TIPPING THE SCALES



SOHA FARBOUD

DONDERS SERIES



PROPOSITIONS

1. Pavlovian biases and instrumental learning serve a shared goal and cooperate more closely than is commonly assumed.

2. The nucleus accumbens acts as a hub that translates motivation into action, influencing both the strength of Pavlovian biases and how we learn from the outcomes of our actions.

3. The amygdala helps determine what is salient and deserves our attention and learning, giving novel experiences a head start in shaping our behavior.

4. Ultrasound neuromodulation nudges ongoing neural processes rather than dictating behavior.

5. The effects of ultrasound neuromodulation are both trait- and state-dependent.

The freedom to write this thesis began with my parents' decision to flee Iran. Those I might have grown up alongside are still risking their lives for that same freedom. I cannot take that lightly. It is in that spirit that I offer my final, and most important, proposition:

6. Freedom comes at a cost, and is never more vulnerable than when we believe it cannot be lost.



TIPPING THE SCALES

Probing the neural substrates of
decision-making with focused ultrasound

Soha Farboud Sheshdeh

Colophon

The work described in this thesis was primarily carried out at the Donders Institute for Brain, Cognition, and Behaviour; Radboud University.

Author: Soha Farboud Sheshdeh
Cover design: Soha Farboud Sheshdeh
Layout: Stefanie van den Herik | HerikMedia | www.herikmedia.nl
Printing: ProefschriftMaken | www.proefschriftmaken.nl

ISBN: 9789462848900

Copyright 2026 ©Soha Farboud.
The Netherlands. All rights reserved. No part of this thesis may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without written permission from the author.

Tipping the Scales

Probing the neural substrates of decision-making with focused ultrasound

Proefschrift

ter verkrijging van de graad van doctor
aan de Radboud Universiteit Nijmegen
op gezag van de rector magnificus prof. dr. J.M. Sanders,
volgens besluit van het college voor promoties
in het openbaar te verdedigen op

woensdag 2 december 2026
om 12:30 uur precies

door

Soha Farboud Sheshdeh
geboren op 4 januari 1995
te Isfahan, Iran

Promotoren:
Prof. Dr. H.E.M. den Ouden
Prof. Dr. W.P. Medendorp

Copromotor:
Dr. L. Verhagen

Manuscriptcommissie:
Prof. Dr. R. Cools
Prof. Dr. K. Roelofs
Prof. Dr. E.F. Fouragnan (University of Plymouth, Verenigd Koninkrijk)

Tipping the Scales

Probing the neural substrates of decision-making with focused ultrasound

Dissertation

to obtain the degree of doctor
from Radboud University Nijmegen
on the authority of the Rector Magnificus prof. dr. J.M. Sanders,
according to the decision of the Doctorate Board
to be defended in public on

Wednesday, December 2, 2026
at 12:30 pm

by

Soha Farboud Sheshdeh
born on January 4, 1995
in Isfahan, Iran

PhD supervisors:
Prof. Dr. H.E.M. den Ouden
Prof. Dr. W.P. Medendorp

PhD co-supervisor:
Dr. L. Verhagen

Manuscript Committee:
Prof. Dr. R. Cools
Prof. Dr. K. Roelofs
Prof. Dr. E.F. Fouragnan (University of Plymouth, United Kingdom)

TABLE OF CONTENTS

PART I

Prologue 10

Chapter 1. General introduction 12

PART II

Interlude I *The Simorgh's Feather* 52

Chapter 2. Rapid modulation of choice behavior by ultrasound on the human frontal eye fields 54

PART III

Interlude II *Jamshid's Cup* 120

Chapter 3. Neural signatures of the nucleus accumbens and amygdala in motivated behavior 122

PART IV

Interlude III *Arash the Archer* 164

Chapter 4. Dissociable causal roles of the nucleus accumbens and amygdala in motivated behavior 166

PART V

Chapter 5. General discussion 228

Epilogue 274

APPENDICES

Dutch summary 278

Research data management 288

About the author 290

PhD portfolio 292

List of publications 295

Acknowledgements 296

Donders graduate school 306



PART I

Prologue
Chapter 1



PROLOGUE

Before my passion for brains, there was my passion for stories.

As a child, I was captivated by the old legends of the Greek, Roman, Norse, and the Persians. Growing up in the Netherlands, in primary school, it was mostly the Western tales that found their way to me. I sat at the edge of my seat while the teacher recounted the wanderings of Odysseus, the hero who took ten years to find his way home – not because the path was long, but because each island, each monster, each enchantress demanded something of him before he could move on. Later, in secondary school, I was even allowed to translate those centuries-old Greek texts myself.

What strikes me now, looking back, is how much these stories echo one another across cultures. Take the twelve labors of Herakles, the trials of strength and cunning imposed upon a hero who must prove himself worthy. In Persian legend, there is Rostam, a warrior of almost mythic invincibility, who faces his own seven trials: demons, deserts, sorcery, each demanding something different before he can continue. These peoples never shared a timeline, yet their stories share a structure – as if the human imagination, when reaching for something essential, keeps arriving at the same shapes. The cunning of Odysseus finds its counterpart in the wisdom of Zāl, raised by a magical bird in the mountains. Prometheus bringing fire to humanity rhymes with Jamshid, the king whose golden age brought crafts, medicine, and even the Persian New Year into being.

This is what I also came to realize during my PhD. The research we do is not so different from the research that came before us. My first study is, quite literally, a translation of animal work into humans – inspired by what others discovered, building upon their methods, hoping to add something small but meaningful. Each of us arrives at a question that feels urgent, works on it for years, and, if we are fortunate, shifts understanding by a few degrees.

The more I thought about it, the more I noticed something else: the old stories do not just mirror how we do science, they sometimes mirror the science itself.

Consider Odysseus, sailing past the Sirens. He knows their song is irresistible – beautiful, deadly – and that no amount of willpower can save him once it begins. So he does not rely on willpower. He has his crew plug their ears with beeswax and orders them to tie him to the mast. He hears the song, he strains against the ropes, and he survives. Does this not reflect what we now call Pavlovian versus instrumental control? The Sirens pull at something automatic, reflexive, hardwired; Odysseus responds with a system that plans, that anticipates, that binds the future self to a wiser course.

Or think of a gift given in Persian legend, that of a single feather, pressed into the hand of a young man before he descends from the mountains to face an uncertain world. It cannot force anything. It does not guarantee victory or control the outcome. But in the right moment, under the right conditions, it tips the scales. A nudge rather than a shove. A small, targeted intervention that makes a different future possible. You will hear this story in full later. For now, it is enough to say: is this not exactly what we hope transcranial ultrasound stimulation can be? Not a tool that commands the brain, but one that modulates it – gently, precisely – when the neural context is ready to be shifted.

The Greek parallels, most of you probably already know. They are part of the shared curriculum of the West.

What I want to offer you is something a little further afield. I want to take you, at moments throughout this dissertation, back to my own roots. To the legends of Iran. To heroes and monsters and great birds with feathers of fire. To stories that may feel unfamiliar, but that carry the same weight as the ones you grew up with and that, in their own way, illuminate the science I have tried to do.

So before each chapter, I will tell you a tale. Because the stories deserve to be heard, and because I have come to suspect that the same impulses that shaped these ancient narratives also shape how we ask questions, how we design experiments, and how we make sense of what we find.

CHAPTER 1

General introduction



RATIONALE

There is an email you have been meaning to send for days. It is not even that complicated, just slightly uncomfortable, and somehow it keeps sliding to tomorrow. Meanwhile, at a party last weekend, you watched your own hand reach into the snack bowl for the third time, as if it belonged to someone else. These small mutinies of the self are so familiar that we barely notice them. But step back, and a pattern emerges. We know what we should do about climate change, about inequality, about the fractured state of public discourse, and yet, collectively, we keep not doing it. The gap between knowing and doing is not a personal failing; it is a feature of how human brains are built[1,2]. And when this gap widens into pathology[3,4], into addiction, compulsion, or paralyzing anxiety, it becomes clear that understanding how we make decisions is not merely an academic puzzle. It may be the key to understanding why we so often act against our own interests, and how we might learn to do otherwise.

So how do we study decision-making? Here is the challenge I faced when I started my PhD: in humans, we have long lacked the tools to establish causal relationships between brain activity and behavior. We could design clever behavioral tasks to probe learning and choice. We could measure brain activity using fMRI, EEG, or MEG while people perform these tasks. Some studies used pharmacological interventions to subtly alter brain chemistry and observe the effects on behavior[5,6], while others studied clinical populations directly[7], patients receiving deep brain stimulation[8], for instance. This collective body of work has enormously advanced our understanding of human decision-making. And yet, something crucial was still missing: the ability to move beyond observation and ask causal questions.

This brings us to the central methodological challenge in human neuroscience. Most human neuroscience studies are correlational: we measure brain activity that co-occurs with behavior. But co-occurrence, i.e., correlation, does not equal causation. Ice cream sales and shark attacks both increase in summer; concluding that ice cream causes shark attacks would be absurd. Similarly, observing that a brain region's activity correlates with behavior does not tell us whether that region plays a role in *causing* that behavior. To answer the question of causation, we need to manipulate neural activity and observe how this changes behavior. Many regions implicated in decision-making, however, lie deep inside the brain, beyond the reach of conventional brain stimulation techniques. Pharmacological interventions, i.e., giving people psychoactive drugs, can reach these deeper regions, but lack spatial specificity as the drug distributes throughout the brain, allowing us to study the causal contribution of neurotransmitter systems more globally, rather than specific regions. We simply could not get precise enough.

This constraint stands in stark contrast to what is possible in animal research, where the toolkit is far more extensive: behavior, neuroimaging, pharmacology, electrophysiology, lesion studies, optogenetics. These techniques have taught us a lot about decision-making circuits. But findings from rodent studies do not always easily translate to humans, as rodent and human brains are not the same. They differ not only in scale and structure, but in the organization of prefrontal systems, the relative contribution of cortical versus subcortical circuits, and the degree to which behavior is shaped by language, culture, and abstract reasoning. Findings from rodent models therefore cannot simply be assumed to generalize to humans — they generate hypotheses that must ultimately be tested in humans. And this is where the field has long faced a hard constraint. The very techniques that allow us to ask causal questions in animals, such as lesions and optogenetics, are invasive and thus unavailable for use in healthy human participants. What we have lacked are non-invasive tools capable of targeting specific brain regions in humans and understanding their causal roles. Without such tools, human neuroscience has been largely confined to correlational observation, measuring brain activity that tracks behavior without being able to establish the causal contribution of that activity to it.

This is where transcranial ultrasound stimulation (TUS) enters the picture. TUS is a new technique that can non-invasively target deep brain structures with remarkable spatial precision; offering, for the first time, the prospect of causal investigations in humans that have long been standard in animal research[9]. TUS is still a young technique, and when I began my doctoral work, it remained unclear whether it could be used successfully to modulate human behavior.

The overarching aim of this thesis is to advance our causal understanding of ***the neural mechanisms underlying motivation-action coupling in humans***. To do so, this thesis addresses two central questions. First: Can we use TUS to bias choice behavior in humans? This establishes whether the tool works in a well-characterized behavioral paradigm and neural circuit. Second: How do deep brain structures causally shape motivational action control in humans?

FROM MOTIVATION TO ACTION

Before we can test causal relations in brain function in human decision making, we need to understand what we are actually studying. What do we mean when we talk about decision-making, and how have scientists managed to study something so central to everyday human functioning in the controlled setting of a laboratory? In the following sections, I will discuss different forms of decision-making, from sensorimotor

transformations to value-based choices, and which brain regions are thought to play a role in these. For each, I will discuss the behavioral phenomena, the neural substrates that have been implicated, and what remains to be understood, particularly in humans.

Sensorimotor choice

At its most fundamental, decision-making is about transforming sensory information into action[10]. You are walking through a park, children are playing nearby, and suddenly a ball comes flying toward you. Without deliberation, you step aside. This seemingly effortless response belies a remarkable computational feat: your visual system detected the ball, estimated its trajectory, and within a fraction of a second your motor system generated an appropriate avoidance movement. This sensory-to-motor transformation lies at the heart of all behavior, from reflexive dodging to the most deliberate choices.

The neural architecture supporting such transformations has been extensively mapped. Sensory information flows from primary sensory cortices through association areas, where it is integrated and progressively linked to motor plans[11]. The frontal eye fields (FEF), for instance, exemplify the tight coupling between perception and action. This prefrontal region is involved in the allocation of visual attention and the preparation of saccadic eye movements, specifically directing gaze toward locations in the contralateral visual field[12–15]. When the ball appears to your left, your right FEF helps rapidly orient your eyes toward it, bringing the ball into foveal view – allowing you to then track its trajectory and coordinate your response. Sensorimotor decisions like these concern the physical, spatiotemporal properties of the world: properties that can be directly observed. How fast is the ball moving? Is it coming towards me? The brain answers such questions by accumulating sensory evidence over time until sufficient information favors one response over another.

Yet sensorimotor transformations do not occur in a motivational vacuum. Your avoidance response is shaped not only by the ball's trajectory but also by the anticipated unpleasantness of getting hit. Beyond physical properties, the brain must also represent value: is this stimulus something to approach or avoid? Two systems are thought to infuse value into this sensory-to-motor pipeline. Pavlovian control automatically biases you toward avoidance in the face of potential punishment[16,17]. But now imagine you are a goalkeeper: the very same ball flying toward you should be intercepted rather than avoided, because blocking a goal is rewarding. Same stimulus, opposite required response (**Figure 1**). This is where instrumental learning comes in. Through experience, you have learned that in this context, approaching the ball leads to a positive outcome[18,19]. While a ball's trajectory follows stable physical laws,

motivational properties must often be inferred and can change: the same stimulus might be rewarding in one context and punishing in another. In the following sections, I examine each of these systems in turn.



Figure 1 | Decision-making ranges from reflexive sensorimotor transformations to value-based action selection. Left: at its most fundamental, decision-making involves transforming sensory input into motor output. Detecting an approaching ball triggers a rapid motor response directed by the physical properties of the stimulus: its trajectory, speed and direction. Right: much of human decision-making, however, is shaped by internally represented value. The same stimulus, an incoming ball, can elicit an opposing response when it carries motivational significance: for a goalkeeper, intercepting the ball means winning, and approach becomes the adaptive choice. This thesis focuses primarily on value-based decision-making, in which actions are selected not only on the basis of sensory evidence, but on the anticipated outcomes they are expected to bring.

Value-based decision-making

While sensorimotor decisions address how to respond to physical properties of stimuli, value-based decisions address whether to respond at all and, if so, with what vigor – that is, they are governed not by the physical features of a stimulus but by its motivational significance[20]. Two complementary systems govern how organisms use value information to guide action. The Pavlovian system generates seemingly hardwired responses to motivationally significant stimuli: invigoration toward rewards, inhibition from threats[1]. These responses are evolutionarily conserved, require no deliberation, and emerge reliably across individuals. The instrumental system, in contrast, supports flexible learning of behavior[18]. Through trial-and-error learning, organisms discover which specific actions produce desired outcomes in specific contexts[18,19]. While Pavlovian responses are elicited by stimuli whose value is innate or acquired through *stimulus-outcome* associations, instrumental values are action values, learned based on observing the consequences of these actions through *action-outcome* contingencies[18].

These Pavlovian and instrumental controllers of behavior are typically aligned but can also conflict (**Figure 2A**). Imagine you are deep into writing your PhD thesis, finally making progress on a difficult section, when your phone buzzes and the screen lights up with a notification. The Pavlovian system responds to this appetitive cue: heightened arousal and an urge to check the message. Perhaps it is something enjoyable, an interesting update from a friend. Meanwhile, the instrumental system, having learned that ignoring distractions improves productivity promotes restraint. When both systems align, behavior is straightforward. But when the Pavlovian pull is strong and the instrumental restraint is weak, perhaps after hours of writing, you may find yourself reaching for the phone despite knowing better. Such conflicts between Pavlovian and instrumental control pervade daily life.

Understanding how these two systems operate and interact is central to explaining both adaptive behavior and its failures[1–4]. The following sections examine each system in detail, tracing their behavioral signatures and historical study to underlying neural mechanisms. I then turn to the interplay between Pavlovian and instrumental control, examining how these systems are balanced in the brain and what happens when this balance is disrupted.

Pavlovian control

The term 'Pavlovian' originates from the name of the behavioral psychologist Ivan Pavlov in the early twentieth century. Pavlov famously demonstrated that dogs could learn to associate a neutral stimulus (such as a bell) with food, eventually eliciting salivation not just when food was presented, but to ringing of the bell alone[21]. This observation revealed a fundamental principle: organisms automatically learn associations between environmental cues and biologically significant outcomes. Crucially, the conditioned response (salivation) was not an arbitrary behavior that the dog chose to perform, but an innate, species-typical response to food that became transferred to the predictive cue. This automaticity distinguishes Pavlovian responses from instrumentally learned responses, where actions are selected based on their consequences (**Figure 2B**).

The Pavlovian system's influence extends far beyond simple reflexes like salivation. When organisms encounter a cue that predicts a biologically significant outcome, this triggers a cascade of preparatory responses that are hardwired into the nervous system[1,2,6,17,22,23]. Appetitive cues, those predicting reward, trigger increased invigoration, approach behavior, and enhanced attention. Aversive cues, those predicting punishment, elicit behavioral inhibition, withdrawal, and vigilance. The evolutionary logic is straightforward: in natural environments, approaching potential food sources while avoiding predators promotes survival and reproduction. These

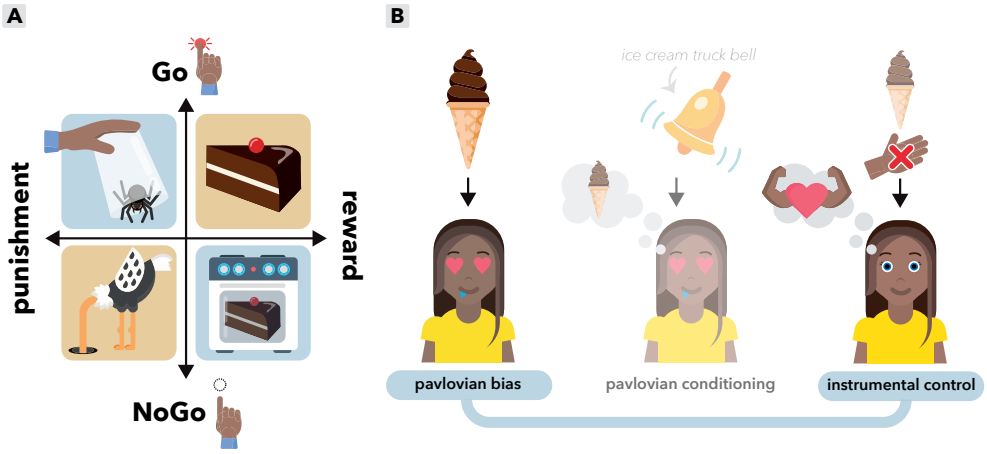


Figure 2 | Pavlovian and instrumental influences on value-based action selection. (A) The prospect of reward naturally invigorates action: seeing a piece of cake triggers an urge to reach for it. Conversely, the prospect of punishment tends to suppress action: faced with a difficult email, the tempting response is to bury one's head in the sand. These stimulus-valence-action mappings reflect Pavlovian biases, evolutionarily conserved tendencies to approach reward and withdraw from punishment. Yet these default responses are not always adaptive. Sometimes punishment calls for action: encountering a spider at home may warrant catching and removing it, lest it ends up somewhere worse. And sometimes reward calls for restraint: when the cake is still in the oven, repeatedly opening the door only slows things down. Optimal behavior thus requires the ability to override Pavlovian biases when context demands it. **(B)** Related but distinct constructs govern how motivation shapes behavior. Pavlovian influences on behavior take different forms. Pavlovian conditioning refers to learned associations between cues and motivationally significant outcomes: through experience, the sound of an ice cream truck comes to predict reward and thereby acquired motivational value. Pavlovian bias refers to the innate, hardwired responses to motivationally significant stimuli such as the mouthwatering at the sight of an ice cream cone. Both operate without explicit action-outcome reasoning, which is what distinguishes them from instrumental control. Instrumental control involves learning that a specific action leads to a specific outcome and can override Pavlovian tendencies when they conflict with goals, such as resisting the reach for ice cream because past experience has taught you it will leave you feeling worse. The present thesis focuses primarily on Pavlovian bias and instrumental control, and how these systems interact in the human brain.

response tendencies are remarkably consistent across species, from fruit flies to humans, reflecting their deep conservation across hundreds of millions of years of evolution[1,24–26]. Because these biases proved adaptive across such diverse ecological niches, they became embedded as seemingly reflexive and automatic responses that require no learning or deliberation to emerge.

A striking demonstration of the inflexibility of Pavlovian approach behavior comes from Hershberger's (1986) experiment with food-deprived baby chicks[27]. Chicks were placed in a runway containing a food cup that moved either with half or twice the speed

of the chick. When the food cup moved at half speed, they needed to move toward it to catch up. When it moved at twice the speed, the chicks needed to move away from it so the food cup would catch up with them. While the chicks readily approached the slow-moving food cup, they were completely unable to learn to move away from the fast-moving food cup, despite hundreds of trials. The sight of the rewarding food cup automatically triggered approach behavior that the chicks could not suppress. This experiment illustrates Pavlovian approach responses so deeply ingrained in behavioral programs that they can persist even when directly in conflict with obtaining desired outcomes.

The neural basis of Pavlovian control involves a network of subcortical and cortical structures (**Figure 3**). In this network, the amygdala is thought to assign motivational significance to stimuli[28,29]. Lesion and neuroimaging studies demonstrate that the amygdala responds both to learned reward and punishment predictors and to novel stimuli of uncertain value[30–33]. This dual sensitivity positions the amygdala as a detector of motivational relevance, whether that relevance is established through learning or merely potential due to novelty. The basolateral amygdala projects to the nucleus accumbens (NAcc), providing a pathway through which stimulus value information can influence striatal activity and modulate behavioral activation[34,35]. The ventromedial prefrontal cortex (vmPFC) then is thought to represent the subjective value of stimuli in a "common currency", integrating information across different reward types and modalities; activity in this region scales with the anticipated value of cues[20,36].

Critically, the motivational value signals – appetitive and aversive significance assigned by the amygdala and represented in vmPFC – must be translated into action. The NAcc is thought to serve as the key hub for this translation, sitting at the interface between motivation and action (**Box 1**). Taylor and Robbins (1984) showed that increasing dopamine in the NAcc enhanced behavioral activation elicited by reward cues[37]. Microinjections of d-amphetamine amplified the invigorating effects of conditioned reinforcers, while this enhancement was attenuated after lesioning of the NAcc. These findings demonstrate dopamine's role in coupling reward prediction to behavioral activation, the neural instantiation of appetitive Pavlovian bias. Importantly, however, the NAcc does not communicate directly with motor output systems; rather, motivational signals are translated into behavioral tendencies through its primary outputs to the ventral pallidum and onward through the mediodorsal thalamus back to vmPFC and orbitofrontal cortex (OFC), forming a cortico-striato-thalamo-cortical loop[38]. Within this loop, dopaminergic signaling modulates the balance between direct and indirect pathway activity in the NAcc: dopamine release to reward-predicting cues potentiates D1-receptor-expressing medium spiny neurons, promoting behavioral activation, while dopamine dips to aversive cues shift this balance toward D2-receptor-

mediated suppression of action[1,39,40]. This architecture provides a neural substrate for the Pavlovian bias: appetitive cues engage Go circuitry, aversive cues disinhibit NoGo circuitry – but both do so by modulating downstream cortical and subcortical targets through the ventral basal ganglia loop, not by directly driving motor output.

Instrumental control

While Pavlovian responses are automatic and stimulus-driven, instrumental learning allows organisms to work out which specific actions produce desired outcomes. The study of instrumental learning traces back to Edward Thorndike's puzzle box experiments in the 1890s[41]. Thorndike observed that cats confined in a box would initially produce random behaviors until they accidentally activated the escape mechanism. With repeated trials, successful actions occurred more quickly while irrelevant behaviors dropped out. Thorndike formalized this as the Law of Effect: actions followed by satisfying consequences become more likely to recur, while actions followed by discomfort become less likely. Importantly, Thorndike's cats were learning that specific actions produced specific outcomes, the hallmark of instrumental control.

But how does instrumental learning work at a computational level? Modern accounts formalize instrumental learning through prediction errors[42]. Rather than simply strengthening actions that are followed by rewards, agents are thought to update beliefs about the value of a particular action (or stimulus) using the difference of actual and expected outcomes. Let's give an example. Suppose you visit your regular coffee place around the corner on your way to work. You order a coffee, and since you have usually enjoyed it, you have high expectations. But the coffee you get turns out to be disappointing. A negative prediction error occurs when the outcome is worse than your expectations which results in weakening the strength of the preceding action ("I will not go there next time"). A positive prediction error occurs when the outcome exceeds expectations and strengthens the value of the preceding action. Maybe the pastry that you ordered with your coffee was even more delicious than you thought it would be ("Will definitely order this again next time"). This framework, proposed by Robert A. Rescorla & Allan R. Wagner[42], explains both acquisition and extinction: once an action no longer produces expected rewards, negative prediction errors drive unlearning.

The question then arises: how does your brain encode these prediction errors? Correlational evidence came from Schultz, Dayan and Montague, who recorded from dopamine neurons in the ventral tegmental area (VTA) and substantia nigra (SN) of macaque monkeys and found that their firing patters looked strikingly like what a prediction error signal would theoretically look like: vigorous firing to unexpected

rewards, and a pause in firing precisely when an expected reward was omitted[43]. This was a landmark discovery – but it remained correlational. It took nearly four decades before causal evidence arrived when Steinberg and colleagues (2013) used optogenetic stimulation of dopamine neurons in mice to show that artificially driving dopamine activity at the time of an outcome was sufficient to drive learning, establishing a causal link between dopaminergic firing and behavioral updating[44].

These dopaminergic prediction error signals from the VTA/SN project to the striatum, which implements action selection based on learned values[43,45] (**Box 1**). The striatum does not operate in isolation, however. The prefrontal cortex – in particular vmPFC, OFC, and anterior cingulate – represents the expected value of outcomes independent of specific actions, allowing values to be used flexibly across contexts[36,46]. Together, these regions form an integrated circuit for learning from experience.

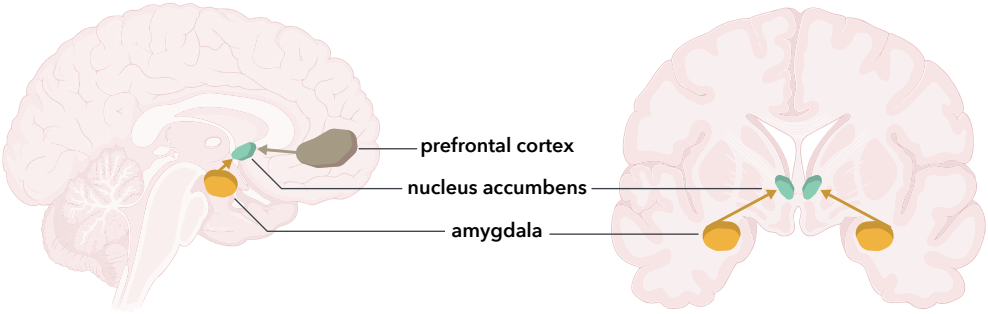


Figure 3 | Key neural substrates of motivational action control. The ventral striatum (nucleus accumbens, NAcc) sits at the interface between motivation and action, receiving value signals from two major inputs: the basolateral amygdala, which assigns motivational significance to stimuli, and the prefrontal cortex – specifically the orbitofrontal cortex (OFC) and ventromedial prefrontal cortex (vmPFC), which represent stimulus value in a common currency. These projections allow both Pavlovian cue values and instrumentally learned action values to shape behavioral activation or suppression via dopaminergic modulation of the NAcc. The present thesis directly targets the NAcc and amygdala using transcranial ultrasound stimulation (TUS) to establish their causal contributions to Pavlovian and instrumental control in humans.

The interplay of Pavlovian and instrumental control

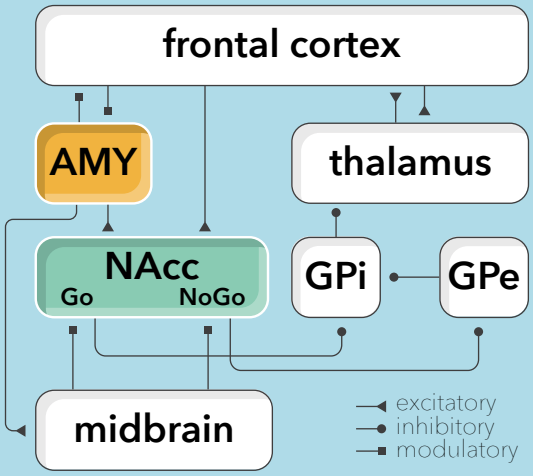
Picture again the PhD student whose phone has just lit up with a notification. The Pavlovian system, responding to this appetitive cue, generates approach tendencies and arousal: the urge to check what could be an interesting message. The instrumental system, drawing on learned consequences of giving in to distractions (lost time, reduced focus, delayed deadlines), promotes restraint. Whether the student reaches

for the phone depends on the relative strength of these competing influences. Real-world behavior rarely reflects pure Pavlovian or instrumental control; rather, it emerges from their dynamic interplay. On some occasions, resisting the impulse feels effortless; on others, perhaps when tired or stressed, the same impulse proves nearly impossible to overcome, even though nothing about the situation has changed and the student "knows better". This variability reveals something important: our actions are not the product of a single system operating in isolation but emerge from the continuous interplay of Pavlovian and instrumental influences – with both systems shaping whether an action is invigorated, or suppressed, all the way down to the motor response.

One paradigm that has been central to studying this interplay is Pavlovian-to-instrumental transfer (PIT). PIT experiments typically involve three stages: first learning an instrumental contingency (e.g., pressing a lever produces a reward), then separately undergoing Pavlovian conditioning (a cue is paired with a motivationally significant outcome), and finally perform the instrumental task while the Pavlovian cues are presented in the background[25,35,49,50]. Crucially, the Pavlovian cues are task-irrelevant – there is no reason to attend to them instrumentally – yet they nonetheless systematically bias behavior[17,51]. Appetitive Pavlovian cues invigorate instrumental responding, while aversive cues suppress it[17,35,52]. Moreover, these effects are action-specific: aversive cues suppress approach actions but can invigorate withdrawal actions, and vice versa for appetitive cues – demonstrating that Pavlovian cues interact with the motivational direction of the instrumental action, not just its overall vigor[16,17]. This body of work, spanning decades and multiple species, established that Pavlovian and instrumental systems are not independent modules but continuously interact, with Pavlovian cue values directly biasing the vigor and direction of ongoing instrumental behavior.

However, PIT paradigms also have a fundamental limitation when it comes to dissecting the computational contributions of each system. Because Pavlovian and instrumental training occur separately, the PIT effect captures the *transfer* of Pavlovian value onto instrumental responding – but the two learning phases are necessarily independent. In many real-world situations, and in the laboratory tasks most directly relevant to psychiatric conditions, Pavlovian and instrumental contingencies are interleaved: the same cue simultaneously predicts an outcome *and* signals what action is required. In such situations, where learning and responding occur together, it becomes difficult to ask which aspects of behavior reflect Pavlovian biases and which reflect instrumental learning, because the two systems are never fully separated. To address this, Guitart-Masip and colleagues, building on theoretical ideas by Cools, Nakamura and Daw (2011)[53], developed the motivational Go/NoGo task[54–57]. Rather than separating

Box 1 | The basal ganglia circuitry for motivated action



The NAcc occupies a key position at the interface of motivation and action. It receives converging inputs from regions encoding value, context, and motivational salience – including the prefrontal regions (orbitofrontal cortex, anterior cingulate cortex), amygdala, hippocampus and midbrain dopamine neurons – placing it at a junction where diverse signals can be integrated into a net motivational drive[34,38]. The amygdala – which maintains reciprocal connections with prefrontal regions – plays an important role in assigning motivational significance to stimuli and conveys this to the NAcc either directly (from the basolateral nucleus) or indirectly (from the central nucleus)[34,47]. The motivational drive is coupled to action through the striatal direct and indirect pathways, Reward-predictive cues elicit dopamine peaks in the NAcc, potentiating D1-receptor-expressing medium spiny neurons of the direct ('Go') pathway, which promotes behavioral activation through disinhibition of thalamic and cortical motor regions[39,40]. Conversely, aversive cues elicit dopamine dips, shifting the balance toward D2-receptor-expressing neurons of the indirect ('NoGo') pathway, which suppresses behavioral activation[48]. Critically, the NAcc does not communicate directly with motor output; rather, motivational signals are relayed via the ventral pallidum and mediodorsal thalamus back to prefrontal cortex, forming a cortico-striato-thalamo-cortical loop[38]. The same circuitry also supports instrumental learning. Dopamine bursts following positive prediction errors potentiate cortico-striatal synapses in the Go pathway, reinforcing recently performed actions, while dopamine dips following negative prediction errors potentiate the NoGo pathway, promoting avoidance[40]. Through the striato-nigro-striatal "ascending spiral", ventral striatal activity can further influence more dorsal striatal regions associated with cognitive and motor control, providing a route by which motivational context shapes action selection more broadly[38,47].

Pavlovian and instrumental phases, this task orthogonalizes them within a single learning context: cue valence (win versus avoid loss) is crossed with required action (Go versus NoGo), allowing the Pavlovian bias and instrumental learning to be teased apart within the same paradigm.

In humans, this interplay is nowadays most commonly studied using the motivational Go/NoGo task. This task systematically probes the interplay between Pavlovian and instrumental systems by orthogonalizing action requirements and outcome valence, which allows Pavlovian biases and instrumental learning to be dissociated from one another in a way that PIT designs cannot fully achieve. Participants must learn, for different cues, whether to invigorate (Go) or withhold responding (NoGo) to obtain rewards or avoid punishments. This yields four conditions: Go-to-Win, NoGo-to-Avoid, NoGo-to-Win, and Go-to-Avoid. Results consistently show a Pavlovian bias: participants find it easier when the required action aligns with the Pavlovian tendency – approaching rewards and withholding from punishments – and harder in conflict conditions, even when they fully understand the contingencies. This asymmetry reveals the limits of instrumental flexibility when opposing Pavlovian tendencies, much like Hershberger's chicks could not override their approach behavior, and much like the PhD student struggles to ignore the phone when the Pavlovian pull is strong.

While for most PhD students, ignoring the phone is merely an occasional struggle, an imbalance between Pavlovian and instrumental control is associated with serious clinical implications in certain disorders. In addiction, for example, Pavlovian cues associated with drug use exert excessive control over behavior, overriding instrumental knowledge that continued use is harmful[3,58]. Drug-associated environments, paraphernalia, and social contexts trigger powerful approach tendencies that recovering individuals struggle to suppress, reflecting a NoGo-to-Win failure: the person knows they should withhold the response, but Pavlovian approach overwhelms instrumental restraint. This is precisely the struggle of a recovering gambler, whose instrumental system *knows* gambling leads to ruin, but whose Pavlovian system responds powerfully to casino cues. In anxiety disorders and depression, the opposite pattern may emerge. Excessive Pavlovian inhibition in response to potentially threatening stimuli can manifest as difficulty with Go-to-Avoid conditions: the person struggles to take action even when doing so would prevent negative outcomes. Social anxiety, for instance, involves excessive behavioral inhibition in social situations, even when engagement would ultimately be rewarding[59]. Depression similarly involves reduced approach behavior, with patients struggling to initiate actions toward rewards[60]. Critically, excessive freezing or withdrawal in the face of potentially threatening environments hinders exploration, preventing the detection of new opportunities, social connections, or sources of reward. The resulting impoverishment of positive experiences may then

maintain or worsen the depressive state. In obsessive-compulsive disorder (OCD), the balance between Go and NoGo pathways may be disrupted differently[61]. Compulsive behaviors persist despite negative consequences and despite the person's explicit knowledge that the behaviors are irrational, suggesting a failure of instrumental control over repetitive action patterns. This may reflect dysfunction in the cortico-striatal circuits that normally allow flexible updating of action values based on outcomes[62,63]. The compulsions continue not because they are rewarding, but because the normal mechanisms for suppressing unrewarded actions are impaired.

Understanding how Pavlovian and instrumental systems interact, and how their balance can be disrupted, thus has direct relevance for understanding psychopathology. However, in humans it remains unclear precisely which brain regions contribute causally to these systems, and in particular what role the NAcc plays in Pavlovian versus instrumental control. Much of the evidence for the neural mechanisms described above comes from animal research involving lesions[37], and single-unit recordings[43], techniques that cannot be applied in healthy humans. The human neuroimaging literature[22,23,45,56,57], while extensive, is fundamentally correlational: observing that a brain region activates during a task does not establish that this region is causally implicated in behavior. This gap between correlation and causation represents a significant limitation in our understanding. To develop more targeted and effective interventions, we need to establish the causal roles of specific brain regions in human motivational control. This requires non-invasive methods capable of directly manipulating neural activity in a spatially precise manner.

TRANSCRANIAL ULTRASOUND STIMULATION: A NEW CAUSAL TOOL

Transcranial ultrasound stimulation (TUS) is a neuromodulation technique with unprecedented spatial precision, capable of targeting even deep brain structures non-invasively. The method uses ultrasound waves – the same type of sound waves that have long been used in diagnostic imaging. The appealing property of sound waves is that they can travel through tissue, which is why they can be used to visualize what we otherwise cannot see. Now, rather than listening to their echoes to create an image, we focus them to impact their targets. Moving from ultrasound imaging to ultrasound intervention. By doing so, we create a focal point of highest intensity that can be placed virtually anywhere in the brain, with millimeter precision.

From ripples to remote control

The story of focused ultrasound in neuroscience took an important turn in the 1950s

with the Fry brothers – an ironic and iconic name, given what high-intensity ultrasound can do to tissue. They successfully ablated subcortical structures by focusing ultrasound waves without damaging surrounding tissue, demonstrating remarkable spatial precision. While their work required craniotomies (i.e., opening of the skull) to achieve adequate ultrasound transmission, they went on to show the clinical efficacy of ablative ultrasound in patients with Parkinson's disease[64,65]. As methods advanced, researchers eventually achieved transcranial focusing, eliminating the need for skull removal.

Ultrasound intensity exists on a spectrum. At its highest, ultrasound can ablate tissue – as the Fry brothers showed. Reducing intensity, much like turning down the volume on music, brings us to the regime of low-intensity focused ultrasound, which has two distinct applications. The first is transient open the blood-brain barrier, enabling spatially targeted drug delivery to specific brain regions[66,67]. This is not achieved through intensity alone, but by combining ultrasound with an intravenous injection of microbubbles: these bubbles, foreign to the body, oscillate at the ultrasound frequency and exert mechanical pressure on the blood-brain barrier, causing it to transiently open. While not the focus of this thesis, this application holds considerable promise for improving the precision of neurotherapeutic interventions. The second application – and the one I work with in this dissertation – uses similar intensities but without microbubbles, directly modulating brain activity non-invasively, reversibly, and with high spatial specificity[9,68].

In the late 2000s, the fundamental question was whether low-intensity ultrasound could influence neurons at all without damaging them through heating. Tyler and colleagues (2008) demonstrated in in vitro and ex vivo hippocampal slice cultures that ultrasound could indeed alter neuronal network activity[69]. This was then shown in vivo preparations by Tufail et al. (2010), who stimulated mouse motor cortex through the intact skull and elicited motor responses, all without significant tissue heating[70]. Yoo and colleagues (2011) extended this to rabbits with MRI-guided targeting, demonstrating region-specific effects and showing that both excitation and suppression of neural activity were possible, depending on stimulation parameters[71]. These early studies highlighted the importance of precise targeting and careful parameter selection.

Over the following decade, the field moved toward more complex behavioral paradigms and larger animal models. Low-intensity TUS targeting the FEF was shown to modulate saccadic eye movements in macaques[12], a finding later replicated by Kubanek and colleagues (2020)[13]. Researchers also began targeting deeper brain structures: Folloni and colleagues (2019) showed that stimulating the amygdala altered functional

connectivity patterns[72], while work by Khalighinejad et al. (2021) demonstrated effects on more complex decision-making behavior in macaques[73]. These primate studies suggested that TUS could be a viable tool for probing the causal role of specific brain regions in cognition.

In parallel, enthusiasm built up for human applications. Legon and colleagues (2014) targeted primary somatosensory cortex in healthy volunteers and reported enhanced performance on sensory discrimination tasks, along with changes in somatosensory evoked potentials recorded with EEG[74]. Fomenko et al. (2020) took a different approach, targeting primary motor cortex and measuring changes in cortical excitability using transcranial magnetic stimulation (TMS) as a readout; they found that TUS reduced motor evoked potential amplitudes, suggesting suppression of corticospinal excitability[75]. As the number of published studies grew rapidly, however, so did methodological concerns. Kop and colleagues demonstrated that several reported effects in early human TUS studies could be attributed to auditory and somatosensory side effects – artifacts arising from the sound and vibration produced by the transducer – rather than direct neuromodulation[76,77]. This work underscored the critical importance of appropriate control conditions and blinding procedures in TUS research.

Fortunately, the field has since responded. At the start of my PhD in 2021, it was still very much an open question whether TUS could effectively modulate human cognition, or whether the effects previously reported were driven by confounds rather than genuine neuromodulation. Indeed, establishing this was one of the aims of my doctoral work. At the time of writing this thesis in 2026, we are now at a point where we can be reasonably confident that TUS works – though that confidence was hard won, and the work presented in this thesis has contributed to building it.

When sound becomes signal: biophysical mechanisms of ultrasound neuromodulation

Focused ultrasound can modulate neural activity through mechanical forces that interact with neuronal tissue in several ways. At higher intensities, ultrasound can produce significant thermal and cavitation effects (the formation and collapse of gas bubbles). However, the low intensities used in neuromodulation make these mechanisms unlikely to play a significant role[68,78]. Instead, the primary mode of action of neuromodulation by low intensity ultrasound appears to be mechanical. This was elegantly demonstrated in C. elegans studies, in which Kubanek and colleagues (2018) showed that ultrasound-evoked behavioral responses were abolished in worms lacking mechanosensitive ion channels, while worms with impaired thermoreceptive

neurons responded normally – directly implicating mechanosensitive pathways rather than thermal ones[79].

Ultrasound can exert mechanical effects on tissue through two complementary physical phenomena[80,81]. First, as an ultrasound wave is a mechanical wave by definition, its propagation through tissue causes local particle displacement: the cyclical back-and-forth movement of tissue elements as compression and rarefaction phases pass through. This oscillatory motion generates normal strain along the axis of propagation, with secondary shear strain caused by relative displacement perpendicular to the direction of propagation. Second, as ultrasound waves propagate through tissue, they also transfer momentum, generating what is known as the acoustic radiation force (ARF). Unlike particle displacement, ARF is a unidirectional net force along the axis of propagation that produces bulk tissue displacement, with shear strain dominant in the regions adjacent to the focal point. The strains induced by both mechanisms, though small in magnitude, are sufficient to activate mechanosensitive cellular machinery.

Brain tissue is particularly susceptible to these mechanical perturbations because it is among the softest tissues in the body[82,83]. Neural tissue has elastic moduli – a measure of how stiff a material is – orders of magnitude lower than bone, muscle, or even arteries, meaning that even modest mechanical forces can produce measurable deformations[83]. Moreover, the brain already employs mechanical signaling in normal function: synaptic activity produces rapid mechanical movements of dendritic spines[84], axons exhibit swelling and shortening during action potentials[85,86], and mechanical membrane waves propagate along neuronal membranes during firing[87,88]. This inherent mechanobiology means that the brain is equipped with molecular machinery for detecting, responding to, and being modulated by mechanical forces, providing the foundation for ultrasound's effects on neural tissue (**Figure 4A**).

Mechanosensitive and mechanomodulated ion-channels are abundantly expressed throughout the human central nervous system: channels that open, close, or modulate their gating in response to mechanical forces. These channels often evolved to detect physical stimuli such as pressure, stretch, and vibration, and they happen to be sensitive to the mechanical forces generated by focused ultrasound (**Figure 4B**). Several classes have been implicated as mediators of ultrasound effects, including transient receptor potential (TRP) channels, two-pore domain potassium (K2P) channels such as TREK1, TREK2, and TRAAK, and Piezo (Piezo1) channels[69,89,90] (**Figure 4C**). Perhaps more surprisingly, even voltage-gated channels classically understood to respond to changes in transmembrane potential, are sensitive to mechanical perturbations, including voltage-gated potassium (Kv), sodium (Nav), and calcium (Cav) channels[69,91]. Beyond channels that are directly mechanosensitive, many ion channels are anchored

to the cytoskeleton, allowing ultrasound-induced tension to influence channel gating indirectly[92] (**Figure 4B**). Notably, even in the absence of gross mechanical displacement or strain, the oscillating pressure waves of ultrasound may modulate ion channel gating directly through hydrostatic pressure changes – suggesting that mechanical flexibility of the cell is not a prerequisite for ultrasound sensitivity.

The neuronal membrane itself may also be a target of ultrasound (**Figure 4B**). Applying mechanical pressure can alter membrane capacitance by inducing displacement and reconfiguration of constituent phospholipids[92,93]. Changes in membrane capacitance affect how action potentials propagate: ultrasound could either facilitate action potential generation by bringing the membrane closer to threshold or inhibit it by pushing the membrane further from threshold. Mechanical strain can also modify the distribution of electrical dipoles on the membrane surface, a phenomenon known as direct flexoelectricity, which results in changes to membrane potential and current flow.

Neurons are not the only cells that respond to ultrasound stimulation. Astrocytes, the most abundant glial cell type in the brain, ensheathe synapses, regulate extracellular ion concentrations, and modulate synaptic transmission[94,95]. Importantly, astrocytes possess mechanosensitive calcium channels of their own[96]. When ultrasound activates these channels, it triggers intracellular calcium signaling cascades that can lead to gliotransmitter release and subsequent activation of neuronal receptors[97]. Given the extensive role of astrocytes in regulating synaptic plasticity and coordinating activity across neuronal populations, ultrasound effects on glia may contribute substantially to observed neuromodulatory effects, particularly those involving longer-lasting changes in circuit function.

Timing is key: studying human cognition using ultrasound

Having established the biophysical mechanisms by which TUS interacts with neural tissue, we can now consider how these mechanisms translate into different modes of neuromodulation. A crucial distinction exists between protocols that modulate neural *activity* during stimulation itself – referred to as online effects – and protocols that modulate neural *plasticity*, inducing longer-lasting changes that persist after stimulation has ended – referred to as offline effects. To understand these different modes of action, it helps to distinguish between neural activity and neural plasticity. Activity reflects the current state of the system: which neurons are firing right now. Plasticity governs how that system changes over time: how activity patterns, encoding, and computation are updated by experience. Online and offline protocols target these different processes, and each comes with distinct advantages and challenges.

Online TUS protocols are designed to modulate neural activity during stimulation. These protocols are typically brief, ranging from a few hundred milliseconds to several seconds, and often use relatively high pulse repetition frequencies (**Figure 5**)[13,76,98,99]. While stimulation is ongoing, the mechanical forces of ultrasound can modulate mechanosensitive ion channels in the targeted region, altering membrane potentials and modulating the probability that neurons will fire. This allows TUS to nudge ongoing neural processes in a particular direction. Crucially, however, such modulation of activity is transient: once stimulation stops, the mechanical perturbation ends and the system returns to its baseline state, with the modulatory influence dissipating rapidly.

The temporal precision of online protocols makes them ideal for probing the role of a brain region at specific moments within a behavioral paradigm. First, within a single session, one can determine whether a particular brain region is causally involved in a behavior by comparing trials with and without stimulation, or by alternating stimulation between the target region and a control site not involved in that behavior (**Figure 5B**). In **Chapter 2**, I demonstrate this approach by alternating TUS between left and right frontal eye fields across trials: because each hemisphere promotes contralateral saccades, this design allowed us to establish spatial specificity of TUS effects within the same session. Second, online protocols allow researchers to dissect which temporal component of a cognitive process is influenced by a brain region. The behavior we measure, such as a button press in a decision-making task, is typically the integrated result of multiple underlying processes: cue processing, action selection, motor execution, and processing of the resulting outcome for updating beliefs. By delivering TUS at different time points within a trial, for instance, during cue versus outcome presentation, one can determine precisely when a brain region contributes to the measured behavior. I will discuss this point in greater detail in the **General Discussion**.

Online protocols come with the significant challenge of potential sensory confounds, however. TUS produces auditory and sometimes tactile sensations. If participants hear or feel something at a specific moment on every trial, this sensory input could influence their behavior independently of any neural effects. This is precisely what Kop and colleagues demonstrated in their critical examination of early human TUS studies[76]. In this study, the presentation of the sound of TUS alone could already induce reductions in motor-evoked potentials attributed to TUS to the motor cortex. For online protocols, rigorous control conditions are essential to disentangle genuine neuromodulatory effects from sensory artifacts[76,77].

Offline TUS protocols are explicitly designed to induce lasting changes in neural function. These protocols use much longer stimulation durations, typically several

minutes, but with often lower pulse repetition frequencies, effectively spreading the total energy delivery over an extended period[100–102]. Rather than modulating activity instantaneously, offline protocols aim to trigger plasticity mechanisms that alter circuit function beyond the stimulation duration.

The mechanisms by which TUS might induce plasticity likely involve the same pathways that underlie learning and memory. The most familiar form is synaptic plasticity, where the strength of a connection between two neurons is modified[103]. In long-term potentiation, synaptic efficacy is persistently increased following specific patterns of activity; in long-term depression, it is decreased[104,105]. By delivering periodic perturbations to a circuit over minutes, TUS may trigger cascades of intracellular signaling, particularly through calcium-dependent pathways[90], that culminate in lasting changes to synaptic strength. The repeated perturbations might signal to the system that something relevant is happening, prompting either an upregulation or downregulation of synaptic efficacy depending on the parameters and context[106,107].

The temporal precision of offline protocols is inherently lower than that of online protocols. This has important implications for experimental design. To determine whether TUS of a brain region influences behavior, offline paradigms require a separate control session (with sham stimulation or active control condition), whereas online paradigms can compare stimulated and non-stimulated trials within the same session (**Figure 5B**). Furthermore, offline protocols cannot dissect which specific component of a cognitive process is modulated by the target region. Any behavioral change, whether in eye movements, button presses, or reaction times, reflects a general modulation of the output, but it remains unclear whether this arose from altered preparatory processes, cue processing, decision computation, or motor execution. The stimulation enhances or diminishes the overall contribution of the target region, but the precise temporal locus of this effect cannot be isolated.

Despite these limitations, offline protocols offer several practical and conceptual advantages. First, because stimulation occurs before behavioral testing, auditory and somatosensory confounds are less likely to contaminate outcome measures. Offline protocols also provide greater experimental flexibility: while online TUS requires the transducer to remain mounted throughout testing, offline TUS allows the transducer to be removed after stimulation, freeing participants for a wider range of physical motion or to undergo neuroimaging without the physical constraints of the stimulation setup. Lastly, because offline protocols interact with the biomolecular mechanisms that underpin neuronal plasticity, they might be conceptually well-suited to study the cognitive process that relies on said plasticity: learning.

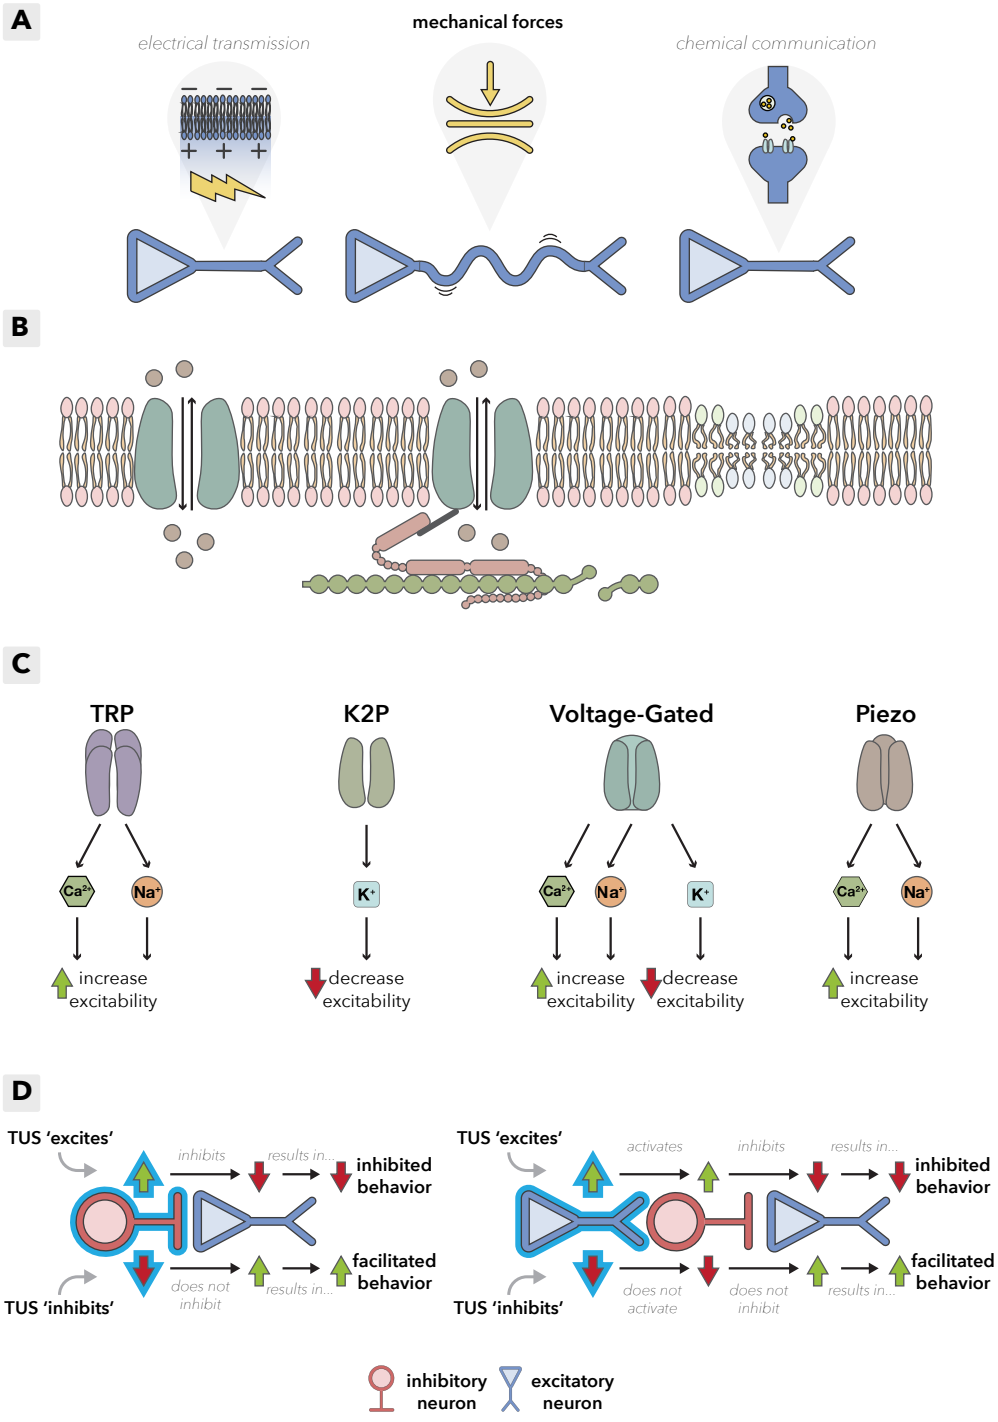


Figure 4 | Proposed biophysical mechanisms of transcranial ultrasound stimulation (TUS). (A) Neurons are physically flexible. Their primary means of communication are electrical and chemical, but neurons also possess mechanical properties (a deformable lipid bilayer, a tension-bearing

cytoskeleton, and mechanically sensitive membrane proteins) that render them susceptible to physical perturbation and pressure. This mechanical flexibility provides a plausible route through which acoustic pressure waves, as delivered by TUS, can influence neural signaling. **(B)** Multiple biophysical mechanisms have been proposed to explain how TUS affects neurons. These include direct activation of mechanosensitive ion channels, indirect mechanosensing via tension propagated through the cytoskeleton (to which ion channels are anchored, even when not themselves mechanosensitive), and changes in membrane capacitance arising from acoustic pressure deforming the lipid bilayer. **(C)** Several ion channel families have been hypothesized as candidate mediators of TUS effects, including TRP, K2P, voltage-gated (K^+, Na^+ and Ca^{2+}), and Piezo channels. These channels differ in their ionic targets and consequently in their effect on neuronal excitability: while Ca^{2+} and Na^+ influx tends to increase excitability, K^+ efflux decreases it (although this does not automatically translate to net inhibition, as suppression of background noise may enhance signal efficacy). **(D)** This mechanistic complexity means that TUS protocols cannot be straightforwardly labeled as simply "facilitatory" or "inhibitory". Because ultrasound does not stimulate neurons in isolation, the behavioral outcome depends on the type of neuron targeted (e.g., excitatory or inhibitory), the channels expressed, and the broader network context. Exciting an inhibitory interneuron, for instance, may ultimately suppress downstream activity. TUS effects are therefore best characterized in terms of net facilitatory or inhibitory outcomes at the circuit or behavioral level, rather than as intrinsic properties of a given stimulation protocol.

Importantly, offline protocols have clear clinical implications. Many neuropsychiatric disorders are proposed as disorders of plasticity, conditions in which aberrant learning has locked the system into an inflexible state[108]. For instance, in addiction, specific cues become tightly bound to powerful and automatic responses: craving, seeking, relapse[58,109]. What once reflected adaptive learning turns into maladaptive, self-perpetuating circuitry. In post-traumatic stress disorder, a cue can trigger an overwhelming stress response[110,111]. In chronic pain, neutral sensory input can become linked to persistent pain perception[112,113]. In each case, the underlying circuit has lost its capacity to update appropriately. This has direct consequences for treatment: if the underlying circuit cannot change, the effects of therapy will not last, no matter how much exposure, training, or repetition the patient partakes in. The goal of plasticity-focused TUS is to intervene at this fundamental level: not merely to modulate neural activity in the moment, but to restore the brain's capacity to adapt. In **Chapter 4**, I show how offline TUS can alter both Pavlovian biases and instrumental learning dynamics, suggesting that TUS modulates the plasticity of motivational circuits rather than steering behavior in a fixed direction. By opening a window of plasticity, TUS could allow behavioral therapies to take hold in circuits that would otherwise resist change. I will return to these clinical possibilities in the discussion.

Online and offline TUS protocols thus represent complementary tools in the neuromodulation toolkit, each suited to different scientific and clinical questions. Online protocols offer temporal precision, enabling researchers to probe the causal role of a brain region at specific moments during behavior, but they require careful

control for sensory confounds. Offline protocols sidestep many confound issues and offer practical flexibility, but they trade temporal precision for the ability to induce longer-lasting changes in circuit function. The choice between approaches depends on the question at hand: is the goal to understand when a region contributes to behavior, or is the goal to open a window for plasticity that helps a circuit rewire? As we will see throughout this thesis, both approaches have their place in building a causal understanding of the neural basis of decision-making.

Nudging the brain

An important distinction between TUS and other brain stimulation techniques concerns what ultrasound can and cannot do. At the intensities used in human research, TUS alone cannot evoke action potentials or initiate behavior[9]. Consider the contrast with transcranial magnetic stimulation (TMS): if a TMS pulse is delivered over the hand area of primary motor cortex, your fingers will twitch whether you want them to or not[114]. TUS does not work that way. It is better conceptualized as a modulatory technique: if neural activity is ongoing, TUS can nudge, bias, or modulate that activity, but it cannot "generate" activity from scratch.

This modulatory nature suggests that TUS may be most effective during particular cognitive states. Emerging evidence points to a "sensitive window" in which TUS effects are most pronounced, specifically moments when we are engaged in a task but not yet fully committed to a course of action. Butler and colleagues (2022) demonstrated this elegantly: TUS effects on behavior were observed only during periods of decision uncertainty, when participants were actively deliberating between options[98]. When the correct response was obvious, TUS had little effect. This finding aligns with the notion that TUS modulates ongoing neural computations rather than imposing new ones. TUS provides a nudge, but the brain must do the work itself. I will return to this idea in **Chapter 2**, where I present evidence for a similar phenomenon in my own data.

The same principle applies to plasticity. Here, we use plasticity to refer to the capacity of a circuit to adapt, not only whether that adaptation strengthens or weakens connections[103]. One can think of plasticity at multiple levels: at the base level, there is the capacity for change; above that, directional plasticity determines whether change manifests as potentiation or depression; and above that, metaplasticity regulates how easily plastic change can occur at all[115]. A useful way to think about plasticity-focused TUS is that its goal is to open a window of plasticity, helping the brain regain the ability to adapt. TUS does not dictate the direction of change; rather, it creates the conditions under which behavior, experience, and therapeutic intervention can drive lasting and directional change. Just as online TUS nudges ongoing activity without prescribing its

outcome, offline TUS nudges the system toward a plastic state without prescribing whether that plasticity will strengthen or weaken particular connections.

Up or down? The directionality of TUS

A natural question to ask is: does TUS excite or inhibit neural activity? The honest answer is that it depends on many factors. Unlike techniques such as optogenetics, where the direction of effect is determined by the opsin expressed in the target cells, TUS interacts with endogenous mechanisms that vary across cell types, brain regions, and even the momentary state of the circuit being stimulated. As the field has matured, we have moved away from labeling TUS protocols as simply "excitatory" or "inhibitory" and toward a more nuanced understanding of how ultrasound interacts with neural tissue (**Figure 4D**).

The direction of TUS effects depends critically on which ion channels are engaged, because different channels serve fundamentally different functions in neural signaling. Potassium conductance strongly shapes resting membrane potential, spontaneous firing rate, overall excitability, and action potential propagation along the axon. Mechanosensitive K2P channels, most notably TREK and TRAAK, are well positioned to modulate these properties. Their activation leads to hyperpolarization and what might be called divisive inhibition: reducing background activity while preserving responsiveness to strong inputs[116,117]. This form of gain modulation can increase signal-to-noise ratio by suppressing low-amplitude, uninformative activity rather than silencing the circuit entirely. Importantly, TREK and TRAAK channels are expressed prominently in axonal domains such as the axon initial segment and nodes of Ranvier, where they regulate resting potential and repolarization[118,119]. This localization implies that ultrasound can modulate not just synaptic processing in gray matter, but also signal conduction in white matter, influencing conduction reliability, timing, or the ability to sustain high-frequency transmission.

Calcium plays a fundamentally different role: it is the central intracellular signal for plasticity induction[120]. Unlike potassium, which primarily shapes excitability and gain, calcium influx signals that meaningful neural activity has occurred. During effective synaptic activation or spiking, calcium enters the neuron through voltage-gated calcium channels and calcium-permeable receptors, binds to calmodulin, and initiates downstream signaling cascades that determine whether synaptic strength is modified. The magnitude and timing of calcium signals are critical: weak activity produces low calcium levels that bias the system toward long-term depression, while strong, well-timed activity produces large calcium transients that bias toward long-term potentiation. Calcium-permeable mechanosensitive channels, including TRP and Piezo family members, provide routes through which mechanical perturbations

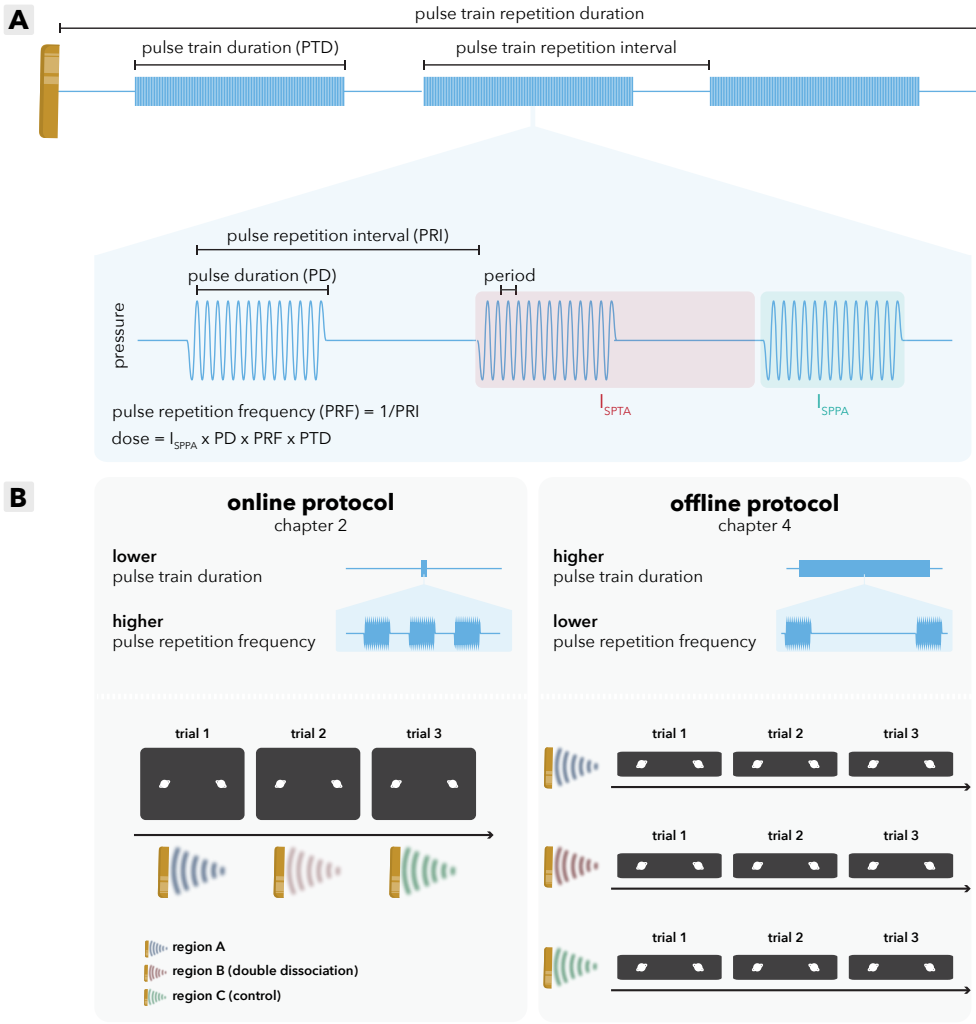


Figure 5 | TUS parameter space and protocol design. (A) A TUS protocol is defined by a hierarchy of temporal parameters. At the finest scale, the fundamental frequency determines the period of individual ultrasound cycles. Multiple cycles are grouped into pulses, characterized by their pulse duration (PD). Pulses repeat at a rate defined by the pulse repetition frequency (PRF; the inverse of the pulse repetition interval, PRI). These pulses are in turn organized into pulse trains, defined by the pulse train duration (PTD). Acoustic intensity is captured by two key measures: the spatial-peak pulse-average intensity (I_{SPPA}), reflecting peak intensity during a pulse, and the spatial-peak temporal-average intensity (I_{SPTA}), which averages intensity over the full pulse train. Together, these parameters determine the total dose delivered: dose = I_{SPPA}× PD × PRF × PTD. **(B)** These parameters can be configured into two broad protocol types with importantly different experimental implications. Online protocols often use shorter pulse train durations and higher PRFs (as I did in Chapter 2), allowing stimulation to be delivered within individual trials. This enables trial-by-trial variation of the stimulation target within a single session — for instance, targeting a region of interest, a second region for double dissociation, and an active control region in an interleaved fashion, all while the participant performs

the task. Offline protocols, by contrast, often apply longer, continuous stimulation blocks prior to task performance (as used in Chapter 4). Because stimulation precedes behavior, different targets cannot be interleaved within a session: each target condition requires a separate session, substantially increasing the experimental burden when multiple brain regions are to be compared.

can engage these plasticity pathways[90]. In contrast to potassium-mediated gain control, which biases activity statistics, calcium signaling determines whether the system enters a plastic state at all.

Different neuron types express different complements of these channels, meaning that the same ultrasound pulse will engage different molecular machinery depending on which cells it encounters[121]. Fast-spiking inhibitory interneurons, excitatory pyramidal cells, and neuromodulatory neurons each have distinct expression profiles of mechanosensitive channels, intrinsic membrane properties, and morphological features[121,122]. A neuron with high expression of TREK channels may respond to TUS primarily through gain modulation, while a neuron with prominent calcium-permeable TRP channels may be more susceptible to plasticity-related effects. Moreover, because different neuron types exert excitatory, inhibitory, or modulatory influences on their targets, engaging one population versus another can have opposite effects on circuit output. This heterogeneity means that the net effect of TUS on any given circuit reflects the weighted sum of its effects across multiple cell types with different channel compositions and functional roles.

Adding another layer of complexity, brain regions differ substantially in their neuronal composition and channel expression. The cerebral cortex is dominated by excitatory pyramidal neurons, which constitute roughly 80% of the neuronal population, with various interneuron classes making up the remainder[122]. The striatum, by contrast, is composed primarily of inhibitory medium spiny neurons[123]. Beyond cellular composition, regions also differ in their expression of mechanosensitive channels: the central amygdala, for example, shows particularly high expression of Piezo1 and perhaps correspondingly a greater sensitivity to ultrasound stimulation[124]. The relative proportions of different cell types, the local circuit architecture, the expression profile of mechanosensitive machinery, and the ongoing activity state of the network will all influence how a given region responds to ultrasound. A protocol that shifts excitatory-inhibitory balance in one direction in cortex might have quite different effects in striatum or amygdala.

The complexity at the cellular and regional level ultimately translates into behavioral effects that can be difficult to predict from first principles. Consider a scenario where TUS preferentially activates inhibitory interneurons in a region that normally suppresses downstream targets. Stimulation would increase local inhibition, but

because that region's output is itself inhibitory, the net effect on downstream structures could be disinhibition, effectively releasing those targets from suppression. At the behavioral level, this might manifest as enhanced performance or increased responding, even though the direct neural effect was inhibitory. Such circuit-level interactions mean that the relationship between neural and behavioral effects is rarely straightforward.

What human TUS studies can tell us with reasonable confidence is the direction of behavioral effects. When we stimulate a region and observe changes in reaction times, accuracy, or decision biases, we can characterize whether stimulation enhanced or impaired performance on a given task. This behavioral readout integrates across all the cellular and circuit-level complexity described above, providing a functional endpoint that is often more interpretable than trying to infer what happened at the level of individual neurons.

The directionality of TUS effects thus depends on many things: the channels being engaged, the neuronal types affected, the stimulation protocol, the brain region targeted, and the state of that region at the moment of stimulation. For these reasons, we refrain from labeling TUS, or even specific TUS protocols, as "excitatory" or "inhibitory". The effect is an emergent property of the interaction between ultrasound parameters and the state of the specific neural substrate being targeted. This complexity is humbling, but it is also what makes TUS such a fascinating tool: by varying parameters, we may eventually be able to selectively engage different circuit elements, rather than simply turning a brain region *on* or *off*.

OUTLINE OF THESIS

The aim of this thesis is to provide a better causal understanding of the neural mechanisms underlying motivation-action coupling in humans. Specifically, I investigate the roles of the ventral striatum (nucleus accumbens, NAcc) and amygdala in the interplay of Pavlovian and instrumental control. To achieve this aim, I employ transcranial ultrasound stimulation (TUS) to causally manipulate these deep brain structures, an approach that has only recently become feasible in humans. The following chapters address two orthogonal yet intertwined questions: first, can TUS effectively bias choice behavior in humans, and second, how do deep brain structures causally shape motivational action control in humans?

At the start of my doctoral work, evidence for TUS efficacy in humans was sparse and often confounded by auditory and somatosensory side effects. While animal studies

had demonstrated robust behavioral effects, it remained unclear whether these findings would translate to humans. In **Chapter 2**, I address this foundational question by translating a protocol from macaque research to humans. I targeted the frontal eye fields (FEF), a well-characterized region for which the causal role in contralateral saccade generation is firmly established. Participants performed a saccadic choice task while TUS was applied to either the left FEF, right FEF, or an active control site. Because neuromodulation effects are known to vary considerably across individuals, I also examined whether baseline inhibitory tone, measured using magnetic resonance spectroscopy, might predict TUS susceptibility. This chapter thus serves a dual purpose: establishing proof of concept for TUS as a behavioral modulation tool in humans and identifying individual-difference factors that may determine its efficacy.

Having established that we can use TUS to successfully modulate behavior in humans, I next sought to identify the neural substrates of motivation-action coupling that would serve as targets for causal intervention. In **Chapter 3**, I used functional magnetic resonance imaging (fMRI) while participants performed a motivational Go/NoGo task. This task allows us to dissociate Pavlovian tendencies from instrumental learning by orthogonalizing cue valence (reward versus loss avoidance) and required action (Go versus NoGo). A central concern in designing this study was that the Pavlovian bias, one of the key behavioral signatures of interest, might diminish with repeated task exposure, rendering the task unsuitable for the multi-session intervention design required by **Chapter 4**. I therefore developed a modified version of the task intended to sustain robust Pavlovian biases across many repetitions. The chapter examines which brain regions track motivational signals during cue processing and outcome evaluation, with particular focus on the NAcc and amygdala as candidate substrates for causal investigation.

In **Chapter 4**, I combine the insights from the preceding chapters to directly test the causal roles of the NAcc and amygdala in motivation-action coupling. Using a within-subject design, participants completed three sessions in which they received TUS to either the NAcc, the amygdala, or sham stimulation (order counterbalanced and randomized) before performing the motivational Go/NoGo task. To characterize the behavioral effects of TUS, we employed computational models that allowed us to quantify Pavlovian biases and instrumental learning dynamics. The central question is whether, and how, perturbing each of these regions alters the balance between Pavlovian and instrumental control – and whether their contributions are shared or dissociable.

Taken together, this thesis provides the first causal evidence for the distinct contributions of the NAcc and amygdala to Pavlovian and instrumental control

in healthy humans. By establishing that TUS can effectively modulate deep brain structures and demonstrating its utility for probing causal brain-behavior relationships, this work opens a new chapter in human decision neuroscience. Where previous research relied on correlational neuroimaging or drew inferences from animal models, we can now directly test whether specific brain regions are necessary for specific aspects of motivated behavior. Beyond advancing our fundamental understanding of how the brain transforms motivation into action, these findings have translational potential: they identify specific neural targets and mechanisms that could inform the development of more precise interventions for disorders of motivation and behavioral control.

REFERENCES

1. Dayan, P., Niv, Y., Seymour, B. & Daw, N. D. The misbehavior of value and the discipline of the will. *Neural Netw. Off. J. Int. Neural Netw. Soc.* ***19***, 1153–1160 (2006).
2. Guitart-Masip, M. *et al.* Differential, but not opponent, effects of l-DOPA and citalopram on action learning with reward and punishment. *Psychopharmacology (Berl.)* ***231***, 955–966 (2014).
3. Garbusow, M. *et al.* Pavlovian-to-instrumental transfer effects in the nucleus accumbens relate to relapse in alcohol dependence. *Addict. Biol.* ***21***, 719–731 (2016).
4. Heinz, A., Schlagenhauf, F., Beck, A. & Wackerhagen, C. Dimensional psychiatry: mental disorders as dysfunctions of basic learning mechanisms. *J. Neural Transm. Vienna Austria 1996* ***123***, (2016).
5. Scholz, V. *et al.* Cortical dopamine reduces the impact of motivational biases governing automated behaviour. *Neuropsychopharmacol. Off. Publ. Am. Coll. Neuropsychopharmacol.* ***47***, 1503–1512 (2022).
6. Swart, J. C. *et al.* Catecholaminergic challenge uncovers distinct Pavlovian and instrumental mechanisms of motivated (in)action. *eLife* ***6***, e22169 (2017).
7. Frank, M. J., Seeberger, L. C. & O'reilly, R. C. By carrot or by stick: cognitive reinforcement learning in parkinsonism. *Science* ***306***, 1940–1943 (2004).
8. Cavanagh, J. F. *et al.* Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nat. Neurosci.* ***14***, 1462–1467 (2011).
9. Darmani, G. *et al.* Non-invasive transcranial ultrasound stimulation for neuromodulation. *Clin. Neurophysiol.* ***135***, 51–73 (2022).
10. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* ***30***, 535–574 (2007).
11. Felleman, D. J. & Van Essen, D. C. Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex N. Y. N 1991* ***1***, 1–47 (1991).
12. Deffieux, T. *et al.* Low-Intensity Focused Ultrasound Modulates Monkey Visuomotor Behavior. *Curr. Biol.* ***23***, 2430–2433 (2013).
13. Kubanek, J. *et al.* Remote, brain region-specific control of choice behavior with ultrasonic waves. *Sci. Adv.* ***6***, eaaz4193 (2020).
14. Paus, T. Location and function of the human frontal eye-field: A selective review. *Neuropsychologia* ***34***, 475–483 (1996).
15. Vernet, M., Quentin, R., Chanes, L., Mitsumasu, A. & Valero-Cabré, A. Frontal eye field, where art thou? Anatomy, function, and non-invasive manipulation of frontal regions involved in eye movements and associated cognitive operations. *Front. Integr. Neurosci.* ***8***, 66 (2014).
16. Geurts, D. E. M., Huys, Q. J. M., den Ouden, H. E. M. & Cools, R. Aversive Pavlovian Control of Instrumental Behavior in Humans. *J. Cogn. Neurosci.* ***25***, 1428–1441 (2013).
17. Huys, Q. J. M. *et al.* Disentangling the Roles of Approach, Activation and Valence in Instrumental and Pavlovian Responding. *PLOS Comput. Biol.* ***7***, e1002028 (2011).

18. Dickinson, A. & Balleine, B. Motivational control of goal-directed action. *Anim. Learn. Behav.* ***22***, 1–18 (1994).
19. Robbins, T. W. & Everitt, B. J. A role for mesencephalic dopamine in activation: commentary on Berridge (2006). *Psychopharmacology (Berl.)* ***191***, 433–437 (2007).
20. Rangel, A., Camerer, C. & Montague, P. R. Neuroeconomics: The neurobiology of value-based decision-making. *Nat. Rev. Neurosci.* ***9***, 545–556 (2008).
21. Pavlov, I. P. *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*. xv, 430 (Oxford Univ. Press, Oxford, England, 1927).
22. Algermissen, J., Swart, J. C., Scheeringa, R., Cools, R. & Den Ouden, H. E. M. Striatal BOLD and midfrontal theta power express motivation for action. *Cereb. Cortex* ***32***, 2924–2942 (2022).
23. Algermissen, J., Swart, J. C., Scheeringa, R., Cools, R. & den Ouden, H. E. M. Prefrontal signals precede striatal signals for biased credit assignment in motivational learning biases. *Nat. Commun.* ***15***, 19 (2024).
24. Corbit, L. H. & Balleine, B. W. The general and outcome-specific forms of Pavlovian-instrumental transfer are differentially mediated by the nucleus accumbens core and shell. *J. Neurosci. Off. J. Soc. Neurosci.* ***31***, 11786–11794 (2011).
25. Rescorla, R. A. & Solomon, R. L. Two-process learning theory: Relationships between Pavlovian conditioning and instrumental learning. *Psychol. Rev.* ***74***, 151–182 (1967).
26. Tempel, B. L., Bonini, N., Dawson, D. R. & Quinn, W. G. Reward learning in normal and mutant Drosophila. *Proc. Natl. Acad. Sci. U. S. A.* ***80***, 1482–1486 (1983).
27. Hershberger, W. A. An approach through the looking-glass. *Anim. Learn. Behav.* ***14***, 443–451 (1986).
28. LeDoux, J. E. Emotion circuits in the brain. *Annu. Rev. Neurosci.* ***23***, 155–184 (2000).
29. Phelps, E. A. & LeDoux, J. E. Contributions of the amygdala to emotion processing: from animal models to human behavior. *Neuron* ***48***, 175–187 (2005).
30. Holland, P. C. & Gallagher, M. Double dissociation of the effects of lesions of basolateral and central amygdala on conditioned stimulus-potentiated feeding and Pavlovian-instrumental transfer. *Eur. J. Neurosci.* ***17***, 1680–1694 (2003).
31. Janak, P. H. & Tye, K. M. From circuits to behaviour in the amygdala. *Nature* ***517***, 284–292 (2015).
32. Paton, J. J., Belova, M. A., Morrison, S. E. & Salzman, C. D. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* ***439***, 865–870 (2006).
33. Pedersen, W. S. *et al.* The effects of stimulus novelty and negativity on BOLD activity in the amygdala, hippocampus, and bed nucleus of the stria terminalis. *Soc. Cogn. Affect. Neurosci.* ***12***, 748–757 (2017).
34. Cardinal, R. N., Parkinson, J. A., Hall, J. & Everitt, B. J. Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neurosci. Biobehav. Rev.* ***26***, 321–352 (2002).

35. Corbit, L. H. & Balleine, B. W. Double dissociation of basolateral and central amygdala lesions on the general and outcome-specific forms of pavlovian-instrumental transfer. *J. Neurosci. Off. J. Soc. Neurosci.* ***25***, 962–970 (2005).
36. Rushworth, M. F., Kolling, N., Sallet, J. & Mars, R. B. Valuation and decision-making in frontal cortex: one or many serial or parallel systems? *Curr. Opin. Neurobiol.* ***22***, 946–955 (2012).
37. Taylor, J. R. & Robbins, T. W. Enhanced behavioural control by conditioned reinforcers following microinjections of d-amphetamine into the nucleus accumbens. *Psychopharmacology (Berl.)* ***84***, 405–412 (1984).
38. Haber, S. N. & Knutson, B. The Reward Circuit: Linking Primate Anatomy and Human Imaging. *Neuropsychopharmacology* ***35***, 4–26 (2010).
39. Collins, A. G. E. & Frank, M. J. Opponent actor learning (OpAL): modeling interactive effects of striatal dopamine on reinforcement learning and choice incentive. *Psychol. Rev.* ***121***, 337–366 (2014).
40. Frank, M. J. Dynamic Dopamine Modulation in the Basal Ganglia: A Neurocomputational Account of Cognitive Deficits in Medicated and Nonmedicated Parkinsonism. *J. Cogn. Neurosci.* ***17***, 51–72 (2005).
41. Thorndike, E. L. Animal intelligence: An experimental study of the associative processes in animals. *Psychol. Rev. Monogr. Suppl.* ***2***, i–109 (1898).
42. Rescorla, R. A. & Wagner, A. R. A Theory of Pavlovian Conditioning Variations in the Effectiveness of Reinforcement and Nonreinforcement. *Class. Cond. Curr. Res. Theory* ***21***, 64–99 (1972).
43. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* ***275***, 1593–1599 (1997).
44. Steinberg, E. E. *et al.* A causal link between prediction errors, dopamine neurons and learning. *Nat. Neurosci.* ***16***, 966–973 (2013).
45. O'Doherty, J. *et al.* Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* ***304***, 452–454 (2004).
46. Rushworth, M. F. S., Walton, M. E., Kennerley, S. W. & Bannerman, D. M. Action sets and decisions in the medial frontal cortex. *Trends Cogn. Sci.* ***8***, 410–417 (2004).
47. Haber, S. N. The primate basal ganglia: parallel and integrative networks. *J. Chem. Neuroanat.* ***26***, 317–330 (2003).
48. DeLong, M. R. & Wichmann, T. Circuits and circuit disorders of the basal ganglia. *Arch. Neurol.* ***64***, 20–24 (2007).
49. Estes, W. K. & Skinner, B. F. Some quantitative properties of anxiety. *J. Exp. Psychol.* ***29***, 390–400 (1941).
50. Holmes, N. M., Marchand, A. R. & Coutureau, E. Pavlovian to instrumental transfer: a neurobehavioural perspective. *Neurosci. Biobehav. Rev.* ***34***, 1277–1295 (2010).
51. Talmi, D., Seymour, B., Dayan, P. & Dolan, R. J. Human pavlovian-instrumental transfer. *J. Neurosci. Off. J. Soc. Neurosci.* ***28***, 360–368 (2008).

52. Cartoni, E., Balleine, B. & Baldassarre, G. Appetitive Pavlovian-instrumental Transfer: A review. *Neurosci. Biobehav. Rev.* ***71***, 829–848 (2016).
53. Cools, R., Nakamura, K. & Daw, N. D. Serotonin and dopamine: unifying affective, activational, and decision functions. *Neuropsychopharmacol. Off. Publ. Am. Coll. Neuropsychopharmacol.* ***36***, 98–113 (2011).
54. Guitart-Masip, M., Beierholm, U. R., Dolan, R., Duzel, E. & Dayan, P. Vigor in the face of fluctuating rates of reward: an experimental examination. *J. Cogn. Neurosci.* ***23***, 3933–3938 (2011).
55. Guitart-Masip, M. *et al.* Action controls dopaminergic enhancement of reward representations. *Proc. Natl. Acad. Sci. U. S. A.* ***109***, 7511–7516 (2012).
56. Guitart-Masip, M. *et al.* Action Dominates Valence in Anticipatory Representations in the Human Striatum and Dopaminergic Midbrain. *J. Neurosci.* ***31***, 7867–7875 (2011).
57. Guitart-Masip, M. *et al.* Go and no-go learning in reward and punishment: interactions between affect and effect. *NeuroImage* ***62***, 154–166 (2012).
58. Kalivas, P. W. & Volkow, N. D. The neural basis of addiction: a pathology of motivation and choice. *Am. J. Psychiatry* ***162***, 1403–1413 (2005).
59. Shin, L. M. & Liberzon, I. The neurocircuitry of fear, stress, and anxiety disorders. *Neuropsychopharmacol. Off. Publ. Am. Coll. Neuropsychopharmacol.* ***35***, 169–191 (2010).
60. Huys, Q. J., Pizzagalli, D. A., Bogdan, R. & Dayan, P. Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis. *Biol. Mood Anxiety Disord.* ***3***, 12 (2013).
61. Burguière, E., Monteiro, P., Mallet, L., Feng, G. & Graybiel, A. M. Striatal circuits, habits, and implications for obsessive-compulsive disorder. *Curr. Opin. Neurobiol.* ***30***, 59–65 (2015).
62. Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A. & Daw, N. D. Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. *eLife* ***5***, e11305 (2016).
63. Robbins, T. W., Gillan, C. M., Smith, D. G., Wit, S. de & Ersche, K. D. Neurocognitive endophenotypes of impulsivity and compulsivity: towards dimensional psychiatry. *Trends Cogn. Sci.* ***16***, 81–91 (2012).
64. Fry, W. J., Fry, F. J., Barnard, J. W., Krumins, R. F. & Brennan, J. F. Ultrasonic Lesions in the Mammalian Central Nervous System. *Science* ***122***, 517–518 (1955).
65. Fry, W. J. & Fry, F. J. Fundamental neurological research and human neurosurgery using intense ultrasound. *IRE Trans. Med. Electron.* ***ME-7***, 166–181 (1960).
66. Hynynen, K., McDannold, N., Sheikov, N. A., Jolesz, F. A. & Vykhodtseva, N. Local and reversible blood-brain barrier disruption by noninvasive focused ultrasound at frequencies suitable for trans-skull sonications. *NeuroImage* ***24***, 12–20 (2005).
67. McDannold, N., Arvanitis, C. D., Vykhodtseva, N. & Livingstone, M. S. Temporary disruption of the blood-brain barrier by use of ultrasound and microbubbles: safety and efficacy evaluation in rhesus macaques. *Cancer Res.* ***72***, 3652–3663 (2012).
68. Blackmore, J., Shrivastava, S., Sallet, J., Butler, C. R. & Cleveland, R. O. Ultrasound Neuromodulation: A Review of Results, Mechanisms and Safety. *Ultrasound Med. Biol.* ***45***, 1509–1536 (2019).

69. Tyler, W. J. *et al.* Remote Excitation of Neuronal Circuits Using Low-Intensity, Low-Frequency Ultrasound. *PLoS ONE* ***3***, e3511 (2008).
70. Tufail, Y. *et al.* Transcranial Pulsed Ultrasound Stimulates Intact Brain Circuits. *Neuron* ***66***, 681–694 (2010).
71. Yoo, S.-S. *et al.* Focused ultrasound modulates region-specific brain activity. *NeuroImage* ***56***, 1267–1275 (2011).
72. Folloni, D. *et al.* Manipulation of Subcortical and Deep Cortical Activity in the Primate Brain Using Transcranial Focused Ultrasound Stimulation. *Neuron* ***101***, 1109-1116.e5 (2019).
73. Khalighinejad, N. *et al.* A Basal Forebrain-Cingulate Circuit in Macaques Decides It Is Time to Act. *Neuron* ***105***, 370-384.e8 (2020).
74. Legon, W. *et al.* Transcranial focused ultrasound modulates the activity of primary somatosensory cortex in humans. *Nat. Neurosci.* ***17***, 322–329 (2014).
75. Fomenko, A. *et al.* Systematic examination of low-intensity ultrasound parameters on human motor cortex excitability and behavior. *eLife* ***9***, e54497 (2020).
76. Kop, B. R. *et al.* Auditory confounds can drive online effects of transcranial ultrasonic stimulation in humans. *eLife* ***12***, (2024).
77. Kop, B. R., de Jong, L., Kim, B. P., den Ouden, H. E. M. & Verhagen, L. Parameter optimisation for mitigating somatosensory confounds during transcranial ultrasonic stimulation. *Brain Stimulat.* ***18***, 1224–1236 (2025).
78. Aubry, J.-F. *et al.* ITRUSST consensus on biophysical safety for transcranial ultrasound stimulation. *Brain Stimulat.* ***18***, 1896–1905 (2025).
79. Kubanek, J., Shukla, P., Das, A., Baccus, S. A. & Goodman, M. B. Ultrasound Elicits Behavioral Responses through Mechanical Effects on Neurons and Ion Channels in a Simple Nervous System. *J. Neurosci.* ***38***, 3081–3091 (2018).
80. Nandi, T., Kop, B. R., Butts Pauly, K., Stagg, C. J. & Verhagen, L. The relationship between parameters and effects in transcranial ultrasonic stimulation. *Brain Stimulat.* ***17***, 1216–1228 (2024).
81. Nandi, T. *et al.* Biophysical effects and neuromodulatory dose of transcranial ultrasonic stimulation. *Brain Stimulat.* ***18***, 659–664 (2025).
82. Gefen, A., Gefen, N., Zhu, Q., Raghupathi, R. & Margulies, S. S. Age-Dependent Changes in Material Properties of the Brain and Braincase of the Rat. *J. Neurotrauma* ***20***, 1163–1177 (2003).
83. Tyler, W. J. The mechanobiology of brain function. *Nat. Rev. Neurosci.* ***13***, 867–878 (2012).
84. Star, E. N., Kwiatkowski, D. J. & Murthy, V. N. Rapid turnover of actin in dendritic spines and its regulation by activity. *Nat. Neurosci.* ***5***, 239–246 (2002).
85. Iwasa, K., Tasaki, I. & Gibbons, R. C. Swelling of nerve fibers associated with action potentials. *Science* ***210***, 338–339 (1980).
86. Tasaki, I. & Iwasa, K. Shortening of nerve fibers associated with propagated nerve impulse. *Biochem. Biophys. Res. Commun.* ***94***, 716–720 (1980).

87. El Hady, A. & Machta, B. B. Mechanical surface waves accompany action potential propagation. *Nat. Commun.* ***6***, 6697 (2015).
88. Kim, G. H., Kosterin, P., Obaid, A. L. & Salzberg, B. M. A mechanical spike accompanies the action potential in Mammalian nerve terminals. *Biophys. J.* ***92***, 3122–3129 (2007).
89. Kubanek, J. *et al.* Ultrasound modulates ion channel currents. *Sci. Rep.* ***6***, 24170 (2016).
90. Yoo, S., Mittelstein, D. R., Hurt, R. C., Lacroix, J. & Shapiro, M. G. Focused ultrasound excites cortical neurons via mechanosensitive calcium accumulation and ion channel amplification. *Nat. Commun.* ***13***, 493 (2022).
91. Morris, C. E. Voltage-Gated Channel Mechanosensitivity: Fact or Friction? *Front. Physiol.* ***2***, (2011).
92. Jerusalem, A. *et al.* Electrophysiological-mechanical coupling in the neuronal membrane and its role in ultrasound neuromodulation and general anaesthesia. *Acta Biomater.* ***97***, 116–140 (2019).
93. Plaksin, M., Shoham, S. & Kimmel, E. Intramembrane Cavitation as a Predictive Bio-Piezoelectric Mechanism for Ultrasonic Brain Stimulation. *Phys. Rev. X* ***4***, 011004 (2014).
94. Allen, N. J. & Eroglu, C. Cell Biology of Astrocyte-Synapse Interactions. *Neuron* ***96***, 697–708 (2017).
95. Perea, G. & Araque, A. Astrocytes potentiate transmitter release at single hippocampal synapses. *Science* ***317***, 1083–1086 (2007).
96. Bhargava, P. *et al.* Bile acid metabolism is altered in multiple sclerosis and supplementation ameliorates neuroinflammation. *J. Clin. Invest.* ***130***, 3467–3482 (2020).
97. Oh, S.-J. *et al.* Ultrasonic Neuromodulation via Astrocytic TRPA1. *Curr. Biol. CB* ***29***, 3386-3401.e8 (2019).
98. Butler, C. R. *et al.* Transcranial ultrasound stimulation to human middle temporal complex improves visual motion detection and modulates electrophysiological responses. *Brain Stimulat.* ***15***, 1236–1245 (2022).
99. Farboud, S. *et al.* Rapid modulation of choice behavior by ultrasound on the human frontal eye fields. *Nat. Commun.* https://doi.org/10.1038/s41467-026-69854-7 (2026) doi:10.1038/s41467-026-69854-7.
100. Fong, P.-Y. *et al.* A double-blind replication attempt of offline 5Hz-rTUS-induced corticospinal excitability. *Imaging Neurosci.* ***3***, IMAG.a.1046 (2025).
101. Yaakub, S. N. *et al.* Non-invasive ultrasonic neuromodulation of the human nucleus accumbens impacts reward sensitivity. *Nat. Commun.* ***16***, 10192 (2025).
102. Zeng, K. *et al.* Effects of different sonication parameters of theta burst transcranial ultrasound stimulation on human motor cortex. *Brain Stimulat.* ***17***, 258–268 (2024).
103. Malenka, R. C. & Bear, M. F. LTP and LTD: An Embarrassment of Riches. *Neuron* ***44***, 5–21 (2004).
104. Bear, M. F. & Malenka, R. C. Synaptic plasticity: LTP and LTD. *Curr. Opin. Neurobiol.* ***4***, 389–399 (1994).

105. Bliss, T. V. & Lomo, T. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* ***232***, 331–356 (1973).
106. Lazzaro, V. D., Ziemann, U. & Lemon, R. N. State of the art: Physiology of transcranial motor cortex stimulation. *Brain Stimulat.* ***1***, 345–362 (2008).
107. Ziemann, U. TMS induced plasticity in human cortex. *Rev. Neurosci.* ***15***, 253–266 (2004).
108. Pittenger, C. & Duman, R. S. Stress, depression, and neuroplasticity: a convergence of mechanisms. *Neuropsychopharmacol. Off. Publ. Am. Coll. Neuropsychopharmacol.* ***33***, 88–109 (2008).
109. Lüscher, C. & Malenka, R. C. Drug-evoked synaptic plasticity in addiction: from molecular changes to circuit remodeling. *Neuron* ***69***, 650–663 (2011).
110. McLean, C. P. & Foa, E. B. State of the Science: Prolonged exposure therapy for the treatment of posttraumatic stress disorder. *J. Trauma. Stress* ***37***, 535–550 (2024).
111. Careaga, M. B. L., Girardi, C. E. N. & Suchecki, D. Understanding posttraumatic stress disorder through fear conditioning, extinction and reconsolidation. *Neurosci. Biobehav. Rev.* ***71***, 48–57 (2016).
112. Apkarian, V. A., Hashmi, J. A. & Baliki, M. N. Pain and the brain: specificity and plasticity of the brain in clinical chronic pain. *Pain* ***152***, S49–S64 (2011).
113. Nijs, J. *et al.* Central sensitisation in chronic pain conditions: latest discoveries and their potential for precision medicine. *Lancet Rheumatol.* ***3***, e383–e392 (2021).
114. Hallett, M. Transcranial magnetic stimulation: a primer. *Neuron* ***55***, 187–199 (2007).
115. Abraham, W. C. Metaplasticity: tuning synapses and networks for plasticity. *Nat. Rev. Neurosci.* ***9***, 387 (2008).
116. Prescott, S. A. & De Koninck, Y. Gain control of firing rate by shunting inhibition: roles of synaptic noise and dendritic saturation. *Proc. Natl. Acad. Sci. U. S. A.* ***100***, 2076–2081 (2003).
117. Sorum, B., Docter, T., Panico, V., Rietmeijer, R. A. & Brohawn, S. G. Tension activation of mechanosensitive two-pore domain K+ channels TRAAK, TREK-1, and TREK-2. *Nat. Commun.* ***15***, 3142 (2024).
118. Escobedo, G. & Rasband, M. N. Potassium channel clustering: mechanisms shaping axonal excitability. *Front. Cell. Neurosci.* ***19***, 1627517 (2025).
119. Kanda, H. *et al.* TREK-1 and TRAAK Are Principal K+ Channels at the Nodes of Ranvier for Rapid Action Potential Conduction on Mammalian Myelinated Afferent Nerves. *Neuron* ***104***, 960-971.e7 (2019).
120. Mateos-Aparicio, P. & Rodríguez-Moreno, A. Calcium Dynamics and Synaptic Plasticity. *Adv. Exp. Med. Biol.* ***1131***, 965–984 (2020).
121. Lemaire, T. *et al.* Microscopic deconstruction of cortical circuit stimulation by transcranial ultrasound. 2024.10.10.617091 Preprint at https://doi.org/10.1101/2024.10.10.617091 (2024).
122. Markram, H. *et al.* Interneurons of the neocortical inhibitory system. *Nat. Rev. Neurosci.* ***5***, 793–807 (2004).

123. Graveland, G. A. & Difiglia, M. The frequency and distribution of medium-sized neurons with indented nuclei in the primate and rodent neostriatum. *Brain Res.* ***327***, 307–311 (1985).
124. Zhu, J. *et al.* The mechanosensitive ion channel Piezo1 contributes to ultrasound neuromodulation. *Proc. Natl. Acad. Sci.* ***120***, e2300291120 (2023).




PART II
Interlude
Chapter 2

THE SIMORGH'S FEATHER

When speaking of old Persian legends, one name rises almost immediately: **Rostam**.

Rostam is remembered for victories that feel larger than life. Most famously for the Seven Labors, a gauntlet of trials in which he fights through wilderness, monsters, deception, and war, as if the world itself keeps placing obstacles in his path and he keeps answering with the same blunt inevitability: he endures, he overcomes, he moves forward. He is so formidable that no one seems to be his equal. And yet even Rostam's story carries its own shadow: in one of the greatest tragedies of the epic tradition, he defeats and kills his own son without knowing who he is until it is too late. Strength, in these stories, is never pure simplicity. It is entwined with consequence.

Rostam feels unstoppable. The question is: what came before him?

Before the hero, there is his father, **Zāl**; and Zāl's story begins not with triumph but with abandonment. As a baby he is left behind by his father, Sām, who is frightened by his appearance and convinced it is an omen or a sickness. Sām carries the child far away, into the mountains, to a place where no one will find him.

But Zāl is found.

A **Simorgh** – the great, radiant bird of Persian myth, phoenix-like in her majesty – comes upon the infant and lifts him gently into her care. With her brilliant, many-colored feathers, she becomes his guardian when no human will. Zāl grows under her protection, from a discarded child into a young man who can stand on his own. He is strong, capable, and marked by a destiny he cannot yet name.

Eventually the day arrives when Zāl must leave the Alborz mountains and return to the world of people. He goes with a quiet confidence that life will open for him. The Simorgh, however, feels the grief of parting. She knows she does not belong among humans; she cannot follow him into courts and cities. And so she gives him one final gift: **a single feather**.

Not a weapon. Not armor. Not a talisman of domination.

A feather – light, delicate, almost nothing.

She tells him: If you ever truly need me, burn this. It will be enough. It will call me back.

Zāl descends. His strength and charm are quickly recognized, and before long he becomes

a warrior of renown. On his path he comes to Kabul, where he sees a woman whose beauty stops him in his tracks: the princess **Rudābeh**. They fall in love. Time passes. Rudābeh becomes pregnant, her belly growing and growing, as if the child inside is already too large for the ordinary world.

Then the day of birth arrives… And soon it becomes clear that something is wrong.

The baby – Rostam, the future hero of Iran – is simply too big. Rudābeh weakens; the child cannot be delivered. It looks, in that moment, as though neither mother nor son will survive. Zāl panics. He is desperate, grieving before the loss has even happened. And then, at the edge of despair, he remembers the small thing he was given long ago.

The feather. He lights it. And as promised, the Simorgh comes.

She does not arrive with spectacle. She arrives with precision, with knowledge, showing Zāl what to do when there seems to be no path forward. In the telling, this becomes the first Caesarean-like birth in the Persian epic imagination: Rudābeh's life is saved, and Rostam is delivered into the world. With that birth, an entire future becomes possible.

And it begins with something as small as a feather.

That feather, on its own, cannot force anything. It is not power in the sense of control. It does not decide the outcome by itself. But in the right moment, under the right conditions, it tips the scales. It is a nudge: a small, targeted intervention that shifts what is possible, and in doing so changes everything that follows.

This is the spirit in which I frame transcranial ultrasound stimulation in this chapter. Not as a tool that determines behavior, not as something that overrides agency or writes commands into the mind, but as a subtle perturbation that can, in the right neural context, modulate a system already in motion. A feather, not a hammer. A tilt, not a shove.

And yet, as the legend reminds us, even a feather – applied at the right time, in the right place – can open the avenue for what is yet to come.

CHAPTER 2

Rapid modulation of choice behavior by ultrasound on the human frontal eye fields

This chapter has been published as:

Farboud, S., Kop, B.R., Koolschijn, R.S., Walstra, S.L.Y., Marques J.P., Chetverikov, A., Medendorp, W.P., Verhagen, L., den Ouden, H.E.M. (2026)

Rapid modulation of choice behavior by ultrasound on the human frontal eye fields

Nature Communications
https://doi.org/10.1038/s41467-026-69854-7

ABSTRACT

A fundamental challenge in neuroscience is establishing causal brain-function relationships with spatial and temporal precision. Transcranial ultrasonic stimulation offers a unique opportunity to modulate deep brain structures non-invasively with high spatial resolution, but temporally precise effects and their neurophysiological foundations have yet to be demonstrated in humans. Here, we develop a temporally precise ultrasound stimulation protocol targeting the frontal eye fields — a well-characterized circuit critical for saccadic eye movements. We demonstrate that ultrasonic stimulation induces robust excitatory behavioral effects. Importantly, individual differences in baseline GABAergic inhibitory tone predict response magnitude. These findings establish ultrasound stimulation as a reliable tool for chronometric circuit interrogation and highlight the importance of neurophysiological state in neuromodulation. This work bridges human and animal research, advancing targeted transcranial ultrasonic stimulation applications in neuroscience and clinical settings.



INTRODUCTION

In recent years, transcranial ultrasonic stimulation (TUS) has emerged as a promising non-invasive technique for brain stimulation, capable of targeting both cortical and subcortical regions with exceptional spatial resolution[1]. This makes TUS highly valuable for studying brain function and offers great potential for therapeutic applications. Much of our current understanding is derived from animal studies[2–6], but there exists a translational gap to human application. The large majority of human studies to date focus on repetitive 'offline' protocols with temporally sustained effects[7,8], while those addressing immediate or 'online' effects remain limited and often marred by confounds[9] (with the exception of work by Butler and colleagues[10]). This leaves questions on the physiological mechanisms and temporal dynamics of TUS unanswered and calls for robust and replicable protocols in humans. In this study, we introduce an effective online TUS protocol for humans with immediate effects, by leveraging a well-established TUS protocol from non-human primates[2]. To this end, we take advantage of an evolutionarily conserved brain circuit with a well-characterized link to readily measurable behavior that acts as a model system for more complex decision-making – the frontal eye fields (FEFs).

The role of the FEFs in the planning and generation of saccadic eye movements has been well established in both humans and non-human primates[11,12], and features a basic topographic representation that encodes both the direction and amplitude of saccades in the opposite visual hemifield[11,12]. FEF's involvement in contralateral saccade generation has been further evidenced by lesion studies[13–16] and transcranial magnetic stimulation experiments[17–24]. This well-characterized role of the FEF in contralateral saccades allows for precise characterization of TUS effects. For instance, in macaques online TUS of the FEF was found to bias saccades towards the contralateral side, which suggests that stimulation has a net excitatory effect[2]. However, it is not known whether these results can be directly translated to humans, i.e., whether online TUS stimulation of the human FEF can induce the same excitatory effect, or whether anatomical, physiological, and behavioral differences between humans and non-human primates would instead result in net inhibitory or perturbatory effects. Indeed, the effect of TUS has been found to vary from excitatory to inhibitory to perturbatory depending on the specific stimulation protocol settings[25], underscoring the need for caution when interpreting TUS-induced behavioral changes.

Considering that the effects of brain stimulation are highly dependent on brain states and traits[26–30], it is expected that TUS effects vary not only between species, but also between individuals. Consequently, it is pertinent to consider the individual neurophysiological state when investigating the mechanisms and consequences of

TUS. This interindividual variability may be influenced by factors such as an individual's cortical inhibitory tone, which has been shown to impact the effects of other noninvasive stimulation methods[31]. Moreover, differences in cortical inhibitory tone in the FEFs have been linked to individual variations in the capacity to resist distractions while generating saccades[32]. Therefore, it seems likely that the neuromodulatory effects of TUS on an individual may have different effects. Given that TUS may modulate both excitatory and inhibitory neuronal populations in the brain, we hypothesize that the net effects of TUS could be shaped by individual differences in the excitation/inhibition balance. To explore whether interindividual differences in the effects of TUS are similarly inhibitory tone-dependent, we measured individual level concentrations of the inhibitory neurotransmitter GABA+ in the FEF using magnetic resonance spectroscopy (MRS).

In the present study, we tested the hypothesis that TUS applied to the human FEF has an immediate, excitatory effect on saccade direction, and that this effect is mediated by local inhibitory tone. Participants completed a simple saccade choice task while receiving TUS during stimulus presentation, applied to either their left or right FEF ('stimulation'), or to the left or right hand motor cortex (M1) ('active control'). FEF TUS induced a significant increase in the selection of contralateral saccades, directly replicating findings from a previous study in macaques[2] and indicating that FEF TUS has net excitatory effects on saccade selection in humans. Notably, participants' characteristic inhibitory tone in FEF was found to predict interindividual differences in the effect of TUS, suggesting that TUS susceptibility is linked to an individual's inhibitory tone. Taken together, our findings pave the way to use TUS as an effective and temporally specific tool to study the functional circuit dynamics of the human brain and offer critical insights into the factors that drive interindividual differences in response to this neuromodulation technique.

RESULTS

Baseline saccade task behavior

Thirty-five right-handed participants (M_{age} = 24.1, SD_{age} = 3.2, range = 20 – 32; 15 females, 20 males) performed a saccadic decision task in which two visual stimuli were presented asynchronously and equidistantly on either side of fixation (**Figure 1A**). Participants were instructed to saccade as quickly as possible to the stimulus that appeared first (i.e., target). We examined the probability of making a rightward saccade across all stimulus onset asynchronies (SOA, i.e., delay between target and distractor). Participants performed well on the task in the baseline (sham) condition: When the

target is on the right, participants were more likely to make a rightward saccade (b = 14.1, 95%-CI [13.0, 15.4], $c^2(1) = 530$, $p < 0.001$, **Figure 1B**). At the group-level, there is a lack of a noticeable rightward or leftward bias in the sham condition, although within participants there is variability in baseline side bias (**Figure 1B**).

In contrast to magnetic or electrical stimulation, which can directly elicit evoked responses (motor-evoked potentials[33]; visual-evoked potentials[34,35]), TUS at the low intensities employed in the current study is thought to act through modulating sub-firing threshold, ongoing activity, thereby gently nudging network dynamics rather than producing immediate, large-scale neural responses[36]. Therefore, we expected TUS effects to bias responses primarily on trials with short SOAs, i.e., high uncertainty about which target appeared first. On these trials, with low sensory evidence for left- or rightward responses, a small perturbation has the potential to bias the outcome, and thus we expected that TUS could tip the scales[2]. In contrast, for long SOAs there was low uncertainty about the correct response and TUS would not be able to overrule the associated robust neural signal. Therefore, we designed the task to oversample trials with shorter SOAs (**Figure 1C**) and focused the primary analysis on trials with SOAs where participants were <75% correct (**Figure 1B**, hereafter referred to as the choice domain).

FEF-TUS shows contralateral bias correlated with GABA+ levels

Ultrasonic stimulation of both the left and right FEFs significantly increased contralateral saccades (**Figure 2C-D**; b = -0.25, 95%-CI [-0.40, -0.10], $c^2(1) = 10.3$, p = 0.001). Left vs right FEF increased the odds of a rightward saccade by 13% (OR=1.13, 95%-CI [1.05–1.22]), corresponding to Cohen's d (logit) = 0.07 (95%-CI [0.03–0.11]). At the participant level, the within-subject SMD of model-predicted probabilities was dz = 1.27 (Hedges' g = 1.24; n=35), indicating a consistent modulation across subjects. This finding aligns with our hypothesis that the protocol induces excitatory behavioral effects, and replicates prior findings observed in non-human primates[2]. This excitatory behavioral effect on contralateral saccades was not observed for stimulation to left versus right M1 (details reported below; for visual comparison of FEF-sham and FEF-M1 refer to **Figure S1**). Comparison of each FEF condition relative to baseline (sham) showed that the sham condition provided a baseline between left and right FEF TUS effects, whereby left FEF TUS increased the proportion of rightward saccades (b = -0.12, p = 0.023), and right FEF showed a trend toward more leftward saccades (b = 0.12, p = 0.08). These results highlight the specificity of the effects to the FEFs and provide converging evidence of direct TUS-induced behavioral changes in humans.

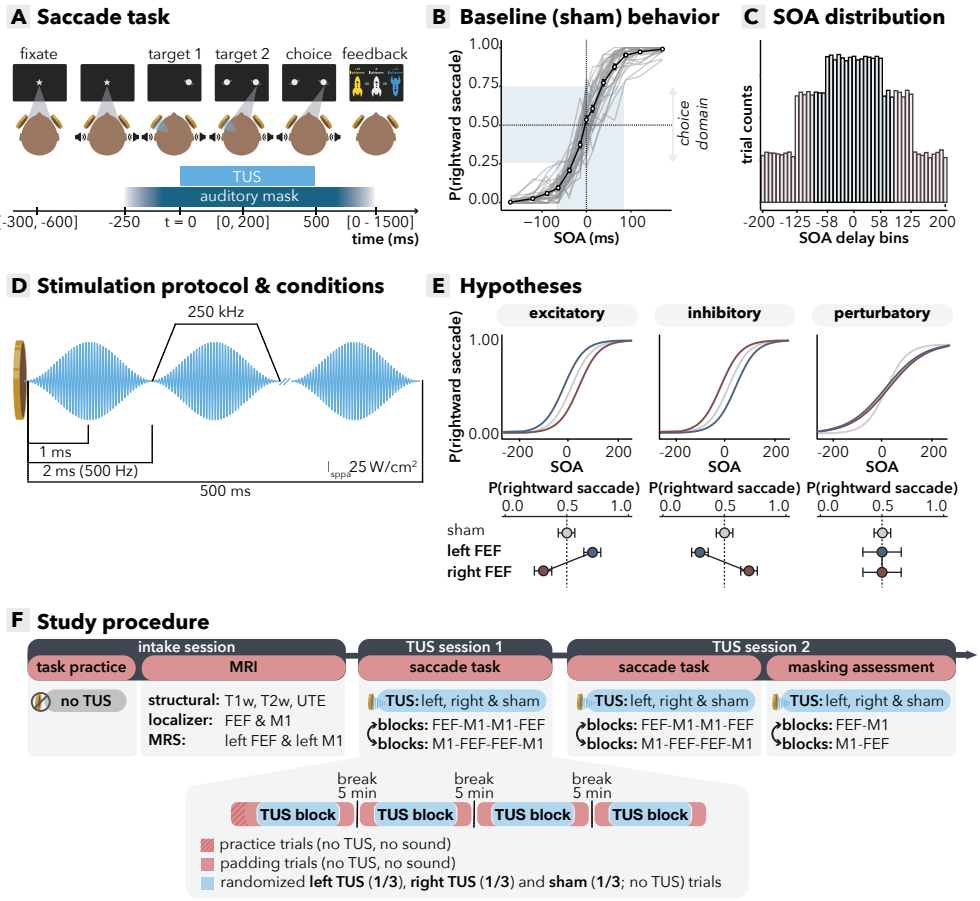


Figure 1 | Study design and baseline behavioral results. (A) Task design. Two lateralized targets appeared, spaced by a variable SOA (0-200 ms). Participants (n = 35) made saccades to the first target and received feedback (correct: +10 points; incorrect: -10; correct but late: 0). TUS was applied to left or right hemisphere, paired with a masking sound; sham trials consisted of auditory mask only. **(B)** Baseline behavioral performance followed a typical psychometric sigmoid function, with lower accuracy at short SOAs. Because TUS effects were expected to be most apparent under low sensory evidence, we defined the choice domain as delays at which performance was <75% correct (blue shading). Black line indicates group mean; error bars indicate standard error of the mean across participants; gray lines represent individual participants. Inferential statistics used continuous SOAs, but for visualization data is binned into approximately equal-trial bins with SOA intervals of [0], [8.3–25.0], [33.3–50.0], [58.3–75.0], [83.3–100.0], [108.3–141.7], [150–200]. **(C)** SOAs frequency, with shorter delays oversampled to increase sensitivity to biasing effects under high uncertainty. Blue bars indicate the empirically defined average choice domain. **(D)** TUS was delivered for 500ms per trial, starting at onset of the first target (fundamental frequency 250kHz; pulse repetition frequency 500Hz; intensity in free water I_{SPPA} = 25 W/cm²), with 1-ms ramp-up and ramp-down per pulse. Conditions comprised left/right FEF, left/right M1, and sham. **(E)** Hypothesized TUS effects. Predicted behavioral patterns if TUS induces a (i) net excitatory bias (more contralateral saccades), (ii) net inhibitory bias

(more ipsilateral saccades), or (iii) perturbatory effect (reduced accuracy/increased noise, but no bias). Note that this panel visualizes predicted patterns and does not represent empirical data. **(F)** Participants completed one intake and two TUS sessions. Intake included task practice, structural MRI for neuronavigation, functional localizers for individual targeting, and baseline GABA+ MRS (left FEF and M1 only). Each TUS session comprised four blocks of the saccade task with counterbalanced FEF/M1 block order. A masking assessment at the end of the final session evaluated whether participants could detect stimulation presence/side under auditory masking.

There was substantial interindividual variability both in baseline (sham) directional bias (**Figure 1B**) as well as in the susceptibility of saccade direction to TUS stimulation (**Figure 2C**). Therefore, we next asked whether the baseline neural inhibitory tone in participants' FEF could explain interindividual differences in TUS susceptibility. Note that we measured only left hemispheric MRS (in FEF and M1, for details see methods). We found that changes in saccade bias induced by left FEF TUS relative to sham were predicted by baseline FEF GABA+ levels (condition (left FEF/sham) x FEF GABA+: b = -0.21, 95%-CI [-0.39, -0.04], $c^2(1) = 5.6$, p = 0.017; **Figure 3A**). Specifically, higher baseline GABA+ levels in the left FEF were associated with a stronger rightward bias on sham trials (sham x FEF GABA+: b = 0.14, 95%-CI [0.04, 0.24], $c^2(1) = 7.0$, $R^2 = 0.096$, p = 0.008; **Figure 3B**, top). Importantly, following TUS stimulation, this relationship of baseline GABA+ and rightward bias disappeared (left FEF x FEF GABA+: b = -0.08, 95%-CI [-0.22, 0.07], $c^2(1) = 1.1$, p = 0.3; **Figure 3B**, bottom). Thus, TUS increased contralateral responding predominantly in participants with lower baseline GABA+ levels in the FEF (voxel placement: **Figure 3C**).

M1 TUS does not affect saccade choices

To ensure that the observed effects of TUS on saccade direction were specific to the FEF, in half of the stimulation trials the left and right M1 (hand area) were stimulated as control regions. Similar to the FEF stimulation, we used an fMRI functional localizer to determine the participant specific target for M1 TUS (**Figure 4A**). Again, post hoc acoustic wave propagation simulations confirmed that we successfully targeted these regions (**Figure 4B**). As expected, left and right M1 TUS did not induce significant differences in contralateral saccades, further supporting the specificity of the observed effects to the FEF and ruling out potential confounds (b = -0.09, 95%-CI [-0.22, 0.04], $c^2(1) = 1.8$, p = 0.12; **Figure 4C**).

However, at delay 0, where both targets appear on the screen simultaneously, we observed a qualitative difference in rightward saccade probability in M1 stimulation (**Figure 4D**) that mirrors the pattern of FEF stimulation. We remark that these zero-delay trials are qualitatively different from the delay trials because of simultaneous target presentation, which may be easily identified and thus trigger different cognitive

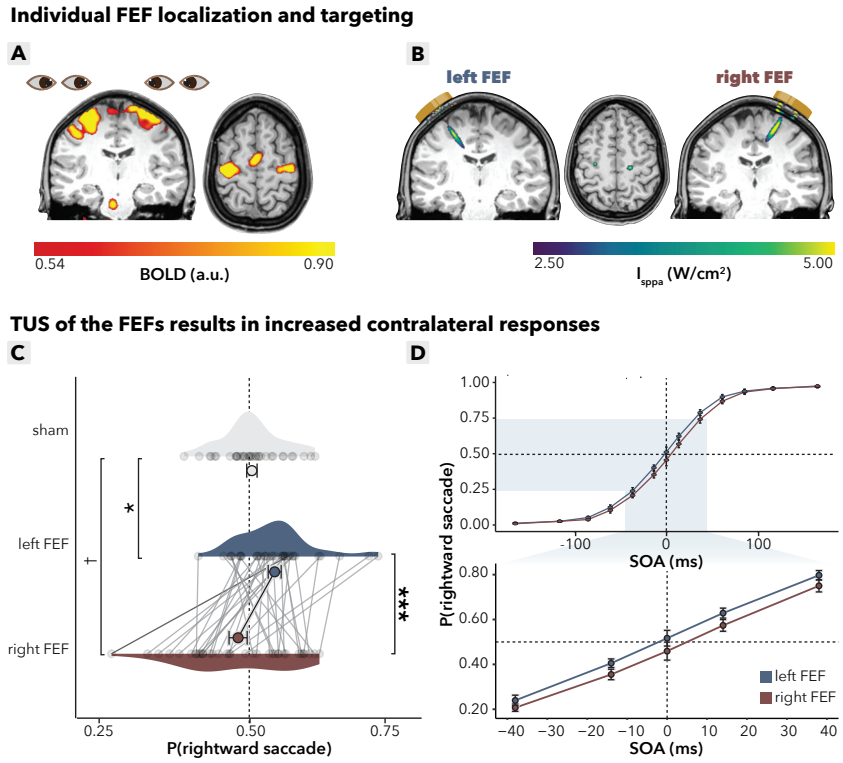


Figure 2 | Main behavioral TUS effects. (A) Individual FEF localization. BOLD responses of left and right FEF for a single subject. Participants performed a functional localizer during the intake session in which they alternated between left/right saccades and fixation, allowing for individual localization of the FEFs. **(B)** Acoustic simulation of TUS wave propagation. Acoustic simulations of left and right FEF for a single subject are shown. The simulation depicts the estimated intracranial intensity (I_{sppa}) with an intensity cutoff at the full width at half maximum (FWHM). **(C)** FEF TUS effects (n = 35). Choice-domain average effects. Gray dots represent individual participants' mean saccadic directions within the choice domain. Colored dots represent group means, error bars indicate the standard error of the mean (S.E.M.). *** p<0.001, * p<0.05, † p<0.1. **(D)** FEF TUS effects (n = 35). Stimulation of the left and right frontal eye fields (FEF) led to increased contralateral saccades, particularly within the choice domain (highlighted in light blue, bottom). Compared to each other, left FEF stimulation produced more rightward saccades, while right FEF stimulation led to more leftward saccades. Data are binned for visualization purposes into bins of approximately equal numbers of trials, resulting in SOA intervals of [0], [8.3 – 25.0], [33.3 – 50.0], [58.3 – 75.0], [83.3 – 100.0], [108.3 – 141.7], and [150 – 200]. Bins are symmetric for negative values. Note that all inferential statistics were performed using continuous SOAs. Dots represent the group mean per bin, and error bars indicate the S.E.M. across participants. **(C)**, **(D)** Exact p values are presented in the main text. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

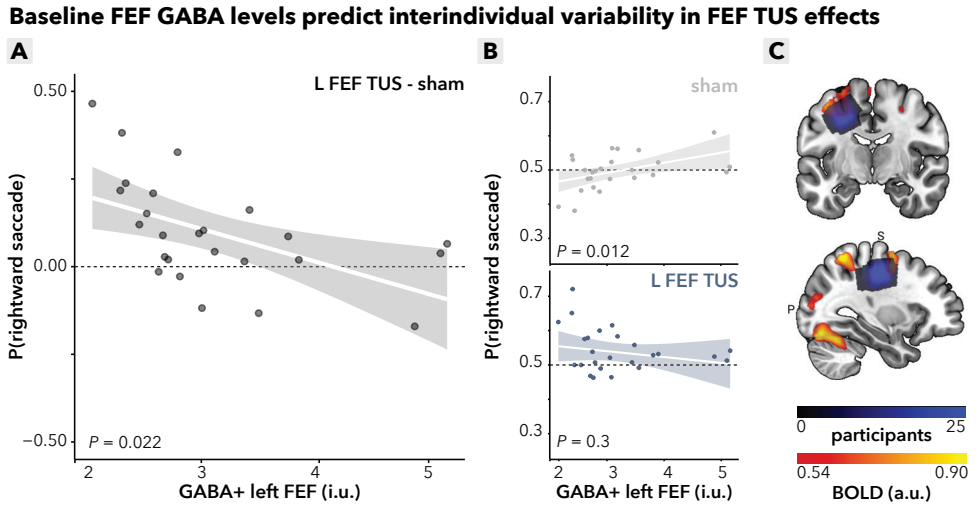


Figure 3 | Interindividual variability in FEF effects. (A) FEF GABA+ predicts FEF TUS effects. The relationship between baseline GABA+ levels in the left FEF and the effect of TUS on saccadic bias, calculated as the difference in probability of making a rightward saccade between left FEF TUS and sham conditions. Higher baseline FEF GABA+ levels correlate with a weaker TUS effect on rightward saccades (p = 0.022). The line is a linear fitted group mean with a 95% confidence interval around the fit; each dot represents a participant. **(B)** FEF GABA+ predicts baseline saccade behavior. *Top:* baseline left FEF GABA+ levels significantly correlate with rightward saccade probability during sham stimulation alone (p = 0.012). Bottom*:* under left FEF TUS, this correlation with rightward saccade probability is absent (p = 0.3). The line is a linear fitted group mean with a 95% confidence interval around the fit; each dot represents a participant; each dot represents a participant. **(C)** MRS voxel placement. Magnetic resonance spectroscopy (MRS) voxel placement for measuring GABA+ concentrations in the left frontal eye field (FEF). Color overlays represent GABA+ concentration distributions in each region. **(A)**, **(B)** Exact p values are presented in the main text. Statistical significance was determined using two-sided binomial linear mixed effects regressions. No multiple comparisons were applied.

processes[37,38]. This observation is further explored in the supplementary documents (**Supplementary Documents S.1**). For consistency across all analyses, we selected the choice domain data without these delay-0 trials.

Critically, a formal side-by-region (FEF/M1) comparison revealed a significant interaction (b = 0.23, 95%-CI [0.03, 0.43], $c^2(1) = 5.1$, p = 0.025) between stimulation side (left/right) and region (FEF/M1). The side-by-region model demonstrated stronger side modulation in FEF than in M1 by 26% (ROR = 1.26, 95%-CI [1.03-1.56]), corresponding to Cohen's d (logit) = 0.13 (95%-CI [0.01-0.24]) indicating a small but reliable interaction. This indicates that the TUS effects influencing choice bias and saccade behavior are specific to FEF stimulation and not to M1 stimulation, and excludes the possibility that these effects are driven by confounds such as auditory or somatosensory stimulation.

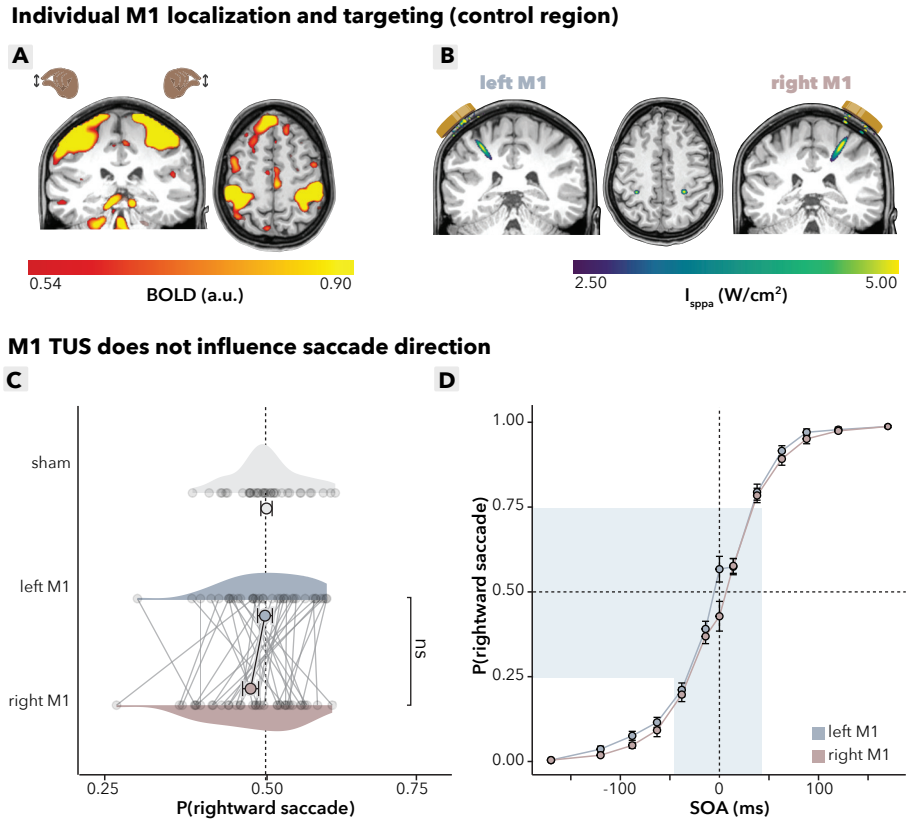


Figure 4 | Control analyses: M1 TUS effects. (A) Individual M1 localization. BOLD responses of left and right M1 for a single subject. Participants performed a functional localizer during the intake session in which they alternated between left- and right-hand tapping movements (index finger and thumb), allowing for individual localization of the hand area in M1. **(B)** Acoustic simulation of TUS wave propagation. Acoustic simulations of left and right M1 for a single subject are shown. The simulation depicts the estimated intracranial intensity (I_{sppa}) with an intensity cutoff at the full width at half maximum (FWHM). **(C)** M1 TUS effects (n = 35). Choice-domain average effects. Gray dots represent individual participants' average saccadic direction within the choice domain. Colored dots represent group means, error bars indicate the standard error of the mean (S.E.M.) with no statistically significant group effects observed. ns p>0.1. **(D)** M1 TUS effects (n = 35). Stimulation of the left and right M1 did not result in significant shifts in contralateral saccades within the choice domain (highlighted in light blue, bottom). Compared to each other, left and right M1 stimulation showed no differences in saccadic direction. Data are binned for visualization purposes into bins of approximately equal numbers of trials, resulting in SOA intervals of [0], [8.3 – 25.0], [33.3 – 50.0], [58.3 – 75.0], [83.3 – 100.0], [108.3 – 141.7], and [150 – 200]. Bins are symmetric for negative values. Note that all inferential statistics were performed using continuous SOAs. Dots represent the group mean per bin, and error bars indicate the S.E.M. across participants. **(C)**, **(D)** Source data are provided as a Source Data file. Exact p values are presented in the main text. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

This finding reinforces the conclusion that TUS selectively modulates behavior via its impact on the FEFs. Finally, between-variability in the effects of TUS in M1 in saccade direction could not be explained by subjects' baseline GABA+ levels in M1 (condition (left M1/sham) x M1 GABA+: b = 0.06, 95%-CI [-0.14, 0.26], $c^2(1) = 0.3$, $R^2 = 0.1$, $p = 0.5$; **Figure 5A-C**). These findings underscore the specificity of the TUS effects to the FEFs and provide additional evidence against potential confounds in the study design.

Control and follow-up analyses

TUS does not disrupt overall performance

In order to assess potential perturbatory effects of FEF TUS on performance, we completed a regression with 'correct response' on TUS and sham trials in the FEF blocks as dependent variable and side, region and delay as independent variables. Note that again zero-delay trials are excluded from this analysis because no correct response can be defined. For a binary choice task, sensory noise is directly reflected in the overall accuracy (i.e., the slope of the psychometric curve is inversely related to the variance of the underlying signal probability distribution). There was no perturbatory effect of TUS on performance (TUS vs. sham: b = 0.07, 95%-CI [-0.13, 0.27], $c^2(1) = 0.5$, $p = 0.5$).

Additionally, we performed supplementary analyses (**Supplementary Documents S.2**) examining estimation of bias, including effect size in decision curve shift (horizontal bias), slope, and lapse rate to confirm that the observed TUS effects were specific to bias and not confounded by changes in slope or lapse rate.

Online TUS effects are immediate and short-lived

Having demonstrated that TUS of FEF has an excitatory effect and that this effect is specific to stimulation of FEF, we next assessed the duration and persistence of TUS effects on saccade direction. This is critical to characterize the temporal dynamics of ultrasonic neuromodulation, in terms of how fast effects arise, and whether they persist into the next trial. Slow and sustained effects suggest early-phase plasticity mechanisms to drive the observed behavior, while fast and temporally precise effects suggest modulation of spiking activity. The latter would enable TUS to be used for cognitive chronometry: to disentangle the functional contributions of brain regions and circuits across time.

First, we examined whether there were any carryover effects of stimulation on sham trials that followed TUS trials. First, when analyzing sham trials during FEF blocks,

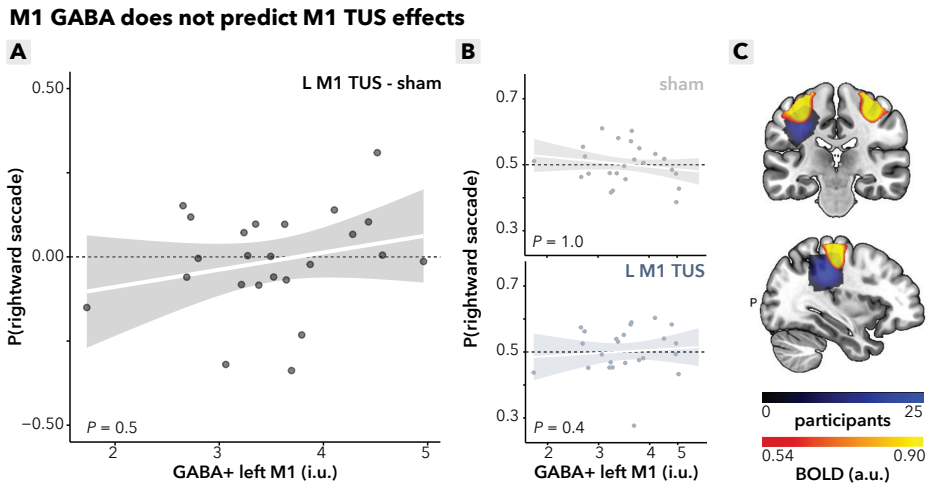


Figure 5 | Control analyses: interindividual variability in M1 TUS effects. (A) M1 GABA+ does not predict TUS effects. Baseline GABA+ levels in the left M1 do not correlate with TUS-induced saccadic bias, calculated as the difference in probability of making a rightward saccade between left M1 TUS and sham conditions (p = 0.5). The line is a linear fitted group mean with a 95% confidence interval around the fit; each dot represents a participant. **(B)** M1 GABA+ does not predict baseline saccade behavior. Top: baseline left M1 GABA+ levels do not significantly correlate with rightward saccade probability during sham stimulation alone (p = 1.0). Bottom: under left M1 TUS, this correlation with rightward saccade probability remains absent (p = 0.4). The line is a linear fitted group mean with a 95% confidence interval around the fit; each dot represents a participant. **(C)** MRS voxel placement. Magnetic resonance spectroscopy (MRS) voxel placement for measuring GABA+ concentrations in the left motor cortex (M1). Color overlays represent GABA+ concentration distributions in each region. **(A)**, **(B)** Source data are provided as a Source Data file. Exact p values are presented in the main text. Statistical significance was determined using two-sided binomial linear mixed effects regressions. No multiple comparisons were applied.

no significant carry-over effects were observed. More specifically, saccade direction was not affected by the side stimulated on the preceding trial (FEF_{t-1} (left/right): b = 0.15, 95%-CI [-0.07, 0.37], $c^2(1) = 1.8$, p = 0.18; BF_{01} = 2). However, when pooling together sham trials during FEF and M1 blocks, participants made significantly more ipsilateral saccades on these trials, directing their saccades toward the side stimulated on the preceding trial ($side_{t-1}$: b = 0.10, 95%-CI [-0.18, -0.02], $c^2(1) = 5.3$, p = 0.021). Crucially, this ipsilateral bias did not differ between FEF and M1 stimulation ($side_{t-1}$ x $region_{t-1}$: b = 0.02, 95%-CI [-0.06, 0.10], $c^2(1) = 0.3$, p = 0.6, BF_{01} = 12; **Figure 6A**), and can therefore not explain the observed specific effects on contralateral saccades following FEF (but not M1) TUS. In **Supplementary Documents S.3**, we will briefly further discuss the non-specific (potentially attentionally driven) after-effects.

Finally, to assess the immediacy of TUS effects relative to stimulation onset, we quantified TUS effects on the fastest saccades, defined as trials with a saccade reaction time below the median of 265 ms. Even on this subset of trials where participants received less than 265 ms of stimulation prior to saccade onset, TUS significantly shifted saccade direction contralaterally (FEF (left/right): b = -0.32, 95%-CI [-0.59, -0.07], $c^2(1)$ = 6.1, p = 0.013). In contrast, no significant saccade bias was observed for left versus right M1 stimulation on fast trials (M1 (left/right): b = -0.15, 95%-CI [-0.38, 0.07], $c^2(1)$ = 1.8, p = 0.19). Taken together, our results highlight the specificity and speed of TUS effects on saccade direction, reinforcing that they are immediate, fast and specific to the FEF.

Masking assessment

To estimate the potential impact of auditory or somatosensory confounds[9,39–42], we included a masking assessment at the end of the second TUS session (**Figure 1F**). This assessment allowed us to verify that potential confounds could not explain the observed dissociation of TUS effects over FEF versus M1. Participants received stimulation (or sham) repeatedly either over FEF or M1 (in blocks), all with an auditory mask, and reported i) whether they perceived stimulation, and ii) on what side (forced choice, left vs. right). First, participants reported perceiving stimulation more frequently on TUS trials compared to sham (stimulation (TUS/sham): b = 2.7, 95%-CI [1.9, 3.4], $c^2(1)$ = 53.6, p < 0.001). Crucially, however, this ability to detect TUS versus sham did not differ between conditions (region (M1/FEF): b = 0.25, 95%-CI [-0.36, 0.86], $c^2(1)$ = 0.7, p = 0.4, BF_{01} = 4.1; **Figure 6B**). Second, on TUS trials, participants were biased to report perceiving stimulation *contralaterally* to the side of actual stimulation (side: b = -1.2, 95%-CI [-2.1, -0.3], $c^2(1)$ = 7., p = 0.008, BF_{01} < 0.001; **Figure 6C**). However, again this contralateral reporting bias was not significantly different between FEF and M1 stimulation (side (left/right) x region (FEF/M1): b = -0.5, 95%-CI [-1.5, 0.5], $c^2(1)$ = 1.1, p = 0.3). In **Supplementary Documents S.3**, we further discuss this contralateral TUS perception in the context of lateralized non-specific aftereffects reported above. Taken together, while putative confounding factors were present in our study, their effects were present in both FEF and M1 conditions, and thus crucially cannot account for our main findings. More broadly, this masking assessment confirms the presence of putative confounding factors and emphasize the importance of active control conditions for online TUS protocols.

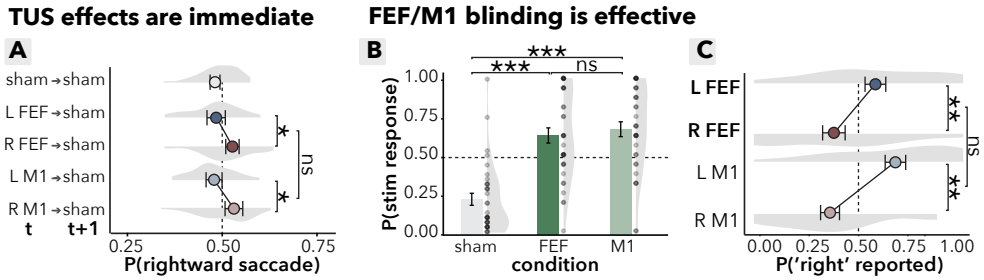


Figure 6 | Effects of online TUS and masking assessment. (A) TUS after-effect assessment on sham trials (n = 35). Each sham trial was labeled based on the preceding trial's stimulation condition (e.g., "L FEF → sham" indicates a sham trial following left FEF stimulation). Density clouds represent participant distributions of saccade bias on sham trials that followed stimulation trials; dots represent group mean; error bars represent the standard error of the mean (S.E.M.). * p<0.05, ns p>0.1. Participants made significantly more ipsilateral saccades on sham trials following stimulation, but importantly and unlike the main TUS effect, this effect was not specific to FEF. **(B)** Masking, perceived stimulation (yes/no). Probability of reported stimulation perception across sham, FEF, and M1 conditions (n = 34). Density clouds represent participant distributions, bars indicate group means, and error bars show the S.E.M. *** p<0.001, ns p>0.1. Participants were significantly more likely to perceive stimulation during TUS trials compared to sham trials, highlighting that sham conditions alone may not fully account for TUS effects. No significant difference was observed between FEF and M1 conditions. **(C)** Masking, perceived stimulation side (left/right). Probability of reported stimulation side perception in the masking assessment (n = 34). Density clouds represent participant distributions, dots represent group means, and error bars show the S.E.M. ** p<0.01, ns p>0.1. Participants were more likely to perceive TUS contralateral to the actual stimulation site. This effect was consistent across FEF and M1 regions, as indicated by the absence of a significant side-by-region interaction, supporting the robustness of the active control design. **(A)**, **(B)**, **(C)** Source data are provided as a Source Data file. Exact p values are presented in the main text. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

DISCUSSION

This study provides evidence for effective online transcranial ultrasound stimulation (TUS) on saccadic decision-making in humans. These results advance our understanding of the underlying neural mechanisms contributing to interindividual differences. We found that short TUS pulse trains (500 ms) to the frontal eye fields (FEF) but not the primary hand motor cortex (M1) have immediate, short lived effects promoting contralateral saccades. Importantly, the effect of FEF TUS is associated with individual inhibitory tone, as indexed with magnetic resonance spectroscopy (MRS). These findings provide behavioral and neurophysiological evidence for the direct effects of TUS on human brain function, establishing its potential as temporally specific neuromodulatory tool for advancing fundamental neuroscience and enhancing our understanding of temporally dynamic brain-behavior relationships.

In recent years, evidence has emerged for sustained and early-phase plasticity effects of TUS in humans[7,8,43]. Despite advances in in-vitro and animal models demonstrating immediate neural effects of TUS[2,4,6], the translation of these findings to humans has remained scarce[10] or potentially marred by confounds[9]. To address this, we adapted a well-established animal TUS protocol for human application, targeting the left and right FEF while participants performed a saccade choice task. This approach allowed us to assess the immediate behavioral effects of TUS on saccadic choices.

Previous lesion and brain stimulation studies have demonstrated the key role of FEF in oculomotor control, specifically mediating contralateral saccade generation[13–24]. Here, we reveal that TUS over FEFs increased the selection of contralateral saccades, particularly during with short SOAs (**Figure 2C-D**). These findings, in conjunction with this past literature, suggest that the saccade bias elicited by TUS reflects a modulation of the oculomotor control by the FEF. An alternative interpretation of the results, inspired by a discussion with a reviewer, may be that the bias effects reflect changes in perception itself. The FEF has an established role in the control of visuospatial attention[17,44] and control of sensory representations in visual cortex[45,46]. Considering that attention affects temporal order judgments[47], a complementary interpretation would be that FEF TUS altered the perceived temporal priority. Regardless, our results emphasize TUS induced changes in the sensorimotor transformations performed by the FEF in the context of eye movements.

Taken together, the directionality and context of the effect are consistent with a net facilitatory effect on FEF-mediated contralateral selection, replicating earlier work in non-human primates[2]. Note that we deliberately avoid characterizing TUS as intrinsically excitatory or inhibitory. Net effects are parameter-, physiology-, and

state-dependent and reflect the local circuitry engaged at the time of stimulation: for example, a net excitatory effect could arise from either excitation of glutamatergic neurons or inhibition of GABAergic neurons. The current literature does not support a one-size-fits-all excitatory TUS recipe independent of the target region and brain state.

To demonstrate the spatial specificity of TUS, and to rule out the possibility that auditory or somatosensory confounds could drive the observed effects, we alternated blocks of FEF stimulation with blocks targeting an active control site – the hand area of the primary motor cortex (M1). Given that on a group-level, participants were able to distinguish TUS from sham trials, we highlight here that control conditions must go beyond a generic sham and explicitly match peripheral confounds and expectancy; a consideration that is often overlooked in non-invasive brain stimulation[9,48]. Importantly, we observed no significant changes in saccadic behavior following M1 stimulation (**Figure 4C**).

Our results show that TUS biases responses on trials with short SOAs, where there was high uncertainty about the correct response (**Figure 2D**). On those trials, which we purposefully oversampled, we hypothesized FEF activity is more sensitive to be nudged by TUS to tip the balance in the opposite direction. This effect supports the hypothesized role of TUS in modulation of ongoing FEF computations, rather than the direct induction of saccades. TUS employs sound waves that are beyond the audible spectrum to mechanically engage with neuronal tissue, influencing characteristics such as membrane capacitance and the activity of mechanosensitive ion channels[6]. Mechanistically, unlike magnetic or electrical stimulation techniques that can directly induce neural firing (motor evoked potentials[33]; visual evoked potentials[34,35]), TUS does not directly evoke neural firing but rather modulates ongoing neural activity through subthreshold modulation and subtly nudges neural activity without inducing immediate or large-scale neural responses[36]. Thus, two principles converge here: low-evidence trials are most susceptible to small motor biases, and temporal-order judgments are sensitive to perceptual and attentional biases[17,44–47]. Together, these predict that TUS effects should peak at short SOAs. Consistent with this, we observe the largest biases in this short SOA regime.

Finally, we established the high temporal specificity of this TUS protocol by showing it has immediate effects that occur only during stimulation trials, and do not persist into follow-up sham trials (**Figure 6A**). This highlights the specificity of TUS effects to the stimulation period itself. Notably, these effects emerge very rapidly: even in trials with the fastest reaction times, where participants made saccades before the TUS duration was complete (i.e., less than 265 ms), the TUS effect was still evident and

specific to the FEF. This high temporal specificity is crucial for using TUS as a tool to probe brain function with millisecond precision, aligning with principles of mental chronometry to better understand the neural dynamics underlying rapid cognitive and motor processes.

The bias toward contralateral saccades suggests a net facilitatory effect of TUS on the FEF. However, this does not necessarily imply that TUS directly excites neuronal tissue or targets specific neuron types. TUS can influence both excitatory and inhibitory neurons by altering action potential thresholds[6,49,50]. The effects of TUS can, therefore, be predicted by the baseline state of the neuronal populations involved[31,51–53,53–55]. Interestingly, our data indicate that such baseline differences in neuronal state indeed modulate the extent of TUS effects.

At baseline, participants exhibited individual biases in saccade direction, with some showing a preference for leftward and others for rightward saccades (**Figure 1B**). These behavioral baselines were mirrored by neural differences, as individuals with a lower inhibitory tone, as quantified by GABA+ levels in the left FEF, tended to have a stronger intrinsic leftward bias (**Figure 3B**, top). While this might appear counterintuitive, it aligns with earlier findings from FEF MRS work[32], which demonstrated that higher GABA+ levels are associated with better suppression of distractors. In this model, cortical regions execute actions via excitatory neurons while inhibiting competing cortical regions — a process regulated by local inhibitory interneurons. Furthermore, in other motor regions, GABA+ has previously been shown to play a critical role in behavior in its relation to motor learning through action plan tuning and long-term potentiation[31,56] and FEF GABA+ predicts saccade behavior[32].

Our findings support the framework proposing that local inhibitory interneurons play a role in distraction suppression and ultimately saccade execution[32,57]. Individuals with higher GABA+ levels in the left FEF were better able to inhibit competing right FEF projections, resulting in a stronger rightward bias at baseline. Furthermore, baseline physiological inhibitory state predicted the magnitude of the TUS-induced behavioral effects, such that participants with higher cortical inhibitory tone show a larger contralateral bias under sham, and exhibit a smaller TUS-induced shift, whereas those with lower inhibitory tone exhibit a larger shift (**Figure 3A**). These findings underscore the state-dependent nature of TUS and highlight the importance of considering baseline neural states when interpreting its effects. Similar to other non-invasive brain stimulation methods, the variability in TUS effects appears to be modulated by the baseline excitatory/inhibitory balance of the targeted neural populations.

Notably, the effects of TUS observed in humans are less pronounced than those documented in earlier animal work[2]. At least two reasons can be conceived. First, human FEF is larger in absolute terms, reducing the relative territory of FEF reached by the TUS focus. This might lead to a reduced efficacy of stimulation, as the average TUS intensity across the entire FEF will be lower in humans than in macaques, even with comparable peak intensities. Second, the FEF circuitry in animals appears more lateralized compared to humans[58,59], which may serve to explain the somewhat weaker effect of particularly right FEF stimulation relative to sham. A left-right asymmetry in lateralization strength is plausible given the well documented right-hemisphere dominance of human attention networks and hemispheric differences in fronto-parietal circuitry[60–63].

TUS is associated with auditory and somatosensory confounds, that could putatively drive behavioral effects masquerading as effects of TUS[9,39]. Therefore, in this study we included a number of careful controls to exclude this possibility. Importantly, the observed effects on contralateral saccades were specific to TUS over FEF, and not observed following M1 TUS. Crucially, however, confound effects were not different between these conditions, as evidenced by a masking assessment following the final session. First, while participants were able to differentiate between sham and stimulation conditions, this did not differ between M1 and FEF TUS (**Figure 6B**). Second, while participants were able to distinguish left versus right TUS stimulation, again this did not differ between M1 and FEF TUS (**Figure 6C**). Taken together, these potential confound effects cannot explain the main finding. Nevertheless, these observations emphasize the importance of including an active control condition rather than relying solely on a sham condition in TUS studies.

A final point of consideration is that while this study investigated baseline GABA+ levels, we did not record GABA+ as an outcome variable of TUS. Instead, our focus was on the immediate effects of TUS and its relationship to individual differences. A key direction for future research would involve assessing MRS-based changes in GABA+ to examine the potential for longer-lasting plasticity effects[8].

In summary, we demonstrated we can bias human choices using fast ultrasonic neuromodulation. Transcranial ultrasound over the frontal eye fields, a model circuit for human decision-making, induced facilitatory effects promoting contralateral saccades. These effects were specific to FEF, immediate, present only during stimulation trials, and emerged rapidly. We showed that interindividual variability in TUS effects could be explained by the inhibitory tone of participants' FEF at baseline, with TUS modulating this balance to bring individuals into a more uniform physiological state. These findings contribute to advancing brain research by demonstrating the

feasibility and efficacy of immediate TUS effects in humans. They emphasize the importance of considering interindividual variability and active control conditions in study designs. Together, this work opens new avenues for future studies aiming to explore causal brain function with high temporal precision and to develop innovative therapeutic interventions.

METHODS

Participants

We preregistered (https://doi.org/10.17605/OSF.IO/K5P2M) a target sample size of 35 participants, based on a small to medium effect size of f ~ 0.35, with an alpha level of 0.05 and a power of 80% (calculated using G*Power 3.1[64]). We restricted inclusion to right-handed adults to reduce variance from hemispheric dominance and lateralization[65–71]. Participants were screened on medical history to exclude putative participants with a history of brain surgery, serious head trauma, epilepsy, convulsion, or seizure, as well as the presence of implanted metal in the head or upper body, diagnosed neurological or psychiatric disorders, and consumption of either more than four alcoholic units within the preceding 24 hours or any recreational drugs within the past 48 hours.

Thirty-nine participants were enrolled in the experiment. Three participants were excluded due to technical eye-tracking failure (frequent signal loss/recalibrations or inability to complete the task), and these sessions were replaced until the target sample was reached. We further excluded one participant as an accuracy outlier, defined by Tukey IQR rule (accuracy < Q1 − 1.5·IQR or > Q3 + 1.5·IQR), |z| > 3 based on the sample mean/SD, and/or robust MAD-z > 3.5) and accompanied by a non-monotone psychometric pattern—indicative of misunderstanding/non-compliance rather than sensory uncertainty (**Figure S4**).

Thirty-five participants (M_{age} = 24.1, SD_{age} = 3.2, range = 20 – 32; 15:20 female:male; right-handed) were included in the final analysis. Written informed consent was obtained from all participants in accordance with the Declaration of Helsinki, and the experimental procedures were approved by the local ethics committee (CMO2022-15953, Commissie Mensgebonden Onderzoek Oost-Nederland. Participants received €90 in total for participation, plus a performance-based bonus of up to €5 per stimulation session. Recruitment was open via university and community channels and did not restrict participation by sex, gender identity, ethnicity, or socioeconomic status. Gender identity was not systematically collected. Sex was recorded by self-

report, and not included as a factor in the study design or statistical analyses because the study was not designed or powered for sex-stratified comparisons and no a priori hypotheses regarding sex differences were specified.

Saccade task results include all 35 participants. Some participants were excluded from the following analyses: For 10 participants, magnetic resonance spectroscopy (MRS) GABA+ acquisition was of poor quality; hence, all MRS results are based on the data of 25 participants (M_{age} = 24.7, SD_{age} = 3.1, range = 20 – 33, 11:14 female:male). One participant did not complete the final stimulation session and was thus excluded from all masking assessment analyses; hence, all masking assessment analyses are based on 34 participants (M_{age} = 24.9, SD_{age} = 3.1, range = 20 – 33, 15:19 female:male). Importantly, since participants experienced all conditions in each TUS session, this participant was still included in the main analyses.

Study overview

The study comprised three double-blind, within-subject sessions, with an interval of approximately one week (and up to three months) between sessions (**Figure 1F**), scheduled at the same times of the day to reduce potential fluctuations in GABA+ induced by circadian rhythm. In the initial intake session, participants engaged in a practice of the saccade task without TUS delivery. Subsequently, they entered the MRI scanner to acquire structural scans. Additionally, participants completed FEF and M1 functional localizers (described below), used to target TUS during the following brain stimulation sessions. Finally, we obtained measures of GABA+ concentration from the left hemispheric stimulation regions using single voxel MRS. MRS measurements were limited to the left hemisphere due to time constraints. The left hemisphere was prioritized over the right, given the stronger lateralization in attentional processes of the right hemisphere in humans[61,72].

The two successive brain stimulation sessions incorporated a mix of sham and TUS trials during the saccade task. Each session started with screening and a 45-minute preparation phase (see Neuronavigation below). Participants started with a practice block without TUS delivery or auditory masking, to reacquaint themselves with the task. For practice blocks where performance was below 60%, participants had to repeat the block. During the main task, no trials were repeated and there was no accuracy threshold. A padding block without TUS nor auditory mask bookended each TUS block. TUS transducer placement (either left and right FEF or left and right M1) was contingent on the blocks within the sequence. Sequence order was counterbalanced across participants in the two stimulation sessions (**Figure S5**). TUS blocks contained three conditions: 1) left and ii) right TUS paired with an auditory masking tone, and iii)

a sham condition with solely the auditory mask. The order of conditions within TUS blocks followed a pseudorandom pattern limited to a maximum of four consecutive trials of the same condition. At the session's conclusion, participants were queried about any symptoms they believed could be associated with TUS. This was only used for debriefing and is not further analyzed. Only after the final stimulation session, the efficacy of blinding was assessed during a short masking assessment.

Tasks

Saccade task

Each trial started with fixation on a star-shaped stimulus (0.25 x 0.25 degrees of visual angle) presented at the center of the screen. After fixation, there was a delay of 300-600 ms (jittered) before the first planet-shaped target (0.5 x 0.5 degrees of visual angle, acceptance window, 3 degrees) briefly appeared in either the left or right hemifield (10 degrees of visual angle left and right from the center of the screen). Simultaneously with the appearance of this first target, TUS was delivered, lasting 500 ms. The auditory mask began 250 ms before the first target appeared and lasted 1 second, fully padding the TUS delivery. The delay between the first and second planet-shaped target ranged from 0 to 200 ms.

Target delays exhibited a non-uniform distribution, with shorter delays clustered around the central peak and the longer delays at the tails. This distribution was designed to optimize the potential for TUS-induced behavioral modulation at relatively short target delays. At the same time, it allowed to make sure that TUS does not simply induce attention lapses, characterized by incorrect responses even with long delays.

Participants were instructed to execute a saccadic eye movement to the first appearing target in either the left or right hemifield. Trial completion was followed by feedback presented for 1 s indicating whether the correct target had been chosen within the designated time window (+10 points for correct responses, -10 points for incorrect responses, and 0 points for correct but late responses after 1500 ms; **Figure 1A**). In this gamified task, participants could earn points and a monetary bonus of up to €5 per stimulation session based on their overall performance. During stimulation sessions, they received mixed sham and TUS trials while an auditory mask was played to blind them to the different stimulation conditions and to prevent auditory confounding. The auditory mask corresponded to the specific condition, either masking or replicating the sound of stimulation (**Figure S6**).

Functional localizers

To prevent the risk of undershooting or missing the target due to the small ultrasound focus, we employed functional localizers to identify each participant's FEF and M1 with high fidelity[73]. The individual coordinates of interest determined using functional localizers were used for neuronavigation in the following brain stimulation sessions.

The FEF localizer[74,75] consisted of alternating 24-second blocks of saccadic eye movements and central fixation (**Figure S7A**). Participants followed and fixated on a target (visual angle, 1 x 1 degrees; white square; duration, 800 ms) presented at randomized screen positions located at the left, right, or center of the screen (target distance, 14 degrees). This eye movement and fixation sequence repeated six times. Assessment of the contrast between active eye movement blocks and baseline fixation blocks allowed for localization of the left and right FEFs.

The M1 localizer[76] consisted of alternating 16-second blocks of left and right finger movement (**Figure S7B**). Specifically, participants repetitively pinched their index finger and thumb together within the 16-second interval, alternating between left and right hands for six blocks per hand. This task enabled the establishment of contrasts between blocks of finger movement for each hand, providing information about left and right M1 activation.

Masking assessment

Following the final stimulation session, participants experienced a shorter series of sham and TUS trials involving both left and right FEF and M1. After each trial, they reported through button presses (up button for yes, down button for no) whether they believed they had received stimulation and on which side (left button for left, right button for right) they believed the stimulation was applied (**Figure S7C**). The order of the three conditions was fully randomized. Additionally, the sequence of stimulation regions was counterbalanced across participants.

Intake session

Task practice

Participants practiced the saccade task outside of the scanner for 15 minutes (198 trials) without delivery of TUS, to acquaint themselves with the task prior to the follow-up brain stimulation sessions.

Structural and functional MRI data acquisition

MRI scanning was performed at the Donders Centre for Cognitive Neuroimaging using a 3 Tesla Magnetom Skyra Scanner (Siemens AG, Erlangen, Germany) equipped with a 32-channel head coil. During structural scan acquisition, participants kept their eyes closed. High-resolution T1w scans were acquired (sagittal plane; repetition time (TR), 2700 ms; echo time (TE), 3.69 ms; flip angle, 9 degrees; voxel size, 0.9 x 0.9 x 0.9 mm; field of view, 230 mm) for MRS voxel placement, co-registration with the functional data, and neuronavigation for TUS delivery during stimulation sessions. To capture detailed skull morphology and tissue properties for acoustic simulations of ultrasonic wave propagation, T2w scans (sagittal plane; TR, 3200 ms; TE, 408 ms; flip angle, T2 var flip angle mode; voxel size, 0.9 x 0.9 x 0.9 mm; field of view, 230 mm), and UTE scans (transversal plane; TR, 3.32 ms; TE, 0.07 ms; flip angle, 2 degrees; voxel size, 0.8 x 0.8 x 0.8 mm; field of view, 294 mm) were acquired.

To functionally localize the stimulation regions, a Multi-Band sequence with an acceleration factor of four (MB4) was used (TR, 995 ms; TE, 32.8 ms; flip angle, 60 degrees; voxel size, 2.5 x 2.5 x 2.5 mm; field of view, 210 x 210 x 130 mm acquired in axial direction). Visual stimuli of the localizer tasks were presented at the rear bore face on a flat panel screen.

MRS data acquisition

Magnetic Resonance (MRS) Single Voxel Spectroscopy (SVS) of the left hemispheric target regions (FEF and M1) allowed for baseline GABA+ measures. For each ROI, after voxel placement based on the participant's T1-weighted scan, shimming was performed using FASTEST map[77,78] and a flip angle calibration process was carried out. For the FEF, the voxel was placed for each participant based on anatomical landmarks, at the intersect of the precentral gyrus, middle frontal gyrus and the superior frontal gyrus in the left hemisphere[11,74]. The M1 voxel was placed at the left hemispheric precentral knob located posterior to the intersection of the superior frontal sulcus that divides the superior from the middle frontal gyrus, and the precentral sulcus[79]. Baseline level of GABA+ was measured using the pulse sequence MEshcher-GArwood Point RESolved Spectroscopy (MEGA-PRESS: TR, 2000 ms; TE, 68 ms; voxel size, 2.0 x 2.0 x 2.0 cm; with VAPOR water suppression[80] 128 averages and water unsuppressed reference 16 averages) as introduced by Mescher and colleagues[81,82]. The baseline level of glutamate and glutamine (Glx) was quantified using the pulse sequence Point RESolved Spectroscopy (PRESS: TR, 20000 ms; TE, 35 ms; voxel size, 2.0 x 2.0 x 2.0 cm; with VAPOR water suppression 64 averages) as described by Marjańska and colleagues[83]. This data was not analyzed in the present paper.

During MRS acquisition, participants were instructed to close their eyes. MRS data acquisition started approximately 40 minutes after the end of the task practice participants completed prior to their MRI scan, thus making it unlikely that the brief task practice biased our baseline GABA+ estimates[84–90].

TUS sessions

Neuronavigation and hair preparation

The transducer was placed at the target location, and monitored throughout the session, using frameless stereotaxic neuronavigation (Localite Biomedical Visualization Systems GmbH, Sankt Augustin, Germany). We used participant specific T1w scans and x-, y-, z-coordinates of the left and right FEF and M1 derived from functional localizers. A reference tracker, five fixed markers (nasian, left and right eye, left and right ear), and 350 – 400 head surface markers were used to register the anatomical image to the participant's physical head. The two TUS transducers were also calibrated using a reference tracker and calibration plate. Transducers positions for the four stimulation regions were registered and quantified for acoustic ultrasonic wave propagation simulations.

Ultrasound gel (Aquaflex Ultrasound Gel, Parker Laboratories) was applied to the participant's head over stimulation regions, followed by placement of gel pads (Aquaflex Ultrasound Gel Pad, Parker Laboratories) between the gelled head and gel-covered transducers to eliminate air bubbles[1]. Refer to **Figure S8** for a schematic set-up.

TUS protocol

Ultrasonic stimulation was delivered using the NeuroFUS PRO system (Brainbox Ltd., Cardiff, UK) with two two-element ultrasound transducers (CTX250-009 and CTX250-014, 45 mm diameter, 250 kHz fundamental frequency, Sonic Concepts Inc., Bothell, WA, USA). We utilized a two-channel transducer to maximize the stimulation focal area. Although the TUS focus is characterized by a cigar-shaped profile that may extend into the white matter, it does not extend into the gray matter territory of neighboring cortical regions. The TUS protocol was adapted from Kubanek et al. (2020) (pulse duration: 2 ms; pulse ramp length: 1 ms; pulse repetition frequency: 500 Hz; pulse train duration: 500 ms; square-wave equivalent I_{sppa} in free water: 25 W/cm^2; effective I_{sppa} in free water: 9.37 W/cm^2; I_{spta} in free water: 9.37 W/cm^2; free-field pressure: 0.9 MPa; MI: 1.76; MI_{tc}: 1.25; **Figure 1D**, **Figure S12**). Relative to the original simulation protocol in non-human primates, we used ramped pulses with an auditory

mask to minimize auditory co-stimulation, increased free-water I_{sppa} to offset skull attenuation, and a longer pulse train duration to match human saccade preparation and execution.

Although squared and sinusoidal ramped pulses have the same integral energy content, it is important to note that squared wave pulses have associated limitations. A squared pulse encompasses a constant intensity peak for a longer duration due to their clear onset and offset, whereas a sinus-shaped pulse exhibits a gradually increasing and decreasing peak that is never fully off. While low-intensity ultrasonic waves are beyond the range of human hearing, the on-offset of the squared pulse is detectable by humans, increasing the likelihood of auditory confounds, and thus contributing to a clearer temporal profile of stimulation[4,91]. Furthermore, since humans have a thicker skull than macaques, a higher free-field I_{sppa} was applied (25 W/cm^2) to match the realized intracranial intensity across species. Moreover, we adjusted the total stimulation duration to the average human saccade duration.

The temperature rise (ΔT) remained below two degrees Celsius and the derated intracranial mechanical index (MI) below 1.9 matching ITRUSST recommendations[92]. During both sham and TUS trials, an auditory mask was played through bone conducting headphones (AfterShokz, New York, US). TUS was delivered during the task through serial commands in a PsychoPy script (PsychoPy 2021.2.3[93]).

Behavioral acquisition

Oculomotor behavior during the saccade task was tracked using Eyelink1000 PLUS (SR Research). Specifically, saccadic eye movements of the dominant eye were tracked from a distance of 80 cm between eye tracker and chinrest (**Figure S8A**). Saccades were detected by the eye tracker's online velocity/acceleration parser using the standard cognitive-task thresholds (velocity 30°/s, acceleration 8000°/s², motion criterion of 0.15°). From the exported raw data file, we used the Python-generated event markers to identify trial boundaries and target onset, and then assigned each saccade to its corresponding trial. Note that this could result in multiple saccades per trial. Only the first qualifying saccade per trial was included in the statistical analyses. The first qualifying saccade after the target was defined by meeting the following criteria: (i) the start position within a central fixation window ($x \in [860, 1060]$ px), (ii) the endpoint within one of two lateral ROIs (Left ROI: $x \in [0, 800]$ px; Right ROI: $x \in [1120, 1920]$ px), and (iii) no blink overlapped the saccade. Direction was defined by the endpoint ROI—landings in the right ROI were labeled right and in the left ROI left. The selection criteria and raw saccadic data are illustrated in **Figure S9**.

Prior to the saccade task, a nine-target calibration and validation process was conducted. Stimuli for the saccade task were programmed using PsychoPy 2021.2.3[93] and displayed on a 24-inch BenQ monitor (resolution, 1920 x 1080; refresh rate, 120 Hz; Qisda Corporation, Taipei, Taiwan).

Data analysis

For all regression analyses reported below, SOA was included as a z-scored covariate, and all factors were sum-to-zero coded. To account for both between and within-subject variability, saccade and masking-assessment data were analyzed with logistic mixed-effects models using the lme4 package in R[94]. We included random effects for all within subject variables. Note that for simplicity, in the equations below, the random effect structure is not included in the syntax. Furthermore, p-values of fixed effects were acquired using Type III conditional F-tests with Kenward-Roger approximation for degrees of freedom, as implemented in the Anova function of the car package[95,96]. Finally, in case of significant fixed effects, post hoc pairwise comparisons were performed using the emmeans function of the emmeans package[97]. For binomial generalized linear mixed models we used simulation-based residual diagnostics (DHARMa) to verify model assumptions (**Supplementary Documents S.4** & **Figure S10**).

Saccade task

Data visualization and analyses were performed using R (version 2021.9.2.382; RStudio Team, 2021). Trials on which participants made double saccades (M = 2.1%, SD = 1.6, range = 0.4% – 7.8%) and where response times exceeded 1 s (M = 2.7%, SD = 2.7, range = 0.1% - 11.0%), which may have indicated failed eyetracking, were excluded. The practice and padding trials were also excluded from the dataset.

Baseline behavior

To evaluate the efficacy of the saccade task by establishing a robust relationship between target delay onsets and the probability of saccades to certain directions. The dependent variable is the probability of making a rightward saccade. The independent variable is target SOA. We hypothesize a higher probability of rightward saccades at larger positive SOA (e.g., target on the right hemifield appeared first) and a lower probability of rightward saccades at more negative target delays (e.g., target on the left hemifield appeared first).

The following lme4 model syntax was used:

$$saccade\ direction \sim 1 + SOA \quad (1)$$

TUS effects

To assess the direction of TUS effects, we looked at the effects of TUS per condition on saccade direction. We hypothesized TUS effects to surface primarily in biasing responses on trials with higher uncertainty (i.e., trials with short SOAs) and therefore focused on the 'choice domain'. Choice domains were defined at the individual level by determining the delay windows, i.e., SOAs, where participants showed a probability of making rightward saccades between 0.25 and 0.75 (see **Supplementary Tables S.1-S.9** for (i) condition (all) SOA, (ii) stimulation side (left/right) by stimulation region (FEF/M1) and SOA, and (iii) MRS effect analyses for this and other choice domain window results). This led to inclusion of on average 455 trials per participant (SD = 127, range = 125 – 733) with an average of 65 trials per TUS condition (SD = 18, range = 15 – 108) and 130 trials per sham condition (SD = 37, range = 40 – 208).

To measure the effects of TUS on saccade behavior, we first present our primary pre-registered comparison of left versus right FEF, for which we expected opposing (lateralized) effects on saccade choices: the lateralized FEFs are each other's most informative control given their contralateral contributions to saccade selection. Next, we tested the left versus right M1 TUS effects. This step allowed us to investigate potential lateralized effects within each stimulated region. In these analyses, we again included target SOA as a continuous predictor. For example, a typical analysis model included predictors for the stimulation condition (e.g., left versus right FEF) and SOA, as follows:

$$saccade\ direction \sim 1 + condition_{leftFEF/rightFEF} + SOA \quad (2)$$

Finally, to further characterize the FEF TUS effects relative to neutral baseline conditions, we compared each FEF condition to sham:

$$saccade\ direction \sim 1 + condition_{leftFEF/rightFEF/sham} + SOA \quad (3)$$

Furthermore, to examine if the effects that we find are specific to FEF modulation, and not a result of any other confounding factors, we looked at the TUS effect of stimulation side (left vs. right) and stimulation region (FEF vs. M1) on saccade direction.

$$saccade\ direction \sim 1 + side_{left/right} * region_{FEF/M1} + SOA \quad (4)$$

Finally, in an exploratory analysis presented in **Supplementary Information S.1**, we examined the effects of TUS during trials where no correct choice could be made based on visual cues alone, specifically when the two targets were presented simultaneously (zero-delay trails). Here, we added the factor of zero-delay into the previous models:

$$saccade\ direction \sim 1 + condition_{leftFEF/rightFEF} * SOA_{zero/notzero} + SOA \quad (5)$$

$$saccade\ direction \sim 1 + condition_{leftM1/rightM1} * SOA_{zero/notzero} + SOA \quad (6)$$

$$saccade\ direction \sim 1 + side_{left/right} * region_{FEF/M1} * SOA_{zero/notzero} + SOA \quad (7)$$

Biasing TUS effects

To ascertain the specificity of TUS effects on saccade biasing rather than general perception, we investigated whether ultrasound modulates the decision curve's characteristics along different axes: horizontal shift (indicating choice bias), slope alteration (indicating impaired target discrimination), or changes in asymptotes/ lapse (indicating a bias beyond the choice domain). These exploratory analyses are discussed in **Supplementary Documents S.2**.

To estimate the slope and bias in milliseconds, we analyzed the interaction effect of condition and delay on saccade direction, focusing specifically on a delay range of -75 to +75 ms to increase sensitivity for detecting any slope effects. This range was selected because it closely approximates the individual choice domain used in other analyses, ensuring consistency and comparability across methods. Unlike previous analyses where individual choice domains were used, we opted for a fixed delay range in this analysis. This decision was made because we aimed to quantify the absolute value of the bias shift (horizontal shift of the curve) in milliseconds. Using the scaled individual choice domains does not provide the opportunity to calculate this fixed bias shift in absolute time units. By including condition and delay as random effects, we were able to estimate the random slopes and biases for each participant. This approach allowed us to determine whether TUS induced horizontal shifts in the decision curve, indicative of a choice bias.

$$saccade\ direction \sim 1 + condition_{sham/leftFEF/rightFEF/leftM1/rightM1} * SOA \quad (8)$$

For lapse estimation, we aimed to understand whether TUS could evoke saccades at longer delays, thus indicating an effect beyond mere biasing. We selected absolute delays between 75 and 200 ms and assessed whether choice accuracy depended on condition and absolute delay. By focusing on absolute delays, rather than distinguishing

between negative and positive delays, we prioritized analyzing overall accuracy rather than side-specific biases. This choice was made because we do not expect side biases to play a role in this context; instead, we are interested in understanding general task performance and accuracy. Therefore, this analysis focused on the asymptotes of the decision curve to determine if TUS influenced saccade behavior even when the delays were long, and the task was easy.

$$saccade\ accuracy \sim 1 + condition_{sham/leftFEF/rightFEF/leftM1/rightM1} + SOA_{abs} \quad (9)$$

Online TUS effects

Moreover, to assess whether this protocol truly functions as a short-lived protocol without producing longer-lasting effects, each sham trial was labeled according to the preceding trial (e.g., left_FEF–sham refers to a sham trial that followed a left FEF trial). However, note that this analysis was conducted with a mean of only 22 trials per condition (SD = 7, range = 6 – 43). We then ran the same side-by-region model to analyze these labeled sham trials. Given the expectation that the protocol only exerts direct, immediate effects, we hypothesized that there would be no significant interaction effect observed. Furthermore, we also performed a Bayesian ANOVA using the same model syntax, as this approach provides a more robust assessment of evidence for the null hypothesis.

$$saccade\ direction \sim 1 + side_{sham(left/right)} * region_{sham(FEF/M1)} + SOA \quad (10)$$

Functional localizers

To accurately target the stimulation sites for each individual, the participants performed a FEF and M1 localizer during the intake session. This data was then pre-processed and analyzed to obtain coordinates for each region per participant.

fMRI pre-processing and analysis were conducted using SPM12 in MATLAB R2023a, along with MRIcroGL for result visualization. The initial steps included excluding the first five fMRI volumes to account for signal steady-state transition, converting IMA files to DICOM compatible format, and visually checking for artefacts. We performed both single subject and group level analyses (N = 35) to establish coordinates within native and standard space, respectively (FEF: **Figure 2A** (single-subject) and **3C** (group-level); M1: **Figure 4A** (single-subject) and **5C** (group-level)).

Realignment and reslicing were performed for both levels, followed by coregistration with the participant's T1w-image for single subject analysis and with Montreal

Neurological Institute (MNI) standard space for group level analysis. Data was smoothing with a six mm FWHM Gaussian kernel, and realignment parameters were inspected. The blocks were convolved with canonical hemodynamic response function, followed by voxel-wise fitting of a general linear model (GLM), resulting in the computation of statistical parametric maps for the comparisons. Subsequently, beta weights for each condition were estimated to create contrast maps, enabling Family-Wise Error corrected cluster-level inferences ($p < .05$). For the FEF localizer, saccades minus fixation blocks were used to obtain coordinates for the left and right FEF. The M1 localizer contrasts involved pinching blocks of right fingers minus left fingers and vice versa to identify the left and right M1, respectively.

To determine FEF targets for TUS delivery, we selected a peak voxel within the significant cluster within FEF, specifically at the junction of the superior precentral sulcus and the superior frontal sulcus. The FEF localizer required reflexive pro-saccades, activating both medial and lateral FEF peaks. The medial peak, linked to higher order cognitive control[75,98,99], was selected for TUS targets. This decision aligns with our hypothesis that TUS holds the highest potential for influencing saccadic behavior at equal preference, requiring the execution of voluntary saccadic eye movements by FEF. The M1 localizer with pinching either the left or right finger elicited a distinctive activation cluster of significant voxels in both left and right M1. Within the activation cluster, the local maximum of peak voxel was selected for the x-, y-, and z-coordinates.

The accuracy of selected coordinates within sulci branches was assessed with FSLeyes by means of visualizing effect sizes modulated by statistical significance with transparent threshold. Once confirmed, established coordinates per stimulation region were entered in the Localite software to plan and monitor TUS delivery. Group level analysis calculated contrast estimates' standard error and mean, determining significance of the average estimate.

MRS analysis

To investigate interindividual differences in TUS susceptibility, we quantified baseline inhibitory tone in the left hemisphere stimulation sites (left FEF and left M1) using MRS.

GABA+ concentrations were quantified using Gannet version 3.1.4[100], with water used as a reference. Gannet's standard preprocessing pipeline was used, which includes frequency and phase correction by spectral registration and line broadening. Edited spectra were generated by subtracting individual edit-ON spectra from edit-OFF spectra. Notably, the editing approach not only targets GABA but also other macromolecules at 3ppm, therefore the concentrations of GABA+ (GABA and macromolecules) are reported. Gray matter, white matter and CSF tissue fractions

for determining tissue-corrected concentrations were obtained for both voxels using SPM12. Metabolite concentrations were then relaxation and tissue-corrected (Gasparovic et al. method).

To ensure data quality, two independent researchers performed visual quality checks of the data. Using the GannetLoad output, water frequency drift was assessed to identify excessive movement artifacts. Next, Cr signal alignment was inspected to evaluate the quality of frequency alignment. For participants with noticeable drift or misalignment, the affected averages were removed, and GABA+ quantification was reprocessed using Gannet. Participants were excluded if more than 50% of their averages had to be removed or if drift and alignment remained insufficient despite reprocessing, as their GABA+ was unreliable or inestimable due to lipid contamination or low signal-to-noise ratio (SNR). Reliable model fits were achieved for 25 out of 35 acquisitions. Data quality was further quantified using the signal-to-noise ratio (SNR) and full-width-at-half-max (FWHM) of N-acetylaspartate (NAA) and fit error of the GABA+ peak provided by Gannet (**Figure S11A-C**).

To assess whether interindividual variability in saccade bias could be explained by baseline inhibitory tone, we examined whether the probability of making a rightward saccade in the left FEF and sham conditions (as well as in the control left M1 and sham conditions) was influenced by baseline GABA+ levels. Given that we only measured the left hemispheric target regions using MRS, we restricted our analyses to the left hemisphere. Specifically, we tested whether the interaction between condition (left FEF vs. sham) and baseline GABA+ levels predicted saccade direction, with target delay included as a separate predictor. We ran the same model for M1 GABA+, to assess if M1 GABA+ levels were predictive of the M1 TUS effects or intrinsic bias.

$$saccade\ direction \sim 1 + condition_{sham/leftFEF} * FEF_{GABA+} + SOA \quad (11)$$

$$saccade\ direction \sim 1 + condition_{sham/leftM1} * M1_{GABA+} + SOA \quad (12)$$

Masking assessment

To evaluate the efficacy of participant blinding to different conditions, we investigated whether participants could distinguish sham from TUS trials by analyzing if stimulation perception (yes/no) depended on the stimulation condition (sham/FEF/M1). Additionally, we assessed whether stimulation and side perception differed between FEF and M1 conditions. Specifically, we analyzed if side perception (left/right) depended on the stimulation side (left/right) and region (FEF/M1) in the following mixed models:

$$stimulation\ perception\ \sim 1 + condition_{sham/FEF/M1} \quad (13)$$

$$side\ perception\ \sim 1 + side_{left/right} * region_{FEF/M1} \quad (14)$$

For both models, we also performed a Bayesian ANOVA to further assess the evidence for the null hypothesis. We hypothesized that there would be no difference in stimulation perception between sham, FEF, and M1 conditions due to the delivery of the auditory mask. Even if differences in stimulation perception were found compared to sham, we expected this not to be problematic due to the inclusion of the active control site (M1), where TUS was also delivered to both the left and right hemispheres. We anticipated no significant differences between FEF and M1 under these conditions. While differing results in side perception might be observed, this would not pose a problem since M1 stimulation is also lateralized.

DATA AVAILABILITY

The anonymized behavioral data generated in this study (trial-level tables), task timing logs, stimulation targets/coordinates, single-subject localizer peaks, and summary spectroscopy outputs with QC (i.e., the dataset required to reproduce the results reported in this paper) have been deposited in the Radboud University Data Repository under DOI https://doi.org/10.34973/drtg-kq58. The raw MRI data are protected and are not available due to data privacy laws (GDPR) and the study's ethics approval (CMO Oost-Nederland, CMO2022-15953). Source data are provided with this paper.

CODE AVAILABILITY

All behavioral and fMRI task code and all analysis scripts are available at https://doi.org/10.34973/drtg-kq58.

REFERENCES

1. Murphy, K. R. *et al.* A Practical Guide to Transcranial Ultrasonic Stimulation from the IFCN-endorsed ITRUSST Consortium. Preprint at https://doi.org/10.48550/arXiv.2407.07646 (2024).
2. Kubanek, J. *et al.* Remote, brain region-specific control of choice behavior with ultrasonic waves. *Sci. Adv.* ***6***, eaaz4193 (2020).
3. Menz, M. D., Oralkan, Ö., Khuri-Yakub, P. T. & Baccus, S. A. Precise Neural Stimulation in the Retina Using Focused Ultrasound. *J. Neurosci.* ***33***, 4550–4560 (2013).
4. Mohammadjavadi, M. *et al.* Elimination of peripheral auditory pathway activation does not affect motor responses from ultrasound neuromodulation. *Brain Stimulat.* ***12***, 901–910 (2019).
5. Murphy, K. R. *et al.* A tool for monitoring cell type-specific focused ultrasound neuromodulation and control of chronic epilepsy. *Proc. Natl. Acad. Sci. U. S. A.* ***119***, e2206828119 (2022).
6. Yoo, S., Mittelstein, D. R., Hurt, R. C., Lacroix, J. & Shapiro, M. G. Focused ultrasound excites cortical neurons via mechanosensitive calcium accumulation and ion channel amplification. *Nat. Commun.* ***13***, 493 (2022).
7. Riis, T. S., Feldman, D. A., Losser, A. J., Okifuji, A. & Kubanek, J. Noninvasive targeted modulation of pain circuits with focused ultrasonic waves. *PAIN* 10.1097/j.pain.0000000000003322 (2022) doi:10.1097/j.pain.0000000000003322.
8. Yaakub, S. N. *et al.* Transcranial focused ultrasound-mediated neurochemical and functional connectivity changes in deep cortical regions in humans. *Nat. Commun.* ***14***, 5318 (2023).
9. Kop, B. R. *et al.* Auditory confounds can drive online effects of transcranial ultrasonic stimulation in humans. *eLife* ***12***, (2024).
10. Butler, C. R. *et al.* Transcranial ultrasound stimulation to human middle temporal complex improves visual motion detection and modulates electrophysiological responses. *Brain Stimulat.* ***15***, 1236–1245 (2022).
11. Paus, T. Location and function of the human frontal eye-field: A selective review. *Neuropsychologia* ***34***, 475–483 (1996).
12. Vernet, M., Quentin, R., Chanes, L., Mitsumasu, A. & Valero-Cabré, A. Frontal eye field, where art thou? Anatomy, function, and non-invasive manipulation of frontal regions involved in eye movements and associated cognitive operations. *Front. Integr. Neurosci.* ***8***, 66 (2014).
13. Gaymard, B., Ploner, C. J., Rivaud-Péchoux, S. & Pierrot-Deseilligny, C. The frontal eye field is involved in spatial short-term memory but not in reflexive saccade inhibition. *Exp. Brain Res.* ***129***, 288–301 (1999).
14. Guitton, D., Buchtel, H. A. & Douglas, R. M. Frontal lobe lesions in man cause difficulties in suppressing reflexive glances and in generating goal-directed saccades. *Exp. Brain Res.* ***58***, 455–472 (1985).

15. Henik, A., Rafal, R. & Rhodes, D. Endogenously generated and visually guided saccades after lesions of the human frontal eye fields. *J. Cogn. Neurosci.* ***6***, 400–411 (1994).
16. Rivaud, S., Müri, R. M., Gaymard, B., Vermersch, A. I. & Pierrot-Deseilligny, C. Eye movement disorders after frontal eye field lesions in humans. *Exp. Brain Res.* ***102***, 110–120 (1994).
17. Grosbras, M.-H. & Paus, T. Transcranial magnetic stimulation of the human frontal eye field: effects on visual perception and attention. *J. Cogn. Neurosci.* ***14***, 1109–1120 (2002).
18. Grosbras, M.-H. & Paus, T. Transcranial magnetic stimulation of the human frontal eye field facilitates visual awareness. *Eur. J. Neurosci.* ***18***, 3121–3126 (2003).
19. Nagel, M. *et al.* Distributed representations of the 'preparatory set' in the frontal oculomotor system: a TMS study. *BMC Neurosci.* ***9***, 89 (2008).
20. Nyffeler, T. *et al.* Repetitive TMS over the human oculomotor cortex: comparison of 1-Hz and theta burst stimulation. *Neurosci. Lett.* ***409***, 57–60 (2006).
21. Ro, T., Henik, A., Machado, L. & Rafal, R. D. Transcranial magnetic stimulation of the prefrontal cortex delays contralateral endogenous saccades. *J. Cogn. Neurosci.* ***9***, 433–440 (1997).
22. Ro, T., Cheifet, S., Ingle, H., Shoup, R. & Rafal, R. Localization of the human frontal eye fields and motor hand area with transcranial magnetic stimulation and magnetic resonance imaging. *Neuropsychologia* ***37***, 225–231 (1999).
23. Ro, T., Farnè, A. & Chang, E. Locating the human frontal eye fields with transcranial magnetic stimulation. *J. Clin. Exp. Neuropsychol.* ***24***, 930–940 (2002).
24. Thickbroom, G. W., Stell, R. & Mastaglia, F. L. Transcranial magnetic stimulation of the human frontal eye field. *J. Neurol. Sci.* ***144***, 114–118 (1996).
25. Nandi, T., Kop, B. R., Butts Pauly, K., Stagg, C. J. & Verhagen, L. The relationship between parameters and effects in transcranial ultrasonic stimulation. *Brain Stimulat.* ***17***, 1216–1228 (2024).
26. Guerra, A. *et al.* Gamma-transcranial alternating current stimulation and theta-burst stimulation: inter-subject variability and the role of *BDNF*. *Clin. Neurophysiol.* ***131***, 2691–2699 (2020).
27. Guerra, A., López-Alonso, V., Cheeran, B. & Suppa, A. Variability in non-invasive brain stimulation studies: Reasons and results. *Neurosci. Lett.* ***719***, 133330 (2020).
28. López-Alonso, V., Cheeran, B., Río-Rodríguez, D. & Fernández-Del-Olmo, M. Inter-individual variability in response to non-invasive brain stimulation paradigms. *Brain Stimulat.* ***7***, 372–380 (2014).
29. Pellegrini, M., Zoghi, M. & Jaberzadeh, S. Cluster analysis and subgrouping to investigate inter-individual variability to non-invasive brain stimulation: a systematic review. *Rev. Neurosci.* ***29***, 675–697 (2018).
30. Pellegrini, M., Zoghi, M. & Jaberzadeh, S. Biological and anatomical factors influencing interindividual variability to noninvasive brain stimulation of the primary motor cortex: a systematic review and meta-analysis. *Rev. Neurosci.* ***29***, 199–222 (2018).
31. Stagg, C. J., Bachtiar, V. & Johansen-Berg, H. The Role of GABA in Human Motor Learning. *Curr. Biol.* ***21***, 480–484 (2011).

32. Sumner, P., Edden, R. A. E., Bompas, A., Evans, C. J. & Singh, K. D. More GABA, less distraction: a neurochemical predictor of motor decision speed. *Nat. Neurosci.* ***13***, 825–827 (2010).
33. Barker, A. T., Jalinous, R. & Freeston, I. L. Non-invasive magnetic stimulation of human motor cortex. *Lancet Lond. Engl.* ***1***, 1106–1107 (1985).
34. Kammer, T. Phosphenes and transient scotomas induced by magnetic stimulation of the occipital lobe: their topographic relationship. *Neuropsychologia* ***37***, 191–198 (1998).
35. Murd, C., Luiga, I., Kreegipuu, K. & Bachmann, T. Scotomas induced by multiple, spatially invariant TMS pulses have stable size and subjective contrast. *Int. J. Psychophysiol.* ***77***, 157–165 (2010).
36. Darmani, G. *et al.* Non-invasive transcranial ultrasound stimulation for neuromodulation. *Clin. Neurophysiol.* ***135***, 51–73 (2022).
37. Mankowska, N. D. *et al.* Critical Flicker Fusion Frequency: A Narrative Review. *Medicina (Mex.)* ***57***, 1096 (2021).
38. Foerster, F. R., Giersch, A. & Cleeremans, A. Spatial but not temporal orienting of attention enhances the temporal acuity of human peripheral vision. *Commun. Psychol.* ***3***, 116 (2025).
39. Braun, V., Blackmore, J., Cleveland, R. O. & Butler, C. R. Transcranial ultrasound stimulation in humans is associated with an auditory confound that can be effectively masked. *Brain Stimulat.* ***13***, 1527–1534 (2020).
40. Guo, H. *et al.* Ultrasound Produces Extensive Brain Activation via a Cochlear Pathway. *Neuron* ***98***, 1020-1030.e4 (2018).
41. Johnstone, A. *et al.* A range of pulses commonly used for human transcranial ultrasound stimulation are clearly audible. *Brain Stimulat.* ***14***, 1353–1355 (2021).
42. Sato, T., Shapiro, M. G. & Tsao, D. Y. Ultrasonic Neuromodulation Causes Widespread Cortical Activation via an Indirect Auditory Mechanism. *Neuron* ***98***, 1031-1041.e5 (2018).
43. Nakajima, K. *et al.* A causal role of anterior prefrontal-putamen circuit for response inhibition revealed by transcranial ultrasound stimulation in humans. *Cell Rep.* ***40***, 111197 (2022).
44. Moore, T. & Fallah, M. Control of eye movements and spatial attention. *Proc. Natl. Acad. Sci.* ***98***, 1273–1276 (2001).
45. Ruff, C. C. *et al.* Concurrent TMS-fMRI and Psychophysics Reveal Frontal Influences on Human Retinotopic Visual Cortex. *Curr. Biol.* ***16***, 1479–1488 (2006).
46. Moore, T. & Armstrong, K. M. Selective gating of visual signals by microstimulation of frontal cortex. *Nature* ***421***, 370–373 (2003).
47. Stelmach, L. B. & Herdman, C. M. Directed attention and perception of temporal order. *J. Exp. Psychol. Hum. Percept. Perform.* ***17***, 539–550 (1991).
48. Duecker, F. & Sack, A. T. Rethinking the role of sham TMS. *Front. Psychol.* ***6***, 210 (2015).
49. Jerusalem, A. *et al.* Electrophysiological-mechanical coupling in the neuronal membrane and its role in ultrasound neuromodulation and general anaesthesia. *Acta Biomater.* ***97***, 116–140 (2019).

50. Yu, K., Niu, X., Krook-Magnuson, E. & He, B. Intrinsic functional neuron-type selectivity of transcranial focused ultrasound neuromodulation. *Nat. Commun.* ***12***, 2519 (2021).
51. Heinen, K., Feredoes, E., Weiskopf, N., Ruff, C. C. & Driver, J. Direct Evidence for Attention-Dependent Influences of the Frontal Eye-Fields on Feature-Responsive Visual Cortex. *Cereb. Cortex N. Y. NY* ***24***, 2815–2821 (2014).
52. Huang, Y.-Z. *et al.* Plasticity induced by non-invasive transcranial brain stimulation: A position paper. *Clin. Neurophysiol.* ***128***, 2318–2329 (2017).
53. Kamke, M. R. *et al.* Visual Spatial Attention Has Opposite Effects on Bidirectional Plasticity in the Human Motor Cortex. *J. Neurosci.* ***34***, 1475–1480 (2014).
54. Massimini, M. *et al.* Breakdown of Cortical Effective Connectivity During Sleep. *Science* ***309***, 2228–2232 (2005).
55. Siebner, H. R. *et al.* Transcranial magnetic stimulation of the brain: What is stimulated? - A consensus and critical position paper. *Clin. Neurophysiol. Off. J. Int. Fed. Clin. Neurophysiol.* ***140***, 59–97 (2022).
56. Kolasinski, J. *et al.* The dynamics of cortical GABA in human motor learning. *J. Physiol.* ***597***, 271–282 (2019).
57. McSorley, E., Haggard, P. & Walker, R. Time Course of Oculomotor Inhibition Revealed by Saccade Trajectory Modulation. *J. Neurophysiol.* ***96***, 1420–1424 (2006).
58. Hutchison, R. M. *et al.* Functional connectivity of the frontal eye fields in humans and macaque monkeys investigated with resting-state fMRI. *J. Neurophysiol.* ***107***, 2463–2474 (2012).
59. Kagan, I., Iyer, A., Lindner, A. & Andersen, R. A. Space representation for eye movements is more contralateral in monkeys than in humans. *Proc. Natl. Acad. Sci.* ***107***, 7933–7938 (2010).
60. Silver, M. A. & Kastner, S. Topographic maps in human frontal and parietal cortex. *Trends Cogn. Sci.* ***13***, 488–495 (2009).
61. Bartolomeo, P. & Seidel Malkinson, T. Hemispheric lateralization of attention processes in the human brain. *Curr. Opin. Psychol.* ***29***, 90–96 (2019).
62. Cazzoli, D. *et al.* The role of the right frontal eye field in overt visual attention deployment as assessed by free visual exploration. *Neuropsychologia* ***74***, 37–41 (2015).
63. Hanning, N. M., Fernández, A. & Carrasco, M. Dissociable roles of human frontal eye fields and early visual cortex in presaccadic attention. *Nat. Commun.* ***14***, 5381 (2023).
64. Faul, F., Erdfelder, E., Buchner, A. & Lang, A.-G. Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behav. Res. Methods* ***41***, 1149–1160 (2009).
65. Vossel, S., Geng, J. J. & Fink, G. R. Dorsal and Ventral Attention Systems. *The Neuroscientist* ***20***, 150–159 (2014).
66. Shulman, G. L. *et al.* Right Hemisphere Dominance during Spatial Selective Attention and Target Detection Occurs Outside the Dorsal Frontoparietal Network. *J. Neurosci.* ***30***, 3640–3651 (2010).

67. Anderson, E. J. *et al.* Cortical Network for Gaze Control in Humans Revealed Using Multimodal MRI. *Cereb. Cortex N. Y. NY* ***22***, 765–775 (2012).
68. Johnstone, L. T., Karlsson, E. M. & Carey, D. P. Left-Handers Are Less Lateralized Than Right-Handers for Both Left and Right Hemispheric Functions. *Cereb. Cortex* ***31***, 3780–3787 (2021).
69. Tomasi, D. & Volkow, N. D. Associations between handedness and brain functional connectivity patterns in children. *Nat. Commun.* ***15***, 2355 (2024).
70. Pool, E.-M., Rehme, A. K., Fink, G. R., Eickhoff, S. B. & Grefkes, C. Handedness and effective connectivity of the motor system. *NeuroImage* ***99***, 451–460 (2014).
71. Serrien, D. J., Ivry, R. B. & Swinnen, S. P. Dynamics of hemispheric specialization and integration in the context of motor control. *Nat. Rev. Neurosci.* ***7***, 160–166 (2006).
72. Heilman, K. M. & Abell, T. V. D. Right hemisphere dominance for attention. *Neurology* ***30***, 327–327 (1980).
73. Sack, A. T. *et al.* Optimizing Functional Accuracy of TMS in Cognitive Studies: A Comparison of Methods. *J. Cogn. Neurosci.* ***21***, 207–221 (2009).
74. Amiez, C., Kostopoulos, P., Champod, A.-S. & Petrides, M. Local Morphology Predicts Functional Organization of the Dorsal Premotor Region in the Human Brain. *J. Neurosci.* ***26***, 2724–2731 (2006).
75. Gagnon, D., O'Driscoll, G. A., Petrides, M. & Pike, G. B. The effect of spatial and temporal information on saccades and neural activity in oculomotor structures. *Brain* ***125***, 123–139 (2002).
76. Tzourio-Mazoyer, N. *et al.* Between-hand difference in ipsilateral deactivation is associated with hand lateralization: fMRI mapping of 284 volunteers balanced for handedness. *Front. Hum. Neurosci.* ***9***, (2015).
77. Gruetter, R. Automatic, localized in vivo adjustment of all first- and second-order shim coils. *Magn. Reson. Med.* ***29***, 804–811 (1993).
78. Gruetter, R. & Tkác, I. Field mapping without reference scan using asymmetric echo-planar techniques. *Magn. Reson. Med.* ***43***, 319–323 (2000).
79. Yousry, T. A. *et al.* Localization of the motor hand area to a knob on the precentral gyrus. A new landmark. *Brain J. Neurol.* ***120 (Pt 1)***, 141–157 (1997).
80. Tkác, I., Starcuk, Z., Choi, I. Y. & Gruetter, R. In vivo 1H NMR spectroscopy of rat brain at 1 ms echo time. *Magn. Reson. Med.* ***41***, 649–656 (1999).
81. Mescher, M., Tannus, A., Johnson, M. O. & Garwood, M. Solvent Suppression Using Selective Echo Dephasing. *J. Magn. Reson. A* ***123***, 226–229 (1996).
82. Mescher, M., Merkle, H., Kirsch, J., Garwood, M. & Gruetter, R. Simultaneous in vivo spectral editing and water suppression. *NMR Biomed.* ***11***, 266–272 (1998).
83. Marjańska, M. *et al.* Brain dynamic neurochemical changes in dystonic patients: A magnetic resonance spectroscopy study. *Mov. Disord.* ***28***, 201–209 (2013).
84. Floyer-Lea, A., Wylezinska, M., Kincses, T. & Matthews, P. M. Rapid modulation of GABA concentration in human sensorimotor cortex during motor learning. *J. Neurophysiol.* ***95***, 1639–1644 (2006).

85. Kolasinski, J. *et al.* The dynamics of cortical GABA in human motor learning. *J. Physiol.* ***597***, 271–282 (2019).
86. Maes, C. *et al.* Task-Related Modulation of Sensorimotor GABA+ Levels in Association with Brain Activity and Motor Performance: A Multimodal MRS–fMRI Study in Young and Older Adults. *J. Neurosci.* ***42***, 1119–1130 (2022).
87. Li, H. *et al.* MRS-assessed brain GABA modulation in response to task performance and learning. *Behav. Brain Funct. BBF* ***20***, 22 (2024).
88. Porges, E. C., Jensen, G., Foster, B., Edden, R. A. & Puts, N. A. The trajectory of cortical GABA across the lifespan, an individual participant data meta-analysis of edited MRS studies. *eLife* ***10***, e62575 (2021).
89. Chen, C. *et al.* Activation induced changes in GABA: Functional MRS at 7 T with MEGA-sLASER. *NeuroImage* ***156***, 207–213 (2017).
90. Schaller, B., Xin, L., O'Brien, K., Magill, A. W. & Gruetter, R. Are glutamate and lactate increases ubiquitous to physiological activation? A 1H functional MR spectroscopy study during motor activation in human brain at 7 Tesla. *NeuroImage* ***93***, 138–145 (2014).
91. Choi, M. H., Li, N., Popelka, G. & Butts Pauly, K. Development and validation of a computational method to predict unintended auditory brainstem response during transcranial ultrasound neuromodulation in mice. *Brain Stimulat.* ***16***, 1362–1370 (2023).
92. Aubry, J.-F. *et al.* ITRUSST Consensus on Biophysical Safety for Transcranial Ultrasonic Stimulation. Preprint at https://doi.org/10.48550/arXiv.2311.05359 (2024).
93. Peirce, J. *et al.* PsychoPy2: Experiments in behavior made easy. *Behav. Res. Methods* ***51***, 195–203 (2019).
94. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting Linear Mixed-Effects Models Using lme4. *J. Stat. Softw.* ***67***, 1–48 (2015).
95. Fox, J., Weisberg, S. & Price, B. car: Companion to Applied Regression. 3.1-3 https://doi.org/10.32614/CRAN.package.car (2001).
96. Fox, J. *et al.* car: Companion to Applied Regression. (2024).
97. Lenth, R. V. *et al.* emmeans: Estimated Marginal Means, aka Least-Squares Means. (2024).
98. Cameron, I. G. M., Riddle, J. M. & D'Esposito, M. Dissociable Roles of Dorsolateral Prefrontal Cortex and Frontal Eye Fields During Saccadic Eye Movements. *Front. Hum. Neurosci.* ***9***, (2015).
99. Neggers, S. F. W. *et al.* A Functional and Structural Investigation of the Human Fronto-Basal Volitional Saccade Network. *PLoS ONE* ***7***, e29517 (2012).
100. Edden, R. A. E., Puts, N. A. J., Harris, A. D., Barker, P. B. & Evans, C. J. Gannet: A batch-processing tool for the quantitative analysis of gamma-aminobutyric acid–edited MR spectroscopy spectra. *J. Magn. Reson. Imaging JMRI* ***40***, 1445–1452 (2014).
101. Neggers, S. F. W., Zandbelt, B. B., Schall, M. S. & Schall, J. D. Comparative diffusion tractography of corticostriatal motor pathways reveals differences between humans and macaques. *J. Neurophysiol.* ***113***, 2164–2172 (2015).

ACKNOWLEDGEMENTS

This experiment was supported by the Dutch Research Council (NWO), awarding VIDI fellowships to L.V. (18919) and H.E.M.d.O. (452-17-016). We would like to acknowledge Edward J. Auerbach, Ph.D., and Małgorzata Marjańska, Ph.D. (Center for Magnetic Resonance Research and Department of Radiology, University of Minnesota, USA) for the development of the pulse sequences for the Siemens platform, which were provided by the University of Minnesota under a C2P agreement.

Additionally, we thank Norbert Hermesdorf, Margely Cornelissen, Hubert Voogd, Sibrecht Bouwstra, Gerard van Oijen, and Pascal de Water from the technical support group at the Donders Centre for Cognition, Faculty of Social Sciences, Radboud University, for their excellent technical assistance and support throughout this study.

Finally, we would also like to thank Marwan Engels (Donders Centre for Cognition, Radboud University) for his substantial role during data acquisition.

AUTHOR CONTRIBUTIONS STATEMENT

S.F., L.V. and H.E.M.d.O. conceptualized and designed the experiment; A.C., S.L.Y.W. and S.F. designed and programmed the behavioral and functional localizer tasks; S.F. and S.L.Y.W. collected the data; J.P.M. set up the fMRI and MRS sequences; S.F., L.V., and H.E.M.d.O. analyzed the behavioral, functional and spectroscopy data; B.R.K. contributed to behavioral data analysis; R.S.K. contributed to spectroscopy analysis; S.L.Y.W. contributed to fMRI analysis; S.F., L.V. and H.E.M.d.O. wrote the manuscript; B.R.K., R.S.K., W.P.M., J.P.M. and A.C. revised the manuscript.

COMPETING INTERESTS STATEMENT

L.V. declares no competing interests relevant to this study. L.V. is a board member of ITRUSST and the Brainbox Initiative. L.V. has received non-financial support from Image Guided Therapy SA (France), Sonic Concepts LLC (US), and Brainbox Ltd (UK), and consulting fees from Nudge LLC (US).

All other authors declare no competing interests.

SUPPLEMENTARY INFORMATION

SUPPLEMENTARY SECTIONS

S.1 Zero-delay trials

Our analyses are tailored to the choice domain, where the SOA provides informative—but not overwhelming—evidence in favor of one saccade direction over the other. When there is no meaningful choice to be made, i.e., when the SOA is long, performance approaches ceiling, and we did not observe any modulatory effect of TUS (**Supplementary Information S.2**).

In contrast, exploratory analyses revealed that when no correct choice can be made (i.e., when the two cues are presented simultaneously), TUS over both FEF and M1 biased saccades. A three-way interaction trend between region (FEF/M1), stimulation side (left/right), and delay (non-zero/zero) suggested that any M1 effect was limited to the zero-delay trials (b = -0.50, 95%-CI [-1.02, 0.02], $\chi^2(1) = 3.6$, $p = 0.058$; **Figure S2**). However, this analysis is underpowered due to the inclusion of multiple interactions. To address this, we split the data into FEF and M1 subsets. These subsequent tests confirmed that M1 TUS effects are specific to zero-delay trials (condition (left M1/right M1) x delay (non-zero/zero): b = -0.57, 95%-CI [-0.94, -0.20], $\chi^2(1) = 9.0$, $p = 0.002$; **Figure S2**). In contrast, FEF TUS effects persisted across both zero- and non-zero delay trials (condition (left FEF/right FEF) x delay (non-zero/zero): b = -0.08, 95%-CI [-0.45, 0.28], $\chi^2(1) = 0.2$, $p = 0.6$; **Figure S2**), indicating that the observed effects of FEF TUS are robust and stable.

The M1 TUS effect on zero-delay trials was significant, both in statistical and absolute terms. When these trials are included in the choice domain, the shared direction of TUS bias across FEF and M1 obscures a putative interaction of stimulation side and region (side (left/right) x region (FEF/M1): b = 0.16, 95%-CI [-0.03, 0.36], $\chi^2(1) = 2.9$, p = 0.09). Instead, it reveals a main effect of stimulation side (side (left/right): b = -0.26, 95%-CI [-0.41, -0.11], $\chi^2(1) = 11.5$, $p = 0.007$), and region (region (FEF/M1): b = -0.14, 95%-CI [-0.26, -0.01], $\chi^2(1) = 4.6$, $p = 0.032$).

The highly specific biasing effect of TUS over M1, observed only when the two visual cues are presented simultaneously, warrants further investigation. In these conditions, no correct choice can be made based on the visual cues alone. One possible explanation is that an M1 TUS bias arises from true neuromodulation of M1. Indeed, M1

circuits anatomically converge with downstream saccade circuits in the basal ganglia to support eye-hand coordination[99]. Perhaps when visual information is absent, motor biases are propagated through these shared effector circuits.

S.2 Estimating the bias, slope and lapse rate

To ascertain the specificity of the TUS effects on biasing saccade direction, we explored whether FEF TUS influences choice bias (a horizontal shift in the decision curve), target discrimination (a change in slope), or bias beyond the choice domain (a change in asymptotes/lapse).

Bias and slope were analyzed using mixed effects logistic regression focusing on trials with short delays (-75 to +75 ms). Choice bias was assessed by examining the main effect of stimulation condition ($\chi^2(1) = 22.2$, $p < 0.001$), revealing that TUS induced a horizontal shift in the decision curve. Specifically, left FEF stimulation shifted the curve by -3.67 ms, whereas right FEF stimulation shifted it by +3.77 ms. Post hoc tests showed that the shift between left and right FEF stimulation was significantly different ($p < 0.001$; **Figure S3A**), unlike the shift between left and right M1 stimulation ($p = 0.09$; **Figure S3A**).

In contrast, the slope of the decision curves, reflecting target discrimination was unaffected by TUS, as indicated by the non-significant interaction between stimulation condition and target delay (condition x SOA: $\chi^2(1) = 2.715$, $p = 0.6$; condition (left FEF) x SOA: $b = 0.001$, 95%-CI [-0.002, 0.003]; condition (right FEF) x SOA: $b = 0.001$, 95%-CI [-0.002, 0.004]; condition (left M1) x SOA: $b = 0.001$, 95%-CI [-0.002, 0.003]; condition (right M1) x SOA: $b = 0.002$, 95%-CI [-0.001, 0.005]; **Figure S3B**).

Finally, to determine if TUS alters bias outside the choice domain, we analyzed trials with longer absolute delays (75 to 200 ms). While there was a trend suggesting a condition effect (condition: $\chi^2(1) = 9.0$, $p = 0.060$; condition (left FEF): $b = 0.23$, 95%-CI [-0.01, 0.0.56]; condition (right FEF): $b = 0.34$, 95%-CI [0.03, 0.64]; condition (left M1): $b = 0.33$, 95%-CI [0.02, 0.64]; condition (right M1): $b = 0.35$, 95%-CI [0.04, 0.65]; **Figure S3C**), post hoc tests confirmed no significant impact of TUS on choice bias these trials (FEF (left/right): $p = 1.0$; M1(left/right): $p = 1.0$; **Figure S3C**). This reinforces the conclusion that TUS selectively biases responses under conditions of response uncertainty and is unlikely to reverse or evoke responses under high certainty.

S.3 Ipsilateral after-effects and stimulation perception

Within each block, participants received pseudorandomized left TUS, right TUS, and sham stimulation. To investigate potential longer-lasting TUS effects beyond the stimulation duration itself, we analyzed sham trials that directly followed a TUS trial. First, as a baseline we tested the sham→sham condition and found no significant bias following sham trials a non-significant intercept on the logit scale (b = −0.07, p = 0.3), corresponding to p = 0.5 (95%-CI [0.455–0.512]) from 0.5.

Interestingly, for sham trials following TUS, we observed a significant increase in ipsilateral responses during these sham trials—for example, if a sham trial followed a left TUS trial, participants were more likely to make a leftward saccade ($side_{t-1}$: b = -0.10, 95%-CI [-0.18, -0.02], $\chi^2(1)$ = 5.3, p = 0.021; **Figure 6A**). This ipsilateral bias on sham trials following TUS was in the opposite direction of the behavioral effects induced by FEF TUS, which increased contralateral saccades. One could conceive that this reversal reflects compensatory carry-over effects of TUS. Importantly, however, these after-effects and stimulation perception biases did not differ between FEF and M1 conditions ($region_{t-1}$: b = 0.8, 95%-CI [-0.08, 0.11], $\chi^2(1)$ = 0.1, p = 0.8; $side_{t-1}$ x $region_{t-1}$: b = 0.02, 95%-CI [-0.06, 0.10], $\chi^2(1)$ = 0.3, p = 0.6), whereas TUS effects were present only for FEF stimulation. Thus, the robust FEF TUS effects cannot be explained by these ipsilateral tendencies, nor do they support the presence of post-TUS compensatory mechanisms.

Thus, while neuromodulatory effects of TUS cannot explain the presence of ipsilateral after-effects, this begs the question what does drive these effects. We speculate that the answer to this question in the similar ipsilateral response pattern that emerged in the masking assessment. Here, participants performed a forced-choice task to report whether they perceived the stimulation as originating from left or right TUS. Participants significantly misattributed stimulation to the ipsilateral side, independent of stimulation region (side: b = -1.2, 95%-CI [-2.1, -0.3], $\chi^2(1)$ = 7.0, p = 0.008; side (left/right) x region (FEF/M1): b = -0.5, 95%-CI [-1.5, 0.5], $\chi^2(1)$ = 1.1, p = 0.3; **Figure 6C**).

This ipsilateral and lateralized perception of TUS likely stems from specific properties of skull morphology. Variations in how flexural waves—vibrations traveling through the skull—are transmitted can cause the highest amplitude near the contralateral cochlea, influencing perceived sound location[35]. We speculate that the regionally non-specific but lateralized after-effects observed in sham trials immediately post TUS may reflect an attention-orienting response. If participants subjectively perceived prior stimulation as originating from the left, they may have been biased toward making leftward saccades afterward.

S.4 Model diagnostics

For binomial generalized linear mixed models we used simulation-based residual diagnostics (DHARMa) to verify model assumptions. We assessed overall model fits by inspecting the uniformity of simulated residuals (DHARMa QQ/uniformity plot with a Kolmogorov–Smirnov test), and tested for (over/under)dispersion and outliers using DHARMa's non-parametric procedures. To evaluate homoscedasticity and potential model mis-specification, we plotted residuals versus fitted values (rank-transformed). Linearity of continuous predictors (e.g., SOA_scaled) was examined both visually via residuals versus predictor plots and formally by comparing the linear specification with a natural-spline alternative (likelihood-ratio test). We also inspected residual distributions across factor levels (e.g., stimulation side) using within-group uniformity checks and a Levene test for equality of variances. Finally, we evaluated multicollinearity of fixed effects via variance inflation factors (VIFs). All diagnostics were implemented in R with the packages DHARMa, performance/see, and influence. ME. Together, these diagnostics indicated adequate fit for both the left vs. right FEF model, and the Side × Region model. Residuals are approximately uniform with no over/under-dispersion or outlier excess, no systematic trend versus fitted values or SOA, and no heteroscedasticity. See **Figure S10** for detailed results.

SUPPLEMENTARY FIGURES

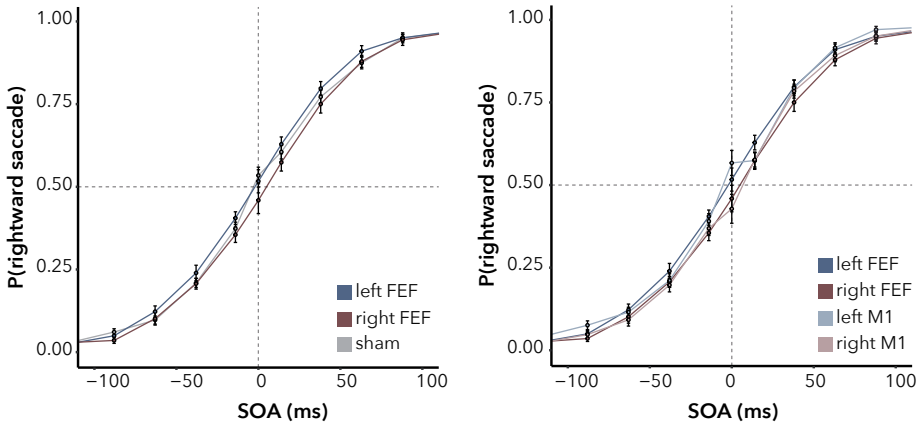


Figure S1 | Empirical psychometric curves of FEF and control conditions. Left panel shows FEF vs sham TUS effects, right panel shows FEF vs M1 TUS effects. Data are binned for visualization purposes into bins of approximately equal numbers of trials, resulting in SOA intervals of [0], [8.3 – 25.0], [33.3 – 50.0], [58.3 – 75.0], [83.3 – 100.0], [108.3 – 141.7], and [150 – 200]. Bins are symmetric for negative values. Note that all inferential statistics were performed using continuous SOAs. Dots represent the group mean per bin, and error bars indicate the S.E.M. across participants (n = 35). Source data are provided as a Source Data file.

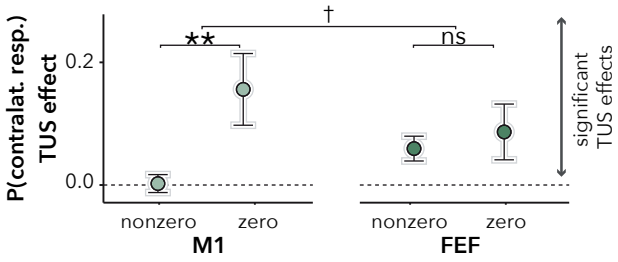


Figure S2 | Differences in TUS effects for zero-delay and nonzero-delay trials in FEF and M1. TUS effects are expressed as the probability of making contralateral saccades (e.g., left hemisphere is stimulated and a rightward saccade is made; n = 35). A contralateral saccade probability greater than 0.5 indicates a TUS effect beyond chance. Left and right FEF and M1 conditions were pooled to form FEF and M1 categories, respectively. For M1, a significant difference in TUS effects was observed between zero-delay and non-zero-delay trials, with TUS effects being present only for zero-delay trials, indicating a delay-dependent modulation of M1. In contrast, FEF stimulation produced significant TUS effects for both zero-delay and non-zero-delay trials, with no significant difference between the delays, suggesting a stable and consistent modulation of FEF regardless of delay. Data are presented as group means with standard error of the mean (S.E.M.). ** $p<0.01$, † $p<0.1$, ns $p>0.1$. Source data are provided as a Source Data file. Exact p values are presented in the supplementary text. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

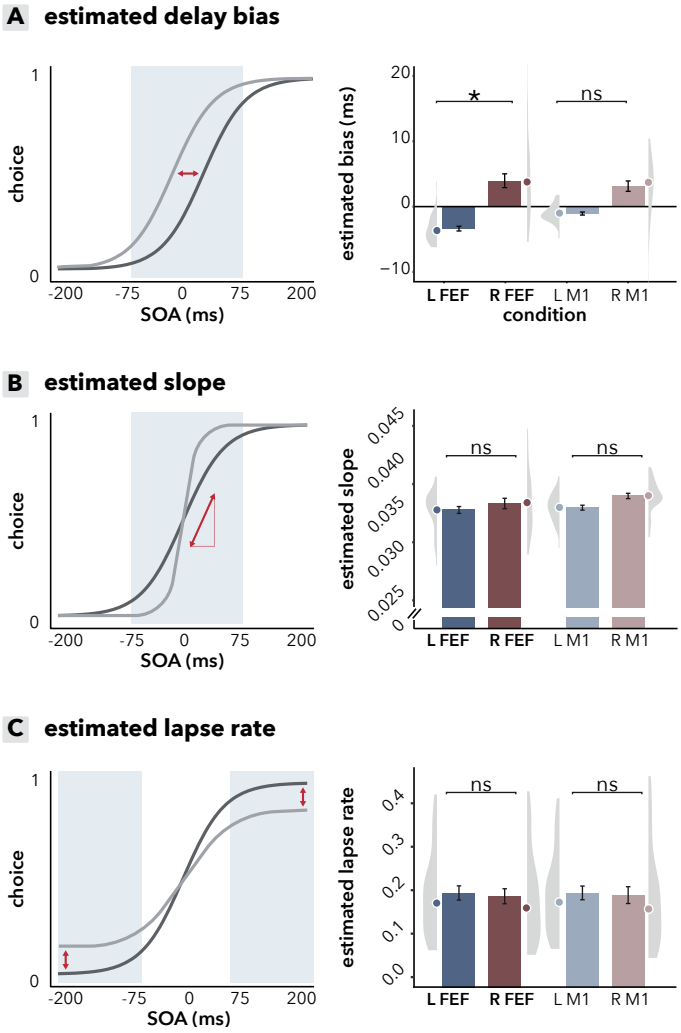


Figure S3 | Estimating bias, slope and lapse rate following FEF and M1 TUS. (A) Choice bias (*n* = 35). Left: a visual representation of a horizontal shift in the decision curve, with the shaded blue area indicating the region of interest for the analysis. Right: estimated bias per participant derived from the mixed-effects model, which includes random effects (individual distributions with cloud and SEM) and fixed effects (group mean represented by the dot). Left FEF stimulation shifted the decision curve by -3.67 ms, while right FEF stimulation shifted it by +3.77 ms. A significant difference was found between left and right FEF ($p < 0.001$), but no significant shift was observed for M1 ($p = 0.09$), suggesting that TUS specifically affects choice bias in FEF. **(B)** Target discrimination (slope; *n* = 35). Left: a visual representation of how changes in slope would look, with the blue shaded area marking the region of interest. Right: estimated slopes per participant derived from the mixed-effects model, which includes random effects (individual distributions with cloud and SEM) and fixed effects (group mean represented by the dot). No significant effect of TUS on slope was observed, as no difference was found between these conditions ($p = 0.6$), indicating that TUS does not significantly impact target discrimination. **(C)** Bias beyond the choice domain (lapse rate; *n* = 35). Left: a visual representation of changes in lapse rate, with the shaded blue area marking the region of interest. Right: estimated lapse

rates per participant derived from the mixed-effects model, which includes random effects (individual distributions with cloud and SEM) and fixed effects (group mean represented by the dot). While a trend toward a condition effect was observed (p = 0.060), post hoc tests confirmed no significant impact of TUS on lapse rate (p = 1.0 for both FEF and M1), reinforcing the conclusion that TUS selectively biases choice behavior without affecting performance in the lapse rate domain. Source data are provided as a Source Data file. Exact p values are presented in the supplementary text. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

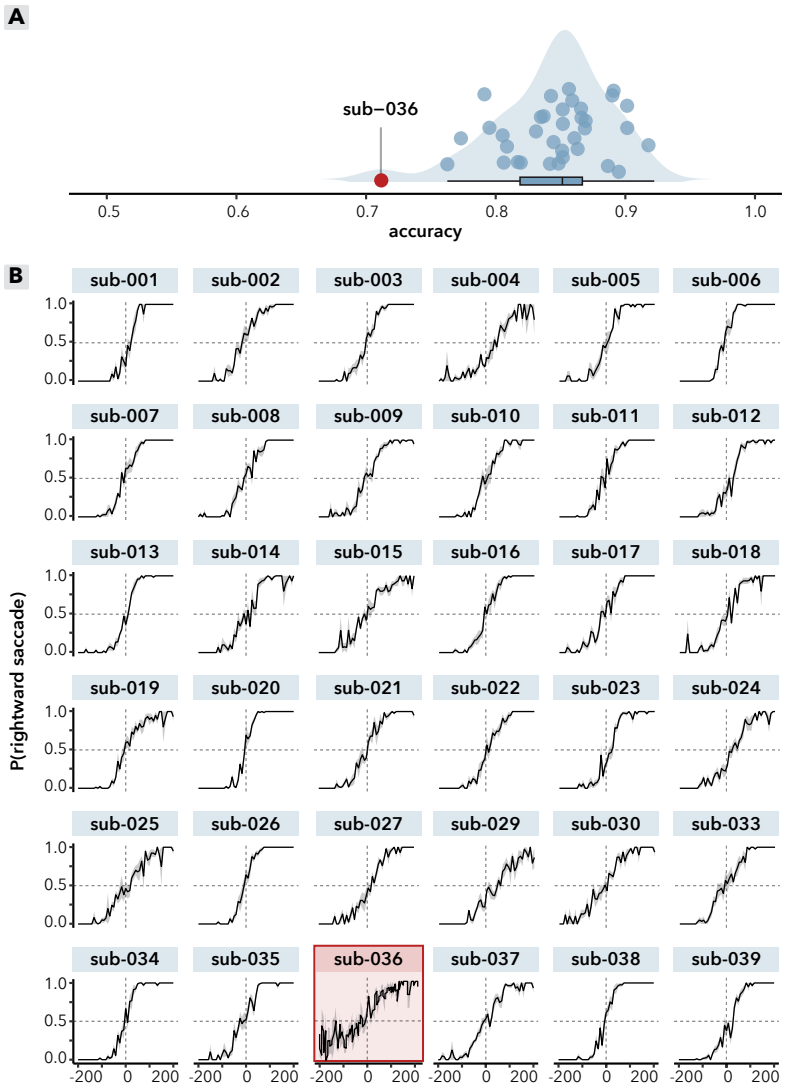


Figure S4 | Participant accuracy and psychometric quality checks. (A) Distribution of overall accuracy across participants. The shaded curve is a kernel-density estimate; the horizontal boxplot shows the median (line), quartiles (box), and whiskers. Participants were flagged as outliers if they met any robust criterion: (i) Tukey IQR rule (accuracy < Q1 − 1.5·IQR or > Q3 + 1.5·IQR), (ii) |z| > 3 based on the

sample mean/SD, and/or (iii) robust MAD-z > 3.5. The red marker highlights sub-036 that exceeded all three criteria and is flagged as an outlier. **(B)** Empirical psychometric curves plotting the probability of a rightward saccade against SOAs (ms; rightward-positive), collapsed across all conditions. Light grey ribbons indicate the S.E.M. across trials per SOA bin. Most participants show the expected monotonic increase from leftward to rightward delays (horizontal dashed line at 0.5 = chance, and vertical dashed line at 0 SOA). In contrast, sub-036 (red panel) exhibits non-monotone, noisy behavior with extended near-chance performance, consistent with non-compliance/unstable responding.

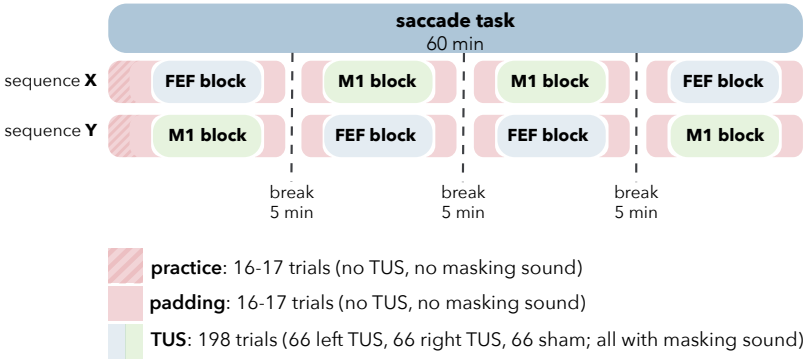


Figure S5 | Schematic overview of the saccade task duration and block sequences. The task consists of two randomized and counterbalanced block sequences: X (FEF, M1, M1, and then FEF) and Y (M1, FEF, FEF, and M1), with the order of the blocks varying across participants. Each block comprises 198 trials (66 left TUS, 66 right TUS, and 66 sham trials, all with masking sound), with additional padding of 16-17 trials before and after the block, where no TUS or masking sound is presented. At the start of each session, participants complete a short practice block for task familiarization. A 5-minute break is provided between blocks for participants to stretch and rest before transducers are recoupled. FEF stimulation is represented in blue and M1 stimulation in green.

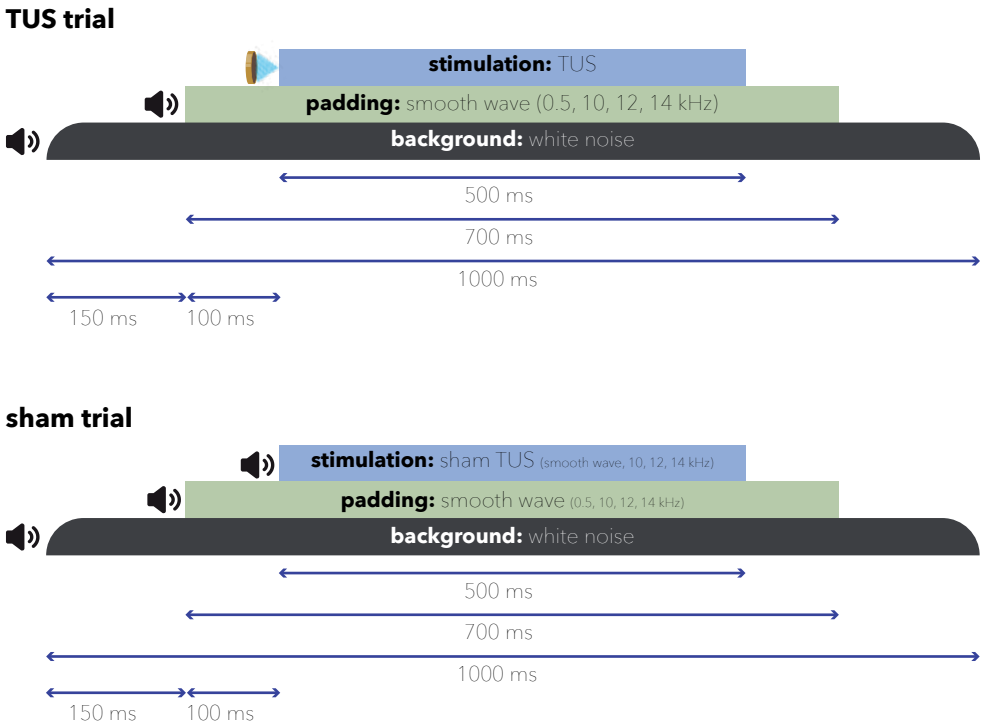


Figure S6 | Schematic overview of masking structure in TUS and sham trials. The trial begins with background white noise, delivered via bone-conducting headphones, lasting 1000 ms. 150 ms after the onset of the white noise, a smooth wave (padding sound) consisting of frequencies 0.5, 10, 12, and 14 kHz is delivered via bone-conducting headphones for 700 ms. For TUS trials (top), the 500 ms ultrasonic stimulation is then delivered 250 ms after the white noise onset (100 ms after the padding sound). For sham trials (bottom), the 500 ms of ultrasonic stimulation is replaced by a smooth wave of 10, 12, and 14 kHz, mimicking the TUS sound.

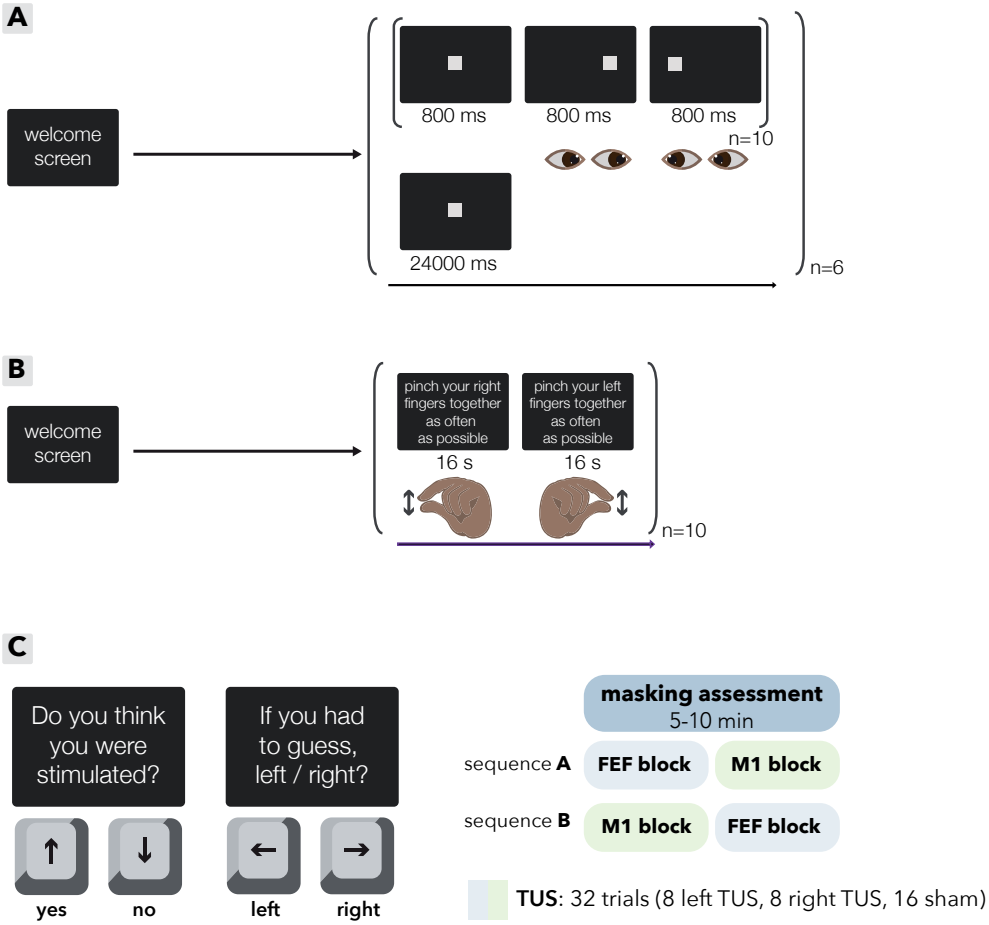


Figure S7 | Schematic overview of functional localizers and masking assessment. (A) FEF functional localizer. After a welcome screen, participants completed blocks of "follow the target" and "fixate on the target". In the 'follow the target' blocks, participants were presented with targets in a random order of left, center, and right positions. This sequence was randomized and repeated 10 times per block, with each target being shown for 800 ms. In the "fixation" block, the target was displayed in the center for 24 seconds. This sequence was repeated six times. The contrast was fixation versus follow the target. **(B)** M1 functional localizer. After a welcome screen, two blocks were presented, repeated 10 times. In each block, participants were instructed to either pinch their right index finger and thumb as often as possible, or their left index finger and thumb as often as possible for 16 seconds. The contrast was left versus right. **(C)** Masking assessment. Participants received TUS on the left/right FEF, left/right M1, and sham stimulation. For each trial, participants were first asked to indicate whether they thought they had received stimulation by pressing the up-arrow key for yes and the down-arrow key for no. Afterward, they were asked to guess whether the stimulation was on the left or right side by pressing the left-arrow key for left and the right-arrow key for right. The order of the conditions/trials was randomized within each block, with the overall sequence determined by the block sequence (right panel). Each block comprises of 8 left TUS, 8 right TUS and 16 sham trials.

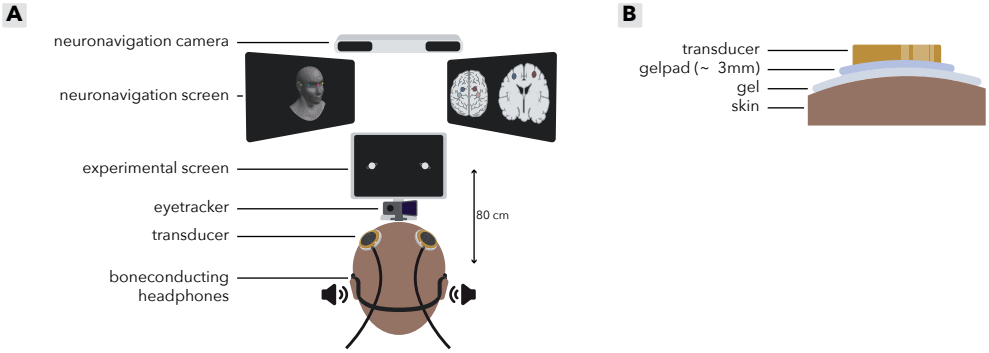


Figure S8 | Overview of the study setup. (A) Schematic overview of the TUS setup. Participants were seated 80 cm from the experimental task screen, with their head stabilized on a chinrest at the center of the screen. Transducers were positioned using a Velcro headcap and guided by neuronavigation for precise targeting (coordinates derived from FEF and M1 functional localizers and entered into Localite software). Bone-conducting headphones for masking were placed on participants, and an eye tracker was positioned below the screen to record eye movements. After neuronavigation, bone-conducting headphones were secured, and transducers were coupled to the participant's head while they remained still. Eye-tracking calibration was performed at the start of each block, prior to initiating the saccade task. **(B)** Schematic overview of coupling materials and layers. The participant's hair and scalp were carefully prepared with ultrasound gel to ensure full coverage of hair follicles and minimize air pockets. A thin (~3 millimeters) gel pad was placed on top of the ultrasound gel layer, allowing for visualization and removal of any remaining air bubbles. Another layer of ultrasound gel was applied to the transducer surface, ensuring no air bubbles were present. The transducer was then carefully positioned at the stimulation site, guided by neuronavigation for precise targeting.

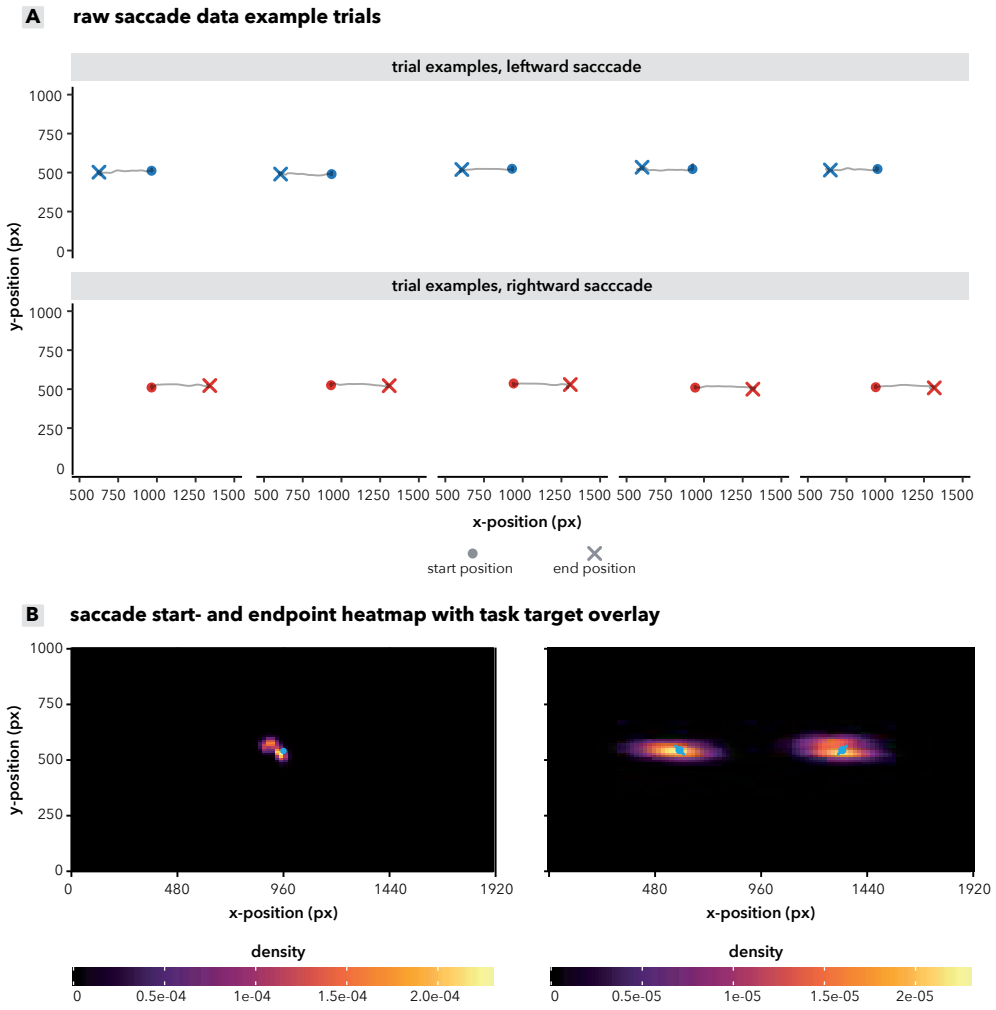


Figure S9 | Raw saccade examples and spatial distributions (example participant; n = 1). (A) *Heatmaps of saccade* start positions *(left) and* endpoints *(right) across all qualifying trials. Bright colors indicate higher density estimates. Blue overlays show the task visual targets:* fixation *at screen center (960 px) and* left/right targets *at ±10° from center (crosses).* **(B)** *Velocity–time traces (°/s) for the same example trials, aligned to* target onset *(t = 0 ms). Shaded rectangles (matching the line colors) indicate the detected saccade duration for each trial; peaks correspond to primary and, when present, corrective saccades. The horizontal dashed line denotes the velocity threshold used for saccade detection.* **(C)** *Overlaid raw examples of the first* qualifying *saccade after target onset from several trials. Filled circles mark saccade onset; × marks the endpoint. Thin lines show the raw gaze trace within each trial. Colors encode saccade direction (blue = leftward; red = rightward). Axes are gaze position in pixels. Source data are provided as a Source Data file.*

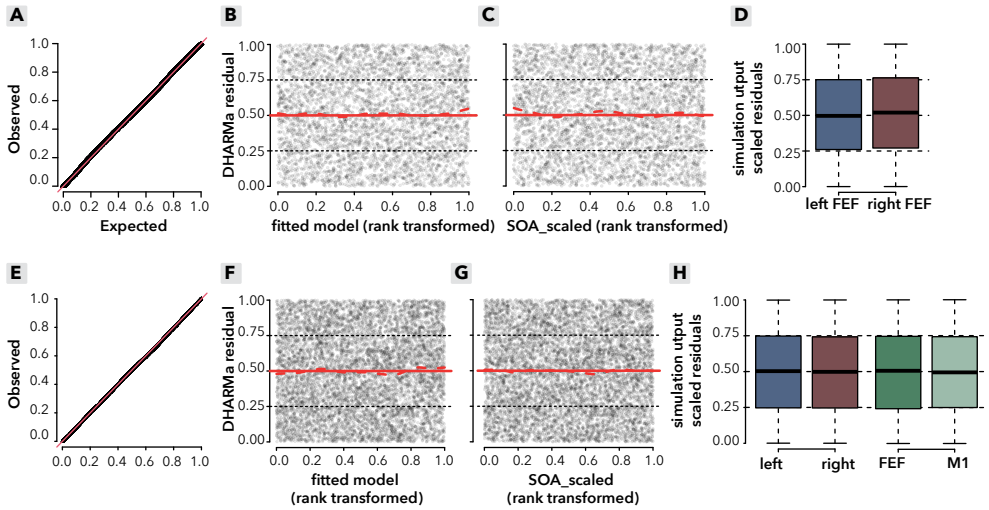


Figure S10 | GLMM diagnostics for two main statistical analyses. (A) FEF-choice domain model: DHARMa quantile–quantile (uniformity) plot of simulation-based residuals; Kolmogorov–Smirnov test: p = 0.13389 (n.s.); dispersion test: p = 0.947 (n.s.); outlier test: p = 1 (n.s.). **(B)** FEF-choice domain model: DHARMa residuals versus model predictions (both rank-transformed). The LOESS smooth (red) is approximately horizontal, indicating no global misfit or excess hetero(dis)persions beyond model expectation. **(C)** FEF-choice domain model: DHARMa residuals versus SOA_scaled (rank-transformed). The near-flat trend indicates no evidence of non-linearity in the logit with respect to SOA; a natural-spline alternative for SOA did not improve fit (LRT $\chi^2(2) = 0.52$, p = 0.77). **(D)** FEF-choice domain model: DHARMa residuals by stimSide (left FEF vs right FEF). Within-group deviation from uniformity: n.s.; Levene test for homogeneity of variance: n.s.; i.e., no heteroscedasticity across levels. **(E)** Side × Region model: DHARMa quantile–quantile (uniformity) plot of simulation-based residuals; Kolmogorov–Smirnov test: p = 0.12922 (n.s.); dispersion test: p = 0.848 (n.s.); outlier test: p = 0.21016 (n.s.). **(F)** Side × Region model: DHARMa residuals versus model predictions (both rank-transformed). The LOESS smooth (red) is approximately horizontal, indicating no global misfit or excess hetero(dis)persions beyond model expectation. **(G)** Side × Region model: DHARMa residuals versus SOA_scaled (rank-transformed). The near-flat trend indicates no evidence of non-linearity on the logit scale; a natural-spline alternative for SOA did not improve fit (LRT $\chi^2(2) = 0.626$, p = 0.731). **(H)** Side × Region model: DHARMa residuals by stimSide (left vs right). Within-group deviation from uniformity: n.s.; Levene test for homogeneity of variance: n.s. **(I)** Side × Region model: DHARMa residuals by stimRegion (FEF vs M1). Within-group deviation from uniformity: n.s.; Levene test for homogeneity of variance: n.s.

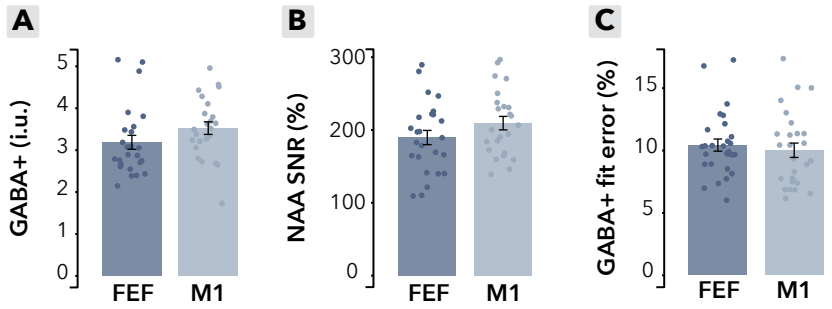


Figure S11 | MRS-derived measurements for the left FEF and left M1 Voxels. (A) Estimated GABA+ levels (in institutional units, i.u.) for each participant in the left FEF (blue) and left M1 (light blue) voxels, derived using Gannet. **(B)** N-Acetylaspartate (NAA) signal-to-noise ratio (SNR) (%) for each participant in the left FEF (blue) and left M1 (light blue) voxels. The NAA SNR reflects the strength and quality of the NAA peak relative to background noise and serves as an indicator of data quality. **(C)** GABA+ fit error (%) for each participant in the left FEF (blue) and left M1 (light blue) voxels. The fit error indicates the quality of the GABA+ estimation, with lower percentages reflecting better fits. Data file. **(A)**, **(B)**, **(C)** Each dot represents an individual participant, with bars showing the group mean and error bars indicating the standard error of the mean (SEM); $n = 25$.

Transducers and drive systems description

	model number, manufacturer	centre frequency	radius of curvature	aperature diameter	number elements	element distribution
transducer	NeuroFUS CTX250-009, Sonic Concepts	250 kHz	63.2 mm	45.5 mm	2	Annular array comprising a circular innter element and an annular outer element of equal area
matching integrated drive system	NeuroFUS TPO-105, Sonic Concepts					

	model number, manufacturer	centre frequency	radius of curvature	aperature diameter	number elements	element distribution
transducer	NeuroFUS CTX250-014, Sonic Concepts	250 kHz	63.2 mm	45.5 mm	2	Annular array comprising a circular innter element and an annular outer element of equal area
matching integrated drive system	NeuroFUS TPO-203, Sonic Concepts					

Drive system settings

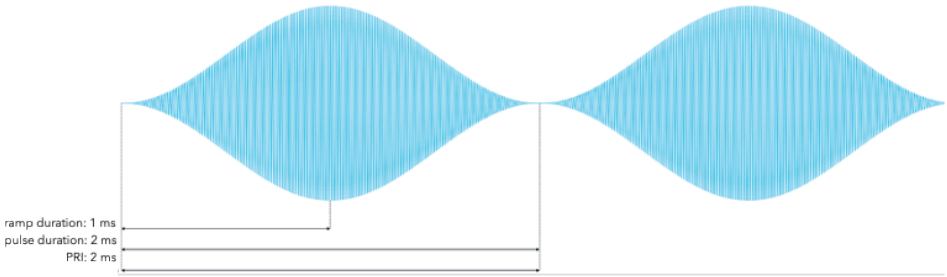
operating frequency	output level setting	focal positioning setting
250 kHz	25 W/cm²(Isppa)	32.5 mm (FLHM)

Free field acoustic parameters

	model number, manufacturer	spatial-peak pressure amplitude	axial position spatial-peak pressure	position of centre of axial -3dB pressure	axial -3dB width	axial -6dB width
transducer	NeuroFUS CTX250-009, Sonic Concepts	0.9 MPa	30 mm	33 mm	25 mm	37 mm
transducer	NeuroFUS CTX250-014, Sonic Concepts	0.9 MPa	31 mm	34 mm	23 mm	34 mm

Pulse timing parameters

	duration	ramp duration	ramp shape	repitition interal / frequency
pulse	2 ms	1000 µs	tukey	2 ms / 500 Hz
pulse train	500 ms			



Intensity parameters

spatial-peak pulse-average intensity	spatial-peak time-average intensities	acoustic impedance
9.37W/cm²	9.37 W/cm²	1.5x10⁶ Rayls

Figure S12 | Transcranial ultrasonic stimulation standardized reporting details

SUPPLEMENTARY TABLES

This section provides a summary of key findings from supplementary statistical analyses:

Table S1 | Additional analyses and statistical outcomes examining the effect of condition and delay on choice behavior. The model was tested across three different choice domains (**25%-75%**, **20%-80%**, and **15%-85%**) to assess the robustness of the effects. **Bold cells** indicate column and row labels, while **shaded blue cells** highlight the main comparison outcomes. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

Model:	*choice ~ condition*$_{(left\ FEF,\ right\ FEF)}$ *+ SOA + (condition*$_{(left\ FEF,\ right\ FEF)}$ *+ SOA | sub)*														
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**
intercept	10	1	0.8	0.14	0.1, 0.2	12	1	<0.001	0.1	0.1, 0.2	13	1	<0.001	0.14	0.1, 0.2
LFEF – RFEF	10	1	0.001	-0.25	-0.4, -0.1	10	1	0.002	-0.2	-0.4, -0.1	8	1	0.005	-0.22	-0.4, -0.1
SOA	185	1	<0.001	1.62	1.4, 1.9	347	1	<0.001	2.2	2.0, 2.5	490	1	<0.001	2.82	2.6, 3.1
Model:	*choice ~ condition*$_{(left\ M1,\ right\ M1)}$ *+ SOA + (condition*$_{(left\ M1,\ right\ M1)}$ *+ SOA | sub)*														
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**
intercept	0	1	0.9	0.0	-0.1, 0.1	0	1	0.6	0.0	-0.1, 0.1	3	1	0.09	0.07	-0.0, 0.1
LM1 – RM1	2	1	0.17	-0.1	-0.2, 0.0	3	1	0.10	-0.1	-0.2, 0.0	4	1	0.044	-0.14	-0.3, -0.0
SOA	187	1	<0.001	1.64	1.4, 1.9	359	1	<0.001	2.3	2.0, 2.5	415	1	<0.001	2.98	2.7, 3.3

Table S2 | Additional analyses and statistical outcomes examining the effect of condition and delay on choice behavior, excluding trials where the target delay was 0 seconds. The model was tested across three different choice domains (**25%-75%**, **20%-80%**, and **15%-85%**) to assess the robustness of the effects. **Bold cells** indicate column and row labels, while **shaded blue cells** highlight the main comparison outcomes. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

Model:	*choice ~ condition$_{(left\ FEF,\ right\ FEF)}$ + SOA + (condition$_{(left\ FEF,\ right\ FEF)}$ + SOA | sub)*														
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**
intercept	7	1	0.007	0.12	0.0, 0.2	9	1	0.002	0.13	0.0, 0.2	11	1	<0.001	0.13	0.1, 0.2
LFEF – RFEF	8	1	0.004	-0.24	-0.4, -0.1	8	1	0.005	-0.22	-0.4, -0.1	6	1	0.012	-0.20	-0.4, -0.0
SOA	175	1	<0.001	1.60	1.4, 1.8	339	1	<0.001	2.20	2.0 – 2.4	481	1	<0.001	2.82	2.6, 3.1
Model:	*choice ~ condition$_{(left\ M1,\ right\ M1)}$ + SOA + (condition$_{(left\ M1,\ right\ M1)}$ + SOA | sub)*														
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**
intercept	1	1	0.4	-0.04	-0.1, 0.1	0	1	0.8	-0.01	-0.1, 0.1	1	1	0.3	0.04	-0.0, 0.1
LM1 – RM1	0	1	0.8	-0.02	-0.1, 0.1	1	1	0.5	-0.05	-0.2, 0.1	2	1	0.2	-0.09	-0.2, 0.0
SOA	175	1	<0.001	1.63	1.4, 1.9	347	1	<0.001	2.28	2.0, 2.5	390	1	<0.001	2.98	2.7, 3.3

Table S3 | Additional analyses and statistical outcomes examining the effect of stimulation side, stimulation region and delay on choice behavior, including an interaction between side and region. The model was tested across three different choice domains (**25%-75%**, **20%-80%**, and **15%-85%**) to assess the robustness of the effects. **Bold cells** indicate column and row labels, while **shaded blue cells** highlight the main comparison outcomes. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

Model:	$choice \sim side_{(left, right)} \times region_{(FEF, M1)} + SOA + (side_{(left, right)} \times region_{(FEF, M1)} + SOA \mid sub)$													
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	p	b	**95%-CI**	χ^2	**Df**	p	b	**95%-CI**	χ^2	**Df**	p	b	**95%-CI**
intercept	9	1	0.003	0.14	0.0, 0.2	8	1	0.004	0.15	0.0, 0.2	6	1	0.011	0.14	0.0, 0.2
side	12	1	<0.001	-0.26	-0.4, -0.1	12	1	<0.001	-0.25	-0.4, -0.1	9	1	0.003	-0.22	-0.4, -0.1
region	5	1	0.032	-0.14	-0.3, -0.0	14	1	0.039	-0.12	-0.2, -0.0	2	1	0.14	-0.08	-0.2, 0.0
SOA	310	1	<0.001	1.63	1.4, 1.8	480	1	<0.001	2.27	2.1, 2.5	557	1	<0.001	2.93	2.7, 3.2
side x region	3	1	0.09	0.16	-0.0, 0.4	3	1	0.11	0.14	-0.0, 0.3	1	1	0.03	0.09	-0.1, 0.3

Table S4 | Additional analyses and statistical outcomes examining the effect of stimulation side, stimulation region and delay on choice behavior, including an interaction between side and region, excluding trials where target delay was 0 seconds. The model was tested across three different choice domains (**25%-75%**, **20%-80%**, and **15%-85%**) to assess the robustness of the effects. **Bold cells** indicate column and row labels, while **shaded blue cells** highlight the main comparison outcomes. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

Model:	*choice ~ side$_{(left, right)}$ x region$_{(FEF, M1)}$ + SOA + (side$_{(left, right)}$ x region$_{(FEF, M1)}$ + SOA | sub)*														
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	*p*	*b*	**95%-CI**	χ^2	**Df**	*p*	*b*	**95%-CI**	χ^2	**Df**	*p*	*b*	**95%-CI**
intercept	7	1	0.009	0.13	0.0, 0.2	7	1	0.009	0.13	0.0, 0.2	6	1	0.018	0.13	0.0, 0.2
side	9	1	0.003	-0.25	-0.4, -0.1	9	1	0.002	-0.24	-0.4, -0.1	7	1	0.008	-0.21	-0.4, -0.1
region	6	1	0.012	-0.18	-0.3, -0.0	5	1	0.020	-0.14	-0.3, -0.0	3	1	0.09	-0.10	-0.2, 0.0
SOA	288	1	<0.001	1.61	1.4, 1.8	468	1	<0.001	2.25	2.0, 2.5	534	1	<0.001	2.91	2.7, 3.2
side x region	5	1	0.025	0.23	0.0, 0.4	4	1	0.038	0.19	0.0, 0.4	2	1	0.15	0.13	-0.0, 0.3

Table S5 | Additional analyses and statistical outcomes examining the effect of stimulation side, stimulation region, presence of zero-delay trials and delay on choice behavior, including an interaction between side and region. The model was tested across three different choice domains (**25%-75%**, **20%-80%**, and **15%-85%**) to assess the robustness of the effects. **Bold cells** indicate column and row labels, while **shaded blue cells** highlight the main comparison outcomes. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

Model:	*choice ~ side$_{(left, right)}$ x region$_{(FEF, M1)}$ x delay0 + SOA + (side$_{(left, right)}$ x region$_{(FEF, M1)}$ + SOA | sub)*														
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**
intercept	7	1	0.008	0.13	0.0, 0.2	10	1	0.002	0.13	0.0, 0.2	5	1	0.020	0.13	0.0, 0.2
side	10	1	0.001	-0.25	-0.4, -0.1	9	1	0.003	-0.24	-0.4, -0.1	8	1	0.006	-0.21	-0.4, -0.1
region	6	1	0.014	-0.17	-0.3, -0.0	4	1	0.040	-0.14	-0.3, -0.0	3	1	0.09	-0.10	-0.2, 0.0
delay 0	0	1	0.5	0.08	-0.2, 0.4	0	1	0.5	0.8	-0.2, 0.3	0	1	0.5	0.08	-0.2, 0.3
SOA	309	1	<0.001	1.63	1.4, 1.8	476	1	<0.001	2.26	2.1, 2.5	554	1	<0.001	2.92	2.7, 3.2
side x region	5	1	0.033	0.22	0.0, 0.4	3	1	0.066	0.19	-0.0, 0.4	2	1	0.15	0.13	-0.0, 0.3
side x del 0	0	1	0.74	-0.06	-0.4, 0.3	0	1	0.6	-0.09	-0.4, 0.3	0	1	0.6	-0.10	-0.5, 0.3
region x del 0	2	1	0.16	0.26	-0.1, 0.6	2	1	0.2	0.23	-0.1, 0.6	1	1	0.3	0.18	-0.2, 0.5
side x region x delay 0	4	1	0.058	-0.50	-1.0, 0.0	3	1	0.09	-0.45	-1.0, 0.1	2	1	0.14	-0.38	-0.9, 0.1

Table S6 | Additional analyses and statistical outcomes examining the interaction between condition and baseline GABA+ levels in the FEF on choice behavior. The model was tested across three different choice domains (**25%-75%**, **20%-80%**, and **15%-85%**) to assess the robustness of the effects. **Bold cells** indicate column and row labels, while **shaded blue cells** highlight the main comparison outcomes. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

Model:	*choice ~ condition$_{(left\ FEF,\ sham)}$ x FEF_GABA + SOA + (condition$_{(left\ FEF,\ sham)}$ + SOA | sub)*														
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	*p*	*b*	**95%-CI**	χ^2	**Df**	*p*	*b*	**95%-CI**	χ^2	**Df**	*p*	*b*	**95%-CI**
condition	6	1	0.011	0.64	0.1, 1.2	9	1	0.003	0.66	0.2, 1.2	9	1	0.003	0.65	0.2, 1.1
FEF GABA	2	1	0.12	0.12	0.0, 0.2	1	1	0.3	0.10	-0.0, 0.2	2	1	0.14	0.16	0.0, 0.3
SOA	269	1	<0.001	1.76	1.5, 2.0	433	1	<0.001	2.23	2.0, 2.4	509	1	<0.001	2.88	2.6, 3.1
condition x FEF GABA	3	1	0.070	-0.15	-0.3, 0.0	4	1	0.050	-0.15	-0.3, -0.0	4	1	0.037	-0.15	-0.3, -0.0
sham x FEF GABA	5	1	0.019	0.12	0.0, 0.2	5	1	0.022	0.10	0.0, 0.2	4	1	0.050	0.15	0.0, 0.3
LFEF x FEF GABA	0	1	0.6	-0.03	-0.2, 0.1	0	1	0.5	-0.05	-0.2, 0.1	0	1	0.9	0.01	-0.2, 0.2

Table S7 | Additional analyses and statistical outcomes examining the interaction between condition and baseline GABA+ levels in the FEF on choice behavior, excluding trials where target delay was 0 seconds. The model was tested across three different choice domains (**25%-75%**, **20%-80%**, and **15%-85%**) to assess the robustness of the effects. **Bold cells** indicate column and row labels, while **shaded blue cells** highlight the main comparison outcomes. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

Model:	*choice ~ condition$_{(left\ FEF,\ sham)}$ x FEF_GABA + SOA + (condition$_{(left\ FEF,\ sham)}$ + SOA | sub)*														
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**	**χ^2**	**Df**	***p***	***b***	**95%-CI**
condition	5	1	0.021	0.84	0.3, 1.4	8	1	0.005	0.84	0.3,1.4	8	1	0.004	0.79	0.3, 1.3
FEF GABA	2	1	0.13	0.14	0.0, 0.2	1	1	0.4	0.11	0.0, 0.2	2	1	0.19	0.16	0.0, 0.03
SOA	274	1	<0.001	1.75	1.5, 2.0	432	1	<0.001	2.21	2.0, 2.4	546	1	<0.001	2.87	2.6, 3.1
condition x FEF GABA	6	1	0.017	-0.21	-0.4, -0.0	6	1	0.011	-0.21	-0.4, -0.0	6	1	0.011	-0.19	-0.3, -0.0
sham x FEF GABA	7	1	0.008	0.14	0.0, 0.2	4	1	0.041	0.11	0.0, 0.2	4	1	0.035	0.16	0.0, 0.3
LFEF x FEF GABA	1	1	0.3	-0.08	-0.2, 0.1	2	1	0.18	-0.09	-0.2, 0.0	0	1	0.7	-0.03	-0.2, 0.1

Table S8 | Additional analyses and statistical outcomes examining the interaction between condition and baseline GABA+ levels in the M1 on choice behavior. The model was tested across three different choice domains (**25%-75%**, **20%-80%**, and **15%-85%**) to assess the robustness of the effects. **Bold cells** indicate column and row labels, while **shaded blue cells** highlight the main comparison outcomes. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

Model:	*choice ~ condition$_{(left\ M1,\ sham)}$ x M1_GABA + SOA + (condition$_{(left\ M1,\ sham)}$ + SOA | sub)*														
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	*p*	*b*	**95%-CI**	χ^2	**Df**	*p*	*b*	**95%-CI**	χ^2	**Df**	*p*	*b*	**95%-CI**
condition	0	1	1.0	-0.24	-0.9, 0.4	0	1	0.7	-0.22	-0.9, 0.5	1	1	0.4	-0.01	-0.6, 0.6
M1 GABA	0	1	0.1	-0.03	-0.2, 0.1	1	1	0.3	-0.06	-0.2, 0.0	1	1	0.4	-0.09	-0.3, 0.1
SOA	191	1	<0.001	1.72	1.5, 2.0	264	1	<0.001	2.24	2.0, 2.5	369	1	<0.001	2.97	2.7, 3.3
condition x M1 GABA	1	1	0.5	0.07	-0.1, 0.3	1	1	0.5	0.07	-0.1, 0.3	0	1	0.8	0.02	-0.1, 0.2
sham x M1 GABA	0	1	0.5	-0.04	-0.2, 0.1	1	1	0.3	-0.06	-0.2, 0.1	1	1	0.4	-0.08	-0.3, 0.1
LM1 x M1 GABA	0	1	0.7	0.03	-0.1, 0.2	0	1	0.9	-0.01	-0.2, 0.2	1	1	0.4	-0.07	-0.2, 0.1

Table S9 | Additional analyses and statistical outcomes examining the interaction between condition and baseline GABA+ levels in the M1 on choice behavior, excluding trials where target delay was 0 seconds. The model was tested across three different choice domains (**25%-75%**, **20%-80%**, and **15%-85%**) to assess the robustness of the effects. **Bold cells** indicate column and row labels, while **shaded blue cells** highlight the main comparison outcomes. Statistical significance was determined using two-sided logistic mixed effects regressions. No multiple comparisons were applied.

Model:	*choice ~ condition$_{(left\ M1,\ sham)}$ x M1_GABA + SOA + (condition$_{(left\ M1,\ sham)}$ + SOA | sub)*														
Data:	choice domain 25% - 75%					choice domain 20% - 80%					choice domain 15% - 85%				
	χ^2	**Df**	*p*	*b*	**95%-CI**	χ^2	**Df**	*p*	*b*	**95%-CI**	χ^2	**Df**	*p*	*b*	**95%-CI**
condition	0	1	0.6	-0.24	-1.0, 0.5	0	1	1.0	0.66	0.2, 1.2	0	1	0.7	0.01	-0.6, 0.6
M1 GABA	0	1	0.8	-0.01	-0.1, 0.1	0	1	0.8	0.10	-0.0, 0.2	1	1	0.5	-0.06	-0.2, 0.1
SOA	185	1	<0.001	1.71	1.5, 2.0	251	1	<0.001	2.23	2.0, 2.4	364	1	<0.001	2.96	2.7, 3.3
condition x M1 GABA	0	1	0.6	0.06	-0.1, 0.3	0	1	0.5	-0.15	-0.3, -0.0	0	1	1.0	0.01	-0.2, 0.2
sham x M1 GABA	0	1	0.9	-0.01	-0.1, 0.1	0	1	0.6	-0.03	-0.2, 0.1	0	1	0.5	-0.06	-0.2, 0.1
LM1 x M1 GABA	0	1	0.5	0.06	-0.1, 0.2	0	1	0.9	0.01	-0.2, 0.2	0	1	0.6	-0.05	-0.2, 0.1




PART III
Interlude
Chapter 3

JAMSHID'S CUP

What it must be like to be all-seeing, to look upon what cannot ordinarily be seen, to glimpse what is hidden.

In a way, we are lucky. We can do something surprisingly close to that. We can watch the living brain as it works, and from those faint signatures begin to understand how decisions are made. That is exactly what this chapter tries to do: to witness the unseen, to observe the neural choreography that unfolds when people choose.

It brings to mind an old Persian legend: **Jamshid's cup**, the Jām-e Jam. This is a vessel said to reveal the world beyond ordinary reach, unbounded by distance, and in some tellings, by time. Jamshid could look into it and see what was happening elsewhere, what might happen later, and what currents were moving beneath the surface. A king, faced with uncertainty, using vision not to dominate the world, but to guide his judgments within it: to see, to observe, to learn – and then, to decide.

In Ferdowsi's imagination and in the wider tradition around it, Jamshid's reign becomes a symbol of flourishing. He is remembered as the ruler who ushered in Iran's first golden age: crafts refined and multiplied – metalworking, weaving – alongside knowledge and care: medicine, scents, the small technologies of everyday life that make a society feel like it is becoming more itself. And he introduced something quieter but just as enduring: the New Year, Nowruz, the "new day"; a celebration not in the dead of winter, but at the hinge of spring. Even now, centuries later, many Iranians around the world feel the year truly turn not on January 1st, but around March 21st, when light returns and the world quite literally begins to bloom again. It is hard not to find that poetic: a calendar anchored to renewal, to growth, to the promise that what was dormant can become alive.

All this, the legend implies, because Jamshid could see beyond the immediate moment.

But the stories also insist on a caution. Because seeing is not the same as controlling. And insight, when it becomes certainty, can curdle into overconfidence. Drunk on his own splendor, Jamshid's vision hardened into arrogance. And under that arrogance, his land entered dark days: the light thinned, the world grew dim, and in that dimness, demons were allowed to rise. In the end, it was one of these forces, born of the darkness his pride had helped invite, that became the downfall of the once-great king.

Let it be known that seeing all is not the same as knowing all. Observation is powerful, however not omniscience. The cup shows; it does not command.

But let us not remember Jamshid only for his darkest chapter. For what he brought into being, is not forgotten to this day, for the flourishing he made possible, for the golden age he ushered in, for the “new day” he gave his people. Not as a warning alone, but as a reminder: that the ability to see further can build something lasting – if it is carried with humility.

CHAPTER 3

Neural signatures of the nucleus accumbens and amygdala in motivated behavior

ABSTRACT

Pavlovian biases and instrumental learning jointly shape motivated behavior: reward-predictive cues tend to invigorate action and punishment-predictive cues promote inhibition (Pavlovian biases), while outcomes are used to learn and update which actions are worth repeating (instrumental learning) – sometimes requiring overcoming those biases. The ventral striatum (nucleus accumbens, NAcc) and amygdala are central to these processes, but in humans their respective contributions remain incompletely understood. Here, we used fMRI to examine how these regions support Pavlovian control and instrumental learning during a motivational Go/NoGo task in 29 healthy adults. Cues signaled either potential reward or potential loss, fully crossed with required action (Go vs. NoGo). Unlike earlier versions of the task, repeated introduction of novel cues throughout the task with probabilistic feedback kept participants in an ongoing learning regime and allowed us to contrast novel versus familiar cues. Behaviorally, participants showed clear instrumental learning (more and faster Go responses when Go was correct) and a robust Pavlovian bias (more and faster responses for Win than Avoid cues). Novelty sensitivity was also evident in behavior: outcome on the very first encounter with a cue had a disproportionately strong impact on subsequent choice behavior compared with later repetitions. At cue onset, BOLD responses in NAcc and ventromedial prefrontal cortex were higher for Win than Avoid cues, whereas amygdala activity did not encode cue valence. At outcome, both NAcc and amygdala responded more strongly to desired than non-desired outcomes, consistent with a shared role in outcome evaluation. Exploratory analyses of novelty further revealed elevated amygdala, but not NAcc, responses for the first presentation of a cue, irrespective of valence. Together, these findings suggest a differentiated but coordinated division of labor between the NAcc and amygdala. The amygdala appears particularly engaged in novelty-sensitive cue processing, boosting fast initial learning. The NAcc, by contrast, supports anticipatory Pavlovian bias and, together with the amygdala, encodes feedback signals that can train future action values.

INTRODUCTION

When seeing a slice of cake on the table, the smell and prospect of this sweet reward can pull you toward it; conversely, a sudden encounter with a spider may make you freeze and hold back. Both reactions reflect Pavlovian control: a seemingly hard-wired tendency to invigorate action when presented with the prospect of reward, and to promote inhibition for punishment- or loss-predictive cues[1–3]. These mechanisms are fast and often adaptive, but also relatively inflexible in everyday life. Sometimes they need to be overridden – for instance, when the cake is still in the oven, repeatedly opening the door only ruins the bake. In that case, the better strategy is to refrain from acting and wait. Such flexibility is provided by instrumental learning, in which actions are selected and flexibly updated based on their experienced consequences[1–3]. Together, Pavlovian control and instrumental learning govern motivated behavior: they cooperate when "go for reward" and "hold back to avoid loss" are optimal but come into conflict when reward requires inhibition or avoidance requires action[1–3]. In several psychiatric conditions, this balance is disrupted[4–8]. For instance, in addiction, where achieving long-term high-value outcomes often requires sustained restraint in the face of a powerful approach trigger[4]. Understanding how these systems shape behavior, and which brain regions support this interaction, is crucial for developing mechanistic models of motivated behavior and for explaining maladaptive patterns observed in psychiatry[4–12]. The present fMRI study contributes to this goal by examining the neural underpinnings of the interplay between Pavlovian control and instrumental learning.

Two key subcortical structures associated with both instrumental and Pavlovian control are the ventral striatum (nucleus accumbens, NAcc) and the amygdala. Previous work has established a central role for the NAcc in instrumental control[13–19], that is, in selecting actions based on their expected consequences and updating those choices when outcomes change. In theories of reinforcement-learning, dopamine projections to the NAcc are proposed to encode reward prediction errors, driving value updates: dopaminergic inputs signal whether an outcome is better or worse than expected, and NAcc uses this information to adjust the encoded value of a given action[17,18,20–22]. Direct evidence for the role of striatal dopamine as a learning signal comes from optogenetic studies in which manipulation of dopaminergic inputs to the NAcc is sufficient to bias instrumental learning, strengthening or weakening specific actions depending on how and when dopamine is released[23]. Human fMRI work converges with this picture: in a range of instrumental learning tasks, ventral striatal BOLD signals covary with outcome value and, specifically, reward prediction errors (the signed difference between expected and obtained outcomes) consistent with a role of NAcc in encoding feedback that trains future choices[10,20,21].

In addition to instrumental learning, NAcc is also implicated in Pavlovian control. Rodent Pavlovian-instrumental transfer (PIT) studies show that reward-paired cues invigorate responding and that NAcc is necessary for this cue-evoked action invigoration[14,16,19]. These paradigms present cues that have acquired motivational value through prior Pavlovian conditioning (i.e., repeated pairing with reward or punishment) during ongoing instrumental performance, allowing for a direct measure of how learned reward or punishment associations energize or suppress actions independently of their instrumental relevance. In humans, however, the picture is more mixed. Monetary incentive delay tasks show that high-reward cues speed up responding and elicit stronger cue-locked NAcc activity, which scales with anticipated reward[24,25]. PIT paradigms similarly demonstrate that appetitive cues energize responding, whereas aversive cues tend to suppress it, with both effects reflected in striatal BOLD signals[26,27]. Together, these findings suggest that the NAcc can link learned cue–outcome associations to both action invigoration for appetitive cues and action suppression for aversive cues. However, studies that explicitly orthogonalize action requirement and valence, such that rewards may require both invigoration and inhibition of action, do not consistently find NAcc involvement in cue valence coding[9,28,29]. In these motivational Go/NoGo learning tasks, cues indicate both the valence of the potential outcome (Win vs. Avoid loss) and whether a Go or NoGo action is required. Instead, cue presentation elicited BOLD responses for those cues that required a response, independent of cue valence[9,28,29]. Thus, while NAcc involvement in instrumental control is well established, there remain discrepancies between animal and human work regarding its role in valence coding and Pavlovian control.

Like the NAcc, the amygdala is a key node in motivational action control. This role is most clearly demonstrated in Pavlovian conditioning, where a neutral cue gradually acquires motivational value because it predicts an appetitive or aversive outcome[27,30–33]. In these paradigms, the cue-outcome association itself needs to be learned. Importantly, once these associations are established, the amygdala contributes to the expression of Pavlovian influences on instrumental choice, for example by biasing which actions are selected in the presence of reward- or loss-predictive cues[14,26,34,35]. Anatomically, the basolateral part of the amygdala (BLA) sends dense projections to the NAcc, forming an amygdalo-striatal circuit in which affective value signals originating in the amygdala are in an excellent position to shape striatal activity and, through that, influence how strongly particular actions are selected[13,14,34].

A second, complementary aspect of amygdala function concerns rapid learning in the context of novelty and change[36–42]. Amygdala neurons respond strongly when motivationally salient cues are encountered for the first time and show heightened sensitivity to the emergence of new, behaviorally relevant stimuli[43,44]. In functional

terms, this pattern supports a cue-processing boost that prioritizes new or important stimuli and facilitates rapid updating. In contrast, the NAcc appears to integrate these teaching signals over longer timescales as cue-action-outcome mappings stabilize, providing a slower but more noise-tolerant substrate for value estimates[45–47]. These studies therefore point to an amygdala contribution to initial, high-impact learning, rapidly tagging new cues with affective value, while the NAcc maintains and updates those values as experience accumulates[17,46,47]. Although this division of labor is well established in rodents and non-human primates, its expression in human motivational action control has been less clear, making it an important target for the present fMRI work.

In this fMRI study, an orthogonalized motivational Go/NoGo (MGNG) task was used in which cue valence (Win vs. Avoid) was crossed with required action (Go vs. NoGo). This task design included probabilistic outcomes and staggered introduction of cues that were each presented for relatively few repetitions. This continuous introduction of new cues keeps behavior in an active learning regime, where cue-action-outcome associations were relatively uncertain. Furthermore, because novel and familiar cues were interleaved throughout the task, we could directly contrast their neural responses within subjects, providing leverage to dissociate rapid, novelty-sensitive amygdala contributions from slower, accumulation-based NAcc updating. At the same time, cue presentation and response execution occurred within a single temporal window, and each cue mapped onto one correct Go or NoGo action, bringing the paradigm closer to classic animal PIT where cues more directly invite a specific action. Together, this design allows us to test (i) whether NAcc shows anticipatory valence signals during Pavlovian cue processing, and (ii) how NAcc and amygdala differentially contribute to learning from novel versus familiar cues and to outcome evaluation in a low-contingency environment.

RESULTS

Twenty-nine healthy participants (15 females and 14 males; age range 20–35 years; mean = 25.1, SD = 4.1 years) performed a MGNG learning task in the MRI scanner (**Figure 1A**). Cues signaled either potential reward (Win) or potential loss (Avoid), and for each cue participants had to discover through trial and error whether pressing (Go) or withholding (NoGo) a response was the correct action. Correct actions yielded reward on Win cues and avoided loss on Avoid cues on 75% of trials. Because cue valence was orthogonalized with required action (**Figure 1B**), some cue-action pairings were congruent with Pavlovian tendencies (G2W, NG2A) while others were incongruent (NG2W, G2A), allowing us to directly measure how motivational value biases action

selection. Crucially, unlike previous implementations of this paradigm[9,10,12,29,48,49], participants remained in a continuous learning regime: they encountered 24 distinct cues that were gradually introduced and removed across the session, and one third of these cues reversed their required action midway through, preventing performance from reaching a stable asymptote (**Figure 1C**, trial structure adapted from Zorowitz and colleauges[50]).

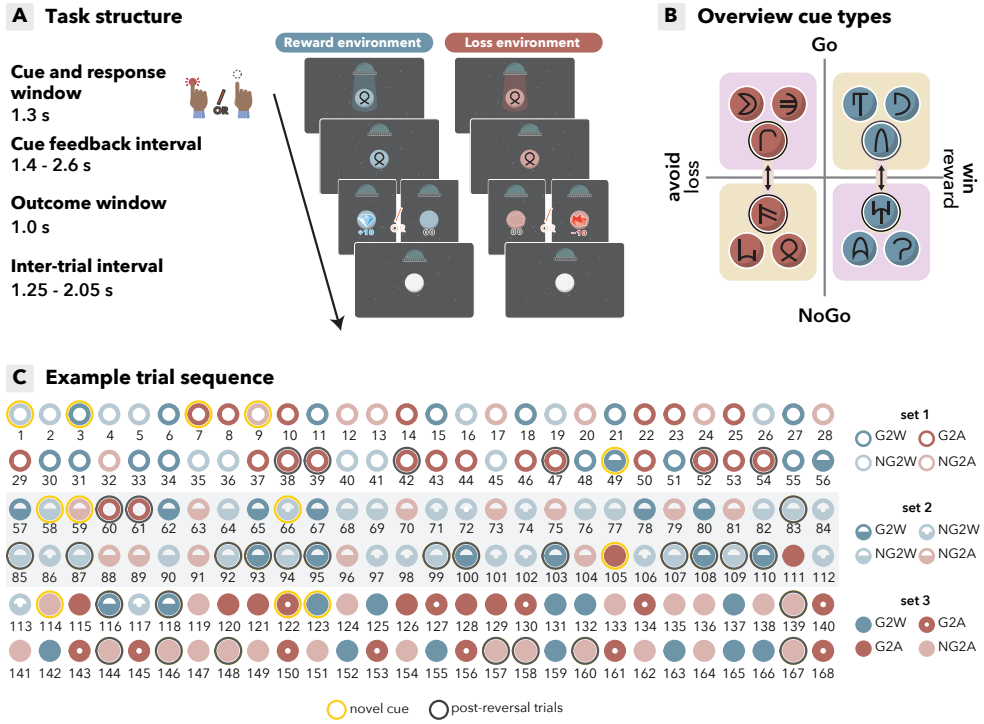


Figure 1 | Task structure and cue schedule in the MGNG learning task. (A) *Trial timeline*. On each trial, a cue ("planet") appeared and participants could either press a button (Go) or withhold their response (NoGo) within a 1.3 s cue/response window. This was followed by a jittered cue–feedback interval (1.4–2.6 s), a 1.0 s outcome window, and a jittered inter-trial interval (1.25–2.05 s). Trials were presented in either a *reward* environment (desired outcome: winning points) or a *loss* environment (desired outcome: avoiding point loss). **(B)** *Cue types*. Four cue types orthogonalized required action and valence: Go-to-Win-reward (G2W), Go-to-Avoid-loss (G2A), NoGo-to-Win-reward (NG2W) and NoGo-to-Avoid-loss (NG2A). In each run, three distinct cues instantiated each condition (12 cues in total). For each cue, outcomes were probabilistic: correct responses produced the desired outcome on ~75% of trials and the neutral/undesired outcome on ~25% of trials, with analogous contingencies for Win and Avoid cues. For the subset of cues presented 16 times, the required action reversed after the 8th presentation. Yellow panels indicate the congruent conditions, which align with our Pavlovian tendencies and are typically easier; purple panels indicate the incongruent, which conflict with our Pavlovian bias and are more difficult. Participants completed two such runs in the scanner. **(C)** *Example trial sequence for one run (168 trials).* The sequence is organized into three "sets" of trials (~56 trials

per set), indicated by ring fill (empty, half-filled, filled). Each set contains four cues (color-coded by condition: dark blue = G2W, light blue = NG2W, light red = G2A, dark red = NG2A). Not all four action–valence combinations necessarily appear in every set, but each set always contains at least one Go and one NoGo cue (i.e., required actions are never all the same within a set). Within a set, dots on the circles indicate different cues of the same condition (e.g. two distinct NG2W cues in set 2). Because cues are presented 12, 14 or 16 times, sets overlap in time, and cues from different sets are interleaved. Cues with a yellow outline mark their first (novel) presentation. For cues that reverse their required action after the 8th presentation, post-reversal trials are shown in grey; reversals are distributed across the run such that each set contains at least one and at most two reversing cues.

Successful instrumental learning and stable Pavlovian bias

To quantify Pavlovian and instrumental learning-related influences on choice, we fitted a logistic mixed-effects regression to the Go/NoGo data with three fixed factors: required action (Go vs. NoGo), cue valence (Win vs. Avoid), and cue-specific trial number (z-scored), plus all two-way interactions.

Participants showed clear instrumental learning: they made more Go responses to Go than NoGo cues (main effect of required action: b = 0.89, 95%-CI [0.70, 1.07], $\chi^2(1)$ = 91.2, p < 2.2e-16; **Figure 2A**), which became stronger with cue repetition (required action × trial number: b = 0.37, 95%-CI [0.24, 0.49], $\chi^2(1)$ = 33.7, p = 6.6e-09; **Figure 2B-C**). Because i) new cues were introduced continuously and ii) the probabilistic nature of the task, group level accuracy did not reach ceiling – participants continued to make errors throughout the task rather than approaching perfect performance – preserving a wide dynamic range for detecting bias-related effects. Crucially, choices also evidenced a robust Pavlovian bias: participants made more Go responses for Win than Avoid cues (main effect of cue valence: b = 0.25, 95%-CI [0.13, 0.36], $\chi^2(1)$ = 18.5, p = 1.7e-05; **Figure 2A**). Mirroring earlier implementations of this task, there was no evidence of significant learning ability differences from Win vs. Avoid cues (required action × cue valence: b = -0.03, 95%-CI [-0.05, 0.12], $\chi^2(1)$ = 0.6, p = .4) nor that Pavlovian bias changed significantly over time (cue valence × trial number: b = -0.03, 95%-CI [-0.11, 0.05], $\chi^2(1)$ = 0.7, p = .4).

Regarding reaction times (RTs), participants made correct Go responses (i.e., to Go cues) more quickly than incorrect responses (i.e., to NoGo cues; main effect of required action: b = -0.02, 95%-CI [-0.02, -0.01], $\chi^2(1)$ = 16.5, p = 4.8e-05; **Figure 2D**), corroborating the learning effect observed in the Go-response model. Further evidence of learning was observed as this difference between Go and NoGo response times increased over trials (required action × trial number: b = -0.01, 95%-CI [-0.02, -0.0003], $\chi^2(1)$ = 4.1, p = .042). Importantly, participants were also faster to respond to Win than Avoid cues (main effect of cue valence: b = -0.02, 95%-CI [-0.03, -0.01], $\chi^2(1)$ = 28.7, p = 8.6e-08; **Figure 2D**), reflecting a robust Pavlovian speeding. Moreover,

RTs shortened over successive presentations of the same cue (main effect of trial number: b = -0.02, 95%-CI [-0.02, -0.01], $\chi^2(1)$ = 31.1, p = 2.4e-08). As for choices, the cue-valence effect did not interact with required action (cue valence × required action: b = 0.002, 95%-CI [-0.01, 0.01], $\chi^2(1)$ = 0.4, p = .5). Thus, we did not find evidence that the Pavlovian boost to motor vigor depended on whether a Go press was correct or erroneous. Finally, there was no evidence of Pavlovian speeding change across trials (absence of a cue valence × trial number interaction: b = -0.002, 95%-CI [-0.01, 0.01], $\chi^2(1)$ = 0.2, p = .7).

BOLD analyses

To link the behavioral and modelling results to BOLD activity, we estimated a GLM at the first and second level. At cue onset, we modelled four regressors corresponding to the experienced action–valence combinations (G2W, G2A, NG2W, NG2A), defined trial by trial based on participants' actual responses rather than the instructed condition (e.g., a Go response to a Win cue contributed to the G2W regressor). At outcome onset, we modelled desired outcomes (winning reward, avoiding loss) versus non-desired outcomes (loss, no reward). All regressors were modelled as a single impulse (zero-duration event) convolved with the canonical hemodynamic response function and entered into a random-effects group analysis; whole-brain results were thresholded at a voxel-wise cluster-defining threshold of p < .001 (uncorrected) combined with cluster-level family-wise error correction at p_FWE < .05; region-of-interest (ROI) definitions and additional trial-specific GLMs are described in the Methods.

Nucleus accumbens is involved in cue processing

At the time of cue presentation, a whole-brain Win > Avoid contrast revealed significant bilateral activity in the ventral striatum (nucleus accumbens, NAcc), a pattern that was corroborated by NAcc region-of-interest (ROI) analyses. In the whole-brain analysis, the same contrast also showed robust activation in ventromedial prefrontal cortex (vmPFC), broadly overlapping the ventromedial / rostral anterior cingulate region previously implicated in cue valence and motivational value signaling in similar tasks (**Figure 3A**; **Table S1**)[9,51]. In contrast, neither the whole-brain analysis nor amygdala ROI analyses revealed any significant cue-valence effects in the amygdala for this contrast (**Figure 3B**). For the reverse contrast Avoid > Win, there was no evidence of significant NAcc or amygdala engagement in either the whole-brain or ROI-based analyses.

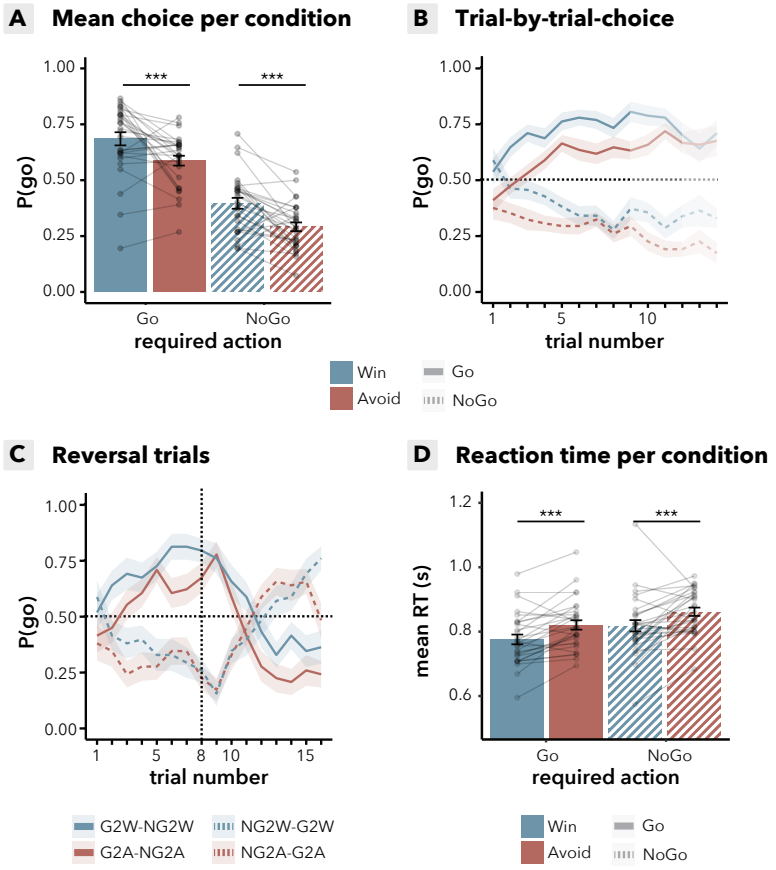


Figure 2 | Behavioral results. (A) *Mean choice per condition*. The probability of making a Go response (P(Go)) as a function of required action (Go vs. NoGo) and cue valence (Win, blue; Avoid, red). Participants made more Go responses when a Go action was required (instrumental learning) and more Go responses for Win than Avoid cues (Pavlovian bias). Bars show group means ± SEM; grey dots are individual participants, connected by lines across valence conditions. **(B)** *Trial-by-trial choice behavior*. P(Go) across cue repetitions for each condition (blue = Win, red = Avoid; solid lines = Go-required cues, dashed lines = NoGo-required cues). Lines show group means and shaded ribbons SEM. Post-reversal trials are omitted, so the number of contributing cues decreases after repetitions 8 and 12, indicated by increasing line translucency. Participants gradually learn the required actions without reaching ceiling, and the Pavlovian bias (Win > Avoid) persists across repetitions for both Go and NoGo cues. **(C)** *Choice behavior for reversal cues*. P(Go) is shown for cues that reversed their required action after the 8th repetition (vertical dashed line). Colors indicate valence (blue = Win, red = Avoid); solid lines = cues that require Go before the reversal and NoGo afterwards; dashed lines = cues that require NoGo before the reversal and Go afterwards. Lines show group means and ribbons SEM, illustrating that participants adjust their actions following the reversal. **(D)** *Mean reaction time per condition*. Mean RT on trials with a button press as a function of required action and cue valence. Participants responded faster on Go than NoGo trials and faster to Win than Avoid cues. Bars show group means ± SEM; grey dots are individual participants connected across valence conditions. Asterisks indicate significant effects (* $p < .05$, ** $p < .01$, *** $p < .001$).

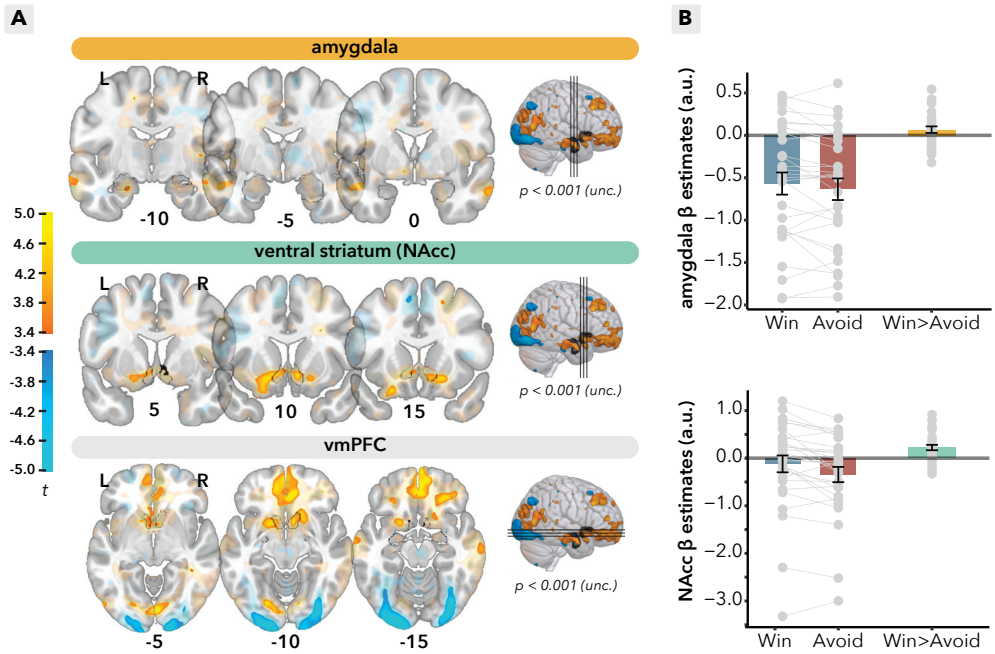


Figure 3 | Cue-valence effects at cue onset. (A) *Whole-brain cue-locked contrast Win > Avoid.* Warm colors show voxels where BOLD responses were stronger for Win than Avoid cues; cool colors show the reverse contrast (Avoid > Win). Maps are thresholded at voxel-wise p < .001 uncorrected (T ≥ 3.4; opaque colors) with subthreshold trends at p < .10 shown transparently. Top row: no reliable valence effects are observed in amygdala (black outline). Middle row: Win cues elicit robust bilateral activation in NAcc (black outline). Bottom row: additional cue-valence effects are present in ventromedial prefrontal cortex (vmPFC). Amygdala and NAcc outlines indicate group-level anatomical masks derived from individual FreeSurfer segmentations normalized to MNI space. Axial and coronal slices are displayed on the MNI152 template; the 3D inset brain indicates slice positions. **(B)** *ROI-based cue-valence effects.* Top: mean beta estimates (± SEM) in amygdala for Win and Avoid cues, and their difference (Win > Avoid). Bottom: corresponding beta estimates in NAcc. Grey dots represent individual participants; dots from the same participant are connected by thin lines across conditions. NAcc shows a clear positive Win > Avoid effect, whereas amygdala does not differentiate between Win and Avoid cues.

Both nucleus accumbens and amygdala process outcome

At the time of outcome, both NAcc (both whole-brain and ROI) and vmPFC (whole-brain) BOLD responses were higher for desired outcomes (obtaining reward or avoiding loss) than for non-desired outcomes (loss or omitted reward; **Figure 4**; **Table S2**), consistent with a role in encoding the value of obtained outcomes and favorable feedback[10,20,21,52,53]. Importantly, and in contrast to the cue-valence analyses, the same Desired > Non-desired contrast also elicited reliable BOLD responses in bilateral

amygdala (whole-brain and ROI; **Figure 4**; **Table S2**). Thus, whereas anticipatory cue processing was selectively expressed in NAcc (and vmPFC), outcome processing engaged NAcc, vmPFC and amygdala alike. No NAcc or amygdala activity in either whole-brain or ROI-based analyses was observed for the reverse contrast Non-desired > Desired outcome.

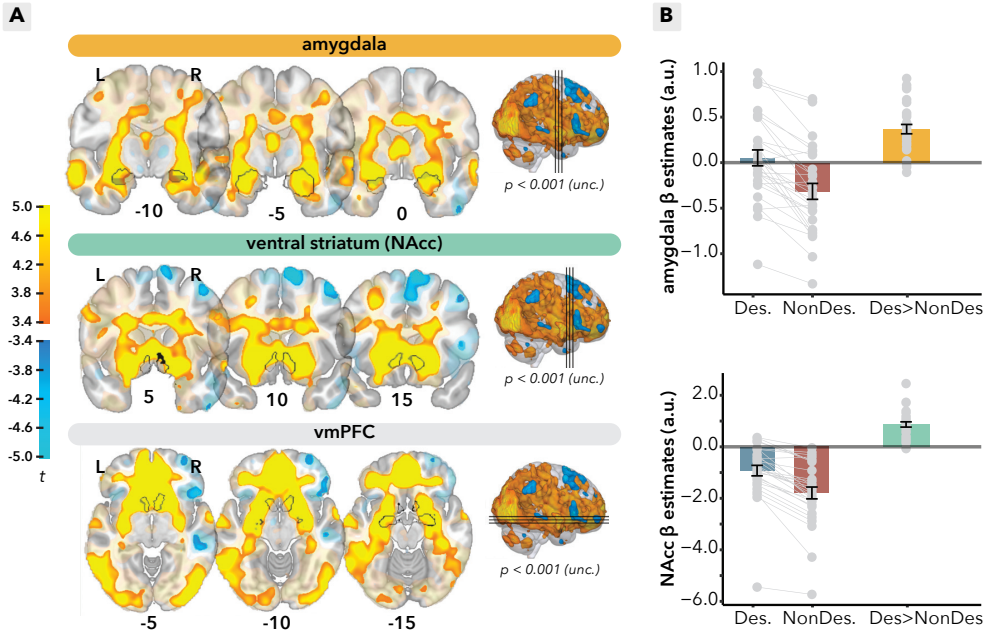


Figure 4 | Outcome-valence effects at outcome. (A) *Whole-brain outcome-locked contrast Desired > Non-desired.* Warm colors show voxels where BOLD responses were stronger for desired outcomes (obtaining reward or avoiding loss) than for non-desired outcomes (loss or omitted reward); cool colors show the reverse contrast (Non-desired > Desired). Maps are thresholded at voxel-wise p < .001 uncorrected (T ≥ 3.4; opaque colors) with subthreshold trends at p < .10 shown transparently. Top row: bilateral amygdala (black outline) shows robust desired > non-desired responses. Middle row: NAcc (black outline) also exhibits strong outcome-valence effects. Bottom row: similar desired > non-desired responses are observed in ventromedial prefrontal cortex (vmPFC). Amygdala and NAcc outlines indicate group-level anatomical masks derived from individual FreeSurfer segmentations normalized to MNI space. Axial and coronal slices are displayed on the MNI152 template; the 3D inset brains indicate slice positions. **(B)** *ROI-based outcome-valence effects.* Top: mean beta estimates (± SEM) in amygdala for Desired and Non-desired outcomes, and their difference (Desired > Non-desired). Bottom: corresponding beta estimates in NAcc. Grey dots represent individual participants; dots from the same participant are connected by thin lines across conditions. Both regions show clear positive outcome-valence effects, with stronger BOLD responses for desired than for non-desired outcomes.

Differential temporal contributions of the NAcc and amygdala to learning

To assess whether novel and familiar cues are processed differently in NAcc and amygdala, cue-locked BOLD responses were analyzed separately for novel and familiar cue presentations. Win and Avoid cues were modelled separately for trial 1 (novel cues), trial 8 (familiar cues), and all remaining trials. Trial 1 was chosen because it is the very first encounter with each cue and therefore captures novelty; trial 8 was selected as a matched "familiar" trial because this is the last repetition before one-third of the cues reversed their required action. Thus, trial 8 is the last trial (i.e., most familiar trial) in which all cues still reflected the original cue–action mapping. Subject-level beta estimates for trial-1 and trial-8 cues were then extracted from orthogonally defined combined functional–anatomical NAcc and amygdala masks (**Figure 5A**; see Methods).

A mixed-effects regression analysis with factors region (NAcc, amygdala) and trial-number (1, 8) showed that cue-locked activity differed significantly between NAcc and amygdala across repetitions (region x trial interaction: b = -0.19, 95%-CI [-0.31, -0.08], $\chi^2(1)$ = 11.3, p = 7.9e-4; **Figure 5B**). There was no evidence that this interaction depended on cue valence (region x trial x valence interaction: b = -0.01, 95%-CI [-0.04, 0.02], $\chi^2(1)$ = 0.6, p = .4; **Figure 5B**), indicating that novelty-related effects were not significantly different for Win and Avoid cues. Post-hoc tests within regions showed that in the amygdala, BOLD responses were clearly elevated for trial-1 relative to trial-8 cues (b = 0.42, 95%-CI [0.30, -0.54], $\chi^2(1)$ = 49.5, p = 1.9e-12; **Figure 5B**). In contrast, NAcc did not show a reliable difference between trial 1 and trial 8 (b = 0.04, 95%-CI [-0.16, 0.23], $\chi^2(1)$ = 0.1, p = .7). Direct between-region comparisons further showed that amygdala responses were higher than NAcc at trial 1 (b = -0.34, 95%-CI [-0.45, -0.24], c2(1) = 38.8, p = 4.6e-10; **Figure 5B**), but did not reliably differ at trial 8 (b = 0.04, 95%-CI [-0.09, -0.17], c2(1) = 0.4, p = 0.5; **Figure 5B**).

To characterize how these putative novelty effects in the amygdala unfold over time, i.e., whether they gradually decline or rather are present only when a completely novel stimulus is present, cue-locked responses were next modelled for each repetition from trial 1 through trial 8. Two mixed-effects regression models were compared: one with a region x trial repetition (1-8) interaction, capturing a gradual trend, and one with a region x novelty (trial-1 vs. trial2-8) factor, capturing a novelty versus everything else step change. Model comparison favored the latter, step-change formulation ($M_{gradual}$ - M_{step}: ΔBIC = 55.4; ΔAIC: 55.3), indicating that the data were better described by a selective peak at the very first presentation than by a smooth decline across repetitions. Consistent with this, amygdala activity showed a pronounced spike on trial 1 and lower, relative stable responses thereafter (**Figure 5C**).

To examine whether NAcc and amygdala also contribute differently to outcome-locked processing (instrumental learning) as cues become familiar, an additional GLM was estimated in which outcome regressors were split by both outcome valence (desired vs. non-desired) and cue repetition (trial 1, trial 8 and remaining trials). For each ROI and trial, we computed a desired–minus–non-desired contrast estimate. This analysis provided no evidence for differences in outcome-locked novelty processing between NAcc and amygdala (region x trial interaction: b = -0.12, 95%-CI [-0.30, 0.05], $\chi^2(1)$ = 2.0, p = .16; **Figure 5D**). There was, however, a main effect of region, such that NAcc showed overall larger outcome-valence responses than amygdala, consistent with a stronger role of NAcc in outcome evaluation across repetitions (b = 0.29, 95%-CI [0.12, 0.47], $\chi^2(1)$ = 10.8, p = .001). A complementary outcome-locked model with separate regressors for each repetition from trial 1 through trial 8 likewise did not reveal clear region-by-repetition interactions (b = 0.01, 95%-CI [-0.03, 0.05], $\chi^2(1)$ = 0.2, p = .6; **Figure 5D**). At the same time, there was evidence for an overall increase in outcome-related BOLD responses across repetitions for both regions (main effect of trial number: b = 0.06, 95%-CI [0.02, 0.10], $\chi^2(1)$ = 7.0, p = .008).

Taken together, the BOLD findings reveal both dissociable and overlapping contributions of the NAcc and amygdala in Pavlovian and instrumental control. At cue onset, the NAcc (together with vmPFC) showed stronger activity for Win than Avoid cues, indicating sensitivity to cue valence at the moment an action can be prepared and consistent with a role in Pavlovian control, where reward-predictive cues preferentially invigorate responding. In contrast, amygdala, but not NAcc, responses were elevated when cues were encountered for the first time, at cue onset and irrespective of valence. This pattern is compatible with a scenario in which amygdala is particularly engaged in novelty-sensitive cue processing, potentially boosting learning of these novel cues. At outcome, both NAcc and amygdala responded more strongly to desired than to non-desired outcomes, where this effect was strongest for the NAcc.

Exploratory and control whole-brain analyses

Exploratory and control whole-brain analyses are reported in **Supplementary Information S.1**, including cue-locked valence contrasts (Avoid > Win), outcome-locked valence contrasts (Non-desired > Desired), and action contrasts (Go > NoGo), along with corresponding brain maps (**Suppl. Figure S1**) and cluster tables (**Suppl. Tables S1 – S3**). Briefly, these analyses showed dorsomedial prefrontal cortex and anterior cingulate cortex engagement for loss-predictive cues, right frontal recruitment following non-desired outcomes, and left motor cortex and supplementary motor area activation for Go responses. Additionally, to examine whether motivational congruency (i.e., the alignment between Pavlovian action tendencies and instrumental

requirements) modulated cue-locked BOLD responses, we computed congruent versus incongruent contrasts (and vice versa). No significant clusters were observed for either direction of this contrast (**Suppl. Information S.2**, **Suppl. Figure S2**).

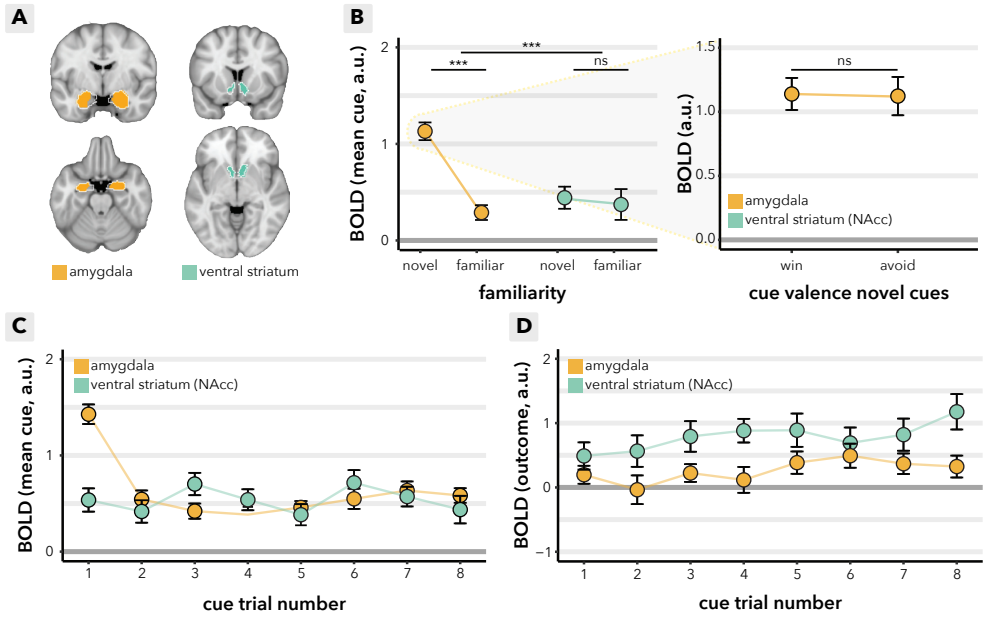


Figure 5 | Differential temporal contributions of nucleus accumbens and amygdala to novelty and familiar-cue processing. (A) Bilateral functional regions of interest (ROIs) for amygdala (orange) and ventral striatum (NAcc; green), defined from the independent outcome-valence contrast Desired > Non-desired (see Methods). **(B)** Left: Cue-locked mean BOLD signal (collapsed across Win and Avoid cues) for novel (trial 1) and familiar (trial 8) cues in amygdala and NAcc. Amygdala responses are strongly elevated for novel relative to familiar cues and exceed NAcc activity at trial 1, whereas NAcc shows no clear novelty-related change. Right: Cue-locked amygdala responses for novel Win versus Avoid cues, indicating no reliable valence dependence of the novelty effect. **(C)** Cue-locked mean BOLD signal across cue repetitions (trials 1–8), showing a pronounced peak in amygdala activity on the first presentation of a cue and lower responses thereafter, while NAcc activity remains comparatively stable across repetitions. **(D)** Outcome-locked Desired > Non-desired effects across cue repetitions (trials 1–8), illustrating an overall gradual increase for amygdala and NAcc activity as cues become familiar. In panels **B–D**, data points depict group means and error bars denote ±SEM.

DISCUSSION

The present study examined how the ventral striatum (nucleus accumbens, NAcc) and amygdala differentially contribute to motivated behavior. We used a motivational Go/NoGo task in which cue valence (Win vs. Avoid) was orthogonalized with required action (Go vs. NoGo), and novel cues were introduced throughout to keep participants in an ongoing learning regime. At cue onset, we observed a clear dissociation: NAcc BOLD responses encoded cue valence, responding stronger to Win than to Avoid cues, whereas amygdala activity did not differentiate between valences but was strongly sensitive to whether cues were novel or familiar. By contrast, both NAcc and amygdala showed increased BOLD responses to desired than to non-desired outcomes at feedback, consistent with a shared role in outcome processing. Together, these findings reveal dissociable but overlapping contributions of NAcc and amygdala to anticipatory valence coding, novelty-driven processing, and outcome evaluation.

Medial fronto-striatal signals during Pavlovian cue processing

Behaviorally, cue valence elicited a robust Pavlovian bias as participants made more and faster Go responses to reward-predictive Win cues than to loss-predictive Avoid cues. Because valence and required action were fully orthogonalized, the same cue type could demand either action (Go or NoGo). This design avoids the common confound in which "reward" is always tied to action execution and "loss" to inhibition and thus allows us to isolate neural signals related to valence per se, rather than value being inseparable from a particular response tendency. Analyzing cue-locked BOLD activity in this context therefore probes the valence signals that can drive Pavlovian biases in action invigoration and inhibition, before feedback is available.

At the neural level, both NAcc and ventromedial prefrontal cortex (vmPFC) showed higher cue-locked BOLD responses for Win than for Avoid cues. The vmPFC effect aligns with previous human MGNG studies[9] and with the broader notion that vmPFC encodes expected value during reward anticipation[52,54,55]. In contrast, ventral striatal cue-valence effects have not been reliably observed in earlier human MGNG fMRI studies using similarly orthogonalized designs[9,28,29]. Notably, however, the presence of cue-locked NAcc valence signals is well-aligned with animal work showing that the NAcc is critical for reward-induced action invigoration and for the impact of motivational cues on responding: when NAcc is lesioned or inactivated, animals no longer show the typical boost in responding to reward-predictive cues[15,16,19]. One plausible explanation for this discrepancy is that, in prior human implementations of the orthogonalized paradigm, cues were less directly "actionable" – either because cue presentation and responding were temporally dissociated[28,29] and/or because

choice sets involved multiple actions[9,10]. In contrast, in the present task the response window is the duration of cue presentation, and responses are limited to only a single Go (or NoGo) option. This tighter coupling between valence prospect and the concrete actions may create precisely the conditions in which NAcc is recruited for both Pavlovian and instrumental control[17,18], following ideas that NAcc will encode values specifically if these can be translated into "actionable" goals.

The fact that vmPFC shows cue-valence effects both here and in prior human fMRI studies[9,51], whereas NAcc cue-valence effects appear more variable and were robust in the present task, suggests a potential functional dissociation. One possibility is that vmPFC tracks expected value in a comparatively context-general manner, while NAcc value representations become most detectable when cue values are tightly coupled to a concrete, available action policy. Future work could test this interpretation more directly by experimentally varying how tightly cues and actions are coupled in time and choice space.

Amygdala involvement in novelty but not valence encoding

Unlike the NAcc, amygdala BOLD activity did not differentiate between Win and Avoid cues. Previous human studies have reported amygdala responses to cue valence[10,26], but in those paradigms cue valence had to be learned from experience before it could bias behavior. In the present task, cue valence was explicitly instructed, removing the need to acquire Pavlovian value from trial-by-trial feedback. This distinction matters: previous work in fear conditioning has indeed shown that the amygdala is especially engaged when Pavlovian value is learned or retrieved from experience, rather than when value is instructed[56]. Animal work supports this view as well. The amygdala, particularly the basolateral complex, is crucial for Pavlovian conditioning and for mediating the motivational impact of appetitive and aversive cues in transfer phases of Pavlovian-Instrumental transfer tasks, where conditioned cue-outcome associations modulate ongoing instrumental behavior[13,14,16,37]. In that light, the absence of valence-dependent amygdala activity at cue onset in the present study may reflect the fact that our contrast primarily indexes value signaling for instructed cues, rather than the acquisition or retrieval of Pavlovian value that typically recruits the amygdala.

While the amygdala was not involved in encoding cue *valence*, amygdala activity at cue *onset* was markedly elevated for novel compared to familiar cues. This effect was present irrespective of valence, and strongly exceeded both NAcc activity for novel cues, and amygdala responses on later trials. Strikingly, this enhancement was specific to the very first encounter with a cue and already diminished by the second presentation, suggesting that the amygdala uniquely boosted processing of

novel stimuli. In ecological terms, such heightened sensitivity to novelty is adaptive, as unfamiliar cues in natural environments may signal potential opportunities or threats and thus warrant additional processing[38,42]. The amygdala may thus amplify learning for novel cues by tagging them with affective significance at first encounter and relaying these value signals — via its projections to NAcc[13,14,34] — into downstream fronto-striatal circuits that shape subsequent action selection.

Nucleus accumbens and amygdala activation for desired outcomes

Outcome-locked activity speaks to how the brain processes feedback at the moment outcomes are revealed. Behaviorally, participants showed clear instrumental learning as their performance improved over repetitions even when the required response ran counter to their Pavlovian action tendencies. In the fMRI data, outcome-locked activity indexes how different regions represent the value of obtained outcomes (here as secondary reinforcers, i.e., points that translate to money). Such value coding is a necessary ingredient for learning from feedback but does not by itself demonstrate that a region is implementing a learning algorithm; a desired > non-desired outcome effect could reflect a value readout, a teaching signal (e.g., prediction-error updating), or a combination of both.

Both NAcc and amygdala showed stronger BOLD responses for desired than for non-desired outcomes. Earlier human fMRI studies using incentive and reinforcement-learning paradigms have repeatedly shown ventral striatal signals that track outcome value and reward prediction errors, supporting a role for this region in updating value estimates based on feedback[10,20,21,54,57]. Animal studies confirm a causal role for these striatal reward-related feedback signals in driving learning, as lesions or inactivation of NAcc impair the ability to adjust behavior when outcome contingencies change[13,15,46]. In line with prior work, the ventral striatal desired > non-desired outcome response in this study may index an outcome-value representation that supports subsequent action selection and incremental behavioral adjustment.

Amygdala outcome responses were also elevated for desired relative to non-desired outcomes, albeit weaker and less widespread than in the NAcc. This is consistent with the view that the amygdala, particularly the basolateral nucleus, encodes the value of stimulus–outcome associations and contributes to reinforcement learning beyond its classical role in threat conditioning[37,46]. Lesion work indicates that amygdala damage slows learning about which stimuli or actions lead to reward[14,46,58]. In that sense, our outcome-locked amygdala effect is compatible with a Pavlovian value-learning role, potentially supporting updates in cue value representations and their influence on behavior.

A similar outcome valence pattern was observed in the vmPFC, consistent with a large body of work linking vmPFC to representing expected and obtained value across decision-making and reinforcement-learning tasks[52,53,55]. The vmPFC and NAcc are often co-activated by rewarding outcomes, fitting models of tightly coupled cortico-striatal loops in which vmPFC carries an integrated value signal and NAcc links this signal to motivational drive and updating[59]. Within this circuitry, amygdala inputs can provide affective value information to both regions, particularly about stimulus-outcome associations[37].

Our data therefore fit a picture in which the same desired vs. non-desired outcome signature may support partly distinct computations across regions. In vmPFC, the effect may primarily reflect a relatively abstract value representation of obtained outcomes. In the amygdala, it may reflect Pavlovian value coding and updating of stimulus-outcome associations. In NAcc, it may reflect value signals that are especially well positioned to support stimulus-action-outcome learning, consistent with broader striatal accounts of instrumental reinforcement learning[20,21,54,59]. Speculatively, the relative strength of the NAcc effect may also be consistent with the task's demands: instrumental learning is required to perform well (actions must be learned and maintained under probabilistic feedback), whereas the Pavlovian value of cues is instructed and thus may not require de novo acquisition. Under this view, outcome signals in NAcc may be particularly engaged because they directly subserve the updating necessary for accurate responding, whereas amygdala signals may contribute more to value tagging and cue-related associations rather than dominating learning per se. More generally, prior work suggests complementary contributions in which amygdala updates can be fast and prominent under rapidly changing contingencies, while NAcc integrates feedback over longer timescales and may be better suited to stable but noisy environments[17,45,47,60].

Conclusion & Future directions

The present findings reveal both shared and dissociable contributions of NAcc and amygdala to motivated behavior. First, we show robust NAcc cue-valence encoding in an orthogonalized Go/NoGo paradigm – an effect absent in earlier human implementations[10,28,29]. This aligns with animal work demonstrating NAcc involvement in anticipatory value coding and reward-cue-driven action invigoration[14,16,19], and suggests that NAcc encodes cue value particularly when motivational cues are immediately actionable. Second, amygdala BOLD responses were elevated for novel cues with striking specificity: the effect was present on the very first cue encounter and had already dissipated by the second presentation, irrespective of valence. This near-instantaneous novelty-tagging points to the amygdala as a rapid salience detector

that flags new stimuli before any learning has occurred, plausibly boosting initial value updating via its direct projections to NAcc[13,14,34]. At outcome, NAcc, amygdala, and vmPFC all responded more strongly to desired than non-desired outcomes, pointing to a shared role in feedback processing. Whether the outcome-locked signals in NAcc and amygdala reflect the same underlying learning principles here remains an open question. More fundamentally, value coding – as the NAcc appears to do here – is necessary but not sufficient for causal involvement: a region may encode outcome value without driving behavioral change. The same logic applies to the amygdala's novelty response – its correlation with first-encounter processing does not establish that it is causally involved for enhanced early learning. Resolving these questions requires targeted causal intervention, which is the focus of **Chapter 4**.

METHODS

Participants

We preregistered (https://doi.org/10.17605/OSF.IO/Q9845) a target sample size of 31 healthy adults (17 females and 14 males; age range 19–35 years; mean = 24.8, SD = 4.1 years), based on a small to medium effect size[61] of d = 0.58 with an alpha level of 0.05 and a power of 80% (calculated using G*Power 3.1)[62]. Participants were screened on medical history to exclude putative participants with a history of brain surgery, serious head trauma, epilepsy, convulsion, or seizure, as well as the presence of implanted metal in the head or upper body, current diagnosis of neurological or psychiatric disorders, as well as recent substance use (i.e., consumption of more than four alcoholic units in the past 24 hours or use of recreational drugs within the past 48 hours). All participants provided written informed consent in accordance with the Declaration of Helsinki. The study has been independently reviewed by the Ethics Committee Social Sciences (ECSS) of the Radboud University (ECSW-2025-021), and there were no formal objections to this study.

Study overview

The present fMRI session formed part of a larger multi-session study investigating transcranial ultrasonic stimulation (TUS) effects on motivational decision-making; the full protocol of the intervention sessions is described in detail in **Chapter 4**. Here, we report the baseline MRI session and associated behavioral and BOLD data.

After screening for MRI and TUS safety, the researchers went over study and task description with the participants using infographics before participants entered

the scanner. They then completed a 1-hour MRI protocol (structural and functional sequences detailed below) and performed a timing-adjusted (with longer inter-stimulus delays and jittering) version of the Go/NoGo-task optimized to acquire BOLD data (see Motivational Go/NoGo-task for more details on the task). Outside the scanner, they practiced a single run of the faster-paced task optimized for behavioral testing to familiarize themselves with auditory feedback and shorter trial structure used in the following TUS sessions. Informed consent was obtained prior to the start of the experiment.

MRI data acquisition

MRI scanning was performed at the Donders Centre for Cognitive Neuroimaging using a 3 Tesla Magnetom Skyra Scanner (Siemens AG, Erlangen, Germany) equipped with a 32-channel head coil. High-resolution T1-weighted (T1w) scans were acquired (sagittal plane; repetition time (TR), 2700 ms; echo time (TE), 3.69 ms; flip angle, 9 degrees; voxel size, 0.9 x 0.9 x 0.9 mm; field of view, 230 mm) for co-registration with the functional data, acoustic/thermal simulations, and neuronavigation for TUS delivery during stimulation sessions. To capture detailed skull morphology and tissue properties for acoustic simulations of ultrasound an Ultra Short Echo Time (UTE) scan (sagittal plane, TR, 3.6 ms; TE, 0.07 ms; flip angle, 1 degree; voxel size, 0.8 x 0.8 x 0.8 mm; field of view, 230 mm) was acquired. Furthermore, a T2w scan was acquired (sagittal plane, TR, 3200 ms; TE, 408 ms; flip angle, T2 var flip angle mode; voxel size, 0.9 x 0.9 x 0.9 mm; field of view, 230 mm).

Functional MRI data for the 34-minute MGNG task (see Motivational Go/NoGo-task for task details) were acquired in two runs (674 volumes each) using a Multi-Band 4 (MB4) sequence (transversal plane, TR, 1500 ms; TE, 33.40 ms; flip angle, 75 deg; voxel size, 2.0 x 2.0 x 2.0 mm; field of view, 213 mm).

To correct for field inhomogeneities in functional imaging, a fieldmap sequence was collected before the BOLD sequence (transversal plane, TR, 567 ms; TE1, 4.30 ms; TE1, 6.67 ms; flip angle, T2 var flip angle mode; voxel size, 2.5 x 2.5 x 2.5 mm; field of view, 210 mm).

Task stimuli were presented on an MRI-compatible screen positioned at the rear of the scanner bore and viewed via a mirror mounted on the head coil. Stimulus presentation was synchronized to the MRI sequence using scanner triggers. Participants responded using MRI-compatible button boxes. Prior to each session, the MRI-compatible equipment (display, response devices) was tested, including verification of screen refresh rate (120 Hz) and correct button box configuration.

Motivational Go/NoGo-task

Participants completed a gamified version of the MGNG task[12,50]. The task was programmed using PsychoPy 2021.2.3.[63]. In the task, participants played as an astronaut whose ship approached successive "planet" stimuli, each distinguished by a unique symbol (**Figure 1A**). When the ship's hatch light illuminated the planet, participants had 1.3 s to decide whether to press the spacebar (Go) or withhold from pressing it (NoGo). They were instructed that their goal was to learn what the best response was for each planet for the most optimal outcome.

Following the response window, a jittered cue–feedback interval was presented before the outcome screen. This delay was uniformly distributed over a discrete set of 14 pre-generated durations (1.43–2.60 s in 0.09 s steps), created prior to each run and then assigned to trials in shuffled order. The outcome screen then displayed probabilistic feedback for 1.0 s, with outcomes delivered according to a predefined validity schedule (predominantly high-probability outcomes, with occasional low-probability outcomes). Planets appeared either in a blue "reward" galaxy or a red "punishment" galaxy, indicated by the spaceship light and planet color. On reward (blue) trials, correct responses yielded a blue diamond and +10 points, whereas incorrect responses yielded an empty diamond (0 points). On punishment (red) trials, correct responses prevented loss and displayed an outlined rock (0 points), whereas incorrect responses yielded a toxic rock and −10 points. Trials were separated by a jittered inter-trial interval, also uniformly distributed across 14 pre-generated durations (1.22–2.00 s in 0.06 s steps). To minimize cumulative timing drift during scanning, the task implemented online drift correction by slightly lengthening or shortening the ITI so that each trial ended at its precomputed scheduled time (i.e., maintaining alignment with the intended timing structure across the run).

Crossing required action (Go vs. NoGo) with outcome domain (Win vs. Avoid) yielded four conditions: Go2Win (G2W), NoGo2Win (NG2W), Go2Avoid (G2A), and NoGo2Avoid (NG2A; **Figure 1B**). Each session comprised 336 trials split into two runs of 168 trials (**Figure 1C**). Within each run, participants encountered 12 distinct cues ("planets"; three per condition), which appeared 12, 14, or 16 times, respectively. For cues presented 16 times, the optimal action reversed after the eighth presentation; participants were informed that reversals could occur but not when or how often.

Trial order was generated per run using constrained semi-randomization and was structured into three consecutive sets of 56 trials, each introducing a new subset of cues (**Figure 1C**). Importantly, set boundaries were not strict: cues from the preceding set could still occur at the start of the next set, reflecting overlap in the cue schedules

("bleeding" across sets). Sets were constructed to maintain response variety (i.e., at least two required actions were represented within each set), and reversals were distributed across the run such that each set contained at least one and at most two reversal cues. Trials were additionally organized in short groups (4-trial units) used for sequence generation/shuffling; however, these units did not enforce unique cues or conditions, such that brief repetitions of the same cue and/or condition could occur (as illustrated in **Figure 1C**). This design—frequent cue switching, staggered cue introductions, and distributed reversals—kept participants in an ongoing learning regime and reduced ceiling performance[50]. The symbols consisted of artificial characters selected from the Brussels Artificial Character Sets (BACS)[64].

Data analysis

Logistic and linear mixed-effects regression models

All behavioral analyses were performed in R (R Core Team) using mixed-effects regression models implemented in the lme4 package, and Type III Wald χ^2 tests from the car package for inference[65]. Unless stated otherwise, all analyses were conducted on a trial-by-trial basis and included all trials with a registered Go or NoGo response. Categorical predictors were coded using effect (sum-to-zero) coding, such that estimated main effects reflect differences averaged across the levels of the other factors. Required action contrasted cues for which a Go response was optimal versus cues for which a NoGo response was optimal (factor required action: Go vs. NoGo). Cue valence contrasted reward-predictive ("Win") versus punishment-predictive ("Avoid") cues (factor cue valence: Win vs. Avoid). To capture learning over repeated encounters with the same cue, we included a z-scored continuous predictor cue-specific trial number.

Logistic mixed-effects regressions on choice

To quantify instrumental learning and motivational biases in choice behavior, we first modelled the probability of emitting a Go response using logistic mixed-effects regression with a binomial link function (glmer, family = binomial). Our primary model included fixed effects of required action (Go vs. NoGo), cue valence (Win vs. Avoid), cue-specific trial number (z-scored), and all two-way interactions among these factors.

All choice models included a full random effects structure, i.e., subject-specific random intercepts and random slopes for all regressors and their interaction terms. Significance of fixed effects was assessed using Type III Wald χ^2 tests (car package), and 95% confidence intervals (CIs) for regression coefficients were obtained from the model's profile likelihood.

Linear mixed-effects regressions on response times

To determine whether the same factors that influenced choice behavior also shaped response speed, we analyzed reaction times (RTs) with linear mixed-effects regression (lmer). Analyses were by necessity restricted to trials with a recorded Go response within the response window. The structure of the regression was identical to the choice regression model. As for the choice models, statistical significance of fixed effects was evaluated using Type III Wald χ^2 tests, and 95%-CIs for regression coefficients were derived from the fitted linear mixed-effects regression models.

fMRI preprocessing

Functional MRI data were preprocessed using fMRIPrep 25.0.0[66], which is based on Nipype 1.9.2[67]. The pipeline combined head-motion correction, susceptibility distortion correction, co-registration to the anatomical image, and normalization to standard space in a single interpolation step.

For each participant, the T1-weighted (T1w) image was corrected for intensity non-uniformity and used as the anatomical reference. The T1w reference was skull-stripped and segmented into grey matter (GM), white matter (WM) and cerebrospinal fluid (CSF). Cortical surfaces were reconstructed with FreeSurfer, and T2-weighted images were integrated to refine pial surfaces where available. The T1w image was then nonlinearly registered to the MNI152NLin2009cAsym template using ANTs, and the resulting transforms were used to bring functional data into standard space.

For each of the two task runs, a BOLD reference was generated and rigidly aligned to the T1w reference. Head motion parameters (6 rigid-body parameters and the corresponding transformation matrices) were estimated with MCFLIRT (FSL). Fieldmaps were co-registered to the BOLD reference and used for susceptibility distortion correction of the functional images. All spatial transformations (motion correction, fieldmap correction, co-registration and normalization) were then combined, and the BOLD time series were resampled into MNI space in one interpolation step (cubic B-spline interpolation).

Several confound time series were derived for later use in the GLM: framewise displacement (FD), DVARS, and global signals from WM, CSF, and whole brain. Anatomical and temporal CompCor components (aCompCor/tCompCor) were computed after high-pass filtering (128 s cut-off); principal components were retained until they explained at least 50% of variance within the respective nuisance masks. Motion and global-signal regressors were expanded with temporal derivatives and

quadratic terms[68]. To further control for motion artefacts, we quantified FD for each volume and run and flagged volumes with FD > 0.9 mm. These high-motion volumes were later modelled with separate nuisance (scrubbing) regressors in the first-level GLM and thus excluded from contributing to the task regressors.

At the participant level, we summarized head motion by computing mean FD (in mm) and counting the number of scrubbed volumes across runs. Using an interquartile-range criterion on these two metrics, we identified participants with clearly outlying motion profiles (**Suppl. Figure S3**). Participants whose mean FD and/or number of scrubbed volumes fell in this outlier range were excluded from the fMRI BOLD analyses. Based on this procedure, two participants were removed, yielding a final fMRI sample of 29 healthy participants (15 females and 14 males; age range 20–35 years; mean = 25.1, SD = 4.1 years). Finally, the normalized preprocessed BOLD images were spatially smoothed with an 8 mm full-width at half-maximum (FWHM) Gaussian kernel using SPM12.

fMRI analyses

All fMRI analyses were performed in SPM12 using event-related general linear models (GLMs). For each participant and each of the two task runs, a first-level design matrix was specified with separate regressors for cue events, outcome events and nuisance regressors.

All task regressors were modelled as stick functions (duration 0 s) at event onset and convolved with the canonical haemodynamic response function. A 128 s high-pass filter was applied to remove low-frequency drifts. To reduce motion- and physiology-related artefacts, the following nuisance regressors were added to each first-level model: 24 motion parameters (the 6 rigid-body parameters, their temporal derivatives, and their squared terms; "Friston-24"), a scrubbing regressor flagging volumes with framewise displacement > 0.9 mm, the first six aCompCor components extracted from white matter and CSF masks, and run-specific intercepts. All task and nuisance regressors entered the design matrix simultaneously.

Main GLM: cue valence, action and outcome value

In the main GLM, cue-locked activity was modelled with four regressors corresponding to the cue valence – response combinations: G2W, NG2W, G2A and NG2A. Trials were sorted according to the action actually taken on each cue (i.e., a Go response to a Win cue contributed to the G2W regressor). Outcome-locked activity at feedback onset was modelled with two regressors reflecting outcome valence: Desired outcomes

(obtaining reward or avoiding loss) and Non-desired outcomes (loss or failing to obtain reward). For each participant, first-level contrast images were computed for: (i) cue valence (Win > Avoid and Avoid > Win); (ii) required action (Go > NoGo and NoGo > Go); (iii) outcome valence (Desired > Non-desired and Non-desired > Desired). Contrast images from the two runs were then entered into second-level random-effects analyses (one-sample t-tests across participants). Whole-brain results were thresholded at a cluster-defining voxel-wise threshold of $p < 0.001$ (uncorrected), and significant clusters were reported using cluster-level family-wise error (FWE) correction at $\alpha = 0.05$.

Trial-specific GLMs for novel vs. learned trials (cue and outcome)

To characterize NAcc and amygdala contributions change as cues progress from novel to familiar, additional trial-specific GLMs were estimated for both cue-locked and outcome-locked activity. These models had the same basic structure as the main GLM (identical cue and outcome regressors, plus the same nuisance regressors), but events were further split by stimulus repetition number. For the novel versus familiar analyses, cue- and outcome-events were each modelled with separate regressors for trial 1 (novel cues), trial 8 (familiar cues), and all remaining repetitions ("rest"). Trial 8 was chosen as the familiar time point because it is the last repetition available for all cues before one-third of the stimuli underwent a reversal in required action and thus provides a late learning stage with matched trial counts across cues. In a finer-grained set of GLMs, events were modelled with separate regressors for each repetition from trial 1 through trial 8, allowing trial-by-trial trajectories to be examined across early learning. For cue-locked GLMs, Win and Avoid cues were split by repetition (trial 1, trial 8, rest; or trial 1–8), while the executed response (Go vs. NoGo) was modelled in separate response regressors at cue onset, as in the main GLM. For outcome-locked GLMs, feedback events were modelled as Desired (reward or avoided loss) and Non-desired (loss or omitted reward), again split either into trial 1, trial 8 and rest, or into separate regressors for each repetition from trial 1 to trial 8. These trial-specific GLMs were used exclusively for ROI-based analyses of novelty/familiarity effects in NAcc and amygdala (see below); no whole-brain statistical inference was performed on these models.

Functional ROIs for NAcc and amygdala were defined from the outcome-valence contrast in the main GLM. For each participant, the Desired > Non-desired outcome contrast was computed at the first level and entered into a group-level one-sample t-test. The resulting statistical map was thresholded at $p < .001$ (uncorrected), and suprathreshold voxels were intersected with bilateral anatomical NAcc and amygdala masks in MNI space. Anatomical masks were derived from individual FreeSurfer segmentations, normalized to MNI space and combined across participants. This

procedure yielded binarized functional–anatomical ROIs for NAcc and amygdala. Within these ROIs, beta estimates were extracted separately for each participant, run, hemisphere (left/right), and regressor of interest from the trial-specific GLMs. For each ROI–hemisphere, the peak voxel was identified as the voxel with the maximum beta value within the corresponding mask. Around this peak coordinate, an 8-mm-radius sphere was defined in MNI space, and the mean beta value was computed within the intersection of this sphere and the ROI hemisphere mask (sphere ∩ ROI-hemi), ensuring that summary estimates remained confined to the functionally defined region.

Statistical analyses on ROI beta estimates were performed with linear mixed-effects models including a random intercept for subject and appropriate fixed-effect factors, and evaluated using likelihood-ratio χ^2 tests between nested models.

For the novelty analysis at cue onset, cue-locked betas were first averaged across Win and Avoid cues to obtain a valence-independent measure of cue responsivity. For trial 1 and trial 8, this yielded one beta per ROI, hemisphere and trial level. These values were entered into a mixed-effects regression model with fixed factors region (NAcc, amygdala), trial (1, 8) and their interaction (region × trial) to test whether novelty sensitivity differed between NAcc and amygdala. To assess possible valence effects, a second model was fit to non-averaged betas, with fixed factors region (NAcc, amygdala), trial (1, 8), valence (Win, Avoid) and their interactions, focusing on the region × trial × valence interaction to test whether novelty-related differences depended on cue valence. To determine whether novelty-related responses declined gradually or were specific to the first presentation, cue-locked betas from the 8-level GLM (trials 1–8) were analyzed in two competing models: (i) a model including region (NAcc, amygdala), trial (1–8, treated as an ordinal variable) and their interaction (region × trial(1–8)), capturing a gradual change over repetitions; and (ii) a model including region and a binary trial factor (trial 1 vs. trials 2–8) and their interaction (region × [trial1 vs. rest]), capturing a one-off novelty peak at the first presentation. Models were compared using Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC), with lower values indicating better fit.

For outcome-locked analyses, trial-specific outcome-valence effects were quantified as the difference between beta estimates for Desired and Non-desired outcomes (Desired − Non-desired) for each trial level and ROI. For the trial 1 versus trial 8 comparison, these difference scores were entered into a mixed-effects regression model with fixed factors region (NAcc, amygdala), Trial (1, 8) and their interaction (region × trial) to test whether outcome-valence coding for novel versus familiar cues differed across regions. Finally, to examine how outcome-valence coding evolved over

repetitions, Desired – Non-desired difference betas from trials 1–8 were analyzed in a mixed-effects regression model including region (NAcc, amygdala), Trial (1–8, ordinal) and their interaction (region × trial(1–8)), testing for region-specific changes in outcome-related activity across early learning.

REFERENCES

1. Boureau, Y.-L., Sokol-Hessner, P. & Daw, N. D. Deciding How To Decide: Self-Control and Meta-Decision Making. *Trends Cogn. Sci.* ***19***, 700–710 (2015).
2. Dayan, P., Niv, Y., Seymour, B. & Daw, N. D. The misbehavior of value and the discipline of the will. *Neural Netw. Off. J. Int. Neural Netw. Soc.* ***19***, 1153–1160 (2006).
3. Moutoussis, M. *et al.* Change, stability, and instability in the Pavlovian guidance of behaviour from adolescence to young adulthood. *PLOS Comput. Biol.* ***14***, e1006679 (2018).
4. Garbusow, M. *et al.* Pavlovian-to-instrumental transfer effects in the nucleus accumbens relate to relapse in alcohol dependence. *Addict. Biol.* ***21***, 719–731 (2016).
5. Huys, Q. J. M. *et al.* The specificity of Pavlovian regulation is associated with recovery from depression. *Psychol. Med.* ***46***, 1027–1035 (2016).
6. Milner, A., Witt, K., LaMontagne, A. D. & Niedhammer, I. Psychosocial job stressors and suicidality: a meta-analysis and systematic review. *Occup. Environ. Med.* ***75***, 245–253 (2018).
7. Mkrtchian, A., Aylward, J., Dayan, P., Roiser, J. P. & Robinson, O. J. Modeling Avoidance in Mood and Anxiety Disorders Using Reinforcement Learning. *Biol. Psychiatry* ***82***, 532–539 (2017).
8. Schad, D. J. *et al.* Dissociating neural learning signals in human sign- and goal-trackers. *Nat. Hum. Behav.* ***4***, 201–214 (2020).
9. Algermissen, J., Swart, J. C., Scheeringa, R., Cools, R. & Den Ouden, H. E. M. Striatal BOLD and midfrontal theta power express motivation for action. *Cereb. Cortex* ***32***, 2924–2942 (2022).
10. Algermissen, J., Swart, J. C., Scheeringa, R., Cools, R. & den Ouden, H. E. M. Prefrontal signals precede striatal signals for biased credit assignment in motivational learning biases. *Nat. Commun.* ***15***, 19 (2024).
11. Cools, R., Nakamura, K. & Daw, N. D. Serotonin and dopamine: unifying affective, activational, and decision functions. *Neuropsychopharmacol. Off. Publ. Am. Coll. Neuropsychopharmacol.* ***36***, 98–113 (2011).
12. Scholz, V. *et al.* Cortical dopamine reduces the impact of motivational biases governing automated behaviour. *Neuropsychopharmacol. Off. Publ. Am. Coll. Neuropsychopharmacol.* ***47***, 1503–1512 (2022).
13. Cardinal, R. N., Parkinson, J. A., Hall, J. & Everitt, B. J. Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neurosci. Biobehav. Rev.* ***26***, 321–352 (2002).
14. Corbit, L. H. & Balleine, B. W. Double dissociation of basolateral and central amygdala lesions on the general and outcome-specific forms of pavlovian-instrumental transfer. *J. Neurosci. Off. J. Soc. Neurosci.* ***25***, 962–970 (2005).

15. Corbit, L. H. & Balleine, B. W. The general and outcome-specific forms of Pavlovian-instrumental transfer are differentially mediated by the nucleus accumbens core and shell. *J. Neurosci. Off. J. Soc. Neurosci.* ***31***, 11786–11794 (2011).
16. Hall, J., Parkinson, J. A., Connor, T. M., Dickinson, A. & Everitt, B. J. Involvement of the central nucleus of the amygdala and nucleus accumbens core in mediating Pavlovian influences on instrumental behaviour. *Eur. J. Neurosci.* ***13***, 1984–1992 (2001).
17. Hazy, T. E., Frank, M. J. & O'Reilly, R. C. Neural mechanisms of acquired phasic dopamine responses in learning. *Neurosci. Biobehav. Rev.* ***34***, 701–720 (2010).
18. O'Reilly, R. C. & Frank, M. J. Making Working Memory Work: A Computational Model of Learning in the Prefrontal Cortex and Basal Ganglia. *Neural Comput.* ***18***, 283–328 (2006).
19. Wyvell, C. L. & Berridge, K. C. Intra-accumbens amphetamine increases the conditioned incentive salience of sucrose reward: enhancement of reward 'wanting' without enhanced 'liking' or response reinforcement. *J. Neurosci. Off. J. Soc. Neurosci.* ***20***, 8122–8130 (2000).
20. Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. Model-based influences on humans' choices and striatal prediction errors. *Neuron* ***69***, 1204–1215 (2011).
21. O'Doherty, J. *et al.* Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* ***304***, 452–454 (2004).
22. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* ***275***, 1593–1599 (1997).
23. Steinberg, E. E. *et al.* A causal link between prediction errors, dopamine neurons and learning. *Nat. Neurosci.* ***16***, 966–973 (2013).
24. Knutson, B., Adams, C. M., Fong, G. W. & Hommer, D. Anticipation of Increasing Monetary Reward Selectively Recruits Nucleus Accumbens. *J. Neurosci.* ***21***, RC159–RC159 (2001).
25. Sescousse, G., Li, Y. & Dreher, J.-C. A common currency for the computation of motivational values in the human striatum. *Soc. Cogn. Affect. Neurosci.* ***10***, 467–473 (2015).
26. Geurts, D. E. M., Huys, Q. J. M., den Ouden, H. E. M. & Cools, R. Aversive Pavlovian Control of Instrumental Behavior in Humans. *J. Cogn. Neurosci.* ***25***, 1428–1441 (2013).
27. Talmi, D., Seymour, B., Dayan, P. & Dolan, R. J. Human Pavlovian–Instrumental Transfer. *J. Neurosci.* ***28***, 360–368 (2008).
28. Guitart-Masip, M. *et al.* Action Dominates Valence in Anticipatory Representations in the Human Striatum and Dopaminergic Midbrain. *J. Neurosci.* ***31***, 7867–7875 (2011).
29. Guitart-Masip, M. *et al.* Go and no-go learning in reward and punishment: interactions between affect and effect. *NeuroImage* ***62***, 154–166 (2012).
30. Ebrahimi, C. *et al.* Elevated Amygdala Responses During De Novo Pavlovian Conditioning in Alcohol Use Disorder Are Associated With Pavlovian-to-Instrumental Transfer and Relapse Latency. *Biol. Psychiatry Glob. Open Sci.* ***3***, 803–813 (2023).
31. Hatfield, T., Han, J.-S., Conley, M., Gallagher, M. & Holland, P. Neurotoxic Lesions of Basolateral, But Not Central, Amygdala Interfere with Pavlovian Second-Order Conditioning and Reinforcer Devaluation Effects. *J. Neurosci.* ***16***, 5256–5265 (1996).
32. LeDoux, J. E. Emotion circuits in the brain. *Annu. Rev. Neurosci.* ***23***, 155–184 (2000).

33. Phelps, E. A. & LeDoux, J. E. Contributions of the amygdala to emotion processing: from animal models to human behavior. *Neuron* ***48***, 175–187 (2005).
34. Holland, P. C. & Gallagher, M. Double dissociation of the effects of lesions of basolateral and central amygdala on conditioned stimulus-potentiated feeding and Pavlovian-instrumental transfer. *Eur. J. Neurosci.* ***17***, 1680–1694 (2003).
35. Prévost, C., Liljeholm, M., Tyszka, J. M. & O'Doherty, J. P. Neural correlates of specific and general Pavlovian-to-Instrumental Transfer within human amygdalar subregions: a high-resolution fMRI study. *J. Neurosci. Off. J. Soc. Neurosci.* ***32***, 8383–8390 (2012).
36. Balderston, N. L., Schultz, D. H. & Helmstetter, F. J. The Effect of Threat on Novelty Evoked Amygdala Responses. *PLOS ONE* ***8***, e63220 (2013).
37. Balleine, B. W. & Killcross, S. Parallel incentive processing: an integrated view of amygdala function. *Trends Neurosci.* ***29***, 272–279 (2006).
38. Blackford, J. U., Buckholtz, J. W., Avery, S. N. & Zald, D. H. A unique role for the human amygdala in novelty detection. *NeuroImage* ***50***, 1188–1193 (2010).
39. Janak, P. H. & Tye, K. M. From circuits to behaviour in the amygdala. *Nature* ***517***, 284–292 (2015).
40. Likhtik, E. & Paz, R. Amygdala-prefrontal interactions in (mal)adaptive learning. *Trends Neurosci.* ***38***, 158–166 (2015).
41. Murray, R. J., Brosch, T. & Sander, D. The functional profile of the human amygdala in affective processing: insights from intracranial recordings. *Cortex J. Devoted Study Nerv. Syst. Behav.* ***60***, 10–33 (2014).
42. Pedersen, W. S. *et al.* The effects of stimulus novelty and negativity on BOLD activity in the amygdala, hippocampus, and bed nucleus of the stria terminalis. *Soc. Cogn. Affect. Neurosci.* ***12***, 748–757 (2017).
43. Belova, M. A., Paton, J. J., Morrison, S. E. & Salzman, C. D. Expectation modulates neural responses to pleasant and aversive stimuli in primate amygdala. *Neuron* ***55***, 970–984 (2007).
44. Paton, J. J., Belova, M. A., Morrison, S. E. & Salzman, C. D. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* ***439***, 865–870 (2006).
45. Averbeck, B. B. & Costa, V. D. Motivational neural circuits underlying reinforcement learning. *Nat. Neurosci.* ***20***, 505–512 (2017).
46. Costa, V. D., Dal Monte, O., Lucas, D. R., Murray, E. A. & Averbeck, B. B. Amygdala and Ventral Striatum Make Distinct Contributions to Reinforcement Learning. *Neuron* ***92***, 505–517 (2016).
47. Neftci, E. O. & Averbeck, B. B. Reinforcement learning in artificial and biological systems. *Nat. Mach. Intell.* ***1***, 133–143 (2019).
48. Guitart-Masip, M. *et al.* Action controls dopaminergic enhancement of reward representations. *Proc. Natl. Acad. Sci. U. S. A.* ***109***, 7511–7516 (2012).

49. Swart, J. C. *et al.* Catecholaminergic challenge uncovers distinct Pavlovian and instrumental mechanisms of motivated (in)action. *eLife* ***6***, e22169 (2017).
50. Zorowitz, S., Karni, G., Paredes, N., Daw, N. & Niv, Y. Improving the Reliability of the Pavlovian Go/No-Go Task for Computational Psychiatry Research. *Comput. Psychiatry* ***9***, (2025).
51. Queirazza, F., Steele, J. D., Krishnadas, R., Cavanagh, J. & Philiastides, M. G. Functional Magnetic Resonance Imaging Signatures of Pavlovian and Instrumental Valuation Systems during a Modified Orthogonalized Go/No-go Task. *J. Cogn. Neurosci.* ***35***, 2089–2109 (2023).
52. Levy, D. J. & Glimcher, P. W. The root of all value: a neural common currency for choice. *Curr. Opin. Neurobiol.* ***22***, 1027–1038 (2012).
53. Rangel, A., Camerer, C. & Montague, P. R. Neuroeconomics: The neurobiology of value-based decision-making. *Nat. Rev. Neurosci.* ***9***, 545–556 (2008).
54. Bartra, O., McGuire, J. T. & Kable, J. W. The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage* ***76***, 412–427 (2013).
55. Klein-Flügge, M. C., Bongioanni, A. & Rushworth, M. F. S. Medial and orbital frontal cortex in decision-making and flexible behavior. *Neuron* ***110***, 2743–2770 (2022).
56. Atlas, L. Y., Doll, B. B., Li, J., Daw, N. D. & Phelps, E. A. Instructed knowledge shapes feedback-driven aversive learning in striatum and orbitofrontal cortex, but not the amygdala. *eLife* ***5***, e15192 (2016).
57. Sescousse, G., Caldú, X., Segura, B. & Dreher, J.-C. Processing of primary and secondary rewards: a quantitative meta-analysis and review of human functional neuroimaging studies. *Neurosci. Biobehav. Rev.* ***37***, 681–696 (2013).
58. Blundell, P., Hall, G. & Killcross, S. Lesions of the basolateral amygdala disrupt selective aspects of reinforcer representation in rats. *J. Neurosci. Off. J. Soc. Neurosci.* ***21***, 9018–9026 (2001).
59. Haber, S. N. & Knutson, B. The Reward Circuit: Linking Primate Anatomy and Human Imaging. *Neuropsychopharmacology* ***35***, 4–26 (2010).
60. Averbeck, B. B. Amygdala and ventral striatum population codes implement multiple learning rates for reinforcement learning. in *2017 IEEE Symposium Series on Computational Intelligence (SSCI)* 1–5 (2017). doi:10.1109/SSCI.2017.8285354.
61. Yaakub, S. N. *et al.* Non-invasive ultrasonic neuromodulation of the human nucleus accumbens impacts reward sensitivity. *Nat. Commun.* ***16***, 10192 (2025).
62. Faul, F., Erdfelder, E., Buchner, A. & Lang, A.-G. Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behav. Res. Methods* ***41***, 1149–1160 (2009).
63. Peirce, J. *et al.* PsychoPy2: Experiments in behavior made easy. *Behav. Res. Methods* ***51***, 195–203 (2019).
64. Vidal, C., Content, A. & Chetail, F. BACS: The Brussels Artificial Character Sets for studies in cognitive psychology and neuroscience. *Behav. Res. Methods* ***49***, 2093–2112 (2017).

65. Bates, D., Mächler, M., Bolker, B. & Walker, S. Fitting Linear Mixed-Effects Models Using lme4. *J. Stat. Softw.* ***67***, 1–48 (2015).
66. Esteban, O. *et al.* fMRIPrep: a robust preprocessing pipeline for functional MRI. *Nat. Methods* ***16***, 111–116 (2019).
67. Gorgolewski, K. *et al.* Nipype: A Flexible, Lightweight and Extensible Neuroimaging Data Processing Framework in Python. *Front. Neuroinformatics* ***5***, (2011).
68. Satterthwaite, T. D. *et al.* An improved framework for confound regression and filtering for control of motion artifact in the preprocessing of resting-state functional connectivity data. *NeuroImage* ***64***, 240–256 (2013).
69. Rushworth, M. F., Kolling, N., Sallet, J. & Mars, R. B. Valuation and decision-making in frontal cortex: one or many serial or parallel systems? *Curr. Opin. Neurobiol.* ***22***, 946–955 (2012).
70. Kolling, N., Behrens, T., Wittmann, M. K. & Rushworth, M. Multiple signals in anterior cingulate cortex. *Curr. Opin. Neurobiol.* ***37***, 36–43 (2016).
71. Shenhav, A., Botvinick, M. M. & Cohen, J. D. The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron* ***79***, 217–240 (2013).
72. Ullsperger, M., Fischer, A. G., Nigbur, R. & Endrass, T. Neural mechanisms and temporal dynamics of performance monitoring. *Trends Cogn. Sci.* ***18***, 259–267 (2014).
73. Aron, A. R., Robbins, T. W. & Poldrack, R. A. Inhibition and the right inferior frontal cortex. *Trends Cogn. Sci.* ***8***, 170–177 (2004).
74. Aron, A. R., Robbins, T. W. & Poldrack, R. A. Inhibition and the right inferior frontal cortex: one decade on. *Trends Cogn. Sci.* ***18***, 177–185 (2014).
75. Holroyd, C. B. & Coles, M. G. H. The neural basis of human error processing: reinforcement learning, dopamine, and the error-related negativity. *Psychol. Rev.* ***109***, 679–709 (2002).
76. Rushworth, M. F. S., Walton, M. E., Kennerley, S. W. & Bannerman, D. M. Action sets and decisions in the medial frontal cortex. *Trends Cogn. Sci.* ***8***, 410–417 (2004).
77. Sallet, J. *et al.* Expectations, gains, and losses in the anterior cingulate cortex. *Cogn. Affect. Behav. Neurosci.* ***7***, 327–336 (2007).
78. Corbetta, M. & Shulman, G. L. Control of goal-directed and stimulus-driven attention in the brain. *Nat. Rev. Neurosci.* ***3***, 201–215 (2002).
79. Menon, V. & Uddin, L. Q. Saliency, switching, attention and control: a network model of insula function. *Brain Struct. Funct.* ***214***, 655–667 (2010).
80. Simmonds, D. J., Pekar, J. J. & Mostofsky, S. H. Meta-analysis of Go/No-go tasks demonstrating that fMRI activation associated with response inhibition is task-dependent. *Neuropsychologia* ***46***, 224–232 (2008).
81. Zhang, R., Geng, X. & Lee, T. M. C. Large-scale functional neural network correlates of response inhibition: an fMRI meta-analysis. *Brain Struct. Funct.* ***222***, 3973–3990 (2017).
82. Swart, J. C. *et al.* Frontal network dynamics reflect neurocomputational mechanisms for reducing maladaptive biases in motivated action. *PLOS Biol.* ***16***, e2005979 (2018).

SUPPLEMENTARY INFORMATION

SUPPLEMENTARY SECTIONS

S.1 Exploratory and control whole-brain analyses

Whole-brain cue-locked Avoid > Win contrast comparison revealed significantly stronger BOLD responses in dorsomedial prefrontal cortex (dmPFC), close to anterior cingulate cortex (ACC; **Figure S1A**; **Table S1**). The ACC/dmPFC region is often described as the interface between valuation and motor systems, integrating information about past outcomes and current choice values to guide future action selection, moving closer to dmPFC when concrete action alternatives (rather than more abstract values) must be evaluated[69,70]. Increased dmPFC/ACC activity for punishment-predictive cues may therefore reflect enhanced control-related processing when action and inaction carry different future consequences[71,72].

At outcome, a whole-brain Non-desired > Desired contrast revealed increased activation in right-lateral and medial frontal regions, including right inferior frontal gyrus, middle and superior frontal gyri, and paracingulate gyrus (**Figure S1B**; **Table S2**). These areas are frequently implicated in cognitive control, response inhibition and performance monitoring, and have been linked to implementing adjustments after errors, losses or unexpected negative feedback[73–77]. The predominance of right-hemispheric activation is consistent with broader evidence for right-lateralized networks involved in inhibitory control and salience detection[78,79]. In the context of the present task, enhanced right frontal responses to non-desired outcomes likely reflect increased control and adjustment demands when feedback indicates that recent actions failed to secure reward or avoid loss.

Finally, as a control analysis, a whole-brain Go > NoGo contrast (collapsed across valence) at the time of response showed expected strong activation in left primary motor cortex, supplementary motor area and associated motor regions (**Figure S1C**; **Table S3**), consistent with right-handed button presses. Conversely, the NoGo > Go contrast revealed bilateral increased activation in dmPFC, compatible with enhanced control- and inhibition-related engagement during successful response withholding[80,81]. Together, these contrasts confirm that our data present canonical motor and inhibitory-control networks.

S.2 Whole-brain BOLD responses for motivational congruency

We examined whether cue-locked BOLD responses differed between motivationally congruent (G2W, NG2A) and incongruent (G2A, NG2W) trials. No significant clusters were observed in prefrontal or midfrontal regions at whole-brain corrected thresholds (**Figure S2**).

The absence of a congruency effect in BOLD is perhaps not surprising given a structural feature of our paradigm. Participants in our task were continuously introduced to novel cues throughout the session, keeping them in an ongoing learning regime. Motivational conflict — in the sense that Swart and colleagues (2018) operationalize it — requires that correct instrumental responses are learned well enough to sometimes stand in opposition to Pavlovian action tendencies[82]. Early in a cue's history, participants have not yet reliably learned its valence or required action; a G2A cue is therefore not yet experienced as conflicting, because the participant does not yet know that this cue calls for suppression of an approach tendency. As a result, the proportion of trials on which participants genuinely experienced motivational conflict may have been relatively low in our paradigm, reducing the magnitude and detectability of any associated prefrontal signal. This is in line with Algermissen and colleagues (2022), who also found no whole-brain significant clusters for the congruency contrast in a MGNG task with simultaneous EEG-fMRI[9]. In their case, however, the likely explanation was poor task performance overall, which similarly limited participants' experience of genuine conflict.

SUPPLEMENTARY FIGURES

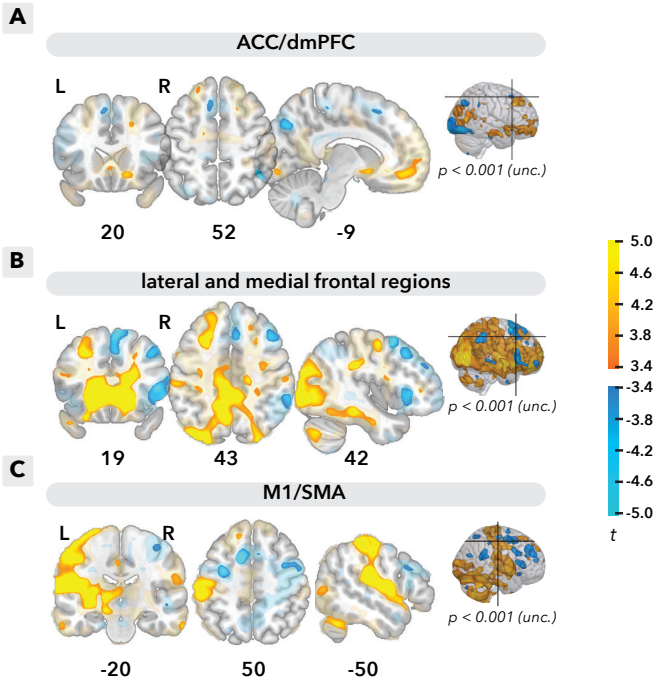


Figure S1 | Exploratory and control whole-brain BOLD effects. (A) *Cue valence.* Whole-brain results for the Avoid versus Win cue contrast (cool colors), showing stronger responses to loss-predictive than reward-predictive cues in dorsal anterior cingulate / dorsomedial prefrontal cortex (ACC/dmPFC). **(B)** *Outcome valence.* Whole-brain results for the Non-desired versus Desired outcome contrast (cool colors), revealing increased activation for negative relative to positive feedback in right-lateral and medial frontal regions, including inferior, middle and superior frontal gyri and paracingulate cortex. **(C)** *Action execution.* Whole-brain results for the Go versus NoGo contrast at response onset (warm colors), showing robust activation in left primary motor cortex and supplementary motor area (M1/SMA), consistent with right-hand button presses; cool colors indicate NoGo > Go and highlight dmPFC/ACC engagement during successful response withholding. **(A)-(C)** Maps are thresholded at voxel-wise $p < .001$ uncorrected ($T \geq 3.4$; opaque colors) with subthreshold trends at $p < .10$ shown transparently. Images are displayed on the MNI152 template; sagittal renderings with crosshairs indicate slice locations.

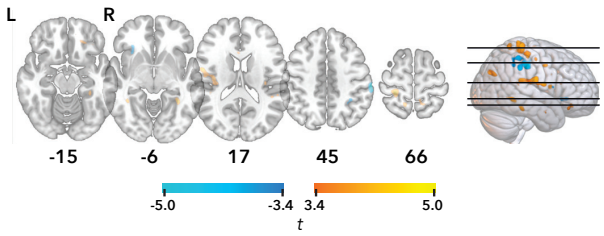


Figure S2 | Whole-brain BOLD effects for motivational congruency. Cue-locked whole-brain results for the incongruent versus congruent contrast (cool colors; G2A and NG2W > G2W and NG2A) and the congruent versus incongruent contrast (warm colors), displayed across axial slices (z = −15, −6, 17, 45, 66) and a sagittal surface rendering. No significant clusters were observed in prefrontal or midfrontal regions for either contrast. Maps are thresholded at T ≥ 3.4; images are displayed on the MNI152 template; sagittal rendering with crosshairs indicates slice locations.

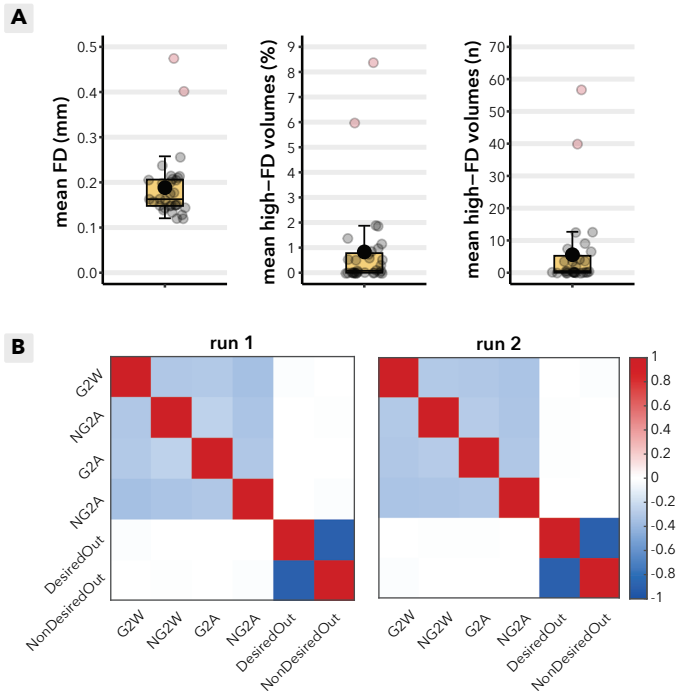


Figure S3 | (A) *Summary of head motion across participants.* From left to right: mean framewise displacement (FD, in mm), proportion of high-FD volumes (%; FD > 0.9 mm), and absolute number of high-FD volumes per participant. Boxplots indicate the group interquartile range (IQR) with the median (horizontal line); the large black circle marks the group mean, and grey circles represent individual participants. Participants falling outside the IQR range on mean FD and/or number of scrubbed volumes (red circles) were classified as motion outliers and excluded from fMRI BOLD analyses (see Methods). **(B)** *Correlation matrices of first-level GLM regressors, averaged across participants, for run 1 (left) and run 2 (right).* Cells show Pearson correlation coefficients between the main task regressors (G2W, NG2W, G2A, NG2A, outcome regressors).

SUPPLEMENTARY TABLES

Table S1 | Cue-locked effects of cue valence (Win vs. Avoid). Significant clusters from whole-brain one-sample *t*-tests on the cue-locked contrasts Win > Avoid and Avoid > Win. For each cluster, the main anatomical region, cluster-level family-wise error–corrected ($p_{FWE\text{-}corr}$) and uncorrected (p_{uncorr}) *p*-values, peak *t*-value, peak-level $p_{FWE\text{-}corr}$ and p_{uncorr}, and MNI peak coordinates (x, y, z) are reported. Clusters are thresholded at voxel-wise $p < .001$ uncorrected and cluster-level *p*FWE-corr < .05. vmPFC = ventromedial prefrontal cortex; NAcc = ventral striatum (nucleus accumbens); ACC/dmPFC = anterior cingulate / dorsomedial prefrontal cortex.

Contrast	Brain region	Cluster-level $p_{FWE\text{-}corr}$	Cluster-level p_{uncorr}	T	Peak-level $p_{FWE\text{-}corr}$	Peak-level p_{uncorr}	Peak coordinates x	y	z
Cue valence (win > avoid)	Visual cortex	0.000	0.000	5.62	0.13	0.000	2	-81	-2
	Visual cortex	0.000	0.000	5.54	0.15	0.000	-24	-81	18
	Visual cortex	0.002	0.000	5.03	0.4	0.000	54	-69	6
	vmPFC	0.000	0.000	5.54	0.16	0.000	4	36	-10
	Right NAcc	0.000	0.000	4.90	0.5	0.000	8	9	-8
	Left NAcc	0.000	0.000	4.72	0.6	0.000	-8	9	-8
	Temporal lobe	0.003	0.000	5.02	0.4	0.000	56	-19	-22
	Left middle frontal gyrus	0.019	0.001	4.67	0.7	0.000	-16	30	42
Cue valence (avoid > win)	Visual cortex (V1)	0.000	0.000	8.88	0.000	0.000	-14	-101	-12
	ACC/dmPFC	0.8	0.111	4.04	0.9	0.000	-8	22	52

Table S2 | Outcome-locked effects of outcome valence (Desired vs. Non-desired). Significant clusters from whole-brain one-sample *t*-tests on the outcome-locked contrasts Desired > Non-desired and Non-desired > Desired. For each cluster, the main anatomical region, cluster-level family-wise error–corrected ($p_{FWE\text{-}corr}$) and uncorrected (p_{uncorr}) *p*-values, peak *t*-value, peak-level $p_{FWE\text{-}corr}$ and p_{uncorr}, and MNI peak coordinates (x, y, z) are reported. Clusters are thresholded at voxel-wise $p < .001$ uncorrected and cluster-level *p*FWE-corr < .05. vmPFC = ventromedial prefrontal cortex; NAcc = ventral striatum (nucleus accumbens).

Contrast	Brain region	Cluster-level $p_{FWE\text{-}corr}$	Cluster-level p_{uncorr}	T	Peak-level $p_{FWE\text{-}corr}$	Peak-level p_{uncorr}	Peak coordinates x	y	z
Outcome (desired > non-desired)	vmPFC	0.000	0.000	15.15	0.000	0.000	6	48	-16
	Left NAcc	0.000	0.000	13.07	0.000	0.000	-16	18	-2
	Right NAcc	0.000	0.000	11.87	0.000	0.000	13	24	-4
	Left Putamen	0.000	0.000	10.88	0.000	0.000	-22	7	-12
	Right putamen	0.000	0.000	12.52	0.000	0.000	23	15	-10
	Right amygdala	0.000	0.000	10.51	0.000	0.000	29	-9	-10
	Left amygdala	0.000	0.000	9.78	0.000	0.000	-28	-11	-10
	Right superior temporal sulcus	0.023	0.000	5.67	0.09	0.000	64	-1	-10
Outcome (non-desired > desired)	Right superior frontal gyrus	0.000	0.000	7.00	0.005	0.000	8	11	66
	Right Paracingulate gyrus	0.000	0.000	4.85	0.4	0.000	4	15	48
	Right middle frontal gyrus	0.020	0.002	5.65	0.09	0.000	41	7	54
	Right inferior frontal gyrus	0.001	0.000	5.00	0.3	0.000	47	18	8
	Right supramarginal gyrus	0.005	0.000	4.82	0.4	0.000	56	-40	38

Table S3 | Action-related effects (Go vs. NoGo). Significant clusters from whole-brain one-sample *t*-tests on the action contrasts Go > NoGo and NoGo > Go. For each cluster, the main anatomical region, cluster-level family-wise error–corrected ($p_{FWE\text{-}corr}$) and uncorrected (p_{uncorr}) *p*-values, peak *t*-value, peak-level $p_{FWE\text{-}corr}$ and p_{uncorr}, and MNI peak coordinates (x, y, z) are reported. Clusters are thresholded at voxel-wise $p < .001$ uncorrected and cluster-level *p*FWE-corr < .05. M1 = primary motor cortex; dACC = dorsal anterior cingulate cortex; dmPFC = dorsomedial prefrontal cortex.

Contrast	Brain region	Cluster-level $p_{FWE\text{-}corr}$	Cluster-level p_{uncorr}	T	Peak-level $p_{FWE\text{-}corr}$	Peak-level p_{uncorr}	Peak coordinates x	y	z
Action (go > nogo)	Left Putamen	0.000	0.000	9.05	0.000	0.000	-37	-1	12
	Left M1	0.000	0.000	8.05	0.000	0.000	-63	-17	18
	Right Putamen	0.000	0.000	6.20	0.033	0.000	25	1	-8
	Visual cortex	0.003	0.000	6.24	0.029	0.000	23	-93	8
	dACC	0.001	0.000	5.48	0.15	0.000	2	-11	36
Action (nogo > go)	Left dmPFC	0.022	0.002	8.14	0.000	0.000	-14	7	50
	Right superior temporal sulcus	0.001	0.000	6.01	0.049	0.000	43	-46	26
	Right anterior insula	0.008	0.001	5.76	0.08	0.000	31	28	8
	Left middle frontal gyrus	0.032	0.003	5.38	0.18	0.000	-38	-5	52
	Right middle frontal gyrus	0.002	0.000	4.83	0.5	0.000	27	46	26




PART IV
Interlude
Chapter 4

ARASH THE ARCHER

Legend has it that in the old Persian realm, Iran had been locked for years in war with Turan. In the mythic geography of the Shahnameh, Iran is the heartland of civilization. It is a kingdom of order, justice, and divine favor, ruled by a lineage of kings stretching back to the dawn of time. Turan, to the north and east, is its shadow: a vast realm of nomadic warriors, fierce and restless, whose lands stretched across what we would now call Central Asia – from the steppes beyond the Oxus River to the edges of the Caucasus.

The enmity between Iran and Turan is ancient, rooted in a family wound. Both peoples descend from the same king: **Fereydun**, who divided his realm among his three sons. To the eldest, **Salm**, he gave the West. The middle son, **Tur**, received Turan: the northern lands, harsh and proud. To the youngest, **Iraj**, he gave Iran – the jewel, the center. But Tur and Salm grew jealous of their brother's inheritance. Together, they murdered Iraj. And from that act of fratricide, a blood feud was born that would echo through generations; war after war, hero after hero, each cycle of violence answered by the next.

This conflict runs through the Shahnameh like a wound that never fully heals. Even Rostam, the great warrior whose story opened this dissertation, would one day unknowingly kill his own son Sohrab on a Turanian battlefield – a tragedy born directly from the endless war between the two realms. The enmity between Iran and Turan shaped the fates of nearly every hero in the Persian epic tradition.

At this moment in the story, the war was going badly for the Persians. Turan was stronger; territory was lost. And then, to deepen the humiliation, the Turanians offered a bargain that sounded like mockery disguised as mercy: If one of your archers can loose a single arrow, you may keep the land up to wherever that arrow falls.

They believed it was a token prize, an insult with a ribbon on it. They did not suspect it could redraw a border.

The Persians were desperate. Who could be entrusted with a task that absurd, and that consequential? Who could carry the future of a nation on one shot?

There was one archer whose name rose above the rest: **Arash**.

Arash understood what was being asked of him. This was not a contest of skill for sport or glory. With one arrow, he would mark the fate of his people and of generations who would never know his name unless he succeeded. He gathered himself completely. Feet planted. Breath measured. Hands steady. All of him – strength, discipline, intent – drawn into a single line.

And then he released.

The arrow flew, deep and far. So far that time seemed to stretch around it. Day turned to night. Night yielded to morning. Another dusk fell. The world kept moving as the arrow kept traveling, until at last it landed.

And when it did, it landed impossibly far.

So far that the border it fixed returned vast land to Iran, even after defeat – an outcome no one had dared imagine when the bargain was offered in scorn.

As for Arash: the story says he had poured everything into that one shot. When the arrow finally came to rest, he did not rise again. They climbed Mount Damavand to find him, but nothing remained — only his bow, lying on the ground where he had stood. He perished – not as a casualty of battle, but as the cost of a single, concentrated act.

And yet he endures: Arash, the archer who – through focus, through one arrow – quite literally moved the boundary of a world. That is the theme of this final chapter: how something narrow, precise, and seemingly small can set far-reaching things in motion.

CHAPTER 4

Dissociable causal roles of the nucleus accumbens and amygdala in motivation-action coupling

ABSTRACT

Motivated choice is shaped by both Pavlovian control and instrumental learning. The first biases actions toward approach under reward and toward inhibition under threat. The second mechanism updates which actions are worth repeating based on their outcomes. Animal work has long implicated the nucleus accumbens (NAcc) and amygdala in these processes, but their separate causal contributions to Pavlovian control and instrumental learning in humans remain unclear. Here, we used transcranial ultrasound stimulation (TUS) to probe the roles of NAcc and amygdala at the interface of motivational and instrumental control. In a within-subject design, healthy participants completed three intervention sessions with offline TUS to NAcc, amygdala, or sham, followed by a continuous-learning motivational Go/NoGo task in which cue valence (Win Reward vs. Avoid Loss) was orthogonalized with required action (Go vs. NoGo). Using computational modelling, we show dissociable effects of NAcc and amygdala stimulation. NAcc stimulation globally enhanced speed of instrumental learning, particularly when outcomes exceeded expectations, and simultaneously strengthened Pavlovian control. This suggests enhanced instrumental learning and stronger Pavlovian bias are not the product of competing systems but both arise from NAcc computations. Amygdala stimulation, by contrast, selectively enhanced learning for novel cues (i.e., the first encounter with a cue), leaving learning on later trials unaffected. Together, these findings provide causal evidence for temporally distinct contributions of NAcc and amygdala to motivated learning and action selection, and highlight transcranial ultrasound as a potentially powerful tool to dissect deep-brain mechanisms underlying maladaptive motivational biases in psychiatric disorders, such as addiction.



INTRODUCTION

Motivated decisions arise from the interplay between two systems: Pavlovian control, by which reward-predictive cues energize approach and loss-predictive cues promote inhibition, and instrumental learning, by which actions are updated based on their consequences[1–3]. How these systems interact, and which neural circuits implement their balance, is central to understanding both adaptive behavior and maladaptive patterns seen in psychiatry[4–11]. Two key nodes of this circuitry are the nucleus accumbens (NAcc) and the amygdala[12–14]. Yet, despite extensive animal work and growing human neuroimaging evidence, the dissociable contributions of NAcc and amygdala to cue-driven motivational biasing (Pavlovian control) versus outcome-driven learning (instrumental control) remains debated. Resolving this requires causal manipulations of these subcortical structures in humans, which has only recently become feasible with non-invasive transcranial ultrasound stimulation (TUS).

Pavlovian control biases actions based on motivational prospect alone: cues predicting reward invigorate action, while cues predicting loss promote inhibition, irrespective of what is instrumentally optimal[1,2,15–17]. Such Pavlovian biases are thought to be particularly useful in novel or uncertain environments, where the optimal action is not yet known and a default strategy is needed[1,2,15–17]. Rodent work has established a central role for the NAcc in translating such Pavlovian value signals into invigoration: in Pavlovian–instrumental transfer paradigms reward-predictive cues boost instrumental responding beyond what learning alone predicts. This effect critically depends on intact NAcc function[18–20]. In human studies, the picture is mixed: some studies report striatal responses tracking Pavlovian invigoration[16,21,22], whereas orthogonalized Go/NoGo studies (which isolate valence from action demands) find robust behavioral biases without consistent NAcc valence coding, observing instead preferential NAcc activity preceding Go responses[4,5,23,24]. Our own fMRI work (**Chapter 3**) showed cue-locked NAcc activity differentiating Win from Avoid cues — underscoring that these inconsistencies persist even across closely related paradigms. These discrepancies highlight two gaps: paradigm differences make it hard to translate circuit-level accounts from animal models to humans, and existing evidence is correlational, leaving open whether NAcc is causally required for Pavlovian control.

While Pavlovian biases provide a useful default, adaptive behavior ultimately requires learning which specific actions reliably produce desired outcomes for each cue encountered. Instrumental learning concerns how actions are selected and updated based on their consequences, driven by prediction errors that signal whether outcomes were better or worse than expected[25,26]. NAcc lesions or inactivation impair this process in animals, blunting sensitivity to outcome devaluation and slowing the acquisition of

action–outcome contingencies[13,15,18]. Human fMRI studies complement this observation by demonstrating NAcc signals that track outcome value and reward prediction errors, consistent with a role in updating action values based on feedback[5,27,28]. The amygdala's contribution to learning has been most clearly established in Pavlovian conditioning – the acquisition of stimulus-outcome associations through experience – which is distinct from Pavlovian bias, the expression of hardwired tendencies triggered by a cue[29–32]. Via projections to striatum and prefrontal cortex, it is well-positioned to relay these associations to circuits governing action selection[13,29,31–34]. Our previous fMRI work (**Chapter 3**) showed that both NAcc and amygdala responded more strongly to desired than to non-desired outcomes, indicating a shared role in outcome evaluation.

Animal studies point to a further specialization within the amygdalo-striatal circuit: a distinction between flexibility and stability. Amygdala neurons respond strongly to novel or abruptly changing cues, whereas NAcc integrates feedback over longer timescales and remains engaged when cue–action–outcome mappings are familiar but embedded in a noisy environment[12,35–37]. Our fMRI data (**Chapter 3**) provide support for this view: amygdala responses were transiently elevated when cues were encountered for the first time, irrespective of outcome prospect, whereas NAcc showed consistently stronger outcome-locked responses for desired relative to non-desired feedback, consistent with a sustained role in feedback-based updating. However, direct evidence for this temporal phase-dependent specialization in learning comes from animal work; causal tests of these distinct contributions of NAcc and amygdala to instrumental learning in humans are still lacking.

Understanding how NAcc and amygdala causally contribute to motivated action selection in humans is a critical gap: these structures are implicated in maladaptive motivational biases across psychiatric disorders[7–11,38]. Here we address this by combining offline TUS to NAcc, amygdala, or sham (within-subject, counterbalanced) with a gamified continuous-learning motivational Go/NoGo (MGNG) task in which cue valence (Win vs. Avoid) was orthogonalized with required action (Go vs. NoGo). TUS offers high spatial precision and the ability to reach deep brain structures non-invasively, and recent human studies have shown that TUS can reliably alter behavior[39–45]. Critically, the design of this task with the introduction of many novel cues while maintaining a steady level of repetition and updating is an innovation over conventional MGNG tasks. This allows robust estimation of learning effects from the very first cue encounter – enabling for the first time in humans a direct test of the novelty learning bias. Using hierarchical reinforcement-learning models, we found dissociable effects of NAcc and amygdala stimulation. NAcc stimulation strengthened Pavlovian control — producing a stronger initial bias – while simultaneously enhancing the overall rate of instrumental learning, such that these biases were more rapidly

overwritten as action-outcome contingencies were acquired. This suggests that these reflect a shared NAcc-driven mechanism rather than competing systems. Amygdala stimulation, by contrast, selectively enhanced learning for novel cues while leaving learning on subsequent trials unaffected. These findings provide the first causal evidence in humans for distinct contributions of these structures to motivated learning and action selection. They further highlight how targeted neuromodulation of this circuit may help to understand — and ultimately ameliorate — maladaptive motivational biases in psychiatric disorders.

RESULTS

Twenty-nine healthy participants (15 females and 14 males; age range 19–35 years; mean = 24.8, SD = 4.1 years) performed a MGNG learning task after stimulation of the NAcc, amygdala, or sham (**Figure 1A**, study design; **Figure 1B**, stimulation protocol). In this task, cue valence (Win Reward vs. Avoid Loss) was orthogonalized to required action (Go vs. NoGo): for each cue, participants had to learn whether to press or withhold a response, such that correct actions yielded reward on Win cues and avoided loss on Avoid cues, on ~75% of trials, while incorrect actions led to these desired outcomes only on ~25% of trials (**Figure 1C-D**). Crucially, unlike previous implementations of this paradigm[4,5,7,17,24,46], participants remained in a continuous learning regime: they encountered 24 distinct cues that were gradually introduced and removed across the session, and for one third of these cues, the required action reversed midway (**Suppl. Figure S1**). This allowed continuous learning with a higher rate of novel and reversed cue-outcome associations than previous studies.

Sham behavior: successful instrumental learning and stable Pavlovian bias

We first characterized choice behavior in absence of stimulation, i.e., in the sham (no-TUS) sessions, using logistic mixed-effects regression to the Go/NoGo choices with required action (Go vs. NoGo), cue valence (Win vs. Avoid), cue-specific trial number (z-scored), and all two-way interactions. Participants showed clear instrumental learning as they emitted more Go responses to Go than to NoGo cues (required action: b = 0.79, 95%-CI [0.62, 0.95], $\chi^2(1) = 87.2$, $p < 2.2e\text{-}16$; **Figure 1E**), which increased over repeated encounters with the same cue (required action × trial number: b = 0.34, 95%-CI [0.27, 0.52], $\chi^2(1) = 39.2$, $p < 3.7e\text{-}10$; **Figure 1F-G**). Choices were additionally shaped by a robust motivational bias: Go responses were more frequent for Win than for Avoid cues (cue valence: b = 0.20, 95%-CI [0.13, 0.28], $\chi^2(1) = 26.6$, $p < 2.5e\text{-}07$; **Figure 1E**). This valence effect did not significantly change across cue repetitions (cue valence ×

trial number: b = 0.04, 95%-CI [-0.02, 0.09], $\chi^2(1)$ = 1.7, p = .19), nor was it significantly different for Go vs. NoGo cues (required action × cue valence: b = 0.06, 95%-CI [-0.003, 0.121], $\chi^2(1)$ = 3.5, p = .06). Analysis of RT effects showed similar evidence of robust learning and motivational bias (see **Supplementary Information S.1**). However, given that RTs are present on only ~50% of trials (i.e., for Go responses), we focus the further analysis of the TUS effects on choice data.

To characterize learning effects in choice as a function of feedback, we analyzed win-stay/lose-shift (WSLS) tendencies. For each trial, we constructed a regressor capturing whether the current response reflected "staying" with the previous choice after a win, or "shifting" away from the previous choice after a loss for the same cue. We then fitted a logistic mixed-effects regression with Go response as the dependent variable and required action, cue-specific trial number, cue valence, and WSLS as fixed factors. This model revealed a robust main effect of WSLS (b = 0.56, 95%-CI [0.49, 0.63], $\chi^2(1)$ = 265.1, p < 2.2e-16; **Figure 1H**), indicating that participants tended to repeat their previous response following a desired outcome and to switch following a non-desired outcome for that cue.

To characterize how experience with the task altered behavior across the study sessions, we examined choice patterns as a function of intervention session number. Participants performed the task five times in total (two runs in the scanner with fMRI-adjusted slower and jittered timings, immediately followed by one run outside the scanner with behavioral task timings during the intake session, and two runs during each of the three TUS sessions). In summary, accuracy increased over intervention sessions (required action × session number: b = 0.10, 95%-CI [0.04, 0.17], χ^2 = 9.5, p = .002), and Pavlovian bias reduced over intervention sessions (cue valence × session number: b = -0.05, 95%-CI [-0.9, -0.01], χ^2 = 6.0, p = .014). Importantly, however, a robust Pavlovian bias was still present even in the final intervention session (cue valence in intervention session 3: b = 0.20, 95%-CI [0.13, 0.27], $\chi^2(1)$ = 28.4, p = 1.0e-07). Full analyses and statistics are reported in supplementary information (**Suppl. Information S.2**).

Computational modelling

To more precisely quantify and dissect Pavlovian biases and the (nonlinear) learning dynamics in behavior, we employed a computational modelling approach. This approach is particularly appropriate for the current paradigm, where the continuous introduction of novel cues as well as presence of reversals keeps participants in a high-learning regime – one in which standard linear approaches are poorly suited to capture the resulting nonlinear choice dynamics. Computational modelling therefore

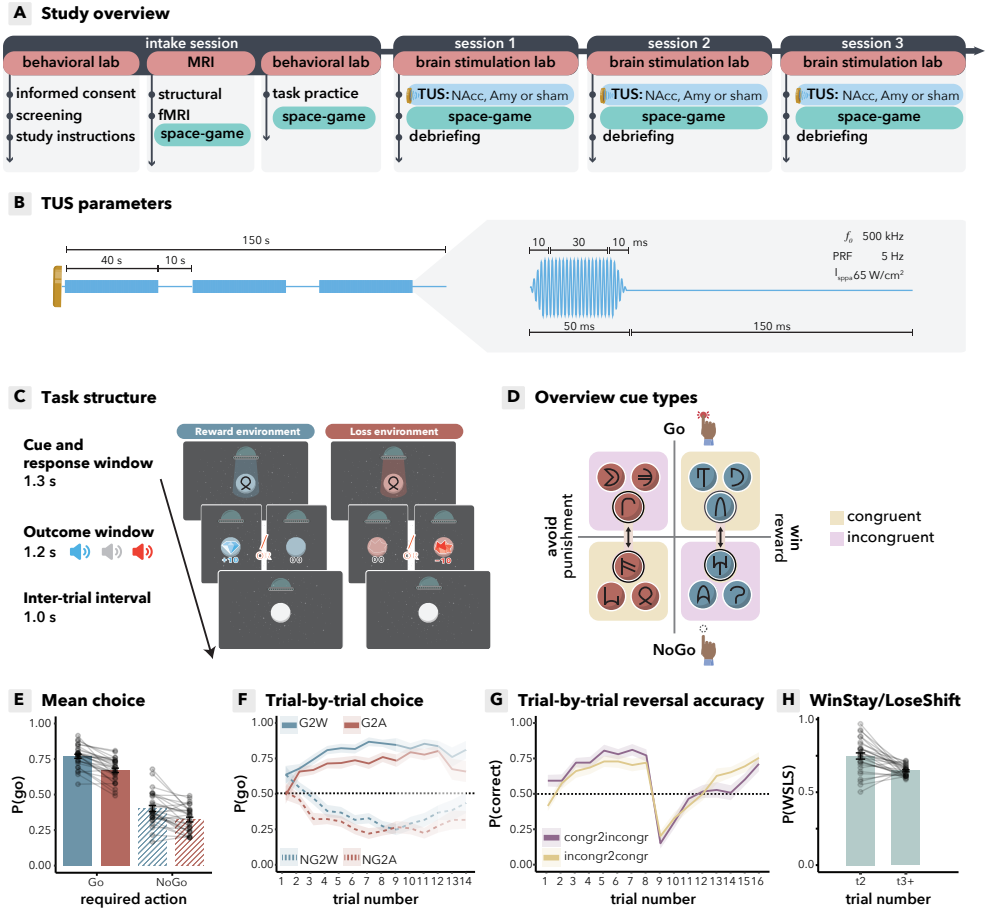


Figure 1 | Task design and baseline behavior across intervention sessions. (A) *Within-subject, four-session design* (n = 29). In the intake session, participants performed the "space game" task in the MRI scanner and then completed a practice block with adjusted trial-timings appropriate for behavioral testing. In three subsequent intervention sessions, they received NAcc, AMY, or sham TUS (counterbalanced) before playing the task. **(B)** *TUS protocol.* Each hemisphere received 150 s of sonication (three cycles of 40 s on, 10 s off). Pulses were delivered at 500 kHz with a pulse repetition frequency of 5 Hz and free-water I_{sppa} of 65 W/cm². Each pulse lasted 50 ms, including 10 ms Tukey ramps at onset and offset, followed by 150 ms off-time within the 200 ms (5 Hz) cycle. **(C)** *Trial structure.* On each trial, a cue (planet) appeared with a 1.3 s cue/response window, followed by outcome (1.2 s) – comprising visual feedback accompanied by a congruent auditory tone – and a 1.0 s inter-trial interval, in reward (win points) or loss (avoid loss) environments. **(D)** *Four cue types orthogonalized action (Go vs. NoGo) and valence (Win vs. Avoid).* Each run contained three cues per condition; one cue per condition reversed its required action halfway through the run. Yellow panels indicate the congruent conditions Go-to-Win-reward (G2W) and NoGo-to-Avoid-loss (NG2A), which align with our Pavlovian tendencies and are typically easier; purple panels indicate the incongruent conditions Go-to-Avoid-loss (G2A) and NoGo-to-Win-reward (NG2W), which conflict with our Pavlovian bias and are more difficult. Participants completed two runs per session. **(E)** *Mean choice behavior pooled over the three intervention sessions.* Participants made more Go responses when Go was required and more Go

responses to Win than Avoid cues, indicating successful instrumental learning and a robust Pavlovian bias. Bars = SEM; dots = participants. **(F)** *Trial-by-trial Go probability pooled over intervention sessions.* Participants increasingly select the correct action over cue repetitions without reaching ceiling, while the Pavlovian bias (Win > Avoid) remains stable for both Go (solid) and NoGo (dashed) trials. Lines show mean P(Go) per condition and cue repetition; shaded ribbons indicate SEM. **(G)** *Trial-by-trial reversal learning.* Accuracy for reversal cues over cue repetitions. The required action reverses after the 8th repetition. Accuracy is higher when mappings are initially congruent (G2W, NG2A; purple) than when they are initially incongruent (NG2W, G2A; yellow). Lines show mean accuracy per trial repetition; shaded ribbons indicate SEM. **(H)** *Win–stay/lose–shift behavior.* Probability of staying after wins or shifting after losses is higher on trial 2 (immediately after the first encounter with a cue) than on later trials, indicating stronger novelty-driven updating. Bars = SEM; dots = participants.

serves as the primary lens through which we interpret TUS effects on learning and bias. We fitted a set of nested Q-learning models to trial-by-trial Go/NoGo choices. All models assumed delta-rule updating of cue–action values and a logistic choice rule, but differed in how they parameterized learning rates, outcome sensitivity, and motivational action biases. This approach provides a clean decomposition of latent computations shaping choice.

Model validation

Model architecture was validated on an independent dataset acquired with the same task (fMRI data reported in **Chapter 3**), using leave-one-out cross-validation (LOO) to select among a set of nested Q-learning models that progressively added a Go bias, valence-specific Pavlovian biases, and asymmetric learning rates. Of these models, M6 — which incorporates valence-specific Pavlovian biases and separate learning rates for positive and negative prediction errors — provided the best account of participants' choices. We then added a novel addition to this model space, extending winning model M6 with an additional distinction between first and subsequent cue presentations. This amended model was motivated by our findings in **Chapter 3**, where we showed enhanced neural activity for novel versus familiar cues, as well as the enhanced WSLS behavior following cues (**Figure 1H**, **Suppl. Figure S4B**). This suggests that with the first time a cue-action-outcome combination is observed, learning is boosted. Model comparison showed that M7 provided the most parsimonious account of participants' choices (**Suppl. Figure S4C**). This winning model captures (i) a separate learning rate for novel cues, (ii) differential subsequent learning sensitive to better-than-expected than to worse-than-expected outcomes (i.e., positive vs. negative prediction errors), and (iii) Go bias modulation by cue valence. Full model comparison results, parameter estimates, and posterior predictive simulations of the winning model are reported in the supplementary information (**Suppl. Information S.3**, **Figure S4**, **Table S1**).

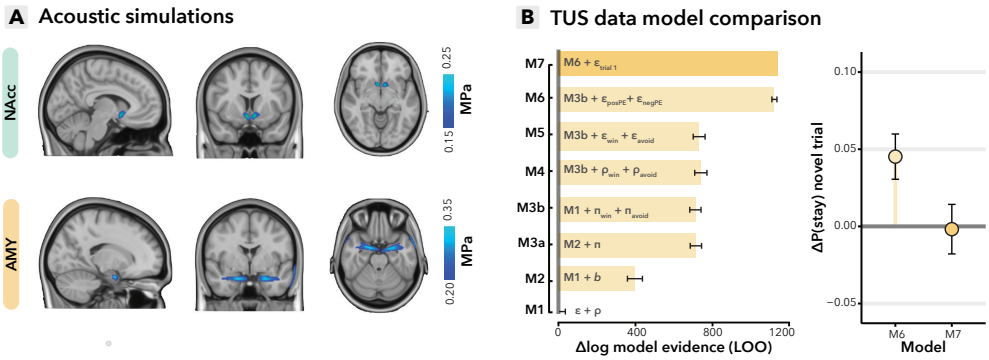


Figure 2 | Acoustic simulations and reinforcement learning model validation. (A) Group-level normalized acoustic simulations of TUS wave propagation for NAcc (top) and amygdala (bottom), shown as estimated peak pressure (MPa) in individual anatomies at the recorded transducer locations during intervention sessions. In each row, the corresponding group-level target mask outline in white is overlaid as an outline (NAcc mask in the top panels, amygdala mask in the bottom panels) to indicate the intended stimulation site. **(B)** Model comparison for the TUS data. Left: ΔLOO relative to the worst model. Models that combined a Pavlovian bias with separate learning rates for positive and negative prediction errors outperformed simpler variants; adding a dedicated novelty learning rate (M7) yielded the best fit despite its extra parameter. Right: difference between model-predicted and empirical stay behavior on novel trials (mean ± SEM ΔP(stay)); M7 captures this novelty-related increase in P(stay) better than M6.

TUS model

To probe the causal contributions of NAcc and amygdala to the interplay of instrumental learning and Pavlovian control, we used TUS to selectively target each region in separate intervention sessions (**Figure 1A**) and verified targeting with individualized acoustic simulations (**Figure 2A**, **Suppl. Figure S5**). We then asked how NAcc and amygdala stimulation affected the latent learning and motivational processes underlying Go/NoGo choices. Building on the model comparison in the independent dataset, we focused on the winning model M7 for the TUS data. For completeness, we also re-fitted the full model space (M1–M7) to the TUS sessions and repeated the LOO-based model comparison; the ranking was essentially identical (M7 > M6 > remaining models; **Figure 2B**). A full overview of posterior parameter estimates for M6 and M7 is provided in the supplementary information (**Table S2**). For all models, TUS condition (sham, NAcc, amygdala) entered at the parameter level as condition-specific offsets around a shared subject-specific baseline, enabling within-subject estimation of TUS-induced changes in each latent parameter.

Because each cue was presented only 12–16 times and the task included reversals, participants remained in a high learning regime throughout, making computational modelling the main analytical approach. Where informative, we complement

model-based findings with regression results to provide converging support. For completeness, the full regression analyses of TUS effects on choice and reaction time are reported in the supplementary information (**Suppl. Information S.4**, **Tables S3–S11**).

Pavlovian bias

Pavlovian biases capture how the mere prospect of reward or loss shapes action, over and above what is instrumentally optimal. Reward-predictive cues tend to invigorate approach and Go-responding, whereas loss-predictive cues promote withholding and avoidance (**Figure 3A**). In our reinforcement learning framework, these effects are summarized by Pavlovian bias parameters that quantify appetitive Go tendencies towards Win cues and aversive Go tendencies towards Avoid cues.

In the full M7 model, the overall Pavlovian bias (π_{win} minus π_{avoid}) was higher under both amygdala and NAcc stimulation than under sham (TUS condition: $\chi^2(2) = 15.8$, $p = 3.7e\text{-}04$; NAcc -sham: $b = 0.17$, SE = 0.05, 95%-CI [0.06, 0.28], $p = .001$, Cohen $D = 0.88$; amygdala – sham: $b = 0.14$, SE = 0.05, 95%-CI [0.25, 0.04], $p = .003$, Cohen $D = 0.72$; NAcc – amygdala: $b = 0.03$, SE = 0.05, 95%-CI [-0.08, 0.13], $p = .6$; **Figure 3B**, **left**), indicating an overall strengthening of valence-dependent biases when either structure was stimulated. To test whether this global effect was symmetric for reward and loss, we decomposed the Pavlovian term into an appetitive Go bias for Win cues (π_{win}) and an aversive Go bias for Avoid cues (π_{avoid}). For π_{win}, there was no evidence for a TUS effect (TUS condition: $\chi^2(2) = 2.1$, $p = .4$; **Figure 3B**, **middle**), indicating that appetitive Go-tendencies towards reward-predictive cues were largely unchanged. By contrast, π_{avoid}, showed a clear TUS effect (TUS condition: $\chi^2(2) = 37.0$, $p = 9.0e\text{-}09$; **Figure 3B**, **right**). Relative to sham, both NAcc and amygdala stimulation reduced the aversive Go bias parameter (stronger π_{avoid}; NAcc - sham: $b = -0.16$, SE = 0.03, 95%-CI [-0.23, -0.10], $p = 1.2e\text{-}07$, Cohen $D = 1.06$; amygdala – sham: $b = -0.09$, SE = 0.03, 95%-CI [-0.15, -0.03], $p = .001$, Cohen $D = 0.93$), such that participants made fewer Go responses on Avoid cues after stimulation (**Figure 3B**, **right**). This reduction was significantly larger for NAcc than for amygdala (NAcc - amygdala: $b = -0.07$, SE = 0.03, 95%-CI -0.14, -0.01], $p = .009$, Cohen $D = 0.42$), suggesting that NAcc exerted a stronger influence on aversive bias.

We followed up these model-based analyses of Pavlovian bias with targeted logistic mixed-effects regressions, focusing on pairwise regression contrasts of stimulation relative to sham. Perhaps surprising given the change in Pavlovian bias in the model-based analysis, there was no significant effect of cue valence for stimulation relative to sham (cue valence x TUS condition [NAcc – sham]: $b = 0.02$, 95%-CI [-0.02, 0.06],

$\chi^2(1) = 0.6$, $p = .4$; [amygdala – sham]: $b = 0.01$, 95%-CI [-0.03, 0.06], $\chi^2(1) = 0.2$, $p = .6$; [NAcc – amygdala]: $b = 0.01$, 95%-CI [-0.03, 0.04], $\chi^2(1) = 0.1$, $p = .8$). However, NAcc stimulation did affect Pavlovian bias differently over time relative to sham: for early trials, the Pavlovian bias was stronger after NAcc TUS than under sham stimulation, while in later trials, participants were less biased after NAcc TUS than following sham (cue valence x trial number x TUS condition [NAcc - sham]: $b = -0.04$, 95%-CI [-0.08, -0.00], $\chi^2(1) = 4.7$, $p = .030$; **Figure 3C**). No analogous interaction emerged for amygdala versus sham (cue valence x trial number x TUS condition [amygdala - sham]: $b = -0.01$, 95%-CI [-0.05, 0.03], $\chi^2(1) = 0.3$, $p = .6$; **Figure 3C**). Breaking down this effect showed that only in the NAcc sessions did cue valence interact significantly with cue-specific trial number ($b = -0.06$, 95%-CI [-0.11, -0.01], $\chi^2(1) = 4.6$, $p = .032$), whereas no such interaction was observed in the sham ($b = 0.03$, 95%-CI [-0.02, 0.08], $\chi^2(1) = 1.4$, $p = .2$) or amygdala sessions ($b = 0.01$, 95%-CI [-0.06, 0.08], $\chi^2(1) = 0.1$, $p = .8$). Together, these results suggest that NAcc stimulation produced an initially heightened Pavlovian bias that was more rapidly overcome with experience: a pattern that points to an interaction between strengthened bias and accelerated outcome-driven updating. We therefore turn next to the direct effects of TUS on learning rates.

Learning rate

Learning rates capture how strongly participants update their beliefs about cue-action values in response to outcomes. In our model, we distinguished the degree of updating for the very first encounter with a cue (ε_1) from subsequent updates following better than expected (ε_{rest_pos}) versus worse than expected (ε_{rest_neg}) outcomes (**Figure 4A**).

TUS broadly boosted learning rates: there was a significant overall effect of TUS condition on both the novelty-sensitive learning rate ε_1 ($\chi^2(2) = 158.9$, $p < 2.2e\text{-}16$) and the learning rate for positive prediction errors on familiar trials ε_{rest_pos} ($\chi^2(2) = 12.9$, $p < 0.002$). Decomposing these effects by region, revealed shared and distinct contributions of NAcc and amygdala. For the novelty learning rate ε_1, both NAcc and amygdala stimulation produced a robust increase relative to sham (NAcc - sham: $b = 0.11$, SE = 0.01, 95%-CI [0.095, 0.14], $p = 7.0e\text{-}15$, Cohen *D* = 4.0; amygdala - sham: $b = 0.12$, SE = 0.01, 95%-CI [0.14, -0.09], $p = 5.1e\text{-}16$, Cohen *D* = 1.9; **Figure 4B**, **left**), indicating that stimulation of either structure accelerates belief updating at the very first encounter with a cue. The two active conditions did not reliably differ from one another (NAcc - amygdala: $b = -0.01$, SE = 0.01, 95%-CI [-0.03, 0.02], $p = .5$), suggesting comparable enhancement of novelty-driven learning across regions. For learning on later trials, only NAcc stimulation increased the positive-PE learning rate ε_{rest_pos} relative to sham (NAcc - sham: $b = 0.02$, SE = 0.01, 95%-CI [0.003, 0.04], $p = .007$, Cohen *D* = 0.70; **Figure 4B**, **middle**), whereas the amygdala effect was small and not credibly

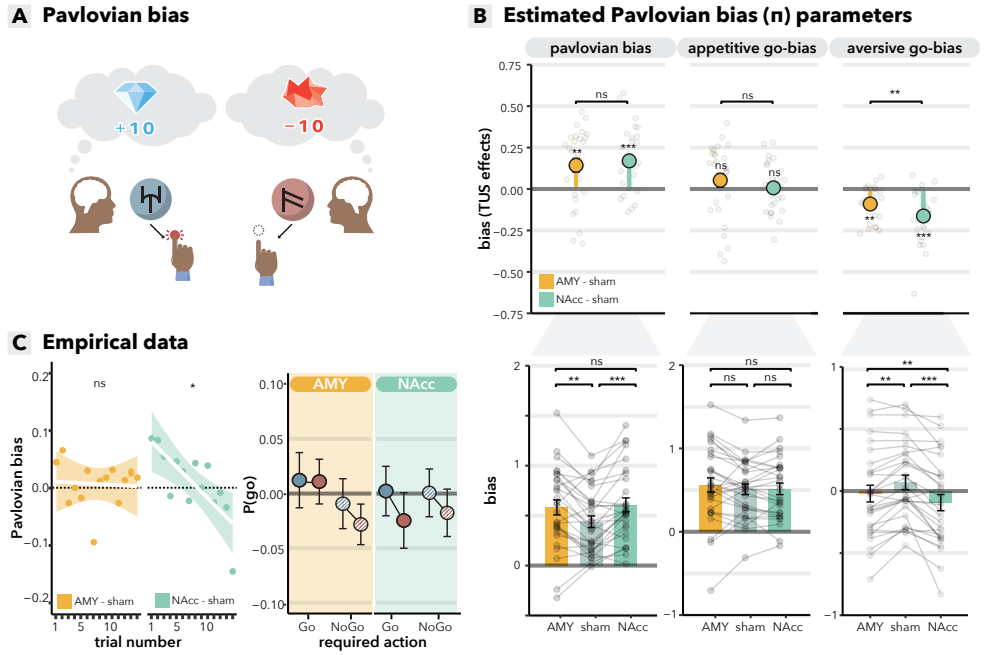


Figure 3 | TUS effects on Pavlovian bias. (A) *Schematic illustration of Pavlovian bias.* Cues predicting reward (left) tend to invigorate Go responses, whereas cues predicting loss (right) tend to inhibit action. **(B)** *Estimated Pavlovian bias (π) parameters.* Top: TUS-induced changes in bias (condition – sham) for amygdala (orange) and NAcc (green); circles = group means ± SEM, faint dots = posterior means per participant. Bottom: corresponding parameter estimates per condition (bars = means ± SEM; dots = participants connected across conditions). Both amygdala and NAcc stimulation increased the overall Pavlovian bias relative to sham (left). Decomposing this effect shows no reliable change in appetitive Go bias (middle), but a clear reduction in aversive Go bias (right): after TUS, participants were less likely to make Go responses to loss-predictive cues, with this effect being strongest for NAcc. Note that for the overall Pavlovian bias and appetitive Go-bias, positive values indicate stronger bias; for the aversive Go-bias, negative values indicate stronger Pavlovian bias (i.e, less Go responding to loss-predictive cues). **(C)** *Empirical data.* Left: TUS effects on trial-wise Pavlovian bias (Win–Avoid difference in P(Go)) for amygdala–sham (orange) and NAcc–sham (green); only NAcc shows a significant decline in bias over repetitions. Right: TUS-induced changes in P(Go) for each cue type (G2W, G2A, NG2W, NG2A). After NAcc stimulation (green panel), both Avoid conditions show reduced Go responding, consistent with a stronger aversive Go bias, whereas amygdala effects are smaller and largely confined to NG2A.

different from zero (amygdala-sham: b = -0.005, SE = 0.01, 95%-CI [0.01, -0.02], p = .6; **Figure 4B**, **middle**). Accordingly, ε_{rest_pos} was significantly higher after NAcc than after amygdala stimulation (NAcc - amygdala: b = 0.03, SE = 0.01, 95%-CI [0.01, 0.05], p = .002, Cohen D = 0.47; **Figure 4B**, **middle**). Neither TUS condition reliably modulated the negative-PE learning rate ε_{rest_neg} ($\chi^2(2)$ = 4.2, p = .12; **Figure 4B**, **right**). Additional analyses examining the specificity of TUS effects across positive and negative prediction errors are reported in supplementary information (**Suppl. Information S.5**).

We again followed up these parameter-level analyses with targeted logistic mixed-effects regressions. We first examined whether overall performance over trials (i.e., the degree to which participants linearly increasingly selected the correct action over repetitions of the same cue) differed between TUS conditions. Pairwise logistic mixed-effects regression contrasts revealed no significant differences either across all conditions ($\chi^2(2) = 0.7$, $p > .9$) or between any pair of conditions (required action x trial number x TUS stimulation [NAcc - sham]: $b = 0.01$, 95%-CI [-0.09, 0.11], $\chi^2(1) = 0.9$, $p = .6$; [amygdala - sham]: $b = -0.01$, 95%-CI [-0.11, 0.08], $\chi^2(1) = 0.7$, $p = .7$; [NAcc - amygdala]: $b = -0.02$, 95%-CI [-0.12, 0.08], $\chi^2(1) = 0.3$, $p = .9$). Note, however, that these linear assumptions about learning are unable to fully capture the highly nonlinear performance trajectory (**Figure 1F**), particularly in this task, as the cues were only presented a limited number of times.

Nonetheless, analyses of direct behavioral updating based on previous outcomes (i.e., WSLS behavior as a function of desired / non-desired outcomes) did mirror the model-based novelty effects. Pairwise comparisons between conditions showed that the difference in WSLS behavior between trial 2 (capturing the immediate consequence of the first encounter with a cue) and trial 3 (when the cue was already familiar) was significantly larger for amygdala than for sham (amygdala - sham: $b = -0.11$, 95%-CI [-0.213, -0.002], $\chi^2(1) = 4.0$, $p = .046$; **Figure 4C**), indicating that amygdala stimulation altered how novelty-driven information was incorporated into subsequent choices. For the NAcc – sham contrast, this difference was not significant (NAcc - sham: $b = -0.05$, 95%-CI [-0.15, 0.05], $\chi^2(1) = 0.9$, $p = .3$; **Figure 4C**).

For later familiar trials, WSLS behavior did not differ significantly between NAcc and sham sessions (WSLS x TUS condition: $\chi^2(2) = 2.8$, $p = .2$; NAcc - sham: $b = 0.02$, 95%-CI [-0.02, 0.07], $\chi^2(1) = 0.8$, $p = .4$). This appears in contrast to the robust NAcc-specific increase in ε_{rest_pos}. A likely explanation is that a learning rate reflects integration of outcomes over multiple past trials, whereas WSLS is strictly one-step-back: it only considers whether the current choice repeats or switches relative to the immediately preceding outcome for that cue. It is therefore appropriate for the novelty effect, but perhaps less so for more graded learning expressed through a learning rate. As a result, the richer, multi-trial learning dynamics captured by ε_{rest_pos} may not be fully expressed in simple first-order stay/switch summaries, underscoring the added sensitivity of the computational modelling approach. Full statistical details and all pairwise comparisons are reported in the supplementary information (**Tables S3**-**S11**).

Together, these results reveal distinct yet partly shared contributions of NAcc and amygdala to motivated behavior. NAcc stimulation produced a broad enhancement of both Pavlovian bias and instrumental learning: Pavlovian

bias was initially stronger but dissipated more rapidly across trials, a pattern consistent with the accompanying increase in the positive-PE learning rate, which would accelerate the updating of cue-action values as experience accumulated. Amygdala stimulation, by contrast, selectively boosted novelty-driven learning, with no reliable effect on Pavlovian bias dynamics or learning from familiar cues. This double dissociation points to a division of labor in which the NAcc links motivational value to action and supports sustained outcome-driven updating, while the amygdala primarily amplifies learning at the first encounter with a new cue.

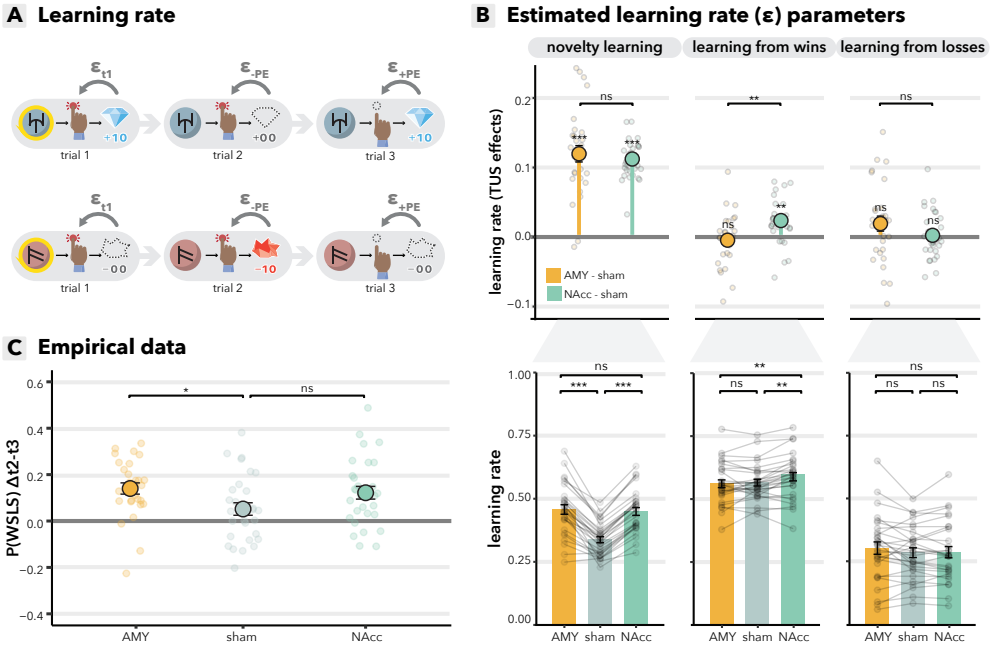


Figure 4 | TUS effects on learning rates and novelty bias. (A) *Schematic illustration of the three learning-rate components in M7 for Win (top) and Avoid (bottom) cues.* Novelty learning on trial 1 (ε_1; yellow outline), updating after negative prediction errors (ε_rest_neg; no reward or loss), and updating after positive prediction errors (ε_rest_pos; reward or avoided loss). **(B)** *Estimated learning-rate (ε) parameters.* Top: TUS-induced changes in learning rates (condition – sham) for amygdala (orange) and NAcc (green); circles = group means ± SEM, faint dots = posterior parameter means per participant. Bottom: corresponding parameter estimates per condition (bars = means ± SEM; dots = participants connected across conditions). Both amygdala and NAcc stimulation increased ε_1 relative to sham, indicating enhanced novelty learning, whereas only NAcc stimulation significantly increased ε_rest_pos; neither stimulation affected ε_rest_neg. **(C)** *Empirical win–stay/lose–shift behavior.* The difference of P(WSLS) between trial 2 (immediately after the first encounter) and trial 3 for amygdala (orange), sham (grey), and NAcc (green). Novel cues show stronger WSLS effects than familiar cues, with the trial-2 vs. later-trial difference being significantly larger for amygdala than for sham. Circle = means ± SEM; dots = participants.

DISCUSSION

Motivated actions are shaped by both hardwired Pavlovian tendencies — invigorating to reward-predictive cues and inhibiting to loss-predictive cues — and instrumental learning, in which actions are updated based on experienced outcomes[1–3]. A key question in this domain is how hardwired tendencies and flexible learning jointly govern action selection. Despite extensive animal work and growing human neuroimaging evidence, the causal contributions of the NAcc and amygdala to this balance have remained unresolved[4,13,23,24]. This gap has direct implications for understanding motivational biases in psychiatric disorders[8–11,38]. Here, we used transcranial ultrasound stimulation to transiently modulate NAcc and amygdala function in healthy participants performing a continuous-learning motivational Go/NoGo task. Our findings show that while both regions contribute to motivational control of action, their roles are distinct: NAcc stimulation modulates goal-directed action by boosting both Pavlovian and instrumental control of action, whereas amygdala stimulation primarily enhances learning for novel stimuli.

NAcc TUS boosts Pavlovian and instrumental control

Our results provide causal evidence that NAcc supports Pavlovian control in humans — a role long established in animal work but never proven causally in humans. NAcc lesions in rodents reduce cue-triggered invigoration for reward-predictive stimuli and blunt the motivational impact of Pavlovian cues on instrumental responding[15,18–20]. However, human neuroimaging evidence has been mixed. Some fMRI studies mirror findings from animal models, observing increased NAcc BOLD to reward-predictive cues[16,21,22,47]. In contrast, fMRI observations in tasks where outcome valence and action requirements are orthogonalized do not show such enhanced NAcc BOLD response[4,5,23,24]. Notably, here, even within such an orthogonalized design, we do find that NAcc TUS increases the Pavlovian bias, aligning with animal studies[15,19,20] as well as our fMRI work in **Chapter 3**.

A possible source of the discrepancy between these findings and previous human MGNG studies is that in the current task design, cues and actions are tightly coupled in both time and action space. Participants act immediately upon cue presentation, and each cue maps into a single action. In prior studies, cues and actions were less clearly coupled. They were either separated in time[23,24] or afforded multiple responses[4,5]. In other words, in the current task design, the cue elicits a goal (win a reward, or avoid a loss) and, at the same time, instructs which response is required to obtain that goal. We therefore speculate that NAcc may drive prospective valence encoding specifically when goals are directionally actionable, emphasizing the role of

the striatum, embedded within cortico-striatal loops, in translating motivational value and goals into motor actions[14,48].

When the Pavlovian bias was decomposed by valence, effects appeared to be primarily driven by Avoid cues (i.e., fewer Go responses for Avoid cues after NAcc stimulation). While we cannot rule out that this enhancement of Pavlovian bias is indeed specific to the Avoid domain, a more likely explanation is that an absence of Pavlovian bias strengthening in the Win domain resulted from a baseline ceiling effect. Across many studies, invigoration of responding to Win cues is extremely robust and much stronger than the Avoid bias, leaving limited room for further increase[4,5,7,17,23,24,46]. Moreover, as a neuromodulatory technique, TUS is thought to be most pronounced when decisions are ambiguous; when behavior is already near ceiling, there is little room for further modulation regardless of intervention[40,41].

Importantly, following NAcc TUS, the facilitated initial Pavlovian bias also diminished more rapidly across cue repetitions. This pattern is explained by a concurrent enhancement of instrumental learning: NAcc TUS increased learning rates for both novel and familiar cues, meaning correct instrumental choices were reinforced more quickly, allowing acquired action values to compete with and eventually override the initial Pavlovian pull. This is consistent with prior animal and human work showing NAcc involvement in outcome-driven choice updating[5,13,18,26–28,49]. Encouragingly, our NAcc TUS-driven enhancement of reinforcement learning is a conceptual replication of previous work[45], and dovetails with a growing number of human TUS studies demonstrating facilitation of learning[39,42,43]. This convergence underscores both the reliability of the current findings and the growing promise of TUS as a tool to facilitate learning in humans.

Taken together, we present evidence that NAcc stimulation boosted both Pavlovian and instrumental control. Pavlovian and instrumental systems are often framed as competing, with Pavlovian tendencies interfering with optimal instrumental choice[2]. Our findings propose that rather than opposing forces, both may reflect complementary branches of goal-directed behavior mediated by the NAcc. Their relative influence is determined dynamically by the reliability of each estimate[50]. Early in learning, when little evidence has accumulated about the best action, Pavlovian biases provide a useful default. As instrumental contingencies are acquired, these hardwired tendencies are progressively overridden by learned action values. NAcc stimulation amplified both phases of this process: boosting the initial bias and accelerating the learning that ultimately corrects it. This suggests that the NAcc does not arbitrate between Pavlovian and instrumental control but rather serves as a common substrate that supports goal-directed behavior.

Amygdala TUS boosts novelty learning

Amygdala TUS robustly increased the learning for novel cues, and – unlike NAcc – left subsequent learning unaffected. This interpretation is supported by our fMRI results, where amygdala cue-locked activity was selectively elevated on the first, but not subsequent presentations of a cue (**Chapter 3**). Consistently, in the empirical TUS data we observed a larger change in win-stay/lose-shift behavior between trial 2 (capturing the immediate consequence of the first encounter with a cue) and trial 3 (when the cue is already familiar) after amygdala stimulation than after baseline sham, indicating that amygdala TUS specifically alters how novelty-driven information is incorporated into the very next choice. Critically, the specificity of the current amygdala effect has not previously been demonstrated. Out of 29 participants, 27 participants showed enhanced novelty learning following amygdala TUS, compared to sham learning increased from 0.35 to 0.46 – a 35% increase with a large effect-size (Cohen $D > 1$). These effect sizes are larger than often observed in other non-invasive neuromodulation studies in healthy humans[51–53] and speak to the amygdala's causal role in initializing value learning for new stimuli.

This pattern aligns with a well-established literature on novelty and saliency signaling in the amygdala. Amygdala neurons respond strongly to stimuli that are new, unexpected, or motivationally significant, and this response habituates rapidly with repetition[34,54–56]. Functionally, this novelty signal is thought to flag stimuli that warrant prioritized processing and may gate plasticity in downstream circuits[57,58]. Complementing this, earlier macaque work directly contrasts the learning dynamics of the amygdala and NAcc, showing that the amygdala drives fast, early learning (particularly when cue-outcome contingencies are new or changing) while the NAcc integrates feedback more slowly over time, providing a stable substrate for value-based action selection in familiar – but noisy – environments[12,37,49]. Our findings extend this framework to humans with causal resolution: NAcc TUS enhanced learning broadly across both novel and familiar cues, whereas amygdala TUS selectively boosted learning at first encounter.

Although amygdala TUS also increased the Pavlovian bias parameter, these effects were notably weaker than those following NAcc TUS and were not clearly supported by the regression models of choice data. The amygdala is classically implicated in Pavlovian conditioning – the acquisition and storage of cue-outcome associations[16,18,30–32,47,59–61] – which is not necessarily the same as the expression of hardwired action tendencies per se. In the current paradigm, cue valence was signaled explicitly from the outset via cue color, such that no new Pavlovian associations needed to be formed. The amygdala may therefore be relatively disengaged from Pavlovian

bias expression when valence is instructed rather than learned[62]. Moreover, in our fMRI data (**Chapter 3**), amygdala BOLD responses were not elevated for cue-locked valence coding either, whereas the NAcc BOLD responses were. Studies reporting reliable amygdala cue-valence signals indeed typically involve paradigms in which participants learn cue valence from experience rather than explicit instruction[5,16,59,62].

It remains an important open question whether the amygdala contributes not only to novelty-driven learning but also to fast updating more broadly – for instance, when contingencies change abruptly mid-task. These two accounts make distinct predictions: a pure novelty signal would be tied to first cue encounters, whereas a fast-updating account would predict amygdala involvement whenever contingencies change, including mid-task reversals. In the current task, reversals were included primarily to maintain a continuous learning regime and keep participants engaged, rather than to formally test reversal learning. Therefore, there are too few reversal trials to reliably distinguish these accounts. Future work explicitly designed around contingency reversals – with sufficient trials per reversal to model amygdala and NAcc contributions separately – would allow a clean dissociation between novelty-driven and change-driven learning signals.

Mechanisms of TUS-induced facilitation

The direction of our TUS effects (i.e., enhanced learning rates and stronger Pavlovian biases rather than impaired ones) suggests that our offline protocol had a net facilitatory effect on these circuits. In lesion and inactivation studies, both NAcc and amygdala damage reliably reduce learning from outcomes and weaken Pavlovian biases[13,18,49]. Our TUS effects run in the opposite direction. Focused ultrasound at low intensities exerts mechanical forces on neural tissue, likely modulating mechanosensitive ion channels, including calcium-permeable TRP and Piezo family members[63,64]. Calcium influx through these channels initiates and modulates intracellular signaling cascades that overlap substantially with those underlying synaptic long-term potentiation[65]. Offline TUS is therefore thought not to impose a fixed pattern of activity, but rather to modulate neuronal excitability and prime circuits for experience-dependent plasticity. The first mechanism affects spontaneous spiking probability, putatively facilitating or inhibiting neural computation. The second mechanism can open a window in which incoming behavioral signals can more readily modify synaptic weights[66]. In the context of the present paradigm, which placed participants in an active learning regime immediately after stimulation, this plasticity-priming effect would be expected to enhance the functional contributions of both targeted regions during subsequent behavior, consistent with the effects we observe.

Conclusion and future directions

In summary, TUS stimulation of the NAcc and amygdala reveals dissociable contributions to motivated action control, in which the amygdala detects salient new cues and boosts learning from these novel cues, putatively through projections to the NAcc "tagging" these novel cues for enhanced learning[29,55–58,67,68]. The NAcc in turn integrates cue–action–outcome signals across both novel and familiar cues, continuing to integrate feedback over time[12,37,49]. Critically, NAcc stimulation enhanced Pavlovian bias early in learning while simultaneously accelerating instrumental learning, such that Pavlovian tendencies were progressively overridden as action–outcome associations became more reliable.

A myriad of psychiatric disorders are characterized by a distorted balance between Pavlovian control and instrumental learning[7–11,38]. In addiction, Pavlovian approach biases toward drug-paired cues can override instrumental knowledge that continued use is harmful[8]; in anxiety, exaggerated avoidance prevents belief updating in the face of safety[11]; and in OCD, difficulty integrating corrective feedback sustains persistent and suboptimal choices[69]. Our findings speak directly to the circuitry that underpins these patterns. This division of labor — NAcc as a hub for cue-driven action control and sustained action-outcome updating, amygdala as a salience detector that flags novel cues — helps refine mechanistic accounts of how motivational biases and learning together shape behavior, and may help explain why some disorders (e.g., those involving responses to novel threat cues) particularly implicate the amygdala, whereas other (e.g., those involving cue-triggered approach) center on the NAcc. More broadly, our work highlights the translational potential of TUS. By transiently modulating deep subcortical circuits non-invasively, TUS may open a window of plasticity during which the brain is more receptive to new learning, suggesting that NAcc- or amygdala-targeted stimulation could be combined with psychological interventions to support learning when Pavlovian biases are maladaptive.

METHODS

Participants

We preregistered (https://doi.org/10.17605/OSF.IO/Q9845) collecting a target sample size of 31 healthy adults, based on a small to medium effect size of d = 0.58[45] with an alpha level of 0.05 and a power of 80% (calculated using G*Power 3.1[70]). Of these, 29 participants completed the TUS sessions and were included in the final analyses (17 females and 14 males; age range 19–35 years; mean = 24.8, SD = 4.1 years). Participants

were screened on medical history to exclude putative participants with a history of brain surgery, serious head trauma, epilepsy, convulsion, or seizure, as well as the presence of implanted metal in the head or upper body, current diagnosis of neurological or psychiatric disorders, as well as recent substance use (i.e., consumption of more than four alcoholic units in the past 24 hours or use of recreational drugs within the past 48 hours). All participants provided written informed consent in accordance with the Declaration of Helsinki. The study has been independently reviewed by the Ethics Committee Social Sciences (ECSS) of the Radboud University (ECSW-2025-021), and there were no formal objections to this study.

Study overview

The study followed a single-blind, within-subject design spanning four sessions (**Figure 1A**). Consecutive sessions were scheduled at least three days apart (maximum three weeks) to limit both ultrasound-induced plasticity carry-over and cumulative task-learning effects. The first visit combined task instruction, acquisition of structural and functional MRI data, and a brief practice round of the standard-timing task. In each of the following (intervention) sessions, participants received bilateral transcranial ultrasound stimulation (TUS; nucleus accumbens (NAcc), amygdala, or sham, order counter-balanced) immediately followed by a 30-min behavioral run of the MGNG task.

Intake session

After screening for MRI and TUS safety, the researchers went over study and task description with the participants using infographics before participants entered the scanner. They then completed a 1-hour MRI protocol (**Chapter 3**) and performed an extended-timing version of the MGNG task to acquire BOLD data (**Chapter 3**). Outside the scanner, they practiced a single run of the faster, intervention-timed task to familiarize themselves with auditory feedback and shorter trial structure used in the following TUS sessions. Informed consent was obtained prior to the start of any procedure.

Intervention sessions

Participants completed three intervention visits, each lasting up to three hours. At every visit, we first conducted TUS safety screening, then performed neuronavigation and hair preparation before administering bilateral TUS. Immediately after sonication, the participant played two successive runs of the MGNG task separated by a brief pause. Target condition (NAcc, amygdala, or sham) order was randomized and

counter-balanced across the sample to control for order, learning and fatigue effects. Laterality was alternated within participant such that the hemisphere stimulated first (left or right) changed from session to session in a predetermined counter-balancing scheme.

At the end of every session, they reported any sensations or symptoms they believed might be related to TUS; these notes served solely for monitoring. A comprehensive debriefing followed the final visit, during which participants were asked to report the difference between the final three sessions, and to guess which sessions, if any, had involved active stimulation.

Neuronavigation and hair preparation

Neuronavigation was performed with a frameless stereotactic system (Localite Biomedical Visualization Systems GmbH, Sankt Augustin, Germany) using each participant's T1w scan and individual NAcc and amygdala masks generated in FreeSurfer. Head registration relied on a reference tracker, five anatomical landmarks (nasion, bilateral eyes, bilateral ears) and between 300 – 400 scalp points. The TUS transducer was calibrated using a reference tracker and a calibration plate prior to each session.

For each session, bilateral transducer positions were defined, targeting the left and right hemisphere of the respective structures (NAcc and amygdala). These poses differed in head position, angle, and focal depth to minimize skull-induced lensing and align the beam axis with each target's long axis, thereby maximizing intracranial pressure at the intended structure. For sham visits, the transducer was placed over the brain region scheduled for the adjacent active visit (or the first target if sham was last) to preserve blinding. The transducer position at each stimulation site was recorded and used for post hoc ultrasound wave propagation simulations.

Ultrasound gel (Aquaflex Ultrasound Gel, Parker Laboratories) was applied to the participant's head over stimulation regions, followed by placement of gel pads (Aquaflex Ultrasound Gel pad, Parker Laboratories) between the gelled head and gel-covered transducer surface to eliminate air bubbles[71]. During stimulation, participants were instructed to sit still using a chinrest and look at a fixation cross on the computer screen.

TUS protocol

Ultrasound stimulation was delivered with a NeuroFUS PRO system (Brainbox Ltd., Cardiff, UK) driving a 500 kHz, four-element annular transducer DPX-500 (64 mm diameter, Sonic Concepts Inc., Bothell, WA, USA). Each hemisphere received three

40 s trains of ramped pulses (50 ms pulse duration, 10 ms ramp up/down), delivered at a 5 Hz pulse-repetition frequency. Trains were separated by 10 s cooling intervals, resulting in a total stimulation time of 150 s per hemisphere. The free-field spatial peak pulse average intensity (I_{sppa}) was set at 65 W/cm^2 (**Figure 1B**). The temperature rise (ΔT) in soft tissue remained below 2 degrees Celsius and the intracranial mechanical index for transcranial application (MI_{tc}) below 1.9, matching ITRUSST recommendations[72]. Ramped pulses were chosen to reduce auditory co-stimulation[73]. Additional masking noise (white noise and 5 Hz tone) was played via speakers during stimulation. Stimulation triggers were sent via serial port commands from a PsychoPy script (PsychoPy v2021.2.3)[74].

During NAcc stimulation, the hand-held transducer was kept stationary. For amygdala stimulation it was translated a few millimeters within the gel layer during each cooling interval to broaden focal coverage while preserving coupling.

Motivational Go/NoGo-task

Participants completed two runs of a gamified version of the MGNG task[7,75] separated by a brief pause, on a BenQ monitor (resolution, 1920 x 1080; refresh rate, 120 Hz; Qisda Corporation, Taipei, Taiwan). The task was programmed using PsychoPy 2021.2.3[74]. In the task, participants played as an astronaut whose ship approached successive "planet" stimuli, each distinguished by a unique symbol (**Figure 1C**). When the ship's hatch light illuminated the planet, participants had 1.3 s to decide whether to press the spacebar (Go) or withhold from pressing it (NoGo). They were instructed that their goal was to learn what the best response was for each planet for the most optimal outcome.

Immediately after the response window, a 1.2 s outcome screen provided probabilistic feedback. Feedback was valid on approximately 75% of trials: specifically, cues appearing 12 times received 3 invalid feedback trials, cues appearing 14 times received 4 invalid feedback trials, and cues appearing 16 times received 4 invalid feedback trials (2 pre- and 2 post-reversal). Planets appeared in either a blue "reward" galaxy or a red "punishment" galaxy, indicated by the spaceship light and planet color. A correct response on a blue planet produced a blue diamond and +10 points, whereas an incorrect action displayed an empty diamond and 0 points. On red planets, a correct action prevented loss and showed an outlined rock (0 points), whereas an incorrect action revealed a red toxic rock and −10 points. Feedback was accompanied by a brief sound (1 s); an ascending chime for reward, a neutral click for 0 points, and a low-frequency buzz for punishment. A 1 second inter-trial-interval separated trials.

Crossing required action (Go vs. NoGo) with outcome domain (Win vs. Avoid) yielded

four conditions: Go-to-Win-reward (G2W), NoGo-to-Win-reward (NG2W), Go-to-Avoid-loss (G2A), and NoGo-to-Avoid-loss (NG2A; **Figure 1D**). Each session comprised 336 trials split into two runs of 168 trials (**Suppl. Figure S1**). Within each run, participants encountered 12 distinct cues ("planets"; three per condition), which appeared 12, 14, or 16 times, respectively. For cues presented 16 times, the optimal action reversed after the eighth presentation; participants were informed that reversals could occur but not when or how often.

Trial order was generated per run using constrained semi-randomization and was structured into three consecutive sets of 56 trials, each introducing a new subset of cues (**Suppl. Figure S1**). Importantly, set boundaries were not strict: cues from the preceding set could still occur at the start of the next set, reflecting overlap in the cue schedules ("bleeding" across sets). Sets were constructed to maintain response variety (i.e., at least two required actions were represented within each set), and reversals were distributed across the run such that each set contained at least one and at most two reversal cues. Trials were additionally organized in short groups (4-trial units) used for sequence generation/shuffling; however, these units did not enforce unique cues or conditions, such that brief repetitions of the same cue and/or condition could occur (as illustrated in **Suppl. Figure S1**). This design—frequent cue switching, staggered cue introductions, and distributed reversals—kept participants in an ongoing learning regime and reduced ceiling performance[75].

The symbols consisted of artificial characters selected from the Brussels Artificial Character Sets (BACS)[76]. In addition, a subset of extra characters was created in Adobe Illustrator, matched in style to the BACS set, to ensure enough unique stimuli so that each cue was used only once across the entire study. Symbol sets were refreshed for each session but were identical across participants for a given session number, ensuring that stimulus identity was counter-balanced with the independently randomized TUS condition.

Acoustic simulations

For all participants, we performed subject-specific acoustic simulations using K-Plan, a user interface for the pseudospectral time-domain solver k-Wave[77]. Transducer locations and orientations, as recorded during each intervention session, were imported into K-Plan for post hoc modelling of the acoustic field. These simulations confirmed that the achieved focal spots closely overlapped the intended NAcc and amygdala targets (**Figure 2A**).

To quantify targeting, we overlaid each participant's simulated pressure distribution

with their individual NAcc and amygdala masks and extracted the maximum pressure within each target region and hemisphere (**Suppl. Figure S5A**). We also computed the mean pressure within each target mask (**Suppl. Figure S5B**). Finally, to assess how much of each target received putatively effective stimulation, we quantified the area within each individual NAcc and amygdala mask exceeding a conservative threshold of 0.2 MPa (**Suppl. Figure S5C**). This threshold was chosen as a plausible lower bound for neuromodulatory relevance.

Data analysis

Logistic and linear mixed-effects regression models

All mixed-effects regression analyses were performed in R (R Core Team) implemented in the lme4 package, and Type III Wald χ^2 tests from the car package for inference. Unless stated otherwise, all analyses were conducted on a trial-by-trial basis and included all trials with a registered Go or NoGo response. Categorical predictors were coded using effect (sum-to-zero) coding, such that estimated main effects reflect differences averaged across the levels of the other factors. Required action contrasted cues for which a Go response was optimal versus cues for which a NoGo response was optimal (factor required action: Go vs NoGo). Cue valence contrasted reward-predictive ("Win") versus punishment-predictive ("Avoid") cues (factor cue valence: Win vs Avoid). TUS condition contrasted sham, NAcc and amygdala as factors. For session effects, session number was included numerically (-1 0 1). To capture learning over repeated encounters with the same cue, we computed a cue-specific trial number for each presentation and standardized this variable before analysis (z-scored continuous predictor cue-specific trial number). Given the dense reciprocal projections between amygdala and NAcc within the amygdalo-striatal circuit[13,78–80], TUS effects on choice were assessed primarily using pairwise models comparing two conditions at a time; three-way models including all TUS conditions were additionally run for completeness.

Logistic mixed-effects regressions on choice

To quantify instrumental learning and motivational biases in choice behavior, we modelled the probability of emitting a Go response using mixed-effects logistic regression with a binomial link (glmer, family = binomial). Our primary model for the TUS sessions included fixed effects of required action (Go vs. NoGo), cue valence (Win vs. Avoid), cue-specific trial number (z-scored), and TUS condition (sham vs. NAcc vs. amygdala), as well as their interactions. Specifically, we included all two-way interactions among required action, cue valence, cue-specific trial number and TUS

condition, and the three-way interactions of trial number with required action × TUS condition and trial number with cue valence × TUS condition. This model was used to test whether TUS modulated instrumental learning and Pavlovian biases over time.

In a complementary analysis, we examined win-stay/lose-shift (WSLS) behavior. For each trial, we computed a WSLS regressor indexing whether the current response was consistent with repeating the previous choice after reward and switching after loss for the same cue. We then fitted a separate logistic mixed-effects regression with Go response as the dependent variable and required action, cue valence, cue-specific trial number, WSLS and TUS condition as fixed effects. This model included the required action x cue-specific trial number and WSLS x TUS condition interaction term, to assess whether TUS modulated basic WSLS tendencies.

To assess more global practice effects across repeated task exposures, we fitted an analogous model in which TUS condition was replaced by TUS session number. In this "session model", fixed effects were required action, cue valence, cue-specific trial number, and session number, together with all two-way interactions and the three-way interactions of trial number with required action × session and trial number with cue valence × session. This model was used to test whether instrumental learning and motivational biases changed systematically across sessions, irrespective of TUS.

All choice models included subject-specific random intercepts and random slopes for the fixed effects (required action, cue valence, cue-specific trial number, WSLS regressor, and, where appropriate, TUS condition or session number) and their interaction terms, using the maximal random-effects structure that allowed convergence. Significance of fixed effects was assessed using Type III Wald χ^2 tests (car package), and 95% confidence intervals (CIs) for regression coefficients were obtained from the model's profile likelihood.

Linear mixed-effects regressions on response times

To determine whether the same factors that influenced choice behavior also shaped response speed, we analyzed reaction times (RTs) with linear mixed-effects regression (lmer). Analyses were restricted to trials with a recorded Go response within the response window. For the TUS sessions, the fixed-effects structure mirrored that of the logistic model: required action (Go vs. NoGo), cue valence (Win vs. Avoid), cue-specific trial number (z-scored), and TUS condition (sham vs. NAcc vs. amygdala), plus all two-way interactions and the three-way interactions involving trial number and TUS condition (trial number × required action × TUS, trial number × cue valence × TUS). This model tested whether TUS affected learning-related speeding (differences in RTs between Go vs. NoGo cues over repetitions) and Pavlovian speeding (RT differences

between Win vs. Avoid cues). To examine session-related changes in RTs, we again fitted a parallel model in which TUS condition was replaced by TUS session number (1–3). Fixed effects were required action, cue valence, cue-specific trial number (z-scored), and session number (factor), along with all two-way interactions and the three-way interactions of trial number with required action × session and trial number with cue valence × session.

Subject-specific random intercepts and random slopes were specified in analogy to the choice models. As for the choice models, statistical significance of fixed effects was evaluated using Type III Wald χ^2 tests, and 95%-CIs for regression coefficients were derived from the fitted linear mixed-effects regression models.

Although these logistic and linear mixed-effects regressions are informative about overall learning and motivational biases, they are not ideally suited to capture the learning dynamics in this specific task. Each cue was presented only 12–16 times (with sporadic contingency reversals after the 8th repetition), so participants remained in an active learning regime and did not approach a stable asymptote for any cue. Fitting a single linear effect in such a setting effectively imposes a straight line on behavior that is changing nonlinearly from trial to trial and therefore fails to capture substantial variance in the data. In earlier versions of the task with ~40 repetitions per cue[5,46], participants' behavior approached a stable phase, and a linear approximation over trials provided a reasonable summary of performance. In contrast, with the limited and dynamically changing cue exposure used here, standard logistic and linear mixed-effects regression models are poorly matched to the underlying reinforcement-learning process. For this reason, we complement these analyses with computational reinforcement-learning models that are specifically designed to capture trial-by-trial learning dynamics and are therefore better suited to quantify how TUS affects value updating and choice.

Computational modelling

To characterize the latent learning and motivational processes underlying choice behavior, we fitted a family of hierarchical reinforcement-learning (RL) models to trial-by-trial Go/NoGo responses. Model development and validation proceeded in two steps. First, we fit the full model space (M1-M7) to an independent dataset acquired with the same task but outside the intervention sessions and used leave-one-out cross-validation (LOO) to identify the best-fitting learning architecture. Second, we applied the same model family to the TUS sessions: for transparency, all models were re-estimated and compared on the TUS data, but statistical inference on TUS effects focused on the winning model (M7), using a hierarchical parameterization with

condition-specific offsets around a sham baseline (see below).

Following previous work using similar MGNG tasks[5,7,24,46,75], in all models the action weights (w) are estimated for each response option (a) for all trials (t) per cue (s). Based on these action weights choice probabilities are computed using a softmax function, as follows:

$$p(a_t|s_t) = \left[\frac{\exp(w(a_t, s_t))}{\sum_{a'} \exp(w(a', s_t))}\right] \quad (1)$$

Model space (M1-M7)

In the simplest model (M1), the action weights are fully determined by the learned action values (Q-values). To compute the action values, we used standard delta-rule learning with two free parameters; a learning rate (ε) scaling the update term, and feedback sensitivity (ρ) scaling the outcome value (comparable to the softmax temperature).

$$Q_t(a_t, s_t) = Q_{t-1}(a_t, s_t) + \varepsilon(\rho r_t - Q_{t-1}(a_t, s_t)) \quad (2)$$

Here outcomes are reflected by r, where r∈{−1,0,1}. In our paradigm, cue valence is instructed, by means of the blue and red planets. Therefore, the initial expected outcome is 0.5 for Win cues and −0.5 for Avoid cues. Initial Q-values (Q_0) are set accordingly to ρ*0.5 for Win cues and ρ*−0.5 for Avoid cues.

M2 extended M1 with a valence-independent Go bias parameter *b* "go", added as a constant shift to the Go weight:

$$w(a, s) = \begin{cases} Q(a, s) + b & if\ a = Go \\ Q(a, s) & otherwise \end{cases} \quad (3)$$

This parameter captures a subject's general tendency to emit Go responses, over and above learned action values.

M3a further extended M2 by incorporating a Pavlovian bias parameter π, which modulates action selection based on cue valence. This parameter captures the tendency of prospective outcome valence to trigger invigoration (for rewards) and inhibition (for losses). Specifically, to invigorate action for reward-predicting stimuli (Win cues) and inhibit responses to loss-predicting stimuli (Avoid cues). The cue valence V(s) was coded as +0.5 for Win cues and −0.5 for Avoid cues. The action weights for Go responses were defined as:

$$w(a, s) = \begin{cases} Q(a, s) + b + \pi V(s) & if\ a = Go \\ Q(a, s) & otherwise \end{cases} \quad (4)$$

This addition allows the model to explain performance asymmetries that are consistent with motivational biases.

M3b decomposed the Pavlovian bias into separate parameters for Win versus Avoid cues:

$$w(a,s) = \begin{cases} Q(a,s) + \pi V_{win}(s) + \pi V_{avoid}(s) & if\ a = Go \\ Q(a,s) & otherwise \end{cases} \quad (5)$$

While being mathematically similar to M3a (**Suppl. Figure S4A**), this allows Pavlovian Go tendencies to differ between reward- and loss-predictive cues.

The next three models further elaborated the learning dynamics while retaining valence-specific Pavlovian biases.

M4 allowed valence-specific feedback sensitivities, such that prediction errors were computed with a valence-dependent outcome gain:

$$\rho_{eff,t} = \begin{cases} \rho_{win}(s) & if\ Win\ cue \\ \rho_{avoid}(s) & if\ Avoid\ cue \end{cases} \quad (6)$$

M5 instead allowed separate learning rates for Win versus Avoid cues while retaining a single feedback sensitivity. The effective learning rate depended on cue valence:

$$\varepsilon_{eff,t} = \begin{cases} \varepsilon_{win}(s) & if\ Win\ cue \\ \varepsilon_{avoid}(s) & if\ Avoid\ cue \end{cases} \quad (7)$$

M6 implemented asymmetric learning from positive versus negative prediction errors. Here, the learning rate depended on the sign of the prediction error:

$$\varepsilon_{eff,t} = \begin{cases} \varepsilon_{pos}(s) & if\ PE_t \geq 0 \\ \varepsilon_{neg}(s) & if\ PE_t < 0 \end{cases} \quad (8)$$

with a single feedback sensitivity. This allowed participants to learn more from better-than-expected than from worse-than-expected outcomes (or vice versa).

Finally, M7 combined PE-sign-specific learning with an additional distinction between first and subsequent cue presentations, motivated by novelty-sensitive learning. For each cue, we tracked a cue-specific repetition counter, indicating how often that cue had been encountered within the run. Prediction errors were computed as in M6, but the learning rate depended both on whether the cue was novel and on PE sign: For the first presentation of a cue a dedicated learning rate was applied, irrespective of the sign of the prediction error. For later presentations, separate learning rates were used for positive versus negative PEs:

$$\varepsilon_{eff,t} = \begin{cases} \varepsilon_1(s) & if\, cue\ count\ =\ 1 \\ \varepsilon_{rest_pos}(s) & if\ cue\ count\ >\ 1\ and\ PE_t\ \geq 0 \\ \varepsilon_{rest_neg}(s) & if\, cue\ count\ >1\ and\ \ PE_t\ <0 \end{cases} \quad (9)$$

M7 thus captured (i) a distinct, novelty-sensitive learning component at first cue exposure, (ii) differentiated learning from better- versus worse-than-expected outcomes thereafter, and (iii) valence-specific Pavlovian influences on Go probability.

Hierarchical Bayesian estimation and TUS parameterization

All RL models were fitted in a hierarchical Bayesian framework. Subject-level parameters were assumed to arise from Gaussian group-level distributions on the latent scale (non-centered parameterization), with weakly informative priors on group-level means and standard deviations to regularize estimation while remaining agnostic about the exact parameter values.

To exploit the within-subject TUS design, we parameterized TUS effects as condition-specific offsets on each RL parameter. For each parameter θ (e.g., ε_1, ε_{pos}, ε_{neg}, ρ, π_{win}, π_{avoid}) and each subject p, we first defined a subject-specific baseline (sham) value

$$\theta_{p,\mathrm{sham}} = \mu_\theta + \kappa_{p,\theta} \quad (10)$$

and then expressed the TUS conditions as deviations from this baseline with both group-level offsets and condition-specific random effects:

$$\begin{aligned} \theta_{p,\mathrm{NAcc}} &= \theta_{p,\mathrm{sham}} + \Delta_\theta^{\mathrm{NAcc}} + \eta_{p,\theta}^{\mathrm{NAcc}}, \\ \theta_{p,\mathrm{AMY}} &= \theta_{p,\mathrm{sham}} + \Delta_\theta^{\mathrm{AMY}} + \eta_{p,\theta}^{\mathrm{AMY}} \end{aligned} \quad (11)$$

Here, $\mu\theta$ is the group-level baseline (sham) mean, $\kappa_{p,\theta}$ is a subject-specific random effect shared across conditions, and $\Delta_\theta^{\mathrm{NAcc}}$ and $\Delta_\theta^{\mathrm{AMY}}$ are group-level offsets capturing NAcc–sham and AMY–sham differences, respectively. The additional terms $\eta_{p,\theta}^{\mathrm{NAcc}}$ and $\eta_{p,\theta}^{\mathrm{AMY}}$ are condition-specific random effects that allow individual subjects to deviate from the group-level TUS effect in the NAcc and AMY sessions. By definition, sham served as the reference condition ($\Delta_\theta^{sham} = 0$). This parameterization ties the condition-specific parameters within subjects together and yields direct within-subject contrasts of NAcc and AMY stimulation against sham for each latent learning and bias parameter.

Sampling was performed using Hamiltonian Monte Carlo (Stan), running multiple chains with random, over-dispersed initial values (Stan's default random initialization on the unconstrained scale). Convergence was assessed by visual inspection of trace plots and by requiring potential scale reduction factors $\hat{R} < 1.01$ for all group- and subject-level parameters, as well as effective sample sizes > 1000 wherever possible. Posterior summaries (means, medians, 95% credible intervals) were obtained from

the joint posterior after discarding warm-up samples. To statistically compare parameter estimates between conditions, paired-samples t-tests were conducted on the participant-level parameter estimates. For effect size reporting, Cohen's D was calculated for paired-samples comparison by dividing the mean difference by the standard deviation of the within-subject difference scores.

To compare models in terms of their ability to generalize to new data, we computed approximate leave-one-out cross-validated (LOO) expected log pointwise predictive density (elpd) for each model using Pareto-smoothed importance sampling (PSIS-LOO). For each fitted model, we extracted trial-wise log-likelihoods from the posterior and used the loo package in R to obtain $elpd_{loo}$ and associated standard errors. Models were ranked by $elpd_{loo}$; pairwise differences in $elpd_{loo}$ and their standard errors were used to quantify the evidence in favor of the winning model relative to simpler alternatives. LOO is a fully Bayesian model selection metric that makes use of the entire posterior distribution of the model parameters, rather than relying on single point estimates. This allows the incorporation of uncertainty in the parameter estimates and any correlations between parameters when evaluating model fit. Unlike AIC or BIC, which count all free parameters equally, LOO computes the effective number of parameters, which reflects how much the parameters affect the model complexity. In hierarchical models, some parameters are partially pooled or strongly correlated, so their individual influence on predictions is smaller than if they were independent. By using the effective number of parameters rather than the raw count, WAIC and LOOIC avoid over-penalizing models for parameters that do not fully act independently, resulting in a more accurate assessment of model fit.

Posterior predictive checks were performed for all models, with a focus on the best-fitting model (M7). For each posterior draw, we simulated synthetic choice data using the participant-specific parameter estimates and the actual sequence of cues and outcomes. From these simulations, we derived predicted Go probabilities and learning curves for each condition (G2W, NG2W, G2A, NG2A), as well as win-stay/lose-shift behavior over cue repetitions. Model fit was assessed by comparing these posterior predictive summaries to the empirical data (**Suppl. Figure S4E-H**). Parameter estimates for all models and full model comparison metrics are reported in the supplementary information (**Tables S1-S2**).

REFERENCES

1. Boureau, Y.-L., Sokol-Hessner, P. & Daw, N. D. Deciding How To Decide: Self-Control and Meta-Decision Making. *Trends Cogn. Sci.* ***19***, 700–710 (2015).
2. Dayan, P., Niv, Y., Seymour, B. & Daw, N. D. The misbehavior of value and the discipline of the will. *Neural Netw. Off. J. Int. Neural Netw. Soc.* ***19***, 1153–1160 (2006).
3. Moutoussis, M. *et al.* Change, stability, and instability in the Pavlovian guidance of behaviour from adolescence to young adulthood. *PLOS Comput. Biol.* ***14***, e1006679 (2018).
4. Algermissen, J., Swart, J. C., Scheeringa, R., Cools, R. & den Ouden, H. E. M. Striatal BOLD and midfrontal theta power express motivation for action. *Cereb. Cortex* ***32***, 2924–2942 (2022).
5. Algermissen, J., Swart, J. C., Scheeringa, R., Cools, R. & den Ouden, H. E. M. Prefrontal signals precede striatal signals for biased credit assignment in motivational learning biases. *Nat. Commun.* ***15***, 19 (2024).
6. Cools, R., Nakamura, K. & Daw, N. D. Serotonin and dopamine: unifying affective, activational, and decision functions. *Neuropsychopharmacol. Off. Publ. Am. Coll. Neuropsychopharmacol.* ***36***, 98–113 (2011).
7. Scholz, V. *et al.* Cortical dopamine reduces the impact of motivational biases governing automated behaviour. *Neuropsychopharmacol. Off. Publ. Am. Coll. Neuropsychopharmacol.* ***47***, 1503–1512 (2022).
8. Garbusow, M. *et al.* Pavlovian-to-instrumental transfer effects in the nucleus accumbens relate to relapse in alcohol dependence. *Addict. Biol.* ***21***, 719–731 (2016).
9. Huys, Q. J. M. *et al.* The specificity of Pavlovian regulation is associated with recovery from depression. *Psychol. Med.* ***46***, 1027–1035 (2016).
10. Milner, A., Witt, K., LaMontagne, A. D. & Niedhammer, I. Psychosocial job stressors and suicidality: a meta-analysis and systematic review. *Occup. Environ. Med.* ***75***, 245–253 (2018).
11. Mkrtchian, A., Aylward, J., Dayan, P., Roiser, J. P. & Robinson, O. J. Modeling Avoidance in Mood and Anxiety Disorders Using Reinforcement Learning. *Biol. Psychiatry* ***82***, 532–539 (2017).
12. Averbeck, B. B. & Costa, V. D. Motivational neural circuits underlying reinforcement learning. *Nat. Neurosci.* ***20***, 505–512 (2017).
13. Cardinal, R. N., Parkinson, J. A., Hall, J. & Everitt, B. J. Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neurosci. Biobehav. Rev.* ***26***, 321–352 (2002).
14. Haber, S. N. & Knutson, B. The Reward Circuit: Linking Primate Anatomy and Human Imaging. *Neuropsychopharmacology* ***35***, 4–26 (2010).
15. Corbit, L. H. & Balleine, B. W. The general and outcome-specific forms of Pavlovian-instrumental transfer are differentially mediated by the nucleus accumbens core and shell. *J. Neurosci. Off. J. Soc. Neurosci.* ***31***, 11786–11794 (2011).

16. Geurts, D. E. M., Huys, Q. J. M., den Ouden, H. E. M. & Cools, R. Aversive Pavlovian Control of Instrumental Behavior in Humans. *J. Cogn. Neurosci.* ***25***, 1428–1441 (2013).
17. Guitart-Masip, M. *et al.* Action controls dopaminergic enhancement of reward representations. *Proc. Natl. Acad. Sci. U. S. A.* ***109***, 7511–7516 (2012).
18. Corbit, L. H. & Balleine, B. W. Double dissociation of basolateral and central amygdala lesions on the general and outcome-specific forms of pavlovian-instrumental transfer. *J. Neurosci. Off. J. Soc. Neurosci.* ***25***, 962–970 (2005).
19. Hall, J., Parkinson, J. A., Connor, T. M., Dickinson, A. & Everitt, B. J. Involvement of the central nucleus of the amygdala and nucleus accumbens core in mediating Pavlovian influences on instrumental behaviour. *Eur. J. Neurosci.* ***13***, 1984–1992 (2001).
20. Wyvell, C. L. & Berridge, K. C. Intra-accumbens amphetamine increases the conditioned incentive salience of sucrose reward: enhancement of reward 'wanting' without enhanced 'liking' or response reinforcement. *J. Neurosci. Off. J. Soc. Neurosci.* ***20***, 8122–8130 (2000).
21. Knutson, B., Adams, C. M., Fong, G. W. & Hommer, D. Anticipation of Increasing Monetary Reward Selectively Recruits Nucleus Accumbens. *J. Neurosci.* ***21***, RC159–RC159 (2001).
22. Sescousse, G., Li, Y. & Dreher, J.-C. A common currency for the computation of motivational values in the human striatum. *Soc. Cogn. Affect. Neurosci.* ***10***, 467–473 (2015).
23. Guitart-Masip, M. *et al.* Action Dominates Valence in Anticipatory Representations in the Human Striatum and Dopaminergic Midbrain. *J. Neurosci.* ***31***, 7867–7875 (2011).
24. Guitart-Masip, M. *et al.* Go and no-go learning in reward and punishment: interactions between affect and effect. *NeuroImage* ***62***, 154–166 (2012).
25. Rescorla, R. A. & Wagner, A. R. A Theory of Pavlovian Conditioning Variations in the Effectiveness of Reinforcement and Nonreinforcement. *Class. Cond. Curr. Res. Theory* ***21***, 64–99 (1972).
26. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* ***275***, 1593–1599 (1997).
27. Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. Model-based influences on humans' choices and striatal prediction errors. *Neuron* ***69***, 1204–1215 (2011).
28. O'Doherty, J. *et al.* Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* ***304***, 452–454 (2004).
29. Balleine, B. W. & Killcross, S. Parallel incentive processing: an integrated view of amygdala function. *Trends Neurosci.* ***29***, 272–279 (2006).
30. Holland, P. C. & Gallagher, M. Double dissociation of the effects of lesions of basolateral and central amygdala on conditioned stimulus-potentiated feeding and Pavlovian-instrumental transfer. *Eur. J. Neurosci.* ***17***, 1680–1694 (2003).
31. LeDoux, J. E. Emotion circuits in the brain. *Annu. Rev. Neurosci.* ***23***, 155–184 (2000).
32. Phelps, E. A. & LeDoux, J. E. Contributions of the amygdala to emotion processing: from animal models to human behavior. *Neuron* ***48***, 175–187 (2005).
33. Lichtenberg, N. T. *et al.* Basolateral Amygdala to Orbitofrontal Cortex Projections Enable Cue-Triggered Reward Expectations. *J. Neurosci.* ***37***, 8374–8384 (2017).

34. Paton, J. J., Belova, M. A., Morrison, S. E. & Salzman, C. D. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* ***439***, 865–870 (2006).
35. Averbeck, B. B. Amygdala and ventral striatum population codes implement multiple learning rates for reinforcement learning. in *2017 IEEE Symposium Series on Computational Intelligence (SSCI)* 1–5 (2017). doi:10.1109/SSCI.2017.8285354.
36. Hazy, T. E., Frank, M. J. & O'Reilly, R. C. Neural mechanisms of acquired phasic dopamine responses in learning. *Neurosci. Biobehav. Rev.* ***34***, 701–720 (2010).
37. Neftci, E. O. & Averbeck, B. B. Reinforcement learning in artificial and biological systems. *Nat. Mach. Intell.* ***1***, 133–143 (2019).
38. Schad, D. J. *et al.* Dissociating neural learning signals in human sign- and goal-trackers. *Nat. Hum. Behav.* ***4***, 201–214 (2020).
39. Algermissen, J. *et al.* Low-intensity focused ultrasound to human amygdala reveals a causal role in ambiguous emotion processing and alters local and network-level activity. 2025.08.14.670358 Preprint at https://doi.org/10.1101/2025.08.14.670358 (2025).
40. Butler, C. R. *et al.* Transcranial ultrasound stimulation to human middle temporal complex improves visual motion detection and modulates electrophysiological responses. *Brain Stimulat.* ***15***, 1236–1245 (2022).
41. Farboud, S. *et al.* Rapid modulation of choice behavior by ultrasound on the human frontal eye fields. *Nat. Commun.* https://doi.org/10.1038/s41467-026-69854-7 (2026) doi:10.1038/s41467-026-69854-7.
42. Koutsoumpari, N. *et al.* Ultrasound neuromodulation reveals distinct roles of the dorsal anterior cingulate cortex and anterior insula in Pavlovian biases. 2025.06.12.659273 Preprint at https://doi.org/10.1101/2025.06.12.659273 (2025).
43. Meijer, S. *et al.* The human amygdala in threat learning and extinction. *Sci. Adv.* https://doi.org/10.1126/sciadv.aea8233 (2026) doi:10.1126/sciadv.aea8233.
44. Nakajima, K. *et al.* A causal role of anterior prefrontal-putamen circuit for response inhibition revealed by transcranial ultrasound stimulation in humans. *Cell Rep.* ***40***, 111197 (2022).
45. Yaakub, S. N. *et al.* Non-invasive ultrasonic neuromodulation of the human nucleus accumbens impacts reward sensitivity. *Nat. Commun.* ***16***, 10192 (2025).
46. Swart, J. C. *et al.* Catecholaminergic challenge uncovers distinct Pavlovian and instrumental mechanisms of motivated (in)action. *eLife* ***6***, e22169 (2017).
47. Talmi, D., Seymour, B., Dayan, P. & Dolan, R. J. Human pavlovian-instrumental transfer. *J. Neurosci. Off. J. Soc. Neurosci.* ***28***, 360–368 (2008).
48. Haber, S. N. & Behrens, T. E. J. The Neural Network Underlying Incentive-Based Learning: Implications for Interpreting Circuit Disruptions in Psychiatric Disorders. *Neuron* ***83***, 1019–1039 (2014).

49. Costa, V. D., Dal Monte, O., Lucas, D. R., Murray, E. A. & Averbeck, B. B. Amygdala and Ventral Striatum Make Distinct Contributions to Reinforcement Learning. *Neuron* ***92***, 505–517 (2016).
50. Daw, N. D., Niv, Y. & Dayan, P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* ***8***, 1704–1711 (2005).
51. Bradley, C., Nydam, A. S., Dux, P. E. & Mattingley, J. B. State-dependent effects of neural stimulation on brain function and cognition. *Nat. Rev. Neurosci.* ***23***, 459–475 (2022).
52. Guerra, A., López-Alonso, V., Cheeran, B. & Suppa, A. Variability in non-invasive brain stimulation studies: Reasons and results. *Neurosci. Lett.* ***719***, 133330 (2020).
53. Huang, Y.-Z. *et al.* Plasticity induced by non-invasive transcranial brain stimulation: A position paper. *Clin. Neurophysiol.* ***128***, 2318–2329 (2017).
54. Belova, M. A., Paton, J. J., Morrison, S. E. & Salzman, C. D. Expectation modulates neural responses to pleasant and aversive stimuli in primate amygdala. *Neuron* ***55***, 970–984 (2007).
55. Blackford, J. U., Buckholtz, J. W., Avery, S. N. & Zald, D. H. A unique role for the human amygdala in novelty detection. *NeuroImage* ***50***, 1188–1193 (2010).
56. Pedersen, W. S. *et al.* The effects of stimulus novelty and negativity on BOLD activity in the amygdala, hippocampus, and bed nucleus of the stria terminalis. *Soc. Cogn. Affect. Neurosci.* ***12***, 748–757 (2017).
57. Janak, P. H. & Tye, K. M. From circuits to behaviour in the amygdala. *Nature* ***517***, 284–292 (2015).
58. Likhtik, E. & Paz, R. Amygdala-prefrontal interactions in (mal)adaptive learning. *Trends Neurosci.* ***38***, 158–166 (2015).
59. Ebrahimi, C. *et al.* Elevated Amygdala Responses During De Novo Pavlovian Conditioning in Alcohol Use Disorder Are Associated With Pavlovian-to-Instrumental Transfer and Relapse Latency. *Biol. Psychiatry Glob. Open Sci.* ***3***, 803–813 (2023).
60. Hatfield, T., Han, J.-S., Conley, M., Gallagher, M. & Holland, P. Neurotoxic Lesions of Basolateral, But Not Central, Amygdala Interfere with Pavlovian Second-Order Conditioning and Reinforcer Devaluation Effects. *J. Neurosci.* ***16***, 5256–5265 (1996).
61. Prévost, C., Liljeholm, M., Tyszka, J. M. & O'Doherty, J. P. Neural correlates of specific and general Pavlovian-to-Instrumental Transfer within human amygdalar subregions: a high-resolution fMRI study. *J. Neurosci. Off. J. Soc. Neurosci.* ***32***, 8383–8390 (2012).
62. Atlas, L. Y., Doll, B. B., Li, J., Daw, N. D. & Phelps, E. A. Instructed knowledge shapes feedback-driven aversive learning in striatum and orbitofrontal cortex, but not the amygdala. *eLife* ***5***, e15192 (2016).
63. Jerusalem, A. *et al.* Electrophysiological-mechanical coupling in the neuronal membrane and its role in ultrasound neuromodulation and general anaesthesia. *Acta Biomater.* ***97***, 116–140 (2019).

64. Yoo, S., Mittelstein, D. R., Hurt, R. C., Lacroix, J. & Shapiro, M. G. Focused ultrasound excites cortical neurons via mechanosensitive calcium accumulation and ion channel amplification. *Nat. Commun.* ***13***, 493 (2022).
65. Mateos-Aparicio, P. & Rodríguez-Moreno, A. Calcium Dynamics and Synaptic Plasticity. *Adv. Exp. Med. Biol.* ***1131***, 965–984 (2020).
66. Malenka, R. C. & Bear, M. F. LTP and LTD: An Embarrassment of Riches. *Neuron* ***44***, 5–21 (2004).
67. Balderston, N. L., Schultz, D. H. & Helmstetter, F. J. The Effect of Threat on Novelty Evoked Amygdala Responses. *PLOS ONE* ***8***, e63220 (2013).
68. Murray, R. J., Brosch, T. & Sander, D. The functional profile of the human amygdala in affective processing: insights from intracranial recordings. *Cortex J. Devoted Study Nerv. Syst. Behav.* ***60***, 10–33 (2014).
69. Robbins, T. W., Vaghi, M. M. & Banca, P. Obsessive-Compulsive Disorder: Puzzles and Prospects. *Neuron* ***102***, 27–47 (2019).
70. Faul, F., Erdfelder, E., Buchner, A. & Lang, A.-G. Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behav. Res. Methods* ***41***, 1149–1160 (2009).
71. Murphy, K. R. *et al.* A practical guide to transcranial ultrasonic stimulation from the IFCN-endorsed ITRUSST consortium. *Clin. Neurophysiol.* ***171***, 192–226 (2025).
72. Aubry, J.-F. *et al.* ITRUSST consensus on biophysical safety for transcranial ultrasound stimulation. *Brain Stimulat.* ***18***, 1896–1905 (2025).
73. Mohammadjavadi, M. *et al.* Elimination of peripheral auditory pathway activation does not affect motor responses from ultrasound neuromodulation. *Brain Stimulat.* ***12***, 901–910 (2019).
74. Peirce, J. *et al.* PsychoPy2: Experiments in behavior made easy. *Behav. Res. Methods* ***51***, 195–203 (2019).
75. Zorowitz, S., Karni, G., Paredes, N., Daw, N. & Niv, Y. Improving the Reliability of the Pavlovian Go/No-Go Task for Computational Psychiatry Research. *Comput. Psychiatry* ***9***, (2025).
76. Vidal, C., Content, A. & Chetail, F. BACS: The Brussels Artificial Character Sets for studies in cognitive psychology and neuroscience. *Behav. Res. Methods* ***49***, 2093–2112 (2017).
77. Treeby, B. E. & Cox, B. T. k-Wave: MATLAB toolbox for the simulation and reconstruction of photoacoustic wave fields. *J. Biomed. Opt.* ***15***, 021314 (2010).
78. Brog, J. S., Salyapongse, A., Deutch, A. Y. & Zahm, D. S. The patterns of afferent innervation of the core and shell in the "Accumbens" part of the rat ventral striatum: Immunohistochemical detection of retrogradely transported fluoro-gold. *J. Comp. Neurol.* ***338***, 255–278 (1993).
79. Groenewegen, H. J., Wright, C. I., Beijer, A. V. & Voorn, P. Convergence and segregation of ventral striatal inputs and outputs. *Ann. N. Y. Acad. Sci.* ***877***, 49–63 (1999).

80. Kelley, A. E., Domesick, V. B. & Nauta, W. J. The amygdalostriatal projection in the rat--an anatomical study by anterograde and retrograde tracing methods. *Neuroscience* ***7***, 615–630 (1982).
81. Cavanagh, J. F., Eisenberg, I., Guitart-Masip, M., Huys, Q. & Frank, M. J. Frontal Theta Overrides Pavlovian Learning Biases. *J. Neurosci.* ***33***, 8541–8548 (2013).
82. Guitart-Masip, M. *et al.* Differential, but not opponent, effects of l-DOPA and citalopram on action learning with reward and punishment. *Psychopharmacology (Berl.)* ***231***, 955–966 (2014).
83. van Nuland, A. J. *et al.* Effects of dopamine on reinforcement learning in Parkinson's disease depend on motor phenotype. *Brain* ***143***, 3422–3434 (2020).

SUPPLEMENTARY INFORMATION

SUPPLEMENTARY SECTIONS

S.1 Regression analyses of reaction times (sham behavior)

To characterize how reaction times tracked learning and motivational biases during sham sessions, we analyzed reaction times (RTs) with a linear mixed-effects regression including required action, cue valence, cue-specific trial number (z-scored), and all two-way interactions. Note that this analysis is restricted to trials on which participants made a Go response; RT results should therefore be interpreted with this reduction in power in mind. Participants responded faster for correct than incorrect Go responses (required action: b = -0.02, 95%-CI [-0.03, -0.01], $\chi^2(1)$ = 13.7, p = 2.0e-04), and this difference in RT increased over cue repetitions (required action × trial number: b = -0.01, 95%-CI [-0.015, -0.002], $\chi^2(1)$ = 7.1, p = .008), again corroborating learning. They were also faster responding to Win than Avoid cues (cue valence: b = -0.03, 95%-CI [-0.03, -0.02], $\chi^2(1)$ = 34.8, p = 3.66e-09), reflecting Pavlovian bias. This motivational speeding did not change significantly across the task (b = -0.003, 95%-CI [-0.01, 0.01], $\chi^2(1)$ = 0.5, p = .5). Finally, there was no evidence that cue valence differentially modulated correct versus incorrect response speed (cue valence × required action: b = 0.00, 95%-CI -0.01, 0.01], $\chi^2(1)$ = 0.0, p > .9).

S.2 Regression analyses of behavior (session number effects)

Because participants completed the MGNG task five times over the course of the study – during the intake session once (two runs) in the scanner with fMRI-adjusted slower and jittered timings, immediately followed by one run outside the scanner with behavioral task timings (**Figure S2A**), and three times during the three TUS sessions – we next examined whether performance changed systematically across sessions. To this end, we fitted the logistic mixed-effects regression on Go responses and the linear mixed-effects regression on RTs, adding a variable session number, coded zero-sum and linearly ([-1 0 1] for sessions 1, 2, 3), which was allowed to interact with all other predictors and interactions.

Accuracy improved across sessions, as participants increasingly aligned their choices with the required action (required action × session number: b = 0.10, 95%-CI [0.04, 0.17], χ^2 = 9.5, p = .002; **Figure S2B**, **C**). Moreover, participants also learned more rapidly within each session (required action × session number × trial number: b = -0.11

(for NoGo) and 0.06 (for Go), χ^2 = 16.7, p < 2.4e-04; **Figure S2B**, **C**), indicating that the build-up of instrumental Go/NoGo discrimination over cue repetitions became steeper in later sessions. At the same time, the Pavlovian bias showed a modest attenuation across sessions (cue valence × session number: b = -0.05, 95%-CI [-0.9, -0.01], χ^2 = 6.0, p = .014; **Figure S2B**, **C**). However, this bias did not decay significantly more quickly within sessions (cue valence × session number × trial number: b = -0.01, 95%-CI [-0.04, 0.06], χ^2 = 0.2, p = .7; **Figure S2B**, **C**), in line with the absence of a significant decay of bias across all sessions (cue valence x trial number = b = -0.03, 95%-CI [-0.14, 0.08], $\chi^2(1)$ = 0.3, p = .6). Thus, across sessions the bias was reduced, but there was no change in within-session dynamic. Crucially, even in the final session participants still expressed a robust Pavlovian bias (main effect of cue valence in session 3: b = 0.20, 95%-CI [0.13, 0.27], $\chi^2(1)$ = 28.4, p = 9.7e-08; **Figure S2B**, **right panel**), confirming that key motivational and learning signatures remained present throughout the experiment and could still be modulated by TUS. Importantly, TUS session order was randomized and counterbalanced across participants, to cancel out any session order effects.

While overall accuracy did improve across sessions, there was no evidence for changes in learning-related speeding across sessions (required action × session number: b = 0.003, 95%-CI [-0.004, 0.010], χ^2 = 0.7, p = .4; **Figure S2D**). Likewise, Pavlovian speeding did not show reliable session dependence (cue valence × session number: b = -0.001, 95%-CI [-0.010, 0.006], χ^2 = 0.1, p = .7; **Figure S2D**). Instead, there was a non-significant trend towards faster responding with increasing session number, irrespective of cue valence or required action (session number: b = -0.01, 95%-CI [-0.029, 0.001], χ^2 = 3,2, p = .08; **Figure S2D**), putatively consistent with nonspecific practice effects.

S.3 Independent validation of reinforcement learning models

We first fitted the full family of reinforcement learning models to an independent dataset acquired with the same task in the same participants, but from a separate session that did not include any TUS or sham stimulation. In the simplest model (M1), choice behavior was driven solely by instrumental action values, with a single learning rate (ε) and feedback sensitivity (ρ) per participant[25]. M2 extended this baseline with a valence-unspecific Go bias (b), capturing a general tendency to respond or hold back. Explicitly signaled cue valence affected Go responses already on the first trial (**Figure 1F**), prior to any experience with a specific cue–outcome combinations. To capture this Pavlovian response bias, cue values were added to the action weights in M3[24,81,82]. In M3a this was implemented as a single Pavlovian bias parameter (π) that increased Go weights for Win cues and decreased them for Avoid cues. In M3b, Pavlovian biases were separated by valence (π_{win}, π_{avoid}), allowing for distinct Go tendencies towards reward-predictive versus punishment-predictive cues. While M3a and M3b capture identical

variance within a session, as they both capture the Go tendencies separately for each valence (**Suppl. Figure S4A**). However, across sessions TUS may selectively affect each of the 4 parameters, i.e., Go responding for Win cues, Go responding for Avoid, their difference, or mean Go responding. M4–M6 then introduced additional specificity in the learning dynamics. M4 allowed valence-specific feedback sensitivities (ρ_{win}, ρ_{avoid}); M5 implemented separate learning rates for Win versus Avoid cues (ε_{win}, ε_{avoid}); and M6 implemented asymmetric learning from positive versus negative prediction errors (ε_{pos}, ε_{neg}).

In this dataset, model comparison based on approximate leave-one-out cross-validation (LOO) showed that the best-fitting model among M1–M6 was M6, which implemented separate learning rates for positive versus negative prediction errors. This pattern differs from several previous applications of this task, where models with valence-specific learning or feedback parameters often performed best[5,7,46,83]. One plausible explanation lies in the learning regime induced by our design. In earlier studies, participants typically learned a small set of cues with many repetitions, such that behavior approached a relatively stable asymptote and learning rates could effectively track cue valence. By contrast, in the present paradigm participants encountered many different cues with comparatively few repetitions and continuous replacement, so that learning remained in a more dynamic, non-asymptotic regime. Under such conditions, the learning rate may be driven primarily by outcome statistics (better- vs. worse-than-expected outcomes), making PE-sign-specific learning particularly important. In this sense, our finding that M6 outperforms valence-based learning models aligns with work on continuous learning tasks such as those studied by [75], from which the present paradigm is derived.

Building on this, we next asked whether distinguishing novelty-related learning from later updates would further improve model fit. Empirical behavior suggested that choice updating differed between the first and subsequent cue encounters (**Suppl. Figure S4B**), motivating an extension of M6 with a separate learning rate for the first encounter with each cue (ε_1; M7). This model indeed provided the best overall account of behavior (M7 > M6 > remaining models; **Suppl. Figure 4C**). Group-level posterior densities for the parameters of M7 are shown in **Suppl. Figure S4D**. At the group level, the learning rate for first cue presentations was moderate (mean $\varepsilon_1 \approx 0.34$), whereas learning from later positive prediction errors was relatively fast (ε_rest_pos ≈ 0.57) and learning from later negative prediction errors slower (ε_rest_neg ≈ 0.26). The comparison between ε_{rest_pos} and ε_{rest_neg} is meaningful, as both are estimated from the same pool of later trials and are conditioned on prediction-error sign. By contrast, ε_1 should not be directly interpreted relative to these "rest" learning rates: it is estimated from substantially fewer updates (one first-presentation trial per cue)

and, by design, pools both positive and negative prediction errors into a single first-exposure parameter. Feedback sensitivity was well above zero ($\rho \approx 3.33$), indicating that outcomes strongly influenced prediction errors and value updates. Importantly, the Pavlovian bias parameters suggested a positive Go tendency for reward-predictive cues ($\pi_{win} \approx 0.21$) and a negative Go tendency for punishment-predictive cues ($\pi_{avoid} \approx -0.29$), consistent with the observed motivational bias in the empirical data. Parameter estimates for all models are reported in **Suppl. Table S1**.

Posterior-predictive checks confirmed that M7 captured the key qualitative features of the behavior. Simulated choice data closely reproduced the overall pattern of Go probabilities across the four conditions, including higher Go rates for Go- than NoGo-required cues and the characteristic Pavlovian advantage for G2W and NG2A over the incongruent conditions (**Suppl. Figure S4E**). The model also recovered the empirical learning curves over trials within each cue type (**Suppl. Figure S4F**), the reversal learning dynamics (**Suppl. Figure S4G**), and the observed win-stay/lose-shift tendencies in choice repetition (**Suppl. Figure S4H**). Together, these results indicate that participants' performance in the task is well described by a reinforcement-learning process with multiple learning rates, including a distinct first-exposure component, and valence-specific Pavlovian control over Go responding.

S.4 Regression analyses of behavior (TUS effects)

To quantify TUS effects of NAcc and amygdala stimulation on motivational and learning-related influences on choice, we first fitted a logistic mixed-effects regression to Go/NoGo responses. The model included four fixed factors: required action (Go vs. NoGo), cue valence (Win vs. Avoid), cue-specific trial number (z-scored), and TUS condition (sham vs. NAcc vs. amygdala), as well as all two-way and relevant three-way interactions among these factors. Given that NAcc and amygdala are densely interconnected within the amygdalo-striatal circuit[13,78–80], and that three-way contrasts may mask pairwise differences, we also ran all models in pairwise contrasts including only two TUS conditions at a time (**Suppl. Table S3**).

Across sessions, participants showed clear instrumental learning. They emitted more Go responses to cues that required a Go response than to cues that required a NoGo response (required action: $b = 0.8$, 95%-CI [0.7, 0.9], $\chi^2(1) = 167.4$, $p < 2.2e\text{-}16$; **Suppl. Figure S3A**), and this advantage increased with repeated presentations of the same cue (required action × trial number: $b = 0.39$, 95%-CI [0.29, 0.50], $\chi^2(1) = 52.5$, $p = 4.2e\text{-}13$; **Suppl. Figure S3B-C**). In this logistic model, however, the magnitude of instrumental learning did not differ between TUS conditions (required action × TUS condition: $\chi^2(2) = 1.7$, $p = .4$), nor did it show TUS modulation of changes in instrumental learning over

time (required action × TUS condition × trial number: $\chi^2(2) = 0.7$, $p > .9$). Choices were also shaped by a robust Pavlovian bias. Go responses were more frequent for Win than for Avoid cues (main effect of cue valence: $b = 0.22$, 95%-CI [0.16, 0.28], $\chi^2(1) = 48.0$, $p = 4.3e\text{-}12$). The magnitude of this motivational bias remained stable across repetitions of the same cue (cue valence × trial number: $b = -0.01$, 95%-CI [-0.04, 0.03], $\chi^2(1) = 0.1$, $p = .8$). Thus, our modified cue-presentation scheme successfully prevented the motivational bias from dissipating once participants had learned the optimal actions. We did not, however, observe TUS effects on motivational bias, either overall (cue valence × TUS condition: $\chi^2(2) = 0.6$, $p = .8$) or in its evolution over time (cue valence × TUS condition × trial number: $\chi^2(2) = 4.5$, $p = .11$).

The pairwise models similarly showed no TUS effect on overall instrumental learning (required action x TUS stimulation [NAcc - sham]: $b = -0.01$, 95%-CI [-0.09, 0.08], $\chi^2(1) = 0.02$, $p = .9$; [amygdala - sham]: $b = 0.03$, 95%-CI [-0.05, 0.12], $\chi^2(1) = 0.12$, $p = .5$; [NAcc - amygdala]: $b = -0.04$, 95%-CI [-0.09, 0.02], $\chi^2(1) = 1.9$, $p = .2$) or its development over time (required action x trial number x TUS stimulation [NAcc - sham]: $b = 0.01$, 95%-CI [-0.09, 0.11], $\chi^2(1) = 0.9$, $p = .6$; [amygdala - sham]: $b = -0.01$, 95%-CI [-0.11, 0.08], $\chi^2(1) = 0.7$, $p = .7$; [NAcc - amygdala]: $b = -0.02$, 95%-CI [-0.12, 0.08], $\chi^2(1) = 0.3$, $p = .9$). Pairwise contrasts also did not reveal a main effect of cue valence as a function of TUS condition (cue valence x TUS condition [NAcc – sham]: $b = 0.02$, 95%-CI [-0.02, 0.06], $\chi^2(1) = 0.6$, $p = .4$; [amygdala – sham]: $b = 0.01$, 95%-CI [-0.03, 0.06], $\chi^2(1) = 0.2$, $p = .6$; [NAcc – amygdala]: $b = 0.01$, 95%-CI [-0.03, 0.04], $\chi^2(1) = 0.1$, $p = .8$), suggesting the gross Win-Avoid difference in Go probability did not differ reliably between conditions in an overall sense. However, pairwise contrasts did reveal that NAcc-induced Pavlovian bias effects were time-varying: under NAcc stimulation, the Win-Avoid difference in Go probability was particularly pronounced early in learning and decreased more steeply across cue repetitions than under sham (cue valence x trial number x TUS condition [NAcc - sham]: $b = -0.04$, 95%-CI [-0.08, -0.00], $\chi^2(1) = 4.7$, $p = .030$). No analogous interaction emerged for amygdala versus sham (cue valence x trial number x TUS condition [amygdala - sham]: $b = -0.01$, 95%-CI [-0.05, 0.03], $\chi^2(1) = 0.3$, $p = .6$), suggesting that this subtle modulation of the Pavlovian bias over time was specific to NAcc TUS. To further characterize this effect, we re-ran the same mixed-effects regression model separately for each TUS condition. Only in the NAcc sessions did cue valence interact significantly with cue-specific trial number ($b = -0.06$, 95%-CI [-0.11, -0.01], $\chi^2(1) = 4.6$, $p = .032$), whereas no such interaction was observed in the sham ($b = 0.03$, 95%-CI [-0.02, 0.08], $\chi^2(1) = 1.4$, $p = .2$) or amygdala sessions ($b = 0.01$, 95%-CI [-0.06, 0.08], $\chi^2(1) = 0.1$, $p = .8$). This indicates that, specifically after NAcc stimulation, the Pavlovian bias starts out particularly strong and is then "unlearned" more rapidly across repetitions.

Guided by the RL results, which suggested that TUS effects were asymmetric across cue valence, we ran these analyses separately for Win and Avoid cues (**Suppl. Tables S4–7**). For Avoid cues, NAcc stimulation produced a trend-level reduction in overall Go probability relative to both sham and amygdala (main effect of TUS condition [NAcc - sham]: b = -0.05, 95%-CI [0.01, -0.10], $\chi^2(1) = 2.7$, p = .098; [NAcc – amygdala]: b = -0.04, 95%-CI [-0.09, 0.003], $\chi^2 = 3.3$, p = .067), consistent with a NAcc-specific strengthening of aversive response inhibition. No such difference was observed between amygdala and sham (main effect of TUS condition [amygdala - sham]: b = -0.01, 95%-CI [-0.06, 0.04], $\chi^2(1) = 0.1$, p = .7), in line with the weaker Pavlovian modulation attributed to amygdala at the parameter level. Importantly, this aversive effect was not time-varying: Go probability on Avoid cues did not change differentially across repetitions for any pairwise contrast (trial number × TUS [NAcc - sham]: b = 0.03, 95%-CI [-0.03, 0.08], $\chi^2(1) = 1.0$, p = .3; [amygdala - sham]: b = -0.002, 95%-CI [-0.06, 0.05], $\chi^2(1) = 0.0$, p = .9; [NAcc - amygdala]: b = 0.02, 95%-CI [-0.05, 0.09], $\chi^2(1) = 0.4$, p = .6). For Win cues, overall Go probability did not differ between any pair of conditions (main effect of TUS condition [NAcc - sham]: b = 0.003, 95%-CI [-0.06, 0.07], $\chi^2(1) = 0.0$, p = .9; [amygdala - sham]: b = 0.03, 95%-CI [-0.04, 0.10], $\chi^2(1) = 0.6$, p = .4; [NAcc - amygdala]: b = -0.02, 95%-CI [-0.09, 0.05], $\chi^2(1) = 0.5$, p = .5). Instead, the NAcc effect on appetitive bias was again time-varying: NAcc stimulation was associated with a stronger initial Go tendency on Win cues that was overwritten more rapidly across repetitions relative to sham (trial number x TUS condition [NAcc - sham]: b = -0.07, 95%-CI [-0.13, -0.02], $\chi^2(1) = 6.2$, p = .013), whereas no such interaction emerged for amygdala versus sham (trial number x TUS condition [amygdala - sham]: b = -0.03, 95%-CI [-0.09, 0.03], $\chi^2(1) = 0.8$, p = .4) or NAcc versus amygdala (trial number x TUS condition [NAcc-amygdala]: b = -0.05, 95%-CI [-0.11, 0.02], $\chi^2(1) = 2.0$, p = .16).

To further characterize learning at the choice level, we additionally analyzed win-stay/lose-shift (WSLS) tendencies. For each trial, we constructed a regressor capturing whether the current response reflected "staying" with the previous choice after a win, or "shifting" away from the previous choice after a loss for the same cue (WSLS). We then fitted a logistic mixed-effects regression with Go response as the dependent variable and require action, cue-specific trial number, cue valence, WSLS, and TUS condition as fixed factors (including their relevant interactions).

This model revealed a robust main effect of WSLS (b = 0.54, 95%-CI [0.49, 0.59], $\chi^2(1) = 381.0$, $p < 2.2e\text{-}16$), indicating that participants tended to repeat their previous response following a reward and to switch following a loss for that cue (**Suppl. Figure S3D**). Critically, there was no evidence that TUS modulated this pattern (WSLS x TUS condition [three-way]: $\chi^2(2) = 2.8$, p = .2; **Suppl. Figure S3D**). The pairwise comparisons likewise showed no evidence of TUS-induced modulation of WSLS behavior for any

pairwise contrast (WSLS x TUS condition [NAcc - sham]: b = -0.02, 95%-CI [-0.07, 0.03], $\chi^2(1)$ = 0.5, p = .5; [amygdala - sham]: b = 0.01, 95%-CI [-0.04, 0.06], $\chi^2(1)$ = 0.2, p = .6; [NAcc - amygdala]: b = -0.03, 95%-CI [-0.07, 0.005], $\chi^2(1)$ = 2.9, p = .089).

To more directly probe novelty effects, we ran an additional model in which we examined how WSLS behavior differed between the first post-novelty trial and the first familiar trial as a function of TUS. Specifically, we compared trial 2 (capturing the immediate consequence of the first encounter with a cue) with trial 3 (when the cue was already familiar) and tested the interaction of this trial index with WSLS and TUS condition in pairwise models. These comparisons between conditions showed that the difference in WSLS behavior from trial 2 to trial 3 was significantly larger for amygdala than for sham (trial number [t2-t3] x TUS condition [amygdala - sham]: b = 0.11, 95%-CI [0.002, 0.213], $\chi^2(1)$ = 4.0, p = .046; **Figure 4C**), indicating that amygdala stimulation altered how novelty-driven information was incorporated into subsequent choices. Although the NAcc – sham contrast was qualitatively similar, this difference did not reach significance (trial number [t2-t3] x TUS condition [NAcc - sham]: b = 0.05, 95%-CI [-0.05, 0.15], $\chi^2(1)$ = 0.9, p = .3; **Figure 4C**).

When focusing exclusively on trial 2, the raw data suggested higher WSLS scores for both amygdala and NAcc than for sham, but in a logistic mixed-effects regression model with Go response as the dependent variable and a WSLS regressor as predictor, the WSLS × TUS condition interaction was not significant (WSLS x TUS condition [three-way]: $\chi^2(2)$ = 1.9, p = .4; [NAcc-sham]: b = 0.03, 95%-CI [-0.13, 0.18], $\chi^2(1)$ = 0.1, p = .7; [amygdala - sham]: b = 0.11, 95%-CI [-0.04, 0.25], $\chi^2(1)$ = 2.1, p = .15; [NAcc - amygdala]: b = -0.09, 95%-CI [-0.25, 0.07], $\chi^2(1)$ = 1.1, p = .3), likely reflecting limited power of this analysis, which relies on a single post-novelty observation per cue and condition.

For later (familiar) trials, empirical win–stay/lose–shift behavior did not differ reliably between NAcc and sham sessions (WSLS x TUS condition [three-way]: $\chi^2(2)$ = 2.8, p = .2; [NAcc - sham]: b = -0.02, 95%-CI [-0.07, 0.02], $\chi^2(1)$ = 0.8, p = .4; [amygdala - sham]: b = 0.004, 95%-CI [-0.05, 0.05], $\chi^2(1)$ = 0.02, p = .9; [NAcc - amygdala]: b = -0.03, 95%-CI [-0.07, 0.01], $\chi^2(1)$ = 1.8, p = .2), even though the RL model inferred a robust NAcc-specific increase in ε_{rest_pos}. A likely explanation is that a learning rate reflects integration of outcomes over multiple past trials, whereas WSLS is strictly one-step-back: it only considers whether the current choice repeats or switches relative to the immediately preceding outcome for that cue. As a result, the richer, multi-trial learning dynamics captured by ε_{rest_pos} may not be fully expressed in simple first-order stay/switch summaries, underscoring the added sensitivity of the computational modelling approach. Full statistical details and all pairwise comparisons are reported in **Suppl.**

Tables S8-S11.

To test whether the same motivational and learning factors that shaped choice also influenced response speed, we analyzed RTs using a linear mixed-effects regression with the same fixed effects and interactions as in the logistic model.

Participants responded more quickly when a Go action was required than when they (incorrectly) pressed during NoGo cues (main effect of required action: b = -0.02, 95%-CI [0.03, 0.02], $\chi^2(1)$ = 37.5, p = 9.3e-10; **Suppl. Figure S3E**), corroborating the instrumental learning effect observed in the Go-response model. RTs also shortened over successive presentations of the same cue (main effect of trial number: b = -0.02, 95%-CI [-0.02, -0.01], $\chi^2(1)$ = 24.1, p = 9.4e-07), consistent with practice-related speeding. However, we did not find evidence for differential learning-related speeding for Go versus NoGo cues (required action × trial number: b = -0.003, 95%-CI [-0.007, 0.002], $\chi^2(1)$ = 1.5, p = .2). Likewise, this linear mixed-effects regression model did not reveal TUS effects on learning-related speeding (required action × TUS condition: b = -0.003, $\chi^2(2)$ = 0.5, p = .8; **Suppl. Figure S3E**) or on their development over time (required action × TUS condition × trial number: b = 0.0003, $\chi^2(2)$ = 6.6, p = .16). RTs were also faster after Win than after Avoid cues (main effect of cue valence: b = -0.02, 95%-CI [-0.03, -0.02], $\chi^2(1)$ = 54.7, p = 1.4e-13), reflecting a robust Pavlovian speeding effect that replicates prior work. The absence of a cue valence × trial number interaction (b = -0.001, 95%-CI [-0.006, 0.003], $\chi^2(1)$ = 0.3, p = .6) indicates that this motivational speeding remained stable throughout the task, in line with the continuous cue-replacement scheme designed to prevent extinction of the bias once cue–action mappings were learned. Again, we found no evidence that TUS modulated Pavlovian speeding, either overall (cue valence × TUS condition [three-way]: b = 0.001, $\chi^2(2)$ = 0.2, p = .9) or over time (cue valence × trial number x TUS condition [three-way]: b = -0.002, $\chi^2(2)$ = 0.3, p = .9).

S.5 TUS effects on learning rate

To test whether the stimulation effect differed between learning from positive versus negative prediction errors, we compared ε_{rest_pos} and ε_{rest_neg} in a single model including the interaction between TUS condition and learning rate valence (positive vs. negative PE). This interaction was not significant ($\chi^2(2)$ = 3.0, p = .2), and there was likewise no main effect of TUS on a general learning rate ($\chi^2(2)$ = 1.0, p = .6), indicating that stimulation effect did not differ significantly between positive and negative general learning rates. Nevertheless, the direction of effects — an increase in ε_{pos} under NAcc stimulation with comparatively little change in ε_{neg} — aligns with recent human TUS work suggesting heightened feedback sensitivity for wins relative to losses following NAcc stimulation[45].

SUPPLEMENTARY FIGURES

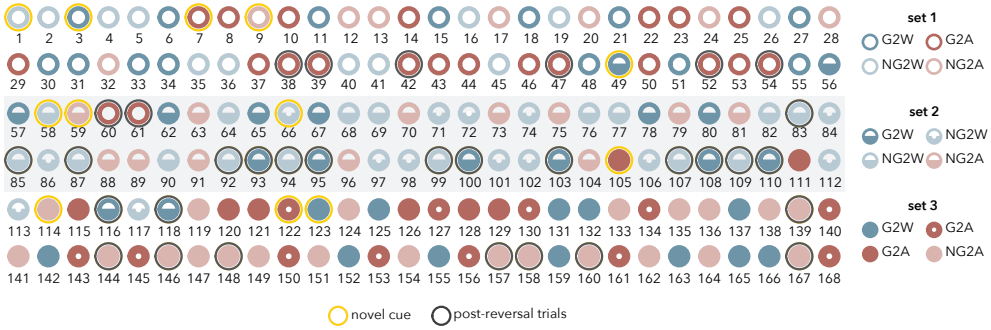


Figure S1 | Example trial sequence for a single run (168 trials). The sequence is organized into three "sets" of trials (~56 trials per set), indicated by ring fill (empty, half-filled, filled). Each set contains four cues (color-coded by condition: dark blue = G2W, light blue = NG2W, light red = G2A, dark red = NG2A). Not all four action–valence combinations necessarily appear in every set, but each set always contains at least one Go and one NoGo cue (i.e., required actions are never all the same within a set). Within a set, dots on the circles indicate different cues of the same condition (e.g. two distinct NG2W cues in set 2). Because cues are presented 12, 14 or 16 times, sets overlap in time and cues from different sets are interleaved. Cues with a yellow outline mark their first (novel) presentation. For cues that reverse their required action after the 8th presentation, post-reversal trials are shown in grey; reversals are distributed across the run such that each set contains at least one and at most two reversing cues. NB each TUS session consists of 2 runs.

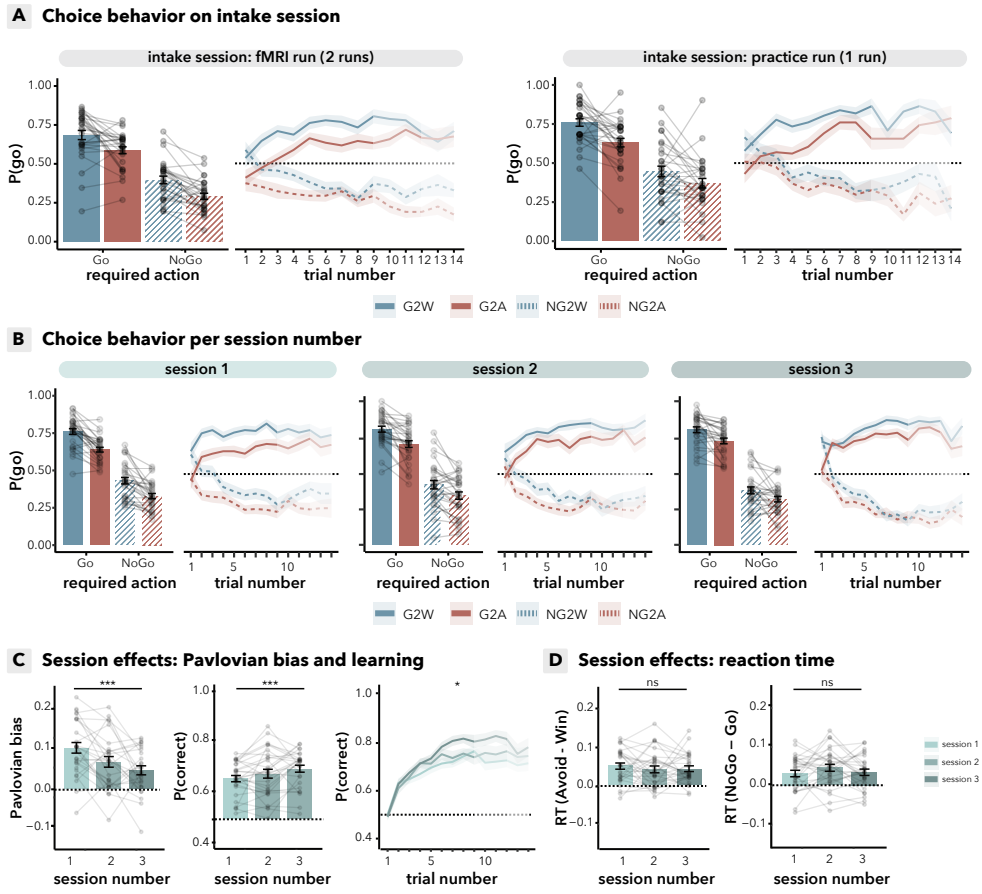


Figure S2 | Session effects on choice, learning, and reaction time. (A) *Choice behavior on intake session for fMRI run (2 runs) and practice run (1 run).* Left panels: mean P(Go) per cue type (bars = mean ± SEM; dots = participants) show robust instrumental learning (more Go when Go is required) and a persistent Pavlovian bias (more Go to Win than Avoid cues) in both fMRI and practice runs. Right panels: trial-by-trial P(Go) for each condition (lines = mean; ribbons = SEM; solid = Go-required, dashed = NoGo-required), illustrating gradual learning and a stable Win > Avoid Go bias across repetitions. **(B)** *Choice behavior per intervention session.* Left panels: mean P(Go) per cue type (bars = mean ± SEM; dots = participants) show robust instrumental learning (more Go when Go is required) and a persistent Pavlovian bias (more Go to Win than Avoid cues) in all three sessions. Right panels: trial-by-trial P(Go) for each condition (lines = mean; ribbons = SEM; solid = Go-required, dashed = NoGo-required), illustrating gradual learning and a stable Win > Avoid Go bias across repetitions. **(C)** *Session effects on Pavlovian bias and learning.* Left: Pavlovian bias (accuracy difference between congruent [G2W, NG2A] and incongruent [G2A, NG2W] cues) decreases across sessions but remains clearly above zero in session 3. Middle: mean accuracy increases with session number (bars = mean ± SEM; dots = participants). Right: trial-by-trial accuracy curves (lines = mean; ribbons = SEM) show that participants learn faster in later sessions. **(D)** *Session effects on reaction time (RT).* Left and middle: RT biases (Avoid – Win and NoGo – Go) remain stable across sessions (bars = mean ± SEM; dots = participants).

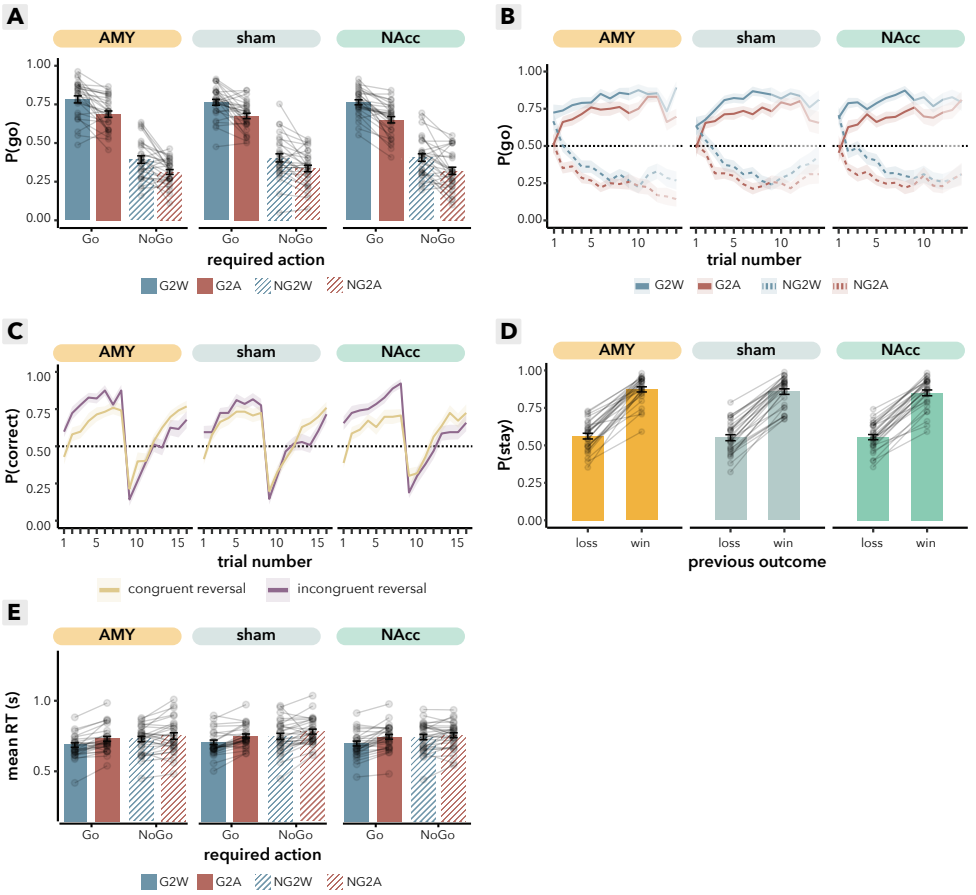


Figure S3 | Empirical behavior across TUS conditions. (A) *Mean choice behavior for amygdala (AMY), sham, and ventral striatum (NAcc) sessions.* Bars show mean probability of Go responses (± SEM) for each cue type (G2W, G2A, NG2W, NG2A); grey dots are individual participants with lines connecting conditions. Participants make more Go responses when Go is required and more Go responses to Win than Avoid cues, indicating robust instrumental learning and Pavlovian bias in all conditions. **(B)** *Trial-by-trial Go probability.* Lines (± SEM ribbons) show P(Go) across cue repetitions for each cue type (solid = Go-required, dashed = NoGo-required; blue = Win, red = Avoid). Learning increases correct responding over trials without reaching ceiling, while the Win > Avoid bias remains present across TUS conditions. **(C)** *Reversal learning.* Accuracy (P(correct)) for reversal cues as a function of cue repetition. Yellow lines show reversals from incongruent→congruent mappings; purple lines show congruent→incongruent reversals. In all conditions, accuracy drops at the reversal point and then recovers, with higher performance when mappings are initially congruent. **(D)** *Win–stay/lose–shift behavior.* Bars (± SEM) show the probability of staying with the previous response after a win and switching after a loss (P(stay)) for each TUS condition; dots are participants. Across conditions, participants are more likely to stay after wins than losses, reflecting outcome-sensitive choice updating. **(E)** *Reaction times.* Mean RTs (± SEM) for each cue type and required action in AMY, sham, and NAcc sessions; dots indicate individual participants. Responses are faster when a Go action is correct and faster after Win than Avoid cues in all TUS conditions.

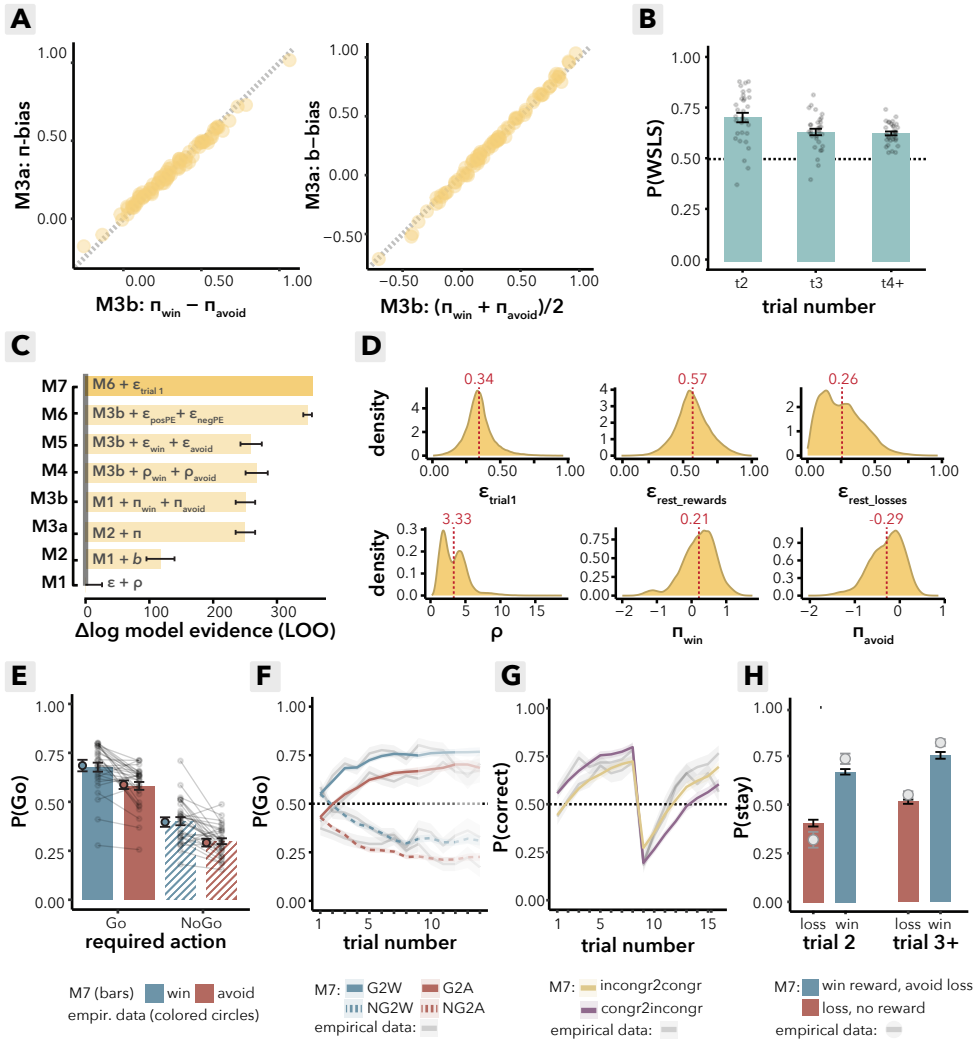


Figure S4 | Model validation in the independent dataset (fMRI behavior). (A) *Relationship between the bias parameters of M3a and M3b.* Scatterplots show that the single Go bias and Pavlovian bias in M3a are mathematically equivalent to linear combinations of the valence-specific win and avoid biases in M3b (left: Pavlovian bias vs. πwin − πavoid; right: Go bias vs. (πwin + πavoid)/2). **(B)** *Win–stay/lose–shift behavior in the independent dataset used for model fitting/selection.* Bars (± SEM) show the probability of staying after wins or shifting after losses (P(WSLS)) on trial 2, trial 3, and later ("rest") trials; grey dots are individual participants. Novelty-driven updating is evident from the higher WSLS on trial 2 than on trial 3 and the remaining trials. **(C)** *Leave-one-out (LOO) model comparison.* Bars show Δlog model evidence relative to the worst model (M1). Models that included Pavlovian bias and separate learning rates for positive vs. negative prediction errors outperformed simpler models; adding a separate learning rate for trial 1 (M7) provided the best fit. **(D)** *Group-level posterior densities for the winning model (M7).* Distributions are shown for the novelty learning rate ε_1, later-trial learning rates after positive and negative prediction errors (ε_rest_pos, ε_rest_neg), feedback sensitivity ρ, and

Pavlovian win and avoid biases (π_win, π_avoid). Vertical dashed lines mark posterior means. **(E–G)** Posterior-predictive checks for M7. **(E)** *Mean choice behavior.* Colored bars (± SEM; blue = Win, red = Avoid; solid = Go-required, hatched = NoGo-required) show model-predicted P(Go); colored circles (± SEM) show empirical data; faint grey dots and connecting lines indicate individual model-predicted participants. **(F)** *Trial-by-trial Go probability.* Colored lines (± SEM ribbons; solid = Go-required, dashed = NoGo-required; blue = Win, red = Avoid) show model predictions across cue repetitions; grey lines and ribbons show empirical data. **(G)** *Reversal learning.* Colored lines (± SEM ribbons) show predicted accuracy for incongruent→congruent (yellow) and congruent→incongruent (purple) reversals; grey lines and ribbons show empirical accuracy. **(H)** *Stay behavior after positive vs. negative outcomes, separated by novelty.* Bars (± SEM) show model-predicted P(stay) after losses (loss or no reward) and wins (reward or avoided loss) on trial 2 and on the remaining trials; grey circles (± SEM) show empirical P(stay).

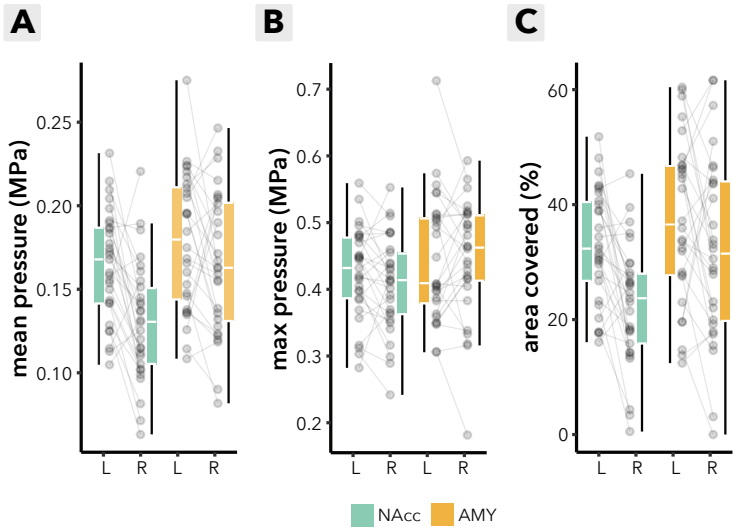


Figure S5 | Acoustic simulations and TUS protocol. (A) Mean pressure was higher in AMY (0.17 MPa, SD = 0.035, range = 0.10–0.23) than in NAcc (0.15 MPa, SD = 0.028, range = 0.09–0.19). Within NAcc, right-hemisphere pressures tended to be slightly lower than left, suggesting somewhat weaker or more variable drive to right NAcc. **(B)** *Maximum pressure within each target.* Colored boxes (orange = AMY, green = NAcc) show the distribution of subject-specific peak pressures (MPa) in left (L) and right (R) targets; grey dots and connecting lines indicate individual participants. Across participants, mean maximum pressure in NAcc was 0.42 MPa (SD = 0.06, range = 0.27–0.53) and in AMY 0.45 MPa (SD = 0.08, range = 0.24–0.63), with slightly higher peaks in AMY, consistent with targeting through the thinner temporal bone. *Mean pressure within each target.* Bars show average pressures (MPa) in NAcc and AMY (orange/green as in A), with grey dots and lines again indicating individuals. **(C)** *Target coverage above the 0.2 MPa threshold.* Bars show the percentage of each anatomical mask exceeding 0.2 MPa, with individual participants in grey. On average, 27.8% of the NAcc mask (SD = 8.0, range = 11.0–44.1) and 34.8% of the AMY mask (SD = 12.8, range = 7.4–57.2) were above threshold. All participants showed bilateral NAcc coverage ≥ 0.2 MPa, and all but one showed bilateral AMY coverage (in that participant, right AMY remained below threshold, yielding predominantly unilateral stimulation).

SUPPLEMENTARY TABLES

Table S1 | Posterior parameter estimates and model comparison for M1–M7 in the fMRI validation data set. Group-level summaries are shown for each RL model fitted hierarchically to the independent fMRI data. Entries give the posterior mean of each group-level parameter with its posterior SD in parentheses. The "Learning rate (ε)" block reports a single learning-rate parameter for M1–M4, separate learning rates for Win and Avoid cues for M5, and sign-specific learning rates for positive and negative prediction errors (ε_posPE, ε_negPE) for M6 and M7. In M7, ε_posPE and ε_negPE apply to later (trial ≥ 2) outcomes, while ε_trial1 captures a separate novelty learning rate for the first presentation of each cue. The "Rho (ρ)" block lists feedback-sensitivity parameters (overall ρ in M1–3 and 5-7, and valence-specific ρ_win and ρ_avoid in M4). The "Go bias (β) / Pavlovian bias (π)" block reports a single Go bias for M2, single Go and Pavlovian bias for M3a, and valence-specific Pavlovian win and avoid biases (π_win, π_avoid) for M3b, M4–M7; the M3 column ("M3 (a & b)") therefore aggregates the common parameters of M3a and M3b and, where relevant, reports either the shared (β, π) or the valence-specific (π_win, π_avoid) versions. The bottom row ("LOO (elpd_loo)") shows the approximate leave-one-out expected log pointwise predictive density for each model (higher values indicate better predictive accuracy), with the standard error of elpd_loo in parentheses.

	M1	M2	M3 (a & b)	M4	M5	M6	M7
Learning rate (ε)	0.38 (0.03)	0.37 (0.03)	0.38 (0.03)	0.38 (0.03)			
ε win					0.37 (0.03)		
ε avoid					0.38 (0.03)		
ε posPE						0.52 (0.01)	0.57 (0.01)
ε negPE						0.25 (0.03)	0.26 (0.03)
ε trial 1							0.34 (0.01)
Rho (ρ)	4.5 (1.00)	4.5 (1.00)	4.53 (1.00)		4.45 (0.90)	3.55 (0.36)	3.33 (0.33)
ρ win				4.89 (1.00)			
ρ avoid				4.18 (0.78)			
Go bias (b)		-0.04 (0.07)	-0.04 (0.07)				
Pavlovian bias (π)			0.28 (0.05)				
π win			0.25 (0.10)	0.27 (0.10)	0.26 (0.10)	0.22 (0.08)	0.21 (0.09)
π avoid			-0.32 (0.07)	-0.32 (0.08)	-0.33 (0.08)	-0.27 (0.06)	-0.29 (0.06)
LOO	-5664 (41.2)	-5547 (43.4)	-5413 (45.2)	-5404 (45.1)	-5396 (45.1)	-5316 (46.2)	-5308 (46.8)

Table S2 | Posterior parameter estimates for M6 and M7 (TUS data models). Group-level summary statistics of the posterior distributions for models M6 and M7, which both include separate learning-rate parameters for positive and negative prediction errors (ε_pos, ε_neg). M7 additionally includes a distinct learning rate for first cue presentations (ε_1), such that ε_pos and ε_neg in M7 apply only to later (non-novel) trials. Columns report estimates for NAcc TUS, sham, and amygdala TUS sessions. For each parameter (learning rate trial 1, learning rate positive PE, learning rate negative PE, feedback sensitivity, Pavlovian win bias, Pavlovian avoid bias), the table lists the posterior mean, standard deviation (SD), and the minimum and maximum posterior values across participants.

		M6			M7		
		NAcc	sham	AMY	NAcc	sham	AMY
Learning rate (ε) trial 1	mean				0.45	0.35	0.46
	SD				0.02	0.01	0.02
	min				0.29	0.23	0.25
	max				0.63	0.49	0.69
Learning rate (ε) positive PE	mean	0.58	0.51	0.56	0.59	0.57	0.56
	SD	0.02	0.02	0.02	0.02	0.01	0.02
	min	0.39	0.33	0.28	0.38	0.44	0.38
	max	0.82	0.77	0.82	0.78	0.76	0.78
Learning rate (ε) negative PE	mean	0.28	0.29	0.30	0.29	0.28	0.30
	SD	0.02	0.02	0.03	0.02	0.02	0.02
	min	0.07	0.08	0.05	0.07	0.08	0.06
	max	0.61	0.48	0.62	0.60	0.50	0.65
Feedback sensitivity (ρ)	mean	4.23	4.49	4.69	3.94	4.30	4.40
	SD	0.44	0.43	0.39	0.39	0.41	0.36
	min	0.62	0.94	0.68	0.62	0.92	0.69
	max	9.46	8.84	8.29	8.24	9.41	7.91
Pavlovian (π) win bias	mean	0.47	0.49	0.51	0.52	0.51	0.56
	SD	0.06	0.05	0.07	0.07	0.06	0.09
	min	-0.20	-0.16	-0.49	-0.17	-0.31	-0.71
	max	1.17	1.10	1.19	1.37	1.34	1.52
Pavlovian (π) avoid bias	mean	-0.08	0.06	-0.03	-0.09	0.07	-0.02
	SD	0.05	0.05	0.05	0.07	0.06	0.07
	min	-0.49	-0.32	-0.48	-0.83	-0.44	-0.71
	max	0.47	0.62	0.57	0.60	0.70	0.74

Table S3 | Pairwise logistic mixed-effects regressions of Go responding as a function of TUS condition. Results from logistic mixed-effects regression models with the log-odds of a Go response (P(Go)) as the dependent variable. Each block shows a separate pairwise contrast between stimulation conditions (NAcc – sham, AMY – sham, NAcc – AMY). Fixed effects were required action (Go vs. NoGo), cue valence (Win vs. Avoid), cue count (cue-specific trial number, z-scored), TUS condition (coded -1/1 within each pairwise model), and their indicated interactions. β gives the regression coefficient, χ^2 the Type-III Wald chi-square statistic with its degrees of freedom (Df), 95-CI the 95% confidence interval for β on the log-odds scale, and *p* the associated *p*-value. Positive β values indicate higher probability of making a Go response. The significant Valence × Cue Count × TUS interaction in the NAcc – sham contrast (bold) reflects a TUS-dependent change in the Pavlovian bias over cue repetitions.

Predictor	**NAcc - sham**					**AMY - sham**					**NAcc - AMY**				
PGO	β	χ^2	**Df**	**95-CI**	ρ	β	χ^2	**Df**	**95-CI**	ρ	β	χ^2	**Df**	**95-CI**	ρ
Intercept	0.18	14.1	1	0.08, 0.27	1.7e-4	0.21	17.6	1	0.11, 0.30	2.8e-05	0.19	14.3	1	0.09, 0.28	1.6e-4
Required Action	0.77	138.2	1	0.64, 0.90	< 2.2e-16	0.81	163.0	1	0.69, 0.93	< 2.2e-16	0.80	150.6	1	0.67, 0.93	< 2.2e-16
Valence	0.22	48.3	1	0.16, 0.28	3.6e-12	0.21	39.0	1	0.14, 0.28	4.3e-10	-0.23	41.8	1	-0.30, -0.16	1.0e-10
Cue Count	-0.04	2.5	1	-0.09,0.01	0.11	-0.03	1.4	1	-0.09, 0.02	0.2	-0.05	9.1	1	-0.09, -0.02	0.003
TUS condition	0.03	1.0	1	-0.03,0.08	0.3	-0.005	0.05	1	-0.05, 0.04	0.8	-0.04	2.2	1	-0.08, 0.01	0.14
Required Action x Cue Count	0.39	42.7	1	0.27, 0.50	6.4e-11	0.39	56.2	1	0.29, 0.50	6.6e-14	0.38	52.6	1	0.28, 0.49	4.1e-13
Valence x Cue Count	-0.01	0.3	1	-0.05, 0.03	0.6	0.02	0.8	1	-0.02, 0.06	0.4	-0.02	1.2	1	-0.06, 0.02	0.3
Required Action x TUS	-0.01	0.02	1	-0.09, 0.08	0.9	0.03	0.5	1	-0.05, 0.12	0.5	-0.04	1.9	1	-0.09, 0.02	0.2
Valence x TUS	0.02	0.6	1	-0.02, 0.06	0.4	0.01	0.2	1	-0.03, 0.06	0.6	0.01	0.1	1	-0.03, 0.04	0.8
Req.Act x Cue Count x TUS	0.01	0.9	2	-0.09, 0.11	0.6	-0.01	0.7	2	-0.11, 0.08	0.7	0.02	0.3	2	-0.08, 0.12	0.9
Valence x Cue Count x TUS	-0.04	4.7	1	-0.08, -0.00	0.030	0.01	0.3	1	-0.05, 0.03	0.6	-0.03	1.7	1	-0.08, 0.02	0.19

Table S4 | Pairwise logistic mixed-effects regressions of Go responding for Avoid cues. Results from logistic mixed-effects regression models with the log-odds of a Go response to Avoid cues (P(Go | Avoid)) as the dependent variable. Each block shows a separate pairwise contrast between stimulation conditions (NAcc – sham, AMY – sham, NAcc – AMY). Fixed effects were required action (Go vs. NoGo), cue count (cue-specific trial number, z-scored), TUS condition (coded -1/1 within each pairwise model), and their indicated interactions. β gives the regression coefficient, χ^2 the Type-III Wald chi-square statistic with its degrees of freedom (Df), 95-CI the 95% confidence interval for β, and *p* the associated *p*-value. Negative β values indicate a lower probability of making a Go response to Avoid cues. The trend-level main effects of TUS condition in the NAcc – sham and NAcc – AMY contrasts reflect numerically fewer Go responses to Avoid cues after NAcc stimulation, consistent with a strengthened aversive Go bias.

Predictor	**NAcc - sham**					**AMY - sham**					**NAcc - AMY**				
PGO (Avoid)	β	χ^2	**Df**	**95-CI**	ρ	β	χ^2	**Df**	**95-CI**	ρ	β	χ^2	**Df**	**95-CI**	ρ
Intercept	-0.04	0.7	1	-0.14,0.06	0.4	-0.001	0.0	1	-0.10, 0.10	0.99	-0.05	0.8	1	-0.15, 0.06	0.4
Required Action	0.73	117.2	1	0.60, 0.86	< 2.2e-16	0.77	162.1	1	0.65, 0.89	< 2.2e-16	0.77	113.1	1	0.62, 0.91	< 2.2e-16
Cue Count	-0.03	1.8	1	-0.08,0.01	0.18	-0.05	1.8	1	-0.12, 0.02	0.18	-0.03	1.1	1	-0.10, 0.03	0.3
TUS condition	0.04	1.7	1	-0.02,0.11	0.19	0.01	0.2	1	-0.04, 0.07	0.6	0.05	3.0	1	-0.01, 0.10	0.083
Required Action x Cue Count	0.36	32.7	1	0.24, 0.49	1.1e-08	0.38	39.0	1	0.26, 0.50	4.2e-10	0.38	39.4	1	0.26, 0.50	3.5e-10
Required Action x TUS	-0.0002	0.0	1	-0.10,0.10	>0.9	0.04	0.9	1	-0.04, 0.12	0.3	-0.03	0.9	1	-0.10, 0.03	0.4
Req.Act x Cue Count x TUS	0.01	0.02	1	-0.07, 0.08	0.9	0.02	0.5	1	-0.04, 0.08	0.5	-0.01	0.1	1	-0.07, 0.05	0.8

Table S5 | Pairwise logistic mixed-effects regressions of Go responding for Avoid cues (reduced model). Results from logistic mixed-effects regression models with the log-odds of a Go response on Avoid trials (P(Go | Avoid)) as the dependent variable. Each block reports a separate pairwise contrast between stimulation conditions (NAcc – sham, AMY – sham, NAcc – AMY). Fixed effects were required action (Go vs. NoGo), TUS condition (coded -1/1 within each pairwise model), and their interaction. β denotes the regression coefficient, χ^2 the Type-III Wald chi-square statistic with its degrees of freedom (Df), 95-CI the 95% confidence interval for β, and *p* the associated *p*-value. In the NAcc – sham and NAcc – AMY contrasts, TUS condition shows a trend toward fewer Go responses after NAcc stimulation, consistent with a numerically stronger aversive Go bias following NAcc TUS; no such trend is evident for AMY – sham.

Predictor	**NAcc - sham**					**AMY - sham**					**NAcc - AMY**				
PGO (Avoid)	β	χ^2	**Df**	**95-CI**	ρ	β	χ^2	**Df**	**95-CI**	ρ	β	χ^2	**Df**	**95-CI**	ρ
Intercept	-0.03	0.4	1	-0.13, 0.07	0.5	0.01	0.03	1	-0.09, 0.11	0.9	-0.03	0.4	1	-0.14, 0.07	0.5
Required Action	0.67	141.3	1	0.56, 0.78	< 2.2e-16	0.72	202.3	1	0.62, 0.81	<2.2e-16	0.71	129.8	1	0.59, 0.83	< 2.2e-16
TUS condition	-0.05	2.7	1	-0.10, 0.01	0.098	-0.01	0.1	1	-0.06, 0.04	0.7	-0.04	3.3	1	-0.09, 0.003	0.067
Required Action x TUS	-0.002	0.003	1	-0.09, 0.08	>0.9	0.04	1.3	1	-0.03, 0.11	0.3	-0.04	1.8	1	-0.10, 0.02	0.18

Table S6 | Pairwise logistic mixed-effects regressions of Go responding for Win cues. Results from logistic mixed-effects regression models with the log-odds of a Go response to Win cues (P(Go | Win)) as the dependent variable. Each block shows a separate pairwise contrast between stimulation conditions (NAcc – sham, AMY – sham, NAcc – AMY). Fixed effects were required action (Go vs. NoGo), cue count (cue-specific trial number, z-scored), TUS condition (coded -1/1 within each pairwise model), and their indicated interactions. β gives the regression coefficient, χ^2 the Type-III Wald chi-square statistic with its degrees of freedom (Df), 95-CI the 95% confidence interval for β, and *p* the associated *p*-value. Negative β values indicate a lower probability of making a Go response to Win cues. The significant Cue Count × TUS interaction in the NAcc – sham contrast (bold) reflects a TUS-dependent change in the appetitive Pavlovian bias over cue repetitions.

Predictor	**NAcc - sham**					**AMY - sham**					**NAcc - AMY**				
PGO (Win)	β	χ^2	**Df**	**95-CI**	ρ	β	χ^2	**Df**	**95-CI**	ρ	β	χ^2	**Df**	**95-CI**	ρ
Intercept	0.41	43.9	1	0.29, 0.53	3.5e-11	0.45	41.4	1	0.31, 0.59	1.3e-10	0.44	37.1	1	0.30, 0.58	1.1e-09
Required Action	0.82	136.4	1	0.69, 0.96	<2.2e-16	0.88	135.8	1	0.73, 1.02	<2.2e-16	0.86	145.9	1	0.72, 0.99	<2.2e-16
Cue Count	-0.04	2.2	1	-0.09, 0.01	0.14	0.01	0.0	1	-0.07, 0.08	0.8	-0.06	4.1	1	-0.12, -0.002	0.042
TUS condition	-0.01	0.2	1	-0.08, 0.05	0.7	0.02	0.4	1	-0.05, 0.10	0.5	-0.04	1.3	1	-0.11, 0.03	0.3
Required Action x Cue Count	0.41	44.2	1	0.29, 0.54	3.0e-11	0.43	59.6	1	0.32, 0.54	1.2e-14	0.41	50.8	1	0.30, 0.52	1.0e-12
Required Action x TUS	-0.01	0.1	1	-0.10, 0.08	0.8	-0.03	0.4	1	-0.14, 0.07	0.5	-0.05	1.7	1	-0.12, 0.02	0.19
Cue Count x TUS	-0.07	5.6	1	-0.13, -0.01	0.018	-0.03	0.8	1	-0.09, 0.03	0.4	-0.05	2.0	1	-0.11, 0.02	0.16
Req.Act x Cue Count x TUS	-0.02	0.4	1	-0.08, 0.04	0.54	-0.01	0.0	1	-0.07, 0.06	0.8	-0.01	0.2	1	-0.08, 0.05	0.7

Table S7 | Pairwise logistic mixed-effects regressions of Go responding for Win cues (reduced model). Results from logistic mixed-effects regression models with the log-odds of a Go response on Win trials (P(Go | Win)) as the dependent variable. Each block reports a separate pairwise contrast between stimulation conditions (NAcc – sham, AMY – sham, NAcc – AMY). Fixed effects were required action (Go vs. NoGo), TUS condition (coded -1/1 within each pairwise model), and their interaction. β denotes the regression coefficient, χ^2 the Type-III Wald chi-square statistic with its degrees of freedom (Df), 95-CI the 95% confidence interval for β, and *p* the associated *p*-value.

Predictor	**NAcc - sham**					**AMY - sham**					**NAcc - AMY**				
PGO (Win)	β	χ^2	**Df**	**95-CI**	π	β	χ^2	**Df**	**95-CI**	π	β	χ^2	**Df**	**95-CI**	π
Intercept	0.40	46.3	1	0.28, 0.51	1.1e-11	0.42	38.2	1	0.29, 0.55	6.5e-10	0.42	40.3	1	0.29, 0.56	2.2e-10
Required Action	0.76	159.3	1	0.65, 0.88	<2.2e-16	0.81	146.1	1	0.68, 0.95	<2.2e-16	0.81	159.2	1	0.68, 0.93	<2.2e-16
TUS condition	0.003	0.0	1	-0.06, 0.07	0.9	0.03	0.6	1	-0.04, 0.10	0.4	-0.02	0.5	1	-0.09, 0.05	0.5
Required Action x TUS	0.005	0.0	1	-0.09, 0.08	0.9	-0.04	0.7	1	-0.06, 0.14	0.4	-0.05	2.0	1	-0.11, 0.02	0.16

Table S8 | Pairwise logistic mixed-effects regressions of win–stay/lose–shift behavior. Results from logistic mixed-effects regression models with the log-odds of a Go response as the dependent variable and win–stay/lose–shift (WSLS) as a trial-wise predictor. Each block reports a pairwise contrast between TUS conditions (NAcc – sham, AMY – sham, NAcc – AMY). Fixed effects were required action (Go vs. NoGo), cue valence (Win vs. Avoid), cue count (cue-specific trial number, z-scored), WSLS (1 = stay after win / shift after loss; 0 = other), TUS condition (coded -1/1 within each pairwise model), required action x cue count, and the WSLS × TUS interaction. β denotes the regression coefficient, χ^2 the Type-III Wald chi-square statistic with its degrees of freedom (Df), 95-CI the 95% confidence interval for β, and *p* the associated *p*-value. Across all contrasts, WSLS showed a strong positive main effect, indicating robust win–stay/lose–shift behavior, whereas the WSLS × TUS interaction was not significant, suggesting that TUS did not reliably modulate single-trial stay/switch tendencies.

Predictor	**NAcc - sham**					**AMY - sham**					**NAcc - AMY**				
WSLS	β	χ^2	**Df**	**95-CI**	π	β	χ^2	**Df**	**95-CI**	π	β	χ^2	**Df**	**95-CI**	π
Intercept	0.19	21.2	1	0.11, 0.27	4.2e-06	0.22	25.3	1	0.13, 0.30	5.0e-07	0.19	20.9	1	0.11, 0.27	4.9e-06
Required Action	0.64	90.1	1	0.50, 0.77	< 2.2e-16	0.68	109.2	1	0.55, 0.81	< 2.2e-16	0.67	99.4	1	0.54, 0.81	< 2.2e-16
Valence	0.21	40.2	1	0.14, 0.27	2.3e-10	0.20	32.2	1	0.13, 0.27	1.4e-08	0.21	34.1	1	0.14, 0.27	5.2e-09
Cue Count	-0.02	0.6	1	-0.07, 0.03	0.4	-0.02	0.2	1	-0.10, 0.06	0.7	-0.02	0.7	1	-0.08, 0.03	0.4
WSLS	0.53	412.9	1	0.48, 0.58	< 2.2e-16	0.55	321.0	1	0.49, 0.61	< 2.2e-16	0.54	285.2	1	0.48, 0.60	5.2e-09
TUS condition	-0.03	1.1	1	-0.08, 0.03	0.3	-0.004	0.03	1	-0.05, 0.04	0.9	-0.03	1.0	1	-0.08, 0.02	0.3
Required Action x Cue Count	0.41	32.6	1	0.27, 0.55	1.2e-08	0.44	42.7	1	0.31, 0.571	6.4e-11	0.43	40.0	1	0.30, 0.56	2.6e-10
WSLS x TUS	-0.02	0.5	1	-0.07, 0.03	0.5	0.01	0.2	1	-0.04, 0.06	0.6	-0.03	2.9	1	-0.07, 0.005	0.089

Table S9 | Pairwise logistic mixed-effects regressions of win–stay/lose–shift on trial 2 (novelty). Results from logistic mixed-effects regression models restricted to trial 2, capturing how the very first outcome for each cue (trial 1) affects the immediately following choice. The dependent variable is the log-odds of a Go response. Fixed effects are required action (Go vs. NoGo), cue valence (Win vs. Avoid), WSLS (1 = stay after win / shift after loss for that cue; 0 = other), TUS condition (coded -1/1 within each pairwise contrast), and the WSLS × TUS interaction. Each block reports one pairwise comparison (NAcc – sham, AMY – sham, NAcc – AMY). β denotes the regression coefficient, χ^2 the Type-III Wald chi-square statistic with its degrees of freedom (Df), 95-CI the 95% confidence interval for β, and *p* the associated *p*-value. WSLS showed a strong positive main effect in all contrasts, indicating robust novelty-driven win–stay/lose–shift behavior, whereas the WSLS × TUS interaction did not reach significance, suggesting that single-trial novelty updating was not reliably modulated by TUS.

Predictor	**NAcc - sham**					**AMY - sham**					**NAcc - AMY**				
WSLS t2	β	χ^2	**Df**	**95-CI**	π	β	χ^2	**Df**	**95-CI**	π	β	χ^2	**Df**	**95-CI**	π
Intercept	0.42	18.7	1	0.23, 0.61	1.5e-05	0.46	21.4	1	0.26, 0.65	3.7e-06	0.41	15.3	1	0.20, 0.62	9.2e-05
Required Action	-0.05	0.4	1	-0.21, 0.11	0.5	-0.13	2.7	1	-0.28, 0.03	0.10	0.05	0.2	1	-0.14, 0.24	0.6
Valence	0.37	27.5	1	0.23, 0.50	1.6e-07	0.34	22.6	1	0.20, 0.49	2.0e-06	0.49	35.4	1	0.33, 0.65	2.7e-09
WSLS	1.31	120.3	1	1.08, 1.54	< 2.2e-16	1.41	165.4	1	1.19, 1.62	< 2.2e-16	1.44	148.4	1	1.21, 1.67	< 2.2e-16
TUS condition	-0.07	0.8	1	-0.24, 0.09	0.4	-0.03	0.2	1	-0.17, 0.11	0.7	-0.04	0.2	1	-0.22, 0.13	0.6
WSLS x TUS	0.03	0.1	1	-0.13, 0.18	0.7	0.11	2.1	1	-0.03, 0.25	0.15	-0.09	1.1	1	-0.25, 0.07	0.3

Table S10 | Pairwise logistic mixed-effects regressions of win–stay/lose–shift on familiar trials (t > 2). Results from logistic mixed-effects regression models restricted to trials after the first post-novelty trial (t > 2), capturing win–stay/lose–shift (WSLS) behavior once cues are familiar. The dependent variable is the log-odds of a Go response. Fixed effects are required action (Go vs. NoGo), cue valence (Win vs. Avoid), WSLS (1 = stay after win / shift after loss for that cue; 0 = other), TUS condition (coded -1/1 within each pairwise contrast), and the WSLS × TUS interaction. Each block reports one pairwise comparison (NAcc – sham, AMY – sham, NAcc – AMY). β denotes the regression coefficient, χ^2 the Type-III Wald chi-square statistic with its degrees of freedom (Df), 95-CI the 95% confidence interval for β, and *p* the associated *p*-value. Across all contrasts, WSLS showed a strong positive main effect, indicating robust stay/switch learning for familiar cues, whereas the WSLS × TUS interaction was non-significant, suggesting that TUS did not reliably modulate win–stay/lose–shift behavior once cue–outcome contingencies were established.

Predictor	NAcc - sham					AMY - sham					NAcc - AMY				
	β	χ^2	Df	95-CI	π	β	χ^2	Df	95-CI	π	β	χ^2	Df	95-CI	π
Intercept	0.17	18.1	1	0.09, 0.25	2.1e-05	0.20	22.5	1	0.11, 0.29	2.1e-06	0.18	18.1	1	0.10, 0.26	2.0e-05
Required Action	0.68	92.9	1	0.54, 0.82	< 2.2e-16	0.74	113.9	1	0.60, 0.87	< 2.2e-16	0.72	101.9	1	0.58, 0.86	< 2.2e-16
Valence	0.20	35.0	1	0.14, 0.27	3.3e-09	0.19	29.4	1	0.12, 0.26	5.8e-08	0.20	28.0	1	0.13, 0.28	1.2e-07
WSLS	0.48	300.8	1	0.43, 0.54	< 2.2e-16	0.50	234.8	1	0.43, 0.56	< 2.2e-16	0.49	227.7	1	0.42, 0.55	< 2.2e-16
TUS condition	-0.03	0.9	1	-0.08, 0.03	0.3	-0.002	0.01	1	-0.05, 0.05	0.9	-0.03	1.2	1	-0.08, 0.02	0.3
WSLS x TUS	-0.02	0.8	1	-0.07, 0.02	0.4	0.004	0.02	1	-0.05, 0.05	0.9	-0.03	1.8	1	-0.07, 0.01	0.18

Table S11 | Pairwise logistic mixed-effects regressions of win–stay/lose–shift behavior on trials 2 and 3 (novelty vs. first familiar trial). Results from logistic mixed-effects regression models restricted to trial 2 and trial 3, capturing how the first outcome after a novel cue (trial 2) versus the first familiar encounter (trial 3) shapes choice. The dependent variable is the log-odds of a Go response (WSLS t2*3). Fixed effects are required action (Go vs. NoGo), cue valence (Win vs. Avoid), cue count (trial 2 vs. 3), WSLS (1 = stay after win / shift after loss for that cue; 0 = other), TUS condition (coded -1/1 within each pairwise contrast), and their interactions, including Cue count × WSLS and Cue count × WSLS × TUS. Each block reports one pairwise comparison (NAcc – sham, AMY – sham, NAcc – AMY). β denotes the regression coefficient, χ^2 the Type-III Wald chi-square statistic with its degrees of freedom (Df), 95-CI the 95% confidence interval for β, and *p* the associated *p*-value. Across contrasts, WSLS shows a strong positive main effect, confirming robust win–stay/lose–shift behavior at the transition from novel to familiar trials. The Cue count × WSLS × TUS interaction is significant only in the AMY – sham comparison, indicating that amygdala stimulation selectively modulates how novelty-driven information (trial 2) versus first-familiar information (trial 3) is incorporated into subsequent choices, whereas no comparable modulation is observed for NAcc.

Predictor	**NAcc - sham**					**AMY - sham**					**NAcc - AMY**				
WSLS t2*3	β	χ^2	**Df**	**95-CI**	π	β	χ^2	**Df**	**95-CI**	π	β	χ^2	**Df**	**95-CI**	π
Intercept	0.32	16.0	1	0.16, 0.47	6.2e-05	0.35	19.0	1	0.20, 0.51	1.3e-05	0.34	24.2	1	0.21, 0.48	8.8e-07
Required Action	0.26	26.5	1	0.16, 0.36	2.6e-07	0.23	11.0	1	0.09, 0.36	9.2e-4	0.35	23.2	1	0.21, 0.49	1.5e-06
Valence	0.35	49.1	1	0.25, 0.44	2.4e-12	0.29	36.7	1	0.20, 0.38	1.4e-09	0.39	59.9	1	0.30, 0.49	10.0e-15
Cue Count	0.09	3.1	1	-0.01, 0.19	0.080	0.12	4.7	1	0.01, 0.22	0.030	0.08	2.0	1	-0.02, 0.19	0.16
WSLS	0.90	169.9	1	0.77, 1.04	< 2.2e-16	0.96	161.4	1	0.81, 1.11	< 2.2e-16	0.92	159.4	1	0.78, 1.06	< 2.2e-16
TUS condition	-0.02	0.1	1	-0.15, 0.11	0.8	0.02	0.1	1	-0.09, 0.12	0.8	-0.04	0.2	1	-0.19, 0.12	0.7
Cue count x WSLS	0.24	17.4	1	0.12, 0.34	3.0e-05	0.28	27.6	1	0.18, 0.39	1.5e-07	0.34	34.0	1	0.23, 0.50	5.4e-09
WSLS x TUS	-0.03	0.3	1	-0.13, 0.07	0.6	0.02	0.1	1	-0.08, 0.11	0.7	-0.06	1.1	1	-0.16, 0.05	0.3
Cue Count x WSLS x TUS	0.05	0.9	1	-0.05, 0.15	0.3	0.11	4.0	1	-0.002, 0.21	0.046	-0.05	0.7	1	-0.16, 0.06	0.4



PART V

Chapter 5
Epilogue



CHAPTER 5

General discussion



Understanding how our motivations shape our actions is central to understanding human behavior. The gap between knowing what we should do and what we actually do is not a personal failing, but a fundamental feature of how the human brain is built. This coupling between motivation and action becomes particularly important when it breaks down: in addiction, cues trigger approach despite intentions to quit; in anxiety, aversive biases promote excessive avoidance. Understanding the neural mechanisms underlying normal motivation-action coupling is therefore essential – both as a fundamental question in Cognitive Neuroscience and as a foundation for eventually developing better interventions when these mechanisms go awry.

The primary aim of this thesis is to **provide a deeper understanding of the neural mechanisms underlying motivation-action coupling**. Specifically, I investigated the causal roles of the ventral striatum (nucleus accumbens, NAcc) and amygdala in Pavlovian control and instrumental learning. To achieve this aim, I employed transcranial ultrasound stimulation (TUS) to causally manipulate these deep brain structures, an approach that has only recently become feasible in humans. However, before I could investigate these subcortical contributions, I first had to establish whether ultrasound stimulation can be used to modulate behavior in humans. Two fundamental questions therefore guided this work: Can TUS effectively modulate human behavior? And, what is the causal role of the amygdalo-striatal circuitry in how motivations shape our actions?

In this thesis, I first demonstrate that TUS can produce rapid and specific modulation of choice behavior in humans, and that these effects depend on baseline inhibitory balance in the brain. Second, I provide causal evidence for distinct contributions of the NAcc and amygdala to reinforcement learning and motivationally driven action bias in healthy humans.

I. SUMMARY OF MAIN FINDINGS

Chapter 2

Before investigating the neural substrates of motivation-action coupling, it was essential to establish whether TUS can effectively modulate human behavior. At the start of my doctoral work, evidence for TUS efficacy in humans was sparse and often confounded by auditory and somatosensory side effects[1,2]. While animal studies had demonstrated robust behavioral effects[3–9], it was unclear whether and how these findings could be translated to humans. In **Chapter 2** I address this foundational question by translating a protocol from macaque research[7] to humans, targeting the

frontal eye fields (FEF), a well-characterized region whose causal role in contralateral saccade generation is firmly established[10–13].

Thirty-five participants performed a saccadic decision task in which two visual stimuli were presented asynchronously on either side of fixation. Participants were instructed to saccade as quickly as possible to the stimulus that appeared first. TUS was applied during stimulus presentation, targeting either the left or right FEF or the left- or right-hand motor cortex M1 (active control condition). FEF TUS induced a significant increase in the selection of contralateral saccades, under conditions of low sensory evidence (short stimulus onset asynchronies), directly replicating findings from a previous study in macaques and indicating that FEF TUS has net excitatory effects on saccade selection in humans. Critically, these effects were specific to trials with short stimulus onset asynchronies, where participants were uncertain about which target appeared first, and sensory evidence was ambiguous. When the correct choice was obvious (long stimulus onset asynchronies), TUS had no effect. This pattern demonstrates that TUS nudges or biases ongoing neural computations rather than generating behavior de novo.

Participants' inhibitory tone in FEF, measured using magnetic resonance spectroscopy (MRS), predicted interindividual differences in TUS response. Specifically, higher baseline GABA+ levels in the left FEF were associated with a stronger rightward bias on sham trials; following TUS, participants converged toward a common behavioral state, with the strongest effects observed in those who initially had lower GABA+ levels.

Importantly, stimulation of the M1 as an active control site did not induce significant differences in contralateral saccades, supporting the specificity of the observed effects to the FEF and ruling out potential confounds. Finally, a masking assessment confirmed that while auditory and somatosensory confounds were present, their effects were equivalent across FEF and M1 conditions and thus could not account for the observed dissociation.

These findings establish that TUS is an effective, spatially and temporally precise method for modulating choice behavior in humans. The demonstration that individual differences in baseline GABAergic inhibitory tone predict response magnitude provides critical insights into why neuromodulation efficacy varies across individuals. Baseline neurochemical state shapes susceptibility to stimulation, transforming interindividual variability from noise into systematic, predictable variation that can potentially be leveraged to personalize protocols. This work bridges human and animal research and paves the way for employing TUS to modulate cognition, specifically choice behavior, in humans.

Chapter 3

Having established that TUS can be used successfully to modulate behavior in humans, I next sought to identify the neural substrates of motivation-action control that would serve as targets for causal intervention. The ventral striatum (nucleus accumbens, NAcc) and amygdala are central to motivational processes[14–16] but in humans their respective contributions to Pavlovian control and instrumental learning remained incompletely understood. In **Chapter 3**, I used functional magnetic resonance imaging (fMRI) while 29 healthy adults performed a motivational Go/NoGo (MGNG) task in which cue valence (Win reward vs. Avoid loss) was orthogonalized with required action (Go vs. NoGo), and repeated introduction of novel cues across conditions with probabilistic feedback kept participants in an ongoing learning regime.

Behaviorally, participants showed clear instrumental learning, emitting more and faster Go responses when Go was correct, and a robust Pavlovian bias, with more and faster responses for Win than Avoid cues. Novelty sensitivity was also evident: the outcome on the very first encounter with a cue had a disproportionately strong impact on subsequent stay/switch choices compared with later repetitions, suggesting heightened learning from initial experiences.

At the neural level, BOLD responses at cue onset were higher for Win than Avoid cues in NAcc and ventromedial prefrontal cortex (vmPFC), consistent with a role in anticipatory value coding that is a prerequisite for Pavlovian biases to emerge[17–19]. In contrast, amygdala activity did not differentiate between Win and Avoid cues at the time of cue presentation, despite extensive prior work linking the amygdala to valence processing[15,16,20–22]. At outcome, NAcc, amygdala, and vmPFC all responded more strongly to desired than non-desired outcomes, consistent with a shared role in outcome evaluation and instrumental learning. Exploratory analyses of novelty further revealed elevated amygdala responses for the first presentation of a cue, irrespective of valence, suggesting that the amygdala is particularly engaged in boosting novelty-sensitive cue processing to support fast initial learning.

Together, these findings suggest a differentiated but coordinated division of labor between the amygdala and NAcc in motivational action control. The amygdala appears to particularly support fast initial learning, tagging new cues with affective value. The NAcc supports both anticipatory Pavlovian bias, as indexed by stronger cue-locked responses to Win versus Avoid cues, and both NAcc and amygdala are involved in outcome processing, putatively supporting learning of action or stimulus values based on feedback. These findings identify the NAcc and amygdala as key nodes in the motivation-action coupling network, with shared and distinct contributions, setting the stage for causal investigation in **Chapter 4**.

Chapter 4

In **Chapter 4**, I combined the insights from the preceding chapters to directly test the causal roles of the NAcc and amygdala in motivation-action coupling. Twenty-nine healthy participants completed three intervention sessions in a within-subject design, receiving offline TUS to either NAcc, amygdala, or sham stimulation before performing the MGNG task. To characterize the behavioral effects of TUS, I employed hierarchical reinforcement-learning models that allowed us to quantify Pavlovian biases and instrumental learning dynamics, including a novelty-learning bonus and separate updating from positive and negative prediction errors.

The results reveal both shared and distinct contributions of NAcc and amygdala to motivational control. Amygdala stimulation specifically boosted novelty-driven learning: it increased the learning rate for the very first encounter with a cue, amplifying the impact of initial outcomes on subsequent choices. This aligns with prior research linking the amygdala to novelty detection and salience processing[23,24]. It also converges with my own fMRI findings in **Chapter 3**, where amygdala BOLD responses were selectively elevated for the first presentation of a cue.

NAcc stimulation, by contrast, enhanced both Pavlovian bias and instrumental learning. This enhanced Pavlovian bias – enhancing action invigoration for reward-predictive cues and inhibition for loss-predictive cues – is consistent with the fMRI findings in **Chapter 3** showing NAcc cue-locked responses for Win versus Avoid cues, and aligns with extensive animal literature demonstrating that NAcc dopamine couples reward prediction to behavioral activation[17,25,26]. Notably, this bias was particularly strong early in learning but declined more steeply over trials. Second, NAcc TUS boosted instrumental learning rates, especially from positive prediction errors, enabling participants to learn faster which actions were optimal. This enhanced learning progressively overrode initial Pavlovian tendencies when they conflicted with task contingencies, such that initial bias was stronger under NAcc TUS relative to sham, but for later cue presentation the bias was weaker. These learning rate effects are consistent with animal lesion studies showing that NAcc damage impairs learning from positive outcomes[27] and with recent human TUS work demonstrating that NAcc stimulation preferentially enhances feedback sensitivity for wins[28].

Together, these findings provide causal evidence for temporally distinct contributions of NAcc and amygdala to motivated learning and action selection in humans. The amygdala plays a specialized role in rapid learning about novel stimuli, flagging new cues for enhanced initial updating. The NAcc, by contrast, supports both Pavlovian control and instrumental learning, translating motivational signals into action biases

while simultaneously enabling faster refinement of action through experience. This dissociation opens new avenues for understanding how dysfunction in these circuits might contribute to psychiatric disorders characterized by altered motivation and learning and suggests that targeted neuromodulation of these specific regions could form the basis for novel therapeutic interventions.

Synthesis

Taken together, this thesis provides the causal evidence for the distinct contributions of the nucleus accumbens and amygdala to Pavlovian and instrumental control in healthy humans. By establishing that TUS can effectively modulate deep brain structures and demonstrating its utility for probing causal brain-behavior relationships, this work opens a new chapter in human decision neuroscience. Where previous research relied on correlational neuroimaging or drew inferences from animal models, we can now directly test whether specific brain regions are causally involved in specific aspects of motivated behavior. Beyond advancing our fundamental understanding of how the brain transforms motivation into action, these findings have translational potential: they identify specific neural targets and mechanisms that could inform the development of more precise interventions for disorders of motivation and behavioral control.

II. MOTIVATION-ACTION COUPLING IN THE AMYGDALO-STRIATAL CIRCUIT

As you are reading this thesis and made it to the final chapter, your phone screen might have just lit up with a notification. In most moments of the day, this would be a welcome cue: the prospect of an interesting message, a moment of connection, a pleasant distraction. The urge to check it is not irrational, it reflects a seemingly hardwired tendency to approach stimuli that have reliably predicted reward. This is a Pavlovian bias: an automatic, valence-driven response to a reward-predictive stimulus that, most of the time, serves you well. But right now, you are nearly at the end of a chapter. Experience has taught you that giving in to distractions at this particular moment will cost you more than it gains – that the message can wait, and that finishing what you started is worth more. This is instrumental control: flexible, learned knowledge about which actions lead to which outcomes in a given context. Whether you reach for the phone depends on the interplay between these two systems.

The intuitive framing is that these two systems operate in tension. Indeed, many

cognitive models of motivated action implicitly treat Pavlovian and instrumental control as competing systems whose outputs conflict, requiring resolution – typically through top-down cognitive control mechanisms that suppress the maladaptive Pavlovian impulse in favor of the instrumentally optimal response[29,30]. One might therefore expect that selectively strengthening one system, would tip the balance toward it at the expense of the other. But what if this tension framing does not tell the whole story? In this thesis, I investigated the interplay between Pavlovian and instrumental control and the neural mechanisms that underlie their contribution to motivation-action coupling.

Thesis

Based on my thesis research, I make three main claims.

First, the NAcc is the key node for goal-directed behavior broadly: it supports both instrumental learning – updating which actions are worth taking in a given context through stimulus-action-outcome associations driven by prediction error signals – and the expression of Pavlovian bias, particularly when cues are directly actionable (**Figure 1A**). In this sense, the NAcc coordinates both learned action values and motivational biases that together shape behavior, receiving distinct value signals via anatomical projections from the basolateral amygdala[14,31–33] and cortical regions such as the ventromedial prefrontal cortex[34,35].

Second, Pavlovian and instrumental control are better understood as collaborative instead of competing contributors to goal-directed behavior. Rather than two systems vying for control, both feed into a shared integration process in the NAcc.

Third, the amygdala plays a distinct but complementary role. Rather than computing and storing action values, it may primarily signal motivational significance or salience to downstream structures such as the NAcc (**Figure 1B**). Here, its contribution to motivation-action coupling lies particularly in adaptively prioritizing novel stimuli – flagging initial encounters with motivationally relevant cues as significant and thereby biasing learning in the NAcc via its projections to the ventral striatum[14,31–33].

In the following sections, I develop and support these claims in turn. I first discuss the evidence for a collaborative rather than competitive relationship between Pavlovian and instrumental control. I then turn to the role of the NAcc and amygdala in instrumental learning, before examining the amygdala's contribution to Pavlovian conditioning versus Pavlovian bias expression. Afterwards, I address the NAcc's role in human Pavlovian bias and the importance of cue actionability. Finally, I discuss the

amygdala's role in novelty-driven learning and its implications for how the amygdala-striatal circuit prioritizes new information.

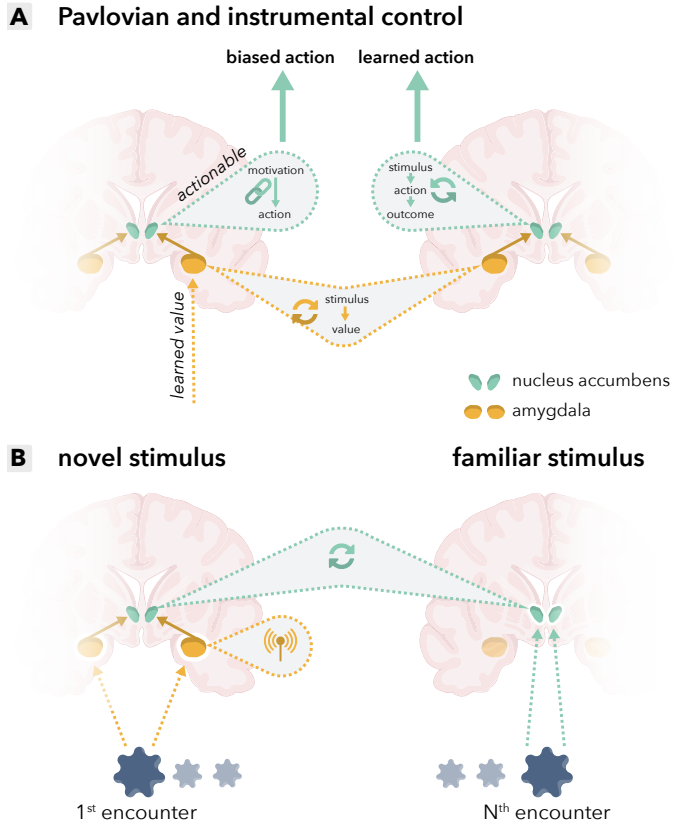


Figure 1 | A proposed amygdalo-striatal circuit model for motivation-action coupling. (A) Left: The NAcc and amygdala (and vmPFC, not depicted here) contribute to value processing but encode different aspects. The basolateral amygdala encodes learned value associations and the NAcc represents actionable value (i.e., motivational signals directly coupled to executable actions). These signals converge on the NAcc, which translates motivation into action through direct and indirect pathways to motor output. Right: the NAcc integrates these value signals with dopaminergic prediction errors from midbrain (VTA/SNc) to update stimulus-action-outcome associations, thereby supporting flexible goal-directed behavior. **(B)** Novelty: When stimuli are novel, the amygdala flags them as salient and sends amplifying signals to the NAcc, biasing it toward enhanced learning from initial outcomes. Once cues become familiar, the amygdala novelty response diminishes and the NAcc takes over stimulus-action-outcome updating, remaining continuously engaged in environments with changing contingencies.

Pavlovian and instrumental control: parallel contributions to action selection

The case for competition

One finding from **Chapter 4** that might appear surprising at first glance is that NAcc stimulation simultaneously enhanced both Pavlovian bias and learning rate. The intuition behind the surprise is familiar: Pavlovian tendencies represent automatic, valence-driven responses that can interfere with flexible goal-directed behavior[36]. In this framing, Pavlovian bias is something of an impulse – the urge to approach reward-predictive cues or avoid threat-predictive ones – that must be overcome when it diverges from the instrumentally optimal action.

This competition framing is not without empirical basis. A broad literature on cognitive control has established that frontal mechanisms, particularly midfrontal cortex, are recruited when competing response tendencies are simultaneously active[37–41]. This conflict-detection architecture has been proposed to extend to the domain of motivated action, where Pavlovian and instrumental response tendencies can diverge. Consistent with this, MGNG paradigms have shown that midfrontal theta power increases specifically on bias-incongruent trials biases[30]. At the pharmacological level, manipulation of prefrontal dopamine reduced Pavlovian bias, further implicating frontal systems in arbitrating between Pavlovian and instrumental control[29]. Together, these findings have been taken to imply that Pavlovian and instrumental systems compete, with frontal mechanisms recruited to resolve that competition.

Towards a collaborative view

Yet a closer look suggests this tension framing may be incomplete. Both Pavlovian and instrumental control ultimately serve the same goal: achieving the best possible outcome. They differ not in aim but in how they operate. Pavlovian control is fast, computationally cheap, and evolutionarily hardwired. It relies on automatic tendencies that generally align with broad ecological principles: act in the face of reward, refrain in the face of punishment[36,42]. Importantly, Pavlovian motivational processes are not entirely divorced from goal-directed behavior: reward-predictive cues invigorate behavior in proportion to the current desirability of the predicted outcome, suggesting that Pavlovian responses reflect representations of valuable outcomes rather than purely reflexive reactions[26,43]. Instrumental control is more flexible but more costly, requiring trial-and-error learning to determine the optimal action in a given context. Early in learning, when instrumental evidence is sparse, Pavlovian tendencies offer a useful default that guides initial behavior. As experience accumulates, instrumental

values become increasingly reliable and their influence grows accordingly[44]. Rather than competing systems that take turns dominating, Pavlovian and instrumental control may thus function as complementary contributors within a broader goal-directed architecture, with their relative weighting shifting fluidly across the course of learning[45] (**Figure 2**).

The data in **Chapters 3** and **4** are consistent with the collaborative view. In **Chapter 4**, NAcc TUS simultaneously increased Pavlovian bias parameters and enhanced instrumental learning parameter rates – amplifying both rather than shifting the balance between them. This is not what one would predict if the NAcc selectively serves one system, but it is precisely what one would expect if the NAcc is a common hub through which both Pavlovian value signals and instrumental prediction errors are translated into action. Notably, a higher learning rate does not straightforwardly imply better or stronger learning – it reflects a greater sensitivity to recent outcomes, which – depending on the volatility of the environment – can have different benefits. Moreover, unlike the conflict-related prefrontal recruitment seen in classic cognitive control tasks such as Stroop and flanker[37,41], and in contrast to what has been observed in related paradigms[30], I did not observe significant prefrontal engagement on correct bias-incongruent trials in **Chapter 3**. This absence suggests that, in our continuous learning paradigm, participants rarely experienced genuine conflict: true conflict requires that the correct instrumental response is known well enough to stand in opposition to automatic tendencies. When participants are still learning what to do, incongruent trials simply reflect uncertainty rather than the suppression of a clear automatic tendency. This is fundamentally different than Stroop or Flanker tasks where there are clear instructions about the correct response, which then can result in clear conflict. Incorrect guesses based on a strong Pavlovian prior may therefore reflect this early-learning dynamic, in which Pavlovian estimates carry more weight because instrumental evidence is still sparse, rather than requiring active top-down conflict resolution.

A shared striatal substrate

This reliability-weighting arbitration can be understood at the neural level through the striatal architecture underlying Go and NoGo responding (**General Introduction: Box 1**). Both Pavlovian and instrumental value signals converge on the same basal ganglia circuitry: appetitive signals potentiate the direct pathway, promoting behavioral activation, while aversive signals engage the indirect pathway, promoting inhibition[36,46]. Critically, both Pavlovian and instrumental controllers provide input to this shared machinery. In this framework, the question is not which system ''wins'', but rather how much each contributes to the net drive toward action or inaction. This is consistent

with evidence accumulation accounts in which Pavlovian value biases the starting point or drift rate of the decision process in favor of the bias-congruent response, with instrumental values modulating that same process in proportion to their current reliability, consistent with uncertainty-based arbitration principles originally formalized for instrumental controllers[44,47]. The output – Go or NoGo – thus reflects the integrated influence of both systems, with their relative contributions shifting dynamically over the course of learning.

I therefore propose that NAcc does not selectively serve either Pavlovian or instrumental control, but rather functions as a common hub for coupling motivational value to action – one through which both Pavlovian value signals and instrumental prediction errors are translated into behavior via overlapping striatal machinery[19]. Against this backdrop, the finding that NAcc stimulation enhanced both Pavlovian bias and instrumental learning parameters becomes more interpretable: TUS broadly increased the sensitivity of this shared motivational-to-action circuit, amplifying both contributors rather than shifting the balance between them. Pavlovian and instrumental control, in this framing, are not rival systems competing for dominance, but complementary contributors to a single goal-directed architecture – one whose output, whether to act or withhold, reflects the continuously weighted influence of both.

Learning from stimuli, actions and outcomes

Instrumental learning refers to the process of learning which actions lead to which outcomes, allowing organisms to adapt their behavior based on experienced consequences. This learning is thought to be driven by prediction errors (mismatches between expected and obtained outcomes) that are encoded by midbrain dopamine neurons and transmitted to striatal targets[48,49]. Extensive animal research using lesions, electrophysiology, and optogenetics has established that the NAcc plays a crucial role in instrumental learning[14,27,50,51]. Human neuroimaging studies converge with these findings, showing that NAcc BOLD responses track prediction errors during instrumental tasks[19,52,53]. The amygdala's involvement in instrumental learning, however, is less well characterized. While the amygdala is clearly critical for Pavlovian conditioning – learning stimulus-outcome associations[15,16,22] – its role in updating action-specific values (stimulus-action-outcome learning) remains debated. Some animal studies suggest that amygdala lesions impair learning from rewarding outcomes, particularly when cues are novel or contingencies change[27,54,55], but these effects could reflect disrupted stimulus-outcome processing rather than direct action-value updating. In **Chapters 3** and **4**, I sought to clarify this division of labor between the NAcc and amygdala by combining fMRI measures of outcome processing

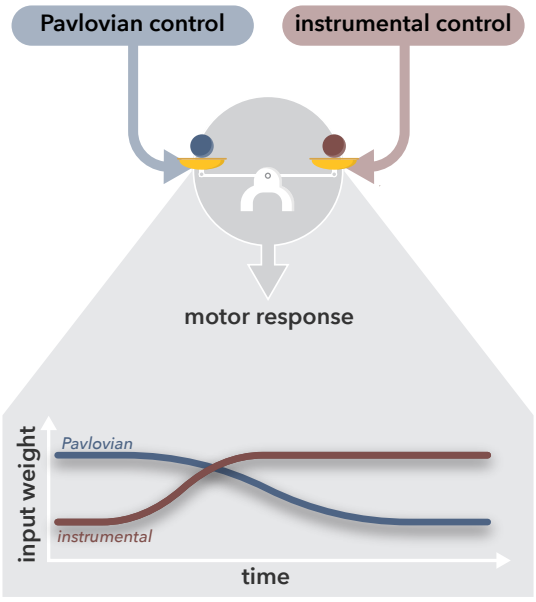


Figure 2 | Pavlovian and instrumental control as parallel contributors to action selection. Both Pavlovian and instrumental control provide input to a shared decision process, here depicted as a balance scale, whose output determines the motor response. The relative weighting of each system shifts dynamically over the course of learning (bottom): early on, Pavlovian input carries more weight because instrumental evidence is sparse; as trial-and-error learning accumulates, instrumental control gains influence and increasingly drives behavior. Rather than competing systems, Pavlovian and instrumental control thus function as complementary contributors within a broader goal-directed architecture, with their balance determined by the relative reliability of each estimate at a given moment in time.

with causal TUS interventions that directly test each structure's role in updating action values.

Together, the fMRI (**Chapter 3**) and TUS (**Chapter 4**) results reveal dissociable roles for NAcc and amygdala on learning dynamics. In **Chapter 3**, both the NAcc and amygdala showed stronger outcome-locked BOLD responses for desired over non-desired outcomes, a pattern shared also with vmPFC. This convergent outcome sensitivity has been observed in previous human fMRI studies of instrumental learning tasks[19,52]. At first glance, this shared recruitment might suggest that the NAcc and amygdala play a comparable role in instrumental learning. However, the causal evidence from **Chapter 4** tells a different story. Following NAcc stimulation, learning rates were enhanced, particularly for positive prediction errors. This pattern of enhanced learning from wins aligns with earlier findings from animal lesion studies showing that NAcc damage impairs learning from positive outcomes[27] and from recent human NAcc TUS work demonstrating enhanced feedback sensitivity for rewards[28]. Amygdala stimulation, by

contrast, affected learning only for novel cues, with no reliable influence on subsequent updating as cues became familiar.

Outcome-locked BOLD activity does not straightforwardly equal instrumental learning. What fMRI captures is neural activity at the moment of feedback processing, encoding the reward value of a presented sensory stimulus. However, this signal can support at least two computationally distinct learning processes. First, reward encoding is a requirement to support instrumental learning, i.e., *stimulus-action-outcome* learning, updating the value of a specific action in a specific context. Second, reward coding is necessary for *stimulus-outcome* learning (updating whether a stimulus is predictive of future rewards, regardless of what action was taken). As I will discuss below, the amygdala appears strongly involved in the latter, which is what the observed outcome coding in the amygdala may reflect. In contrast, the NAcc outcome coding may rather support instrumental learning of action values. The present findings highlight precisely how causal interventions in **Chapter 4** complement the fMRI findings of **Chapter 3**: where neuroimaging reveals that both NAcc and amygdala respond to outcomes, TUS allows us to ask which of these responses eventually drives subsequent learning, and of what kind.

Amygdala's involvement in Pavlovian conditioning versus Pavlovian bias

One initially surprising finding from **Chapter 3** was the absence of cue-valence BOLD encoding in the amygdala – a result that diverges from earlier human fMRI studies in which robust amygdala valence responses at cue onset were observed[21,56]. To understand this discrepancy, it is useful to distinguish between two related but separable processes: Pavlovian conditioning – the process through which neutral stimuli acquire motivational significance through repeated pairing with rewarding or punishing outcomes – and Pavlovian bias – the seemingly hardwired tendency to invigorate action in the face of reward and suppress it in the face of punishment (see **General Introduction**). The amygdala's role may be specific to the former.

In animal literature, the amygdala, and particularly its basolateral complex, is well established as critical for Pavlovian conditioning[15,16,22]. Lesions of the basolateral amygdala impair the formation of new cue-outcome associations and disrupt the capacity of conditioned cues to influence ongoing behavior[14,20,50,55]. Human neuroimaging studies broadly converge with this framework, showing amygdala responses to conditioned reward and threat cues[21,56–59]. Moreover, recent causal evidence in humans shows that ultrasound stimulation of the amygdala enhanced learning of conditioned responses[60]. Crucially, the key feature shared across these

paradigms is that cue-outcome associations had to be learned – the amygdala was engaged precisely because new Pavlovian value was being acquired or retrieved from experience.

This is where our paradigm differs fundamentally. In our task, cue valence was *instructed* through the color of the stimuli on every trial, even on first encounter, removing any need to track or retrieve Pavlovian value from *experience*. Direct evidence for this *instructed-versus-experienced* distinction comes from Atlas et al. (2016), who demonstrated that the amygdala is specifically engaged when Pavlovian values are acquired through direct experience, rather than through explicit instruction[61]. Using a fear conditioning reversal paradigm, they showed that when participants were *instructed* about a contingency reversal, skin conductance responses and NAcc activity updated immediately (before any new reinforcement had occurred) whereas amygdala responses did not. The amygdala continued to track the old, experientially learned contingencies until participants had directly *experienced* the new pairings. This reveals a clear dissociation: instructed values, implemented via prefrontal cortex and NAcc, can override the amygdala's experientially learned representations to guide behavior, while the amygdala itself updates only through experience. The absence of amygdala valence coding in **Chapter 3** likely reflects precisely this: when instructed information continuously and unambiguously signals cue valence, the amygdala's experiential tracking is rendered redundant and may be actively down-regulated.

This also raises a broader possibility worth considering: it may be that the amygdala's engagement depends not just on the *source* of value information, but on whether top-down instructed and bottom-up experiential signals are *aligned*. When instruction and experience converge, the amygdala may disengage, as the experiential tracking it provides offers no additional information. When they conflict – as at the moment of reversal in the Atlas paradigm[61] – the amygdala may conversely be recruited precisely because the discrepancy signals a need to update experiential representations. This alignment account would generate the prediction that amygdala responses should be largest when instruction and experience mismatch, and smallest when they are fully redundant, as in our task.

Notably, this alignment account may also offer a parsimonious explanation for the outcome-locked amygdala responses observed in **Chapter 3**, where the amygdala responded more strongly to desired than non-desired outcomes. The increased outcome-locked activity may reflect a residual conditioning process occurring in parallel. Even when participants already know which cues predict reward or loss through instruction, experiencing actual outcomes may still engage Pavlovian learning mechanisms that update or reinforce these associations. However, because the

instructed information already provides a strong prior, this experiential learning may be overridden for generating appropriate cue-evoked responses or Pavlovian biases.

Taken together, these results suggest that the amygdala's contribution to Pavlovian control is present for situations where valence is learned from experience, but that such learning is overridden when explicitly signaled. To causally test this distinction, future studies could combine TUS interventions with paradigms that systematically vary whether cue values are learned or instructed, directly testing whether amygdala stimulation differentially affects Pavlovian bias under these two conditions.

The role of nucleus accumbens in motivated action

The findings reported in this thesis provide converging evidence that the NAcc plays a causal role in human Pavlovian control. **Chapter 3** revealed robust valence encoding at cue onset, with stronger BOLD responses to Win than Avoid cues. **Chapter 4** provided causal evidence that NAcc TUS amplified Pavlovian bias, enhancing action invigoration for reward cues and action suppression for loss cues. These results align with the animal literature[26,49,51,62–67]. Yet they appear to diverge from several previous human fMRI studies using MGNG paradigms[52,56,68,69]. What could explain this discrepancy?

The discrepancies

In the animal literature, the role of the NAcc in Pavlovian bias has been predominantly established using Pavlovian-instrumental transfer (PIT) paradigms. In PIT, animals are first trained to associate cues with reward or punishment outcomes (Pavlovian conditioning), and separately trained to perform instrumental actions for reward. When the reward-paired cues are then presented during ongoing instrumental responding, they invigorate action – demonstrating that motivational value encoded by a cue can directly influence behavior, independently of any learned cue-action contingency (see **General Introduction**). When the NAcc was lesioned, this invigorating effect of reward cues on instrumental action was abolished[26,50,51,63]. Animals continued to perform the instrumental response, but reward-predictive cues no longer enhanced their responding, demonstrating that the NAcc is necessary for translating motivational value into behavioral activation[26,50,51,63]. Electrophysiological recordings show that NAcc neurons respond to reward-predictive cues and encode the motivational value that drives this action invigoration[49,65]. More recently, optogenetic studies have demonstrated that activating NAcc dopamine inputs is sufficient to enhance action vigor in response to motivational cues[64,66,67]. Together, these findings establish the NAcc as a critical interface through which motivational value can influence action selection and behavioral activation – the neural implementation of Pavlovian bias.

In humans, however, the evidence for NAcc involvement in Pavlovian responses has been mixed. Several neuroimaging studies show robust NAcc responses to reward-predictive cues, consistent with the animal literature. For example, in monetary incentive delay tasks, NAcc BOLD activity scales with anticipated reward magnitude and correlates with reaction time speeding[18,70], and in PIT paradigms NAcc activity increases when reward-conditioned cues are presented in the background of instrumental choices[21,59]. However, in tasks more directly comparable to the one I have used in this thesis – where action requirements and cue valence are orthogonalized – the picture is less clear. Studies using versions of the MGNG paradigms have generally not observed NAcc valence encoding at cue onset, despite robust behavioral Pavlovian biases[52,56,68,69]. Instead, striatal activity at cue presentation in these paradigms tends to reflect the action that participants will subsequently perform rather than the valence of the cue itself - encoding cue-evoked action values rather than pure Pavlovian valence signals.

Actionability

A fundamental difference between our MGNG task and earlier human versions is that in our paradigm, cues were directly *actionable*. Before elaborating on the experimental distinction, it is worth pausing on what "actionability" means in a broader sense. Consider the goal of addressing climate change. Most people care about this deeply, yet the goal itself is vast and diffuse – it is not clear what single action to take right now to stop climate change. As a result, the motivational pull of the goal rarely translates into immediate behavior. Contrast this with a concrete, daily choice about reducing one's carbon footprint: when heading to work, take the bicycle or the car? This decision is immediately actionable – it is tied to a specific moment, a specific action, and a specific outcome.

In our MGNG paradigm in **Chapters 3** and **4**, participants had to prepare and execute their response while the cue was visible, and responses were no longer possible once the cue disappeared. This created a tight temporal coupling between the motivational prospect evoked by the cue (e.g., winning a reward or avoiding a loss) and the action required to achieve that outcome. In contrast, several earlier versions of the task presented cues in one phase and only allowed responses after the cue disappeared or even allowing responses during a separate, jittered release cue[68,69], creating a temporal separation between motivational prospect and the action required to achieve it. Additionally, our task featured a single response option (one Go button or one NoGo withhold), whereas other versions offered multiple Go response alternatives[56], potentially distributing action preparation across options and diluting the motivation-action coupling signal.

Anatomically, the NAcc occupies a central position within cortico-basal ganglia networks that link motivational information to action selection (**General Introduction: Box 1**). What was observed in earlier MGNG studies[68,69] – that striatal signals encode action values rather than raw Pavlovian valence – may reflect precisely this function: the NAcc integrating motivational value with available action options. When no immediate action is available, however, this transformation may not be triggered in the same way.

This leads to a speculative hypothesis: NAcc involvement in Pavlovian control depends critically on cues being directly actionable. When participants must deliberate across multiple action options[52,56], or when cue presentation and action execution are temporally separated[68,69], the motivational-to-motor transformation may rely less heavily on the NAcc, perhaps recruiting more prefrontal or deliberative circuits instead. This interpretation aligns with the NAcc's established role in action invigoration and effort allocation[71], processes most relevant when a clear and immediate action opportunity is available.

In contrast to the NAcc findings, the vmPFC results in **Chapter 3** are fully consistent with earlier human MGNG fMRI studies: both our paradigm and previous versions of the task show stronger vmPFC BOLD responses to Win than Avoid cues at cue onset[56,68,72]. This result aligns with its established role in representing stimulus value across large body of literature[73–77]. This convergence suggests that vmPFC may subserve a more general form of value encoding, evaluating the value of a stimulus without necessarily considering actions that need to be taken to attain this value/goal. In this account, vmPFC encodes "how good is this stimulus?" while the NAcc encodes "how will my actions lead to this valuable goal?".

The hypothesis that the NAcc preferentially supports motivational control when goals are directly actionable remains speculative. Future work could test this prediction by directly manipulating whether cues are immediately actionable or only prospectively relevant for later decisions. For example, tasks could include both actionable and non-actionable cue conditions within the same paradigm, allowing researchers to determine whether NAcc activity and causal perturbations of this region differentially influence behavior depending on whether motivational cues can be immediately translated into action. Such experiments would clarify the boundary conditions under which the NAcc contributes to motivated action and help reconcile apparent discrepancies between animal findings and human neuroimaging studies.

The amygdala's specialized role in novelty-driven learning

Prioritization of learning for novel stimuli

A core function of the amygdala across species is to prioritize attentional and learning resources toward stimuli that are potentially important for survival. This function is not limited to any single category of stimulus: rewarding, threatening, novel, and otherwise salient events all engage the amygdala, forming a broad class of behaviorally relevant inputs that warrant preferential processing[15,23,24,78–80]. Signaling the motivational or predictive value of cues, as discussed in the preceding sections, is one expression of this broader prioritization function. But the amygdala's sensitivity to novelty and salience more generally reflects the same underlying principle: any stimulus whose consequences are unknown or uncertain constitutes a potential opportunity or threat, and the brain's default response is to treat it as important until proven otherwise. It is this general-purpose salience and novelty detection function, and its causal consequences for learning, that the findings in **Chapters 3** and **4** speak to directly.

A large body of work demonstrates that the amygdala responds more strongly to novel than to familiar stimuli, regardless of whether those stimuli are rewarding, threatening, or neutral[23,24,79,80]. Moreover, when outcome probabilities shift or stimulus-outcome mappings reverse, amygdala lesions slow behavioral updating[81–84], suggesting that the same rapid detection mechanism that flags novelty also supports fast learning when things change. Both reflect a common capacity: detecting when a stimulus carries important, uncertain information that warrants enhanced attention and learning.

The findings in **Chapters 3** and **4** provide converging and causal evidence that the amygdala boosts learning specifically for novel stimuli. In **Chapter 3**, amygdala BOLD responses were selectively elevated when participants first encountered a cue, regardless of its valence – consistent with prior human neuroimaging showing enhanced amygdala responses to novel compared to familiar stimuli[23,79,80]. Critically, this novelty response was highly specific to the very first presentation: by the second encounter, amygdala activity had already returned to baseline. Moreover, it occurred at cue onset rather than at outcome, suggesting that the amygdala rapidly detects novelty and flags new stimuli as salient before any learning has occurred. In **Chapter 4**, amygdala TUS specifically increased learning for the first encounter with a cue, amplifying the impact of the initial outcome on subsequent choice – an effect also evident in the empirical data, where amygdala stimulation selectively altered win-stay/lose-shift behavior between the first and second encounter with a cue. Behaviorally, the consequence is a strengthened primacy effect: the outcome from the very first trial with a novel cue exerts a disproportionately large influence on subsequent choices relative to later encounters with the same cue[85].

Novelty-driven learning through the amygdalo-striatal circuit

The present findings thus offer a window into a broader principle: the amygdala implements adaptive prioritization of attention and learning based on behavioral relevance, and novelty is one prominent trigger of this mechanism. The same prioritization likely extends to other forms of salience, such as reward-predictive cues, stimuli associated with threat, or any event that carries uncertain but potentially significant consequences[15,23,24,79,80]. At the circuit level, the implementation of this novelty-driven learning boost runs likely through the well-characterized dense projections from the basolateral amygdala to the NAcc[14,31–33]. Via these projections, a novelty or saliency signal from the amygdala could prime the NAcc for enhanced updating from the immediately following outcome – effectively biasing the NAcc toward heightened plasticity when a stimulus is encountered for the first time.

Critically, the present data provide indirect but compelling evidence that this is precisely how the amygdala's novelty contribution is implemented. The key observation is that NAcc TUS, like amygdala TUS, increased the novelty-specific learning rate – but NAcc TUS additionally increased the learning rate for positive prediction errors throughout the task, an effect not seen after amygdala stimulation. This pattern is most parsimoniously explained by two dissociable processes: the amygdala provides a novelty bias signal to the NAcc, and the NAcc itself performs the actual learning rate updating. On this account, amygdala TUS enhances the novelty signal fed into the NAcc, boosting initial learning. NAcc TUS, by contrast, amplifies the learning computation itself – affecting both novelty-gated and ongoing updating, because both depend on NAcc mechanisms. Crucially, the amygdala TUS effect on the novelty learning rate does not imply that the amygdala itself computes the update; rather, it implies that the amygdala biases the NAcc, which then implements it. A crossed disconnection design – in which TUS is applied unilaterally to the NAcc and amygdala on the same versus opposite hemispheres – could directly test whether the interaction between these two regions causally contributes to the novelty-driven learning rate enhancement. If the amygdala's novelty contribution is implemented through its projection to the ipsilateral NAcc, modulating same-hemisphere connectivity should selectively attenuate the novelty learning rate boost, while sparing the NAcc ongoing learning effects.

In my task, continued NAcc updating is particularly necessary because the environment is probabilistic and includes occasional reversals. The NAcc is well-suited to maintain ongoing learning in such volatile, uncertain environments[27,83,86]. Even after participants form initial stimulus-action-outcome associations, the stochastic feedback and reversal contingencies ensure there is always more to learn, keeping the NAcc

engaged throughout. In contrast, animal studies suggest that in stable, deterministic environments the NAcc may gradually cede control to more habitual, dorsal striatal circuits once associations are well established[87,88]. The noisy, dynamic environment in the present task likely prevents this handoff. Future studies using deterministic outcomes without reversals could test this directly, by examining whether NAcc TUS effects on learning decay once the task environment stabilizes.

Taken together, these empirical and theoretical considerations converge on a coherent picture. The data in **Chapters 3** and **4** suggest that the amygdala does not simply encode valence in a static manner. Rather, it dynamically modulates learning by signaling when a stimulus is important – because it is new, uncertain, or otherwise salient – and directing learning resources accordingly. This biases the NAcc toward heightened learning from the immediately following outcome, boosting the initial update of the stimulus-action-outcome association (**Figure 1C**).

III. TUS EFFECTS

State dependency of neuromodulation effects

Brain stimulation does not affect everyone in the same way. Repetitive transcranial magnetic stimulation (rTMS) as a treatment method yields response rates of only 50-60% with full remission in just one third of patients[89,90]. The same holds for pharmacological interventions: only about one third of patients with major depression achieve remission following first-line antidepressants, although remission rates increase after multiple treatment steps[91,92]. For patients, this is not merely a statistical inconvenience – each failed treatment trial erodes patient expectations of recovery, reinforces negative cognitions, and may condition the brain not to respond: remission rates decline with every additional treatment step, and the number of prior failed attempts is itself a negative predictor of subsequent response[91,93–95]. This interindividual variability reflects a fundamental reality: each person's brain has a unique neurochemical profile, anatomical structure, and baseline physiological state. Understanding the sources of this variability is therefore crucial for both basic cognitive neuroscience research and clinical translation. Without the ability to predict who will respond to a given intervention and tailor protocols accordingly, therapeutic applications will remain limited by their inconsistency. In **Chapter 2** of this thesis, I took a first step toward addressing this challenge by demonstrating that baseline inhibitory tone predicts individual differences in TUS response, providing proof-of-principle that interindividual variability is not random noise but rather systematic variation that can be measured, predicted, and potentially leveraged to optimize treatment.

Interindividual variability

As comprehensively reviewed by Huang and colleagues (2017)[96], individuals exposed to identical stimulation protocols – including transcranial magnetic and electrical stimulation – can show markedly different, or even opposite, behavioral and physiological outcomes. This variability arises from multiple sources: differences in neurochemical balance, age, sex, skull morphology, brain anatomy, and critically, the baseline physiological state of the targeted neural circuits[97–102]. In **Chapter 2** of this thesis, I provide the first demonstration in humans that baseline inhibitory tone predicts individual differences in TUS efficacy. Using magnetic resonance spectroscopy to measure GABA+ concentrations in the FEF, we found that participants' baseline inhibitory tone was associated with their intrinsic saccadic bias: higher GABA+ levels in the left FEF correlated with stronger rightward bias during sham (no stimulation) trials. Critically, following TUS, this relationship disappeared. Rather than amplifying pre-existing individual differences, TUS brought participants toward a common behavioral state, with the strongest effects observed in individuals who initially showed the weakest intrinsic bias.

This normalizing pattern is reminiscent of homeostatic plasticity: just as biological systems maintain physiological variables within functional bounds through compensatory feedback[103,104], TUS appears to modulate neural activity within a physiological range rather than extending beyond it. Participants already showing a strong intrinsic bias were not further biased following stimulation, suggesting a ceiling effect consistent with the idea that the excitatory-inhibitory balance is maintained within bounds that TUS can shift but not exceed. This stands in contrast to deep brain stimulation, which can drive activity well outside the normal physiological range through sustained high-frequency electrical current. TUS, by virtue of its more subtle mechanosensitive mechanism of action, may be inherently constrained by the same homeostatic limits that govern normal neural dynamics.

Shifting the balance

Emerging evidence from studies directly measuring GABA+ following TUS demonstrate that TUS can indeed shift the excitatory/inhibitory balance by reducing inhibitory tone[102,105]. These studies report decreased GABA+ concentration following TUS, increasing network excitability, providing a plausible neurochemical basis for the facilitative behavioral effects observed in some studies. A direct next step would be to combine MRS measurements before and after TUS in the paradigm used in **Chapter 4**, testing whether the facilitative effects on Pavlovian bias and learning rate are indeed accompanied by a measurable shift in excitatory-inhibitory balance in the targeted regions.

My findings in **Chapter 2** suggest that TUS functions, at least in part, by normalizing circuit dynamics rather than simply enhancing or suppressing activity uniformly across individuals. From a clinical perspective, this has important implications. First, it provides insight into why different patients might respond differently to the same therapeutic protocol: pre-existing differences in neurochemical balance could determine the magnitude and even direction of treatment effects. Second, it opens the possibility of personalized neuromodulation, in which stimulation parameters are tailored based on individual neurophysiological measurements such as MRS-derived GABA+ concentrations. Third, and perhaps most intriguingly, it suggests that TUS acts as a restorative intervention, helping to restore normal circuit function in individuals whose baseline state deviates from the typical range[106,107] – representing a form of circuit *rebalancing* rather than indiscriminate excitation or inhibition.

State dependency

Another critical determinant of neuromodulation outcomes is what the participant is doing immediately before and during stimulation. Cognitive state influences TUS efficacy through at least two distinct but related mechanisms[108]. First, active engagement with a task drives moment-to-moment changes in neural excitability: sustained depolarization in task-relevant circuits places neurons closer to firing threshold, making them more susceptible to perturbation. This proximity to firing threshold can be modulated independently of slower neurochemical changes like GABA concentration and may explain why the same stimulation protocol produces different effects depending on whether a region is actively engaged or at rest at the time of stimulation. Second, homeostatic plasticity mechanisms mean that a circuit that has recently been strongly activated may enter a refractory state, attenuating subsequent stimulation effects[103,104]. A circuit at low baseline firing rates, by contrast, may show stronger responses to the same perturbation. Together, these mechanisms mean that cognitive state – what a person is doing, attending to, or anticipating – shapes both the immediate excitability of targeted populations and their capacity for plasticity induction, determining how effectively a given TUS protocol can shift neural processing and behavior.

Converging evidence from magnetic and electrical stimulation studies demonstrates this powerful state-dependency effects directly[108]. The same TMS protocol can produce facilitation or suppression depending on the prior activation state of the targeted cortex – a phenomenon termed state-dependent TMS[109,110]. Similarly, preceding inhibitory electrical stimulation reverses the typical facilitatory effects of high-frequency repetitive TMS, while preceding excitatory electrical stimulation produces the opposite pattern[104]. State dependency is not limited to differences in

prior stimulation history: the profound difference in global brain state between sleep and wakefulness shapes TMS-evoked responses, with sleep producing more focal and wakefulness producing broader cortical recruitment, reflecting how background neural activity constrains the propagation of a perturbation[111]. Even within the waking state, attention powerfully gates plasticity induction: lasting changes only occur when participants attend to the stimulated region or relevant stimulus, and directing attention elsewhere abolishes plasticity entirely despite identical physical stimulation – an effect that operates both across body parts and across stimulus categories within the same anatomical region[112,113].

For experimental rigor, standardizing cognitive state during stimulation is as important as standardizing the physical parameters. In my studies (**Chapters 2** and **4**), I took several steps to promote consistent brain states across participants and sessions: during hair preparation and transducer positioning, participants listened to the same calming playlist; experimenters provided consistent levels of interaction and reassurance. For the online protocol in **Chapter 2**, participants actively performed the saccade task during stimulation itself, ensuring FEF engagement. In **Chapter 4**'s offline protocol, participants viewed a fixation stimulus to maintain a similar attentional state during stimulation. These seemingly minor details may substantially reduce variability and enhance replicability across sessions.

A related and emerging question is whether actively engaging the targeted brain region during stimulation further enhances neuromodulation efficacy. Yaakub and colleagues (2023) found differential effects of TUS on GABA+ levels depending on the stimulated region: posterior cingulate cortex (PCC) stimulation decreased GABA+, while anterior cingulate cortex (ACC) stimulation did not[102]. Notably, stimulation was delivered during "rest", when PCC is typically active as part of the default mode network, but ACC is relatively quiescent. This pattern suggests that a region's activity level at the time of stimulation may be a key determinant of its susceptibility to neuromodulation. Evidence from TMS is convergent: the neural and behavioral impact of stimulation depends on the activation state of the targeted region and network at the moment of stimulation. Accordingly, several studies suggest that functionally engaging the target network during stimulation can yield more selective or more consistent effects than stimulation during rest[110,114,115,115–117].

My findings in **Chapter 4**, along with other studies showing cognitive modulation following offline TUS administered at rest, demonstrate that engagement is not strictly necessary: participants who receive TUS while at rest can still show behavioral effects when later performing a task. However, it remains possible that effect sizes would be larger and variability smaller if the targeted region were actively engaged

during stimulation. A direct test of this hypothesis would be to vary engagement within the same TUS study – delivering stimulation in one condition while participants perform a task that recruits the target region, and in other while they rest or perform an independent task. This would allow a within-subject comparison of whether active engagement amplifies neuromodulatory efficacy for TUS as well. If such an effect were confirmed, it would have immediate practical implications for protocol design. For future decision-making paradigms, having participants perform a brief version of the task immediately before or during TUS, would serve the dual purpose of standardizing brain state across participants and sessions, and of pre-activating the circuits that are the intended targets of neuromodulation.

TUS directionality: from channels to circuits to behavior

Lessons from TMS

In TMS, the field long operated with a seemingly tidy distinction between excitatory and inhibitory protocols – high-frequency repetitive TMS was labeled excitatory, low-frequency inhibitory. This framing shaped decades of experimental design and interpretation. As the field matured, it became clear that this dichotomy was an oversimplification: the direction of a TMS effect depends not only on the protocol but on the targeted region, the local excitatory-inhibitory balance at the time of stimulation, individual differences in cortical anatomy and neurochemistry, and the state of the network being engaged[118]. TUS is a younger field, and there is both an opportunity and a reason to avoid repeating the same mistake. Rather than assigning labels of "excitatory" or "inhibitory" to protocols, the goal should be to develop a mechanistic understanding of how ultrasound interacts with neural tissue at multiple levels of organization, and how that interaction produces behavioral change in a given circuit and context[119–121].

At the molecular level, TUS engages mechanosensitive ion channels whose expression varies across neuron types and brain regions, with different channel families mediating gain modulation versus plasticity induction depending on which neural populations are engaged[122–124] (see **General Introduction** for details). Critically, because cellular composition differs substantially across brain regions – from excitatory-dominated cortex to inhibitory-dominated striatum to the Piezo1-enriched amygdala – predictions about TUS directionality cannot simply be carried over from one target to another[125–127].

Lessons from the FEF: interpreting Chapter 2

The FEF provides a useful starting point for reasoning about TUS directionality, precisely

because its functional architecture and behavioral outputs are so well characterized. The FEF controls contralateral saccades and is shaped by mutual interhemispheric inhibition: each FEF suppresses its counterpart, with local inhibitory interneurons further modulating this competition[128,129]. This circuit architecture offered a direct behavioral assay for TUS directionality: increased contralateral saccades indicate facilitation of the circuit's canonical output; decreased contralateral saccades indicate suppression; random errors would suggest nonspecific perturbation. Previous work by Sumner and colleagues (2010) showed that individual GABA+ concentrations in the FEF correlate with distractor suppression – higher GABA+ was associated with better suppression of competing hemifield targets, consistent with local inhibitory control of interhemispheric competition[130]. A qualitatively similar relationship emerged in **Chapter 2**: participants with higher left FEF GABA+ showed a stronger intrinsic rightward bias during sham, suggesting that stronger local inhibition more effectively suppressed leftward distractors. Following TUS, this relationship disappeared.

A plausible mechanism is that TUS preferentially engaged FEF inhibitory interneurons, suppressing interhemispheric inhibition from the contralateral FEF and thereby releasing ipsilateral saccade generation – a circuit-level disinhibition manifesting as facilitation at the behavioral level. Alternatively, TUS might have directly modulated excitatory saccade neurons, boosting contralateral saccade generation through direct facilitation of pyramidal output rather than via disinhibition. These possibilities cannot currently be distinguished without direct neural recordings. What the behavioral data do support is a net facilitative effect on the FEF's canonical output, consistent with earlier online TUS work in macaques yielding net excitatory behavioral effects with comparable parameters[7].

Lessons from lesion studies: interpreting Chapter 4

The net behavioral direction of the **Chapter 4** TUS effects can be inferred from comparison with lesion studies. The observed pattern of enhanced learning rates and strengthened Pavlovian biases is the opposite of what lesion and inactivation studies produce: damage to these regions reduces precisely these behavioral signatures[27,50,63]. This supports a facilitative interpretation at the functional level, whereby TUS amplified the computations these regions normally perform.

What TUS does at the neural level, however, is harder to specify without direct measurements. One possibility is that offline TUS produced two temporally dissociable neurophysiological effects: an initial facilitation of circuit excitability that boosted Pavlovian bias expression, and a subsequent opening of a plasticity window that enhanced learning rates through increased synaptic modifiability. The cellular

architecture of both regions – the NAcc and amygdala – adds further interpretive complexity, as detailed in the **General Introduction**, with the NAcc's predominantly inhibitory medium spiny neurons and the amygdala's multiple functionally distinct subnuclei – although TUS was primarily targeted at the basolateral amygdala – making it difficult to specify which neural populations drove the observed effects[125,131,132].

Importantly, the 5 Hz pulsed protocol used in **Chapter 4** is not inherently facilitative: while facilitative effects have been reported in some subcortical targets[28], mixed and suppressive effects have been observed in cortical ones[133–135] – a reminder that protocol-level labels carry little predictive value without specifying the target region and study paradigm.

Taken together, these considerations highlight that the behavioral consequences of TUS cannot be reduced to a single directional effect of stimulation. Instead, stimulation outcomes likely depend on how a given set of acoustic parameters interacts with the physiological state and circuit architecture of the targeted region. While this makes the interpretation of TUS effects more complex, it also opens an important methodological opportunity. By systematically probing how different parameters influence different circuit nodes, TUS may ultimately allow researchers to map how different stimulation parameters interact with specific circuit elements to shape neural processing and behavior.

Nudging the brain: TUS as a modulatory rather than prescriptive intervention

Not all predicted effects in **Chapter 4** emerged with equal clarity. Although NAcc stimulation strengthened aversive Pavlovian biases and enhanced learning, we did not observe a comparably strong appetitive bias effect on Go-to-Win trials. One likely explanation is a ceiling effect: participants were already performing very well in this condition, even better than in the NoGo-to-Avoid condition, because a general Go bias makes the Go-to-Win condition particularly easy. When behavior is already near ceiling, the relevant neural populations have effectively committed to a course of action, leaving little residual uncertainty for TUS to exploit and to push performance further. In seesaw terms, the evidence had already tipped overwhelmingly to one side.

This ceiling effect points to a fundamental property of TUS as a brain stimulation technique: unlike TMS, which can evoke action potentials and initiate behavior independent of ongoing neural state – a TMS pulse over primary motor cortex will make your fingers twitch whether you want them to or not[136] – TUS at the intensities used in human research cannot generate activity de novo[137]. More precisely, TUS appears

to require that relevant neural populations are already sufficiently depolarized, close enough to threshold that a small additional perturbation can tip the balance. TUS is therefore better conceptualized as a neuromodulatory technique: it nudges, biases, or amplifies ongoing activity, but it cannot prescribe behavior from a quiescent circuit. TUS may therefore be most effective when neural populations are engaged but not yet committed – when competing representations are near-balanced and a small perturbation can determine the outcome[138].

In **Chapter 2**, I leveraged this principle through careful task design. By systematically varying stimulus onset asynchronies (SOAs) between target and distractor, I created trials that ranged from unambiguous (long SOA, where one stimulus clearly appeared first) to highly ambiguous (short SOA, where participants were uncertain which appeared first). This design allowed us to directly examine TUS's "sensitivity window": the range of decision uncertainty within which neuromodulation can influence behavior. FEF TUS influenced saccade choices only during trials with short SOAs, where the decision was ambiguous and participants were uncertain which target appeared first. When the delay between targets was large and the correct choice was obvious, TUS had no effect. To use an analogy: if a seesaw is already heavily weighted to one side, adding a small weight to the other side will not tip the balance. But when the seesaw is nearly balanced, that same small weight can determine the outcome. TUS provides the nudge, but only when there is uncertainty to exploit (**Figure 2**). This pattern demonstrates that TUS is fundamentally a modulatory technique: it shifts ongoing computations rather than creating them.

Returning to **Chapter 4**, this framing also explains why the MGNG task was generally well suited for TUS intervention. The probabilistic feedback, occasional reversals, and continuous introduction of novel cues kept participants in a state of ongoing uncertainty throughout the session, with no stable asymptote where performance plateaued. These are precisely the conditions under which a modulatory nudge can shift behavior – and the absence of an effect on Go-to-Win trials, where performance had effectively reached ceiling, is the exception that proves the rule.

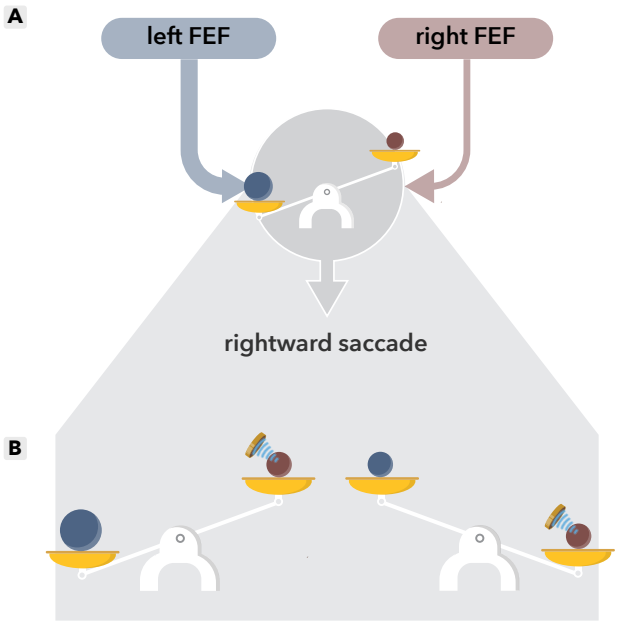


Figure 3 | Evidence accumulation in the FEF circuit and the modulatory nature of TUS. (A) FEF circuit as evidence accumulation: The FEF circuit for saccade selection can be understood through the same parallel evidence framework proposed for Pavlovian and instrumental control. The left and right FEF mutually inhibit each other, creating competitive dynamics in which each hemisphere accumulates evidence for contralateral saccades. The saccade direction is determined by which hemisphere accumulates sufficient evidence first (in this example the left FEF, leading to a rightward saccade). This is not purely a winner-take-all competition, but rather a continuous weighting of evidence from both sides, with local inhibitory circuits modulating the balance. **(B)** TUS as a nudge to tip the balance: TUS modulates ongoing evidence accumulation rather than evoking behavior de novo. When evidence is strongly weighted toward one side (left), adding a small TUS-induced bias to the other side has little effect – the decision is already determined. However, when evidence is approximately balanced (right), the same TUS input can tip the scale and determine the outcome. This seesaw analogy captures the nudging principle: TUS amplifies what is already there but cannot create activity from scratch. This explains why TUS effects are strongest during periods of decision uncertainty and why our probabilistic, volatile task environment, which maintains ongoing uncertainty, was well suited for detecting TUS modulation.

IV. FUTURE PERSPECTIVES

From abstract choices to embodied decisions: towards more ecologically valid study designs

The MGNG task employed in **Chapter 3** and **4** provides tight experimental control over stimulus timing, action contingencies, and outcome probabilities, and affords theoretical clarity by orthogonalizing cue valence and required action. However, it necessarily abstracts away from much of the richness of real-world decision-making. In everyday life, we do not choose between arbitrary cues on a computer screen; we navigate complex spaces, orient toward rewarding objects, and coordinate movements to obtain goals or avoid threats. An important next step, therefore, is to extend these findings from binary choice paradigms to more ecologically valid, richer continuous action spaces where motivation is expressed through the full complexity of movement. Understanding how the brain transforms motivational signals into graded, coordinated action may reveal computational principles that remain hidden in highly controlled laboratory tasks, such as how the NAcc integrates value signals to modulate action vigor continuously over time, or how motivational context shapes the spatiotemporal dynamics of motor preparation before a commitment is made.

One limitation of the MGNG task is that behavioral output is measured through simple button presses. While this provides a clear, quantifiable response, it forces what is likely a continuous, graded decision process into a binary outcome. This is particularly consequential for interpreting NoGo responses: a NoGo is recorded whenever no button press occurs, but the internal states underlying this outcome may differ substantially across conditions. On NoGo-to-Avoid trials, withholding a response may reflect a genuine Pavlovian freeze – a motivated suppression of action in response to threat – consistent with findings from Algermissen, who used pupillometry to show that aversive cues elicit autonomic arousal signatures associated with freezing even in the absence of overt movement[139]. On NoGo-to-Win trials, by contrast, participants must actively suppress what may already be an initiated Go response driven by Pavlovian approach motivation. In this case, the NoGo outcome conceals not an absence of motor activation, but its successful inhibition. Our binary measure collapses these mechanistically distinct states into a single category, making it impossible to determine, for instance, whether NAcc TUS modulated the vigor of an initiated but suppressed Go-response on incongruent trials. Extending the behavioral repertoire beyond binary press or no-press decisions is therefore essential for capturing the motivational dynamics that currently go undetected.

A first step toward richer behavioral measurement would be to adopt more sensitive

response measures within paradigms structurally like the current one. Pressure-sensitive response pads could reveal hesitant, partial responses – earlier, more tentative contact that develops into a full press only gradually or fails to reach threshold altogether. Eye-tracking offers another window into sub-threshold motor preparation: gaze dwell time on motivationally relevant cues predicts response bias even before overt action[139]. Body posture provides yet another channel: freeze responses to aversive stimuli are expressed in postural adjustments that precede and predict behavioral outcomes[140,141]. This is a measure that is entirely invisible to button-press paradigms.

Another question is whether NAcc motivational signals operate only at the level of discrete action selection, or whether they also modulate the spatial and temporal unfolding of movement itself. In natural behavior, motivation is rarely expressed as a binary; we lean toward rewarding stimuli, adjust our trajectory around threats, and modulate grip force according to the value of a held object[142,143]. Paradigms that embed choice within continuous navigation, such as virtual environments with dynamic force feedback varying with motivational context, would allow investigation of how NAcc output coordinates with motor cortex, cerebellum, and parietal cortex to produce fluid, goal-directed movement. Examining whether TUS of the NAcc shifts not only the probability but also the kinematics of action would directly test whether its motivational role extends into the continuous domain of motor execution.

These richer paradigms would not replace the rigorous, well-controlled tasks used in this thesis; rather, they would complement them. Simple tasks isolate specific computational processes and enable precise hypothesis testing, while more naturalistic tasks assess whether those processes generalize to the full dimensionality of real-world behavior. Together, they can build a more complete understanding of how motivation shapes action across the full spectrum of human decision-making.

From chronometric tool to clinical augmentation: the future of TUS

Over the course of my PhD, a clear shift has occurred in how TUS is used in cognitive neuroscience. Offline protocols – in which stimulation is delivered before a behavioral task rather than during it – have become the dominant approach in the field[28,102,144–147]. Their appeal is understandable: offline protocols have produced considerable effect sizes, are logistically and technically more straightforward to implement, and because their modulatory effects persist long enough to influence behavior across an entire task block, they more easily allow for complementary outcome measures such as fMRI and MRS. Moreover, offline protocols are less susceptible to auditory and somatosensory confounds, as stimulation occurs before the task rather than during

critical decision moments when such confounds could directly influence behavior. In **Chapter 4**, I leveraged these advantages by employing an offline protocol to investigate TUS effects on Pavlovian bias and instrumental learning. Yet in **Chapter 2**, I make a complementary case for the value of online TUS, in which stimulation is delivered concurrently with the cognitive process under investigation.

A chronometric tool

Online TUS offers something that offline protocols cannot: temporal resolution. A behavioral outcome like a button press or a saccade appears deceptively simple, but it is the endpoint of a cascade of interactive processes – cue perception, value retrieval, evidence accumulation, action preparation, and motor execution – that unfold in parallel and continuously influence each other. Offline TUS can tell us that a region is causally involved in a behavior; online TUS can tell us what computational aspect of this sensorimotor transformation it plays a role in. By delivering stimulation at precisely timed moments relative to task events, we can ask whether a region is critical during cue processing, during the deliberation phase, during outcome evaluation, or during motor output. This kind of chronometric interrogation is fundamental to understanding not just which regions are causally involved in a behavior, but also what specific computational role they play and at what point in the processing stream they do so.

In fact, the speed with which online TUS can influence neural processing is demonstrated directly in **Chapter 2**. TUS delivered to the FEF shifted saccadic behavior immediately during stimulation, with no carryover effects onto subsequent sham trials – confirming that the effect was tightly locked to the stimulation period itself. Crucially, this effect emerged within less than 265 ms: even on the fastest trials, where participants initiated a saccade before stimulation was complete, TUS already significantly biased saccade direction contralaterally. This rapidity suggests that TUS modulates ongoing spiking activity rather than inducing slower plasticity mechanisms and establishes it as a chronometric tool capable of probing neural computations with millisecond precision.

This chronometric perspective connects directly to the broader conceptual themes of this thesis. If Pavlovian and instrumental systems collaborate toward a shared goal by contributing different computations at different moments in time, then online TUS is uniquely positioned to dissect that temporal architecture. Stimulating the NAcc selectively during cue presentation (to probe its role in encoding actionable value) versus during outcome delivery (to probe its role in prediction-error-based updating) would test whether its contributions to Pavlovian bias and instrumental learning are

temporally separable computations or a single integrated process. For the amygdala, stimulating only at cue onset, but not during outcome presentation – i.e., the precise moment when the novelty BOLD response peaks in **Chapter 3** – would allow a direct causal test of whether the amygdala's tagging of novel cues is sufficient for the enhanced first-encounter learning we observed.

These are not idle empirical questions. Understanding whether the amygdala's contribution is restricted to a brief novelty window (a phasic modulation), or extends across the full trial (a more tonic modulation), would determine whether it is feasible to target that window therapeutically – for example, by stimulating the amygdala at cue onset during exposure therapy to boost the encoding of new safety associations at the moment when the novel contingency is first encountered. Similarly, knowing whether NAcc contributions to Pavlovian bias and instrumental updating are temporally dissociable would inform whether the same offline protocol can serve both functions, or whether separating them requires targeting different moments in the trial. Online TUS thus moves our questions from "which regions matter?" to "which computation does each region perform, and when is that computation modifiable?" – a fundamentally richer characterization of the amygdalo-striatal circuit with direct implications for intervention design.

Clinical augmentation

While online TUS holds great promise for characterizing neurocomputational processes in cognitive neuroscience, offline TUS may prove equally transformative in clinical translation. The temporal precision of online TUS allows us to dissect moment-by-moment brain function; the sustained plasticity effects of offline TUS open therapeutic possibilities. The core therapeutic potential of plasticity-focused offline TUS is that it may help the brain regain the ability to adapt – either by increasing regional excitability such that plasticity induction follows consequently, or by directly opening a window of heightened synaptic plasticity within which experience can drive lasting change.

The findings of this thesis suggest two concrete instantiations of this principle. The NAcc results from **Chapter 4** raise the possibility that TUS-induced plasticity in accumbal circuits could be leveraged in addiction treatment. A core challenge in addiction is that drug-associated cues have acquired powerful Pavlovian approach tendencies. Pavlovian control has come to dominate at the expense of goal-directed instrumental control – and these tendencies persist even after a person has explicitly decided to abstain[148]. If TUS can transiently open a plasticity window in NAcc circuits, this window could be used to rebalance Pavlovian and instrumental contributions:

training responses to trigger cues through cognitive reappraisal or structured avoidance practice, re-weighting the system toward goal-directed control at precisely the moment when the circuit is most amenable to change.

The amygdala findings from **Chapters 3** and **4** suggest a parallel possibility in anxiety and trauma-related disorders. Extinction learning – the process by which a previously threatening stimulus loses its aversive valence – is fundamentally a form of new learning and depends critically on the amygdala encoding a novel safety association at first encounter with the altered contingency[149]. If TUS applied to the amygdala can boost the encoding of new information at first encounter, as the **Chapter 4** data suggest, then stimulating the amygdala during exposure therapy could enhance the formation of new safety memories at the precise moment they are being created. More broadly, by temporarily enhancing plasticity in circuits relevant to emotional regulation, extinction learning, or cognitive reappraisal, TUS might make patients more responsive to therapeutic interventions that would otherwise require many more sessions or might not work at all. The therapist provides the content; TUS provides the neural conditions under which that content can be more effectively encoded and integrated.

The promise of TUS has generated considerable enthusiasm – and rightly so. Over the course of this work, I have come to a more measured, though no less optimistic, perspective. Many existing therapies work very well for a large number of people. Approximately two-thirds of patients with depression eventually achieve remission with sequential antidepressant trials[91,92], and cognitive behavioral therapy remains one of the most effective first-line treatments across a range of psychiatric conditions. Rather than replacing these interventions, TUS may function best as an augmentation method. This view is not defeatist – it is arguably more powerful. Helping the brain rather than controlling it is safer, because it works within the physiological range by design; more achievable, because we only need to provide opportunity rather than dictate outcome; and more powerful, because the brain's own plasticity mechanisms do much of the work. Indeed, everyday life itself may already have therapeutic effects, as engagement with meaningful activities, social connection, and healthy lifestyle choices all drive neural plasticity. If TUS lowers the threshold at which the brain can learn from these experiences, it may augment not only formal therapy but the patient's own daily life. TUS thus may achieve its greatest impact not by replacing existing interventions, but by enabling them, and nudging the brain toward states where learning and change become possible.

V. SUMMARY

The overarching aim of this thesis is to advance our understanding of how the brain transforms motivation into action – a question critical not only for fundamental neuroscience, but also for addressing psychiatric conditions in which this coupling is disrupted, such as addiction, depression, and anxiety[148,150–153]. To establish causal evidence, I first validated transcranial ultrasound stimulation (TUS) as a tool for modulating human behavior. Across three empirical chapters, I combined neuroimaging with targeted neuromodulation to dissect the contributions of the nucleus accumbens (NAcc) and amygdala to Pavlovian control and instrumental learning.

This work makes several key contributions. First, I demonstrated that TUS can reliably modulate the human brain with sufficient precision to influence cognition. **Chapter 2** showed that oculomotor decisions can be rapidly biased by TUS, whereby individual baseline GABA+ concentrations predict TUS efficacy, transforming interindividual variability from noise into systematic variation that can be leveraged for personalized protocols. Second, through fMRI (**Chapter 3**) and causal TUS (**Chapter 4**), I showed that the NAcc serves as a hub coupling motivation to action, contributing to both Pavlovian bias (invigorated action for reward cues, inhibition for threat cues) and instrumental learning (updating stimulus-action-outcome contingencies based on experienced feedback), supporting the gradual shift toward goal-directed control. Third, the amygdala signals and acts upon novelty, putatively to drive enhanced learning from initial cue encounters – a mechanism likely reflecting a broader role in prioritizing processing of and learning from salient and motivationally relevant stimuli. Embedding these findings in the broader literature, I propose a circuit model in which the NAcc integrates value and motivational signals from the amygdala to guide actionable decisions. The amygdala acts as an early salience detector that prioritizes initial learning, while the NAcc translates the resulting value representations into action – supporting both Pavlovian invigoration and suppression and the ongoing refinement of goal-directed instrumental control. Understanding these circuits is essential for addressing maladaptive decision-making in psychiatric disorders, where TUS may ultimately augment existing therapies by opening windows of plasticity during which associations can be retrained, extinction strengthened, or reward sensitivity restored.

REFERENCES

1. Kop, B. R. *et al.* Auditory confounds can drive online effects of transcranial ultrasonic stimulation in humans. *eLife* ***12***, (2024).
2. Kop, B. R., de Jong, L., Kim, B. P., den Ouden, H. E. M. & Verhagen, L. Parameter optimisation for mitigating somatosensory confounds during transcranial ultrasonic stimulation. *Brain Stimulat.* ***18***, 1224–1236 (2025).
3. Deffieux, T. *et al.* Low-Intensity Focused Ultrasound Modulates Monkey Visuomotor Behavior. *Curr. Biol.* ***23***, 2430–2433 (2013).
4. Folloni, D. *et al.* Manipulation of Subcortical and Deep Cortical Activity in the Primate Brain Using Transcranial Focused Ultrasound Stimulation. *Neuron* ***101***, 1109-1116.e5 (2019).
5. Khalighinejad, N. *et al.* A Basal Forebrain-Cingulate Circuit in Macaques Decides It Is Time to Act. *Neuron* ***105***, 370-384.e8 (2020).
6. Kim, H., Chiu, A., Lee, S. D., Fischer, K. & Yoo, S.-S. Focused Ultrasound-mediated Non-invasive Brain Stimulation: Examination of Sonication Parameters. *Brain Stimulat.* ***7***, 748–756 (2014).
7. Kubanek, J. *et al.* Remote, brain region-specific control of choice behavior with ultrasonic waves. *Sci. Adv.* ***6***, eaaz4193 (2020).
8. Tufail, Y. *et al.* Transcranial Pulsed Ultrasound Stimulates Intact Brain Circuits. *Neuron* ***66***, 681–694 (2010).
9. Wattiez, N. *et al.* Transcranial ultrasonic stimulation modulates single-neuron discharge in macaques performing an antisaccade task. *Brain Stimulat.* ***10***, 1024–1031 (2017).
10. Paus, T. Location and function of the human frontal eye-field: A selective review. *Neuropsychologia* ***34***, 475–483 (1996).
11. Schall, J. D. The neural selection and control of saccades by the frontal eye field. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* ***357***, 1073–1082 (2002).
12. Tehovnik, E. J., Sommer, M. A., Chou, I. H., Slocum, W. M. & Schiller, P. H. Eye fields in the frontal lobes of primates. *Brain Res. Brain Res. Rev.* ***32***, 413–448 (2000).
13. Vernet, M., Quentin, R., Chanes, L., Mitsumasu, A. & Valero-Cabré, A. Frontal eye field, where art thou? Anatomy, function, and non-invasive manipulation of frontal regions involved in eye movements and associated cognitive operations. *Front. Integr. Neurosci.* ***8***, 66 (2014).
14. Cardinal, R. N., Parkinson, J. A., Hall, J. & Everitt, B. J. Emotion and motivation: the role of the amygdala, ventral striatum, and prefrontal cortex. *Neurosci. Biobehav. Rev.* ***26***, 321–352 (2002).
15. LeDoux, J. E. Emotion circuits in the brain. *Annu. Rev. Neurosci.* ***23***, 155–184 (2000).
16. Phelps, E. A. & LeDoux, J. E. Contributions of the amygdala to emotion processing: from animal models to human behavior. *Neuron* ***48***, 175–187 (2005).
17. Stuber, G. D. *et al.* Reward-predictive cues enhance excitatory synaptic strength onto midbrain dopamine neurons. *Science* ***321***, 1690–1692 (2008).

18. Knutson, B., Adams, C. M., Fong, G. W. & Hommer, D. Anticipation of Increasing Monetary Reward Selectively Recruits Nucleus Accumbens. *J. Neurosci.* ***21***, RC159–RC159 (2001).
19. O'Doherty, J. *et al.* Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science* ***304***, 452–454 (2004).
20. Balleine, B. W. & Killcross, S. Parallel incentive processing: an integrated view of amygdala function. *Trends Neurosci.* ***29***, 272–279 (2006).
21. Geurts, D. E. M., Huys, Q. J. M., den Ouden, H. E. M. & Cools, R. Aversive Pavlovian Control of Instrumental Behavior in Humans. *J. Cogn. Neurosci.* ***25***, 1428–1441 (2013).
22. Paton, J. J., Belova, M. A., Morrison, S. E. & Salzman, C. D. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* ***439***, 865–870 (2006).
23. Blackford, J. U., Buckholtz, J. W., Avery, S. N. & Zald, D. H. A unique role for the human amygdala in novelty detection. *NeuroImage* ***50***, 1188–1193 (2010).
24. Pedersen, W. S. *et al.* The effects of stimulus novelty and negativity on BOLD activity in the amygdala, hippocampus, and bed nucleus of the stria terminalis. *Soc. Cogn. Affect. Neurosci.* ***12***, 748–757 (2017).
25. Taylor, J. R. & Robbins, T. W. Enhanced behavioural control by conditioned reinforcers following microinjections of d-amphetamine into the nucleus accumbens. *Psychopharmacology (Berl.)* ***84***, 405–412 (1984).
26. Wyvell, C. L. & Berridge, K. C. Intra-accumbens amphetamine increases the conditioned incentive salience of sucrose reward: enhancement of reward 'wanting' without enhanced 'liking' or response reinforcement. *J. Neurosci. Off. J. Soc. Neurosci.* ***20***, 8122–8130 (2000).
27. Costa, V. D., Dal Monte, O., Lucas, D. R., Murray, E. A. & Averbeck, B. B. Amygdala and Ventral Striatum Make Distinct Contributions to Reinforcement Learning. *Neuron* ***92***, 505–517 (2016).
28. Yaakub, S. N. *et al.* Non-invasive ultrasonic neuromodulation of the human nucleus accumbens impacts reward sensitivity. *Nat. Commun.* ***16***, 10192 (2025).
29. Scholz, V. *et al.* Cortical dopamine reduces the impact of motivational biases governing automated behaviour. *Neuropsychopharmacol. Off. Publ. Am. Coll. Neuropsychopharmacol.* ***47***, 1503–1512 (2022).
30. Swart, J. C. *et al.* Frontal network dynamics reflect neurocomputational mechanisms for reducing maladaptive biases in motivated action. *PLOS Biol.* ***16***, e2005979 (2018).
31. Brog, J. S., Salyapongse, A., Deutch, A. Y. & Zahm, D. S. The patterns of afferent innervation of the core and shell in the "Accumbens" part of the rat ventral striatum: Immunohistochemical detection of retrogradely transported fluoro-gold. *J. Comp. Neurol.* ***338***, 255–278 (1993).
32. Groenewegen, H. J., Wright, C. I., Beijer, A. V. & Voorn, P. Convergence and segregation of ventral striatal inputs and outputs. *Ann. N. Y. Acad. Sci.* ***877***, 49–63 (1999).

33. Kelley, A. E., Domesick, V. B. & Nauta, W. J. The amygdalostriatal projection in the rat--an anatomical study by anterograde and retrograde tracing methods. *Neuroscience* ***7***, 615–630 (1982).
34. Haber, S. N. & Behrens, T. E. J. The Neural Network Underlying Incentive-Based Learning: Implications for Interpreting Circuit Disruptions in Psychiatric Disorders. *Neuron* ***83***, 1019–1039 (2014).
35. Haber, S. N. & Knutson, B. The Reward Circuit: Linking Primate Anatomy and Human Imaging. *Neuropsychopharmacology* ***35***, 4–26 (2010).
36. Dayan, P., Niv, Y., Seymour, B. & Daw, N. D. The misbehavior of value and the discipline of the will. *Neural Netw. Off. J. Int. Neural Netw. Soc.* ***19***, 1153–1160 (2006).
37. Botvinick, M. M., Cohen, J. D. & Carter, C. S. Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn. Sci.* ***8***, 539–546 (2004).
38. Cavanagh, J. F., Zambrano-Vazquez, L. & Allen, J. J. B. Theta lingua franca: A common mid-frontal substrate for action monitoring processes. *Psychophysiology* ***49***, 220–238 (2012).
39. Cavanagh, J. F. & Frank, M. J. Frontal theta as a mechanism for cognitive control. *Trends Cogn. Sci.* ***18***, 414–421 (2014).
40. Cohen, M. X. A neural microcircuit for cognitive conflict detection and signaling. *Trends Neurosci.* ***37***, 480–490 (2014).
41. Shenhav, A., Botvinick, M. M. & Cohen, J. D. The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron* ***79***, 217–240 (2013).
42. Breland, K. & Breland, M. The misbehavior of organisms. *Am. Psychol.* ***16***, 681–684 (1961).
43. Dickinson, A. & Balleine, B. Motivational control of goal-directed action. *Anim. Learn. Behav.* ***22***, 1–18 (1994).
44. Dorfman, H. M. & Gershman, S. J. Controllability governs the balance between Pavlovian and instrumental action selection. *Nat. Commun.* ***10***, 5826 (2019).
45. Bach, D. R. & Dayan, P. Algorithms for survival: a comparative perspective on emotions. *Nat. Rev. Neurosci.* ***18***, 311–319 (2017).
46. Collins, A. G. E. & Frank, M. J. Opponent actor learning (OpAL): modeling interactive effects of striatal dopamine on reinforcement learning and choice incentive. *Psychol. Rev.* ***121***, 337–366 (2014).
47. Daw, N. D., Niv, Y. & Dayan, P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* ***8***, 1704–1711 (2005).
48. Rescorla, R. A. & Wagner, A. R. A Theory of Pavlovian Conditioning Variations in the Effectiveness of Reinforcement and Nonreinforcement. *Class. Cond. Curr. Res. Theory* ***21***, 64–99 (1972).
49. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* ***275***, 1593–1599 (1997).

50. Corbit, L. H. & Balleine, B. W. Double dissociation of basolateral and central amygdala lesions on the general and outcome-specific forms of pavlovian-instrumental transfer. *J. Neurosci. Off. J. Soc. Neurosci.* ***25***, 962–970 (2005).
51. Corbit, L. H. & Balleine, B. W. The general and outcome-specific forms of Pavlovian-instrumental transfer are differentially mediated by the nucleus accumbens core and shell. *J. Neurosci. Off. J. Soc. Neurosci.* ***31***, 11786–11794 (2011).
52. Algermissen, J., Swart, J. C., Scheeringa, R., Cools, R. & den Ouden, H. E. M. Prefrontal signals precede striatal signals for biased credit assignment in motivational learning biases. *Nat. Commun.* ***15***, 19 (2024).
53. Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. Model-based influences on humans' choices and striatal prediction errors. *Neuron* ***69***, 1204–1215 (2011).
54. Blundell, P., Hall, G. & Killcross, S. Lesions of the basolateral amygdala disrupt selective aspects of reinforcer representation in rats. *J. Neurosci. Off. J. Soc. Neurosci.* ***21***, 9018–9026 (2001).
55. Holland, P. C. & Gallagher, M. Double dissociation of the effects of lesions of basolateral and central amygdala on conditioned stimulus-potentiated feeding and Pavlovian-instrumental transfer. *Eur. J. Neurosci.* ***17***, 1680–1694 (2003).
56. Algermissen, J., Swart, J. C., Scheeringa, R., Cools, R. & den Ouden, H. E. M. Striatal BOLD and midfrontal theta power express motivation for action. *Cereb. Cortex* ***32***, 2924–2942 (2022).
57. Büchel, C. & Dolan, R. J. Classical fear conditioning in functional neuroimaging. *Curr. Opin. Neurobiol.* ***10***, 219–223 (2000).
58. Seymour, B. *et al.* Opponent appetitive-aversive neural processes underlie predictive learning of pain relief. *Nat. Neurosci.* ***8***, 1234–1240 (2005).
59. Talmi, D., Seymour, B., Dayan, P. & Dolan, R. J. Human pavlovian-instrumental transfer. *J. Neurosci. Off. J. Soc. Neurosci.* ***28***, 360–368 (2008).
60. Meijer, S. *et al.* The human amygdala in threat learning and extinction. *Sci. Adv.* https://doi.org/10.1126/sciadv.aea8233 (2026) doi:10.1126/sciadv.aea8233.
61. Atlas, L. Y., Doll, B. B., Li, J., Daw, N. D. & Phelps, E. A. Instructed knowledge shapes feedback-driven aversive learning in striatum and orbitofrontal cortex, but not the amygdala. *eLife* ***5***, e15192 (2016).
62. Corbit, L. H., Muir, J. L. & Balleine, B. W. The Role of the Nucleus Accumbens in Instrumental Conditioning: Evidence of a Functional Dissociation between Accumbens Core and Shell. *J. Neurosci.* ***21***, 3251–3260 (2001).
63. Hall, J., Parkinson, J. A., Connor, T. M., Dickinson, A. & Everitt, B. J. Involvement of the central nucleus of the amygdala and nucleus accumbens core in mediating Pavlovian influences on instrumental behaviour. *Eur. J. Neurosci.* ***13***, 1984–1992 (2001).
64. Hamid, A. A. *et al.* Mesolimbic dopamine signals the value of work. *Nat. Neurosci.* ***19***, 117–126 (2016).
65. Matsumoto, M. & Hikosaka, O. Two types of dopamine neuron distinctly convey positive and negative motivational signals. *Nature* ***459***, 837–841 (2009).

66. Saunders, B. T., Richard, J. M., Margolis, E. B. & Janak, P. H. Dopamine neurons create Pavlovian conditioned stimuli with circuit-defined motivational properties. *Nat. Neurosci.* ***21***, 1072–1083 (2018).
67. Steinberg, E. E. *et al.* A causal link between prediction errors, dopamine neurons and learning. *Nat. Neurosci.* ***16***, 966–973 (2013).
68. Guitart-Masip, M. *et al.* Action Dominates Valence in Anticipatory Representations in the Human Striatum and Dopaminergic Midbrain. *J. Neurosci.* ***31***, 7867–7875 (2011).
69. Guitart-Masip, M. *et al.* Go and no-go learning in reward and punishment: interactions between affect and effect. *NeuroImage* ***62***, 154–166 (2012).
70. Sescousse, G., Li, Y. & Dreher, J.-C. A common currency for the computation of motivational values in the human striatum. *Soc. Cogn. Affect. Neurosci.* ***10***, 467–473 (2015).
71. Salamone, J. D. & Correa, M. The mysterious motivational functions of mesolimbic dopamine. *Neuron* ***76***, 470–485 (2012).
72. Queirazza, F., Steele, J. D., Krishnadas, R., Cavanagh, J. & Philiastides, M. G. Functional Magnetic Resonance Imaging Signatures of Pavlovian and Instrumental Valuation Systems during a Modified Orthogonalized Go/No-go Task. *J. Cogn. Neurosci.* ***35***, 2089–2109 (2023).
73. Bartra, O., McGuire, J. T. & Kable, J. W. The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage* ***76***, 412–427 (2013).
74. Clithero, J. A. & Rangel, A. Informatic parcellation of the network involved in the computation of subjective value. *Soc. Cogn. Affect. Neurosci.* ***9***, 1289–1302 (2014).
75. Klein-Flügge, M. C., Bongioanni, A. & Rushworth, M. F. S. Medial and orbital frontal cortex in decision-making and flexible behavior. *Neuron* ***110***, 2743–2770 (2022).
76. Levy, D. J. & Glimcher, P. W. The root of all value: a neural common currency for choice. *Curr. Opin. Neurobiol.* ***22***, 1027–1038 (2012).
77. Rushworth, M. F., Kolling, N., Sallet, J. & Mars, R. B. Valuation and decision-making in frontal cortex: one or many serial or parallel systems? *Curr. Opin. Neurobiol.* ***22***, 946–955 (2012).
78. Ranganath, C. & Rainer, G. Neural mechanisms for detecting and remembering novel events. *Nat. Rev. Neurosci.* ***4***, 193–202 (2003).
79. Weierich, M. R., Wright, C. I., Negreira, A., Dickerson, B. C. & Barrett, L. F. Novelty as a dimension in the affective brain. *NeuroImage* ***49***, 2871–2878 (2010).
80. Wright, C. I. *et al.* Novelty responses and differential effects of order in the amygdala, substantia innominata, and inferior temporal cortex. *NeuroImage* ***18***, 660–669 (2003).
81. Averbeck, B. B. Amygdala and ventral striatum population codes implement multiple learning rates for reinforcement learning. in *2017 IEEE Symposium Series on Computational Intelligence (SSCI)* 1–5 (2017). doi:10.1109/SSCI.2017.8285354.
82. Averbeck, B. B. & Costa, V. D. Motivational neural circuits underlying reinforcement learning. *Nat. Neurosci.* ***20***, 505–512 (2017).

83. Neftci, E. O. & Averbeck, B. B. Reinforcement learning in artificial and biological systems. *Nat. Mach. Intell.* ***1***, 133–143 (2019).
84. Roesch, M. R., Calu, D. J., Esber, G. R. & Schoenbaum, G. Neural Correlates of Variations in Event Processing during Learning in Basolateral Amygdala. *J. Neurosci.* ***30***, 2464–2471 (2010).
85. Murdock Jr., B. B. The serial position effect of free recall. *J. Exp. Psychol.* ***64***, 482–488 (1962).
86. Hazy, T. E., Frank, M. J. & O'Reilly, R. C. Neural mechanisms of acquired phasic dopamine responses in learning. *Neurosci. Biobehav. Rev.* ***34***, 701–720 (2010).
87. Balleine, B. W. & O'Doherty, J. P. Human and Rodent Homologies in Action Control: Corticostriatal Determinants of Goal-Directed and Habitual Action. *Neuropsychopharmacology* ***35***, 48–69 (2010).
88. Yin, H. H. & Knowlton, B. J. The role of the basal ganglia in habit formation. *Nat. Rev. Neurosci.* ***7***, 464–476 (2006).
89. Berlim, M. T., van den Eynde, F., Tovar-Perdomo, S. & Daskalakis, Z. J. Response, remission and drop-out rates following high-frequency repetitive transcranial magnetic stimulation (rTMS) for treating major depression: a systematic review and meta-analysis of randomized, double-blind and sham-controlled trials. *Psychol. Med.* ***44***, 225–239 (2014).
90. Lefaucheur, J.-P. *et al.* Evidence-based guidelines on the therapeutic use of repetitive transcranial magnetic stimulation (rTMS). *Clin. Neurophysiol. Off. J. Int. Fed. Clin. Neurophysiol.* ***125***, 2150–2206 (2014).
91. Rush, A. J. *et al.* Acute and longer-term outcomes in depressed outpatients requiring one or several treatment steps: a STAR*D report. *Am. J. Psychiatry* ***163***, 1905–1917 (2006).
92. Trivedi, M. H. *et al.* Evaluation of outcomes with citalopram for depression using measurement-based care in STAR*D: implications for clinical practice. *Am. J. Psychiatry* ***163***, 28–40 (2006).
93. Gaynes, B. N. *et al.* What did STAR*D teach us? Results from a large-scale, practical, clinical trial for patients with depression. *Psychiatr. Serv.* ***60***, 1439–1445 (2009).
94. Hunter, A. M., Cook, I. A. & Leuchter, A. F. Does prior antidepressant treatment of major depression impact brain function during current treatment? *Eur. Neuropsychopharmacol. J. Eur. Coll. Neuropsychopharmacol.* ***22***, 711–720 (2012).
95. Hunter, A. M. *et al.* Antidepressant treatment history and drug-placebo separation in a placebo-controlled trial in major depressive disorder. *Psychopharmacology (Berl.)* ***232***, 3833–3840 (2015).
96. Huang, Y.-Z. *et al.* Plasticity induced by non-invasive transcranial brain stimulation: A position paper. *Clin. Neurophysiol.* ***128***, 2318–2329 (2017).
97. Farboud, S. *et al.* Rapid modulation of choice behavior by ultrasound on the human frontal eye fields. *Nat. Commun.* https://doi.org/10.1038/s41467-026-69854-7 (2026) doi:10.1038/s41467-026-69854-7.

98. Guerra, A., López-Alonso, V., Cheeran, B. & Suppa, A. Variability in non-invasive brain stimulation studies: Reasons and results. *Neurosci. Lett.* ***719***, 133330 (2020).
99. López-Alonso, V., Cheeran, B., Río-Rodríguez, D. & Fernández-Del-Olmo, M. Inter-individual variability in response to non-invasive brain stimulation paradigms. *Brain Stimulat.* ***7***, 372–380 (2014).
100. Pellegrini, M., Zoghi, M. & Jaberzadeh, S. Biological and anatomical factors influencing interindividual variability to noninvasive brain stimulation of the primary motor cortex: a systematic review and meta-analysis. *Rev. Neurosci.* ***29***, 199–222 (2018).
101. Stagg, C. J., Bachtiar, V. & Johansen-Berg, H. The Role of GABA in Human Motor Learning. *Curr. Biol.* ***21***, 480–484 (2011).
102. Yaakub, S. N. *et al.* Transcranial focused ultrasound-mediated neurochemical and functional connectivity changes in deep cortical regions in humans. *Nat. Commun.* ***14***, 5318 (2023).
103. Karabanov, A. *et al.* Consensus Paper: Probing Homeostatic Plasticity of Human Cortex With Non-invasive Transcranial Brain Stimulation. *Brain Stimulat.* ***8***, 442–454 (2015).
104. Siebner, H. R. *et al.* Preconditioning of Low-Frequency Repetitive Transcranial Magnetic Stimulation with Transcranial Direct Current Stimulation: Evidence for Homeostatic Plasticity in the Human Motor Cortex. *J. Neurosci.* ***24***, 3379–3385 (2004).
105. Jung, J., Atkinson-Clement, C., Kaiser, M. & Lambon Ralph, M. A. Transcranial focused ultrasound stimulation enhances semantic memory by modulating brain morphology, neurochemistry and neural dynamics. *Nat. Commun.* https://doi.org/10.1038/s41467-026-69579-7 (2026) doi:10.1038/s41467-026-69579-7.
106. Anticevic, A. *et al.* The role of default network deactivation in cognition and disease. *Trends Cogn. Sci.* ***16***, 584–592 (2012).
107. Marsman, A. *et al.* Glutamate in schizophrenia: a focused review and meta-analysis of ^{1}H-MRS studies. *Schizophr. Bull.* ***39***, 120–129 (2013).
108. Bradley, C., Nydam, A. S., Dux, P. E. & Mattingley, J. B. State-dependent effects of neural stimulation on brain function and cognition. *Nat. Rev. Neurosci.* ***23***, 459–475 (2022).
109. Silvanto, J., Muggleton, N. & Walsh, V. State-dependency in brain stimulation studies of perception and cognition. *Trends Cogn. Sci.* ***12***, 447–454 (2008).
110. Silvanto, J. & Pascual-Leone, A. State-dependency of transcranial magnetic stimulation. *Brain Topogr.* ***21***, 1–10 (2008).
111. Massimini, M. *et al.* Breakdown of Cortical Effective Connectivity During Sleep. *Science* ***309***, 2228–2232 (2005).
112. Kamke, M. R. *et al.* Visual Spatial Attention Has Opposite Effects on Bidirectional Plasticity in the Human Motor Cortex. *J. Neurosci.* ***34***, 1475–1480 (2014).
113. Heinen, K., Feredoes, E., Weiskopf, N., Ruff, C. C. & Driver, J. Direct Evidence for Attention-Dependent Influences of the Frontal Eye-Fields on Feature-Responsive Visual Cortex. *Cereb. Cortex N. Y. NY* ***24***, 2815–2821 (2014).

114. Grosshagauer, S. *et al.* Chronometric TMS-fMRI of personalized left dorsolateral prefrontal target reveals state-dependency of subgenual anterior cingulate cortex effects. *Mol. Psychiatry* ***29***, 2678–2688 (2024).
115. Jackson, J. B., Feredoes, E., Rich, A. N., Lindner, M. & Woolgar, A. Concurrent neuroimaging and neurostimulation reveals a causal role for dlPFC in coding of task-relevant information. *Commun. Biol.* ***4***, 588 (2021).
116. Sack, A. T. *et al.* Target Engagement and Brain State Dependence of Transcranial Magnetic Stimulation: Implications for Clinical Practice. *Biol. Psychiatry* ***95***, 536–544 (2024).
117. Taylor, S. F. *et al.* Lateral Prefrontal Stimulation of Active Cortex With Theta Burst Transcranial Magnetic Stimulation Affects Subsequent Engagement of the Frontoparietal Network. *Biol. Psychiatry Cogn. Neurosci. Neuroimaging* ***9***, 235–244 (2024).
118. Siebner, H. R. *et al.* Transcranial magnetic stimulation of the brain: What is stimulated? - A consensus and critical position paper. *Clin. Neurophysiol. Off. J. Int. Fed. Clin. Neurophysiol.* ***140***, 59–97 (2022).
119. Bestmann, S. & Walsh, V. Transcranial electrical stimulation. *Curr. Biol.* ***27***, R1258–R1262 (2017).
120. Dayan, E., Censor, N., Buch, E. R., Sandrini, M. & Cohen, L. G. Noninvasive brain stimulation: from physiology to network dynamics and back. *Nat. Neurosci.* ***16***, 838–844 (2013).
121. Polanía, R., Nitsche, M. A. & Ruff, C. C. Studying and modifying brain function with non-invasive brain stimulation. *Nat. Neurosci.* ***21***, 174–187 (2018).
122. Lemaire, T. *et al.* Microscopic deconstruction of cortical circuit stimulation by transcranial ultrasound. 2024.10.10.617091 Preprint at https://doi.org/10.1101/2024.10.10.617091 (2024).
123. Sorum, B., Docter, T., Panico, V., Rietmeijer, R. A. & Brohawn, S. G. Tension activation of mechanosensitive two-pore domain K+ channels TRAAK, TREK-1, and TREK-2. *Nat. Commun.* ***15***, 3142 (2024).
124. Yoo, S., Mittelstein, D. R., Hurt, R. C., Lacroix, J. & Shapiro, M. G. Focused ultrasound excites cortical neurons via mechanosensitive calcium accumulation and ion channel amplification. *Nat. Commun.* ***13***, 493 (2022).
125. Graveland, G. A. & Difiglia, M. The frequency and distribution of medium-sized neurons with indented nuclei in the primate and rodent neostriatum. *Brain Res.* ***327***, 307–311 (1985).
126. Markram, H. *et al.* Interneurons of the neocortical inhibitory system. *Nat. Rev. Neurosci.* ***5***, 793–807 (2004).
127. Zhu, J. *et al.* The mechanosensitive ion channel Piezo1 contributes to ultrasound neuromodulation. *Proc. Natl. Acad. Sci.* ***120***, e2300291120 (2023).
128. Buschman, T. J. & Miller, E. K. Top-Down Versus Bottom-Up Control of Attention in the Prefrontal and Posterior Parietal Cortices. *Science* ***315***, 1860–1862 (2007).
129. Sommer, M. A. & Wurtz, R. H. Brain circuits for the internal monitoring of movements. *Annu. Rev. Neurosci.* ***31***, 317–338 (2008).

130. Sumner, P., Edden, R. A. E., Bompas, A., Evans, C. J. & Singh, K. D. More GABA, less distraction: a neurochemical predictor of motor decision speed. *Nat. Neurosci.* ***13***, 825–827 (2010).
131. Gerfen, C. R. & Surmeier, D. J. Modulation of striatal projection systems by dopamine. *Annu. Rev. Neurosci.* ***34***, 441–466 (2011).
132. Janak, P. H. & Tye, K. M. From circuits to behaviour in the amygdala. *Nature* ***517***, 284–292 (2015).
133. Bao, S., Kim, H., Shettigar, N. B., Li, Y. & Lei, Y. Personalized depth-specific neuromodulation of the human primary motor cortex via ultrasound. *J. Physiol.* ***602***, 933–948 (2024).
134. Curtin, D., Walinjinkar, J., Fouragnan, E. F., Hendrikse, J. J. & Coxon, J. P. Transcranial ultrasound stimulation induces white matter plasticity in the human corticospinal tract. Preprint at https://doi.org/10.1101/2025.11.10.687582 (2025).
135. Fong, P.-Y. *et al.* A double-blind replication attempt of offline 5Hz-rTUS-induced corticospinal excitability. *Imaging Neurosci.* ***3***, IMAG.a.1046 (2025).
136. Hallett, M. Transcranial magnetic stimulation: a primer. *Neuron* ***55***, 187–199 (2007).
137. Darmani, G. *et al.* Non-invasive transcranial ultrasound stimulation for neuromodulation. *Clin. Neurophysiol.* ***135***, 51–73 (2022).
138. Butler, C. R. *et al.* Transcranial ultrasound stimulation to human middle temporal complex improves visual motion detection and modulates electrophysiological responses. *Brain Stimulat.* ***15***, 1236–1245 (2022).
139. Algermissen, J. & den Ouden, H. E. M. Pupil dilation reflects effortful action invigoration in overcoming aversive Pavlovian biases. *Cogn. Affect. Behav. Neurosci.* ***24***, 720–739 (2024).
140. Ly, V., Huys, Q. J. M., Stins, J. F., Roelofs, K. & Cools, R. Individual differences in bodily freezing predict emotional biases in decision making. *Front. Behav. Neurosci.* ***8***, (2014).
141. Hashemi, M. M. *et al.* Human defensive freezing: Associations with hair cortisol and trait anxiety. *Psychoneuroendocrinology* ***133***, 105417 (2021).
142. Eder, A. B., Krishna, A., Sebald, A. & Kunde, W. Embodiment of approach-avoidance behavior: Motivational priming of whole-body movements in a virtual world. *Motiv. Sci.* ***7***, 133–144 (2021).
143. Pessiglione, M. *et al.* Subliminal Instrumental Conditioning Demonstrated in the Human Brain. *Neuron* ***59***, 561–567 (2008).
144. Algermissen, J. *et al.* Low-intensity focused ultrasound to human amygdala reveals a causal role in ambiguous emotion processing and alters local and network-level activity. 2025.08.14.670358 Preprint at https://doi.org/10.1101/2025.08.14.670358 (2025).
145. Clarke, S. *et al.* Multi-focal ultrasound neuromodulation to the dorsal anterior cingulate cortex disrupts behavioural and neural pain processing. 2025.09.04.674195 Preprint at https://doi.org/10.1101/2025.09.04.674195 (2025).
146. Koutsoumpari, N. *et al.* Ultrasound neuromodulation reveals distinct roles of the dorsal anterior cingulate cortex and anterior insula in Pavlovian biases. 2025.06.12.659273 Preprint at https://doi.org/10.1101/2025.06.12.659273 (2025).

147. Nakajima, K. *et al.* A causal role of anterior prefrontal-putamen circuit for response inhibition revealed by transcranial ultrasound stimulation in humans. *Cell Rep.* ***40***, 111197 (2022).
148. Garbusow, M. *et al.* Pavlovian-to-instrumental transfer effects in the nucleus accumbens relate to relapse in alcohol dependence. *Addict. Biol.* ***21***, 719–731 (2016).
149. Quirk, G. J. & Mueller, D. Neural Mechanisms of Extinction Learning and Retrieval. *Neuropsychopharmacology* ***33***, 56–72 (2008).
150. Burguière, E., Monteiro, P., Mallet, L., Feng, G. & Graybiel, A. M. Striatal circuits, habits, and implications for obsessive-compulsive disorder. *Curr. Opin. Neurobiol.* ***30***, 59–65 (2015).
151. Heinz, A., Schlagenhauf, F., Beck, A. & Wackerhagen, C. Dimensional psychiatry: mental disorders as dysfunctions of basic learning mechanisms. *J. Neural Transm. Vienna Austria 1996* ***123***, (2016).
152. Huys, Q. J., Pizzagalli, D. A., Bogdan, R. & Dayan, P. Mapping anhedonia onto reinforcement learning: a behavioural meta-analysis. *Biol. Mood Anxiety Disord.* ***3***, 12 (2013).
153. Kalivas, P. W. & Volkow, N. D. The neural basis of addiction: a pathology of motivation and choice. *Am. J. Psychiatry* ***162***, 1403–1413 (2005).

EPILOGUE

I began this dissertation with stories of heroes and trials, with feathers and cups and arrows, because I suspected that the way we tell stories and the way we do science are not so different. Both are attempts to make sense of what we observe. Both reach for patterns, for causes, for meaning.

The questions I asked were for the most part handed to me by decades of prior work. The tools I used were largely built by others. The hands that guided me – supervisors, colleagues, friends, family – were always present. Still, there is something in these pages that did not exist before. A small contribution. A single arrow, loosed with whatever focus I could muster, landing somewhere in the distance. In the legend of Arash, the archer pours everything into a single shot and does not rise again. There were moments, I admit, while writing this dissertation, when I felt I might not either – but that is just my somewhat dramatic nature. Luckily, I get to continue, to ask new questions, to learn new things.

What I found, I have tried to share in these pages. Not as final answers – there are none – but as honest attempts.

And if there is one thing I have learned, it is this: the scales are always in motion. Pavlovian and instrumental, emotion and reason, impulse and intention – they do not settle into a fixed balance. They shift with experience, with context, with the slow accumulation of evidence. Our task, as scientists, is not to force the scales to one side, but to understand what tips them.

This dissertation has been my attempt at tipping the scales.

APPENDICES




Dutch summary
Research data management
About the author
PhD portfolio
List of publications
Acknowledgements
Donders graduate school

NEDERLANDSE SAMENVATTING

Waarom doen we wat we doen, zelfs als het eigenlijk niet slim is?

Een e-mail die bedoeld is voor je baas over dat uitgelopen project staat al een week in je conceptenmap – nog steeds niet verstuurd. Eigenlijk helemaal niet ingewikkeld, maar toch schuif je het elke dag weer voor je uit. Of je staat op een verjaardag en voordat je het doorhebt, pak je voor de derde keer een handje chips, terwijl je dat echt niet meer van plan was.

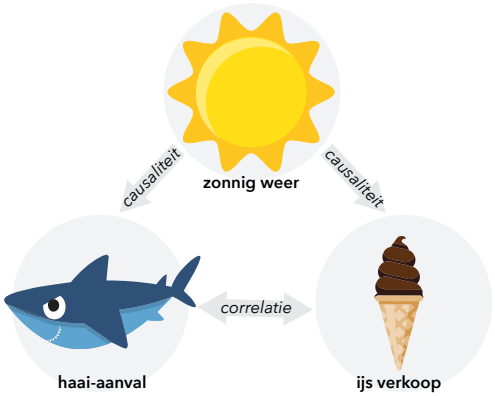
Dit soort momenten voelen bijna automatisch aan. Dat is ook precies wat ze zijn. Ons brein is geen perfecte rationele machine die elke beslissing zorgvuldig afweegt. Het heeft snelkoppelingen ingebouwd die ons helpen snel te handelen, zonder dat we er bewust over na hoeven te denken. Meestal zijn die snelkoppelingen handig, maar soms werken ze ons tegen.

Zodra die snelkoppelingen echt uit de hand lopen, herkennen we ze in aandoeningen zoals verslaving, dwangmatig gedrag of angststoornissen. Mensen met een verslaving weten heel goed dat het middel slecht voor hen is, maar de automatische drang om toch te gebruiken overheerst. Hetzelfde geldt voor mensen met een angststoornis: zij weten dat hun angst overdreven is, maar vermijden toch nieuwe situaties. In al deze gevallen doet het brein iets wat de persoon eigenlijk niet wil.

Om dat beter te begrijpen, is het belangrijk om te weten welke hersengebieden betrokken zijn bij dit soort keuzegedrag en wat die gebieden concreet doen. Dat is waar dit proefschrift over gaat.

Wat hersenscans ons wél en niet kunnen vertellen

Neurowetenschappers maken al tientallen jaren gebruik van hersenscans zoals *MRI*: een techniek waarmee je zowel gedetailleerde foto's als filmpjes van de hersenen kunt maken. Met die filmpjes – die in real time laten zien welke hersengebieden actief zijn – kunnen we zien welke gebieden oplichten op het moment dat iemand een beslissing neemt. Maar er zit een fundamenteel probleem aan die aanpak: het laat alleen een verband zien. Je kunt zien dat een bepaald gebied oplicht als iemand iets doet, maar dat betekent niet dat datzelfde gebied daadwerkelijk een rol speelt in dat gedrag. Dit is het concept van correlatie versus causaliteit (**Figuur 1**).



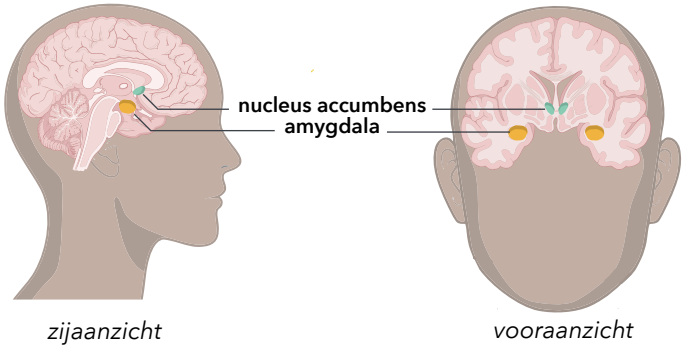
Figuur 1 | Correlatie hoeft geen causaliteit te betekenen. Op warme dagen, gaan meer mensen zwemmen én worden er meer ijsjes verkocht. Daardoor correleren ijsverkoop en haaiaanvallen, maar de een veroorzaakt de ander niet. Zonnig weer is de gemeenschappelijke oorzaak van beide. Om te achterhalen of ijsjes de boosdoener zijn, zou je ze bijvoorbeeld uit de winkels moeten halen en kijken of de haaiaanvallen daardoor ook afnemen. Dit illustreert waarom een verband tussen twee dingen nog niet betekent dat het een het ander veroorzaakt.

Om te bewijzen dat een hersengebied daadwerkelijk betrokken is bij een bepaald gedrag, moet je de activiteit in dat gebied tijdelijk verstoren om te zien wat er dan verandert. Dat wil je doen op een veilige en omkeerbare manier, zonder operatie. Maar dat bleek tot voor kort onmogelijk voor de hersengebieden die het meest interessant zijn voor keuzegedrag en motivatie zoals de *nucleus accumbens* en de *amygdala* (**Figuur 2**). Deze gebieden liggen diep in de hersenen en ver buiten het bereik van bestaande niet-invasieve technieken zoals magnetische of elektrische stimulatie. De niet-invasieve technieken kunnen namelijk alleen de oppervlakkige hersenschors bereiken. Medicijnen kunnen die diepe gebieden wel bereiken, maar verspreiden zich door het hele brein, waardoor je niet kunt bepalen welk specifiek gebied het effect veroorzaakt.

We hadden dus een methode nodig die diep, precies en veilig kon stimuleren. Die methode bestaat nu, namelijk: *transcraniële ultrageluidsstimulatie*, afgekort *TUS*.

Ultrageluid: van echo naar hersenstimulatie

Ultrageluid is gefocust geluid dat door weefsel heen gaat zonder het te beschadigen en is misschien wel het meest bekend van de echo waarmee je een baby in de buik kunt bekijken. Dezelfde technologie kan ook op de hersenen worden gericht. Door een klein apparaatje op het hoofd te plaatsen en gerichte ultrageluidstrillingen te sturen, kunnen we hersenactiviteit *veilig* van *buitenaf* beïnvloeden, tot *diep* in de hersenen,

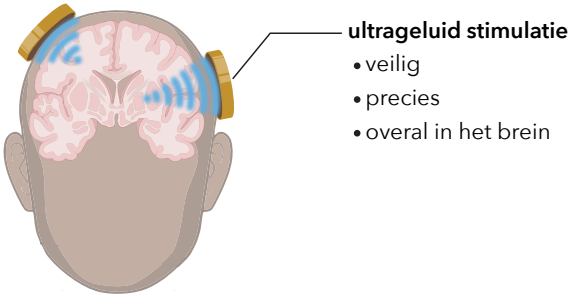


Figuur 2 | Hersengebieden die betrokken zijn bij keuzegedrag, zoals de nucleus accumbens (groen) en amygdala (oranje), zitten vaak diep in de hersenen. Zowel vanuit zij- als vooraanzicht is te zien hoe diep deze structuren verscholen zitten.

met een *precisie* van slechts een paar millimeter (**Figuur 3**).

Dieronderzoek liet al mooie resultaten zien, maar aan het begin van mijn promotieonderzoek was er nog maar weinig bewijs of het ook bij mensen werkte. Dat was een probleem, want zonder dat bewijs kon TUS niet worden ingezet om grote wetenschappelijke vragen over het menselijk brein te beantwoorden.

Daarom had dit proefschrift twee doelen. Eerst: laat zien dat TUS in mensen gedrag kan beïnvloeden. Dan pas: gebruik TUS om te onderzoeken welke hersengebieden betrokken zijn bij hoe wij keuzes maken.



Figuur 3 | Met TUS kunnen we zowel de oppervlakkige als diepe hersengebieden stimuleren, veilig van buitenaf en met hoge precisie. Door een klein apparaatje op het hoofd te plaatsen worden gerichte ultrageluidtrillingen door de schedel gestuurd.

Een duwtje in de juiste richting

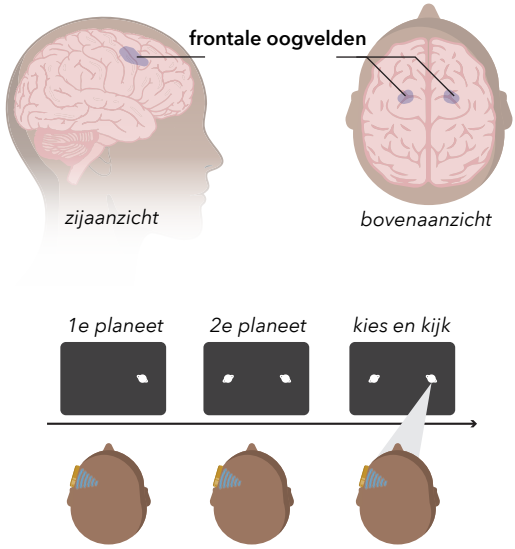
Om te testen of TUS werkelijk gedrag kan veranderen bij mensen, is het slim om een situatie te kiezen waarvan we al precies weten hoe het werkt in de hersenen. Want als je twee onbekenden tegelijk test, weet je nooit zeker wat je vindt. Stel dat je een nieuw medicijn test bij een ziekte die je nog niet begrijpt: als je dan iets ziet, weet je niet of dat door het medicijn komt of gewoon de gewone gang van zaken is bij die ziekte. Je hebt dus een goed bekend systeem nodig als uitgangspunt.

Voor mijn onderzoek in **hoofdstuk 2** was dat de aansturing van oogbewegingen. We weten namelijk heel goed welke hersengebieden betrokken zijn bij de keuze om je ogen naar links of rechts te bewegen: de *oogvelden* in de frontale hersenschors (**Figuur 4**, **boven**). Eerder dieronderzoek met mensapen had bovendien al aangetoond dat TUS van de oogvelden de keuze welke kant de aap opkijkt daadwerkelijk kon beïnvloeden. De vraag was nu: werkt dit ook bij mensen?

Vijfendertig deelnemers speelden een computerspel waarmee we de invloeden van TUS op hun keuzegedrag konden bestuderen (**Figuur 4**, **onder**). De resultaten lieten zien dat – ook in mensen – TUS instaat is om het keuzegedrag te beïnvloeden. Stimulatie van de linker oogvelden zorgde voor meer oogbewegingen naar rechts en stimulatie van de rechter oogvelden zorgde voor meer oogbewegingen naar links. Het effect was specifiek en niet gedreven door placebo: als we dezelfde techniek toepasten op de handmotorschors als controlemeting, veranderde er niets aan de oogbewegingen. Bovendien trad het effect al op binnen een fractie van een seconde na het begin van de stimulatie.

Elk mens is anders en dat bleek ook hier duidelijk: zowel de hersenchemie als het gedrag verschilden sterk tussen deelnemers. We maten bij elk van hen de concentratie van GABA, een remmende neurotransmitter, een stof waarmee hersencellen met elkaar communiceren, in de oogvelden. Mensen met hogere GABA-concentraties hadden van nature een sterkere voorkeur om naar rechts te kijken. Na TUS verdween dit verschil: iedereen bewoog naar hetzelfde niveau. TUS bracht de hersenen als het ware terug in balans, ongeacht hoe je hersenen er van nature voor stonden.

Deze bevindingen waren een van de eerste overtuigende aanwijzingen dat TUS direct en specifiek gedrag bij mensen kan sturen. Daarmee was de weg vrij voor de hoofdvraag van dit proefschrift: welke hersengebieden spelen een rol bij keuzegedrag en motivatie?



Figuur 4 | De frontale oogvelden zijn betrokken bij het steuren van oogbewegingen naar links en rechts, en vormden het doelwit van TUS in hoofdstuk 2. Boven: de ligging van de frontale oogvelden in de hersenschors, zichtbaar vanuit zij- en bovenaanzicht. Onder: De deelnemers speelden een computerspel waarbij links en rechts van hun scherm planeten verschenen. De deelnemer moest zo snel mogelijk naar de planeet kijken die als eerste in beeld verscheen. Een oogcamera registreerde hun oogbewegingen. Terwijl ze dit speelden, werd TUS toegepast op de frontale oogvelden in hun hersenen. De afbeelding laat het verloop van een beurt in het ruimtecomputerspel zien. Deelnemers doorliepen dit honderden keren, waardoor we nauwkeurig konden meten of TUS de keuze naar welke planeet ze keken beïnvloedde.

Twee systemen die keuzes sturen

Voordat ik de resultaten bespreek, is het goed om te begrijpen welke twee mechanismen centraal staan in dit proefschrift.

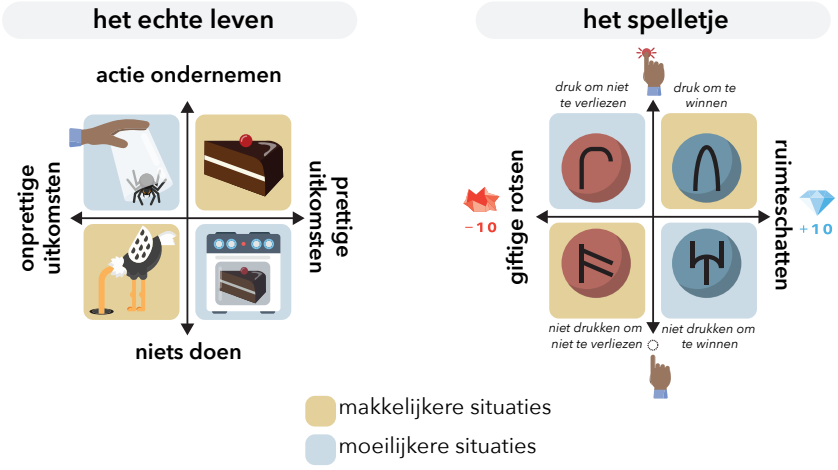
Het eerste mechanisme is *leren van ervaringen*: we leren door te doen en te ontdekken wat werkt. Als je iets probeert en het heeft een goede uitkomst, onthoud je dat en doe je het de volgende keer weer. Dit is het rationele, flexibele systeem dat zich aanpast op basis van ervaring.

Het tweede mechanisme is de *autopiloot* van ons brein: als iets goeds wordt verwacht, worden we automatisch geneigd om in actie te komen. Als we iets slechts verwachten, zijn we juist geneigd om stil te blijven. Dit is snel, automatisch en evolutionair oud. Je ziet het ook bij andere dieren: een prooidier dat een roofdier ruikt, bevriest spontaan.

Meestal werken deze twee systemen samen en versterken ze elkaar. Denk aan

iemand die voor het eerst een mango ziet op de markt: de frisse geur en het felle geel roepen meteen een positief gevoel op, en dat zet aan tot kopen. Thuis blijkt de mango heerlijk, en de volgende keer op de markt is de beweging naar de fruitkraam bijna vanzelfsprekend. De autopiloot zette de eerste stap, en het leren van ervaring ('dit was lekker, dat herhaal ik') verankerde het gedrag.

Maar soms werken ze elkaar tegen. Denk terug aan die chips op de verjaardag. Je hebt geleerd dat je er niet te veel van moet eten, maar de aanblik van de schaal roept automatisch een pakbeweging op. Of iemand in herstel van een verslaving die langs een kroeg loopt: rationeel weet hij dat hij niet naar binnen moet, maar de automatische drang is er toch. Juist in die momenten, wanneer de twee systemen niet hetzelfde willen, wordt duidelijk hoe sterk de autopiloot kan zijn (**Figuur 5**, **links**). Om dit te meten ontwikkelde ik een computer gebaseerde gedragstaak, waarmee ik zowel de invloed van de autopiloot als het leren van de gevolgen van keuzes kon meten (**Figuur 5**, **rechts**).



Figuur 5 | In het echte leven worden onze keuzes gestuurd door een samenspel van onze autopiloot en wat we leren van ervaringen. Links: onze autopiloot zet ons aan tot actie wanneer we iets prettigs verwachten (zoals een taart pakken) en niets te doen bij onprettige uitkomsten (zoals een struisvogel die zijn kop in het zand steekt). Maar soms vereist de situatie juist het tegenovergestelde: geduldig wachten bij een prettige uitkomst (de taart moet eerst nog afbakken in de oven) of actie ondernemen bij een onprettige uitkomst (de spin vangen voordat die later in je bed terechtkomt). In die gevallen stuurt onze autopiloot ons de verkeerde kant op en moeten we leren dat een andere reactie meer oplevert. Rechts: in het computerspel bootsten we deze vier situaties na. Deelnemers moesten per planeet leren of drukken of niets doen de beste actie was, om ruimteschatten te verdienen of giftige rotsen te vermijden. De gele vakken tonen de makkelijkere situaties, waarin de autopiloot aanzet tot de juiste actie. De blauwe vakken tonen de moeilijkere situaties, waarin de autopiloot juist tegenwerkt. Aan de hand van de keuzes en het leerproces konden we onderzoeken hoe sterk de autopiloot en het leren van ervaringen bijdroegen aan het gedrag van iedere deelnemer.

De hersenen in actie: wat gebeurt er tijdens keuzes?

In **hoofdstuk 3** richtte ik me eerst op de vraag welke hersengebieden oplichten tijdens dit soort keuzes. In de MRI scanner speelden 29 deelnemers het ruimteavontuur (**Figuur 5**, **rechts**), terwijl ik ondertussen hun hersenactiviteit mat. Welke gebieden lichten op als een deelnemer een planeet ziet? Wat gebeurt er in de hersenen op het moment van de keuze? Wat als ze positieve of negatieve uitkomsten te zien krijgen?

Wat we zagen dat de nucleus accumbens (**Figuur 2**), een kleine kern diep in de hersenen die al lange tijd wordt geassocieerd met beloning en motivatie, reageerde sterker op blauwe planeten dan op rode planeten. Dit past bij een rol in onze autopiloot: het gebied lijkt beloningssignalen op te pikken en ons alvast in de actiestand te zetten.

De amygdala, een ander diepgelegen hersengebied dat klassiek bekendstaat als het 'angstcentrum' maar in werkelijkheid veel breder betrokken is bij het verwerken van emotioneel relevante informatie (**Figuur 2**), liet een ander patroon zien. Bij het zien van planeten reageerde ze bijzonder sterk op planeten die deelnemers voor de allereerste keer zagen. Al bij de tweede aanbieding was die sterke reactie weg. Het lijkt erop dat de amygdala fungeert als een soort alarmsysteem voor het onbekende: 'nieuw symbool, nog geen idee wat je ermee moet doen, let op!'.

Beide gebieden reageerden bovendien sterk op het moment dat de uitkomst verscheen: het moment waarop deelnemers zagen of hun keuze goed was geweest. Ze zijn dus beide betrokken bij de verwerking van uitkomsten.

Dit waren echter correlationele bevindingen. De ijsjes en de haaiaanvallen. Tijd om de ijsjes uit de schappen te halen.

Oorzaak of gevolg? De nucleus accumbens en amygdala onder de loep

In het laatste empirische hoofdstuk, **hoofdstuk 4**, speelden dezelfde 29 deelnemers het ruimtespelletje opnieuw, maar nu drie keer, en elke keer kregen ze van tevoren TUS op een van drie condities: stimulatie van de nucleus accumbens, stimulatie van de amygdala, of schijnstimulatie waarbij alles hetzelfde was maar het ultrageluid niet daadwerkelijk in de hersenen terechtkwam. De volgorde was willekeurig zodat we eerlijk konden vergelijken. Door het gedrag na de drie condities te vergelijken, konden we achterhalen wat er veranderde in hun keuzes en leervermogen als we deze hersengebieden tijdelijk beïnvloedden.

Na stimulatie van de nucleus accumbens zagen we een tweeledig effect: zowel de invloed van de autopiloot als de snelheid van leren nam toe. Deelnemers begonnen met een sterkere neiging om te handelen bij beloningssignalen en om niets te doen bij strafsignalen, maar leerden vervolgens ook sneller. Het resultaat was dat de invloed van de autopiloot aanvankelijk groter was, maar ook sneller afnam naarmate deelnemers meer leerden.

Dat klinkt misschien tegenstrijdig, maar eigenlijk zijn deze twee processen minder elkaars tegenpolen dan vaak gedacht wordt. Ze hebben hetzelfde doel: optimaal gedrag. Wanneer je iets nog niet kent, dan biedt je autopiloot een handig vertrekpunt: doe iets als er een beloning te halen valt, doe niks als er straf dreigt. Snel, automatisch, en voor een beginnerssituatie vaak goed genoeg. Maar zodra je meer ervaring hebt, gaat wat je hebt geleerd zwaarder meewegen en kun je genuanceerder reageren. De nucleus accumbens lijkt een centrale rol te spelen in allebei die processen.

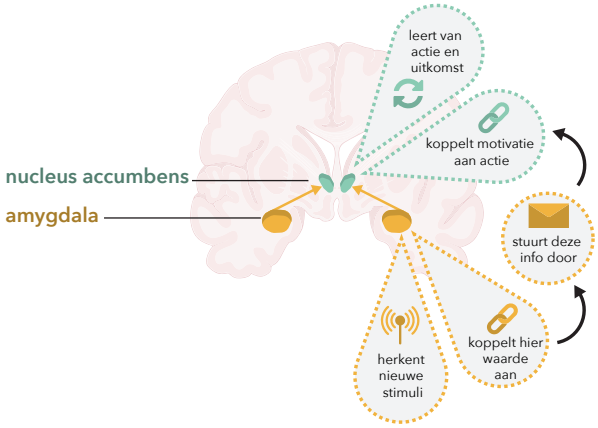
Voor de amygdala zagen we een andere specifieke bijdrage: na amygdalastimulatie leerden deelnemers sneller van planeten die ze voor het eerst zagen. De amygdala lijkt dus inderdaad betrokken bij de snelle eerste verwerking van nieuwe prikkels, precies wat we ook in de scanner hadden gezien.

Samen vormen deze twee gebieden een belangrijk circuit. De amygdala detecteert nieuwe prikkels en kent er snel een waarde aan toe, en stuurt die informatie vervolgens door naar de nucleus accumbens. Die fungeert als de eigenlijke hub: het is daar waar motivatie signalen – van zowel de amygdala als andere hersengebieden – worden samengebracht en omgezet in concrete actie. De amygdala geeft als het ware het startsein, de nucleus accumbens bepaalt wat je ermee doet. Deze bevindingen zijn uniek: voor het eerste hebben we causaal aangetoond dat de nucleus accumbens en amygdala bij mensen daadwerkelijk een bepalende rol spelen in hoe motivatie wordt omgezet in actie (**Figuur 6**).

Wat betekent dit voor de praktijk?

Deze bevindingen zijn niet alleen theoretisch interessant, maar ze helpen ons ook begrijpen waarom keuzegedrag bij bepaalde aandoeningen ontspoort, en misschien hoe we dat kunnen herstellen.

Denk bijvoorbeeld aan verslaving. Een van de kernsymptomen is een overweldigende, automatische drang om te gebruiken, zelfs als je weet dat het slecht voor je is. De autopiloot kan dan zo sterk worden dat wat je hebt geleerd onvoldoende tegenwicht biedt aan die drang. Beide processen, de te sterke invloed van de autopiloot én het



Figuur 6 | De nucleus accumbens en amygdala spelen elk een eigen rol in hoe motivatie wordt omgezet in actie. Uit de hersenscan- en stimulatie-experimenten bleek dat de nucleus accumbens betrokken is bij zowel gedrag op de autopiloot als het leren van ervaringen, en dat de amygdala nieuwe situaties herkent en betrokken is bij het verwerken van uitkomsten. Samen vormen ze een circuit: de amygdala koppelt waarde aan nieuwe prikkels en stuurt die informatie door naar de nucleus accumbens, die vervolgens fungeert als de centrale hub waar motivatie wordt omgezet in concrete actie.

verminderde vermogen om te leren van negatieve uitkomsten, zijn terug te voeren op de nucleus accumbens. Mijn bevindingen ondersteunen dat beeld en laten zien dat TUS van de nucleus accumbens dat patroon kan beïnvloeden.

Angststoornissen laten een ander patroon zien. Mensen met angststoornissen vermijden nieuwe situaties en beschouwen het onbekende als bedreigend. Dat wordt in verband gebracht met een overactieve amygdala. Interessant genoeg zagen wij dat amygdalastimulatie het leren van nieuwe prikkels juist versnelde, wat suggereert dat gerichte beïnvloeding van de amygdala mensen kan helpen om het onbekende sneller een plek te geven.

Net als in hoofdstuk 2, waar TUS de hersenen als het ware terugbracht naar een gemeenschappelijk niveau ongeacht hoe iedereen van nature in elkaar zat, biedt TUS ook in de context van keuzegedrag de mogelijkheid om op maat te werken. Niet iedereen heeft dezelfde hersenen, en niet iedereen zal op dezelfde manier reageren op een behandeling. TUS biedt de precisie om juist dat individuele verschil te kunnen adresseren.

Precies daar ligt de klinische belofte van TUS: niet als hersenbeheerser, maar als hulpmiddel. De hersenen moeten nog steeds zelf het werk doen. Maar een gerichte ultrageluidsbehandeling zou de hersenen een helpende hand kunnen bieden bij het

opnieuw leren, bij het doorbreken van oude automatische patronen. Een duwtje in de juiste richting.

Tot slot

Dit proefschrift laat zien dat ultrageluid een effectieve technologie is om menselijk gedrag te beïnvloeden, dat de nucleus accumbens en amygdala ieder op hun eigen manier bijdragen aan hoe wij keuzes maken en leren, en dat die bijdragen ook werkelijk een oorzakelijk en niet slechts een correlationeel verband hebben. Samen brengen deze inzichten ons een stap dichter bij een beter begrip van de menselijke hersenen en bij nieuwe mogelijkheden voor de behandeling van aandoeningen waarbij het keuzegedrag is ontspoord.

RESEARCH DATA MANAGEMENT

This research followed the applicable laws and ethical guidelines. Research Data Management was conducted according to the FAIR principles. The paragraphs below specify in detail how this was achieved.

Ethics

This thesis is based on the results of human studies, which were conducted in accordance with the principles of the Declaration of Helsinki. The medical research ethics committee (METC) region Arnhem/Nijmegen provided ethical approval for the study samples acquired in Chapter 2 (CMO2022-15953/NL82309.091.22). For Chapter 3 and 4, The Ethical Committee of the faculty of Social Sciences (ECSS) has given a positive advice to conduct these studies to the Dean of the Faculty, who formally approved the conduct of these studies (ECSW-2025-021). This research is funded by the DCC PhD Fellowship.

Findable Accessible

The table below details where the data and research documentation for each chapter can be found on the Donders Repository (DR). All data archived as a Data Sharing Collection remain available for at least10 years after termination of the studies.

Chapter	DAC	RDC	DSC	DSC License
2	di.dcc. DAC_2023.00021_006	di.dcc. RDC_2023.00021_883	di.dcc. DSC_2025.00018_488	CC-BY-4.0
3	di.dcc. DAC_2024.00062_488	di.dcc. RDC_2025.00035_313	In progress	CC-BY-4.0
4	di.dcc. DAC_2025.00035_671	di.dcc. RDC_2025.00035_313	In progress	CC-BY-4.0

DAC = Data Acquisition Collection, RDC = Research Documentation Collection, DSC = Data Sharing Collection

For chapters 2, 3, 4 and 5 research data have also been stored on the Donders project drive. These data were accessible to all members involved in the project. After finalization of chapter 2, data were removed from the Donders project drive. The publication of chapters 3 and 4 is still under preparation. The data will be shared with the reviewers in a DSC and will be made publicly available, once the article has been published. Prior to public release, all data will be anonymized/de-identified to remove any personally identifiable information. It will then be shared

under the CC-BY-4.0 license.
Informed consent was obtained on paper following the Centre procedure. The forms are archived in the central archive of the Centre for 10 years after termination of the studies.

Interoperable, Reusable

The raw data are stored in the DAC in their original form. For RDC and DSC long-lived file formats (e.g. .sav, .csv, .tif,) have been used ensuring that data remains usable in the future. The data of the RDC and DSC are organized according to the BIDS standards, with concomitant readme files. Results are reproducible by providing a description of the experimental setup, raw data (DAC and DSC), analysis scripts or pipelines (RDC and DSC). Also, the used software including version numbers is specified.

Privacy

The privacy of the participants in this thesis has been warranted using random individual subject codes. A pseudonymization key linked this random code with the personal data. This pseudonymization key was stored on a network drive that was only accessible to members of the project who needed access to it because of their role within the project. The pseudonymization key was stored separately from the research data. The pseudonymization keys of chapters 2, 3 and 4 were destroyed within one month after finalization of these projects. Data in chapter 2 are not identifiable and shared without restrictions. MRI-data of chapter 3 will be defaced and anonymized and shared under CC-BY-4.0.

ABOUT THE AUTHOR

Soha Farboud was born in 1995 in Isfahan, Iran. After moving to the Netherlands with her parents at the age of four, she grew up across various cities before settling in Nijmegen, where she completed her secondary education at Het Stedelijk Gymnasium. She went on to pursue a Bachelor of Science in Biology at Radboud University Nijmegen, during which she participated in an interdisciplinary Honours Academy program and completed a six-month research internship at the Hospital of San Raffaele in Milan, Italy. There, she conducted diffusion tensor imaging (DTI) analyses to examine the impact of birth season on bipolar disorder severity — her first hands-on experience with neuroimaging and clinical research. She also served as a student representative on the Faculty Education Committee and taught practicals and workgroups to first-year students before graduating in 2017.

Soha continued her training with a Master of Science in Medical Biology, specializing in Neurobiology, at Radboud University Nijmegen. During her first internship at the Max Planck Institute for Psycholinguistics, supervised by Prof. Dr. Simon Fisher, Dr. Cleo Smeets, and Dr. Bernhard Englitz, she investigated the spatial dynamics of courtship vocalization behavior in Foxp2 knockdown mice using chronic neural recordings, immunohistochemistry, and genetic analyses — work she continued as a Research Assistant. Her second internship took her to the Sleep & Memory Lab at the Donders Institute, supervised by Prof. Dr. Martin Dresler and Dr. Hong-Viet Ngo, where she investigated the influence of sleep rhythms on memory consolidation using EEG and contributed to an fMRI project on emotional memory. Alongside this work, she co-authored a review paper on closed-loop auditory stimulation of sleep slow oscillations. She graduated cum laude in 2021.

That same year, Soha commenced her PhD in cognitive neuromodulation at the Donders Institute for Brain, Cognition and Behaviour, supervised by Prof. Dr. Hanneke den Ouden, Dr. Lennart Verhagen, and Prof. Dr. Pieter Medendorp. Her doctoral research combined two lines of inquiry: the development and validation of transcranial ultrasonic stimulation (TUS) as a non-invasive brain stimulation tool, and the investigation of causal neural mechanisms underlying motivational decision-making. She established foundational knowledge on how TUS applied to the frontal eye fields influences oculomotor behavior, and used TUS to probe the contributions of the nucleus accumbens and amygdala to value-based learning and decision-making in humans. Beyond her research, Soha played a central role in building a brain stimulation laboratory at the Donders Institute and helped place focused ultrasound on the map at Radboud University, efforts further strengthened through her role as a Business Developer at the FUS Initiative. She shared her expertise through workshops

and lectures for the Brainbox Initiative, NeuroTechEU, and the Donders NIBS Toolkit, consulted as a clinical trial consultant, and saw her research covered by national and international media.

PHD PORTFOLIO

Name PhD candidate	Soha Farboud
Department	Donders Centre for Cognition
PhD supervisor:	Prof. Dr. Hanneke E.M. den Ouden
PhD supervisor:	Prof. Dr. W. Pieter Medendorp
PhD co-supervisor:	Dr. Lennart Verhagen
Research period	September 2021 – March 2026
PhD job	Teaching (Psychopharmacology course)
	Neuromodulation Research Support

Conferences & Workshops

Name	Year	Contribution	Location
FUN Conference (iTRUSST)	2021	Educational talk	Online
Brainbox Initiative Conference	2022	Workshop (lectures and practicals)	Online
NVP Conference	2022	Poster*	Egmond aan Zee, NL
Donders NIBS Toolkit	2022	Workshop (lectures and practicals)	Nijmegen, NL
FUN Conference	2022	Poster	Mainz, GER
Donders NIBS Toolkit	2023	Workshop (lectures and practicals)	Nijmegen, NL
Brainbox Initiative	2023	Workshop (lectures and practicals)	London, UK
BeNe Brain Stimulation Conference	2023	Talk	Nijmegen, NL
NVP Conference	2023	Poster*	Egmond aan Zee, NL
BRIC TUS Workshop	2024	Workshop (lectures and practicals)	Plymouth, UK
Donders NIBS Toolkit	2024	Workshop (lectures and practicals)	Nijmegen, NL
Donders Day	2024	Poster	Nijmegen, NL
FUN Conference	2024	Talk	Toronto, CA
Brainbox Initiative Conference	2024	Talk	London, UK
Brainbox Initiative Workshop	2024	Workshop (lectures and practicals)	London, UK
Brainbox Initiative Workshop	2025	Workshop (lectures and practicals)	London, UK
NVP Conference	2025	Talk	Egmond aan Zee, NL

* Received Award

Teaching

Course	Year	Activity
Psychopharmacology	2022	Exam questions & slides
Psychopharmacology	2023	Video assignment
NeurotechEU	2023	Online lectures and practicals
Psychopharmacology	2024	Video assignment

Invited Talks & Outreach

Name	Year	Contribution	Location
Donders Sessions	2023	Talk	Nijmegen, NL
Radboud University Executive Board (CvB)	2024	Talk and demonstration	Nijmegen, NL
Donders Science Advisory Board	2024	Talk	Nijmegen, NL
Neurology Department, Radboudumc	2024	Talk	Nijmegen, NL
Oxford University (team meeting)	2024	Talk	Online
Rathenau NeurotechEU	2025	Talk	Nijmegen, NL
Sensorimotor Department Day	2025	Talk	Nijmegen, NL

Media

Outlet	Year	Activity	Location
NPO Radio1: met het oog op morgen	2024	Radio interview	Hilversum, NL
QUEST magazine	2025	Interview	Nijmegen, NL
KIJK magazine	2025	Interview	Nijmegen, NL
NEMO Science Museum	2026	Interview	Nijmegen, NL
GGZ Nieuws, Earth.com, Time.News myScience.org MedicalXpress.com	2026	Research Coverage	Online

Supervision

Name	Program	Duration
Solenn L.Y. Walstra	Cognitive neuroscience	Academic year
Christian Bauer	Cognitive neuroscience	Academic year
Floris van der Werff	Cognitive neuroscience	Academic year
Mihaela Belacurencu	Cognitive neuroscience	Academic year

Other activities		
Activity	**Year**	**Organization**
Business Developer	2024	FUS Initiative
Clinical trial consultant	2025	The Science Behind, UK
Scientific Jury	2025	Brainbox Initiative Conference, UK
Co-author, Handbook of Transcranial Ultrasound Stimulation (Ch. 5.3, *Confounds*)	2025/2026	Oxford University Press

Certificates
Article 9
Good Clinical Practice (GCP)

LIST OF PUBLICATIONS

Soha Farboud, Benjamin R Kop, Renée S Koolschijn, Solenn LY Walstra, José P Marques, Andrey Chetverikov, W Pieter Medendorp, Lennart Verhagen, Hanneke EM den Ouden (2026). Rapid modulation of choice behavior by ultrasound on the human frontal eye fields.

Nature Communications,
https://doi.org/10.1038/s41467-026-69854-7

Kop, B. R., Shamli Oghli, Y., Grippe, T. C., Nandi, T., Lefkes, J., Meijer, S. W., **Farboud, S.**, Engels, M., Hamani, M., Null, M., Radetz, A., Hassan, U., Darmani, G., Chetverikov, A., den Ouden, H. E., Bergmann, T. O., Chen, R., & Verhagen, L. (2024). Auditory confounds can drive online effects of transcranial ultrasonic stimulation in humans.

eLife, Volume 12,
https://doi.org/10.7554/eLife.88762

Soha Farboud, Mahdad Jafarzadeh Esfahani, Hong-Viet V. Ngo, Jules Schneider, Frederik D. Weber, Lucia M. Talamini, Martin Dresler (2023). Closed-loop auditory stimulation of sleep slow oscillations: Basic principles and best practices.

Neuroscience & Biobehavioral Reviews, Volume 153,
https://doi.org/10.1016/j.neubiorev.2023.105379

Soha Farboud, Gabriel Oliveira-Stahl, Max L. Sterling, Jesse J. Heckman, Bram van Raalte, Dionne Lenferink, Amber van der Stam, Cleo J.L.M. Smeets, Simon E. Fisher, Bernhard Englitz (2023). High-precision spatial analysis of mouse courtship vocalization behavior reveals sex and strain differences.

Scientific Reports, Volume 13,
https://doi.org/10.1038/s41598-023-31554-3

ACKNOWLEDGEMENTS

It takes a team. It really does. This thesis belongs not only to the people who helped me think, experiment, analyze and write, but also to those who made it all possible through their love and support. There are more people to thank than I can fit here, and more to say than these pages allow. But this is a start, and I mean every word.

Hanneke, do you remember asking for an extra call during my interview to make sure I understood the risks of this PhD? Well, I can't say I wasn't warned. But I am so happy I took that risk, and even happier that I took it with you. Beyond everything you have taught me, your warmth and care have meant so much to me throughout these years. Your knowledge still amazes me. Whenever I thought I had finally reached your level of understanding, you would ask a question that made me realize how much more there was to learn. And you always took the time to help me get there, with such patience and kindness. I have learned so much from you, and I will always be grateful. In some ways, we are very alike, especially in our ability to spend far too long perfecting the smallest detail of a figure. We tried to protect each other from this, with limited success (oops). In other ways, we balance each other out: when I rushed ahead, you knew when to put on the brakes. Somewhere in the middle, we found a way of working together that I look back on with so much pride and affection. You have been my supervisor, teacher and mentor, and I am so happy I now also get to call you my friend. I'm thankful that through you, I also got to know **Peter**, **Alec** and, of course, my tiny best friend **Mirre**.

Lennart, you do realize you're to blame for all of this, right? On a bleak day during COVID, I came across your research shortly after you returned to the Netherlands. I think I emailed you that same day about a master's project. You introduced me to ultrasound and the wonderful people working on it, and somehow, as a rookie PhD student, I found myself travelling abroad to give lectures and workshops. Your energy is contagious. You think big, make people believe they can do more than they thought, and give them opportunities to try. You are a teacher, mentor and a true visionary, and your belief in me has meant so much. You encouraged me to take on things I would not have imagined doing, while always making me feel supported. I have grown so much through working with you. Working together felt natural during my master's, and these past years have only confirmed that. I'm so happy our work continues. Our visions are aligned, and we need few words to understand each other, although our hours-long meetings might suggest otherwise. I look forward to everything we have planned and, as you would say, conquering the world. I'm also happy that through you I got to know **Luke**, **Quinten** and **Inge**, who has become a mentor and someone I look up to herself.

Pieter, you were a steady presence alongside my two very, and sometimes a little too, enthusiastic supervisors (Hanneke and Lennart, please look away). I admired how you combined running a lab, teaching and developing a new bachelor's program, yet always found time to respond or have a quick chat. Your pragmatic approach helped me keep things in perspective. I'm glad we found a shared interest in the frontal eye fields and would happily have done more TUS-FEF projects together. Thank you for being my co-promotor.

Benjamin, apart from this thesis, which I assume I'm supposed to mention, my friendship with you is the best thing to come out of my PhD. On my second day, you stepped out of your uncle's law-firm car in sweatpants and a shirt with holes in it, carrying three Starbucks coffees: one decaf, one vegan latte and one regular latte (in case I was normal). An excellent introduction. Since then, you have become part of not just my life, but Sebas's and my friends' and family's lives too. I honestly don't think I could have done this without you. You paved the way, taught me so much and were there whenever I needed you. Mostly, though, we have so much fun. Even now that we live worlds apart, we still FaceTime almost every day. You are ridiculously smart, funny, kind and easy to love. You excel at everything you do, and I have looked up to you since the day we met. I will continue to do so until we open our café in the Ooijpolder and grow old together. Thank you for being my partner in crime, one of my best friends and my paranymph. I love you so much.

Johannes, my unofficial supervisor. For a while, during those late COVID days, it was mostly just the two of us in the office at the MM. I'm so grateful I had you there. You are one of the smartest and most patient people I know, and you always made time to help me understand something, however long it took. In the time it took me to finish this PhD, you finished yours, did a postdoc in Oxford, secured grants and moved to Zurich. I can't wait to see you set up your own lab. To me, you're already a PI. I'm in awe of everything you do, but what means the most to me is how kind you have always been to me. I feel very lucky that sharing a lab and an office also gave me such a dear friend.

Elena, you beautiful Italian social butterfly. You welcomed me from day one and became my friend and, most importantly, my ultimate spinning partner. I love our conversations, whether we're complaining or laughing. See you Monday and Tuesday at six, or five with **Egbert** completing our spinning trio. Preferably not Wednesdays.

Floortje, data collection kept us apart for a while, but I'm so happy we grew closer towards the end. I have always been so impressed by how smart and thoughtful you are as a scientist. Getting to know you as a wonderful mum to the sweetest little boy has given me even more to admire. Please give **Orion** all the blueberries he wants on my behalf.

Renée, my MRS queen. One of my favorite memories is going for drinks with you and Ben in your first week, then ending up at mine making dinner. From patiently explaining things to listening when I needed it, you always had my back, even when you had plenty going on yourself. Thank you for your kindness, creativity and wonderfully dry humor. I'm so happy we still see each other often, and that I now also get to spend time with little **Thomas** and not so little **Stephen**.

Bertalan, I miss our coffee breaks and conversations about politics, analyses and general nonsense. I also miss our Vertoef days with Renée and Ben, and all the dinner parties. Your cooking alone would be a good reason to want you back, but mostly I just miss your company. Nijmegen is not quite the same without you. That vineyard tour in Hungary still needs to happen.

Bob, you are exactly the sort of person I would love to keep working with in academia. You are grounded, thoughtful and critical, ask good questions and take a real interest in other people's work. From my first weeks to my last, I always felt welcome to come to you with questions, however small or complicated. Thank you for making so much time for me and for being someone I could always turn to. I have really valued working with you.

To the rest of the **Learning & Decision-Making lab**, thank you for making it such a good place to work. **Julie**, people say nobody is perfect, but I still haven't found the evidence in your case. **Rares**, thank you for your humor and kindness. I'm glad my negative reinforcement improved your bowling. **Helena**, you already seemed like a professional during your internship, and now you're approaching the end of your own PhD. I miss your baking. **Fleur**, I admire how calmly you take things in your stride. **Maikel**, you always made me laugh, although I did worry about which bone you would break next. **Rijk**, your enthusiasm was contagious. **Rico**, I'm so impressed by how you managed a double degree and still don't know how you made it look so easy.

Sjoerd, I'm convinced you were born a PI: bald, in khaki trousers, probably wearing a statement blouse and a denim jacket. Managing a team of assistants and interns seemed to come naturally to you from the start. I've always admired your ability to step back, see the bigger picture and work out what actually matters. The only role that has suited you even better is being Hannah's dad. She's growing into such a lovely little girl with so much personality, and I love seeing you together. I'm so glad we got to do our PhDs alongside each other.

Marwan, we grind, we struggle, we eat zure matten, we suffer. Nine hours a day, five days a week, with Tyrone Taylor's **Sufferation** as our soundtrack. You made those long

data-collection days bearable and, somehow, fun. We spent so much time together that we even started dressing alike without planning it. Thank you for the laughs and all your hard work. Chapter 2 would not have been the same without you.

Andrew, I knew we would be friends when you felt comfortable enough to mock my mason-jar lunch the second time we met. My only complaint is that you waited until last year to move to the Netherlands. Did you even consider my feelings? Think of all the brunches and laughs we missed. Please make up for it by staying a very long time. I love being your least-Dutch Dutch friend.

Carys, we hit it off at a conference in 2022, and I was so happy when you joined the group soon afterwards. You are always willing to help, whether something falls within your expertise or you need to figure it out alongside someone. I love that about you. You bring so much kindness and humor to the people around you, and I'm grateful that meeting you at a conference turned into getting to work with you and call you my friend.

Tulika, I first met you at your virtual poster during FUN21, in the first week of my PhD. I was nervous to ask questions because I understood approximately nothing. From the moment we started talking, you were so welcoming that I soon forgot to be nervous. You patiently walked me through your work and made me feel comfortable asking questions. That combination of intelligence and kindness is something I really admire in you. I have remembered that kindness ever since.

Julian, your efficiency seems to extend to everything: science, reading an improbable number of books and cycling in the perilous streets of Groesbeek, no less. But my favorite thing about you is how funny you are, sometimes without even trying, and especially after one (or two) beers.

And of course, to the rest of the **Cognitive Neuromodulation lab**, thank you for all the help, laughs and good company over the years. **Margely**, I have loved working with you. Your technical knowledge and professionalism impress me, and I admire the confidence with which you do your work. A true boss! **Joana**, please never change. You are warm, passionate, creative and so bright. **Vivek**, your kindness and enthusiasm for Friday drinks gave my final year a welcome boost. **Martin**, you were more organized on your first internship day than I was on my last PhD day. I'm glad you stayed for a PhD. **Kenneth**, what started with a message through Maartje turned into years of working together. I'll miss your dry humor. **Maie**, thank you for being so thoughtful and kind. **Javier**, thank you for always making me laugh, and for bringing your fan. **Jaël**, I like to think your work with Christian, and perhaps a little with me, helped lead you towards

your research with Marwan. You make a wonderful team. **Judith**, **Colette**, **Meike**, **Tinke**, **Jesse**, **Davide** and **Britt**, you were the first generation of Neuromod members and set the standard for those who followed. I'm glad I got to work with you. **Camilla**, I miss you. Please come back to Nijmegen. **Eleonora**, **Sebastian**, **Lutz**, **Carla**, **Ruggero**, **Sanne**, **Lena**, **Kate**, **Nico** and **Radovan**, thank you for the conversations, company and good moments along the way.

To **my interns**, supervising you was one of the best parts of this PhD. You contributed so much to the research, but just as importantly, you made the work more enjoyable.

Solenn, from our first interaction, it was clear you would go far. Luckily for me, you stayed for a lab rotation, a student assistantship and an internship. You work incredibly hard, and I'm proud to have been there at the start. Thank you for everything you contributed to Chapter 2. I can't wait to see what you do next.

Christian, you brought questions, initiative and a lot of energy to our meetings and the lab. Not many interns end up supervising other students, but you took on that responsibility with confidence. The groundwork you laid made Chapters 3 and 4 possible. I'm so proud that your determination brought you to your PhD in Hamburg.

Floris, I only ever had to explain something once. You are sharp, analytical and willing to question the obvious approach. I admire that you applied that same independence to your own future and chose your path after your master's.

Mihaela, I can easily picture you running a research group one day, doing something brilliant that I won't quite understand. Just as easily, I can picture you running an entire town somewhere in Iceland. Whatever you choose, I have no trouble imagining you making it happen. You are dedicated, smart and so kind. Thank you for choosing to do your internship with me. I enjoyed it so much.

Floris and Mihaela, Chapters 3 and 4 owe a great deal to your hard work.

The **Brainbox** has been such an important part of my PhD. **Andrew**, **Dan**, **Tonia**, **Soffía** and **Mark**, thank you for trusting me with workshops and talks, and for all the opportunities to work together. More than anything, I'm grateful that our shared interest in brain stimulation became a friendship that goes well beyond it.

Kim, **Elsa**, **Miriam**, **Til** and **Elly**, thank you for welcoming me into the ultrasound community and making conferences something to look forward to. **Kim**, you have been both an external mentor and a friend. The months you spent in Europe on sabbatical

were some of the best of my PhD.

Roshan, **Karin** and **Elsa**, thank you for being role models throughout my PhD, and for agreeing to be on my manuscript committee. I have learned so much from your work and our conversations, and it means a lot to me that the people I have looked up to over these years have now taken the time to read and assess this thesis.

Andrey and **José**, thank you for the expertise you brought to our work. Andrey, your help in setting up the study was absolutely essential. José, now that we eventually got it working, I can fondly remember our time at the scanner figuring out the MRS sequence.

To all the wonderful people at the **Donders Institute**, and especially my colleagues at the **DCC**, thank you for making me feel so at home over these years.

Vanessa, **Karin**, **Aleixa** and **Nicola**, you really are the heart of the DCC. Thank you for letting me wander into your offices to talk, laugh or complain. The warmth you brought to my working days mattered more than you probably realize. Each of you means so much to me, and I will always remember how welcome and cared for you made me feel.

Mieke, I have always admired your independence and determination, and the career you have built. You have been a role model to me and will continue to be. **Jiska**, I miss our Thursday morning meetings! **Miriam** and **Pauline**, thank you so much for everything, and for your patience whenever I emailed yet another question I probably should have known the answer to. **Paul**, thank you for your support and all the laughs at the MRI scanner. **Stein**, I'm so glad you joined the FUS team. **Norbert**, **Gerard**, **Pascale**, **Hubert**, **Sibrecht**, **Maarten** and **Twan**, building a new lab around a new method with a group of rookie PhD students cannot always have been easy. We could not have done it without you. Thank you so much.

Cleo and **Hongi**, you were my first mentors in academia, and both helped me decide to pursue a PhD. Cleo, you gave me clear guidance, trusted me with real independence and stood up for me when it mattered. I have never forgotten that. Hongi, supervising me during lockdown while moving to another country cannot have been easy, but you remained patient, kind and committed to helping me. Thank you both, from the bottom of my heart, for the care and trust you gave me at the very start.

Mariana, although you were never formally my supervisor, I have always considered you a mentor. I admire how you have handled the difficulties of academia, and I'm grateful for your encouragement through mine, from early stages during my internship

to the final stages of this PhD.

To **all the participants**, this work quite literally would not exist without you. Thank you for coming back for multiple sessions, letting us ruin your hair and continuing to do your best, even when the tasks became repetitive. I tried to tell you in person how much we appreciated your time and effort, and I hope you felt it. I also truly enjoyed getting to know you. Our conversations during preparation, about everything from chemistry and history to Eurovision, taught me so much and made those long days more enjoyable. Thank you for everything you brought to this research and to my time doing it.

Marjorie, I acknowledge you and thank you. Maar even serieus: zeggen dat jij het verschil hebt gemaakt voor mijn PhD doet je tekort. Je hebt het verschil gemaakt in mijn leven. Je bent eerder een zus dan een vriendin, inclusief het gekibbel. Bij jou kan ik alles zeggen, meestal omdat jij precies hetzelfde al denkt. *The blind leading the blind*, of gewoon *two dumb bitches*. Ik bewonder je loyaliteit, je kracht en hoe warm, slim en mooi je bent. Ik snap oprecht niet waar je de energie vandaan haalt: je werk, je sport, je sociale leven, en ondertussen ben je er ook altijd voor de mensen van wie je houdt. Ook je familie is me zo dierbaar geworden. **Lizet**, ik wacht nog steeds tot je bij me komt wonen. **Tim**, mijn lang verloren broer, je bent zo'n warm en lief mens, en ik moet nog een onderwerp vinden waar jij nίét iets over weet. Ik ben zo blij dat jullie elkaar hebben gevonden. Mar, We weten nog steeds niet wie destijds de ander op Facebook heeft toegevoegd. Jij. Maar dat maakt niet uit. Ruim tien jaar later weet ik nog steeds dat jij voor me opkomt en achter me staat. En als ik straks dit boekje moet verdedigen, sta je als mijn paranimf ook nog eens naast me. Ik kan me niemand voorstellen die beter op die plek past. Ik hou van je.

Buiten kantooruren waren er natuurlijk **mijn lieve vrienden**. **Anne**, **Bram**, **Emile**, **Jelle**, **Jolien**, **Jordi**, **Koen**, **Milou**, **Quint**, **Remy**, **Rosa**, **Rik**, **Tim**, **Twan** en **Veerle**, van de studententijd tot waar we nu zijn: iets volwassener, gelukkig niet per se serieuzer. We zijn allemaal bij de bowling begonnen en inmiddels op heel verschillende plekken terechtgekomen. Juist dat maakt het zo leuk om elkaars leven te blijven volgen. Ondanks mijn vriendschap met jullie heb ik het toch gehaald. Grapje. Al maakten jullie het me wel heel moeilijk om af en toe een borrel of feestje over te slaan, omdat het zo gezellig was. Dank jullie wel dat jullie daar altijd begrip voor hadden en me eraan herinnerden dat er meer was dan mijn PhD.

Mijn lieve **WGW**, jullie verdienen nog een apart plekje. **Milou**, ik bewonder hoe je alles voor elkaar hebt, samen met Jordi, en toch zo lief en betrokken blijft. **Jolien**, wat leuk dat we zelfs nog een tijd collega's zijn geweest! **Veerle**, hoe lang we elkaar soms ook niet zien, het voelt altijd alsof het gisteren was. **Rosa**, mijn oude huisgenootje, zullen

we weer dichter bij elkaar gaan wonen? Brakkenstein lijkt me prima, ook handig voor Riks voetbal. **Anne**, ik kan zo om je lachen. Onze NEC-postergirl met haar lieve PSV'er Bram.

Alexander, **Coen**, **Jan-Brecht**, **Jan-Luther**, **Maarten**, **Mees** en **Sjors**, Sebas heeft de allerliefste vrienden, en wat heb ik geluk dat jullie ook de mijne zijn geworden. Inmiddels kennen we elkaar al jaren, maar ik had hooguit twee ontmoetingen nodig om jullie in mijn hart te sluiten. Ik vind het zo mooi hoe hecht jullie zijn, sommigen al sinds de basisschool, en hoe vanzelfsprekend ik erbij mocht horen. En gelukkig hebben jullie ook nog eens een uitstekende smaak. **Roos**, met jou is het altijd fijn, of we nu serieus praten of gewoon wat borrelen. **Kim**, wat ben jij een leuke toevoeging aan de groep. En mijn **TPTW**, **Hannah**, **Maartje**, **Marianne** en **Nienke**: onze bijna maandelijkse afspraken maakten de mannen soms jaloers. "Jullie zien elkaar vaker dan wij." Klopt, en daar ben ik heel blij mee. Dank jullie wel voor jullie aandacht, gezelligheid en vriendschap.

Tamana, met niemand ben ik al zo lang bevriend als met jou. Het gevaar daarvan is dat we werkelijk alle fases samen hebben meegemaakt. Welke dat waren, zetten we maar niet zwart op wit. Jij bent familie en je ouders voelen ook een beetje als de mijne. Ik ben zo blij dat je **Wouter** hebt gevonden, en dat hij daardoor ook in mijn leven is gekomen. Ik hou van je.

Noortje, mijn trouwe labpartner, huisgenoot en een van mijn liefste vriendinnen. Microbiologiepractica na een avond stappen en dieren ontleden op een veel te lege maag: sommige geuren vergeet je helaas nooit. Inmiddels zijn we blijkbaar volwassen, en hebben jij en Olivier zelfs een prachtig kindje op de wereld gezet. **Olivier**, met jou kan ik uren praten, over politiek of werkelijk wat dan ook. Wat fijn dat jullie in mijn leven zijn.

Lieve **Jolanda** en **Erwin**, vanaf het eerste moment hebben jullie me als een eigen kind behandeld. Ik voel me bij jullie thuis, geliefd en helemaal onderdeel van het gezin. Dat zit in zoveel kleine dingen, in jullie aandacht, zorg en de vanzelfsprekendheid waarmee ik er altijd bij hoor. Ik hoop dat jullie weten hoeveel dat voor me betekent en hoeveel ik van jullie hou. Jullie hebben zo'n warm gezin, en ik ben ontzettend dankbaar dat ik daar deel van uitmaak. Dank jullie wel voor jullie lieve zoons, vooral die ene, en de schoonzussen die ik er daardoor bij kreeg.

En dan nu in mijn lievelingstaal, in het Nimweegs: **Job**, **Lieuwe**, jullie sien de bruurs die ik nooit heb gehad en (a'k sie hoe jullie onderling met mekuir umgaon soms) misschien oek nooit had wille hebbe? Grepke, ik hou fan jullie. **Amber**, **Lonneke**, ik bin so blij da ik jullie as schoonsus ien m'n lefe heb en da we onderdeel fan deselfde femilie meuge

sien. Ellek moment met de kids en de ouwe erbij fuult as 'n feest. Lieve schoonfamilie, ik hou echt van jullie.

Liefste mama, papa, Awa en mijn lieve Sebastiaan: dit proefschrift is voor jullie.

Awa, mijn liefste zusje, op Nala na. Het mooie aan een stuk ouder zijn, is dat ik je bewust heb zien opgroeien. Van de baby die ik, verblind door zussenliefde I guess, oprecht de mooiste ooit vond, tot de kleuter die zei: "Angelina Jolie is bijna net zo mooi als ik." Op de basisschool vond je het verschil tussen een 3 en een E nog niet zo interessant. De basisschool was voor plezier, leren zou je op de middelbare wel doen. En dat deed je. School, bachelor, master met magna cum laude en op je vierentwintigste be**ë**digd als advocaat. Gelukkig zit daar nog steeds mijn zusje in dat af en toe bruto en netto door elkaar haalt. Vroeger keek jij naar mij op. Inmiddels zijn de rollen omgedraaid. Je werk, je sociale leven, je kledingkast: ik kijk graag bij je af. We zijn in zoveel opzichten tegenpolen, maar juist daardoor staan we zo dicht bij elkaar. Ik ben zo dankbaar dat we zo'n goede band hebben, dat we alles kunnen delen en dat ik altijd op je kan rekenen. Jou als zusje hebben is een van de fijnste dingen in mijn leven. Ik hou ontzettend veel van je en ben zo trots op wie je bent. Dat je als mijn paranimf naast me staat en we dit samen mogen meemaken, maakt deze dag voor mij nog mooier.

Mam en **pap**, alles wat ik ben en alles wat ik heb, begint bij jullie. We hebben als gezin moeilijke tijden meegemaakt, maar ik heb nooit hoeven twijfelen aan jullie liefde of aan mijn plek bij jullie. Wat er ook speelde, jullie zorgden ervoor dat wij ons veilig en geliefd voelden. Hoe ouder ik word, en hoe dichter ik kom bij de leeftijd die jullie toen hadden, hoe beter ik begrijp hoeveel kracht en liefde dat heeft gevraagd. Ik zou willen dat ik in woorden kon uitdrukken hoe dankbaar ik daarvoor ben. Mama, je zorgzaamheid lijkt geen einde te kennen, en je creativiteit blijft me verbazen. Papa, je wilskracht, intelligentie en humor zijn dingen waar ik altijd naar heb opgekeken. Met maar een fractie van jullie kwaliteiten heb ik al zo'n mooi leven kunnen opbouwen. Dat zegt voor mij alles over hoeveel jullie in je hebben. In mijn ogen kunnen jullie werkelijk alles, en ik blijf me afvragen hoe jullie het allemaal hebben gedaan. Niets wat ik hier schrijf kan genoeg zijn om jullie te bedanken. Voor alles wat jullie hebben gedaan, alles wat jullie me hebben meegegeven en de liefde waarop ik altijd kan rekenen. Ik hoop dat jullie in dit boekje ook iets van jezelf herkennen, want zonder jullie was het er niet geweest. Alles wat ik doe, is ook voor jullie. Altijd. Ik hou van jullie.

Sebastiaan, ik kan een hele thesis over je schrijven en nog steeds niet genoeg hebben gezegd over wat je voor me betekent. Daarom houd ik het hier 'kort', want niets voelt ooit als genoeg. Met jou deel ik alles, de grote momenten en de gewone dagen waar ik minstens zo gelukkig van word. Ik ben zo dankbaar dat ik alles bij je kwijt kan en

dat ik me zo geliefd voel bij jou. Soms kan ik na al die jaren nog steeds niet helemaal geloven dat ik jou heb, en vraag ik me af waar ik zoveel liefde aan verdiend heb. Als ik terugkijk op onze jaren samen, ben ik trots op alles wat we hebben meegemaakt en opgebouwd. Nu geniet ik van het leven dat we samen hebben, van ons thuis en van het simpele geluk dat jij er bent. En als ik aan de toekomst denk, heb ik zo veel zin in alles wat we nog samen gaan meemaken. Er is niemand anders die ik naast me zou willen hebben. Dank je wel voor je liefde en voor de talloze manieren waarop jij mijn leven mooier maakt. Ik hou zo ontzettend veel van je.

DONDERS GRADUATE SCHOOL

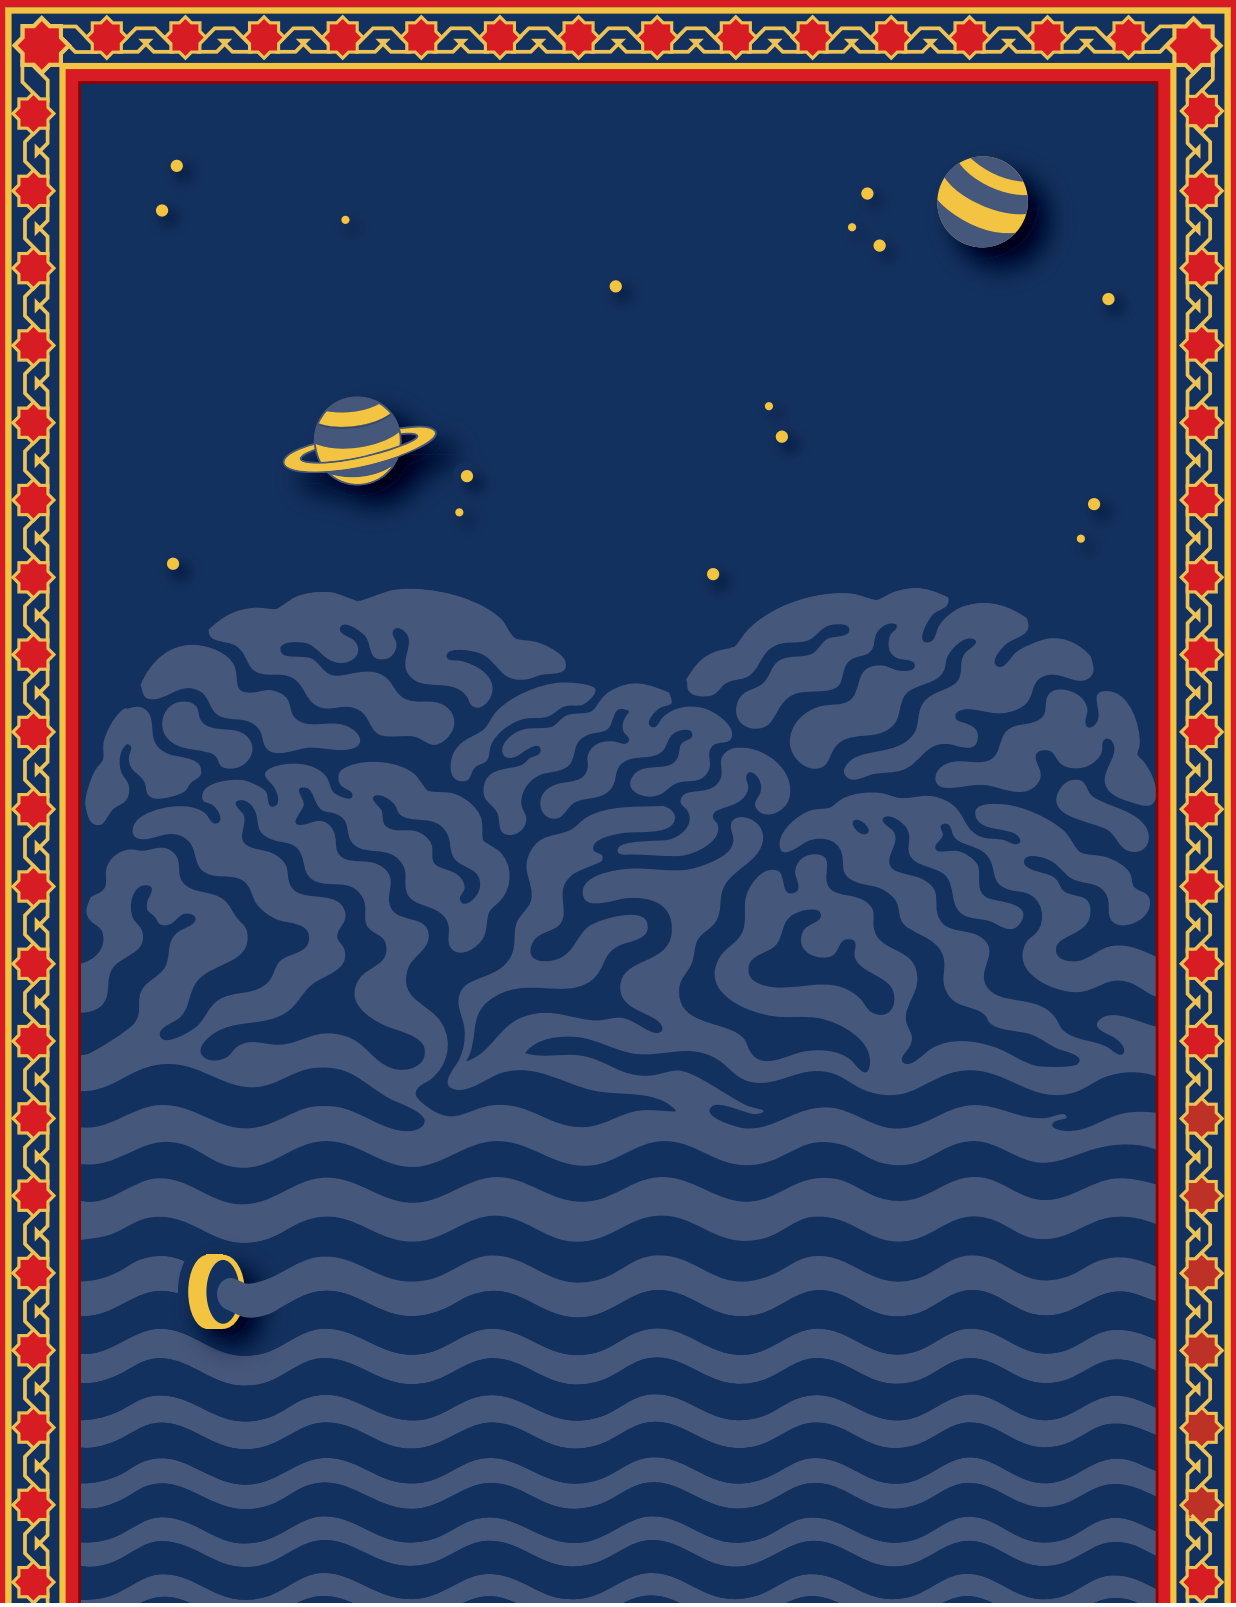
For a successful research Institute, it is vital to train the next generation of scientists. To achieve this goal, the Donders Institute for Brain, Cognition and Behaviour established the Donders Graduate School in 2009. The mission of the Donders Graduate School is to guide our graduates to become skilled academics who are equipped for a wide range of professions. To achieve this, we do our utmost to ensure that our PhD candidates receive support and supervision of the highest quality.

Since 2009, the Donders Graduate School has grown into a vibrant community of highly talented national and international PhD candidates, with over 500 PhD candidates enrolled. Their backgrounds cover a wide range of disciplines, from physics to psychology, medicine to psycholinguistics, and biology to artificial intelligence. Similarly, their interdisciplinary research covers genetic, molecular, and cellular processes at one end and computational, system-level neuroscience with cognitive and behavioural analysis at the other end. We ask all PhD candidates within the Donders Graduate School to publish their PhD thesis in de Donders Thesis Series. This series currently includes over 750 PhD theses from our PhD graduates and thereby provides a comprehensive overview of the diverse types of research performed at the Donders Institute. A complete overview of the Donders Thesis Series can be found on our website: https://www.ru.nl/donders/donders-series

The Donders Graduate School tracks the careers of our PhD graduates carefully. In general, the PhD graduates end up at high-quality positions in different sectors, for a complete overview see https://www.ru.nl/donders/destination-our-former-phd. A large proportion of our PhD alumni continue in academia (>50%). Most of them first work as a postdoc before growing into more senior research positions. They work at top institutes worldwide, such as University of Oxford, University of Cambridge, Stanford University, Princeton University, UCL London, MPI Leipzig, Karolinska Institute, UC Berkeley, EPFL Lausanne, and many others. In addition, a large group of PhD graduates continue in clinical positions, sometimes combining it with academic research. Clinical positions can be divided into medical doctors, for instance, in genetics, geriatrics, psychiatry, or neurology, and in psychologists, for instance as healthcare psychologist, clinical neuropsychologist, or clinical psychologist. Furthermore, there are PhD graduates who continue to work as researchers outside academia, for instance at non-profit or government organizations, or in pharmaceutical companies. There are also PhD graduates who work in education, such as teachers in high school, or as lecturers in higher education. Others continue in a wide range of positions, such as policy advisors, project managers, consultants, data scientists, web- or software developers, business owners, regulatory affairs specialists, engineers, managers, or IT

architects. As such, the career paths of Donders PhD graduates span a broad range of sectors and professions, but the common factor is that they almost all have become successful professionals.

For more information on the Donders Graduate School, as well as past and upcoming defences please visit:
http://www.ru.nl/donders/graduate-school/phd/



DONDERS INSTITUTE

Radboud University

Radboudumc