


Caution Influences Avoidance and Approach Behaviors Differently

Ji Zhou, Sebastian Hormigo, Muhammad S. Sajid, and  Manuel A. Castro-Alamancos

Department of Neuroscience, University of Connecticut School of Medicine, Farmington, Connecticut 06001

While conflict between incompatible goals has well-known effects on actions, in many situations the same action may produce harmful or beneficial consequences during different periods in a nonconflicting manner, e.g., crossing the street during a red or green light. To avoid harm, subjects must be cautious to inhibit the action specifically when it is punished, as in passive avoidance, but act when it is beneficial, as in active avoidance or active approach. In mice of both sexes performing a signaled action to avoid harm or obtain reward, we found that addition of a new rule that punishes the action when it occurs unsignaled delays the timing of the signaled action in an apparent sign of increased caution. Caution depended on task signaling, contingency, and reinforcement type. Interestingly, caution became persistent when the signaled action was avoidance motivated by danger but was only transient when it was approach motivated by reward. Although caution is represented by the activity of neurons in the midbrain, it developed independent of frontal cortex or basal ganglia output circuits. These results indicate that caution disrupts actions in different ways depending on the motivational state and may develop from unforeseen brain circuits.

Key words: approach; avoidance; basal ganglia; frontal cortex; midbrain

Significance Statement

Actions, such as crossing the street at a light, can have benefits during one light signal (getting somewhere) but can be harmful during a different signal (being run over). Humans must be cautious to cross the street during the period marked by the appropriate signal. In mice performing a signaled action to avoid harm or obtain reward, we found that addition of a new rule that punishes the action when it occurs unsignaled, delays the timing of the signaled action in an apparent sign of increased caution. Caution became persistent when the signaled action was motivated by danger, but not when it was motivated by reward. Moreover, the development of caution did not depend on prototypical frontal cortex circuits.

Introduction

Animals, including humans, can produce the same action with different goals (Thorndike, 1898; Skinner, 1938; Mowrer, 1960; Dickinson et al., 1980; Mackintosh, 1983; Rescorla, 1987; Sutton and Barto, 1998). For example, humans may cross the street to meet a friend on the other side (active approach) or to avoid a foe on the same side (active avoidance). Conflict arises when an action has incompatible goals, such as when both the friend and foe are on the other side of the street. These conflicts cause well-known effects on responding because of the confluence of Pavlovian and operant learning mechanisms, and related motivational states (Miller, 1944; Mowrer, 1960; Rescorla and Solomon,

1967; Church et al., 1970; Bolles et al., 1980; Blanchard and Blanchard, 1988; Gray and McNaughton, 2000; Dickinson and Balleine, 2002; McNaughton and Corr, 2004; Dayan et al., 2006; Boyd et al., 2011). In other nonconflicting situations, the same action (crossing the street at a light) may be beneficial during a period signaled by a stimulus (green light) but harmful during another period (red light) in a nonoverlapping manner. In these cases, subjects must be cautious to produce the action during the beneficial period (active approach or active avoidance) but inhibit the action during the harmful period (passive avoidance).

Caution, defined as taking care to avoid errors or danger, has been studied in humans and monkeys using psychophysics in the context of speed-accuracy trade-off tasks and evidence-accumulation decision-making models (Gold and Shadlen, 2007; Bogacz et al., 2010; van Maanen et al., 2011; Heitz and Schall, 2012). Typically, in these studies, caution is set by a cue that instructs the subject to respond rapidly (low caution) or accurately (high caution) to obtain reward; caution to avoid errors is represented by delayed response timing. Under these circumstances, human frontal cortical areas projecting to striatum become activated and have been proposed to be involved in updating the level of caution (van Maanen et al., 2011). Much

Received Sep. 17, 2021; revised June 9, 2022; accepted June 13, 2022.

Author contributions: M.A.C.-A. designed research; J.Z., S.H., M.S.S., and M.A.C.-A. performed research; J.Z., S.H., and M.A.C.-A. analyzed data; M.A.C.-A. wrote the first draft of the paper; J.Z., S.H., M.S.S., and M.A.C.-A. edited the paper; M.A.C.-A. wrote the paper.

This work was supported by National Institutes of Health grants (M.A.C.-A.). Additional information at castro-lab.org. We thank Natan Busel for technical assistance.

The authors declare no competing financial interests.

Correspondence should be addressed to Manuel A. Castro-Alamancos at mcastro@uchc.edu.

<https://doi.org/10.1523/JNEUROSCI.1892-21.2022>

Copyright © 2022 the authors

less is known about the neural circuits and response adjustments that occur when subjects must avoid punishment (Guitart-Masip et al., 2012; Millner et al., 2018). Moreover, caution resulting from the avoidance of harmful punishment in dangerous situations may be represented by different response adjustments and neural circuits than caution resulting from the avoidance of errors associated with the omission or loss of reward. Here, we studied how mice adjust the timing of signaled actions when they must be cautious to avoid danger.

Materials and Methods

Experimental design and statistical analysis

All procedures were reviewed and approved by the institutional animal care and use committee and conducted in adult (more than eight-week-old) male and female C57BL/6J mice (Jax 00664) unless otherwise indicated. The results from both sexes were combined since there is no sex difference in the basic active avoidance behavior measured for the strains used (Hormigo et al., 2019).

Unless otherwise stated, experiments involved a repeated measures design in which the mice serve as their own controls. We tested for a main effect using a repeated measures ANOVA followed by comparisons with Tukey's test. Tukey's test was conducted for the repeated measures factor when the within-subjects effect (F value) was statistically significant at a level of $p < 0.01$. To compare different groups of animals we used a linear mixed-effects model with the lme4 package in R. The model had two fixed effects (group and task). The group effect had two levels (e.g., AA and AR) while the task effect had different levels corresponding to different task periods (e.g., AA1/AR1, AA2/AR2[1–4], and AA2/AR2[5–10]). The sessions were part of the random effects nested within the subjects. Using the model, we performed pairwise comparisons to determine whether there were differences between the groups for different task periods. To enable rigorous approaches, we maintain a local server with a central database accessed through a wiki that logs all details and metadata related to the experiments, including all information about animals and details about surgical procedures, behavioral sessions, histology, scripts used for analyses, etc. Moreover, during daily behavioral sessions, computers run experiments automatically using pre-set parameters logged for reference during analysis. Analyses are performed using scripts that automate all aspects of data analysis from access to metadata and data files to population statistics and graph generation (scripts and metadata will be accessible through our website or by request).

Surgeries

Frontal cortex lesions were performed by aspirating the cortex bilaterally under isoflurane anesthesia. The cavities were filled with sterile Gelfoam. The lesions targeted the dorsal frontal association cortex or the medial prefrontal cortex, starting ~1 mm anterior to bregma. Histologic processing of the tissue revealed the extent of these lesions which were reconstructed in 3D using NeuroLucida (Microbrightfield). Sham-lesion mice underwent the same surgical procedure without cortical aspiration.

Active avoidance task

Mice were trained in a basic signaled active avoidance task using procedures similar to those described previously for rats and mice (Cohen and Castro-Alamancos, 2007, 2010; Hormigo et al., 2016, 2019). During an active avoidance session, mice are placed in a standard shuttle box (16.1" × 6.5") that has two compartments separated by a partition with side walls forming a doorway that the animal has to traverse to shuttle between compartments. A trial consists of a 7-s avoidance interval followed by a 10-s escape interval. During the avoidance interval, an auditory conditioned stimulus (CS; 8 kHz, 80 dB) is presented for the duration of the interval or until the animal produces a conditioned response (avoidance or avoid response) by moving to the adjacent compartment, whichever occurs first. If the animal avoids by moving to the next compartment, the CS ends, the escape interval is not presented, and the trial terminates. However, if the animal does not avoid, the escape interval ensues which presents an unconditioned stimulus (US) consisting of

white noise and a mild scrambled electric foot-shock (0.3 mA) delivered through the grid floor of the occupied half of the shuttle box. This US readily drives the animal to move to the adjacent compartment (escape response), at which point the US terminates, and the escape interval and the trial ends. Thus, an avoidance response will eliminate the imminent presentation of a harmful stimulus. An escape response is driven by presentation of the harmful stimulus to eliminate the harm it causes. Successful avoidance warrants the absence of harm. Each trial is followed by an intertrial interval (duration is randomly distributed; 25- to 45-s range), during which the animal awaits the next trial. We employed four variations of the basic signaled active avoidance procedure termed AA1, AA2, AA3, and Yoked.

In AA1, mice are free to cross between compartments during the intertrial interval; there is no consequence for intertrial crossings (ITCs).

In AA2, mice receive a 0.2-s foot-shock (0.3 mA) and white noise for each ITC. Therefore, in AA2, mice must passively avoid during the intertrial interval by inhibiting their tendency to shuttle between trials. Thus, during AA2, mice perform both signaled active avoidance during the signaled avoidance interval (like in AA1) and unsignaled passive avoidance during the unsignaled intertrial interval.

In AA3, mice are subjected to a CS discrimination procedure in which they must respond differently to a CS1 (8-kHz tone at 80 dB) and a CS2 (4-kHz tone at 70 dB) presented randomly (half of the trials are CS1). Mice perform the basic signaled active avoidance to CS1 (like in AA1 and AA2), but perform signaled passive avoidance to CS2, and ITCs are not punished. In AA3, if mice shuttle during the CS2 avoidance interval (7 s), they receive a 0.5-s foot-shock (0.3 mA) with white noise and the trial ends. If animals do not shuttle during the CS2 avoidance interval, the CS2 trial terminates at the end of the avoidance interval (i.e., successful signaled passive avoidance).

In the Yoked procedure, mice receive the same punishment (US presentations) during the intertrial interval than animals subjected to AA2 but these punishments are noncontingent to ITCs. We calculated the number of ITCs per trial in mice subjected to AA2 and then the same rate of punishments were presented randomly. This was done in two groups of yoked mice. In one group, we assured that the random shocks never occurred contingent on the occurrence of an ITC, within 3 s after the occurrence of ITCs. This prevented these yoked mice from ever experiencing (by chance) the same contingency as animals performing AA2, but they received the same punishment as the AA2 mice to which they are yoked. In another yoked group, this assurance was not in place. As noted below, there was no significant difference between these two yoked groups in any measured parameter of the avoidance task.

In anxiolytic drug experiments, mice are trained in the AA1 procedure for four sessions, receiving saline injections in the last two of these sessions. This is followed by 15 drug-injection sessions, 5 in AA1 and 10 in AA2.

Active approach task

The basic active approach task (AR1) was conducted in the same shuttle boxes as the active avoidance tasks. The only difference is that at the end of each side of the shuttle box an opening allowed the mouse access to a liquid dipper that provided a fixed amount of water (0.02 μ l per reinforcement). The mice were water-restricted and maintained at 85–90% of their original body weight. During task performance, mice received 1.5–2 ml of water per day as a function of body weight. Water-restricted mice were initially placed in the shuttle box for 2 d to learn to shuttle for water; water was provided for 5 s, at the far end of the compartment that was entered, starting immediately after the mouse crossed between compartments. On the third day, and thereafter, the mice only received water when they crossed during the presentation of the CS (8 kHz, 80 dB) during the approach interval (7 s). Like in AA1, crossing in between trials had no consequence during AR1. Thus, AR1 is identical to AA1, except that animals shuttled to obtain water (AR1) instead of to avoid a foot-shock (AA1) during the CS signaled interval. As per AA2, during AR2, ITCs were punished.

Behavioral measures and video tracking in the shuttle box

There are three main variables representing task performance in the shuttle box. The percentage of active avoidance or approach responses

(% avoids or approaches) represent the trials in which the animal actively avoided the US or approached the water reward in response to the CS, respectively. The avoidance or approach latency represents the time (seconds) at which the animal enters the safe or reward compartment after the CS onset only for successful trials. The number of crossings during the intertrial interval (ITCs) represents either random shuttling because of locomotor exploratory activity in the AA1, AA3, Yoked, and AR1 procedures, or failures to passively avoid in the AA2 and AR2 procedures. Animals are also video tracked (30 FPS) during most active avoidance and approach sessions. The tracking followed the animal's contour or color markers located on the head connector above the nose and between the ears. Several movement (tracking) measures are derived and converted to metric units using calibrations (Hormigo et al., 2019). Onset response latencies were derived from the video tracking by calculating the first derivative of the movement and using a threshold approach to detect when the movement leading to a completed avoidance or approach was initiated.

Open field

At the corresponding times after drug or saline injections, naive (unhandled) mice are placed in a brightly lit square arena (16.5" × 16.5") for 15 min. Mice are automatically video tracked and the amount of time the mice spend in the periphery or center of the arena is calculated. The center area was defined by a centered virtual square with sides 2/3 the size of the full arena.

Histology

Mice were deeply anesthetized with an overdose of ketamine. Upon losing all responsiveness to a strong tail pinch, the animal was decapitated and the brain was rapidly extracted and placed in fixative. The brain was sectioned (100- μ m sections) in the coronal or sagittal planes. Sections were mounted on slides, cover-slipped with DAPI mounting media, and photographed using a fluorescent microscope. The lesions were traced in 3D using NeuroLucida (MicroBrightfield).

Results

Response timing during signaled active avoidance reflects action caution

In signaled active avoidance, mice learn to avoid an aversive US (foot-shock plus white noise) by shuttling between two compartments in a cage during an avoidance interval (7 s) signaled by the presentation of an auditory tone CS (8 kHz, 80 dB). In the basic signaled active avoidance procedure, termed AA1, mice are free to shuttle between the cage compartments during the intertrial interval; spontaneous ITCs produce no consequence. In contrast, during a procedure termed AA2, ITCs are punished; a short US occurs contingent on entering the other compartment during the intertrial interval. Otherwise, AA2 is identical to AA1. Figure 1A shows a group of mice ($n = 32$) trained first in AA1 and then subjected to training in AA2. During AA2, mice virtually abolish ITCs (Tukey $t_{(31)} = 16.3$, $p < 0.0001$) with little change in the percentage of active avoidance responses (<5%). In other words, during AA2, mice only shuttle between compartments during the CS signaled avoidance interval. Interestingly, concomitant with the suppression of ITCs during AA2, there was a consistent increase in the latency of the active avoidance responses compared with AA1 (Tukey $t_{(31)} = 18.5$, $p < 0.0001$). To evaluate this delayed latency shift, we used a nonlinear exponential Gaussian to fit the distribution of avoidance latencies for each procedure (Fig. 1B). Comparison of these curves (F test) revealed that the latency distribution was different between AA1 and AA2 (F test, $F_{(5,58)} = 185.53$, $p < 0.0001$) showing a distinct rightward, delayed shift during AA2.

Mice can adjust their movement in several ways to produce longer latency active avoidance responses during CS presentation.

For instance, mice may start moving later during the CS signaled avoidance interval (i.e., longer onset-latency), move slower, etc. We measured instantaneous speed using video tracking from CS onset (Fig. 1C) or from avoidance or escape response occurrence (Fig. 1D). Compared with AA1, during AA2 mice reduced their baseline speed measured before the CS (Fig. 1E; -1 to 0 -s window). This effect was significant for avoidance responses (Tukey $t_{(31)} = 7.21$, $p < 0.0001$) but not for escape responses (Tukey $t_{(29)} = 2.7$, $p = 0.06$); note that escape responses are uncommon (<5% of the total) compared with active avoidance responses. Upon CS presentation, mice reacted with a typical orienting movement visible as a small change in peak speed (Fig. 1F; Δ Peak speed; 0 to 1 -s window, baseline corrected) that was maximal at ~ 0.5 s (Fig. 1F; Time to peak). The orienting movement was not different between AA1 and AA2 (Tukey $t_{(31)} = 1.38$, $p = 0.33$) and was mostly absent in trials that led to escape responses. After the initial orienting movement, mice produced the avoidance response (1 - to 6 -s window), which during AA2 produced a larger change in peak speed (Tukey $t_{(31)} = 10.3$, $p < 0.0001$) and a longer time to peak speed (Tukey $t_{(31)} = 10.1$, $p < 0.0001$) compared with AA1. The difference in peak speed between AA1 and AA2 was prominent when the speed of avoidance responses was plotted from response occurrence (Fig. 1D,F; -1 to 1 -s window; Tukey $t_{(31)} = 13.6$; $p < 0.0001$), which aligns the responses as opposed to the CS onset. Importantly, the change in peak speed and time to peak speed for the escape responses driven by the US did not differ between AA1 and AA2 when measured after US onset (i.e., 7 – 9 s after CS onset; Tukey $t_{(29)} = 1.1$; $p = 0.43$) or around the escape response occurrence (Fig. 1C,D,F; Tukey $t_{(29)} = 1.8$; $p = 0.2$). Thus, during AA2, mice avoid with a higher speed that peaks later compared with AA1. The faster speed during AA2, concomitant with the longer avoidance response latencies, indicates that mice start the movement to avoid later during AA2. Indeed, estimation of the active avoidance onset-latency from the start of the CS revealed that mice begin to move later during AA2 compared with AA1 (Fig. 1G, Tukey $t_{(31)} = 5.7$, $p = 0.0003$). One possibility is that during AA2, mice freeze to the CS between the orienting and avoidance responses leading to longer avoidance response latencies. This would be visible in the speed traces as a stronger dip in speed between the CS-evoked orienting and avoidance responses. This effect was not obvious; CS-evoked speed did not dip much below the baseline speed (before the CS) during either AA1 or AA2 (Fig. 1C,D). We measured this directly by comparing the CS-evoked speed reduction between AA1 and AA2 (before the avoidance response), which were not significantly different (Tukey $t_{(31)} = 0.27$, $p = 0.84$ AA1 vs AA2).

In conclusion, punishing the action when it occurs spontaneously, un signaled outside of the CS signaled active avoidance interval (ITCs during AA2) leads mice to transform their behavior by reducing spontaneous movement, delaying the action during the CS, and compensating the delayed onset with a faster action. Mice postpone the action but move faster. During AA2, mice become more cautious in an apparent effort to maximize certainty that acting is the appropriate decision.

Action caution develops regardless of active avoidance interval duration

The basic signaled active avoidance procedure uses a 7 -s avoidance interval during which the CS is presented. One possibility is that the delay of the active avoidance response observed during AA2 is related to a particular avoidance interval duration. To address this, we trained three groups of mice ($n = 27$; 9 animals

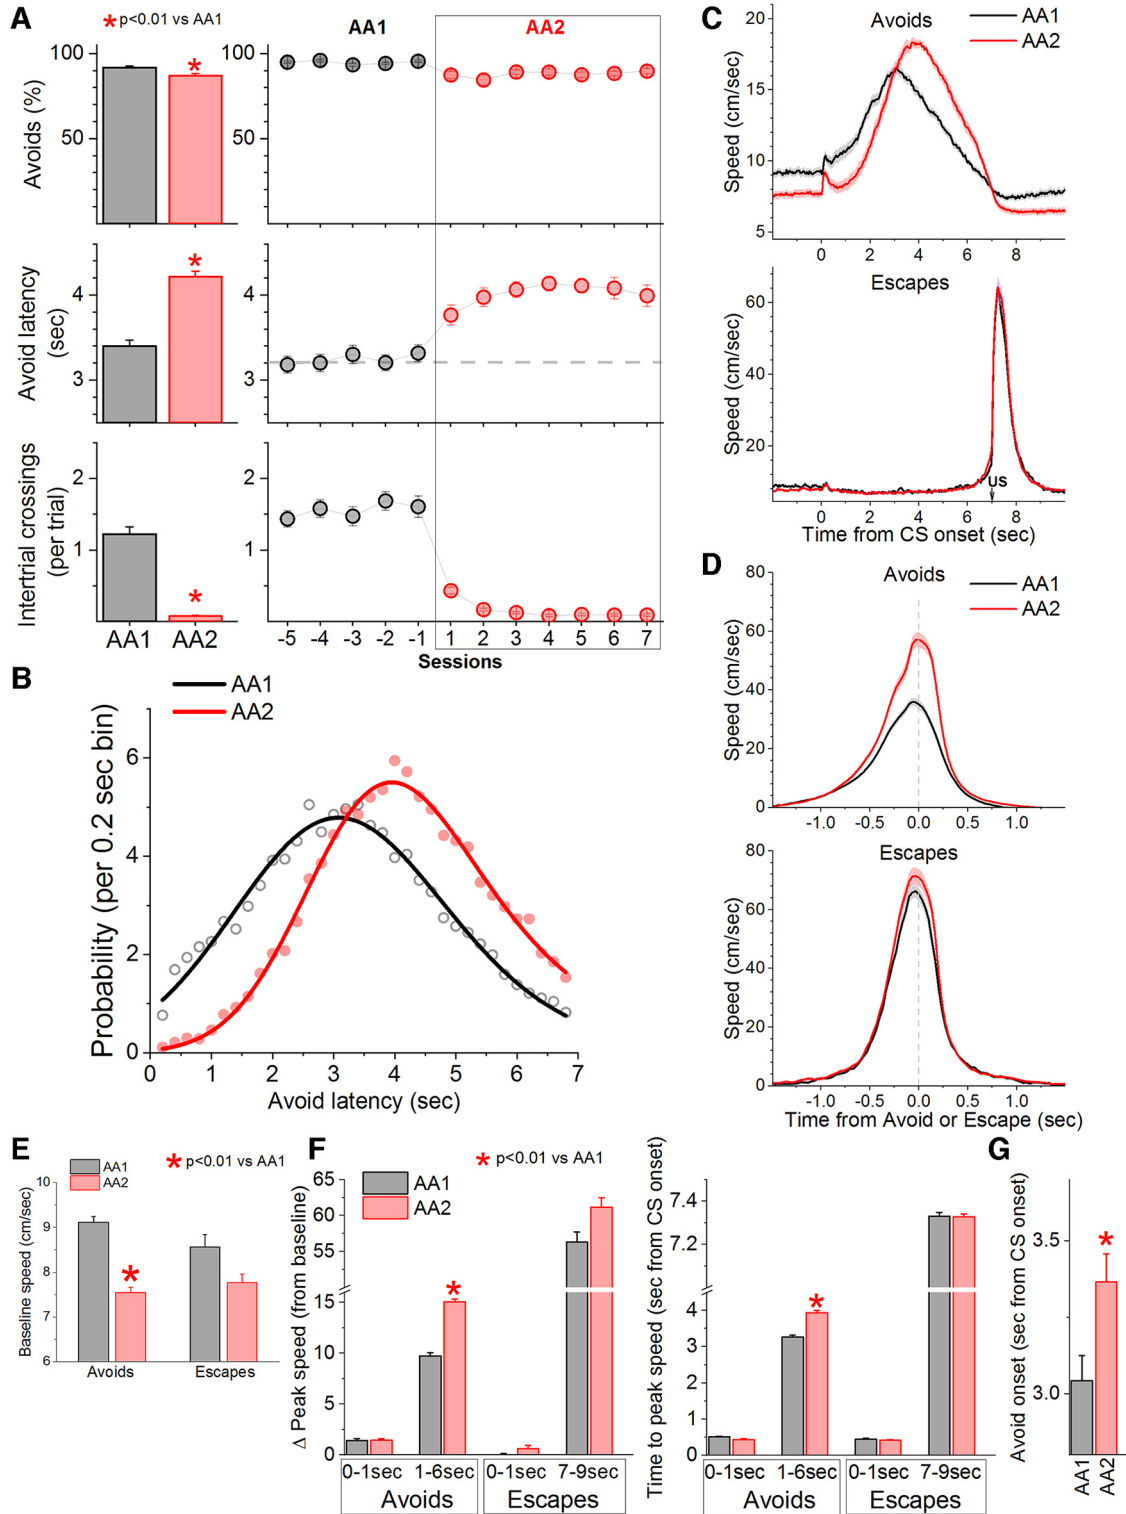


Figure 1. Punishing an action when it occurs unsignaled leads to caution about producing the signaled action. **A**, Performance of signaled active avoidance during two different procedures (AA1 and AA2) that vary only with respect to the consequence of producing ITCs. In AA1, ITCs have no consequence. In AA2, ITCs are punished. Percentage of active avoidance responses (upper), avoidance latency (middle), and ITCs (lower) of a group of mice during AA1 followed by AA2 on consecutive daily sessions. Punishing ITCs virtually abolishes these responses but also delays the timing of the active avoidance latencies with little effect on the percentage of avoidance responses. **B**, Probability histogram (%) of active avoidance latencies during AA1 and AA2 fitted with an exponential Gaussian. Note the rightward shift of the latencies indicating that the mice delayed their action in a sign of caution. **C**, Speed traces (mean \pm SEM) of active avoidance responses (upper) and escape responses (bottom) during AA1 and AA2 procedures aligned by the CS onset. Escapes are responses driven by the US when mice failed to avoid. Note the faster avoidance responses during AA2 despite starting at a lower baseline. **D**, Same as C but speed traces are aligned by the response occurrence (crossing into the safe compartment) and baseline speed (before trial onset) is subtracted to show the change in speed (Δ Speed). **E, F**, Comparison of the peak baseline speed (**E**) and Δ Speed (**F**) for avoidance and escape responses during different windows in relation to CS onset. Δ Speed was faster for avoidance response and peaked later during AA2 compared with AA1. **G**, Avoidance response time onset estimated from the speed traces. Mice begin moving later to avoid in AA2 compared with AA1.

per group) in the AA1 procedure followed by the AA2 procedure using either a 4 s, 7 s, or 15 s active avoidance interval. The number of ITCs were virtually abolished during AA2 with little effect on active avoidance performance (Fig. 2A). The active avoidance latencies adapted to the particular avoidance interval duration; shorter CS intervals produced shorter active avoidance response latencies (Gallistel and Gibbon, 2000). However, in the three groups of mice, the active avoidance latencies increased during AA2 compared with AA1 (Fig. 2A; Tukey $t_{(8)} = 11.09$, $p < 0.0001$ for 4 s; $t_{(8)} = 17.2$, $p < 0.0001$ for 7 s; $t_{(8)} = 8.7$, $p = 0.0002$ for 15 s). The fitted curves of the avoidance latency distributions were also different showing a rightward, delayed shift during AA2 for the three groups (Fig. 2B; F test, $F_{(5,28)} = 22.1$, $p < 0.0001$ for 4 s; $F_{(5,58)} = 46$, $p < 0.0001$ for 7 s; $F_{(4,140)} = 38$, $p < 0.0001$ for 15 s).

We also trained mice ($n = 9$) with a different CS auditory tone (4 kHz at 70 dB) using a standard avoidance interval (7 s) to determine whether the delay of the active avoidance response during AA2 was related to the auditory tone employed. The results obtained with the 4-kHz tone replicated those observed with the 8-kHz tone (data not shown). Thus, during AA2, the number of ITCs were virtually abolished without detriment in active avoidance performance, and the active avoidance latencies shifted longer (Tukey $t_{(8)} = 15$, $p < 0.0001$). In conclusion, during AA2, animals become more cautious by shifting their active avoidance response latencies regardless of the duration of the avoidance interval or the CS auditory tone.

Noncontingent punishment does not produce action caution

ITCs are the same movement as the learned CS signaled active avoidance response but produced spontaneously during the unsignaled intertrial interval. It is possible that mice become cautious about producing the active avoidance response (reflected by the delayed timing) because that same action is punished when it occurs unsignaled (ITC). To test this possibility, we used a Yoked procedure during which mice receive the same punishment (US presentations) during the intertrial interval than animals subjected to AA2 but these punishments are noncontingent to ITCs. Mice were trained in AA1 and then subjected to the yoked procedure. In one yoked group (Yoked1 $n = 9$), we assured that none of the random punishments coincided with a spontaneous ITC to prevent the mice from experiencing the same contingency as the AA2 mice. In another yoked group (Yoked2 $n = 8$), the random punishment did not include this assurance. The results from the two yoked groups (mixed ANOVA Group \times Procedure) did not differ either during the AA1 procedure or the subsequent yoked procedure in avoidance rate (Yoked1 group vs Yoked2 group in AA1: 94.5 ± 0.7 vs 90.3 ± 2.4 Tukey $t_{(15)} = 3.4$, $p = 0.12$; Yoked1 group vs Yoked2 group in yoked procedure: 97.1 ± 0.5 vs 95.3 ± 0.9 ; $t_{(15)} = 1.3$, $p = 0.8$), avoidance latency (Yoked1 group vs Yoked2 group in AA1: 2.93 ± 0.1 vs 3.1 ± 0.2 ; Tukey $t_{(15)} = 1.3$, $p = 0.8$; Yoked1 group vs Yoked2 group in yoked procedure: 2.7 ± 0.1 vs 3.2 ± 0.1 ; $t_{(15)} = 2.1$, $p = 0.3$) or number of ITCs (Yoked1 group vs Yoked2 group in AA1: 1.8 ± 0.2 vs 1.3 ± 0.3 ; Tukey $t_{(15)} = 1.7$, $p = 0.6$; Yoked1 group vs Yoked2 group in yoked procedure: 2.4 ± 0.3 vs 1.7 ± 0.3 , $t_{(15)} = 2.8$, $p = 0.23$). Figure 3 shows the data from the yoked group with the noted assurance (Yoked1). Interestingly, despite being subjected to the same punishment rate as mice in the AA2 procedure, the yoked mice did not delay their CS signaled active avoidance response latencies (Fig. 3A; Tukey $t_{(16)} = 1.9$, $p = 0.36$), and there was no detriment in active avoidance rate. In fact, the distribution of active avoidance latencies in yoked mice tended to become faster, although this was not significant (Fig. 3B; F test, $F_{(4,60)} = 1.2$,

$p = 0.3$). Moreover, the yoked mice did not decrease the number of ITCs but actually tended to increase them (Tukey $t_{(16)} = 3.29$, $p = 0.01$). This is a consequence of mice crossing when they are yoke-shocked, since this is what they learned to do during AA1, escape the US by crossing.

Importantly, subsequent training of the yoked mice in AA2 virtually abolished the ITCs (Tukey $t_{(16)} = 6.3$, $p < 0.0001$ vs AA1) and resulted in the typical delayed signaled active avoidance latencies (Tukey $t_{(16)} = 6.2$, $p < 0.0001$ vs AA1), which was also evident in the latency distributions (Fig. 3B; F test, $F_{(4,60)} = 9.2$, $p < 0.0001$ vs AA1). In conclusion, mice become cautious about generating the CS signaled action when the unsignaled action (ITC) is punished, not when the same amount of punishment occurs unrelated to the action. The mice become cautious about producing the action when the action is punished in other periods.

Punishing the action during a discriminable signal does not cause action caution

The previous results indicate that mice become cautious about generating the signaled action when occurrence of the action unsignaled is punished. Punishing ITCs during the unsignaled intertrial interval may generate uncertainty (because of the lack of predictability) about when the action can occur without harmful consequences (e.g., the subject cannot sample all moments of the unsignaled period to be certain about the consequences). Signaling provides predictability about when harmful consequences for producing the action occur. This certainty may focus caution on the specific signaled period instead of the action per se. Thus, we next asked whether signaling the period when the action is punished, with a discriminable cue, would have a different effect on action caution. This procedure (AA3) is similar to a typical Go/No-Go task (Guitart-Masip et al., 2012). Mice ($n = 9$) were first subjected to AA1 followed by AA3, during which mice continue performing signaled active avoidance during presentation of CS1 (8 kHz, 80 dB; same used during AA1) but must also perform signaled passive avoidance during presentation of CS2 (4 kHz, 70 dB). In AA3, mice are punished if they shuttle during the passive avoidance interval (7 s) signaled by CS2, but not if they shuttle during the intertrial interval. We found that during AA3, mice continue to actively avoid the US during CS1 at the same rate as during AA1, and passively avoid during CS2. The number of ITCs during AA3, which are not punished, were only reduced during the intertrial interval after CS2 trials when they passively avoided (Tukey $t_{(32)} = 12$, $p < 0.0001$ vs AA1) but not after CS1 trials when they actively avoided (Tukey $t_{(32)} = 1.6$, $p = 0.7$ vs AA1). Interestingly, active avoidance latencies to CS1 did not change during AA3 (Fig. 3C; Tukey $t_{(32)} = 2.2$, $p = 0.5$ vs AA1). Moreover, the distribution of the latencies did not differ greatly compared with AA1, although they tended to shift slightly (Fig. 3D; F test, $F_{(4,60)} = 3.03$, $p = 0.02$ vs AA1). The results indicate that punishing the action during a discriminable signaled period (CS2) does not lead mice to become cautious about producing the action when it is required (CS1). In other words, punishing the action during a specific signaled interval (CS2 in AA3), instead of the whole unsignaled intertrial interval (AA2), alleviates caution about generating the action when it is required.

We further trained these mice in a challenging reversal of the AA3 procedure. In this procedure, termed AA3rev, the auditory tones are swapped (CS1 becomes 4 kHz, 70 dB and CS2 becomes 8 kHz, 80 dB) while keeping the same contingencies (CS1 continues to require active avoidance and CS2 passive

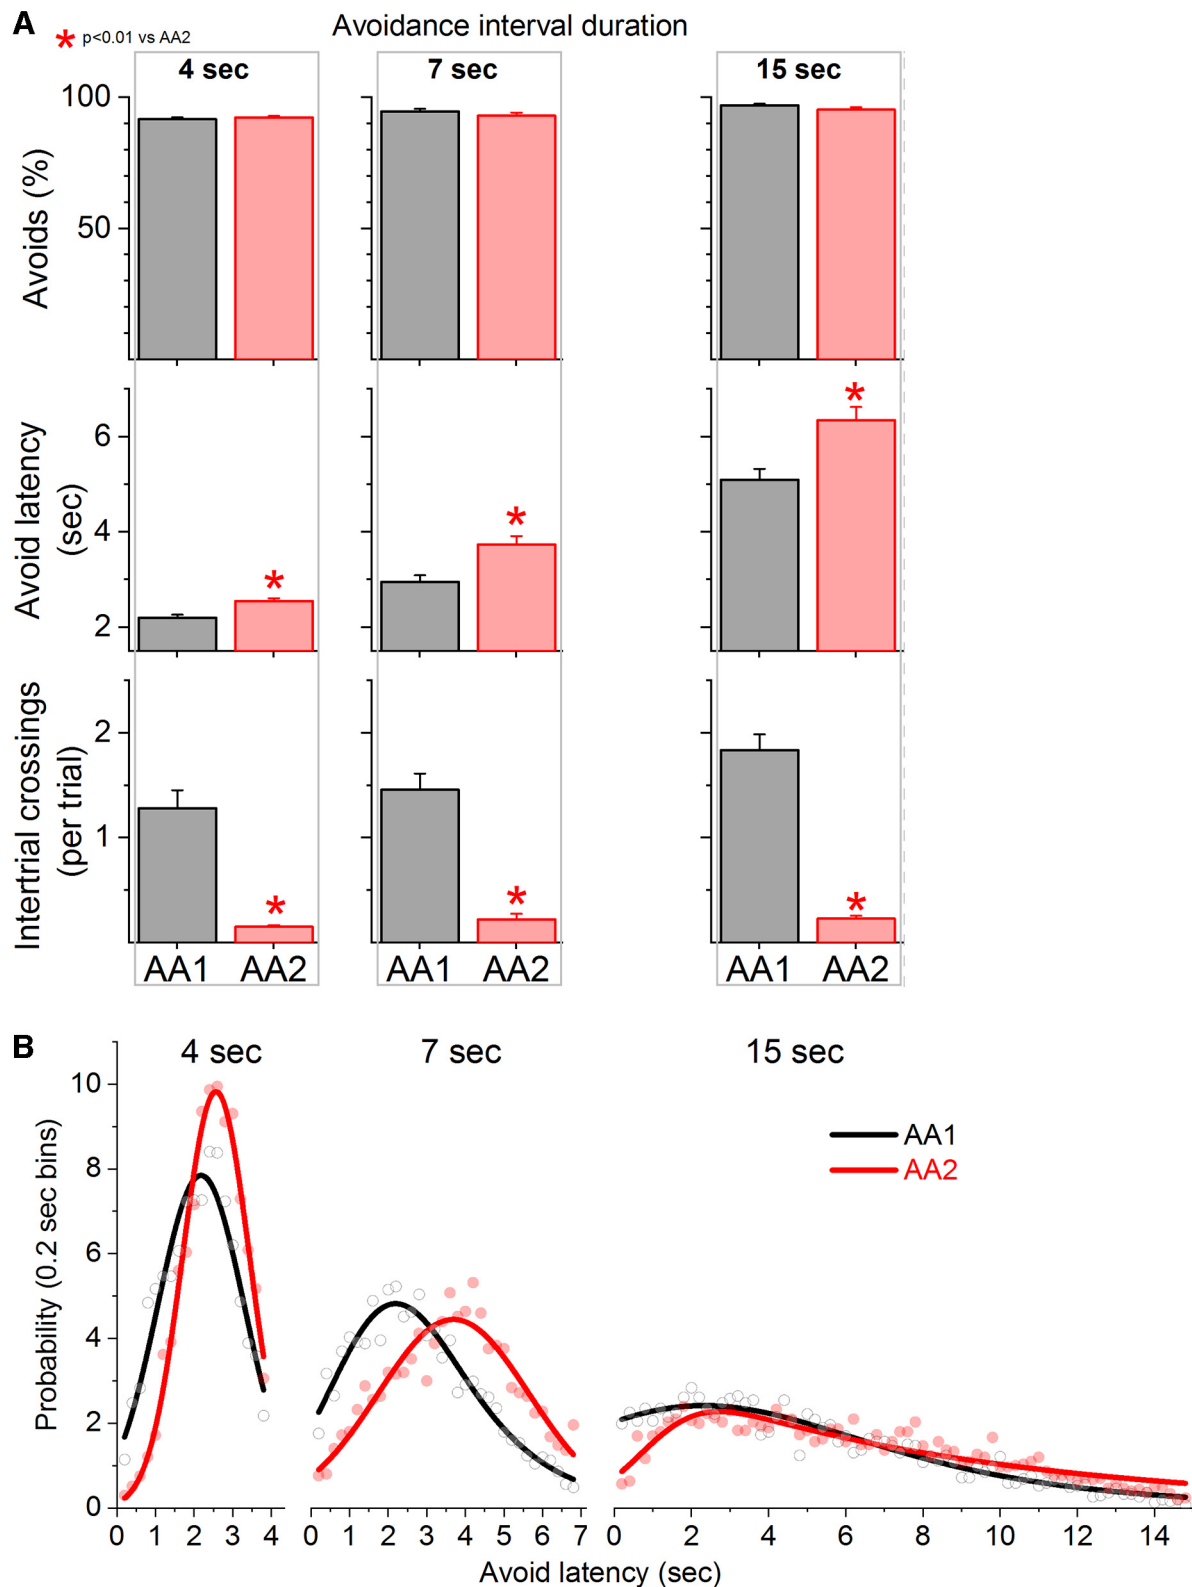


Figure 2. Punishing the unsigned action leads to caution about producing the signaled action regardless of the duration of the avoidance interval. **A**, Performance of signaled active avoidance during three different AA1 procedures in which the avoidance interval lasted 4, 7, or 15 s. Addition of AA2, which punishes ITCs, in any of these procedures caused delayed avoidance latencies, regardless of the duration of the avoidance interval. **B**, Probability histogram (%) of avoidance latencies during AA1 and AA2 fitted with an exponential Gaussian for the procedures in **A**. Note the rightward shift of the latencies indicating the mice delayed the signaled action in a sign of caution.

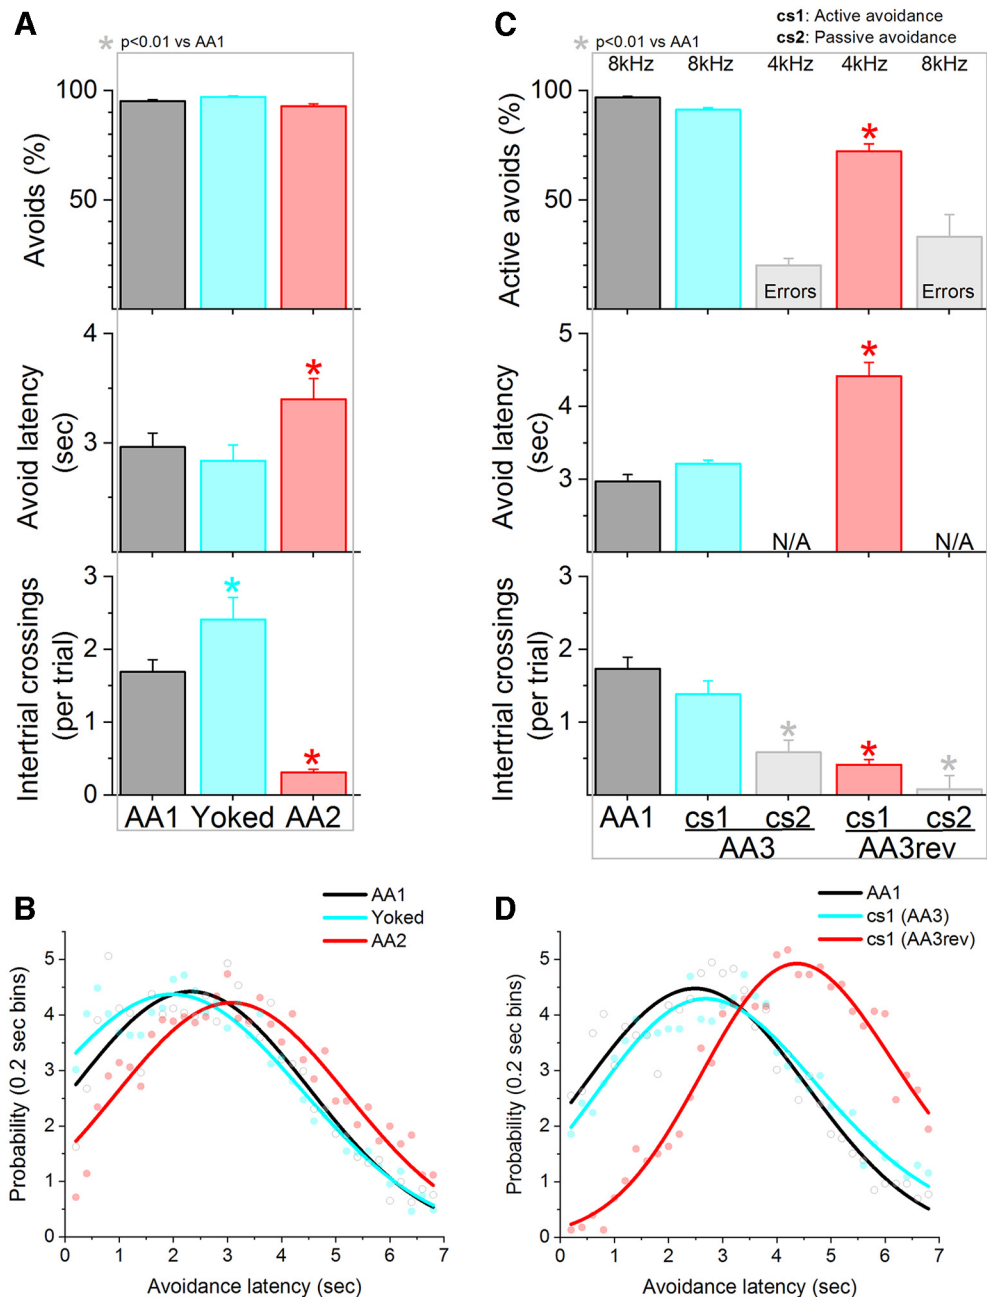


Figure 3. Random punishment noncontingent to the unsignaled action (Yoked) or signaling the interval when the action is punished does not lead to caution about producing the signaled action. **A**, Performance of signaled active avoidance during the AA1 procedure followed by Yoked procedures. Yoked mice receive random presentations of punishment (US) during the intertrial interval at the same rate experienced by animals performing AA2, but noncontingent to ITCs. During Yoked, the number of ITCs increased and the signaled avoidance latency was not delayed. Subsequent training in AA2 lead to virtual abolishment of ITCs and a concomitant delay of avoidance latencies. Thus, random punishment does not lead to caution about producing the signaled action. **B**, Probability histogram (%) of avoidance latencies during AA1, Yoked, and AA2 fitted with an exponential Gaussian. Note the rightward shift of the latencies during AA2 indicating the mice delayed their action in a sign of caution only when the unsignaled action was punished, not when the same amount of punishment was delivered unrelated to the action. **C**, Performance of signaled active avoidance during AA1 followed by AA3 and AA3rev. AA3 is a discrimination procedure that requires mice to continue active avoidance during presentation of CS1 (8-kHz tone) but crossings are punished during CS2 (4-kHz tone) only, not during the intertrial interval. As in AA1, during AA3 mice must actively avoid during CS1 but passively avoid during CS2 only (not during the whole intertrial period as in AA2). Signaling the period when the action is punished (AA3) did not lead to the development of caution about producing the signaled action. Subsequently, reversal of the tones that signal CS1 and CS2 contingencies (AA3rev) led to the development of strong caution about generating the signaled action. This occurred concomitant with worse performance and a reduction of ITCs, although ITCs are not punished. **D**, Probability histogram (%) of avoidance latencies during AA1, AA3, and AA3rev fitted with an exponential Gaussian. Note the rightward shift of the latencies during AA3rev indicating the mice delayed their action in a sign of caution only when the meaning of the signaling was reversed.

avoidance). We expected this to lead to greater caution because the signaling is completely altered (e.g., a simile in humans would be that suddenly, the meaning of streetlights is reversed). While mice are able to learn the new rule, this procedure is challenging, and mice perform at a lower level compared with the initial AA3; producing fewer active avoids to CS1 and more

errors (shuttling) to CS2 (~70/40% shuttling to CS1/CS2 during AA3rev compared with ~90/20% during AA3). AA3rev was associated with a sharp delay in active avoidance latencies to CS1 (Fig. 3C; Tukey $t_{(32)} = 10.4$, $p < 0.0001$ vs AA1) and a large shift of the latency distributions (Fig. 3D; F test, $F_{(4,60)} = 137.6$, $p < 0.0001$ vs AA1). Moreover, there was a reduction in ITCs occurring after

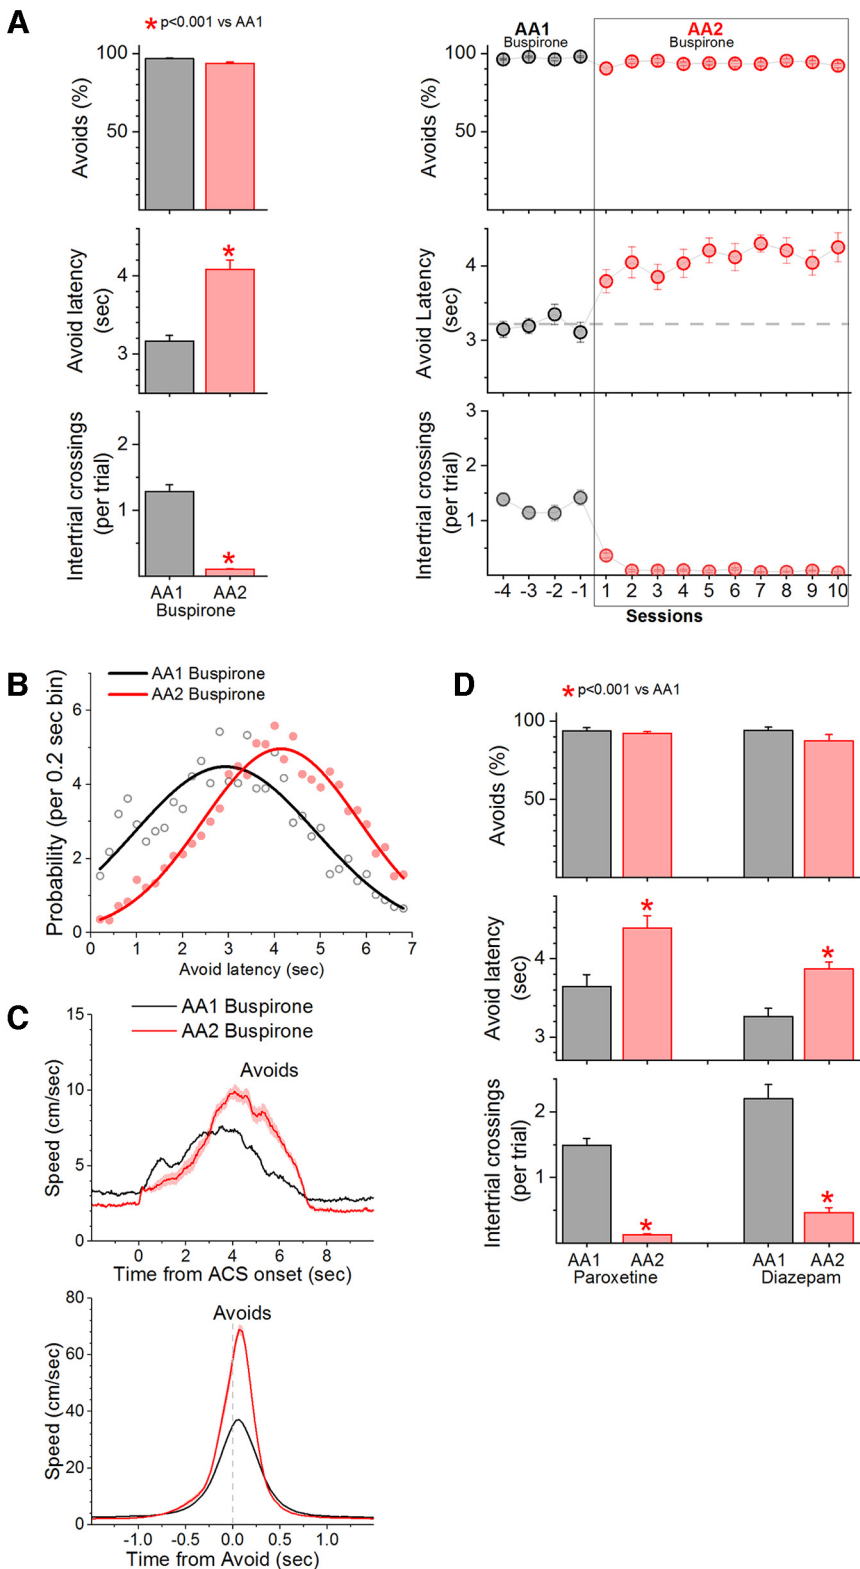


Figure 4. Anxiolytics do not abolish action caution. **A**, Performance of signaled active avoidance during the AA1 and AA2 procedures in mice injected with buspirone (2 mg/kg, i.p.). Training in AA2 led to the usual abolishment of ITCs and a sharp delay in avoidance latencies without impaired performance. Being cautious about producing the signaled action when the unsignaled action is punished is not alleviated by an anxiolytic. **B**, Probability histogram (%) of avoidance latencies during AA1 and AA2 fitted with an exponential Gaussian for mice injected with buspirone. Note the rightward shift of the latencies indicating the mice delayed their action in a sign of caution. **C**, Speed traces (mean \pm SEM) of active avoidance responses aligned by the CS onset (upper) and baselined-corrected avoidance responses aligned by the response occurrence (bottom) during AA1 and AA2 procedures. Note the faster avoidance responses during AA2 despite starting at a lower baseline. **D**, Performance of signaled active avoidance during the AA1 and AA2 procedures in mice injected with paroxetine (10 mg/kg,

CS1 (Tukey $t_{(32)} = 15.2, p < 0.0001$ vs AA1), which are not punished and had not changed during the AA3 procedure. Thus, changing the signaling rule in AA3rev makes the mice very cautious about producing the action. Signaling the period when the action is punished reduces the need to be cautious about generating the action when it is required, but changing the well-established signaling rules, increases caution greatly concomitant with impaired performance.

Action caution is not because of increased anxiety

An intriguing theoretical framework postulates that conflict between passive avoidance and active avoidance brain circuits leads to activation of a brain inhibition system that produces anxiety, and related behavioral inhibition (Gray and McNaughton, 2000; McNaughton and Corr, 2004). Increased anxiety would delay the generation of the signaled action when it is required. To test this possibility, we subjected mice ($n = 9$) to the AA1 and AA2 procedures in the presence of an anxiolytic (Buspirone, 2 mg/kg, i.p., injected 45 min before each daily session) during both procedures. During AA1, buspirone increased the number of ITCs compared with saline (Tukey $t_{(8)} = 5.7, p = 0.003$) without affecting performance or avoidance latency. In addition, we found that AA2 caused a delay in the signaled active avoidance latencies compared with AA1 (Fig. 4A; Tukey $t_{(8)} = 14.6, p < 0.0001$). The fitted curves of the avoidance latency distributions were also different showing the typical rightward, delayed shift during AA2 (Fig. 4B; F test, $F_{(5,58)} = 33.5, p < 0.0001$). This occurred without detriment in signaled active avoidance performance and with virtual abolishment of ITCs. Moreover, as in controls, Δ peak speed was faster during AA2 compared with AA1 (Fig. 4C; Tukey $t_{(8)} = 24.8, p < 0.0001$).

To verify that buspirone had an anxiolytic effect, we placed naive mice in an open field and compared the percentage of time saline-injected ($n = 9$ mice) or buspirone-injected ($n = 9$ mice) mice spent in the center of the open field compared with the periphery. Consistent with an anxiolytic effect, buspirone significantly increased

i.p.) or diazepam (1 mg/kg, i.p.). Training in AA2 led to the usual abolishment of ITCs and a sharp delay in avoidance latencies without impaired performance. Being cautious about producing the signaled action when the unsignaled action is punished is not alleviated by two different additional anxiolytics.

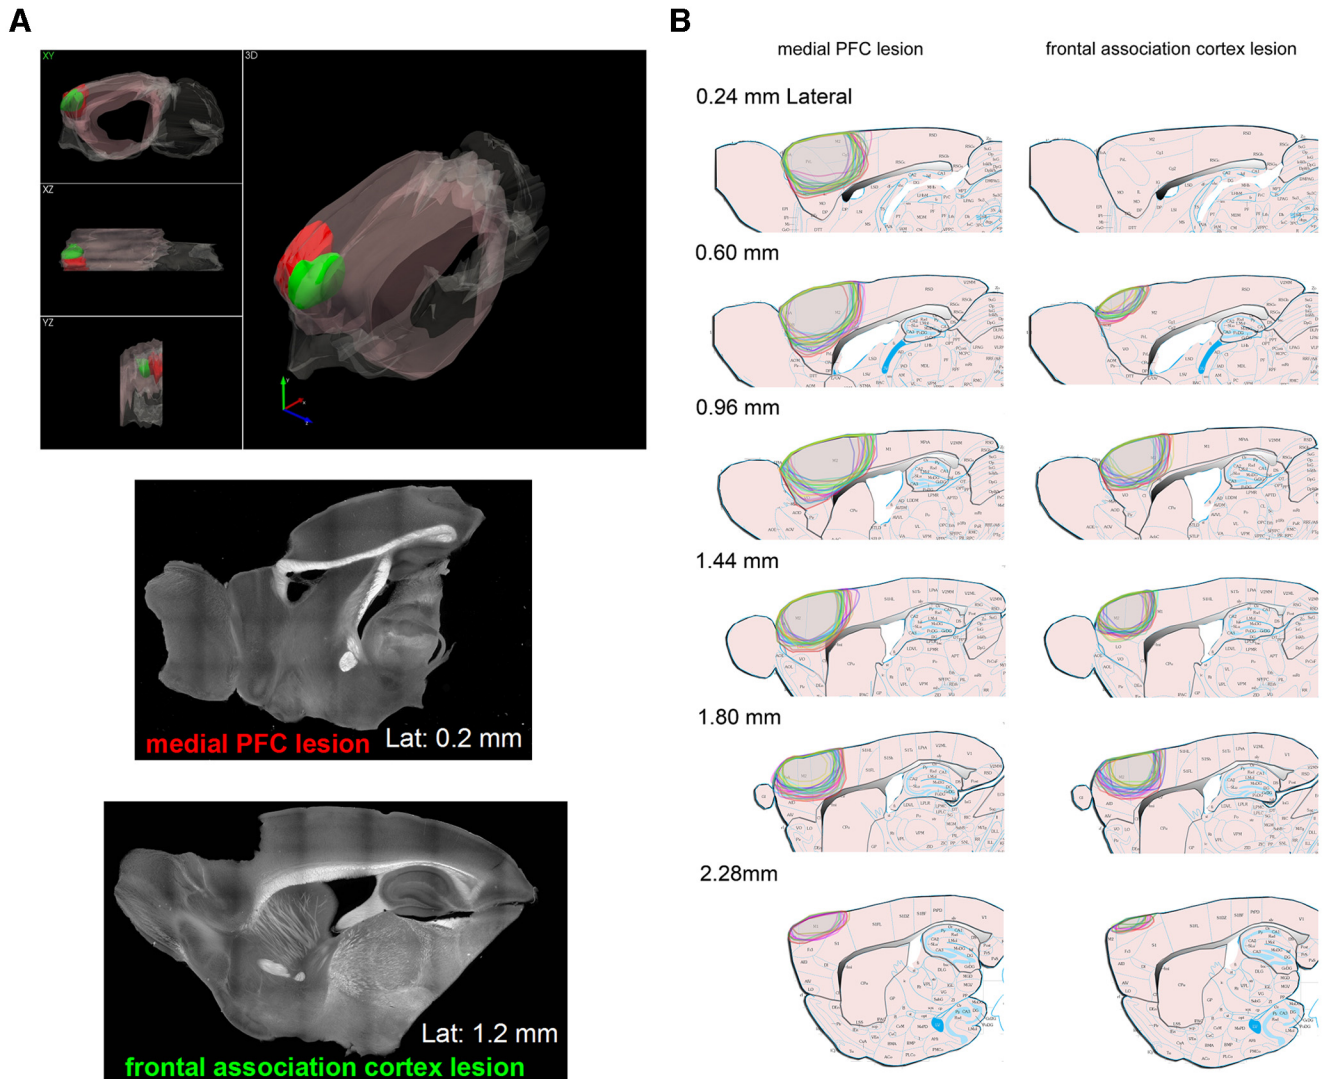


Figure 5. Lesions of the frontal cortex. **A**, Examples of frontal cortex lesions centered in the frontal association cortex (green volume and lower right) or the medial prefrontal cortex (red volume and upper right; mPFC). The 3D plot (left) overlays the tracings of lesions from two animals (one from each group) on one side of the brain. The 3D reconstruction is shown in [Movie 1](#) (top panel). The colored volumes depict the minimal lesion area for each group because there was always damage extending into the adjacent area. According to atlas nomenclature (Franklin and Paxinos, 2008), the frontal association cortex lesion eliminated FrA and rostral part of M2, while the medial prefrontal cortex lesion eliminated PrL, IL, and the rostral part of Cg1 and Cg2. The *x*, *y*, *z* arrows point in the posterior, dorsal, and lateral directions, respectively, and represent 1 mm. The bottom panels show dark-field images taken from sagittal sections of both lesion types. **B**, Lesions reconstructed on the atlas for the medial prefrontal cortex (PFC) and frontal association cortex groups. The traced areas show the lesions for all animals in sagittal plane sections at different distances from the midline.

the time mice spent in the center of the open field (saline vs buspirone, $88 \pm 1\%$ vs $81 \pm 2\%$ of time in the periphery of the open field; unpaired *t* test $t_{(16)} = 2.7$, $p = 0.01$).

In addition to buspirone, we tested the effects of two additional anxiolytics, paroxetine (10 mg/kg, i.p., injected 60 min before each daily session; $n = 9$ mice) and diazepam (1 mg/kg, i.p., injected 30 min before each daily session; $n = 15$ mice). While diazepam is well known to have hypnotic/sedative effects, we used a dose that minimizes these effects in mice (Pádua-Reis et al., 2021). As per buspirone, during AA1, both drugs increased the number of ITCs compared with saline (Tukey $t_{(8)} = 6.9$, $p = 0.001$ for paroxetine; Tukey $t_{(14)} = 7.3$, $p = 0.0001$ for diazepam) without affecting performance or avoidance latency. In addition, we found that AA2 caused a delay in the signaled active avoidance latencies compared with AA1 in both paroxetine (Fig. 4D; Tukey $t_{(8)} = 12.4$, $p < 0.0001$) and diazepam (Tukey $t_{(14)} = 10.2$, $p < 0.0001$) injected mice. The fitted curves of the avoidance

latency distributions were also different showing the typical rightward, delayed shift during AA2 (*F* test, $F_{(5,58)} = 30.6$, $p < 0.0001$ for paroxetine; *F* test, $F_{(5,58)} = 18.5$, $p < 0.0001$ for diazepam). The avoidance latency shift occurred without detriment in signaled active avoidance performance and with virtual abolishment of ITCs. Moreover, in the presence of these drugs, Δ peak speed was faster during AA2 compared with AA1, recapitulating the effects of buspirone (Tukey $t_{(8)} = 11.6$, $p < 0.0001$ for paroxetine; Tukey $t_{(14)} = 9.5$, $p < 0.0001$ for diazepam). In conclusion, a state of anxiety is not causing the action caution resulting from punishing the unsignaled action. Indeed, it seems logical to infer that mice can be cautious without being anxious.

Action caution occurs without frontal cortex or during basal ganglia output deactivation

The frontal cortex together with its projections to basal ganglia have been proposed as a circuit that controls response timing in

relation to caution in speed-accuracy trade-off situations (van Maanen et al., 2011), and generally responding during conflict situations (Botvinick et al., 2004). The term avoidance is used in many contexts. Behavioral tasks that include this term in their name (e.g., Sidman avoidance, Platform avoidance, etc.) have studied the involvement of the prefrontal cortex, but those procedures are very different from the classical signaled active avoidance procedure used in our study. For example, one study (Capuzzo and Floresco, 2020) found effects of prefrontal cortex inactivation on lever-press avoidance in rats, but rats are well-known to acquire and perform this behavior very poorly if at all (Bolles, 1970). A more recent study showed that prefrontal cortex inactivation using muscimol and optogenetics produced deficits in mice performing signaled active avoidance, but post-training lesions were not tested (Jercog et al., 2021); perhaps technical issues related to drug diffusion and well-known unwanted consequences of optogenetic inhibition (post-inhibition rebound excitation) may explain those results. One study in rats investigated the effects of post-training prefrontal cortex lesions on signaled active avoidance (Castro-Alamancos and Borrell, 1992), which found little effect on performance. Moreover, little is known about the neural circuits that generate caution when the signaled action is punished if it occurs unsignaled in a nonconflicting manner (as in AA2). We bilaterally aspirated two different parts of the frontal cortex in two groups of mice that had been previously trained in AA1. The lesions (Fig. 5) included the dorsal frontal association cortex ($n = 11$ mice) or the medial prefrontal cortex ($n = 9$ mice). The lesions, which were reconstructed on an atlas (Franklin and Paxinos, 2008), eliminated the targeted frontal cortex areas but extended significantly into adjacent areas (Fig. 5). In the medial prefrontal group, the lesion eliminated most of the prefrontal cortex, including anterior cingulate, and the medial dorsal, prelimbic, and infralimbic areas. The results of the two lesion groups were combined into a single frontal cortex lesion group ($n = 20$ mice) because there were no significant differences between them. An additional sham group ($n = 8$ mice) underwent the surgical procedure without any lesion. After the lesion (7–10 d after), the mice were placed back in the AA1 procedure. The lesion caused little change on the ability of the mice to perform AA1 (Fig. 6A), as previously shown (Castro-Alamancos and Borrell, 1992). Mixed ANOVAs of Group \times Lesion interaction revealed that there were no significant effects on the percentage of avoidance

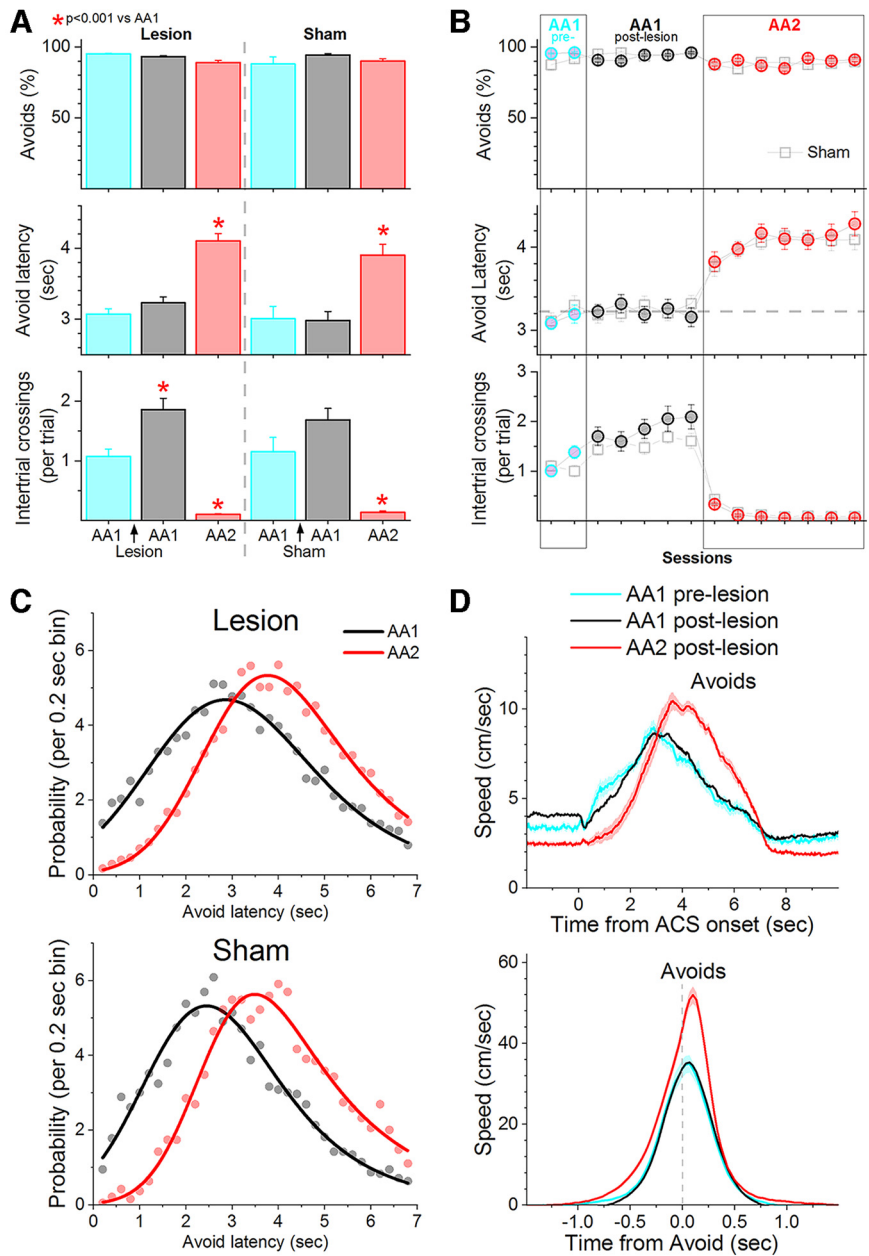


Figure 6. Lesions of the frontal cortex do not interfere with the development of caution when the unsignaled action is punished. **A, B**, Performance of signaled active avoidance during the AA1 procedure before and after frontal cortex lesions (with both lesion groups shown in Fig. 5 combined), and subsequent training of lesioned mice in AA2. During AA1, frontal cortex lesions had no effect on the percentage of avoidance responses or avoidance latencies, but increased the number of ITCs. Subsequent training in AA2 led to abolishment of the ITCs and a sharp delay in avoidance latencies without impaired performance. Being cautious about producing the signaled action when the unsignaled action is punished does not require the frontal cortex. **C**, Probability histogram (%) of avoidance latencies during AA1 and AA2 fitted with an exponential Gaussian for frontal cortex lesion and sham lesion mice. Note the rightward shift of the latencies indicating the mice delayed their action in a sign of caution. **D**, Speed traces (mean \pm SEM) of active avoidance responses aligned by the CS onset (upper) and baselined-corrected avoidance responses aligned by the response occurrence (bottom) during AA1 (prelesion and postlesion) and AA2 procedures. Note the faster avoidance responses during AA2 despite starting at a lower baseline.

responses for the lesion group (Fig. 6A,B; Tukey $t_{(54)} = 1.8$, $p = 0.77$ vs prelesion) or the sham group (Fig. 6A,B; Tukey $t_{(54)} = 2.5$, $p = 0.48$ vs prelesion), which did not differ significantly between each other before or after the lesion (Fig. 6A,B; Tukey $t_{(54)} = 1.03$, $p = 0.97$ Lesion vs Sham AA1 after the lesion). The lesion did not affect the avoidance latencies (Fig. 6A; Tukey $t_{(54)} = 2.16$, $p = 0.64$ vs prelesion), and these were not different between the lesion and sham mice (Fig. 6A,B; Tukey $t_{(54)} = 2.5$,

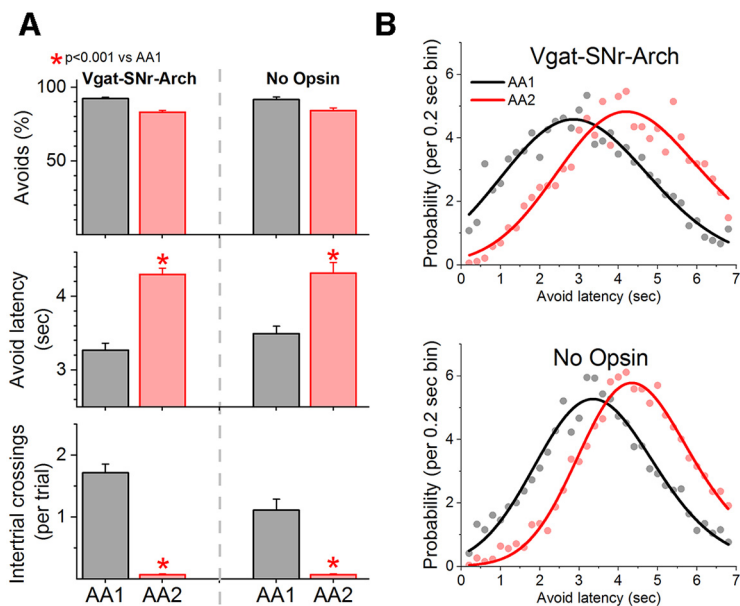


Figure 7. Deactivating the basal ganglia output via the SNr does not interfere with the development of caution when the unsignaled action is punished. **A**, Performance of signaled active avoidance during the AA1 and AA2 procedures when SNr GABAergic neurons are inhibited during the avoidance interval and a period preceding it. The data are taken from Hormigo et al. (2021b), which did not consider avoidance latencies. Training in AA2 led to abolishment of the ITCs and a sharp delay in avoidance latencies without impaired performance indicating that basal ganglia output via SNr is not required for being cautious about producing the signaled action. A No opsin group of mice that underwent the same optogenetic procedures but did not express Arch is included for comparison. **B**, Probability histogram (%) of avoidance latencies during AA1 and AA2 fitted with an exponential Gaussian for the data in **A**. Note the rightward shift of the latencies indicating the mice delayed their action in a sign of caution in both the Opsin and No Opsin groups.

$p = 0.49$ Lesion vs Sham AA1 after the lesion). The lesion caused an increase in the number of ITCs (Fig. 6A; Tukey $t_{(54)} = 6.29$, $p < 0.0001$); this tended to occur to some extent in the sham animals but not significantly (Fig. 6A; Tukey $t_{(54)} = 2.7$, $p = 0.39$). However, when the mice were subjected to the AA2 procedure, the ITCs were virtually abolished in both the lesion (Fig. 6A,B; Tukey $t_{(54)} = 14.1$, $p < 0.0001$ AA1 vs AA2 postlesion) and sham (Fig. 6A,B; Tukey $t_{(54)} = 7.9$, $p < 0.0001$ AA1 vs AA2 postlesion) groups with no detriment in performance. Moreover, the avoidance latencies were sharply delayed in both the lesion (Fig. 6A,B; Tukey $t_{(54)} = 9.6$, $p < 0.0001$ AA1 vs AA2 postlesion) and sham (Tukey $t_{(54)} = 6.39$, $p = 0.0004$ AA1 vs AA2 postlesion) groups, as were the latency distributions (Fig. 6C; F test, $F_{(5,58)} = 119.4$, $p < 0.0001$ AA1 vs AA2 postlesion; lesion group). During AA2, there were no significant differences between the lesion and sham groups in these measures (avoids, avoidance latency, and ITCs; Tukey $p > 0.94$ lesion vs sham). Furthermore, in the lesion mice, the speed recapitulated the effects observed in normal mice, including faster peak speeds during AA2 (Fig. 6D; Tukey $t_{(40)} = 12.9$, $p < 0.0001$ vs AA1 postlesion; lesion group). Thus, like normal mice, frontal cortex lesion mice are cautious about generating the signaled action when occurrence of the action unsignaled is punished (as during AA2); the frontal cortex is not required for mice to be cautious.

In a recent paper, we used fiber photometry to compare calcium signals of CaMKII-expressing neurons in the pedunculo-pontine tegmentum region (PPT) of mice trained in AA1 followed by AA2 (Hormigo et al., 2021a). These neurons show robust activation during performance of signaled active avoidance and are critically required for signaled active avoidance

because avoidance is abolished when they are inhibited (Hormigo et al., 2019). Moreover, we found that the activity of these neurons reflects both the delayed latency and the higher speed observed during AA2 compared with AA1 (Hormigo et al., 2021a, see their Fig. 6A,B). Thus, cells in PPT robustly represent action caution that occurs when the unsignaled action is punished. Since the PPT region is under the control of the basal ganglia, we reasoned that the latency shift that occurs during AA2 might be driven by the output of GABAergic neurons in the substantia nigra pars reticulata (SNr). Although, deactivating SNr neurons does not impair signaled active avoidance (Hormigo et al., 2021b), indicating that SNr does not drive active avoidance responses, SNr output may modulate PPT during AA2 in a way that it causes the delayed response latency reflecting caution. If this were the case, deactivating SNr would eliminate the rightward latency shift observed during AA2 compared with AA1 without impairing performance. Using previously published data (Hormigo et al., 2021b), we performed a new analysis to compare active avoidance latencies of mice performing AA1 followed by AA2 when the output of the basal ganglia via GABAergic neurons (expressing eArch3.0 or IC+ +; Opsin group, $n = 8$) in the SNr was inhibited with optogenetics (Fig. 7A, B). GABAergic SNr cells were inhibited with green or blue light during the avoidance interval and during a random period preceding the avoidance interval ($\sim 1/4$ of the total intertrial interval) in both AA1 and AA2. This deactivates the SNr during the interval when mice produce the signaled active avoidance response; extensive validation of these methods is included in (Hormigo et al., 2021b). In that previous study, the active avoidance response latencies (i.e., successful signaled active avoidance responses) were not compared between AA1 and AA2 during SNr deactivation. Thus, we include this new analysis here. We also included a No Opsin group (with additional mice, $n = 9$) that underwent the same procedures as the Opsin mice but did not express opsins. This analysis showed that the distribution of active avoidance response latencies during SNr deactivation ($n = 8$ mice) have a robust rightward shift during AA2 compared with AA1 (Fig. 7B; F test, $F_{(5,58)} = 48.5$, $p < 0.0001$ vs AA1) concomitant with longer mean latencies (Fig. 7A; Tukey $t_{(7)} = 6.77$, $p = 0.002$ vs AA1). These results were similar to the effects observed in the No Opsin controls ($n = 9$; Fig. 7A,B). There were no differences (mixed ANOVA Group \times AA interaction) in avoidance latency between the Opsin and No Opsin groups during either AA1 (Tukey $t_{(15)} = 1.6$, $p = 0.6$ Opsin vs No Opsin during AA1) or AA2 (Tukey $t_{(15)} = 0.46$, $p = 0.98$ Opsin vs No Opsin during AA2). The main difference between the Opsin and the No Opsin mice was that during AA1 the number of ITCs were higher for the Opsin mice (Fig. 7A; Tukey $t_{(15)} = 4.86$, $p = 0.017$ Opsin vs No Opsin during AA1). This is a logical consequence of the effect of inhibiting SNr cells during the ITI (preceding the ACS), which increases the motor activity in Opsin mice (Hormigo et al., 2016, 2021b) and hence the number of ITCs. However, when ITCs are punished during AA2, ITCs are abolished in both the Opsin and No Opsin mice, and they are no longer different (Tukey $t_{(15)} = 0.04$, $p = 0.99$ Opsin vs No Opsin during AA2). Thus, disruption of the output of the

basal ganglia via SNr does not alter the ability of mice to display action caution when the unsignaled action is punished.

In conclusion, while action caution is reflected in the activity of PPT neurons that are required for signaled active avoidance, it is unlikely that the development of action caution is mediated by frontal cortex and basal ganglia output via SNr. More likely, these brain regions may monitor the occurrence of the action but are not causing the latency shifts that reflect caution.

Action caution is only transient when the action is an appetitive active approach

An important question is whether the delay of the signaled action reflecting caution, which occurs when the unsignaled action is punished, depends on the motivational state that drives the action. Thus, we trained mice in a signaled active approach procedure (AR1) that mimics signaled active avoidance (AA1). During AR1, mice perform the same action as in AA1 with the goal of obtaining reward instead of avoiding harm. Thus, water-restricted mice must shuttle during the approach interval (7 s) signaled by a CS (8 kHz, 80 dB) to obtain water at the opposite end of the compartment they enter. During AR1, water is not dispensed during the intertrial interval and ITCs are not punished.

We found that mice ($n = 10$) perform active approach responses during the approach interval at high rates (>90%; Fig. 8), but there are two major differences between the AR1 and AA1 behaviors. First, the response latencies of the signaled active approach responses are much faster (~1 s) than those of the signaled active avoidance responses (~3.5 s). Second, the number of ITCs during AR1 are generally higher than during AA1. Next, we determined the effect of punishing ITCs (AR2). During the first AR2 session, the number of ITCs were virtually abolished, but some mice also inhibited their overall active approach responses during the CS. This is understandable since it is the first time these mice are exposed to the US. The initial AR2 sessions were associated with a large increase in the latency of active approach responses, denoting action caution. However, within two to three AR2 sessions, the rate and latencies of active approach responses returned to control (AR1) levels. Thus, we separated the AR2 sessions into four early (AR2 one to four sessions) and six proper sessions (AR2 5–10 sessions). Comparison of active approach latencies during AR1 versus AR2, excluding the early AR2

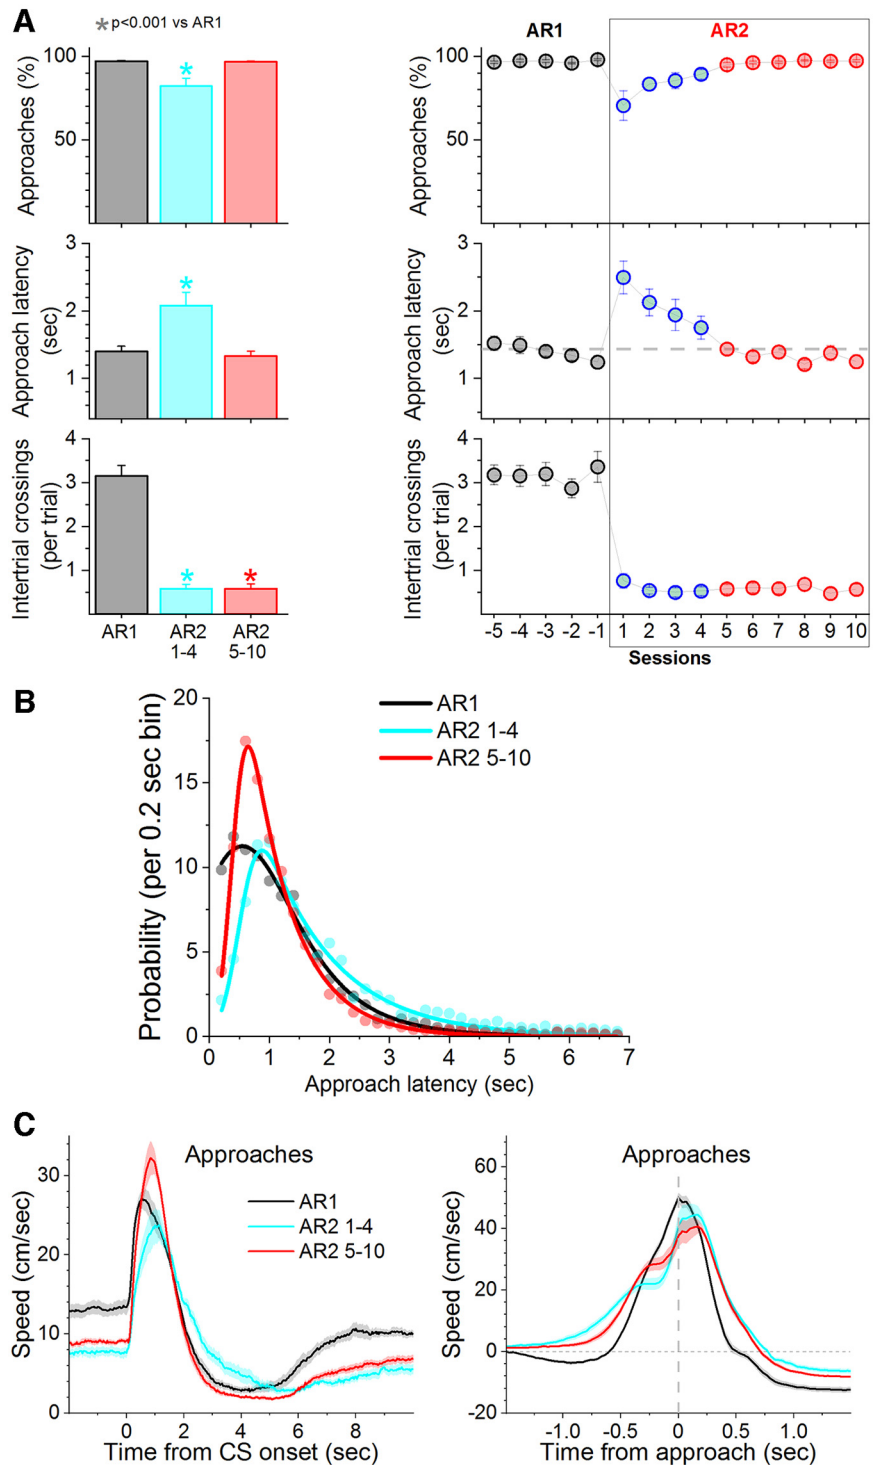


Figure 8. When the signaled action is an active approach, punishing the unsignaled action leads only to transient caution about producing the signaled action. **A**, Performance of signaled active approach (AR1) to obtain water (reward) in water-restricted mice. The task is similar to AA but mice shuttle during the CS presentation (approach interval) to obtain water instead of avoiding the US. Subsequent training in AR2, which punishes ITCs, led only to brief caution about producing the signaled action (initial 1–4 AR2 sessions; blue) but thereafter (5–10 AR2 sessions; red) caution was not evident in response timing. **B**, Probability histogram (%) of approach latencies during AR1 and AR2 fitted with an exponential Gaussian. Note the slight rightward shift of the latencies during the initial one to four sessions but subsequently (5–10 sessions) latencies shift left producing a sharp short-latency peak. During AR2, when the action is an approach, mice only display caution transiently. **C**, Speed traces (mean ± SEM) of active approach responses aligned by the CS onset (left) and baseline-corrected approach responses aligned by the response occurrence (right) during AR1 and AR2 procedures.

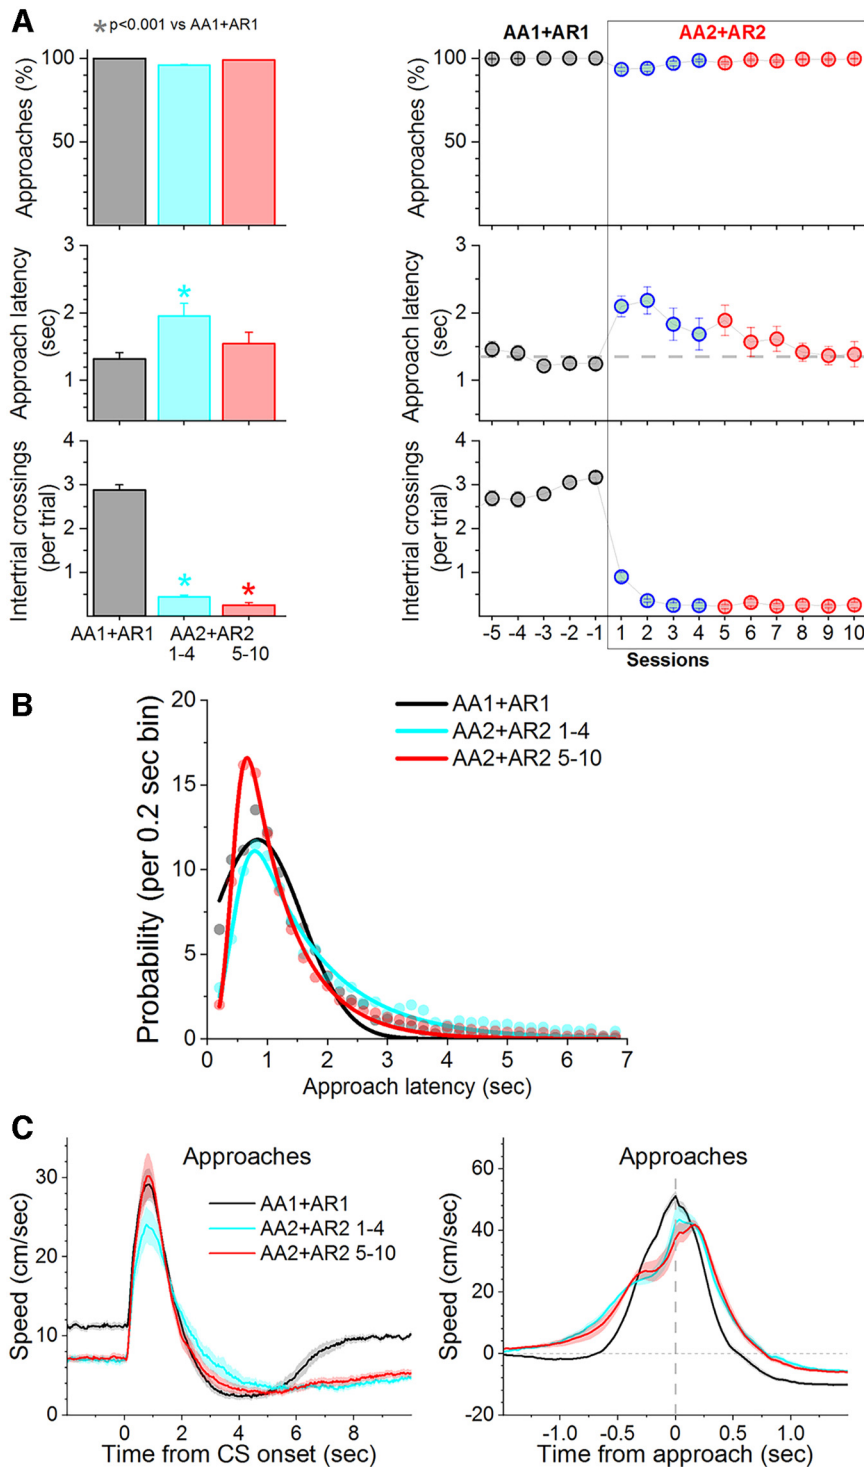


Figure 9. When the signaled action is both an active approach and an active avoidance, mice behave as they are performing active approach. Punishing the unsignaled action leads only to transient caution about producing the signaled action. **A**, Performance of a combined signaled active approach and active avoidance (AA1+AR1) in water-restricted mice. Subsequent training in AA2+AR2, which punishes ITCs, led only to brief caution about producing the signaled action (initial 1–4 AA2 sessions; blue), but thereafter (5–10 AA2 sessions; red) caution was not evident in response timing. **B**, Probability histogram (%) of avoidance latencies during AA1+AR1 and AA2+AR2 fitted with an exponential Gaussian. Note the slight rightward shift of the latencies during the initial one to four sessions but subsequently (5–10 sessions) latencies shift left producing a sharp short-latency peak. During AA2+AR2, mice behave as if they are only performing AR2 and only display caution transiently. **C**, Speed traces (mean \pm SEM) of active approach responses aligned by the CS onset (left) and baselined-corrected responses aligned by the response occurrence (right) during AA1+AR1 and AA2+AR2 procedures.

sessions, revealed no difference in approach latencies (Fig. 8A; Tukey $t_{(27)} = 0.8$, $p = 0.92$ AR1 vs AR2 5–10). However, the latency distributions were different because during AR2 the responses peaked sharply at shorter latencies denoting a leftward shift (Fig. 8B; F test, $F_{(4,60)} = 13.12$, $p < 0.0001$ AR1 vs AR2 5–10). Thus, there is no evidence of a delayed approach action when the first few AR2 sessions are excluded, which is in contrast with the delayed active avoidance action during AA2 sessions.

Video tracking revealed that mice reduced their baseline speed during the intertrial interval during AR2 compared with AR1 (Fig. 8C; Tukey $t_{(18)} = 12.13$, $p < 0.0001$ AR1 vs AR2 1–4; Tukey $t_{(18)} = 9.27$, $p < 0.0001$ AR1 vs AR2 5–10), which is similar to what occurs during AA2 compared with AA1. Upon CS presentation, mice immediately began producing the active approach response (Fig. 8C; Speed from CS onset), which contrasts with the much longer latencies of active avoidance responses. The baseline corrected peak speed (Δ Peak speed) of the response from CS onset (0- to 6-s window) was faster during AR2 compared with AR1 (Tukey $t_{(18)} = 5.09$, $p = 0.005$ AR1 vs AR2 1–4; Tukey $t_{(18)} = 9$, $p < 0.0001$ AR1 vs AR2 5–10), while the time to peak speed was not significantly different. Measurement of the baseline corrected approach response speed around the response occurrence (Fig. 8C; –1–1 s; Speed from approach occurrence) did not reveal a difference indicating that the Δ speed at the time of entering the new compartment is not different between AR1 and AR2. Interestingly, during AR2 there is an obvious indent in the speed trace only during AR2 (early and late) sessions starting at about ~ 300 ms before the response occurrence. Perusal of the videos revealed that during this period the mice briefly (~ 200 ms) slow down at the door between the compartments, as if hesitating to cross, before completing the approach response by entering the new compartment. The hesitation is evident when the speed is aligned by the response occurrence, not by the CS onset. This significantly delayed the time to peak speed during AR2 around the response occurrence (Tukey $t_{(18)} = 6.8$, $p = 0.00038$ AR1 vs AR2 1–4; Tukey $t_{(18)} = 5.6$, $p = 0.002$ AR1 vs AR2 5–10). In other words, because of the hesitation, the peak speed around response occurrence (which is not different between AR1 and AR2 at response occurrence) occurs later during AR2 than during AR1 (Fig. 8C). Intriguingly, comparison of

the approach response onset latencies revealed that during early sessions of AR2, mice delayed their onset (seconds: 1.13 ± 0.07 AR1 vs 1.48 ± 0.15 AR2 1–4; Tukey $t_{(27)} = 3.9, p = 0.04$). However, during AR2 proper, the approach response onset latencies were faster than during AR1 (seconds: 1.13 ± 0.07 AR1 vs 0.74 ± 0.04 AR2 5–10; Tukey $t_{(27)} = 4.49, p = 0.01$). In conclusion, after a few AR2 sessions, during which mice display caution by delaying their approach responses, mice show no evidence of caution reflected in their response latencies. Compared with AR1, during AR2 proper, mice begin moving earlier and run faster to produce the approach action, but the faster onset latencies and speed are not translated into shorter response latencies because mice appear to hesitate, by slowing down briefly, before crossing the door. This hesitation may be considered evidence of action caution but is not directly reflected in the response latencies.

We also trained mice ($n = 9$) in a combined signaled active avoidance and approach procedure (AA1+AR1) that is a merger of both procedures. In AA1+AR1, the interval (7 s) signaled by the CS is both an active avoidance and an active approach interval. If mice shuttle during this interval, they avoid the presentation of the US and receive water in the compartment they entered; ITCs have no consequence. Subsequently, mice were subjected to AA2+AR2, during which ITCs are punished. Interestingly, the results from this procedure (Fig. 9) recapitulated the results obtained in the AR1 and AR2 procedures, not the results obtained in the AA1 and AA2 procedures, as if the animals chose to perform active approach instead of active avoidance. Thus, comparison of active approach latencies during AA1+AR1 versus AA2+AR2, excluding the initial four AR2 sessions, revealed no difference in approach latencies (Fig. 9A; Tukey $t_{(24)} = 2.4, p = 0.3$ AA1+AR1 vs AA2+AR2 5–10). However, the latency distributions were different because during AA2+AR2 responses peaked more sharply at the shorter latencies (Fig. 9B; F test, $F_{(4,60)} = 6.8, p = 0.00013$ AA1+AR1 vs AA2+AR2 5–10), which is similar to the results obtained in the AR1 and AR2 procedures.

Video tracking revealed that during AA2+AR2 mice reduced their baseline speed during the intertrial interval compared with AA1+AR1 (Tukey $t_{(16)} = 19.8, p < 0.0001$ AA1+AR1 vs AA2+AR2 1–4; Tukey $t_{(16)} = 18.4, p < 0.0001$ AA1+AR1 vs AA2+AR2 5–10). Upon CS presentation, mice immediately began producing the approach response (Fig. 9C; Speed from CS onset) which is similar to AR1. The baseline-corrected peak speed (Δ Peak speed) of the response from CS onset (0- to 6-s window) was not faster during the early AA2+AR2 one to four sessions compared with AA1+AR1 (Tukey $t_{(16)} = 0.8, p = 0.8$), which is different from what occurs during AR2. However, this speed was faster for the AA2+AR2 5–10 sessions compared with AA1+AR1 (Tukey $t_{(16)} = 6.22, p = 0.0012$). Measurement of the baseline-corrected approach response speed around the response occurrence (Fig. 9C; -1- to 1-s window; Speed from approach occurrence) did not reveal a difference, although the time to peak speed was delayed during AA2+AR2 and the hesitation indent observed during AR2 was evident during AA2+AR2 (Fig. 9C). Thus, AA2+AR2 mostly recapitulated the results observed during AR2, not those observed during AA2,

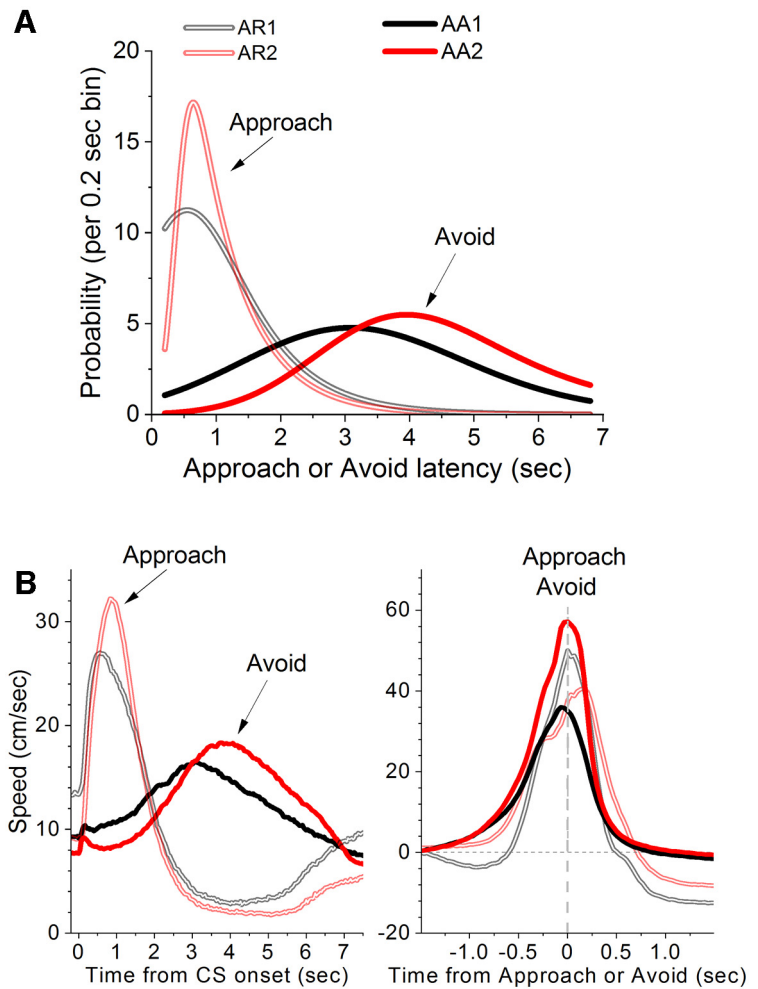



Figure 10. Comparison of signaled active avoidance and approach responses. **A**, Probability histogram (%) of response latencies for two groups of mice fitted with an exponential Gaussian. One group performed AR1 and AR2 (proper), the other group performed AA1 and AA2. **B**, Mean speed traces for the data in **A** aligned by the CS onset (left) or by the response occurrence (right; baselined-corrected).

indicating that animals selected to perform active approach, not active avoidance. Consequently, since animals seemed to be motivated by reward, not danger, caution was only transient.

Comparison of active approach and active avoidance

We employed a linear mixed-effects model to compare the behavioral measures between the mice performing active avoidance (AA1 and AA2) and active approach (AR1 and AR2) for the three periods of the tasks; first period (AA1/AR1) when ITCs are not punished, second (AA2/AR2[1–4]) and third (proper) periods (AA2/AR2[5–10]) when ITCs are punished (Fig. 10 shows the first and the third periods). Note that the second period is when caution is transient during active approach (AR2). The model revealed that response latencies (Fig. 10A) were longer during active avoidance than active approach for all three periods of the task (AA1/AR1 $t = 17.9, p < 0.0001$; AA2/AR2[1–4] $t = 20.3, p < 0.0001$; AA2/AR2[5–10] $t = 23.6, p < 0.0001$). The number of ITCs were greater for active approach than for active avoidance during the first period before ITCs were punished (AA1/AR1 $t = 11.3, p < 0.0001$), but not during the second period (initial sessions) when ITCs were punished (AA2/AR2[1–4] $t = 2.1, p = 0.2$). During the third period (later sessions), the



Movie 1. Movie of the 3D plot in Figure 5A. [View online]

number of ITCs tended to be somewhat larger during active approach compared with active avoidance (AA2/AR2[5–10] $t = 2.9$, $p = 0.03$) indicating that during AA2/AR2 proper, mice performing signaled active approach tended to produce more passive avoidance errors than mice performing signaled active avoidance. During this third period, when passive avoidance errors occurred in mice performing active approach, the percentage of approach responses were slightly larger than the percentage of avoidance responses (96.7% vs 88.9%; AA2/AR2[5–10] $t = 4.4$, $p < 0.001$). This difference did not occur in the other periods suggesting that during AR2 proper the mice performing active approach were not being cautious about their responding by trying to maximize delivery of the reward (water) at the expense of making more errors during the unsignaled period.

Finally, Figure 10B shows the absolute speeds from CS onset (Fig. 10B, left) and the baseline-corrected speeds from response occurrence (Fig. 10B, right) for the first period (AA1/AR1) and the third (proper) period (AA2/AR2[5–10]). We compared the peak speeds from response occurrence (baseline-corrected) between active avoidance and approach for the three task periods. When ITCs were not punished, approach peak speeds were faster than avoidance peak speeds (AA1/AR1 $t = 4.24$, $p < 0.0001$). However, when ITCs were punished, avoidance peak speeds were faster than approach peak speeds during the later sessions (AA2/AR2[5–10] $t = 4.7$, $p < 0.0001$), but not during early sessions (AA2/AR2[1–4] $t = 1.8$, $p < 0.6$). Thus, the faster peak speeds occur during active avoidance when ITCs are punished (i.e., AA2); this is when the mice postpone their responses but compensate by running faster to avoid. Moreover, the time to peak speeds from response occurrence were always longer during active approach than active avoidance (for all three periods of the task; $p < 0.0001$), because in active approach, the mice continue running toward the water after crossing the door, which marks the response occurrence.

Discussion

We found that during performance of a signaled action, the addition of a nonconflicting, nonoverlapping rule requiring action to be withheld when unsignaled, causes a delay in the timing of the signaled action in an apparent reflection of caution. Importantly,

caution depended on the reinforcement type motivating the action. When the signaled action was motivated by positive reinforcement (delivery of reward during active approach), caution was transient. However, when the action was motivated by negative reinforcement (omission of punishment during active avoidance), caution was persistent. Thus, the motivational state determines how the need to be cautious disrupts action timing. We also found that caution was not caused by anxiety, since an anxiolytic did not alleviate it, and it did not require frontal cortex or basal ganglia output circuits that are engaged by caution and other forms of inhibitory control in humans.

Action timing is different between avoidance and approach

Despite the similarity of the procedures, the timing of signaled active approach and signaled active avoidance responses was very different. During active approach, on CS onset, the mice respond with very short response latencies (~ 1 s) to obtain the reward as soon as possible. In contrast, during active avoidance, the mice respond with latencies that are three to four times longer compared with active approach. In basic forms of operant (e.g., fixed interval schedules of reward) or Pavlovian (e.g., eye-blink conditioning) conditioning, animals time conditioned responses to occur when the reward is available or just before the US is expected, respectively (Gallistel and Gibbon, 2000). This explains the rapid approach responses; mice respond as soon as the opportunity to obtain reward is signaled. However, the timing of signaled active avoidance response scales as a function of the interval between the CS and the US (Gallistel and Gibbon, 2000). Active avoidance responses do not peak immediately before the US (Pavlovian) nor do they occur as soon as the opportunity to avoid is signaled (operant). Instead, during our typical 7-s avoidance interval, active avoidance responses occur about half-way between the CS and US (3–4 s). The faster conditioned responses compared with basic forms of Pavlovian conditioning minimizes the possibility of failures (punishment). However, why are not these active avoidance responses even faster? After all, the opportunity for negative reinforcement during signaled active avoidance is present from the onset of the CS, just like the opportunity for positive reinforcement is present from the onset of the CS during signaled active approach. Likely, the timing of active avoidance responses reflects the confluence of Pavlovian and operant conditioning as a function of distinct motivational states imposed by the CS when it predicts a reward versus when it predicts punishment (Mowrer, 1960; Rescorla and Solomon, 1967; Dickinson and Balleine, 2002). For instance, in signaled active avoidance, the CS has traditionally been considered to elicit a central state of fear acquired via Pavlovian conditioning that can evoke a myriad of defensive reactions, some of which are incompatible with instrumental active avoidance (e.g., freezing). From an ethological perspective, the particular defensive reaction (freeze, flight, fight) depends on the spatial distance to the danger as ascertained by the subject (Blanchard and Blanchard, 1969, 1988). This defensive distance may be temporal; at intermediate temporal distances ascertained by the subject (roughly half-way between the CS and US at 7 s) is when flight

(active avoidance) occurs. Nevertheless, in well-trained animals that show little evidence of fear, response timing may also reflect the level of vigilance and arousal (Castro-Alamancos, 2004a,b). Interestingly, we found that the subject selects the motivational state that drives the action because when the procedure involved simultaneous active avoidance and active approach, the mice behaved as if they were performing active approach, not active avoidance. In a sense, the mice chose to interpret the CS as signaling the opportunity for reward instead of the imminence of danger. Consequently, mice behaved as if they were alleviating their thirst, by seeking the reward, instead of alleviating their fear by avoiding the danger.

Unsignaled passive avoidance causes action caution

A main finding is that addition of an unsignaled passive avoidance rule in a nonconflicting, nonoverlapping manner persistently delays the timing of signaled active avoidance responses, but only transiently delays the timing of signaled active approach responses. When the reinforcer is reward (positive reinforcement), punishing the unsignaled action leads to transient caution about producing the signaled action, but when the reinforcer is the omission of punishment (negative reinforcement), punishing the unsignaled action leads to persistent caution about producing the signaled action. One way to view these findings is as a situation where there is conflict between goals (Miller, 1944; Blanchard and Blanchard, 1988; Blanchard et al., 1991, 2011; Gray and McNaughton, 2000; Graeff and Zangrossi, 2002; McNaughton and Corr, 2004). For instance, conflict between a freeze-fight-flight brain system engaged by passive avoidance and a brain approach system engaged by both active avoidance and approach is purported to activate a brain inhibition system that has the primary effect of producing anxiety (Gray and McNaughton, 2000; McNaughton and Corr, 2004). In other words, addition of the unsignaled passive avoidance rule causes anxiety, which causes caution about producing the signaled action. However, we found that anxiety is not causing caution under our conditions since caution was insensitive to an anxiolytic. The prefrontal cortex is part of a neural network that represents anxiety, which can disrupt goal-directed behavior and cognitive functioning (Park and Moghaddam, 2017). Consistent with the idea of a link between prefrontal cortex and anxiety, our results show that action caution does not require either anxiety or the frontal cortex.

Although conflict between incompatible goals is a critical consideration, conflict does not formally exist in our conditions because the signaled action (active avoidance or active approach) and passive avoidance are ascribed to different periods. However, when the passive avoidance rule is unsignaled, there is likely to be more uncertainty or conflict within the subject about when the action must be withheld. This uncertainty would lead to caution about producing the signaled action. Indeed, we found that signaled passive avoidance did not lead to caution about producing the signaled action. Thus, signaling is a critical variable controlling the level of perceived conflict between the purported brain freeze-flight system for passive avoidance and the brain approach subsystem for active avoidance.

While perceived conflict (where it does not exist) may lead to caution, the perceived conflict should be identical during both the active avoidance and the active approach procedures in our conditions, but caution only becomes persistent during active avoidance (when the signaled action is motivated by danger). A simple explanation would be that the anxiety created by the

perceived conflict becomes persistent (chronic) because it interacts with the fear motivational state generated by the CS during active avoidance. However, as already noted, an anxiolytic has no effect on caution development during signaled active avoidance, when fear is purported to occur. From a brain systems perspective, if active avoidance and active approach engage the same brain system (McNaughton and Corr, 2004), it is difficult to explain how the addition of unsignaled passive avoidance leads to persistent caution during active avoidance but not during active approach. More likely, the brain subsystems engaged by active avoidance and active approach are different, which seems plausible considering the differences in the expression of these behaviors. In this scenario, the interaction of the brain inhibition system engaged by passive avoidance with each of the action subsystems would be distinct. Signaling allows the brain systems activated during active and passive avoidance to operate independently, without conflict. However, without explicit signaling these systems may conflict, perhaps because the motivational system they engage is the same. In contrast, even without signaling and after some experience, the active approach and passive avoidance systems can operate without conflict, perhaps because their motivational systems are distinct.

A form of inhibitory control, termed *proactive*, involves a preparatory step before the signaled action is triggered and engages frontal cortex and basal ganglia circuits (Aron, 2011; Meyer and Bucci, 2016; Hardung et al., 2017). A typical example, is the addition of no-go trials to go trials, which is manifested by go trial caution in the go/no-go situation (Bogacz et al., 2010). However, we found that the addition of signaled no-go trials (i.e., CS2 in AA3), does not produce caution about responding to the signaled go trials (CS1 in AA3). An important consideration is that proactive inhibitory control (go/no-go) procedures involve reward (positive reinforcement) and its omission, but we found that caution develops and becomes persistent only when punishment and its omission (negative reinforcement) are involved. Thus, in addition to the frontal cortex and basal ganglia circuits engaged by inhibitory control when reward and its omission are involved, there must be other forms of inhibitory control that engage other brain circuits, when punishment and its omission are involved. A basis for deciphering these circuits would be to consider the basic circuits that are essential for performance during signaled active avoidance (Hormigo et al., 2019, 2021a,b), the circuits engaged by passive avoidance (e.g., amygdala, hippocampus; Grossman et al., 1975; Nagel and Kemble, 1976; Gray and McNaughton, 2000; McNaughton and Corr, 2004; Levita et al., 2012), and how the latter can influence the former. A connection between these circuits could provide the distinct inhibitory control that produces caution when punishment and its omission are involved.

References

- Aron AR (2011) From reactive to proactive and selective control: developing a Richer model for stopping inappropriate responses. *Biol Psychiatry* 69:e55–e68.
- Blanchard RJ, Blanchard DC (1969) Passive and active reactions to fear-eliciting stimuli. *J Comp Physiol Psychol* 68:129–135.
- Blanchard DC, Blanchard RJ (1988) Ethoexperimental approaches to the biology of emotion. *Annu Rev Psychol* 39:43–68.
- Blanchard DC, Blanchard RJ, Rodgers RJ (1991) Risk assessment and animal models of anxiety. In: *Animal models in psychopharmacology* (Olivier B, Mos J, Slangen JL, eds). Basel: Birkhäuser.
- Blanchard DC, Griebel G, Pobbe R, Blanchard RJ (2011) Risk assessment as an evolved threat detection and analysis process. *Neurosci Biobehav Rev* 35:991–998.

- Bogacz R, Wagenmakers EJ, Forstmann BU, Nieuwenhuis S (2010) The neural basis of the speed-accuracy tradeoff. *Trends Neurosci* 33:10–16.
- Bolles RC (1970) Species-specific defense reactions and avoidance learning. *Psychol Rev* 77:32–48.
- Bolles RC, Holtz R, Dunn T, Hill W (1980) Comparisons of stimulus learning and response learning in a punishment situation. *Learn Motiv* 11:78–96.
- Botvinick MM, Cohen JD, Carter CS (2004) Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn Sci* 8:539–546.
- Boyd RL, Robinson MD, Fetterman AK (2011) Miller (1944) revisited: movement times in relation to approach and avoidance conflicts. *J Exp Soc Psychol* 47:1192–1197.
- Capuzzo G, Floresco SB (2020) Prelimbic and infralimbic prefrontal regulation of active and inhibitory avoidance and reward-seeking. *J Neurosci* 40:4773–4787.
- Castro-Alamancos MA (2004a) Absence of rapid sensory adaptation in neocortex during information processing states. *Neuron* 41:455–464.
- Castro-Alamancos MA (2004b) Dynamics of sensory thalamocortical synaptic networks during information processing states. *Prog Neurobiol* 74:213–247.
- Castro-Alamancos MA, Borrell J (1992) Facilitation and recovery of shuttle box avoidance behavior after frontal cortex lesions is induced by a contingent electrical stimulation in the ventral tegmental nucleus. *Behav Brain Res* 50:69–76.
- Church RM, Wooten CL, Matthews TJ (1970) Discriminative punishment and the conditioned emotional response. *Lear Motiv* 1:1–17.
- Cohen JD, Castro-Alamancos MA (2007) Early sensory pathways for detection of fearful conditioned stimuli: tectal and thalamic relays. *J Neurosci* 27:7762–7776.
- Cohen JD, Castro-Alamancos MA (2010) Neural correlates of active avoidance behavior in superior colliculus. *J Neurosci* 30:8502–8511.
- Dayan P, Niv Y, Seymour B, Daw ND (2006) The misbehavior of value and the discipline of the will. *Neural Netw* 19:1153–1160.
- Dickinson A, Balleine B (2002) The role of learning in the operation of motivational systems. In: *Stevens' handbook of experimental psychology*. New York: John Wiley & Sons.
- Dickinson A, Gelder M, Gray J (1980) *Contemporary animal learning theory*. Cambridge: Cambridge University Press.
- Franklin KBJ, Paxinos G (2008) *The mouse brain in stereotaxic coordinates*, Ed 3. Amsterdam; Boston: Elsevier/Academic Press.
- Gallistel CR, Gibbon J (2000) Time, rate, and conditioning. *Psychol Rev* 107:289–344.
- Gold JJ, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
- Graeff FG, Zangrossi JH (2002) Animal models of anxiety disorders. In: *Biological Psychiatry*, (D'Haenen H, den Boer, JA, Willner P eds), pp 96–103. Hoboken: John Wiley & Sons Ltd.
- Gray JA, McNaughton N (2000) *The neuropsychology of anxiety: an enquiry into the functions of the septo-hippocampal system*, Ed 2. Oxford: Oxford University Press.
- Grossman SP, Grossman L, Walsh L (1975) Functional organization of the rat amygdala with respect to avoidance behavior. *J Comp Physiol Psychol* 88:829–850.
- Guitart-Masip M, Huys QJM, Fuentemilla L, Dayan P, Duzel E, Dolan RJ (2012) Go and no-go learning in reward and punishment: interactions between affect and effect. *Neuroimage* 62:154–166.
- Hardung S, Epple R, Jäckel Z, Eriksson D, Uran C, Senn V, Gibor L, Yizhar O, Diester I (2017) A functional gradient in the rodent prefrontal cortex supports behavioral inhibition. *Curr Biol* 27:549–555.
- Heitz RP, Schall JD (2012) Neural mechanisms of speed-accuracy tradeoff. *Neuron* 76:616–628.
- Hormigo S, Vega-Flores G, Castro-Alamancos MA (2016) Basal ganglia output controls active avoidance behavior. *J Neurosci* 36:10274–10284.
- Hormigo S, Vega-Flores G, Rovira V, Castro-Alamancos MA (2019) Circuits that mediate expression of signaled active avoidance converge in the pedunculo-pontine tegmentum. *J Neurosci* 39:4576–4594.
- Hormigo S, Shanmugasundaram B, Zhou J, Castro-Alamancos MA (2021a) A signaled locomotor avoidance action is fully represented in the neural activity of the midbrain tegmentum. *J Neurosci* 41:4262–4275.
- Hormigo S, Zhou J, Chabbert D, Shanmugasundaram B, Castro-Alamancos MA (2021b) Basal ganglia output has a permissive non-driving role in a signaled locomotor action mediated by the midbrain. *J Neurosci* 41:1529–1552.
- Jercog D, Winke N, Sung K, Fernandez MM, Francioni C, Rajot D, Courtin J, Chaudun F, Jercog PE, Valerio S, Herry C (2021) Dynamical prefrontal population coding during defensive behaviours. *Nature* 595:690–694.
- Levita L, Hoskin R, Champi S (2012) Avoidance of harm and anxiety: a role for the nucleus accumbens. *Neuroimage* 62:189–198.
- Mackintosh NJ (1983) *Conditioning and associative learning*. Oxford: Clarendon Press.
- McNaughton N, Corr PJ (2004) A two-dimensional neuropsychology of defense: fear/anxiety and defensive distance. *Neurosci Biobehav Rev* 28:285–305.
- Meyer HC, Bucci DJ (2016) Neural and behavioral mechanisms of proactive and reactive inhibition. *Learn Mem* 23:504–514.
- Miller NE (1944) Experimental studies of conflict. In: *Personality and the behavioural disorders* (Huny JM, ed), pp 431–465. New York: Ronald Press.
- Millner AJ, Gershman SJ, Nock MK, den Ouden HEM (2018) Pavlovian control of escape and avoidance. *J Cogn Neurosci* 30:1379–1390.
- Mowrer OH (1960) *Learning theory and behavior*. New York: Wiley.
- Nagel JA, Kemble ED (1976) Effects of amygdaloid lesions on the performance of rats in four passive avoidance tasks. *Physiol Behav* 17:245–250.
- Pádua-Reis M, Nôga DA, Tort ABL, Blunder M (2021) Diazepam causes sedative rather than anxiolytic effects in C57BL/6J mice. *Sci Rep* 11:9335.
- Park J, Moghaddam B (2017) Impact of anxiety on prefrontal cortex encoding of cognitive flexibility. *Neuroscience* 345:193–202.
- Rescorla RA (1987) A Pavlovian analysis of goal-directed behavior. *Am Psychol* 42:119–129.
- Rescorla RA, Solomon RL (1967) Two-process learning theory: relationships between Pavlovian conditioning and instrumental learning. *Psychol Rev* 74:151–182.
- Skinner BF (1938) *The behavior of organisms: an experimental analysis*. Oxford: Appleton-Century.
- Sutton RS, Barto AG (1998) *Reinforcement learning: an introduction*. Cambridge: The MIT Press.
- Thorndike EL (1898) *Animal intelligence: an experimental study of the associative processes in animals*. *Psychol Rev* 2:i–109.
- van Maanen L, Brown SD, Eichele T, Wagenmakers EJ, Ho T, Serences J, Forstmann BU (2011) Neural correlates of trial-to-trial fluctuations in response caution. *J Neurosci* 31:17488–17495.