice previously experienced with conditional discriminations (CDT) selected “Yes” biased responding (Figure 3 and 4). Surprisingly, selecting biased response strategies during the pre-solution period appears to be associated with a higher rate at which subjects acquire asymptotic discriminability. Naïve and CDT mice had higher rates of learning than CCRT mice (Figure 6 Panel C, Supplemental Table 2). This latter effect is at odds with existing theories of learning.

We discuss the implications of these findings for both signal detection theory and learning theories in turn.

A Challenge for Signal Detection Theories

Although signal detection theories are largely silent on the nature of acquisition, most signal detection theories argue that the criterion but not discriminability is sensitive to feedback (Killeen et al., 2018; Mill et al., 2014; Mistry et al., 2018; Mueller & Weidemann, 2008; but see Turner et al., 2011). Indeed, criterion learning is thought to be sensitive to manipulations of signal probability and payoffs which in turn affect the feedback animals receive about their performance (Killeen et al., 2018; McCarthy & Davison, 1981; Mill et al., 2014; Teichert &

Ferrera, 2010). Most models implementing these ideas, however, have been developed to describe criterion adjustment during manipulation of steady state behavior (e.g., Mill et al., 2014). This leaves open the question how well these models generalize to acquisition. A more general implication of the claim that criterion but not discriminability are sensitive to feedback is that discriminability should not change systematically over the course of acquisition. It may fluctuate in a stochastic manner, but it should not be altered as subjects learn the stimulus-response-outcome mapping. However, this is in contrast to the continued observation and acknowledgment that discriminability improves with training (Ball & Sekuler, 1982; D'Amato & Jagoda, 1961; Killeen et al., 2018; Kurylo et al., 2017).

Likewise, we found that discriminability increases as a function of session, with such dynamics well described by a cumulative Weibull function (Eq. 3, Figure 3B, Supplemental Figure 8; also see Supplemental Figures 3-7; Gallistel et al., 2004). This effect indicates that discriminability is sensitive to feedback and may thus be considered acquired. Following the theorizing of Lawrence (1950a, 1950b), we suspect that acquired discriminability reflects a process wherein the intrinsic properties of stimuli (in this case, tonal qualities such as amplitude and frequency) become more salient because of their differential relationship to differential response-outcome associations rather than the inherent differences between stimuli. Importantly, the notion of acquired discriminability is consistent with the observation that tuning in the auditory cortex is experience dependent (Shepard et al., 2013; Trachtenberg et al., 2002; Yanagihara & Yazaki-Sugiyama, 2016), showing sensitivity to both training regimen and response strategy (Bieszczad & Weinberger, 2010; David et al., 2012). Moreover, despite the dearth of signal detection acquisition data, recent models of human signal detection suggest that humans have to learn the distributions associated with both the noise and noise plus signal state

(Turner et al., 2011). Thus, converging data suggests that discriminability is acquired, indicating that discriminability does not describe just the inherent difference between a noise and signal state but rather the degree to which differing stimulus qualities become relevant in capturing attention and thus guiding behavior.

Even if discriminability is acquired, it does not rule out that criterion are acquired as well. In our data, however, criterion was not obviously sensitive to feedback during training. Criterion did not appreciably change in a systematic manner as mice mastered the auditory signal detection task (Figure 3 Panel C, Supplemental Figure 4, 10-14). Indeed, criterion largely appear to be mean reverting, shock sensitive random samples from a normal distribution whose variance decreases as mice move out of the pre-solution period and into the solution period. At a minimum, this suggests that the influence of feedback on criterion is not through the operation of any prediction-error algorithm on the criterion per se (as in Mill et al., 2014). Instead, criterions are much less predictable, likely reflecting the kind of hypothesis testing suggested by Heinemann (1983) in which subjects sample and test different response strategies until settling on the response strategy appropriate to the task at hand (but see Witt et al., 2015 for caution in this interpretation). It is likely that criterion dynamics alone do not adequately capture such hypothesis testing. While criterion dynamics can indicate “No” and “Yes” biased responding unequivocally, criterion dynamics are unable to distinguish between random and conditional responding as both take on a criterion value close to zero. Alternatively, as mice sample and test different response strategies, it seems possible that what is updated via feedback is not the criterion per se but the probabilities with which response strategies are sampled. If these probabilities are updated via feedback, the nature of such feedback is unclear because it seems that while some mice largely persist with one strategy until adopting the conditional responding

strategy other mice sample all strategies. This suggests that the underlying criterion dynamics and idiosyncrasies observed here and elsewhere (Ashwood et al., 2020; Bathellier et al., 2013; Heinemann, 1983; Roy et al., 2021; Stüttgen et al., 2011) reflect some yet unclarified process.

A Challenge for Learning Theories

It is clear that for signal detection theories to adequately describe discriminability and criterion dynamics, fundamental assumptions about discriminability and criterion need to be revisited. Alternatively, associative and reinforcement learning theories may provide a more direct path by which to explain the present findings. Such theories focus on describing raw choice behavior, letting signal detection parameters emerge from calculations based on said behavior. Most of these theories are consistent with Heinemann's (1983) suggestion that acquisition be divided into a pre-solution and a solution period. Reinforcement learning models tend to account for acquisition periods by assuming that subjects may start out with a bias for one response over another and/or that learning rates evolve throughout training (Bathellier et al., 2013; Kuchibhotla et al., 2019). Although such modifications along with others (e.g., the effect of previous trials on performance in the next trial, Roy et al., 2021), descriptively capture the periods of acquisition they do not provide a *mechanistic psychological account* built upon Heinemann's theorizing. In other words, these models do not explain how subjects learn the appropriate task structure or stimulus-response-outcome mappings in which values are learned and action policies executed via response strategy sampling. This question of how mice learn the task structure has begun to receive direct attention from those interested in learning processes (e.g., Niv, 2019). We believe that addressing this question will allow learning models to capture a greater range of acquisition phenomena, including those reported in the present paper.

First, prior experience, even in the absence of stimulus-response-outcome values, shortens the pre-solution period (Table 1, Figure 5, Figure 6 Panel B); that is, the time it takes to sample, test, and deploy the conditional responding strategy. Mice previously experienced with correct choices (response mappings consistent with hits and correct rejections) in the auditory signal detection task (CCRT mice) and mice previously experienced with conditional responding (CDT mice) had shorter pre-solution periods than Naïve mice. To the extent that similar behavioral patterns imply similar mechanisms, the similarity of these effects despite their different experiences suggests shorter pre-solution periods are *not* the result of learning stimulus-response-outcome values. Whereas CCRT mice previously learned the mapping between the value of choosing (for example) the left lever after the tone plus white-noise for reward and the right lever after just white-noise for reward, CDT mice did not learn any stimulus-response-outcome values. Instead, CDT mice entered the task having previously learned that conditional responding is a viable response strategy, which implies a stimulus-response-outcome mapping consistent with that imposed by the signal detection task. This suggests prior experience shortens the pre-solution period predominantly because it facilitates discovery of the correct stimulus-response-outcome mapping not because of stimulus-response-outcome value accrual. Further reinforcing this conclusion is the observation that most CCRT mice did not make only correct responses as choice trials were introduced (see Supplemental Figure 3).

Second, conditional responding was mostly acquired through selection of biased response strategies rather than random sampling of all response strategies. Naïve mice and CDT mice selected to sample and test the “No” and “Yes” biased response strategies during the pre-solution period, respectively (Table 1, Figure 5). Such selections were also reflected in criterion estimates: criterion was most negative for Naïve mice and most positive for CDT mice (Figure

3C, Figure 7, and Figure 8A and 8C). In contrast, although CCRT mice also selected to sample the “No” biased response strategy, the depth of their bias was shallower than observed in Naïve mice; CCRT mice were also more likely to test and sample the conditional response strategy during the pre-solution period. Surprisingly, biased mice (Naïve and CDT) also transitioned more abruptly from low discriminability to asymptotic discriminability indicating a higher learning rate than CCRT mice (Figure 6 Panel C). This suggests that knowledge of stimulus-response-outcome-values (as trained in CCRT mice) may actually interfere with acquisition of asymptotic discriminability. Such interference appears to be due to a reduced sampling and testing of “No” and “Yes” biased response strategies during the pre-solution period. This indicates that when sampling and testing response strategies to learn the stimulus-response-outcome mappings, sampling and testing biased response strategies may be a more effective way by which to master conditional discriminations.

To explain these findings, we hypothesize that learning of stimulus-response-outcome mappings and values (for evidence of their independence, see St. Claire-Smith & MacLaren, 1983) occur in parallel via response strategy sampling but can be facilitated via an insight process in which if subjects recognize the appropriate response strategy and task structure, they can bootstrap their learning via a mixture of generalization and inference. We refer to this model as the response-sampling insight (RSI) model. The RSI model is intended as a model of the pre-solution period; otherwise, we include the basic assumptions of signal detection when processing incoming sensory input and reinforcement learning when updating response-outcome values. Our goal is not to adjudicate between the RSI and alternatives, but to show that the RSI is indeed a tenable alternative. Thus, what follows is an initial formalization of the RSI model, including proof-of-concept fits and simulations to illustrate its viability.

A Sketch of the Response-Sampling Insight (RSI) Model

Signal Detection Module

The proposed model is composed of four different modules: a signal detection module, a learning and choice module, a response-sampling module, and an insight module. The signal detection module allows the RSI to determine whether the current trial is a Tone or No-Tone trial. This module assumes that auditory stimuli like the white-noise and tone used in the present study are best represented by gaussian likelihoods. As such, sensory input can be modeled on each trial as the difference in these likelihoods as follows,

$$y_t \sim N(\text{Tone dB}_t - \text{White-noise dB}_t, \sigma)$$

$$s_i \leftarrow \begin{cases} 1, & \text{if } y_t \geq \theta \\ 0, & \text{else} \end{cases} \quad (4)$$

where t indexes trials, i indexes the obtained trial type, y_t indicates the sensory input on the t -th trial, Tone dB and White-noise dB indicate the total dB level of auditory stimuli on that trial, and σ is the standard deviation or degree of noise in sensory input. To determine if the t -th trial is a Tone or No-Tone trial, the sensory input must surpass the sensory threshold θ . If $y_t \geq \theta$, $s_t = 1$ indicating a Tone trial; otherwise, $s_t = 0$ indicating a No-Tone trial; $\sigma > 0$. Note that in Eq. 4, σ and θ may be interpreted as equivalent to inherent (not acquired) discriminability and criterion.

Learning and Choice Module

At the end of each trial, the learning and choice module updates the value of the chosen response (r_j) given the obtained trial (s_i) and outcome (o_k) following a prediction error algorithm expressed in nomenclature typical of associative learning models as follows,

$$\Delta V_t(r_j | s_i) = V_t(o_k | r_j) - V_t(r_j | s_i)$$

$$V_{t+1}(r_j|s_i) \leftarrow V_t(r_j|s_i) + \alpha_{o_k} \Delta V_t(r_j|s_i) \quad (5)$$

where j indexes the response ($r_0 = 0$ = “No” Tone or $r_1 = 1$ = “Yes” Tone), k indexes the outcome ($o_0 = 0$ = no-reward, $o_1 = 1$ = reward); $V_t(o_k|r_j)$ is the value of the obtained outcome given the chosen response, $V_t(r_j|s_i)$ is the value of the chosen response given the obtained trial; and α is a learning rate indexed by o_k indicating different learning rates when $o_k = 0$ and $o_k = 1$. Note, that increments and decrements in the value of a response given the obtained trial is a proxy for acquired discriminability.

It is assumed that when the auditory stimuli are terminated and levers extended, the RSI chooses the response with the greatest value with some noise given the obtained trial as follows,

$$p_t(r_1|s_i) = \frac{1}{1 + e^{\tau[V_t(r_1|s_i) - V_t(r_0|s_i)]}} \quad (6)$$

which is a softmax function of $V_t(r_1|s_i) - V_t(r_0|s_i)$ with inverse noise parameter τ ; $\tau \geq 0$. As τ approaches infinity, choice becomes noiseless; as τ approaches 0, choice is drowned in noise. Note that $p_t(r_0|s_i)$ is the compliment of $p_t(r_1|s_i)$.

Modifying Acquisition with Response Sampling and Insight

Eq. 4-6 ensures a relatively optimal solution to a conditional discrimination task that can capture asymptotic performance without capturing acquisition trajectories, including the sensitivity of those trajectories to prior experience. To capture these effects, the RSI model assumes a response-sampling module in which different response strategies and task structures are tested during acquisition, and an insight module in which after some threshold is reached, learning is bootstrapped through weighted aggregation of conditional response values across all possible task structures.

Response-Sampling Module

Prior work suggests that subjects fluctuate in and out of task engagement and/or test different response strategies even after reaching asymptotic performance (Ashwood et al., 2020; Brackney et al., 2011; Daniels et al., 2018; Gallistel et al., 2004; Killeen et al., 2009). Similarly, in the response sampling module, it is assumed that at the start of each trial, the RSI samples one of 4 possible response strategies to test on each trial as follows,

$$Q = (q_{l=0}=\text{Random}, q_{l=1}=\text{Biased "No"}, q_{l=2}=\text{Biased "Yes"}, q_{l=3}=\text{Conditional})$$

$$z_t \sim P_t Q \tag{7}$$

where l indexes the response strategies represented as $q_0 \dots q_3$ in vector Q , P_t represents the vector of weights on t -th trial indicating the probability with which each response strategy is sampled, and z_t is the response strategy sampled on the t -th trial.

In contrast to previous work, each response strategy is assumed to correspond to a hypothesis about the structure of the task. Figure 9A shows the task structures associated with each response strategy. The task structure corresponding to the correct response strategy, the conditional responding strategy, is a 2-by-2 matrix row indexed by s_i and column indexed by r_i . The other task structures are subsets of this structure. For example, Biased “Yes” and Biased “No” responding corresponds to a task structure composed of a single column indexed by s_i but with r_i fixed. This implies that despite being biased subjects still learn about the value of that particular response given the sensory input. In contrast, Random Responding corresponds to a task structure composed of a single row indexed by r_i but not s_i . This implies that when random responding, subjects learn the value of each response regardless of the sensory input. These assumptions imply the following modification of Eq. 5,

$$\Delta V_t(r_j|s_i, z_t) = V_t(o_k|r_j, z_t) - V_t(r_j|s_i, z_t)$$

$$V_{t+1}(r_j|s_i, z_t) \leftarrow V_t(r_j|s_i, z_t) + \alpha(o_k) \Delta V_t(r_j|s_i, z_t) \quad (8)$$

And Eq. 6,

$$p_t(r_1|s_i, z_t) = \frac{1}{1 + e^{\tau[V_t(r_1|s_i, z_t) - V_t(r_0|s_i, z_t)]}} \quad (9)$$

where the value of responses and outcomes, and the probability of choosing to indicate the presence of a tone now depends on the obtained trial or chosen response *and* the sampled response strategy and thus task structure. If z_t is equal to q_1 then $p_t(r_1|s_i, z_t=q_1) = 0.5$; if z_t is equal to q_2 then $p_t(r_1|s_i, z_t=q_2) = 0 + \gamma$ and if z_t is equal to q_3 then $p_t(r_1|s_i, z_t=q_3) = 1 - \gamma$; γ is a smudge factor that allows for some noise in choice when subjects are biased; $\gamma \leq 0.3$. Note that when the model chooses a biased response strategy and corresponding task structure but makes the alternative response the model does not update the value of that response. This assumption is tantamount to assuming the subject is completely disengaged from the task (as in Ashwood et al., 2020; Daniels & Sanabria, 2017a).

Figure 9. Schematic of the Response Sampling and Insight Modules of the Response-Sampling Insight (RSI) Model. Panel A shows each response strategy and its corresponding task structure in the form of a vector or matrix. Response strategies are indexed by q_i ; responses are indexed by r_i ; and trial types are indexed by s_i . Panel B shows how at the moment of insight (π), the model generalizes value from q_1 to q_3 with the same s_t and/or r_t as q_3 , and infers value from q_1 to q_3 with orthogonal s_t and r_t to q_3 .

The vector of weights P_t are not expected to be static. Feedback likely allows the probability with which each response strategy is chosen to evolve with training. This evolution probably converges on the conditional discrimination strategy and task structure, with low probabilities for the other strategies and task structures. However, the function describing how such probabilities evolve remains unclear. To get the RSI model running, we empirically estimated the weights by running our two-stage algorithm at a trial block level rather than session level (more details are provided in our viability test of the RSI).

Insight Module

In the insight module, it is assumed that once the RSI reaches some threshold (π) it makes a one-time/in the moment generalization and inference across task structures and associated stimulus-response-outcome mappings and values to update the Conditional Responding task structure (for functionally equivalent theorizing see Frey & Sears, 1979). Figure 9B visually represents both generalization and inference at the moment of insight. Generalization across task structures only occurs for like cells sharing the same s_t and/or r_i . For example, a mouse generalizing value to $V(r_0|s_0, q_3)$ will transfer a proportion of the value from $V(r_0|q_0)$ and $V(r_0|s_0, q_1)$. In contrast, inference across task structures occurs only for diagonal cells with orthogonal s_t and r_i . For example, a mouse inferring value to $V(r_0|s_0, q_3)$ will transfer some proportion of value from $V(r_1|s_1, q_2)$. Taken together, these processes are represented as follows for each cell/combination of responses and trial types of the Conditional Responding task structure

$$\begin{aligned}
 V_{t+1}(r_0|s_0, q_3) &\leftarrow V_t(r_0|s_0, q_3) + \eta V_t(r_0|q_0) + \eta V_t(r_0|s_0, q_1) + \psi V_t(r_1|s_1, q_2) \\
 V_{t+1}(r_0|s_1, q_3) &\leftarrow V_t(r_0|s_1, q_3) + \eta V_t(r_0|q_0) + \eta V_t(r_0|s_1, q_1) + \psi V_t(r_1|s_0, q_2) \\
 V_{t+1}(r_1|s_0, q_3) &\leftarrow V_t(r_1|s_0, q_3) + \eta V_t(r_1|q_0) + \eta V_t(r_1|s_0, q_2) + \psi V_t(r_0|s_1, q_1) \\
 V_{t+1}(r_1|s_1, q_3) &\leftarrow V_t(r_1|s_1, q_3) + \eta V_t(r_1|q_0) + \eta V_t(r_1|s_1, q_2) + \psi V_t(r_0|s_0, q_1)
 \end{aligned} \tag{10}$$

where parameters are as before, η is the generalization coefficient, and ψ is the inference coefficient; $0 \leq \eta, \psi \leq 1$. Note that π is an estimate of the pre-solution period similar to pre-solution estimates from Eq. 3 and the two-stage algorithm. At this time, however, it is unclear what processes (not including prior experience) set the value of π .

Testing the RSI Model

Eq. 4 and 7-10 comprise the sketch of the RSI model of conditional discrimination acquisition; parameters, their interpretation, and estimates for the model selected for each representative mouse can be found in Table 2. Although there are missing components (i.e., the feedback function affecting the vector of weights of response-strategies and thus task structures; and the algorithm that sets the insight threshold), it is still possible to test the viability of the RSI model. To do this, we focused on the data of our representative mice.

Table 2. RSI Model Parameters, Interpretations, and Estimates

Module	Parameter	Description	Estimates		
		Interpretation (unit)	Naïve (25)	CCRT (26)	CDT (66)
Signal	σ	Discriminability (dB)		1	
Detection	θ	Criterion (dB)		3	
Learning & Choice	$\alpha(o_0)$	Learning rate when $r = 0$ (unitless)	.008	.004	.006
	$\alpha(o_1)$	Learning rate when $r = 1$ (unitless)	.004	.01	.004
	τ	Noise in conditional choice (unitless)	7	10	7
Response Sampling	P	Vector of weights associated with response strategies and task structures (unitless)	Empirically estimated		
	γ	Probability of non-biased response (unitless)	0.1	0.25	0.1
Insight	π	Insight threshold/pre-solution period length (session)	17	-	7
	η	Generalization coefficient (unitless)	.14	-	.15
	ψ	Inference coefficient (unitless)	.34	-	.37

Note. Although P is estimated empirically, we treat P as representing four free parameters, one for the weight of each response strategy and corresponding task structure. Thus, the full RSI has 13 free parameters, the RSI without insight has 10 free parameters, and the RSI without insight or response-sampling has 5 free parameters

First, we estimated the empirical probability with which each response strategy was sampled by having our two-stage algorithm parse the data over blocks of trials rather than sessions. Sessions were divided into 10 trial blocks, giving approximately 10-12 blocks per session. The session-by-session weights were calculated as the number of blocks corresponding to each strategy divided by the total number of blocks within that session. With these weights

calculated, we simulated acquisition of the representative Naïve subject by running the RSI 100 times with parameters reasonably similar to those published by other researchers (Daniels & Sanabria, 2018; Luksys et al., 2009; Marshall & Kirkpatrick, 2017). From these 100 runs, we calculated the probability of making a hit and false alarm as a function of session, which allowed us to calculate the residual sum of squares (RSS) from the log odds of observed hits and false alarms minus the log odds of simulated hits and false alarms. We adjusted parameters until the RSS did not appreciably improve. We then moved onto the CDT subject and the CCRT subject, aiming to keep parameters constant across subjects and only revisiting parameter estimates if there was a clear misfit. After settling on a set of parameters, we simulated the RSI 1000 times under three conditions (all modules, no insight, no response-sampling or insight module). Once all of these simulations were done, we again calculated the RSS and converted into Akaike information criterion (AIC). We then selected the model with a $\Delta\text{AIC} \geq 4$ (Anderson et al., 1998; for a practical example see Daniels & Sanabria, 2018) that was robust to removal of outlier residuals; in particular, outlier residuals coming from asymptotic performance in which model identifiability is more ambiguous. Note that a model with a $\Delta\text{AIC} \geq 4$ is at least e^2 times more likely than other competing models¹.

¹ We chose to not estimate parameters using maximum likelihood or Bayesian approaches precisely because of the two missing components: feedback function for updating P as a function of trial and the processes that inform π . Without knowing how to specify these components, misfits could be due to the inadequacy of our approximations rather than some in-principal failure of the model. Nevertheless, we took a principled approach in which we tried to minimize RSS by slowly adjusting parameters of each module and then asking the relative import of each module.

Figure 10. Proof-of-concept Fits of the Response-Sampling Insight Model (RSI) of Conditional Discriminations to Data of Representative Mice. $P(\text{Hit} = \text{Blue})$ and $P(\text{False Alarm} = \text{Orange})$ as a function of session (obtained data = colored circles, model fits = lines) for representative mice from each training group (Naïve= darkest color; CCRT= middle color; CDT = lightest color). Three fits are depicted. RSI without insight or response-sampling (dotted lines), RSI without response-sampling (dashed lines), and the full RSI (solid line).

As shown in Figure 10, whereas the full RSI model and the RSI model without insight captures acquisition trajectories, the RSI model without both the insight and response-sampling modules fails to capture acquisition trajectories. Inclusion of insight, however, improved the ability of the RSI to capture the acquisition trajectories of Naïve and CDT but not CCRT mice. In support of this visual impression, AIC indicated that in Naïve and CDT the full RSI was at least $e^{6.98}$ (i.e., $\Delta\text{AIC} = 13.96$) times more likely than the RSI without insight and at least e^{135} more times likely than the RSI without both insight and response-sampling. This remained true even when removing a single session during asymptotic performance for each mouse (when the models with and without the insight modules are functionally equivalent) with undue influence on AIC (at least $\Delta\text{AIC} = 8.12$ for RSI without insight; $\Delta\text{AIC} = 277$ for RSI without both insight and response-sampling). In contrast, although AIC selected the full RSI for the CCRT mouse, this was largely the result of the undue influence of hits and false alarms from session 23 ($\Delta\text{AIC} = 18.91$ for RSI without insight when session 23 is included; $\Delta\text{AIC} = 2.54$ for RSI without insight when session 23 is excluded). When that session was removed, the full RSI was only $e^{1.27}$ times more likely than the RSI without insight. Naïve and CDT but not CCRT mice appear to recruit the insight process. Thus, our proof-of-concept fits show that the machinery of the RSI can at least in principle account for acquisition trajectories, while also identifying potential differences between our three groups of mice.

Specifically, model selection indicates that whereas all mice may recruit the response-sampling module as hypothesized by Heinemann, only Naïve and CDT mice may recruit the insight module hypothesized in the present paper. An implication of this is that whereas acquisition for some mice may be characterized by two periods, for other mice it may be characterized by one period. Although two periods appear to be the norm, inspection of

individual data suggests one period learning is particularly dominant in CCRT (cf. Figure 4B and 10B with Supplemental Figure 3) followed by a couple of mice in CDT and almost no mice in Naïve. Such a conclusion could not have been reached with traditional analyses or modeling. Of course, this raises the question of how does prior experience modulate whether mice recruit the insight module?

We suspect that prior experience affects whether and how quickly mice can effectively deploy an inference to bootstrap learning. Specifically, if mice enter training *without* knowledge of task-specific stimulus-response-outcome mapping and values (i.e., Naïve and CDT) then those mice are more likely to deploy a biased response-sampling strategy from which they can bootstrap their learning. For example, a mouse may learn a subset of the stimulus-response-outcome mapping and values: Tone $\rightarrow$ Left Lever $\rightarrow$ Reward and No-Tone $\rightarrow$ Left Lever $\rightarrow$ No-Reward. This subset can be leveraged to make the inference that Tone $\rightarrow$ Right Lever $\rightarrow$ No-Reward and No-Tone $\rightarrow$ Right Lever $\rightarrow$ Reward. CDT mice make the inference sooner than Naïve mice because they are already aware that stimulus-response-outcome mapping and values consistent with a conditional discrimination strategy are possible. In contrast, if mice enter training *with* knowledge of task-specific stimulus-response-outcome mapping and values, whether they can make an inference depends on the content of that prior training. Mice whose prior training content is biased towards the left or right lever will be able to make an inference much like Naïve and CDT mice. However, mice whose prior training content is focused on correct (as in CCRT mice; or incorrect choices) cannot make an inference. This is because they learned the following subset: Tone $\rightarrow$ Left Lever $\rightarrow$ Reward and No-Tone $\rightarrow$ Right Lever $\rightarrow$ Reward. Because in this subset there is no experience with No-Reward, it is not possible to infer the incorrect stimulus-response-outcome mapping and values; likewise, if subjects were given

prior experience with only incorrect responses, then they could not infer the correct responses. Although this hypothesis needs to be tested explicitly in future experience, the RSI model suggests that whether the insight module is recruited depends in part on the content of prior experience.

Although parameter estimates at this point are just suggestions, inspection of the parameters selected for our representative mice (see Table 2) indicate correspondence with the effects reported from our analysis of discriminability. Estimates of the pre-solution period from the discriminability analysis (L , Eq. 3) indicated that prior experience shortens the pre-solution period. Likewise, the estimate of the insight threshold (π) was more than twice as long for the representative Naïve mouse compared to the representative CDT mouse. Although not shown in Table 2 because the insight module was not required for the representative CCRT mouse, the insight threshold for the CCRT mouse was also substantially shorter ($\pi = 4$) than the insight threshold of the Naïve mouse.

Additionally, estimates of the rate of transition to asymptotic performance in the solution period from the discriminability analysis (s , Eq. 3) indicated that biased mice (Naïve and CDT) learn faster than less biased mice once they exit the pre-solution period (CCRT). Parameter s could be related to the inference and generalization coefficients of the insight module. However, this comparison is made difficult by the fact that the representative CCRT mouse did not require the insight module; even when the insight module was included, the CCRT mouse had similar estimates for the inference (ψ) and generalization (η) coefficients compared to the Naïve and CDT mouse. Alternatively, s may be related to learning rates when animals make incorrect choices; that is when the outcome is no-reward rather than reward. Estimates of α_{o0} scaled such that Naïve $\approx$ CDT $>$ CCRT, which is similar to the scaling observed for estimates of s . This

could indicate that the abrupt transitions reflect sensitivity to incorrect choices rather than learning in general.

Limitations

The present work is characterized by two distinct limitations. The first has already been acknowledged: there are two unspecified components to our sketch of the RSI model. Despite this limitation, it is worth highlighting that both the response-strategy sampling and insight modules are supported by previous research. Whereas response-strategy sampling has been relatively well documented (Heinemann, 1983; Leeuwen et al., 2013; Reader, 2015; Santos & Sanabria, 2020), insight has, until recently, long been studied independent of conditional discriminations (for a review see Shettleworth, 2012). Nonetheless, the assumption that biased responding affords an insight process composed of both inference and generalization is supported by the observation that non-human animals are capable of some form of inferential processing (Aust et al., 2008; Daniels et al., 2014; DeVito et al., 2010; Jensen et al., 2019; Kornell, 2009; Pineño & Miller, 2007). It is also consistent with data indicating that inferential reasoning can facilitate (McCormack et al., 2013) or at a minimum may be involved in learning (Guez & Stevenson, 2011; Killeen, 1981; Pineño & Miller, 2007). Thus, the RSI is theoretically sound; and as demonstrated a viable alternative to similar, potentially functionally equivalent models (Bathellier et al., 2013; Frey & Sears, 1979).

Second, a limitation of the present data is that CCRT mice differed from Naïve and CDT mice on more than one dimension. In addition to differing in prior experience, their choice trials were gradually rather than abruptly introduced. The gradual introduction of choice trials could explain both why CCRT mice did not respond as randomly during the pre-solution period as expected by our model and why CCRT mice had lower rates of acquisition. Specifically, gradual

introduction of choice trials likely reduced the predicted interference effect of no-reward trials, thus reducing recourse to random responding (Bouton, 2019). Additionally, gradual introduction of choice trials may have reduced the salience of the transition from training with correct choices to choice trials, which may have in turn kept the learning rate low or prevented mice from realizing the structure of the environment had changed. Previous work indicates that learning rates or inference of a new environmental structure may increase following more easily detected, that is cued or abrupt, transitions (Bouton, 2004; Gershman et al., 2013; Redish et al., 2007). Note, it is important to keep in mind that these two potential effects of gradually introducing choice trials counteract each other, meaning this limitation can only be resolved by future research.

Conclusion

In sum, the RSI model as offered here is a sketch; more data is required to work out the missing components. Nonetheless, the major findings of the present study are that during the pre-solution period, mice selectively sample biased response strategies; that the degree and direction of bias is affected by prior experience; and that while prior experience generally shortens the pre-solution period, only when prior experience enhances bias during the pre-solution period does it result in a facilitative effect on acquisition. These findings were reflected in common signal detection parameters of discriminability and criterion. Discriminability was sensitive to feedback and prior experience, reflecting both the shortening of the pre-solution period and the facilitative effect of biased responding on acquisition; criterion was insensitive to feedback but sensitive to prior experience, reflecting the biases developed during the pre-solution period. These findings are inconsistent with common instantiations of signal detection theory but as was shown can be accommodated by learning theories if they are expanded to include response-strategy sampling

and a moment of insight (i.e., the RSI model). Future research will focus on testing the core assumptions of the RSI model.

References

- Alves-Pinto, A., Sollini, J., & Sumner, C. J. (2012). Signal detection in animal psychoacoustics: Analysis and simulation of sensory and decision-related influences. *Neuroscience*, 220, 215–227. <https://doi.org/10.1016/j.neuroscience.2012.06.001>
- Anderson, D. R., Burnham, K. P., & White, G. C. (1998). Comparison of Akaike information criterion and consistent Akaike information criterion for model selection and statistical inference from capture-recapture studies. *Journal of Applied Statistics*, 25(2), 263–282. <https://doi.org/10.1080/02664769823250>
- Andrzejewski, M. E., & Holder, N. (2016). Chapter 11—Amphetamine and Signal Detection. In V. R. Preedy (Ed.), *Neuropathology of Drug Addictions and Substance Misuse* (pp. 117–125). Academic Press. <https://doi.org/10.1016/B978-0-12-800212-4.00011-X>
- Ashwood, Z. C., Roy, N. A., Stone, I. R., The International Brain Laboratory, Churchland, A. K., Pouget, A., & Pillow, J. W. (2020). *Mice alternate between discrete strategies during perceptual decision-making* [Preprint]. *Neuroscience*. <https://doi.org/10.1101/2020.10.19.346353>
- Aust, U., Range, F., Steurer, M., & Huber, L. (2008). Inferential reasoning by exclusion in pigeons, dogs, and humans. *Animal Cognition*, 11(4), 587–597. <https://doi.org/10.1007/s10071-008-0149-0>
- Balci, F., Freestone, D., Simen, P., deSouza, L., Cohen, J. D., & Holmes, P. (2011). Optimal Temporal Risk Assessment. *Frontiers in Integrative Neuroscience*, 5. <https://doi.org/10.3389/fnint.2011.00056>
- Ball, K., & Sekuler, R. (1982). A Specific and Enduring Improvement in Visual Motion Discrimination. *Science*. <https://www.science.org/doi/abs/10.1126/science.7134968>
- Bathellier, B., Tee, S. P., Hrovat, C., & Rumpel, S. (2013). A multiplicative reinforcement learning model capturing learning dynamics and interindividual variability in mice. *Proceedings of the National Academy of Sciences*, 110(49), 19950–19955. <https://doi.org/10.1073/pnas.1312125110>
- Benovitski, Y. B., Blamey, P. J., Rathbone, G. D., & Fallon, J. B. (2014). An automated psychoacoustic testing apparatus for use in cats. *Hearing Research*, 309, 1–7. <https://doi.org/10.1016/j.heares.2013.11.002>
- Bieszczad, K. M., & Weinberger, N. M. (2010). Learning strategy trumps motivational level in determining learning-induced auditory cortical plasticity. *Neurobiology of Learning and Memory*, 93(2), 229–239. <https://doi.org/10.1016/j.nlm.2009.10.003>
- Bouton, M. E. (2004). Context and Behavioral Processes in Extinction. *Learning & Memory*, 11(5), 485–494. <https://doi.org/10.1101/lm.78804>
- Bouton, M. E. (2019). Extinction of instrumental (operant) learning: Interference, varieties of context, and mechanisms of contextual control. *Psychopharmacology*, 236(1), 7–19. <https://doi.org/10.1007/s00213-018-5076-4>

- Brackney, R. J., Cheung, T. H. C., Neisewander, J. L., & Sanabria, F. (2011). The Isolation of Motivational, Motoric, and Schedule Effects on Operant Performance: A Modeling Approach. *Journal of the Experimental Analysis of Behavior*, 96(1), 17–38. <https://doi.org/10.1901/jeab.2011.96-17>
- Brown, S. W., Newcomb, D. C., & Kahrl, K. G. (1995). Temporal-Signal Detection and Individual Differences in Timing. *Perception*, 24(5), 525–538. <https://doi.org/10.1068/p240525>
- Carandini, M., & Churchland, A. K. (2013). Probing perceptual decisions in rodents. *Nature Neuroscience*, 16(7), 824–831. <https://doi.org/10.1038/nn.3410>
- Chen, L., Wang, X., Ge, S., & Xiong, Q. (2019). Medial geniculate body and primary auditory cortex differentially contribute to striatal sound representations. *Nature Communications*, 10(1), 1–10. <https://doi.org/10.1038/s41467-019-08350-7>
- D'Amato, M. R., & Jagoda, H. (1961). Analysis of the role of overlearning in discrimination reversal. *Journal of Experimental Psychology*, 61(1), 45–50. <https://doi.org/10.1037/h0047757>
- Daniels, C. W., Laude, J. R., & Zentall, T. R. (2014). Six-term transitive inference with pigeons: Successive-pair training followed by mixed-pair training. *Journal of the Experimental Analysis of Behavior*, 101(1), 26–37. <https://doi.org/10.1002/jeab.65>
- Daniels, C. W., Overby, P. F., & Sanabria, F. (2018). Between-session memory degradation accounts for within-session changes in fixed-interval performance. *Behavioural Processes*, 153, 31–39. <https://doi.org/10.1016/j.beproc.2018.05.004>
- Daniels, C. W., & Sanabria, F. (2017a). Interval timing under a behavioral microscope: Dissociating motivational and timing processes in fixed-interval performance. *Learning & Behavior*, 45(1), 29–48. <https://doi.org/10.3758/s13420-016-0234-1>
- Daniels, C. W., & Sanabria, F. (2017b). About bouts: A heterogeneous tandem schedule of reinforcement reveals dissociable components of operant behavior in Fischer rats. *Journal of Experimental Psychology: Animal Learning and Cognition*, 43(3), 280–294. <https://doi.org/10.1037/xan0000144>
- Daniels, C. W., & Sanabria, F. (2018). An associability decay model of paradoxical choice. *Journal of Experimental Psychology: Animal Learning and Cognition*, 44(3), 258–271. <https://doi.org/10.1037/xan0000179>
- David, S. V., Fritz, J. B., & Shamma, S. A. (2012). Task reward structure shapes rapid receptive field plasticity in auditory cortex. *Proceedings of the National Academy of Sciences*, 109(6), 2144–2149. <https://doi.org/10.1073/pnas.1117717109>
- Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience*, 8(4), 429–453. <https://doi.org/10.3758/CABN.8.4.429>
- DeVito, L. M., Lykken, C., Kanter, B. R., & Eichenbaum, H. (2010). Prefrontal cortex: Role in acquisition of overlapping associations and transitive inference. *Learning & Memory*, 17(3), 161–167. <https://doi.org/10.1101/lm.1685710>

- Eck, K. O., & Thomas, D. R. (1970). Discrimination learning as a function of prior discrimination and nondifferential training: A replication. *Journal of Experimental Psychology*, 83(3, Pt.1), 511–513. <https://doi.org/10.1037/h0028874>
- Fay, R. R. (1988). Comparative psychoacoustics. *Hearing Research*, 34(3), 295–305. [https://doi.org/10.1016/0378-5955\(88\)90009-3](https://doi.org/10.1016/0378-5955(88)90009-3)
- Frey, P. W., & Sears, R. J. (1979). Model of conditioning incorporating the Rescorla-Wagner associative axiom, a dynamic attention process, and a catastrophe rule. *Psychological Review*, 85(4), 321. <https://doi.org/10.1037/0033-295X.85.4.321>
- Gallistel, C. R., Fairhurst, S., & Balsam, P. (2004). The learning curve: Implications of a quantitative analysis. *Proceedings of the National Academy of Sciences*, 101(36), 13124–13131. <https://doi.org/10.1073/pnas.0404965101>
- Gallun, F. J., & Kampel, S. D. (2016). Using detection theory to analyze human decision making and sensory processing. *Proceedings of Meetings on Acoustics*, 26(1), 050004. <https://doi.org/10.1121/2.0000379>
- Gershman, S. J. (2015). A Unifying Probabilistic View of Associative Learning. *PLOS Computational Biology*, 11(11), e1004567. <https://doi.org/10.1371/journal.pcbi.1004567>
- Gershman, S. J., Jones, C. E., Norman, K. A., Monfils, M.-H., & Niv, Y. (2013). Gradual extinction prevents the return of fear: Implications for the discovery of state. *Frontiers in Behavioral Neuroscience*, 7. <https://doi.org/10.3389/fnbeh.2013.00164>
- Gold, J. M., Sekuler, A. B., & Bennett, P. J. (2004). Characterizing perceptual learning with external noise. *Cognitive Science*, 28(2), 167–207. https://doi.org/10.1207/s15516709cog2802_3
- Green, D. M. (1960). Psychoacoustics and Detection Theory. *The Journal of the Acoustical Society of America*, 32(10), 1189–1203. <https://doi.org/10.1121/1.1907882>
- Guez, D., & Stevenson, G. (2011). Is reasoning in rats really unreasonable? Revisiting recent associative accounts. *Frontiers in Psychology*, 2, 277. <https://doi.org/10.3389/fpsyg.2011.00277>
- Hangya, B., Ranade, S. P., Lorenc, M., & Kepecs, A. (2015). Central Cholinergic Neurons Are Rapidly Recruited by Reinforcement Feedback. *Cell*, 162(5), 1155–1168. <https://doi.org/10.1016/j.cell.2015.07.057>
- Hanks, T. D., & Summerfield, C. (2017). Perceptual Decision Making in Rodents, Monkeys, and Humans. *Neuron*, 93(1), 15–31. <https://doi.org/10.1016/j.neuron.2016.12.003>
- Heffner, H. E., & Heffner, R. S. (2007). Hearing Ranges of Laboratory Animals. *Journal of the American Association for Laboratory Animal Science*, 3.
- Heinemann, E. G. (1983). The Presolution Period and the Detection of Statistical Associations. In *Quantitative Analyses of Behavior: Discrimination Processes* (Vol. 4, pp. 21–36). Ballinger. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.536.1978&rep=rep1&type=pdf>
- Heinemann, E. G., & Chase, S. (1970). Conditional stimulus control. *Journal of Experimental Psychology*, 84(2), 187. <https://doi.org/10.1037/h0029152>

- Horga, G., & Abi-Dargham, A. (2019). An integrative framework for perceptual disturbances in psychosis. *Nature Reviews Neuroscience*, 20(12), 763–778. <https://doi.org/10.1038/s41583-019-0234-1>
- Jaramillo, S., & Zador, A. M. (2014). Mice and rats achieve similar levels of performance in an adaptive decision-making task. *Frontiers in Systems Neuroscience*, 8. <https://doi.org/10.3389/fnsys.2014.00173>
- Jarosz, A. F., & Wiley, J. (2014). What Are the Odds? A Practical Guide to Computing and Reporting Bayes Factors. *The Journal of Problem Solving*, 7(1). <https://doi.org/10.7771/1932-6246.1167>
- JASP Team. (2019). *JASP (Version 0.11.1)[Computer software]*. <https://jasp-stats.org/>
- Jensen, G., Alkan, Y., Ferrera, V. P., & Terrace, H. S. (2019). Reward associations do not explain transitive inference performance in monkeys. *Science Advances*, 5(7), eaaw2089. <https://doi.org/10.1126/sciadv.aaw2089>
- Jones, M., & Canas, F. (2010). Integrating Reinforcement Learning with Models of Representation Learning. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 32(32). <https://escholarship.org/uc/item/88x6f84q>
- Karim, M., Harris, J. A., Morley, J. W., & Breakspear, M. (2012). Prior and Present Evidence: How Prior Experience Interacts with Present Information in a Perceptual Decision Making Task. *PLOS ONE*, 7(5), e37580. <https://doi.org/10.1371/journal.pone.0037580>
- Killeen, P. (1981). Quantitative analysis of behavior. In *Learning as causal inference* (Vol. 1, pp. 321–340). Ballinger. https://www.academia.edu/21780437/Learning_as_causal_inference
- Killeen, P. R. (2016). Signal Detection Theory. In *The SAGE Encyclopedia of Theory in Psychology*. SAGE Publications, Inc. <https://doi.org/10.4135/9781483346274.n292>
- Killeen, P. R., Sanabria, F., & Dolgov, I. (2009). The dynamics of conditioning and extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, 35(4), 447–472. <https://doi.org/10.1037/a0015626>
- Killeen, P. R., Taylor, T. J., & Treviño, M. (2018). Subjects adjust criterion on errors in perceptual decision tasks. *Psychological Review*, 125(1), 117–130. <https://doi.org/10.1037/rev0000056>
- Kim, Y.-H., Kang, D.-W., Kim, D., Kim, H.-J., Sasaki, Y., & Watanabe, T. (2015). Real-Time Strategy Video Game Experience and Visual Perceptual Learning. *Journal of Neuroscience*, 35(29), 10485–10492. <https://doi.org/10.1523/JNEUROSCI.3340-14.2015>
- Klink, K. B., Bendig, G., & Klump, G. M. (2006). Operant methods for mouse psychoacoustics. *Behavior Research Methods*, 38(1), 1–7. <https://doi.org/10.3758/BF03192744>
- Kornell, N. (2009). Metacognition in Humans and Animals. *Current Directions in Psychological Science*, 18(1), 11–15. <https://doi.org/10.1111/j.1467-8721.2009.01597.x>
- Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017). Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron*, 93(3), 480–490. <https://doi.org/10.1016/j.neuron.2016.12.041>

- Krechevsky, I. (1932). "Hypotheses" in rats. *Psychological Review*, 39(6), 516–532. <https://doi.org/10.1037/h0073500>
- Kuchibhotla, K. V., Sten, T. H., Papadoyannis, E. S., Elnozahy, S., Fogelson, K. A., Kumar, R., Boubenec, Y., Holland, P. C., Ostojic, S., & Froemke, R. C. (2019). Dissociating task acquisition from expression during learning reveals latent knowledge. *Nature Communications*, 10(1), 1–13. <https://doi.org/10.1038/s41467-019-10089-0>
- Kurt, S., & Ehret, G. (2010). Auditory discrimination learning and knowledge transfer in mice depends on task difficulty. *Proceedings of the National Academy of Sciences*, 107(18), 8481–8485. <https://doi.org/10.1073/pnas.0912357107>
- Kurylo, D. D., Waxman, R., Kidron, R., & Silverstein, S. M. (2017). Visual training improves perceptual grouping based on basic stimulus features. *Attention, Perception, & Psychophysics*, 79(7), 2098–2107. <https://doi.org/10.3758/s13414-017-1368-8>
- Lany, J., Gómez, R. L., & Gerken, L. A. (2007). The Role of Prior Experience in Language Acquisition. *Cognitive Science*, 31(3), 481–507. <https://doi.org/10.1080/15326900701326584>
- Lawrence, D. H. (1950a). Acquired distinctiveness of cues: I. Transfer between discriminations on the basis of familiarity with the stimulus. *Journal of Experimental Psychology*, 39(6), 770. <https://doi.org/10.1037/h0058097>
- Lawrence, D. H. (1950b). Acquired distinctiveness of cues: II. Selective association in a constant stimulus situation. *Journal of Experimental Psychology*, 40(2), 175. <https://doi.org/10.1037/h0063217>
- Leeuwen, E. J. C. V., Cronin, K. A., Schütte, S., Call, J., & Haun, D. B. M. (2013). Chimpanzees (Pan troglodytes) Flexibly Adjust Their Behaviour in Order to Maximize Payoffs, Not to Conform to Majorities. *PLOS ONE*, 8(11), e80945. <https://doi.org/10.1371/journal.pone.0080945>
- Limpert, E., Stahel, W. A., & Abbt, M. (2001). Log-normal Distributions across the Sciences: Keys and Clues. *BioScience*, 51(5), 341. [https://doi.org/10.1641/0006-3568\(2001\)051\[0341:LNDATS\]2.0.CO;2](https://doi.org/10.1641/0006-3568(2001)051[0341:LNDATS]2.0.CO;2)
- Love, J., Selker, R., Marsman, M., Jamil, T., Dropmann, D., Verhagen, J., Ly, A., Gronau, Q. F., Smíra, M., Epskamp, S., Matzke, D., Wild, A., Knight, P., Rouder, J. N., Morey, R. D., & Wagenmakers, E.-J. (2019). **JASP**: Graphical Statistical Software for Common Statistical Designs. *Journal of Statistical Software*, 88(2). <https://doi.org/10.18637/jss.v088.i02>
- Luksys, G., Gerstner, W., & Sandi, C. (2009). Stress, genotype and norepinephrine in the prediction of mouse behavior using reinforcement learning. *NATURE NEUROSCIENCE*, 12(9), 10.
- MacMillan, N. A. (2002). Signal Detection Theory. In *Stevens' Handbook of Experimental Psychology*. American Cancer Society. <https://doi.org/10.1002/0471214426.pas0402>
- Marshall, A. T., & Kirkpatrick, K. (20170706). Reinforcement learning models of risky choice and the promotion of risk-taking by losses disguised as wins in rats. *Journal of Experimental Psychology: Animal Learning and Cognition*, 43(3), 262. <https://doi.org/10.1037/xan0000141>
- Mccarthy, D., & Davison, M. (1981). Towards a behavioral theory of bias in signal detection. *Perception & Psychophysics*, 29(4), 371–382. <https://doi.org/10.3758/BF03207347>

- McCormack, T., Simms, V., McGourty, J., & Beckers, T. (2013). Encouraging Children to Think Counterfactually Enhances Blocking in a Causal Learning Task. *Quarterly Journal of Experimental Psychology*, 66(10), 1910–1926. <https://doi.org/10.1080/17470218.2013.767847>
- McNamara, J. M., & Trimmer, P. C. (2019). Sequential choices using signal detection theory can reverse classical predictions. *Behavioral Ecology*, 30(1), 16–19. <https://doi.org/10.1093/beheco/ary132>
- Meier, G. W., & McGee, R. K. (1959). A re-evaluation of the effect of early perceptual experience in discrimination performance during adulthood. *Journal of Comparative and Physiological Psychology*, 52(4), 390–395. <https://doi.org/10.1037/h0045596>
- Mill, R. W., Alves-Pinto, A., & Sumner, C. J. (2014). Decision Criterion Dynamics in Animals Performing an Auditory Detection Task. *PLoS ONE*, 9(12), e114076. <https://doi.org/10.1371/journal.pone.0114076>
- Mistry, P. K., Skewes, J., & Lee, M. D. (2018). An Adaptive Signal Detection Model Applied to Perceptual Learning. *CogSci*, 6.
- Mueller, S. T., & Weidemann, C. T. (2008). Decision noise: An explanation for observed violations of signal detection theory. *Psychonomic Bulletin & Review*, 15(3), 465–494. <https://doi.org/10.3758/PBR.15.3.465>
- Niv, Y. (2019). Learning task-state representations. *Nature Neuroscience*, 22(10), 1544–1553. <https://doi.org/10.1038/s41593-019-0470-8>
- Odic, D., Hock, H., & Halberda, J. (2014). Hysteresis affects approximate number discrimination in young children. *Journal of Experimental Psychology: General*, 143(1), 255–265. <https://doi.org/10.1037/a0030825>
- Pineño, O., & Miller, R. R. (2007). Comparing associative, statistical, and inferential reasoning accounts of human contingency learning. *Quarterly Journal of Experimental Psychology*, 60(3), 310–329. <https://doi.org/10.1080/17470210601000680>
- Rabin, M. D. (1988). Experience facilitates olfactory quality discrimination. *Perception & Psychophysics*, 44(6), 532–540. <https://doi.org/10.3758/BF03207487>
- Ratcliff, R., Voskuilen, C., & McKoon, G. (2017). Internal and external sources of variability in perceptual decision-making. *Psychological Review*, 125(1), 33. <https://doi.org/10.1037/rev0000080>
- Reader, S. M. (2015). Causes of Individual Differences in Animal Exploration and Search. *Topics in Cognitive Science*, 7(3), 451–468. <https://doi.org/10.1111/tops.12148>
- Redish, A. D., Jensen, S., Johnson, A., & Kurth-Nelson, Z. (2007). Reconciling reinforcement learning models with behavioral extinction and renewal: Implications for addiction, relapse, and problem gambling. *Psychological Review*, 114(3), 784–805. <https://doi.org/10.1037/0033-295X.114.3.784>
- Ritchie, J. B., & Carlson, T. A. (2016). Neural Decoding and “Inner” Psychophysics: A Distance-to-Bound Approach for Linking Mind, Brain, and Behavior. *Frontiers in Neuroscience*, 10. <https://doi.org/10.3389/fnins.2016.00190>

Roy, N. A., Bak, J. H., Akrami, A., Brody, C. D., & Pillow, J. W. (2021). Extracting the dynamics of behavior in sensory decision-making experiments. *Neuron*, 109(4), 597-610.e6. <https://doi.org/10.1016/j.neuron.2020.12.004>

Saksida, L. M. (1999). Effects of similarity and experience on discrimination learning: A nonassociative connectionist model of perceptual learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 25(3), 308–323. <https://doi.org/10.1037/0097-7403.25.3.308>

Santos, C., & Sanabria, F. (2020). Past outcomes and time flexibly exert joint control over midsession reversal performance in the rat. *Behavioural Processes*, 171, 104028. <https://doi.org/10.1016/j.beproc.2019.104028>

Sewell, D. K., Jach, H. K., Boag, R. J., & Van Heer, C. A. (2019). Combining error-driven models of associative learning with evidence accumulation models of decision-making. *Psychonomic Bulletin & Review*, 26(3), 868–893. <https://doi.org/10.3758/s13423-019-01570-4>

Shepard, K. N., Kilgard, M. P., & Liu, R. C. (2013). Experience-Dependent Plasticity and Auditory Cortex. In Y. E. Cohen, A. N. Popper, & R. R. Fay (Eds.), *Neural Correlates of Auditory Cognition* (Vol. 45, pp. 293–327). Springer New York. https://doi.org/10.1007/978-1-4614-2350-8_10

Shettleworth, S. J. (2012). Do animals have insight, and what is insight anyway? *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale*, 66(4), 217. <https://doi.org/10.1037/a0030674>

St. Claire-Smith, R., & MacLaren, D. (1983). Response preconditioning effects. *Journal of Experimental Psychology: Animal Behavior Processes*, 9(1), 41–48. <https://doi.org/10.1037/0097-7403.9.1.41>

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

Stüttgen, M. C., Yildiz, A., & Güntürkün, O. (2011). Adaptive Criterion Setting in Perceptual Decision Making. *Journal of the Experimental Analysis of Behavior*, 96(2), 155–176. <https://doi.org/10.1901/jeab.2011.96-155>

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning, second edition: An Introduction*. MIT Press.

Teichert, T., & Ferrera, V. (2010). Suboptimal Integration of Reward Magnitude and Prior Reward Likelihood in Categorical Decisions by Monkeys. *Frontiers in Neuroscience*, 4, 186. <https://doi.org/10.3389/fnins.2010.00186>

Trachtenberg, J. T., Chen, B. E., Knott, G. W., Feng, G., Sanes, J. R., Welker, E., & Svoboda, K. (2002). Long-term in vivo imaging of experience-dependent synaptic plasticity in adult cortex. *Nature*, 420(6917), 788–794. <https://doi.org/10.1038/nature01273>

Turner, B. M., Van Zandt, T., & Brown, S. (2011). A dynamic stimulus-driven model of signal detection. *Psychological Review*, 118(4), 583. <https://doi.org/10.1037/a0025191>

Wagenmakers, E.-J., Love, J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Selker, R., Gronau, Q. F., Dropmann, D., Boutin, B., Meerhoff, F., Knight, P., Raj, A., van Kesteren, E.-J., van Doorn, J., Šmíra, M., Epskamp, S., Etz, A., Matzke, D., ... Morey, R. D. (2018). Bayesian inference for psychology. Part II:

Example applications with JASP. *Psychonomic Bulletin & Review*, 25(1), 58–76.
<https://doi.org/10.3758/s13423-017-1323-7>

Wagenmakers, E.-J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., Selker, R., Gronau, Q. F., Šmíra, M., Epskamp, S., Matzke, D., Rouder, J. N., & Morey, R. D. (2018). Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin & Review*, 25(1), 35–57.
<https://doi.org/10.3758/s13423-017-1343-3>

Wickens, T. D. (2002). *Elementary Signal Detection Theory*. Oxford University Press.

Witt, J. K., Taylor, J. E. T., Sugovic, M., & Wixted, J. T. (2015). Signal Detection Measures Cannot Distinguish Perceptual Biases from Response Biases. *Perception*, 44(3), 289–300.
<https://doi.org/10.1068/p7908>

Wixted, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*, 114(1), 152. <https://doi.org/10.1037/0033-295X.114.1.152>

Wixted, J. T. (2019). The forgotten history of signal detection theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 46(2), 201. <https://doi.org/10.1037/xlm0000732>

Xiong, Q., Znamenskiy, P., & Zador, A. M. (2015). Selective corticostriatal plasticity during acquisition of an auditory discrimination task. *Nature*, 521(7552), 348–351. <https://doi.org/10.1038/nature14225>

Yanagihara, S., & Yazaki-Sugiyama, Y. (2016). Auditory experience-dependent cortical circuit shaping for memory formation in bird song learning. *Nature Communications*, 7(1), 11946.
<https://doi.org/10.1038/ncomms11946>

Zeng, F.-G., Fu, Q.-J., & Morse, R. (2000). Human hearing enhanced by noise¹¹Published on the World Wide Web on 23 May 2000. *Brain Research*, 869(1), 251–255. [https://doi.org/10.1016/S0006-8993\(00\)02475-6](https://doi.org/10.1016/S0006-8993(00)02475-6)

Znamenskiy, P., & Zador, A. M. (2013). Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature*, 497(7450), 482–485. <https://doi.org/10.1038/nature12077>