# Rewired by time: a dopaminergic account of cognitive flexibility across the mouse lifespan

A thesis presented

by

Xiang Yi Sean Meng

Presented to

The Faculty of the Committee on Degrees in Neuroscience

in partial fulfillment of the requirements for the degree of

Bachelor of Arts

Harvard College

Cambridge, Massachusetts

March 13, 2026

# The Harvard College Honor Code

Members of the Harvard College community commit themselves to producing academic work of integrity – that is, work that adheres to the scholarly and intellectual standards of accurate attribution of sources, appropriate collection and use of data, and transparent acknowledgement of the contribution of others to their ideas, discoveries, interpretations, and conclusions. Cheating on exams or problem sets, plagiarizing or misrepresenting the ideas or language of someone else as one's own, falsifying data, or any other instance of academic dishonesty violates the standards of our community, as well as the standards of the wider world of learning and affairs.

# The Neuroscience Concentration AI Policy

Completing a thesis develops valuable skills in analytical thinking and scientific writing. **Students are strictly forbidden from including AI-generated content in any part of their thesis manuscript, including preliminary drafts**. Incorporating AI-produced text into your thesis constitutes an Honor Code violation at Harvard College and will be considered academic misconduct.

Certain applications of generative AI tools are acceptable during other research phases. For example, these tools may be valuable for broadening or deepening your knowledge of specific concepts, understanding methodologies, or troubleshooting analyses. However, exercise caution when using these resources: current AI models regularly produce inaccurate information. This is especially true for scientific research, as journals are often paywalled and thus under-represented in training data.

**You bear full responsibility for the accuracy of all assertions in your thesis**, making it essential to rigorously verify any AI contributions. You should be prepared to answer questions about any part of your thesis during the oral conference.

If you are unsure as to whether a particular use is allowed, contact your concentration advisor promptly.

By signing below, I confirm that I have adhered fully to all components of the Honor Code and AI Policy in the production of this thesis.

*Digital Signature:* 

# Acknowledgments

I would like to start by extending my most heartfelt gratitude to Dr. Bernardo Sabatini for welcoming a wide-eyed freshman into his lab and allowing him to pursue science to his heart's extent. Bernardo, the freedom and trust you gave me so early changed the course of not only my time at Harvard but a whole career and life ahead. The way you think about science, the standards you hold, and the mentor you are have left a mark on me that I will always carry.

To Dr. Kevin Mastro, the mentor who I owe it all to, thank you for taking a chance on me. You took me under your wing, you believed in me before I believed in myself, and you have shaped how I think, how I question, and how I approach the curiosities of the world in ways I discover more of every day. I am so deeply grateful for you and for everything you have given me.

To Wen Chun Lee, Nathan Mallet, Maxwell Kounga, Khalila Lord, and Arond, thank you for being such wonderful colleagues and friends. We have done some exciting science together, and through all of it you have all been generous with your time, your ideas, and your support, and working alongside you has been an immense pleasure.

To my blockmates Max, Tyler, Jeffrey, Alex, and Corwin, you are the most inspiring group of people I have ever had the pleasure of calling my best friends. You motivate me every single day, you make me laugh until it hurts, and I know that even at the most difficult moments we are never more than a phone call away. and I am so lucky to have done Harvard with you. I look forward to our lifetime of friendship.

To my family, who are the foundation of everything. Mom and Dad, I am so lucky to be your son. You have supported every dream I have ever had without hesitation, and the fact that I get to pursue science freely is because you have built a nest of safety, stability and love through years and years of hard work and devoted parenting. I hope you know how much that means to me and how grateful I am for the life you have made possible. Philip, watching you grow has been one of the greatest joys of my life. You inspire me with who you are and who you are becoming, and I look forward to the day we take on the world together. Joel, thank you for being my mentor, my friend, and my advocate. You have believed in me with a generosity I am not sure I deserve, and that belief has carried me further than you know. Having you in my corner has meant everything

And to Doug Melton and Steven Hyman, thank you for showing me what it looks like to dedicate a life to science with vision and genuine purpose. I aspire to be the scientists that you are.

# List of Contributions

The research presented in this thesis was conducted in the laboratory of Professor Bernardo Sabatini at Harvard University. I joined the Sabatini Lab as a freshman and have worked in the lab across four years, supported by funding from the KRANIUM program, KURE (term-time), the Herchel Smith Undergraduate Research Fellowship, the Yun Family Fellows program, the TIME Fellowship, and the Sabatini Lab. I have received course credit via NEURO99 for Spring (2025-2026) and Fall (2026).

The project was conceived and designed jointly by Dr. Kevin Mastro and myself. All resources, animals, and workspaces were provided by the Sabatini Lab.

The full aging behavioral cohort for the two-armed bandit task was run jointly by myself, Wen Chun Lee, Khalila Lord-Arond, and Maxwell Kounga. The two-armed bandit task pipeline was developed by Dr. Kevin Mastro, who trained me in its use. All behavioral analysis and computational modeling were performed by myself with guidance and support from Dr. Kevin Mastro. The recursively formulated logistic regression model was originally developed by Dr. Celia Beron and adapted here for application to the aging cohort.

Fluorescence lifetime photometry recordings in freely moving mice were conducted jointly by myself and Dr. Kevin Mastro. Head-fixed fiber photometry recordings were conducted jointly by myself and Nathan Mallet. The FLIPR system and associated recording approach were developed by Dr. Bart Lodder in the Sabatini Lab.

Progressive ratio behavioral experiments were designed and conducted by myself.
The single-nucleus RNA sequencing dataset analyzed in Chapter 5 is a publicly available dataset from Linker et al. (2025), Nature Communications. All reanalysis of this dataset was performed independently by myself.

All figures and written work in this thesis are my own unless otherwise stated.

# Table of Contents

# Abstract

The ability to flexibly update decisions based on changing outcomes is fundamental to adaptive behavior and is known to decline across neurological and psychiatric conditions. Whether this capacity deteriorates during normative aging, and what circuit mechanisms underlie any such change, remains poorly understood. Using a two-arm bandit probabilistic reversal task in freely-moving mice, we show that aging progressively impairs reinforcement-guided updating across the lifespan. Older mice accumulated more unrewarded outcomes and were slower to re-establish preferences following contingency reversals, despite intact task acquisition, preserved exploration, and maintained motivational effort. Computational modeling revealed that aging selectively reduced sensitivity to recent reward outcomes while prolonging the persistence of prior evidence, together producing a signature of rigid, history-bound decision making. Tracking dopamine dynamics in the nucleus accumbens across the lifespan using fluorescence lifetime photometry, we found that tonic dopamine baseline declined progressively with age while phasic release amplitudes were preserved, compressing the absolute dynamic range of reward signaling without altering its relative structure. Reanalysis of published single-nucleus RNA sequencing data from young and aged striatum revealed coordinated transcriptional changes across cell types, including reactive microglial activation, loss of astrocytic glutamate transporters, and bioenergetic stress in neurons, consistent with degraded fidelity of reinforcement signals and impaired translation into behavioral updating. Together, these findings establish a normative trajectory of cognitive aging and correlatively link progressive impairment in outcome-guided decision making to the erosion of dopaminergic tone within a broadly remodeled striatal environment.

# I. Introduction

## A. Plasticity as the Foundation of Learning Across the Lifespan

The familiar saying that “you can’t teach an old dog new tricks” reflects a widely observed phenomenon: young individuals acquire new skills rapidly and adapt readily to changing circumstances, whereas older individuals often learn more slowly and exhibit diminished flexibility.

This intuition has a clear neurobiological basis. Cognitive flexibility depends on the ability of neural circuits to convert transient patterns of activity into persistent modifications (Citri & Malenka, 2008; Pan & Monje, 2020). Synaptic plasticity provides the substrate for these adjustments, allowing recently active connections to be strengthened, weakened, or reorganized according to environmental demands (Oberman & Pascual-Leone, 2013). Through these processes, organisms form predictive relationships between cues and outcomes, form expectations about value, refine their behavioral policies, and revise those policies when contingencies change (Whyte et al., 2019).

A large body of evidence demonstrates that this flexibiity is not preserved uniformly across the lifespan. (Oberman & Pascual-Leone, 2013; Voglewede et al., 2019; Marzola et al., 2023). Older humans exhibit characteristic impairments in reinforcement learning, reversal learning, and other tasks requiring flexible updating (Breton et al., 2015). Rodents display parallel deficits, including slower adaptation to changing reward contingencies and reduced capacity to shift strategies (Schoenbaum et al., 2002; Zhang et al., 2022; Tait et al., 2013). These behavioral changes align with convergent findings at the level of electrophysiology and molecular shifts. Aging is

associated with reduced induction and maintenance of long-term potentiation (LTP), increased susceptibility to long-term depression (LTD), diminished dendritic spine density and turnover, and changes in activity-dependent transcriptional programs required for late-phase plasticity (Norris et al., 1996; Foster & Kumar, 2007; Ducote et al., 2024; Boric et al., 2008; Wang, 2008; Chepkova et al., 2015)

Notably, plasticity does not disappear in older individuals. Learning remains possible, and new skills can still be acquired. The difference lies in the probability and magnitude of plastic changes. With age, the system becomes more strongly biased toward stability rather than modification. This bias may confer advantages by protecting established structure and accumulated experience, but it reduces the capacity to adjust when demands shift. (Pauwels et al., 2018; James et al., 2024) The proverb about older learners, therefore, reflects a quantitative shift in the responsiveness of plasticity mechanisms, rather than a categorical loss of learning capacity.

It is tempting to attribute these age-related learning deficits solely to changes within synapses. Yet synaptic plasticity is not governed exclusively by local patterns of activity. Experience-dependent modification of synaptic strength is tightly regulated by neuromodulatory systems that specify when plasticity should occur and which representations should be updated (Speranza et al., 2021). Among these systems, dopamine plays a particularly influential role in reinforcement learning. Dopaminergic signals determine when synapses become eligible for long-term modification, how strongly recent outcomes shape synaptic weights, and which components of an internal model are revised following unexpected events (Pignatelli & Bonci, 2015; Vidal-Gadea & Pierce-Shimomura, 2012). Alterations in dopaminergic signaling,

therefore, offer a compelling explanation for how plasticity can remain structurally intact yet functionally weakened across aging.

## B. Dopamine as a Neuromodulatory Gate for Plasticity and Learning

The role of dopaminergic signalling in plasticity and synapse modification emerges from its anatomical organization, signalling dyanmics, and computational function.

Dopamine neurons in the ventral tegmental area (VTA) and substantia nigra pars compacta (SNr) project and receive input broadly across the striatum and frontal cortex, among many other regions. (Menegas et al., 2015; Trutti et al., 2019; Cai & Tong, 2022) These projections allow dopamine to modulate both excitability and the induction of long-term synaptic changes in the striatal and cortical circuits heavily responsible for action selection, value estimation, and behavioral adaptation (Haber, 2011).

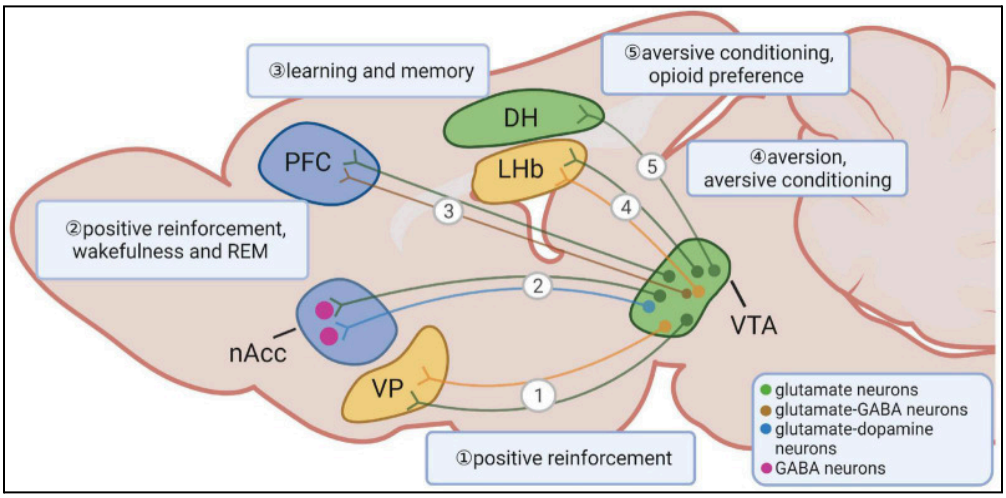


***Schematic 1:*** *Long-range projections of VTA glutamate and dopamine neurons and their major targets (Cai & Tong, 2022).*

Dopamine acts primarily through two families of G protein–coupled receptors. D1-class receptors couple to Gs and Golf proteins and stimulate the production of cyclic AMP, thereby activating protein kinase A (PKA) (Zhuang et al., 2021; Teng et al., 2022). D2-class receptors couple to Gi and Go proteins and inhibit adenylyl cyclase, reducing cyclic AMP and suppressing PKA activity (Gerfen et al., 1990). Once engaged, these cascades exert powerful control over the cellular machinery of plasticity. PKA phosphorylates ion channels and NMDA receptor subunits to increase calcium influx, enhances AMPA receptor trafficking and synthesis, and activates transcriptional regulators such as CREB that consolidate long-term changes (Skeberdis et al., 2006; Iino et al., 2020; Lau et al., 2004; Nayak et al., 1998). Through these PKA-dependent modifications, dopamine converts transient synaptic activity into lasting changes in synaptic strength, determining whether recently active inputs undergo long-term potentiation (LTP), long-term depression (LTD), or remain unchanged.

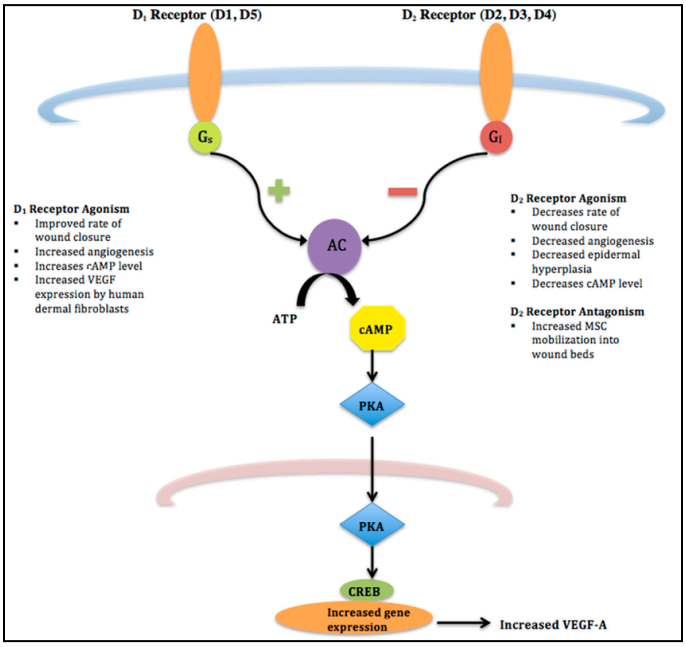


***Schematic 2:*** *Divergent G-protein signaling cascades downstream of D1 and D2 receptors (Vaughn et al., 2018)*

A defining computational property of midbrain dopamine neurons is their encoding of reward prediction errors (RPEs): transient increases in firing when outcomes exceed expectations and brief pauses when outcomes fall short (Schultz, 2016; Rescorla & Wagner, 1972). This bidirectional pattern directly parallels the error term in temporal-difference learning, where the brain updates value estimates in proportion to the mismatch between predicted and actual outcomes (Dayan & Daw, 2008; Katahira, 2015). In vivo, this error signal is implemented by fast phasic fluctuations in extracellular dopamine concentration.

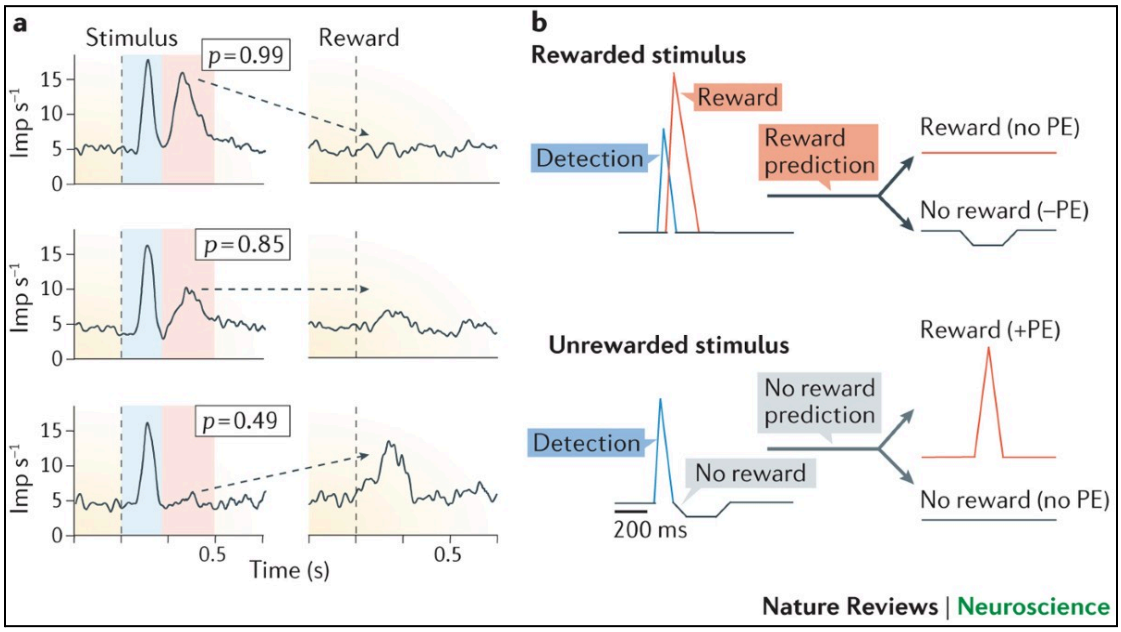


***Schematic 3:*** *Sustained value coding in dopamine prediction-error signals (Schultz, 2016).*

RPEs guide learning by providing the evaluative signal that determines whether recent neural activity should be strengthened or weakened. When an animal makes a choice, the associated neural populations generate a short-lived biochemical “eligibility trace”, which serves as a molecular footprint marking recently active synapses as temporarily modifiable (Yamaguchi et

al., 2022). On their own, these traces do not specify whether plasticity should occur. Instead, the subsequent RPE supplies the critical information: a positive RPE elevates dopamine above baseline, flagging the preceding synaptic activity as behaviorally beneficial, whereas a negative RPE lowers dopamine below baseline, indicating that those same synapses contributed to an unfavorable outcome (Shouval & Kirkwood, 2025; Yagishita et al., 2014).

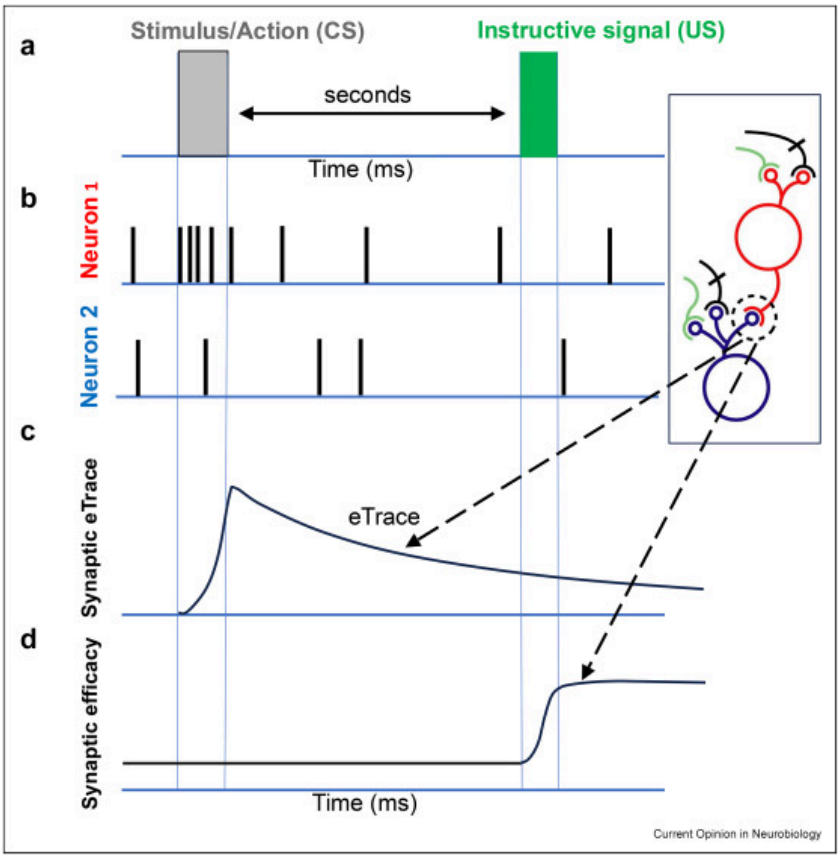


***Schematic 4:*** *Eligibility traces bridge delayed reinforcement in synaptic learning. (Shouval & Kirkwood, 2025)*

In this framework, dopamine transients do not simply encode reward, but *gate* when plasticity is permitted and assign a direction to that plasticity. By linking the timing of neural activity to the

evaluative signal carried by RPEs, dopamine enables the brain to incrementally refine its internal estimates of action value and improve future decisions. This coupling between eligibility traces and phasic dopamine thus forms a core mechanism through which prediction errors are converted into lasting changes in circuit function.

To understand how effectively these phasic teaching signals can drive plasticity, it is also essential to distinguish them from the slower tonic dopamine levels on which they are superimposed. Phasic dopamine transients represent rapid, millisecond-to-second scale deviations from a background concentration that changes far more slowly (Floresco et al., 2003; Grace, 1991). This tonic level sets the dynamic range through which phasic signals operate and therefore influences the effective amplitude of a prediction error (Floresco et al., 2003; Grace, 1991). A phasic rise that begins from a low baseline may not reach the biochemical threshold required to activate downstream signaling, whereas the same rise from a higher baseline may produce a much larger intracellular response (Romero Pinto & Uchida, 2025). In this sense, tonic dopamine establishes the gain with which phasic fluctuations are converted into molecular events that permit synaptic change. Variations in tonic dopamine can therefore modulate how strongly prediction errors shape learning without altering the structure of the phasic signal itself.

These principles give dopamine a dual role in reinforcement learning circuits. At the computational level, dopamine transients carry the reward prediction error that determines the direction of learning. At the cellular level, dopamine controls the probability and magnitude with which recently active synapses undergo long-term modification.

## C. The Nucleus Accumbens as a Circuit Substrate for Dopamine-Dependent Plasticity

The nucleus accumbens (NAc) occupies a central position in the basal ganglia networks that support reward-guided behavior. As the ventral extension of the striatum, it forms part of cortico-basal ganglia loops that link prefrontal, limbic, and midbrain regions (Xu et al., 2024; Klawonn & Malenka, 2019; Yin et al., 2008). The medial prefrontal cortex (mPFC) and orbitofrontal cortex (OFC) send inputs that represent action plans, task rules, and value estimates (Hikida et al., 2020; Lai et al., 2022). The hippocampus contributes contextual and spatial information, while the basolateral amygdala conveys learned associations between cues and outcomes (Ibrahim et al., 2024; Klawonn & Malenka, 2019). Additional signals related to arousal, interoception, and salience arrive from thalamic and insular regions (Xu et al., 2024). On top of this rich convergence is dense dopaminergic innervation from the VTA, which delivers the moment-to-moment prediction error signals described above (Castro & Bruchas, 2019). This anatomical arrangement positions the accumbens as a site where information about what the animal did, what it expected, and what actually happened can be jointly evaluated and encoded.

The principal neurons of the accumbens are spiny projection neurons (SPNs), which comprise the vast majority of the local neuronal population and act as the biological substrates through which dopamine signals are transformed into changes in synaptic strength (Kravitz & Kreitzer, 2012; Shen et al., 2008). SPNs are broadly divided into two populations based on dopamine receptor expression and projection targets. D1-expressing, striatomesencephalic SPNs project along the direct pathway and tend to facilitate selected actions. D2-expressing, striatopallidal SPNs project along the indirect pathway and tend to suppress competing or maladaptive actions (Gerfen et al., 1990). Although this categorical distinction is not absolute, it captures a

fundamental computational feature: the accumbens contains parallel output channels that can differentially promote or inhibit behavior.

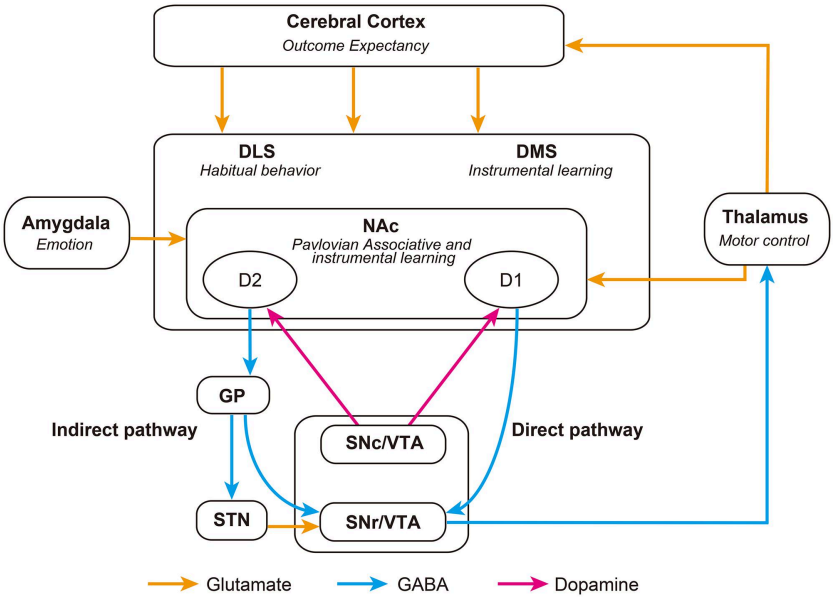


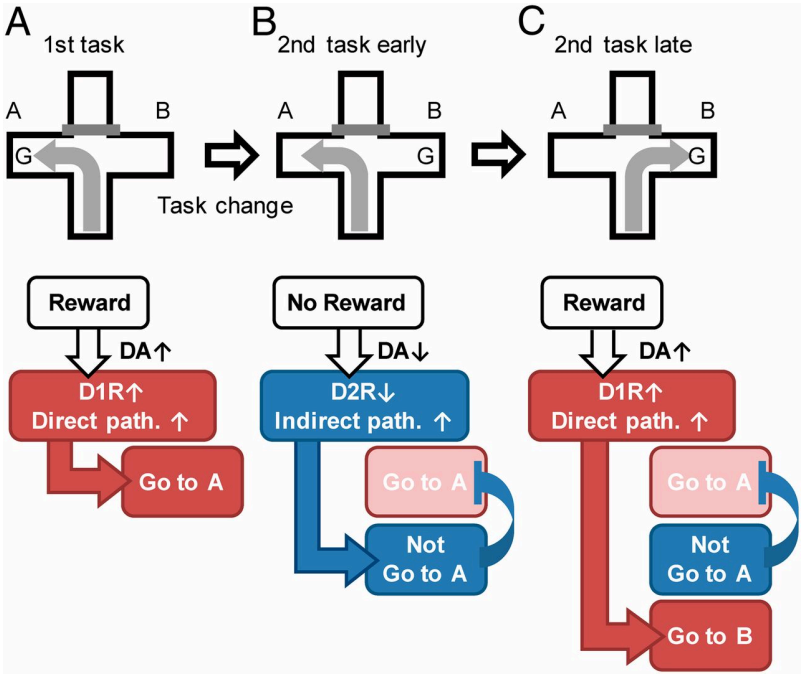


***Schematic 5:*** *Representations of striatal-associated decision-making neurocircuitry (Macpherson et al., 2014; Yawata et al., 2012)*

Transient increases in dopamine, corresponding to positive prediction errors, preferentially engage D1-receptor signaling in recently active neurons and promote forms of long-term potentiation at their excitatory inputs (Lee et al., 2021; Steinberg et al., 2014). Conversely, transient dips below baseline, corresponding to negative prediction errors, reduce D2-receptor occupancy and bias intracellular pathways toward synaptic weakening at inputs linked to actions that yielded worse-than-expected outcomes (Lee et al., 2021; Romero Pinto & Uchida, 2025). Iterated over many trials, this bidirectional modulation provides a mechanism by which the accumbens can strengthen inputs that predict rewarding consequences and weaken those that signal disappointing ones. Recent in vivo imaging offers a proof-of-concept demonstration that this canonical model operates during natural behavior. Optical measurements that combine dopamine sensors with reporters of PKA activity reveal that brief elevations in dopamine drive PKA activation predominantly in D1-expressing neurons, whereas brief dopamine dips activate PKA primarily in D2-expressing neurons. These cell-type-specific patterns show that the sign of the prediction error is preserved in the polarity of intracellular kinase activity: positive errors produce biochemical states conducive to synaptic strengthening, while negative errors yield states that favor synaptic weakening. Together, these findings validate the idea that dopamine fluctuations are transformed into complementary, population-specific plasticity signals that refine action tendencies within the accumbens (Lee et al., 2021).

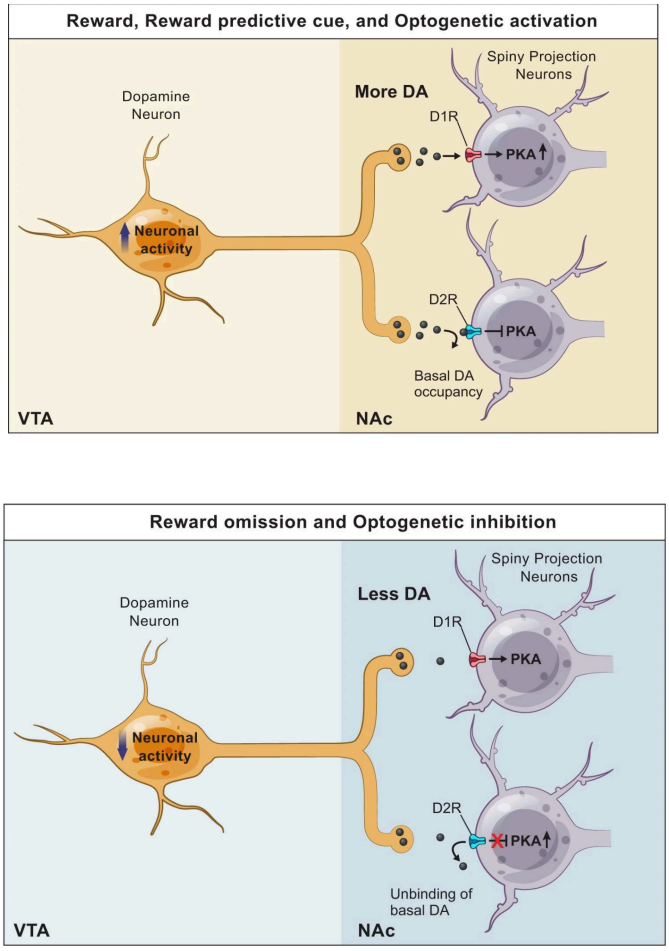


***Schematic 6:*** *Opposing dopamine control of PKA signaling in NAc D1 and D2 SPNs (Lee et al., 2021)*

Against the anatomical and functional backdrop, dopamine-dependent plasticity provides the mechanism through which reinforcement learning is implemented. The convergence of cortical, limbic, and hippocampal inputs onto SPNs supplies eligibility traces that mark recently active synapses, while dopaminergic prediction errors provide the teaching signals that determine whether these traces are strengthened or weakened. As a result, the accumbens functions as a circuit-level substrate where experience-dependent updating of action values is instantiated.

## D. Aging-Related Alterations in Dopamine and Consequences for Learning

Dopaminergic function has been shown to decline across healthy aging, though the form of this decline varies and no single mechanistic account explains all observations (Morrison & Hof, 1997). Several studies report steady reductions in striatal dopamine content from early adulthood into late life, with dopamine and its metabolites falling to a fraction of their peak levels (Goldman-Rakic and Brown 1981; Kish et al., 1992). These reductions frequently occur even when presynaptic terminals remain structurally present and when markers of dopamine synthesis, vesicular packaging, and reuptake show relatively modest or region-specific change (Kish et al., 1992). This finding suggests that aging can limit the amount of transmitter available for release while leaving much of the underlying anatomical scaffold intact.

Other observations emphasize alterations in postsynaptic sensitivity. In several rodent models, total striatal dopamine concentrations remain stable well into advanced age, yet the availability of D1 and D2 receptors declines, and intracellular signaling responses become weaker (Karrer et al., 2017). Various human imaging studies echo this pattern, but also reveal substantial individual variability in the extent and spatial distribution of receptor and transporter reductions (Dang et al., 2017). These datasets indicate that aging often affects the ability of postsynaptic neurons to detect and respond to dopamine even when presynaptic stores of transmitter are relatively preserved. Under this view, dopaminergic decline reflects changes in both the supply of dopamine and the sensitivity of the neurons depending on it.

Age-related vulnerability in midbrain dopamine neurons adds another dimension to this process. Structural imaging studies report reduced integrity within the substantia nigra in older adults, and histological work confirms modest neuronal loss over the lifespan (Fearnley and Lees, 1991).

Functional measurements in animal models show that dopamine release becomes smaller and less temporally precise before substantial structural degeneration becomes apparent. Age-related shifts in tonic dopamine levels and alterations in corticostriatal modulation further indicate that the regulatory influence of dopamine on circuit function weakens significantly during aging (Rodriguez et al., 2014). These impairments reveal that dopaminergic decline can take the form of disrupted release dynamics even when tissue architecture appears largely preserved.

Taken together, these findings suggest that dopamine function weakens with age, yet the underlying components do not change in a uniform or predictable manner. Presynaptic markers, postsynaptic receptors, release dynamics, tonic dopamine levels, and midbrain microstructure each exhibit age-related alterations, but they do so with different magnitudes and timelines and with considerable variability across individuals. Some older organisms show pronounced reductions in dopamine content, others display prominent postsynaptic decline, and others exhibit early functional impairments in release despite relatively stable biochemical measures (Berry et al., 2019). This variability has prevented the field from establishing a single mechanistic framework for dopaminergic aging. The most consistent conclusion is that the amplitude, timing precision, and reliability of dopamine signaling diminish across the lifespan, although the pathways that lead to this deterioration differ across studies and individuals.

These biological changes are paralleled by alterations in reward-related neural processing. Functional neuroimaging studies show that older human adults often exhibit reduced ventral striatal engagement during reward anticipation and altered timing of reward responses in both striatum and midbrain (Eppinger et al., 2013; Dreher et al., 2008). Aging rodents display comparable reductions in dopamine-dependent responses to reward-predictive cues (Tryon et al., 2020; Lei et al., 2025).

Behaviorally, these neural changes manifest as well-characterized deficits in reinforcement learning. Older adults and aged rodents learn more slowly from positive outcomes, update value estimates less effectively, and struggle to flexibly adjust choices when reward contingencies shift (Cools et al., 2009). In dynamic or probabilistic tasks, many older individuals show reduced trial-by-trial sensitivity to feedback and a greater tendency toward perseverative strategies (Rutledge et al., 2009). These patterns resemble the behavioral consequences of experimentally reducing dopaminergic signaling, further supporting the idea that age-related reinforcement learning deficits reflect changes in dopamine biology.

Pharmacological studies strengthen this link. In older adults, enhancing dopaminergic transmission with levodopa can increase learning rates and restore more complete reward prediction error signals in the nucleus accumbens (Chowdhury et al., 2013). The degree of behavioral improvement varies across individuals, but in those who benefit, dopamine enhancement reinstates the canonical relationship between expected value, reward outcome, and striatal prediction-error responses (Chowdhury et al., 2013; Moustafa, 2011). These findings offer direct evidence that dopaminergic insufficiency contributes to impaired reward-guided learning in aging.

## E. Central Knowledge Gap: How Does Aging Alter Dopamine-Dependent Plasticity in the Nucleus Accumbens?

Although aging is known to disrupt dopamine signaling and impair reinforcement learning, the mechanisms that link these changes remain poorly understood. Dopamine shapes the plasticity

rules through which the NAc integrates reward history, yet it is not known how aging alters the mapping from dopamine signals to synaptic plasticity, or from plasticity to behavior. This thesis addresses this gap by employing a behavioral assay of cognitive flexibility and *in vivo* fluorescence-lifetime measurements of dopamine. The two-armed bandit task (2-ABT) provides a readout of trial-by-trial reward learning, outcome sensitivity, and adaptability when contingencies change (Mastro et. al, 2024). Lifetime dopamine photometry (FLIP-R) enables direct quantification of extracellular dopamine fluctuations independent of intensity-based distortions, following recent advances that demonstrate stable lifetime-based dopamine readouts in behaving animals (Lodder et al., 2025). By tracking these signals across the mouse lifespan, this work seeks to investigate the potential role of age-dependent neuromodulator dynamics in neural computation and behavior.

The central hypothesis is that aging reduces dopamine in the nucleus accumbens, thereby altering the functional impact of dopamine signaling on reinforcement-dependent behavioral updating. As a consequence, decision-making becomes less sensitive to recent outcomes, reinforcement learning slows, and behavioral flexibility declines.

**Aim 1. Define how reinforcement learning and cognitive flexibility change across the lifespan:** Mice at adolescent, adult, and aged timepoints will perform a 2-ABT that requires continuous updating of reward expectations and rapid adaptation when contingencies reverse. Computational modeling previously developed in the lab will be utilized to quantify learning rates, outcome sensitivity, exploration tendencies, and perseveration (Beron et al., 2022). I expect that aging will slow reward-based updating, reduce flexibility after reversals, and increase reliance on habitual or perseverative policies. This aim establishes the behavioral phenotype and

provides interpretable parameters that directly connect to neuromodulatory changes measured in subsequent aims.

**Aim 2. Quantify lifespan changes in tonic and phasic dopamine signals using fluorescence-lifetime dopamine photometry:** To measure both basal dopamine tone and fast reward-evoked dopamine fluctuations in the nucleus accumbens, I will use fluorescence lifetime dopamine photometry (FLIP-R). Lifetime-based readouts provide a stable and quantitative measure of neuromodulatory dynamics across age. I predict that aged animals will exhibit reduced dopamine levels and thus diminished reward/information-evoked phasic dopamine responses. These changes are expected to weaken the effective teaching signal that underlies updating.

**Aim 3. Identify age-dependent structural and molecular correlates of dopamine decline using histology and transcriptomic profiling:** To determine mechanisms underlying age-related dopamine reductions, I will perform transcriptomic profiling via a re-analysis of a public mouse striatum scRNA-seq dataset at 3 month and 18 month age points. I will seek to identify candidate genes and age-dependent transcriptional changes. I predict that aging will be associated with structural or molecular adaptations with transcriptional basis in dopamine-relevant circuits.

# II. Materials and Methods

All experimental procedures were conducted in accordance with the National Institutes of Health guidelines and were approved by the Harvard Medical School Institutional Animal Care and Use Committee (IACUC).

**Animals**

Only Male mice were used across experiments. Wild-type C57BL/6J mice (n=19) used for lifetime dopamine photometry and FR/PR experiments were obtained from the National Institute on Aging (NIA) Aged Rodent Colonies, which maintain C57BL/6J mice across a broad age range (1–30 months). D1-Cre C57BL/6J mice (n=6) used for dopamine and D1-SPN/D2-SPN PKA (not included in thesis) experiments were obtained from The Jackson Laboratory. Wild-type C57BL/6J mice (n= 22) used for head-fixed dopamine photometry were obtained from The Jackson Laboratory. These animals formed the primary cohort for experiments examining age-dependent dopamine signaling and behavior.

Wild-type C57BL/6J mice (n=80 total; n=76 progressed to 80/20) used for 2-arm bandit behavioral experiments were obtained from The Jackson Laboratory.

Mice were housed in a temperature- and humidity-controlled vivarium on a 12 h light/dark cycle with ad libitum access to food. Behavioral experiments were conducted during the dark phase of the light cycle. At the time of surgery, animals spanned early adulthood through advanced aging, enabling cross-sectional and longitudinal analyses across the lifespan. All animals were

group-housed following surgical implantation, singly housed following the onset of behavioral experiments.

**Surgeries**

Mice were anesthetized with isoflurane (induction at 3–4%, maintenance at 1–2%) and placed in a stereotaxic frame. Body temperature was maintained using a closed-loop heating pad, and depth of anesthesia was continuously monitored. Following scalp incision and skull exposure, bregma and lambda were aligned in the dorsoventral plane. Craniotomies were drilled above target regions using stereotaxic coordinates referenced to bregma.

For nucleus accumbens core (NAcc) targeting, viral injections were performed at the following coordinates: anterior–posterior (A/P) +1.3 mm, medial–lateral (M/L) ±1.3 mm, dorsal–ventral (D/V) −4.1 mm. Injection pipettes were lowered beyond the target depth (−4.15 mm) and retracted to the final depth prior to infusion to minimize backflow. Viruses were delivered using pulled glass pipettes at a rate of 80 nL/min, with total volumes of 300 nL per site unless otherwise specified. Following infusion, pipettes were left in place for an additional 5 min before slow withdrawal.

Following viral injection, the scalp was then retracted to fully expose the skull, which was cleaned, lightly etched, and dried to promote adhesion. Skin edges were secured to the skull perimeter using tissue adhesive (Vetbond) to prevent soft tissue encroachment. Pre-cut optical fibers (200/230 μm core/cladding, NA 0.37, 4.5 mm length; Doric Lenses) were mounted in a stereotaxic implant holder and lowered slowly to a position 100 μm dorsal (-4.0 mm) to the injection site. Fibers were positioned to avoid contact with the injection tract and allowed to settle prior to fixation.

Fiber implants were secured to the skull using a thin layer of cyanoacrylate adhesive (Loctite) applied at low temperature to prevent premature curing and carefully distributed around the fiber–skull interface while avoiding the fiber tip. Adhesive was rapidly polymerized using an accelerator (ZipKicker), after which the implant holder was gently removed. A stainless steel headbar was positioned posterior to the fiber implant and affixed using the same cyanoacrylate-based procedure.

To reinforce implant stability, a dental cement (Metabond; powder:base:catalyst at a 1:3:1 ratio) was mixed immediately prior to use and applied in three successive layers around the base of the fiber implant, headbar, and exposed skull. Each layer was allowed to partially cure before application of the next, forming a rigid, sealed headcap. Implant integrity and alignment were verified before recovery.

Postoperatively, mice received analgesia and were allowed to recover for a minimum of 2–3 weeks prior to behavioral testing and photometry recordings to ensure stable viral expression and full recovery.

**Handling and Water Restriction**

Prior to 2-arm bandit behavioral training, mice were gradually water restricted and handled daily to acclimate them to the experimenter and behavioral procedures. Water access was limited to approximately 1–3 mL per mouse per day, and body weight was monitored daily to maintain animals at no less than 80% of their baseline weight. Mice were handled for 5 consecutive days before the start of behavioral experiments. Tail handling was avoided, and mice were instead transferred using gentle cupping or tunnel handling. Water restriction and handling procedures were maintained throughout behavioral training and testing.

**Food restriction**

For FR/PR experiments requiring food-motivated responding, mice were gradually food restricted and handled daily to acclimate them to experimental procedures. Food access was limited to approximately 1–3 g of standard chow per mouse per day, and body weight was monitored daily to maintain animals at no less than 80% of their baseline weight. The precise amount of food provided was adjusted as needed to maintain stable body weight within this range. For experiments involving fluorescence lifetime photometry and related physiological measurements, food restriction protocols were modified as necessary to ensure stable physiological conditions and consistent lifetime signal measurements while maintaining animal health and weight above the minimum threshold. Food restriction was maintained throughout behavioral testing.

**Viruses**

Recombinant adeno-associated viruses (AAVs) were used to deliver the sensors that allowed for monitoring of extracellular dopamine dynamics and downstream intracellular signaling. Viral vectors were delivered either in a Cre-independent manner or using cell-type–restricted enhancer or Cre-dependent FLEX configurations as indicated. AAVs were obtained from commercial vector cores or in-house sources. All viral aliquots were stored at −80 °C upon arrival and thawed on ice immediately prior to use.

Extracellular dopamine dynamics were measured using genetically encoded dopamine sensors, including red-shifted dLight variants (AAV-hSyn-dLight[r]) and GRAB-DA3.0 (AAV PHP.eb-hSyn-DA3.0). For dLight experiments targeting the nucleus accumbens (NAc), viruses were used at an initial titer of ~1.1 × 10^13 genomic copies (gc)/mL and diluted to a working

concentration of ~7 × 10^12 gc/mL. Typical injection volumes were 300 nL per site, infused at 50 nL/min. For GRAB-DA3.0 experiments, working titers were ~1.3 × 10^13 gc/mL with injection volumes of 250–300 nL per site.

Intracellular PKA signaling was measured using fluorescence lifetime–based reporters (FLIM-AKAR). Both enhancer-driven and Cre-dependent FLEX FLIM-AKAR constructs were employed to achieve cell-type specificity within the striatum. Enhancer-based FLIM-AKAR variants selective for D1- or D2-expressing striatal projection neurons were injected at titers ranging from ~5.0 × 10^12 to 9.2 × 10^13 gc/mL, with injection volumes of 150–300 nL per site. In D1-Cre mice, FLEX FLIM-AKAR was additionally used to selectively target D1 receptor–expressing neurons. In several experiments, dLight and FLIM-AKAR viruses were co-injected to enable simultaneous measurement of dopamine release and downstream intracellular signaling.

Across all experiments, injection volumes ranged from 150–300 nL per site and infusion rates were carefully controlled to minimize tissue damage and backflow. The FLEX FLIM-AKAR reporter used in this study is available through Addgene (AAV-FLEX-FLIM-AKART391A, #60446). Dopamine sensor constructs (dLight variants and GRAB-DA3.0) were obtained through the Sabatini laboratory or UNC Neurotools and are available upon request.

**Pharmacology**

The D1 dopamine receptor antagonist SCH23390 (Tocris Bioscience, #0925) was dissolved in sterile 0.9% saline and administered via intraperitoneal injection at a dose of 10 mg/kg (prepared at 1 mg/mL). Injections were performed using a fine-gauge syringe (BD 324911, BD Biosciences) and were typically completed within 5–20 s of the onset of scruffing to minimize

handling-related variability. Photometry recordings were conducted following drug administration as specified in the corresponding experimental protocols.

**FLIPR Lifetime Photometry**

Fluorescence lifetime photometry was performed using a fluorescence lifetime photometry at high temporal resolution (FLIPR) system set up in the Sabatini laboratory based on published designs (Lodder, 2025). Optical excitation light from a pulsed 473 nm laser (50 MHz repetition rate) was delivered through an optical patch cord to an implanted multimode optical fiber stub. Emitted fluorescence was collected through the same fiber, spectrally separated from excitation light, and detected by a photomultiplier tube (PMT). Fluorescence lifetime was computed in real time using frequency-domain analog phase comparison between the PMT emission signal and the laser reference (Lodder, 2025), yielding continuous measurements of fluorescence lifetime in nanoseconds (ns) alongside fluorescence intensity.

The analog processing unit low-pass filtered the laser reference and PMT signals to isolate the 50 MHz components and computed phase-dependent intermediate-frequency voltages using four phase-shifted reference channels (0, 0.5π, π, 1.5π). These intermediate-frequency signals and the DC intensity component were digitized using a data acquisition (DAQ) device and converted to lifetime and phasor components using custom MATLAB acquisition software (Lodder, 2025).

*Calibration*

Calibration of the FLIPR system was performed using the fluorescence lifetime standard Coumarin 6 dissolved in ethanol, which has a known fluorescence lifetime of approximately 2.40 ns. Calibration measurements were acquired through optical fibers of identical configuration to those used in vivo. Phase alignment, voltage offset correction, and system response calibration

were performed as described previously (Lodder, 2025). Calibration was performed prior to experimental recordings to ensure accurate and stable lifetime measurements across sessions.

*Quantification*

Tonic baseline dopamine was quantified as the modal fluorescence lifetime value within each session. Phasic dopamine events were quantified by detecting local maxima in the fluorescence lifetime trace using MATLAB findpeaks with a minimum peak prominence threshold of 0.075 ns and extracting the resulting peak times and peak values. Peak detection outputs were visually inspected to confirm that detected peaks corresponded to transient lifetime excursions rather than noise. Relative phasic magnitude was computed as peak-minus-baseline and peak-to-baseline ratio.

**Behavioral Setup**

*Fixed Ratio/ Progressive Ratio Sugar Pellet Task*

In fixed ratio (FR) sessions, mice were required to perform a fixed number of responses to obtain a reward. Task parameters were adjusted to probe effort-independent dopamine signaling and establish baseline response–outcome contingencies.

In progressive ratio (PR) sessions, the number of responses required for reward delivery increased systematically across trials. Breakpoints were defined as the final completed ratio prior to cessation of responding and were used as a behavioral index of motivational state and effort valuation.

Mice were trained on operant tasks using a staged protocol designed to progressively shape instrumental responding and assess effort-based motivation. Prior to training, animals underwent mild food restriction to motivate task engagement while maintaining health, and were maintained at approximately 80% of free-feeding body weight with daily monitoring. Training proceeded sequentially through free-feeding exposure, fixed-ratio schedules, and finally a progressive ratio schedule, with advancement contingent on behavioral criterion rather than elapsed time.

Training began with free-feeding exposure to familiarize animals with the reward pellets and feeding device. Mice were then trained on a fixed-ratio 1 (FR1) schedule, in which each response resulted in delivery of a single sugar pellet. Upon meeting criterion performance (earning 30 pellets within a session), animals were advanced to a fixed-ratio 3 (FR3) schedule, in which three responses were required per reward, again with a criterion of 30 pellets per session.

Daily training sessions during free-feeding, FR1, and FR3 stages were standardized to 3 h in duration. Once animals met FR3 criterion, they were transitioned to a progressive ratio (PR) schedule to assess effort-based motivation. In PR sessions, the response requirement increased systematically after each earned reward, and sessions were not capped by pellet number. PR sessions ran for 6 h in duration.

All operant training and testing was performed using the FED3 (Feeding Experimentation Device version 3), which automatically delivered 20 mg sugar pellets contingent on behavioral responses and recorded trial-by-trial behavioral data. Training schedules were staggered to allow the introduction of new animals while maintaining consistent task structure for previously trained cohorts.

*2-Arm Bandit Lick Task*

The two-armed bandit task (2-ABT) was implemented using the same behavioral apparatus and task logic as previously described studies, with minor adaptations for aging and photometry experiments. Briefly, mice were trained to perform a center-initiated choice between two side ports whose reward probabilities reversed in an uncued manner across blocks. The task was designed to probe reinforcement learning, behavioral flexibility, and reward history integration.

Behavior was conducted in custom-built operant chambers (4.9″ × 6″ acrylic) equipped with three nose ports (center, left, right), each monitored by infrared beam-break sensors. Rewards consisted of liquid sucrose solution delivered via solenoid-controlled valves. Behavioral control and data acquisition were implemented using custom software, and all training and testing occurred during the dark phase of the light cycle under no ambient illumination.

*Task Acquisition and Shaping:* To facilitate acquisition of the task, mice progressed through a series of structured training stages in which the center-to-side action sequence and decision timing were gradually shaped (Supplementary Table 1). Across stages, key task parameters, including the allowable decision window between center and side pokes, reward probabilities associated with each side port, and the length of blocks over which contingencies remained constant, were adjusted.

In Stage 1, mice were rewarded for pokes into either side port to establish the basic action–reward contingency. In Stage 2, trials required initiation via a center nose poke followed by a side port selection, establishing the center-to-side action sequence. This association was progressively strengthened across Stages 2–5 by reducing the decision window (the maximum time permitted between the center poke and side poke), thereby encouraging rapid, deliberate

choices. In Stage 6, block lengths were shortened to promote sensitivity to contingency changes. Advancement through Stages 1–5 required acquisition of 100 rewards within a single session, while advancement from Stage 6 required acquistion of 200 rewards.

*Probabilistic Execution Phase:* Upon successful completion of the shaping stages, mice advanced to Stage 7, which constituted the execution phase of the two-armed bandit task. In this final stage, one side port was designated as the high-reward option ($p = 0.8$) and the other as the low-reward option ($p = 0.2$). Reward contingencies reversed unpredictably in uncued blocks of variable length (10–30 trials), requiring animals to continuously update action values based on recent reward history.

For initial task execution experiments, mice completed 10 sessions of the probabilistic task at the 80/20 contingency. For retraining experiments, mice completed 15 execution sessions at the same 80/20 contingency without repeating earlier shaping stages. These execution-only sessions were used to assess stable performance, learning dynamics, and age- or manipulation-dependent differences in reinforcement learning.

All behavioral sessions were conducted during the dark phase of the sleep–wake cycle. Only completed trials, defined as trials in which a center poke was followed by a valid side poke, were included in behavioral analyses.

*Head-Fixed*

For head-fixed recordings, stainless steel head bars were secured to a custom-designed adjustable rigid holder to stabilize the skull while allowing the body to remain supported within a custom cylindrical platform.

Mice performed the two-arm bandit task using two lick ports positioned within reach of the animal. Licks were detected using capacitive sensors, and rewards consisted of water delivery contingent on probabilistic reward schedules identical to those used in the freely moving configuration. Behavioral events and fluorescence intensity signals were recorded simultaneously.

**Recursive Formulation of Logisic Regression (RFLR) Model**

2-ABT behavior was modeled using the recursive formulation of logistic regression (RFLR), as described in Beron et al., 2022. This model provides a compact and biologically plausible description of probabilistic decision-making by integrating recent choice–outcome history with an explicit bias toward repeating previous actions.

On each trial t, the model computes a latent decision variable $\psi_{t+1}$ representing the log-odds of selecting one port over the alternative:

$$\psi_{t+1} = \alpha\, c_t + \varphi_t$$

where $c_t$ represents the animal's choice on trial t, encoded as +1 for rightward choices and −1 for leftward choices. The parameter α captures a choice history bias, also referred to as stickiness, which reflects a tendency to repeat the previous action independent of outcome. The accumulated evidence term $\varphi_t$ is updated recursively according to:

$$\varphi_t = \beta\,(c_t\, r_t) + e^{(-1/\tau)}\, \varphi_{t-1}$$

where $r_t$ denotes trial outcome, encoded as +1 for rewarded trials and 0 for unrewarded trials. The interaction term $c_t\, r_t$ represents signed evidence favoring or opposing the chosen action. The parameter β determines the strength of evidence integration, and τ controls the timescale over which past evidence decays.

The probability of selecting the right port on the subsequent trial is given by the logistic function:

$$P(\text{choice}_{t+1} = \text{right}) = 1 / (1 + \exp(-\psi_{t+1}))$$

Model parameters α, β, and τ were estimated by maximizing the likelihood of observed trial-by-trial choices. This recursive formulation enables efficient computation by maintaining a single internal state variable that represents exponentially weighted choice–outcome history.

**scRNA-seq Analysis**

*Data acquisition and preprocessing.* Single-nucleus RNA sequencing data were obtained from two publicly available datasets: GSE225741 (2-month-old mice, n = 2 biological replicates) and GSE226138 (18-month-old mice, n = 4 biological replicates), both generated from whole striatum using the 10x Genomics Chromium platform. Raw count matrices were loaded using Scanpy (v1.11.1) via sc.read_10x_mtx. All samples were concatenated into a single AnnData object with sample and age group labels retained.

*Quality control.* Mitochondrial gene content was calculated for each nucleus (mouse mitochondrial genes identified by Mt- prefix). Nuclei were retained if they expressed between 600 and 3,500 genes and had fewer than 5% mitochondrial reads. Genes expressed in fewer than 3 nuclei were excluded. These thresholds were matched to those reported in the original data publications.

*Doublet detection.* Putative doublets were identified independently per sample using Scrublet (expected doublet rate = 6%), with doublet scoring performed on 20 principal components. Predicted doublets were removed prior to downstream analysis.

*Normalization and feature selection.* Raw counts were normalized to 10,000 counts per nucleus and log-transformed (log1p). Top 6,000 highly variable genes were selected using the Seurat v3 variance-stabilizing transformation with sample as a batch covariate (sc.pp.highly_variable_genes, flavor='seurat_v3', batch_key='sample'). Expression values were scaled to zero mean and unit variance, clipping at a maximum value of 10.

*Dimensionality reduction and batch integration.* Principal component analysis (PCA) was performed on the highly variable genes (30 components). To correct for sample-level batch effects while preserving biological variation, Harmony integration was applied to the PCA embedding using harmonypy (v0.2.0), with a diversity penalty $\theta = 2.0$ and a maximum of 30 iterations (batch_key='sample'). A nearest-neighbor graph was constructed on the Harmony-corrected embedding using 15 neighbors and 10 principal components (sc.pp.neighbors). UMAP coordinates were computed for visualization.

*Clustering and cell type annotation.* Leiden clustering was performed at a resolution of 0.08, yielding 10 clusters. Cell types were assigned manually by inspection of canonical striatal marker gene expression, visualized as dot plots and UMAP feature plots. Marker genes used for annotation included: Drd1, Pdyn (D1-SPN); Drd2, Penk (D2-SPN); Chat, Slc18a3 (cholinergic interneurons); Pvalb, Sst, Npy (GABAergic interneurons); Aldoc, Aqp4, Sox9 (astrocytes); Plp1, Mbp (oligodendrocytes); Pdgfra (OPCs); Cx3cr1, P2ry12 (microglia); Cldn5 (endothelial); Acta2, Rgs5 (mural cells/pericytes). Cluster-level marker genes were also computed using Wilcoxon rank-sum tests (sc.tl.rank_genes_groups) to confirm annotations.

*Differential expression analysis.* Aging-associated differential expression was assessed within each annotated cell type independently. Nuclei from 18-month-old animals were compared to 2-month-old animals using the Wilcoxon rank-sum test (sc.tl.rank_genes_groups,

method='wilcoxon'), applied to log-normalized counts. Cell types with fewer than 10 nuclei per age group were excluded. P-values were corrected for multiple comparisons using the Benjamini-Hochberg procedure. Genes were considered differentially expressed at adjusted $p < 0.05$ and $|\log_2 FC| > 0.25$ (summary analyses) or $|\log_2 FC| > 0.5$ (GO enrichment analyses).

*Statistics.* Comparative analyses between age group within clusters were performed using the Mann-Whitney U test (two-sided). Statistical significance was assessed at $p < 0.05$.

*Software.* All analyses were performed in Python 3.11 using Scanpy (v1.11.1), AnnData (v0.11.4), harmonypy (v0.2.0), Scrublet, gseapy, NumPy, pandas, SciPy, and Matplotlib.

*Reanalysis of Kilfeather et al. (2024) TRAP sequencing data*

Translating Ribosome Affinity Purification (TRAP) sequencing data from Kilfeather et al. (2024) were reanalyzed from publicly deposited raw count matrices from young and aged male mice across ventral striatum and midbrain/VTA. Raw count matrices were loaded per sample and merged with associated metadata. Library-size normalization was performed by converting counts to counts per million (CPM) followed by $\log_2$ transformation ($\log_2(\text{CPM} + 1)$). Quality control metrics including library size and number of detected genes were assessed per sample prior to analysis. Principal component analysis was performed on the top 2,000 most variable genes to assess sample clustering by region and age.

Differential expression analysis was performed separately per region. For the midbrain, where $n = 8$ per age group drawn from two cohorts, a per-gene Welch's t-test was applied comparing old versus young log-CPM values, with p-values corrected for multiple comparisons using the Benjamini-Hochberg procedure (adjusted $p < 0.05$, $|\log_2 FC| \geq 1.5$). For the ventral striatum and bulk total fractions where $n = 1$ per condition, statistical testing was not performed. Genes were

ranked by fold change and thresholded at $|\log_2 FC| \geq 1.5$. Gene symbols were mapped from Ensembl IDs using the mygene package. Gene Ontology Biological Process enrichment was performed on significantly differentially expressed gene lists using gseapy (Enrichr interface, GO_Biological_Process_2023, mouse). Results were filtered to adjusted $p < 0.05$ and visualized as bubble plots. All analyses were performed in Python 3.11 using NumPy, pandas, SciPy, statsmodels, and Matplotlib.

**Histology and Imaging**

Following behavioral and photometry experiments, mice were deeply anesthetized with isoflurane and transcardially perfused with phosphate-buffered saline (PBS), followed by fixation with 4% paraformaldehyde (PFA). Brains were extracted and post-fixed overnight in 4% PFA at 4 °C, then transferred to 30% sucrose in PBS for cryoprotection until fully equilibrated. Brains were embedded in optimal cutting temperature compound (O.C.T.; Tissue-Tek), rapidly frozen on dry ice, and sectioned coronally at 50 μm using a cryostat. Free-floating sections were mounted using a DAPI-containing antifade mounting medium (Vectashield, H-1200) to preserve fluorescence. Fluorescent reporter expression and optical fiber placements were verified using widefield fluorescence microscopy with a 10× objective. Native fluorescence from viral reporters was sufficient for visualization, and no immunohistochemical amplification was performed. Animals with incorrect viral targeting or fiber placement were excluded from further analysis.

**Ethics**

All experimental procedures were conducted in accordance with the guidelines set by the Institutional Animal Care and Use Committee (IACUC).

# III. Results

## Acquisition of the two-arm bandit task across the lifespan

Water-restricted mice across six age groups were trained on a two-armed bandit task in which trials began with initiation at a center port followed by selection of one of two side ports, resulting in delivery or omission of water reward (Fig. 1a; see Materials & Methods). Mice first learned under deterministic conditions in which correct choices reliably produced reward, enabling acquisition of the basic action-outcome contingency. Task difficulty increased across subsequent stages as the decision window progressively shortened and reward contingencies organised into blocks that reverse in an uncued manner. Advancement between stages required predefined performance criteria, ensuring that progression reflected task acquisition rather than passive exposure.

To determine whether aging alters acquisition of the task, we examined performance during the deterministic training stages across age groups. We quantified acquisition rate, reward outcomes, switching behavior, and trial execution dynamics.

Acquisition speed declined with age, with older mice requiring progressively more sessions to reach criterion (Fig. 1b). The effect was largest in the oldest age group, with 18-24 month mice requiring on average approximately 9 more sessions than 0-2 month mice to reach the final probabilistic stage.

To examine how reinforcement outcomes differed during acquisition, behavior was aligned by normalised learning progress and divided into early, middle, and late phases. During early

acquisition, older mice accumulated disproportionately more unrewarded outcomes and had markedly lower reward rates than younger animals, with the 18-24 month group showing the largest deficit (Figs. 1c, d). These differences were absent by the late phase, consistent with eventual acquisition of task structure by all age groups.

The increased frequency of unrewarded outcomes in aged mice could reflect differences in exploration strategy, such as increased action switching. However, switch probability systematically declined with acquisition and did not differ significantly across age groups at any learning phase (Fig. 1e), indicating that excess unrewarded sampling was not explained by increased exploration. Trial duration similarly showed no systematic increase with age after early adulthood, with older age groups not differing consistently from one another despite their elevated error rates (Fig. 1f), indicating that aged mice were adequately engaged and capable of task execution.

Together, these findings demonstrate that aging selectively impairs the efficiency of reinforcement-guided learning, with aged mice requiring substantially more experience to establish the action-outcome contingency despite preserved exploration and task engagement.

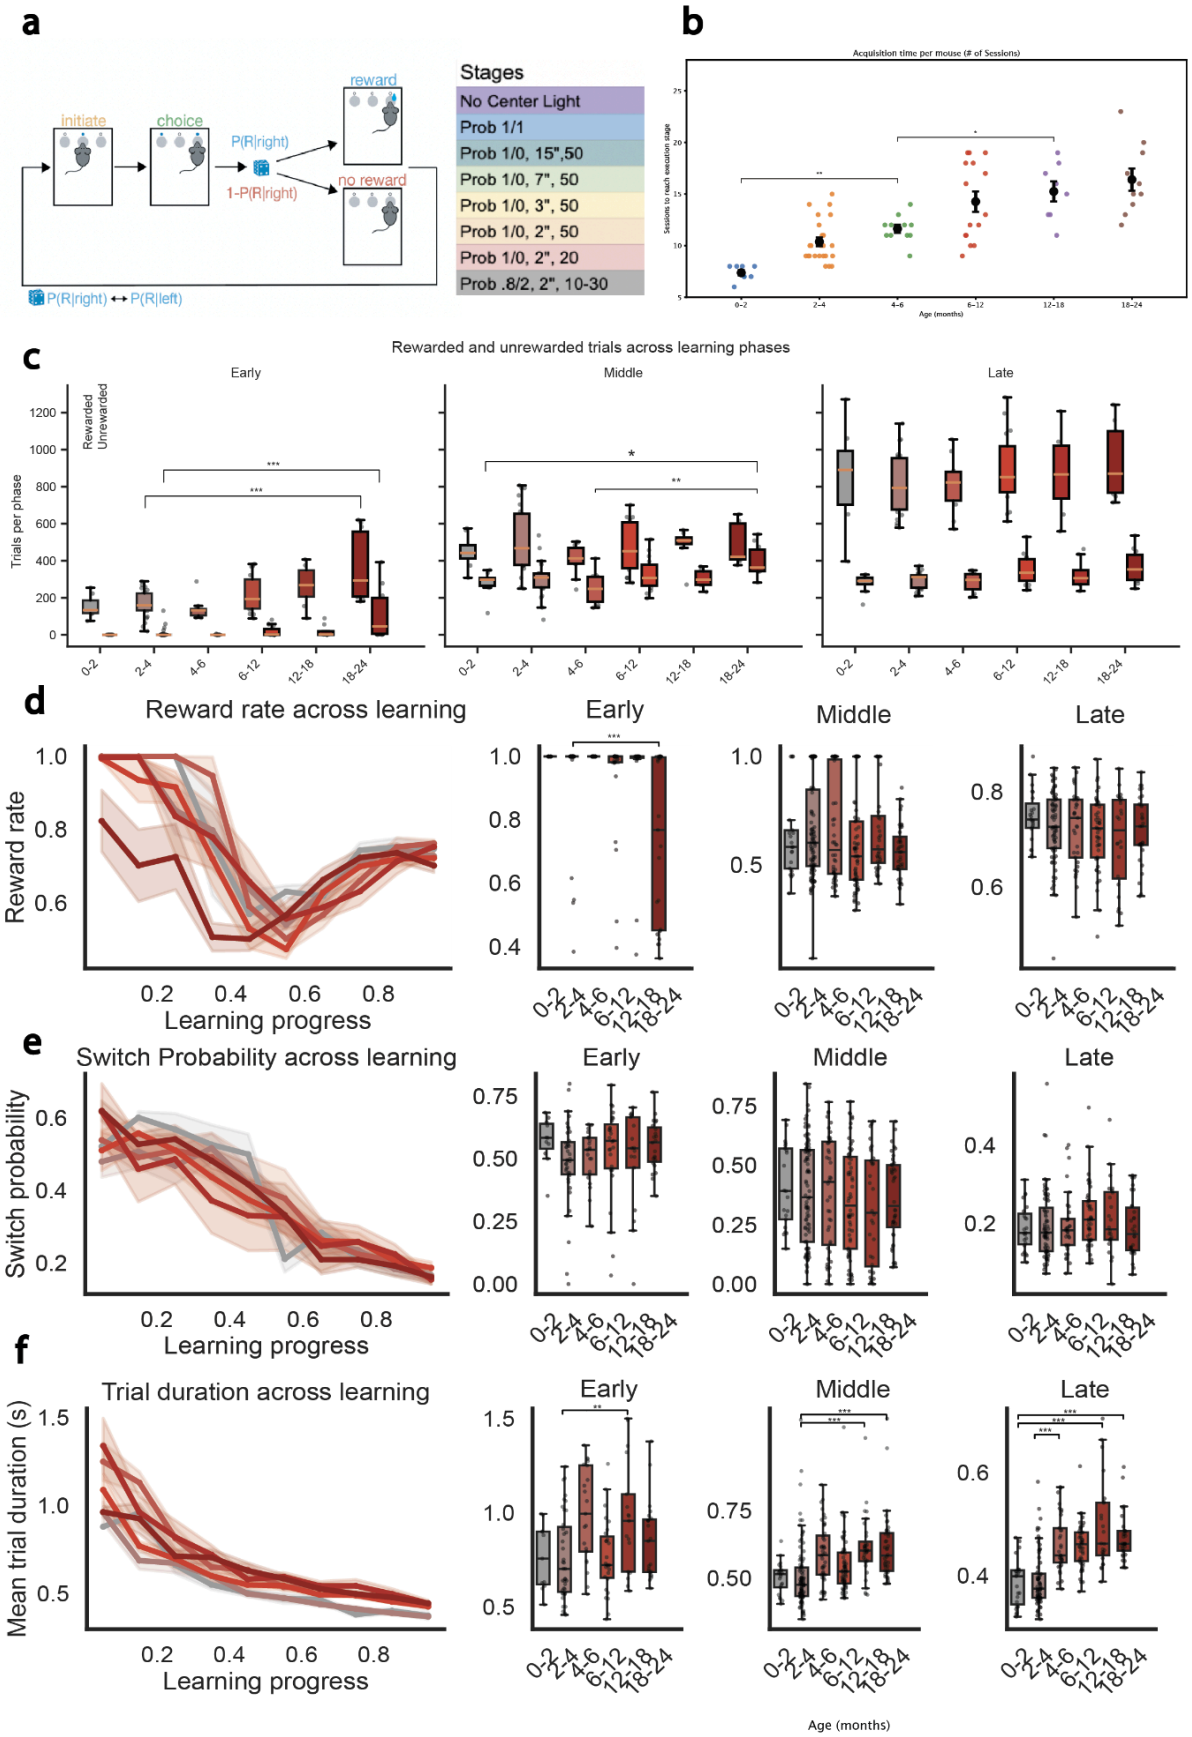
a
initiate
choice
P(R|right)
reward
no reward
1-P(R|right)
P(R|right) ↔ P(R|left)
Stages
No Center Light
Prob 1/1
Prob 1/0, 15",50
Prob 1/0, 7", 50
Prob 1/0, 3", 50
Prob 1/0, 2", 50
Prob 1/0, 2", 20
Prob .8/2, 2", 10-30
b
Acquisition time per mouse (# of Sessions)
Sessions to reach execution stage
Age (months)
c
Rewarded and unrewarded trials across learning phases
Early
Middle
Late
Rewarded
Unrewarded
Trials per phase
d
Reward rate across learning
Reward rate
Learning progress
e
Switch Probability across learning
Switch probability
f
Trial duration across learning
Mean trial duration (s)
Age (months)

**Figure 1. Aging impairs the efficiency of reinforcement-guided task acquisition without disrupting exploration or task engagement.**

*n = 8, 27, 9, 17, 8, 11 mice for age groups 0–2, 2–4, 4–6, 6–12, 12–18, and 18–24 months, respectively. *p < 0.05, **p < 0.01, ***p < 0.001, Tukey HSD post-hoc following one-way ANOVA unless otherwise stated.*

**a, Schematic of the two-armed bandit task.** Mice initiated trials at a center port and selected one of two side ports, resulting in delivery or omission of water reward. Reward contingencies reversed periodically in an uncued manner. Training progressed through stages of increasing temporal and probabilistic difficulty, culminating in an 80/20 probabilistic contingency with block lengths of 10–30 trials.
**b, Sessions required to reach the final probabilistic stage**, plotted per mouse (colored dots) as mean ± SEM. One-way ANOVA: $F(5,74) = 17.06$, $p < 0.001$. Selected pairwise comparisons shown (***$p < 0.001$).
**c, Rewarded (left) and unrewarded (right) trial counts across learning phases (early: normalized progress < 0.33; middle: 0.33–0.66; late: > 0.66)**, plotted as box plots (median, IQR, whiskers = 1.5× IQR) with individual animals overlaid. Trial counts are mouse-balanced across phases. Unrewarded trials, early phase: $F(5,70) = 6.37$, $p < 0.001$. Rewarded trials, early phase: $F(5,70) = 7.80$, $p < 0.001$. Unrewarded trials, middle phase: $F(5,70) = 3.62$, $p = 0.006$. Rewarded trials, middle and late phases and unrewarded late phase: not significant or no individual comparisons reached significance. Selected pairwise comparisons shown (***$p < 0.001$).
**d, Reward rate across learning.** Left, mean ± SEM per age group plotted against normalized learning progress. Right, reward rate at early, middle, and late phases plotted as in c. Early phase: $F(5,132) = 7.48$, $p < 0.001$. Middle and late phases: not significant ($p = 0.110$ and $p = 0.330$, respectively). Selected pairwise comparisons shown (***$p < 0.001$).
**e, Switch probability across learning phases**, plotted as in d. One-way ANOVA at each phase: early $F(5,132) = 0.97$, $p = 0.439$; middle $F(5,284) = 0.75$, $p = 0.585$; late $F(5,221) = 0.86$, $p = 0.510$. No significant differences detected at any phase.
**f, Trial duration across learning phases**, plotted as in d. Early phase: $F(5,132) = 4.74$, $p < 0.001$. Middle phase: $F(5,284) = 9.46$, $p < 0.001$. Late phase: $F(5,221) = 24.95$, $p < 0.001$. Significant differences at early and middle phases are driven primarily by the 0–2 and 2–4 month groups; older age groups do not differ consistently from one another. Selected pairwise comparisons shown (*$p < 0.05$, ***$p < 0.001$).

## Behavior during probabilistic execution of the two-armed bandit task across the lifespan

Having established that mice across the lifespan acquired the task structure and remained engaged, we next examined behavior under probabilistic reward contingencies (80/20), which require continuous reinforcement-guided updating.

Unrewarded trials increased with age, with the largest difference observed between 18-24 month mice and younger adults (Fig. 2a). Rewarded trial counts did not increase in parallel, indicating

that older mice accumulated more errors without compensating through increased correct selections.

Trial duration did not track error rates across age groups. Duration increased between 2-4 and 4-6 months but remained stable across older age groups (Fig. 2b), indicating that the elevated error rates in aged mice reflect altered action selection rather than slowed execution or reduced task engagement.

Consistent with this, high-port selection probability was reduced in the oldest age group, with 18-24 month mice showing lower selection of the high-probability port and greater inter-animal variability than younger adults (Fig. 2c).

To determine whether age-related performance deficits reflected impaired detection of contingency change or impaired post-reversal adaptation, behavior was aligned to block transitions (Fig. 2d). All age groups exhibited an immediate reduction in high-port selection following reversal, indicating intact detection of contingency change. Pre-reversal baseline high-port selection showed a numerical decline across age groups, from approximately 0.77 in younger mice to 0.66 in the oldest group, and trials required to recover to baseline showed a corresponding numerical increase from approximately 7.6 to 12.7 in the oldest group. However, neither difference reached significance ($F(5,69) = 2.14$, $p = 0.071$ and $F(5,68) = 0.70$, $p = 0.623$, respectively). Post-reversal switching probability similarly showed a numerical reduction in 18–24 month mice without reaching significance (Fig. 2e; $F(5,69) = 1.20$, $p = 0.319$).

To isolate outcome-dependent updating, we quantified lose-switch and win-stay probabilities (Figs. 2f, 2g). Lose-switch probability showed a numerical decline with age, from a median of approximately 0.25 in younger groups to approximately 0.20 in the oldest group, while win-stay

probability remained high across all groups at approximately 0.95 to 0.975. Neither difference reached significance (lose-switch: $F(5,69) = 1.58$, $p = 0.177$; win-stay: $F(5,69) = 0.86$, $p = 0.512$), though the directional dissociation suggests a selective impairment in updating following negative outcomes rather than a global reduction in choice consistency.

These results indicate that aging reduces choice optimality under probabilistic reward contingencies, with older mice accumulating more unrewarded outcomes and showing reduced capacity to stably select the high-probability port. Directional trends toward impaired post-reversal recovery and reduced lose-switch probability are consistent with a deficit in reinforcement-driven updating, though these did not reach significance and warrant further investigation with larger cohorts.

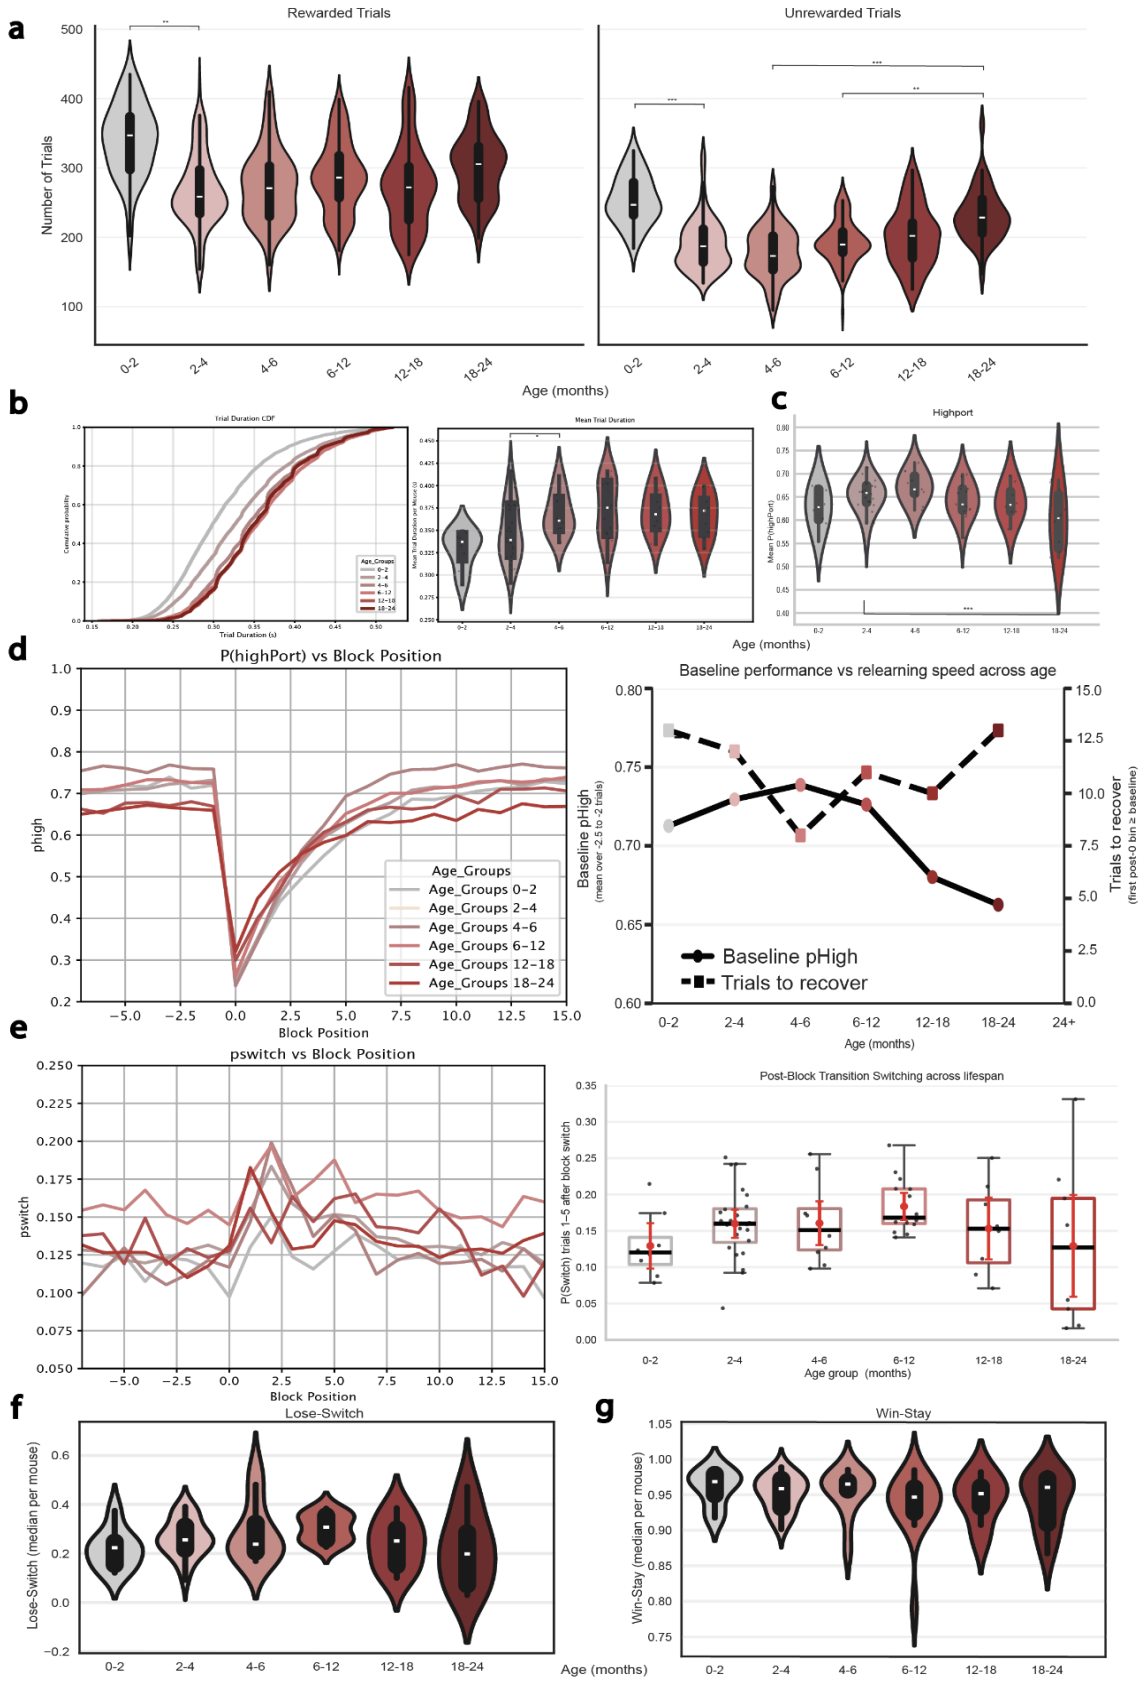
a
Rewarded Trials
Unrewarded Trials
Number of Trials
Age (months)
0-2
2-4
4-6
6-12
12-18
18-24
b
Trial Duration CDF
Mean Trial Duration
c
Highport
Age (months)
d
P(highPort) vs Block Position
phigh
Block Position
Age_Groups
Age_Groups 0-2
Age_Groups 2-4
Age_Groups 4-6
Age_Groups 6-12
Age_Groups 12-18
Age_Groups 18-24
Baseline performance vs relearning speed across age
Baseline pHigh
(mean over -2.5 to -2 trials)
Trials to recover
(first post-0 bin ≥ baseline)
Baseline pHigh
Trials to recover
Age (months)
e
pswitch vs Block Position
pswitch
Block Position
Post-Block Transition Switching across lifespan
P(Switch) trials 1–5 after block switch
Age group (months)
f
Lose-Switch
Lose-Switch (median per mouse)
Age (months)
g
Win-Stay
Win-Stay (median per mouse)

**Figure 2. Aging impairs reinforcement-dependent updating during probabilistic decision-making.**

*n = 8, 24, 11, 15, 8, 9 mice for age groups 0–2, 2–4, 4–6, 6–12, 12–18, and 18–24 months, respectively. *p < 0.05, **p < 0.01, ***p < 0.001, Tukey HSD post-hoc following one-way ANOVA unless otherwise stated.*

**a, Rewarded (left) and unrewarded (right) trial counts per session across age groups**, plotted as violin plots with embedded box plots (median, IQR, whiskers = 1.5× IQR). Rewarded trials: $F(5,69) = 3.76$, $p = 0.005$; significant difference driven by the 0–2 month group. Unrewarded trials: $F(5,69) = 12.01$, $p < 0.001$. Selected pairwise comparisons shown (**p < 0.01, ***p < 0.001).

**b, Trial duration across age groups. Left**, cumulative distribution functions per age group. **Right**, mean trial duration per mouse plotted as in a. One-way ANOVA: $F(5,69) = 3.99$, $p = 0.003$; differences driven by the 0–2 month group being slower than all others, with no consistent differences among groups aged 4 months and older. Selected pairwise comparisons shown (**p < 0.01, ***p < 0.001).

**c, High-port selection probability per mouse across age groups**, plotted as in a. One-way ANOVA: $F(6,69) = 3.99$, $p = 0.003$. Selected pairwise comparisons shown (**p < 0.01).

**d, High-port selection aligned to contingency reversal. Left**, mean high-port selection probability per age group aligned to reversal (block position 0). **Right**, pre-reversal baseline high-port selection (solid line) and trials required to recover to baseline (dashed line) per age group, plotted as in a. Baseline high-port selection: $F(5,69) = 2.14$, $p = 0.071$. Post-reversal recovery trials: $F(5,68) = 0.70$, $p = 0.623$. No significant pairwise differences detected.

**e, Switch probability aligned to contingency reversal. Left**, mean switch probability per age group aligned to reversal. **Right**, switch probability in trials 1–5 following reversal per mouse, plotted as in a. One-way ANOVA: $F(5,69) = 1.20$, $p = 0.319$. No significant differences detected.

**f, Lose-switch probability per mouse across age groups,** plotted as in a. One-way ANOVA: $F(5,69) = 1.58$, $p = 0.177$. No significant differences detected.

**g, Win-stay probability per mouse across age groups**, plotted as in a. One-way ANOVA: $F(5,69) = 0.86$, $p = 0.512$. No significant differences detected.

## Interpretable modelling of probabilistic updating with a recursively formulated logistic regression (RFLR) model

To identify the computational processes underlying age-dependent changes in reinforcement-driven behavior, we fitted trial-by-trial choices using a compact reinforcement learning framework previously developed for this task and shown to accurately capture both port selection and switching dynamics (RFLR; Beron et al., 2022; Fig. 3a). In this model, choice probability is generated by applying a stochastic policy to a latent decision variable that is updated recursively based on recent action-outcome history. This formulation provides an

interpretable decomposition of behavior into distinct components governing choice persistence, sensitivity to reinforcement, and temporal integration of evidence.

The model maintains a scalar latent state representing accumulated evidence favoring one port over the other. Following each trial, this state decays toward baseline and is updated according to the most recent choice and its outcome, such that rewarded actions reinforce the selected port, whereas unrewarded outcomes shift the decision variable toward the alternative. This recursive update captures both reinforcement-dependent updating and stochastic exploration observed in behavior.

The model is defined by three parameters: the persistence parameter $\alpha$ captures an intrinsic tendency to repeat the previous choice independent of reinforcement; the reinforcement sensitivity parameter $\beta$ scales the influence of the immediately previous action-outcome history on subsequent choice; and the integration timescale parameter $\tau$ governs the temporal persistence of accumulated evidence across trials.

The model accurately reproduced key features of behavior across animals, including switching probabilities, block-aligned choice dynamics, and outcome-conditioned transitions (Fig. 3b). Model fit quality remained consistent across age groups, with log likelihood per trial ranging from approximately -0.30 to -0.36, indicating that the model provided a comparably good description of behavior across the lifespan (Fig. 3f).

We next examined age-dependent changes in model parameters. The persistence parameter $\alpha$ remained stable across age groups (Fig. 3c), indicating preserved baseline choice persistence.

In contrast, β declined significantly with age, with the most pronounced reduction observed in mice aged 18-24 months relative to the 4-6 month group (Fig. 3d). This reduction reflects diminished influence of recent reward evidence on subsequent choices.

The integration timescale τ increased significantly with age, with 18-24 month mice showing substantially longer integration timescales than younger adults (Fig. 3e). This prolonged persistence of prior evidence would be expected to slow the incorporation of new reinforcement signals, compounding the effect of reduced β on behavioral updating.

Together, these parameter changes define a distinct computational signature of aging characterized by reduced reinforcement sensitivity and increased persistence of prior evidence. Reduced β limits the magnitude of reinforcement-driven updating, whereas increased τ prolongs the influence of previously accumulated evidence. This combination could reasonably produce slower behavioral updating and increased stability of prior choices, in line with the observed reduction in lose-switch probability, diminished steady-state high-port selection, and the trend toward delayed adaptation following contingency reversals.

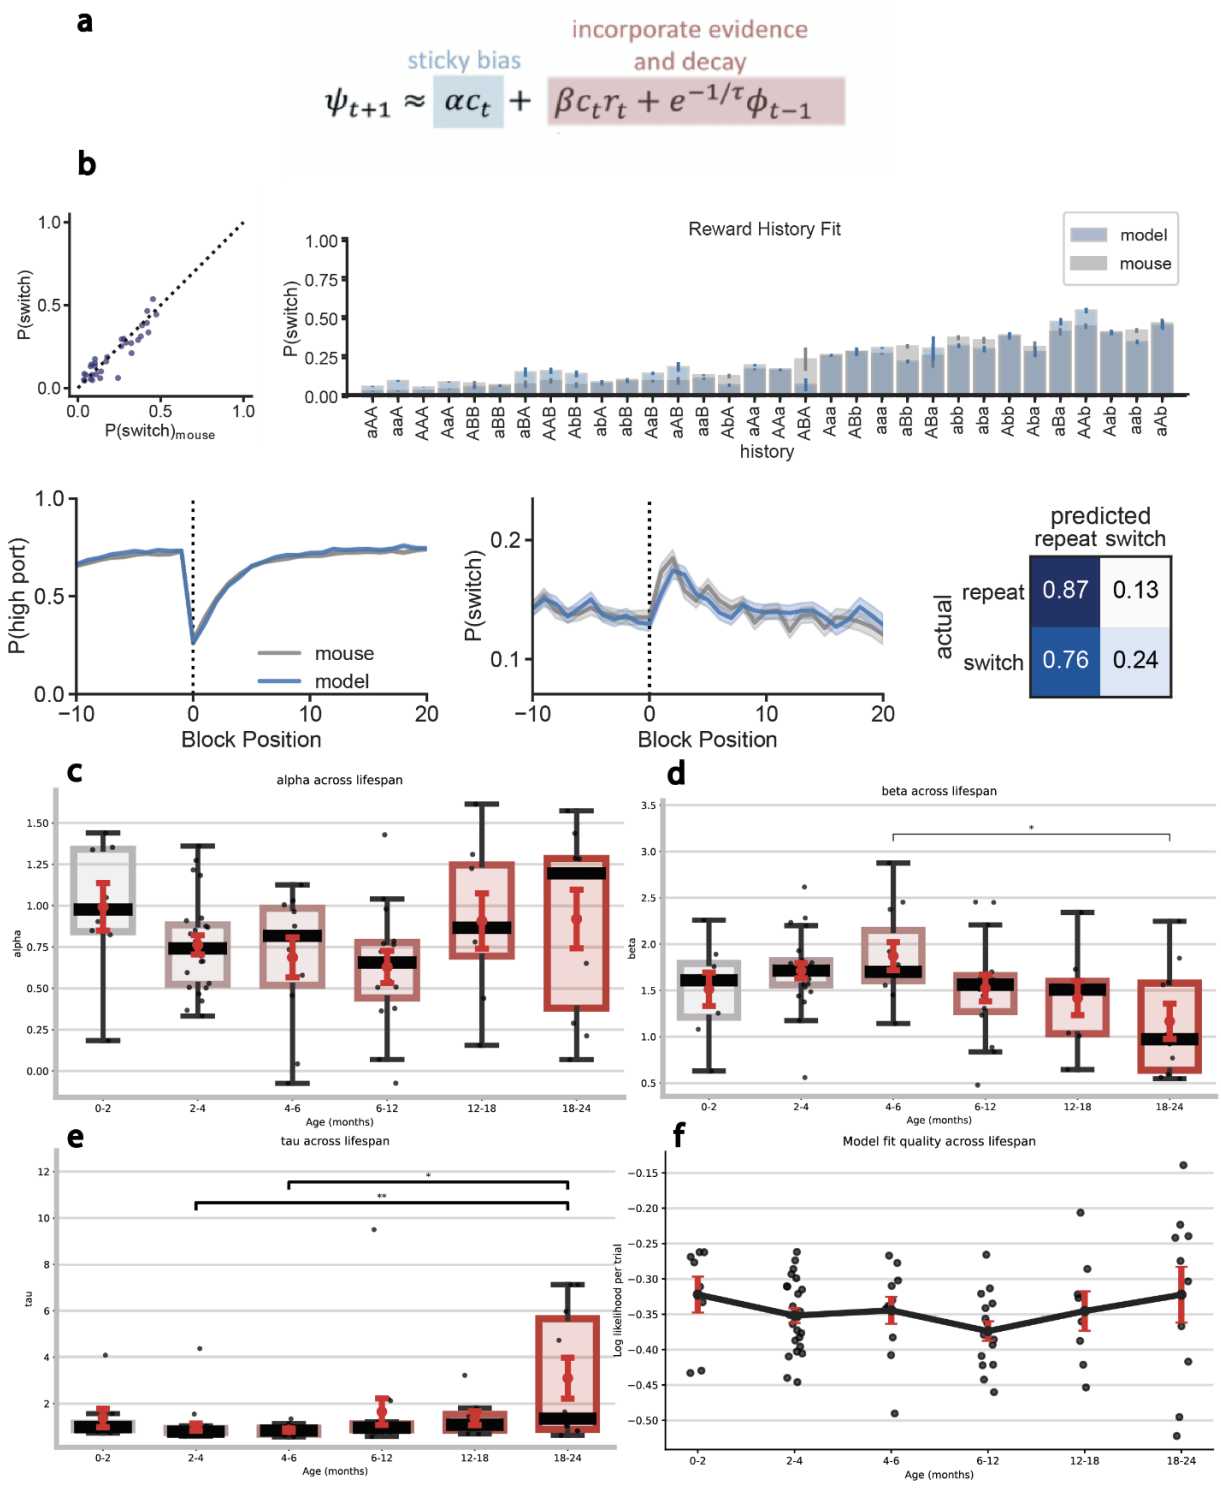


**Figure 3. Model reveals age-dependent reduction in reinforcement sensitivity and prolonged persistence of prior evidence.**

*n = 8, 24, 11, 15, 8, 10 mice for age groups 0–2, 2–4, 4–6, 6–12, 12–18, and 18–24 months, respectively. *p < 0.05, **p < 0.01, Tukey HSD post-hoc following one-way ANOVA unless otherwise stated.*

**a, RFLR model schematic.** The decision variable ψ is updated recursively on each trial as the sum of a sticky bias term (blue) and a reinforcement and decay term (red), where α governs choice persistence, β scales reinforcement sensitivity, and τ controls the timescale of evidence integration.
**b, Model validation. Top left**, predicted versus observed switch probability per mouse across all data. **Top right**, switch probability conditioned on the three most recent action-outcome histories (model, blue; mouse, grey), plotted as mean ± SEM. **Bottom**, mean high-port selection probability (left) and switch probability (right) aligned to contingency reversal for model and mouse. **Right**, confusion matrix showing predicted versus actual repeat and switch choices across all data.
**c, Persistence parameter (α) per mouse across age groups**, plotted as box plots (median, IQR, whiskers = 1.5× IQR) with individual animals overlaid. One-way ANOVA: $F(5,70) = 1.43$, $p = 0.224$. No significant differences detected.
**d, Reinforcement sensitivity parameter (β) per mouse across age groups**, plotted as in c. One-way ANOVA: $F(5,70) = 2.68$, $p = 0.029$. Selected pairwise comparisons shown (*$p < 0.05$).
**e, Evidence integration timescale (τ) per mouse across age groups**, plotted as in c. One-way ANOVA: $F(5,70) = 3.10$, $p = 0.014$. Selected pairwise comparisons shown (*$p < 0.05$, **$p < 0.01$).
**f, Model fit quality (log likelihood per trial) per mouse across age groups**, plotted as in c, with mean ± SEM overlaid (red) and polynomial fit (blue). One-way ANOVA: $F(5,70) = 0.87$, $p = 0.504$; Spearman $\rho = -0.037$, $p = 0.753$. No significant differences detected.

## Dopamine transients encode reward history in deterministic and probabilistic versions of the two-armed bandit task.

We next sought to examine a candidate neural substrate through which the age-dependent reduction in reinforcement sensitivity might be expressed. Nucleus accumbens (NAc) dopamine is a well-established encoder of reward prediction errors and is thought to provide the teaching signal driving reinforcement-dependent updating of action selection. If dopamine amplitude carries information about recent reward history, age-related changes in this signal could contribute to the diminished sensitivity to outcomes observed in older animals.

To test whether NAc dopamine encodes recent reward history and whether this encoding differs across early adulthood, we recorded dopamine transients using fiber photometry in head-fixed mice performing the task at P60 and P120 (Fig. 4a, b). This comparison spans the developmental

window in which differences in reward outcome accumulation and trial duration first emerge (Fig. 2a, b).

Dopamine transients were tightly time-locked to reward consumption and were systematically scaled by recent reward history (Fig. 4c). Rather than responding uniformly to each reward, transient amplitude decreased progressively across consecutive wins in both age groups, indicating that dopamine encodes the position of a reward within a recent outcome sequence and not merely its occurrence. P120 mice showed larger absolute responses at each sequence position than P60 mice, but the directional attenuation across wins was preserved in both groups.

This history-dependent encoding extended beyond the moment of consumption. At cue onset, dopamine responses increased across consecutive win sequences before plateauing (Fig. 4d), indicating that prior reward history shapes anticipatory dopamine release as well. Following losses, signals were suppressed below baseline in a manner that also reflected recent outcome history, with P60 mice showing deeper initial suppression than P120 mice (Fig. 4e). Dopamine therefore tracks both positive and negative outcome histories in a graded manner, providing a continuously updated reinforcement signal across the trial.

Response magnitudes differed between age groups across both deterministic and probabilistic contingencies. Under the 80/20 contingency, peak responses at the first reward in a sequence were larger in P120 than P60 mice, and the attenuation across subsequent rewards was more pronounced in P60 animals (Fig. 4f). Because dopamine amplitude encodes recent outcome history in a manner consistent with a role in guiding subsequent choices, these differences in signaling across early adulthood provide a plausible neural basis for the age-dependent reduction in reinforcement sensitivity identified by the RFLR model (Fig. 3).

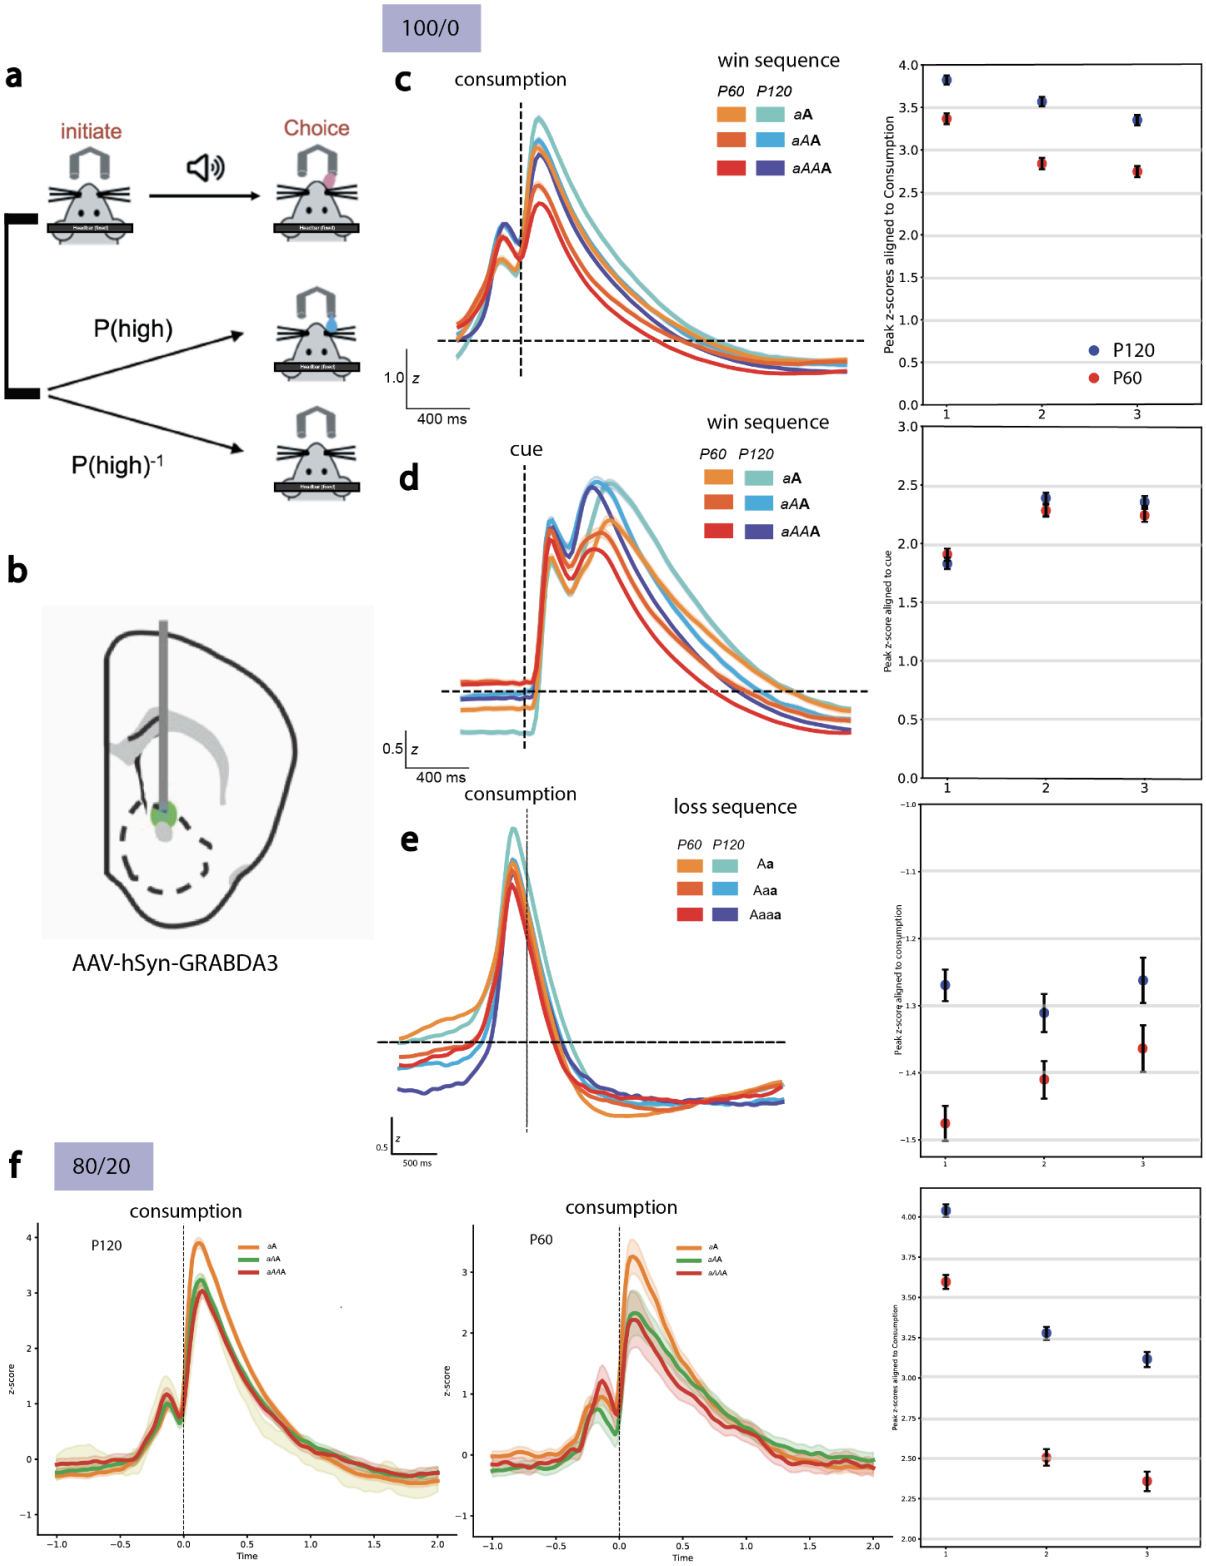
100/0
a
initiate
Choice
P(high)
P(high)-1
b
AAV-hSyn-GRABDA3
c
consumption
win sequence
P60 P120
aA
aAA
aAAA
1.0 z
400 ms
d
cue
win sequence
P60 P120
aA
aAA
aAAA
0.5 z
400 ms
e
consumption
loss sequence
P60 P120
Aa
Aaa
Aaaa
0.5 z
500 ms
Peak z-scores aligned to Consumption
P120
P60
Peak z-score aligned to cue
Peak z-score aligned to consumption
f
80/20
consumption
P120
consumption
P60
z-score
Time

**Figure 4. Dopamine transients encode reward history in deterministic and probabilistic versions of the two-armed bandit task.**

*n = 12, 10 mice for P60 and P120 respectively.*

**a, Head-fixed task schematic.** Head-fixed mice initiated trials by licking the center spout, followed by a tone cue prompting choice of one of two side spouts. Reward was delivered with probability P(high) at the chosen port or withheld with probability 1−P(high).

**b, Viral delivery strategy.** AAV-hSyn-GRABDA3 was injected into the nucleus accumbens for fluorescence-based dopamine imaging.

**c, Reward-evoked dopamine responses across win sequences (100/0 contingency). Left**, z-scored photometry signals aligned to reward consumption for increasing win sequence lengths (aA, aAA, aAAA) in P60 (warm colours) and P120 (cool colours) mice, plotted as mean ± SEM. **Right**, peak z-scores at each sequence position for P60 (red) and P120 (blue), plotted as mean ± SEM. P120 peak responses declined from 3.82 ± 0.05 at the first reward in a sequence to 3.34 ± 0.06 at the third; P60 responses declined from 3.36 ± 0.06 to 2.74 ± 0.07.

**e, Consumption-aligned dopamine responses across loss sequences (100/0 contingency). Left**, z-scored photometry signals aligned to reward consumption for increasing loss sequence lengths (Aa, Aaa, Aaaa) in P60 and P120 mice, plotted as mean ± SEM. **Right**, peak z-scores at each sequence position plotted as c. Peak suppression following the first loss in a sequence was −1.47 ± 0.03 in P60 mice and −1.27 ± 0.02 in P120 mice.

**f, Reward-evoked dopamine responses across win sequences (80/20 contingency).** Z-scored photometry signals aligned to reward consumption for P60 (left) and P120 (right) mice across win sequence lengths (aA, aAA, aAAA), plotted as mean ± SEM. Right, peak z-scores at each sequence position plotted as c. Peak responses at the first reward in a sequence were 4.05 ± 0.05 in P120 mice and 3.60 ± 0.05 in P60 mice.

## Aging reduces tonic dopamine baseline in the nucleus accumbens

We then asked whether dopaminergic signaling itself undergoes broader change across the lifespan. To address this, we used fluorescence lifetime photometry in freely moving mice expressing the genetically encoded dopamine sensor chrimsonDlight (Fig. 5a). Unlike intensity-based photometry, fluorescence lifetime provides a concentration-dependent readout that is largely independent of motion artifacts, photobleaching, and variability in expression or optical coupling, enabling stable comparison of absolute dopamine levels across animals and timepoints. Prior work has further validated that lifetime signals track dopamine transients detected in intensity recordings, with transient increases in lifetime corresponding to dopamine release events (Lodder et al., 2025).

To verify that the fluorescence lifetime signal reflects biologically meaningful dopamine sensor dynamics, we pharmacologically perturbed dopaminergic signaling via systemic administration of the D1 receptor antagonist SCH23390. Antagonist administration produced a rapid and sustained reduction in fluorescence lifetime to a stable lower bound of approximately 1.7 ns, consistent with previously described sensor floor conditions (Fig. 5b). This reduction was accompanied by catatonic behavior, confirming effective pharmacological suppression.

Across recordings, fluorescence lifetime signals exhibited a stable tonic baseline punctuated by discrete phasic transients, which were clearly visible as upward deflections in representative traces and differed markedly in their absolute range between young and aged animals (Fig. 5c).

Quantification of baseline lifetime using the modal signal value revealed a significant age-dependent decline, with values remaining stable across early and mid-adulthood before falling substantially in mice aged 24 months and older (Fig. 5d). Peak lifetime values showed a parallel reduction in the oldest group (Fig. 5e).

When phasic responses were quantified as the difference between peak and baseline lifetime, amplitude remained stable across all age groups (Fig. 5f), indicating that while the absolute level of dopamine decreases with age, the capacity for phasic release is preserved. The age-related reduction therefore reflects a downward shift in the tonic dopamine baseline rather than a deterioration of phasic signaling mechanisms.

To exclude the possibility that lifetime measurements were confounded by tissue autofluorescence or optical artifacts, we quantified fluorescence intensity as a function of excitation power across animals. Fluorescence intensity scaled linearly with excitation power

across the tested range (Fig. 5g), consistent with fluorophore-dependent emission rather than nonlinear autofluorescence or hemodynamic contributions.

Together, these findings demonstrate that aging is associated with a selective reduction in tonic dopamine baseline in the nucleus accumbens, such that phasic transients occur from a lower absolute starting point while retaining their relative amplitude. Given that dopamine encodes reward history in a manner that scales with reinforcement information (Fig. 4), a reduction in the absolute range of dopamine signaling provides a plausible mechanism through which aged animals exhibit diminished sensitivity to recent outcomes, linking the neural and computational signatures of aging identified here.

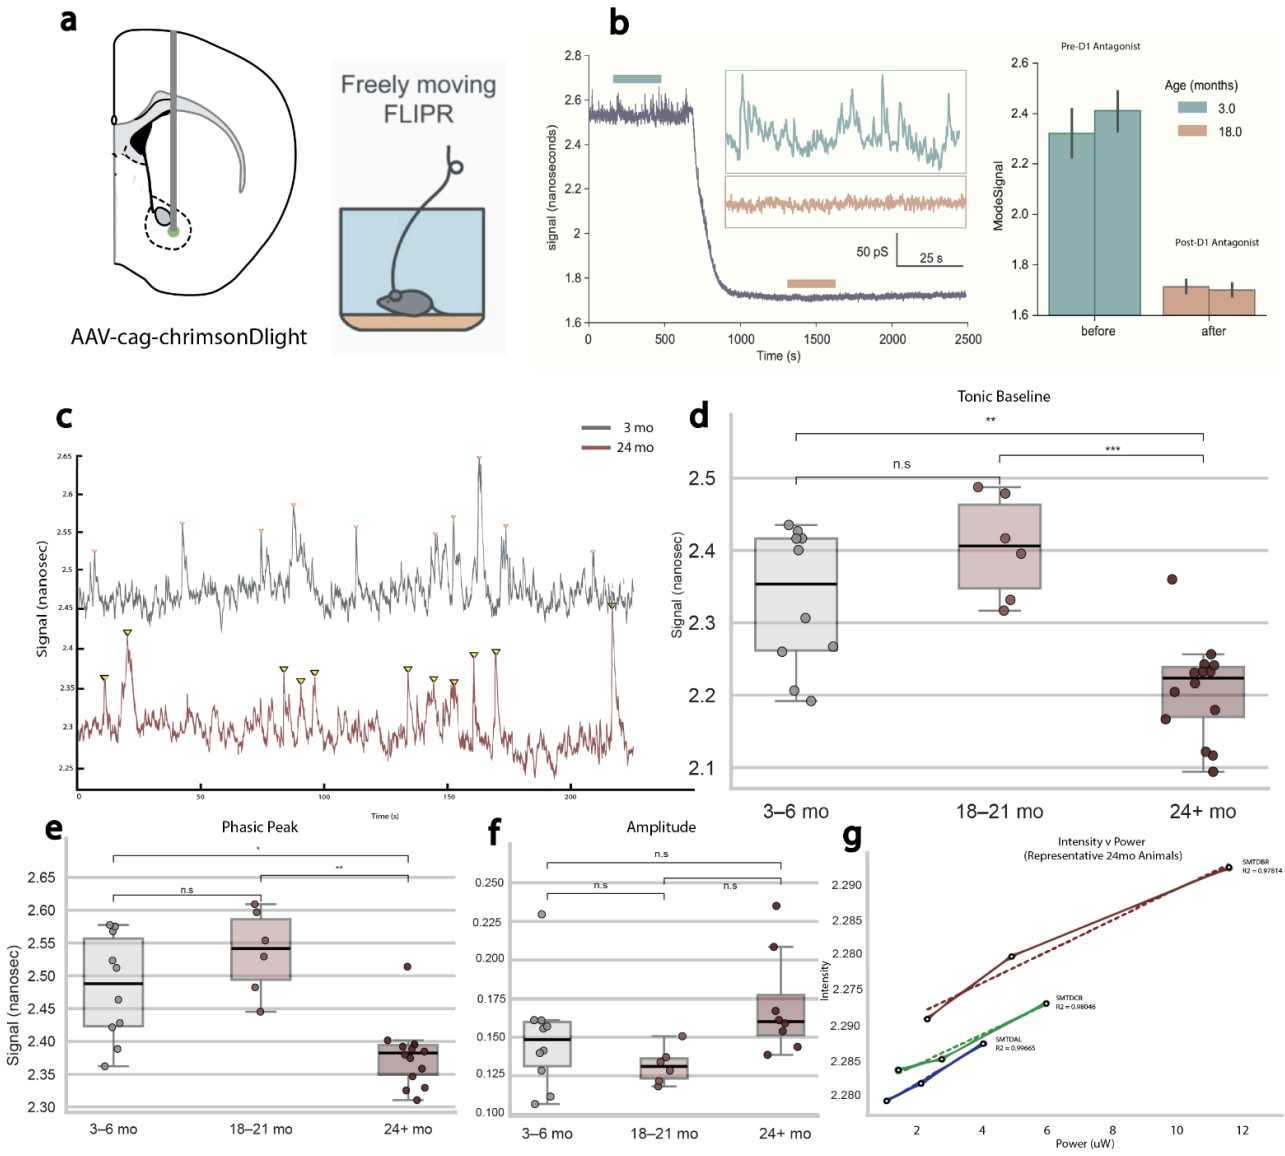


**Figure 5. Aging produces a downward shift in tonic dopamine baseline while preserving relative phasic signaling.**

*n = 10, 6, and 14 recordings from 4, 3, and 5 mice at 3–6, 18–21, and 24+ months respectively. Statistical unit is recording. Several animals contributed recordings at multiple timepoints. *p < 0.05, **p < 0.01, ***p < 0.001, Tukey HSD post-hoc following one-way ANOVA unless otherwise stated.*

**a, Surgical and recording schematic.** AAV-cag-chrimsonDlight was injected into the nucleus accumbens and fluorescence lifetime was recorded in freely moving mice using a FLIPR fiber photometry system.

**b, Pharmacological validation. Left**, representative fluorescence lifetime trace during systemic administration of the D1 receptor antagonist SCH23390 (teal bar, 3 months; orange bar, 18 months). Inset, expanded traces showing signal dynamics in each age group before and after antagonist administration. **Right**, modal lifetime signal before and after antagonist administration in 3 month (teal) and 18 month (orange) mice, plotted as mean ± SEM.

**c, Representative lifetime traces across age**. Overlay of fluorescence lifetime recordings from a 3 month (grey) and 24 month (red) mouse during free behaviour. Triangles denote detected phasic transient events.
**d, Baseline lifetime across age groups.** Modal baseline fluorescence lifetime per recording plotted as box plots (median, IQR, whiskers = 1.5× IQR) with individual recordings overlaid. One-way ANOVA: $F(2,21) = 9.80$, $p < 0.001$. Baseline lifetime was significantly reduced in 24+ month recordings relative to 3–6 months (*$p = 0.014$, Tukey HSD); 18–21 and 3–6 month groups did not differ significantly ($p = 0.269$). Selected pairwise comparisons shown (*$p < 0.05$, ***$p < 0.001$).
**e, Peak lifetime across age groups.** Peak fluorescence lifetime per recording plotted as in d. One-way ANOVA: $F(2,21) = 9.47$, $p = 0.001$. Peak lifetime was significantly reduced in 24+ month recordings relative to both 18–21 months (**$p = 0.001$, Tukey HSD) and 3–6 months (*$p = 0.013$, Tukey HSD); 18–21 and 3–6 month groups did not differ significantly ($p = 0.329$). Selected pairwise comparisons shown (*$p < 0.05$, **$p < 0.01$).
**f, Phasic amplitude across age groups.** Peak minus baseline fluorescence lifetime per recording plotted as in d. One-way ANOVA: $F(2,21) = 2.97$, $p = 0.073$. No significant differences detected.
**g, Autofluorescence control.** Fluorescence intensity plotted as a function of excitation power for representative 24 month animals, with linear regression fits and $R^2$ values shown ($R^2 = 0.978$, 0.980, 0.997).

## Tonic dopamine baseline does not covary with motivational state or effort expenditure

A key alternative interpretation of the age-related reduction in tonic dopamine baseline is that it reflects a generalised change in motivational state as opposed to a specific change in reinforcement signalling. Tonic dopamine has been implicated in regulating motivation and effort expenditure, and a reduction in motivational drive could independently account for some of the behavioural changes observed in aged animals.

To assess motivation independently of reinforcement learning, mice were trained on a progressive ratio schedule in which the number of responses required per pellet increased incrementally across the session (Fig. 6a). Breakpoint declined with increasing percentage of initial body weight, confirming that the task was sensitive to motivational state, with animals at higher body weights achieving lower breakpoints (Fig. 6b). The mixed effects model similarly showed a negative relationship between body weight and breakpoint that approached but did not

reach significance ($p = 0.060$), consistent with a modest but reliable effect of hunger state on effort expenditure.

Breakpoint did not differ across age groups (Fig. 6c; $F(3,30) = 1.38$, $p = 0.268$), indicating that motivational effort is broadly preserved across the lifespan.

We next asked whether individual differences in tonic dopamine baseline predicted motivational effort. Neither baseline dopamine lifetime nor peak amplitude was associated with PR breakpoint across animals (Fig. 6d; baseline: Pearson $r = 0.005$, $p = 0.991$; peak: Pearson $r = -0.183$, $p = 0.665$), indicating that tonic dopamine level does not track effort expenditure at the individual level.

To determine whether tonic dopamine baseline reflects acute physiological motivational state, we compared baseline lifetime between hungry and sated sessions within the same animals across age groups. Baseline lifetime did not differ significantly between hunger states in any age group (Fig. 6e), indicating that the tonic dopamine signal is not acutely modulated by food motivation.

Together, these findings demonstrate that the age-related reduction in tonic dopamine baseline is dissociable from both motivational effort and acute hunger state, suggesting that the decline in tonic dopamine more plausibly shapes the fidelity of reinforcement signalling than it reflects a global reduction in motivational drive.

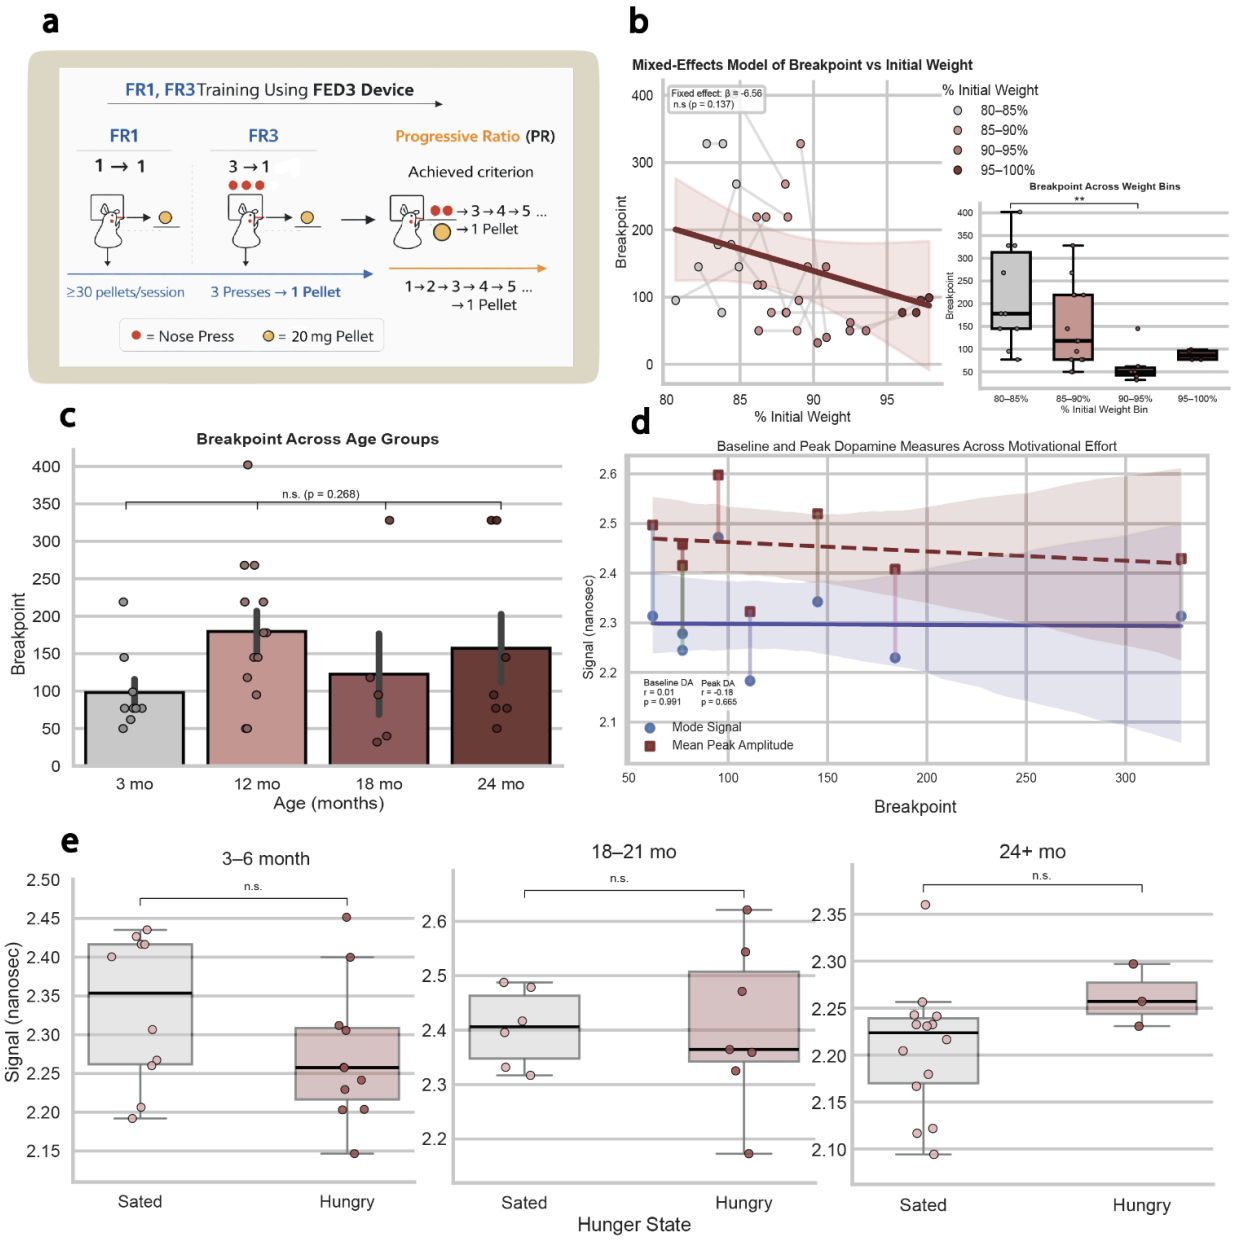


**Figure 6. Tonic dopamine baseline does not predict motivation or hunger state across aging.**

**p < 0.05, **p < 0.01, ***p < 0.001, Tukey HSD post-hoc following one-way ANOVA unless otherwise stated.*

**a, Progressive ratio task schematic.** Mice were trained on fixed ratio schedules (FR1, FR3) using a FED3 device before progressing to a progressive ratio schedule in which the number of nose presses required per pellet increased incrementally across the session. Breakpoint was defined as the highest ratio completed before cessation of responding.

**b, Breakpoint as a function of body weight. Left**, PR breakpoint plotted against percentage of initial body weight per recording, with mixed effects model fit and individual animal trajectories shown. **Right**, breakpoint across percentage initial weight bins plotted as box plots (median, IQR, whiskers = 1.5× IQR) with individual recordings

overlaid. One-way ANOVA: $F(3,30) = 4.65$, $p = 0.009$. Selected pairwise comparisons shown (**$p < 0.01$). n = 34 recordings from 13 mice.
**c, Breakpoint across age groups.** PR breakpoint per recording plotted as box plots with individual recordings overlaid across 3, 12, 18, and 24 month age groups. One-way ANOVA: $F(3,30) = 1.38$, $p = 0.268$. No significant differences detected. n = 9, 13, 5, 7 recordings at 3, 12, 18, and 24 months respectively.
**d, Dopamine lifetime measures as a function of PR breakpoint.** Baseline dopamine lifetime (mode signal, blue) and mean peak amplitude (dark red) plotted against PR breakpoint per animal, with linear regression fits and 95% confidence intervals. Neither baseline (Pearson $r = 0.005$, $p = 0.991$) nor peak dopamine (Pearson $r = -0.183$, $p = 0.665$) predicted breakpoint. n = 8 animals.
**e, Baseline dopamine lifetime in sated and hungry states across age groups.** Modal baseline lifetime per session plotted as box plots with individual sessions overlaid for 3–6 month (left), 18–21 month (middle), and 24+ month (right) animals in sated and hungry conditions. No significant difference between hunger states was detected in any age group (3–6 months: Mann–Whitney $U = 39.0$, $p = 0.176$; 18–21 months: $U = 45.0$, $p = 0.364$; 24+ months: $U = 29.0$, $p = 0.136$). n = 12 hungry and 10 sated sessions at 3–6 months; n = 7 hungry and 10 sated sessions at 18–21 months; n = 3 hungry and 12 sated sessions at 24+ months.

## Age-dependent transcriptional remodeling of the striatum at single-cell resolution

To situate the dopaminergic and behavioral changes identified above within a broader striatal context, we performed a hypothesis-free re-analysis of single-nucleus RNA sequencing data deposited by Linker et al. (2025), examining age-dependent transcriptional changes across all major striatal cell types (Fig. 7).

Unsupervised clustering identified ten transcriptionally distinct populations in both age groups (Fig. 7a). Two patterns in the compositional and burden analyses motivated deeper examination of specific cell types. Astrocytes showed the largest proportional decline with age yet carried the lowest differential expression burden, indicating state compression of the surviving population rather than transcriptional activation. Microglia demonstrated the opposite, featuring a modest proportional increase accompanied by the greatest transcriptional disruption of any cell type, exceeding 2,000 differentially expressed genes (Fig. 7b, c). Striatal interneurons, which

increased proportionally and showed substantial differential expression burden, were examined given their role in regulating striatal circuit gain and dopamine release probability.

Microglia exhibited coordinated loss of homeostatic markers P2ry12, Tmem119, Sall1, Hexb, and Cx3cr1 with concurrent gain of DAM markers Trem2, Apoe, Lpl, Cst7, and Axl (Fig. 7d), indicating widespread reactive transition in the aged striatum. Among surviving astrocytes, Vim and S100b were significantly elevated while Slc1a2, encoding the primary astrocytic glutamate transporter GLT-1, was significantly reduced (Fig. 7e), suggesting that aged astrocytes are both fewer and functionally compromised in their capacity to buffer extracellular glutamate.

Spiny Projection neurons and striatal interneurons displayed convergent upregulation of mitochondrially encoded transcripts including Lars2, Gm42418, mt-Co2, mt-Nd3, and mt-Atp8, alongside downregulation of genes involved in epigenetic regulation, proteostasis, and synaptic signaling (Fig. 7f, g).

Together, these findings describe an aged striatal environment characterized by reactive glial remodeling, impaired glutamate homeostasis, and convergent neuronal bioenergetic stress across cell types. This cellular context would be expected to degrade the precision of reinforcement signals and impair the synaptic machinery required to translate dopamine transients into behavioral updating, providing a systems-level account of the age-related impairment in reinforcement learning that complements the physiological and behavioral findings reported above.

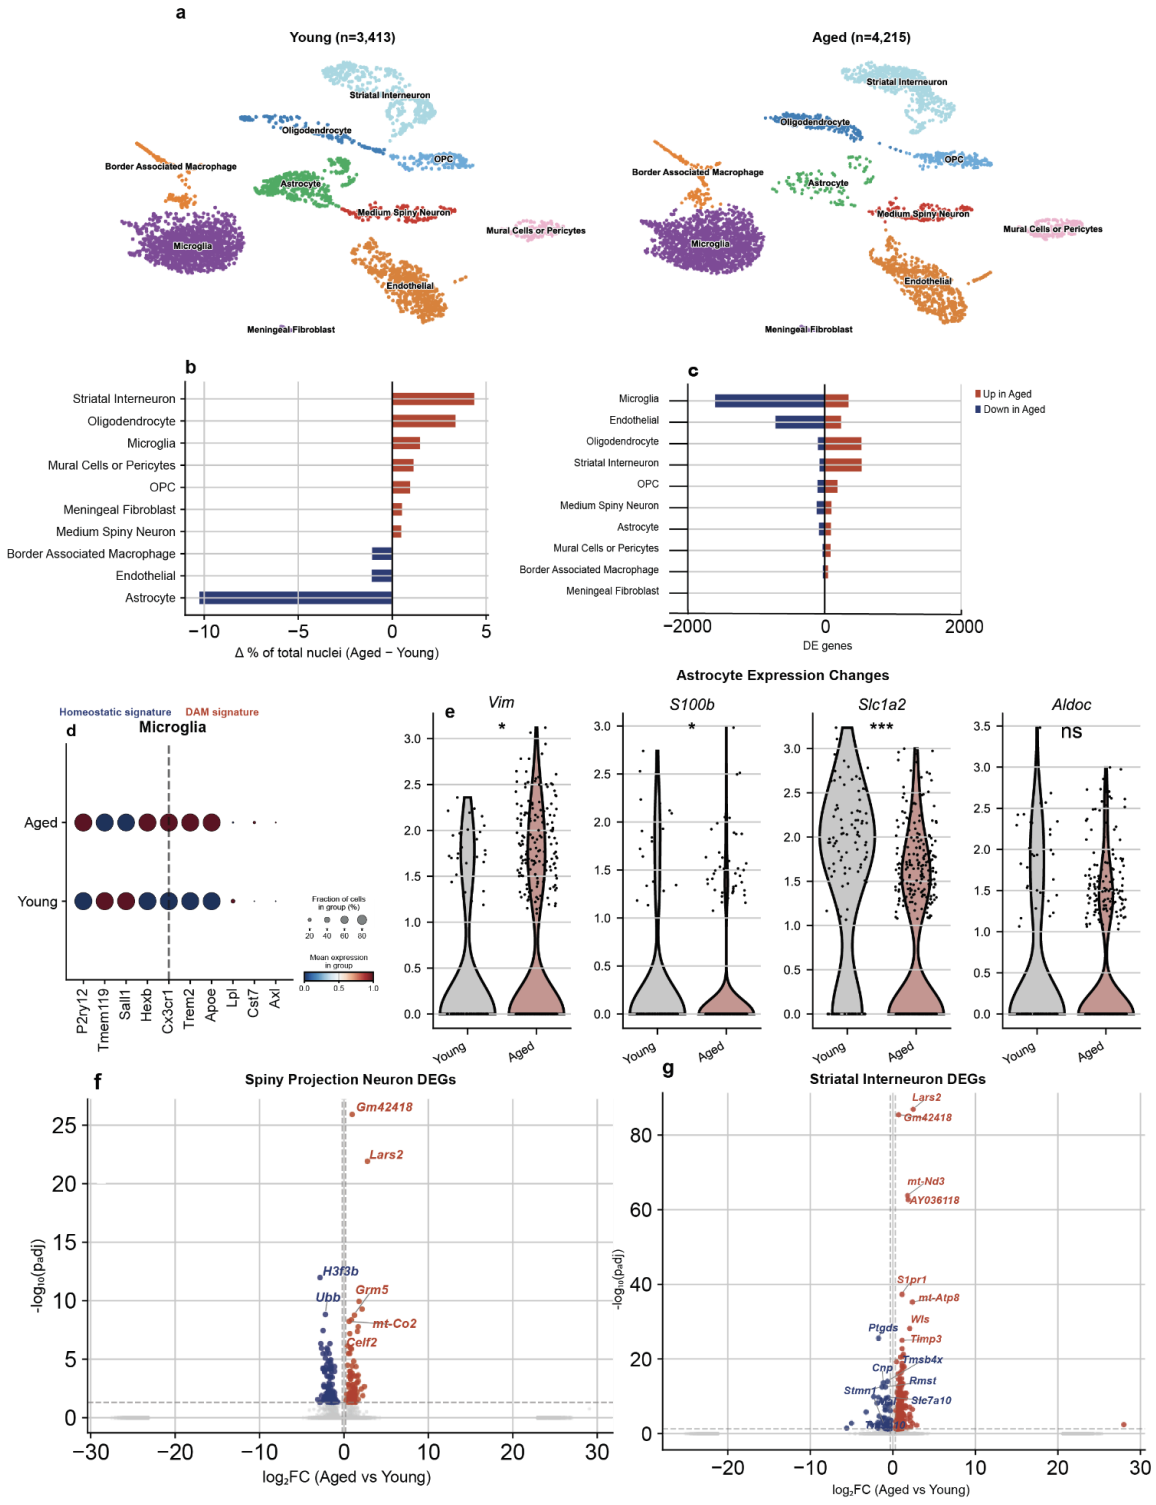


**Figure 7. Aging drives coordinated transcriptional remodeling across striatal cell types.**

*Data derived from reanalysis of deposited single-nucleus RNA sequencing data from Linker et al. (2025), Nature Communications 16, 8496. Young: 2 months; aged: 18 months. n = 4 mice per group. *p < 0.05, ***p < 0.001, Mann-Whitney U test unless otherwise stated.*

**a, Single-nucleus transcriptomic atlas of young and aged striatum.** UMAPs of nuclei from young (left, n = 3,413) and aged (right, n = 4,215) mouse striatum colored by cell type. Ten transcriptionally distinct clusters were identified across both age groups.
**b, Cell type proportional composition shifts with age.** Change in percentage of total nuclei per cell type (Δ% = aged minus young). Astrocytes show the largest proportional decline; striatal interneurons show the largest proportional increase.
**c, Differential expression burden by cell type.** Number of upregulated (red) and downregulated (blue) DEGs per cell type (padj < 0.05, $|\log_2 FC| > 0.2$). Microglia carry the greatest transcriptional burden; astrocytes the least.
**d, Microglia homeostatic-to-DAM transition.** Dot plot of homeostatic (P2ry12, Tmem119, Sall1, Hexb, Cx3cr1) and DAM (Trem2, Apoe, Lpl, Cst7, Axl) marker expression in young and aged microglia. Dot size reflects fraction of cells expressing the gene; color reflects mean expression level.
**e, Astrocyte marker gene expression across age.** Violin plots of Vim, S100b, Slc1a2, and Aldoc in young (grey) and aged (salmon) astrocytes (young n = 436, aged n = 105). Vim and S100b are significantly elevated in aged astrocytes ($p = 0.019$ and $p = 0.024$ respectively); Slc1a2 is significantly reduced ($p < 0.001$); Aldoc is unchanged. Selected comparisons shown (*$p < 0.05$, ***$p < 0.001$, Mann-Whitney U test).
**f, Spiny Projection neuron differential expression.** Volcano plot of SPN DEGs (aged vs young). Top upregulated transcripts include Gm42418, Lars2, and mt-Co2; top downregulated include H3f3b, Ubb, Grm5, and Celf2.
**g, Striatal interneuron differential expression.** Volcano plot of interneuron DEGs (aged vs young, 539 upregulated, 74 downregulated). Top upregulated transcripts include Lars2, Gm42418, mt-Nd3, and mt-Atp8.

# IV. Discussion

## Aging is associated with reduced reinforcement sensitivity and a downward shift in tonic dopamine baseline

Adaptive behavior depends on reinforcement signals that enable organisms to update action policies in response to changing outcomes. Here, we show that aging is associated with reduced reinforcement sensitivity during adaptive choice and a concurrent downward shift in tonic dopamine baseline in the nucleus accumbens, while preserving the relative structure of phasic signaling dynamics. These findings identify a physiological correlate of age-dependent changes in reinforcement learning and establish tonic dopaminergic baseline as a candidate mechanism through which aging may alter reinforcement-dependent behavioral updating.

A central observation is that tonic dopamine baseline declines with age while phasic signaling dynamics remain preserved. This dissociation suggests that the capacity of dopaminergic neurons to generate transient release events remains relatively intact, while the baseline conditions governing extracellular dopamine tone are altered. Multiple mechanisms could contribute to such a selective shift. Aging-related cellular changes may alter intrinsic excitability of dopaminergic neurons, for example through changes in ion channel expression or membrane properties that reduce baseline firing rates without impairing burst firing. Cholinergic interneurons exert powerful control over dopamine release probability and could modulate basal dopamine tone independently of phasic dynamics. Changes in dopamine clearance mechanisms, including altered expression or function of the dopamine transporter, could further reduce extracellular dopamine levels by accelerating reuptake. Global neuromodulatory or hormonal changes across aging may additionally alter the excitability landscape of dopaminergic neurons or their afferent inputs. Critically, distinguishing between these possibilities requires molecular resolution at the level of the dopaminergic neuron and terminal itself. The transcriptional changes we describe in the aged striatum reflect the aggregate cellular landscape of a target region, and cannot speak to what is occurring within dopaminergic neurons upstream.

## Exploratory molecular evidence points toward bioenergetic and cholinergic mechanisms at dopaminergic terminals

To preliminarily address this, we performed an exploratory reanalysis of translatomic data from Kilfeather et al. (2024), which employed a DAT-TRAP strategy to affinity-purify actively translating ribosomes selectively from dopaminergic axonal terminals in ventral striatum, providing a compartment-specific molecular readout unavailable from other publicly deposited

datasets. Given the limited replicate numbers in this dataset, findings should be interpreted as preliminary and hypothesis-generating. Nonetheless, several signals are noteworthy.

Mitochondrially encoded respiratory chain subunits including mt-Nd1, mt-Nd2, mt-Nd4, mt-Nd5, mt-Co1, and mt-Co2 were among the most consistently upregulated transcripts in aged dopaminergic terminals, suggesting bioenergetic stress within the terminal compartment that is consistent with the high energetic demands imposed by sustained vesicle cycling and dopamine reuptake. Among downregulated transcripts was Slc5a7, encoding the high-affinity choline transporter, raising the possibility that cholinergic regulation of dopamine release probability at striatal terminals is altered with age. Additionally, Pkib, encoding the beta isoform of the protein kinase A inhibitor, was upregulated in aged ventral striatum, providing tentative transcriptomic support for the possibility that postsynaptic PKA signaling gain is reduced in aged striatum. These exploratory findings converge with the broader pattern of neuronal mitochondrial stress identified across striatal cell types in Figure 7, and together motivate future work directly interrogating dopaminergic terminal bioenergetics, cholinergic coupling, and intracellular signaling fidelity across the lifespan.

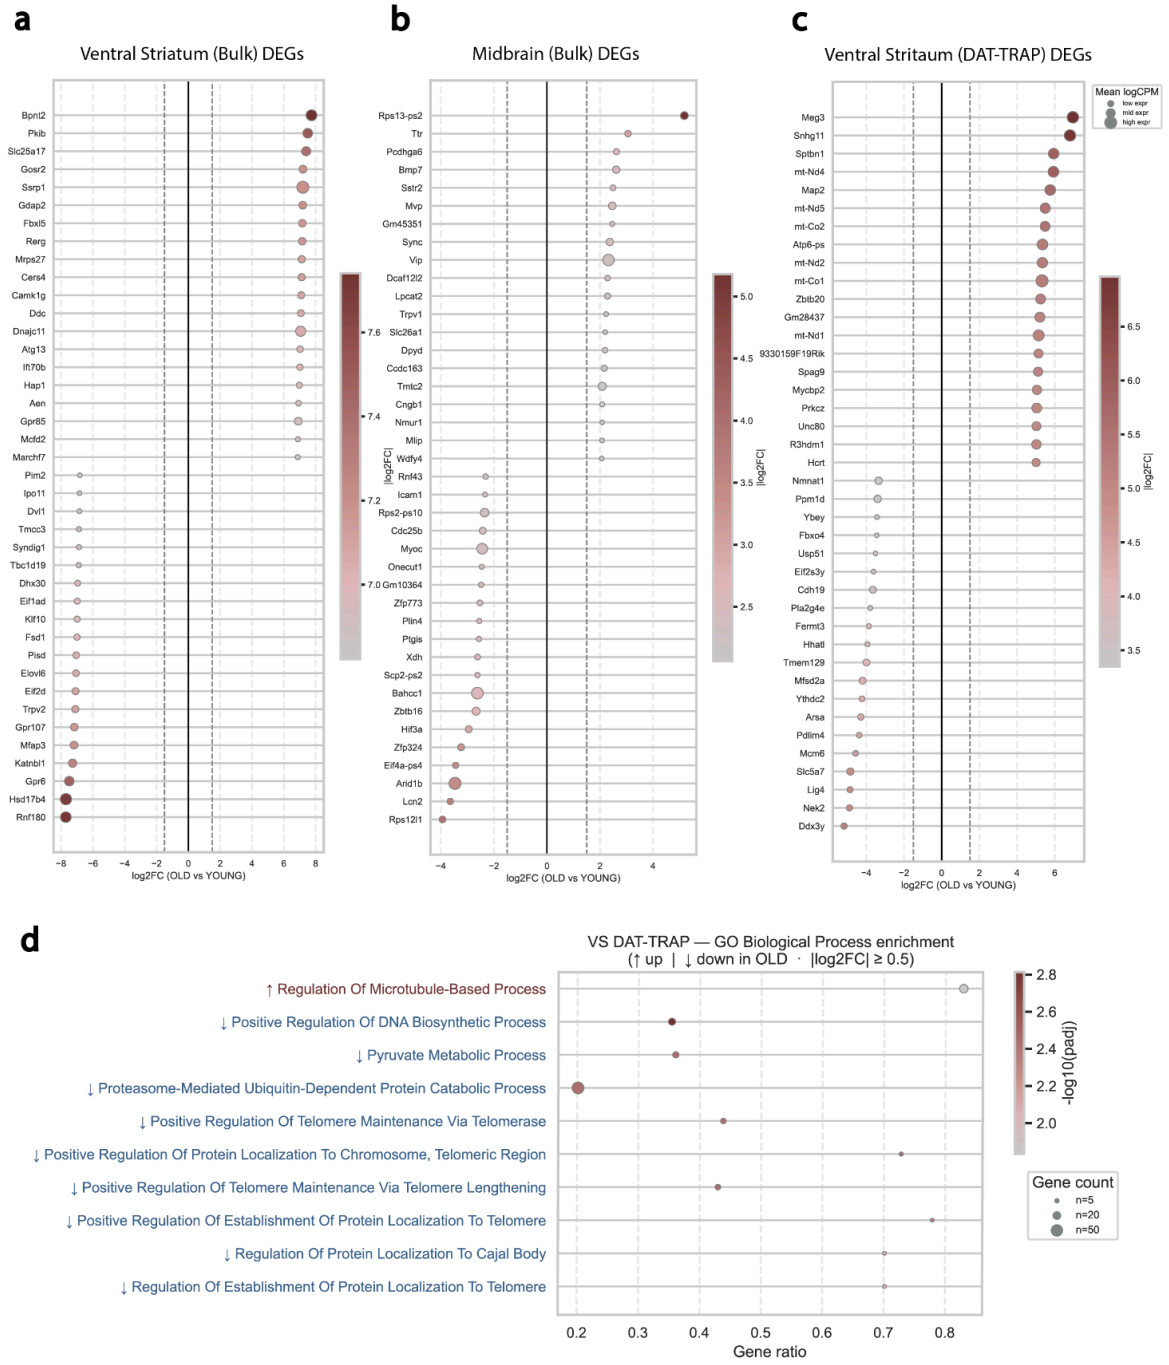


**Supplemental Figure 1. Reanalysis of Kilfeather et al. (2024) TRAP sequencing data reveals age-dependent transcriptional remodeling in ventral striatum and midbrain dopaminergic neurons.**

*Data reanalyzed from Kilfeather et al. (2024), Cell Reports 43, 113784. Young: 3 months; aged: 18+ months.*

**a, Ventral striatum bulk differential gene expression.** Top differentially expressed genes in ventral striatum bulk tissue ranked by $|\log_2 FC|$ (old vs young), thresholded at $|\log_2 FC| \geq 1.5$. Dot size reflects mean logCPM; color intensity reflects $|\log_2 FC|$. Note n = 1 per condition; results are exploratory.
**b, Midbrain bulk differential gene expression.** Plotted as a for ventral striatum bulk tissue.
**c, DAT-TRAP ventral striatum differential gene expression.** Top differentially expressed genes in DAT-TRAP ventral striatum, reflecting translatomic changes specific to the dopaminergic axonal terminal compartment, thresholded at $|\log_2 FC| \geq 1.5$. Mitochondrially encoded transcripts are prominent among genes upregulated in aged animals; Slc5a7 is among the most consistently downregulated. Note n = 1 per condition; results are exploratory.
**d, Gene Ontology Biological Process enrichment in aged DAT-TRAP ventral striatum terminals.** Top enriched GO Biological Process terms from DAT-TRAP differentially expressed genes ($|\log_2 FC| \geq 0.5$, adjusted $p < 0.05$). Arrow direction indicates whether the gene set is upregulated (↑) or downregulated (↓) in aged animals. Dot size reflects gene count; color reflects $-\log_{10}(\text{padj})$. Proteasome-mediated ubiquitin-dependent protein catabolism and pyruvate metabolic process are among the most significantly downregulated pathways; regulation of microtubule-based processes is upregulated.

## Functional implications of reduced tonic dopamine baseline for striatal reinforcement signaling

Reinforcement-dependent plasticity depends not only on the presence of phasic dopamine transients, but also on their ability to engage intracellular signaling pathways within striatal projection neurons. Our findings raise the possibility that reduced tonic dopamine baseline in aging alters the effective gain of dopamine signaling, such that phasic transients produce weaker activation of downstream D1-SPN and D2-SPN intracellular plasticity pathways (Supplementary Fig. 2a,b). Under this model, reinforcement signals remain present but may be less effective in driving plasticity due to reduced postsynaptic signaling engagement. The tentative upregulation of Pkib in aged ventral striatum described above lends independent transcriptomic support to this possibility. We are actively conducting direct measurement of dopamine dynamics alongside striatal SPN protein kinase A signaling to determine whether a causative relationship exists.

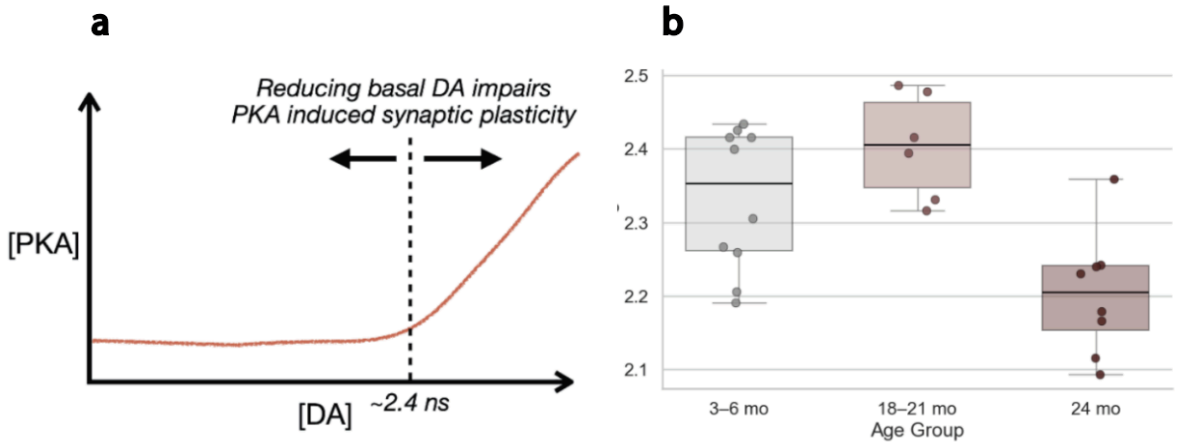


**Supplemental Figure 2: Age-dependent reduction in tonic dopamine baseline shifts signaling further from PKA activation threshold**

**a, Schematic of striatal dopamine concentration and PKA activation relationship.** The PKA activation curve as a function of dopamine
concentration is understood to be sigmoidal, with a steep activation region centered around approximately 2.4 ns fluorescence lifetime. A reduction in basal dopamine tone shifts the resting state leftward along this curve, increasing the distance between tonic baseline and the activation threshold and thus requiring larger phasic excursions to drive equivalent PKA-dependent plasticity.
**b, Age-dependent decline in tonic dopamine baseline.** Modal fluorescence lifetime per animal across age groups (3 to 6, 18 to 21, and 24+ months), plotted as box plots (median, IQR, whiskers = 1.5× IQR) with individual animals overlaid. Baseline lifetime remains stable between 3 to 6 and 18 to 21 months before declining substantially at 24+ months, falling below the approximate PKA activation threshold of 2.4 ns.

An alternative, non-mutually exclusive possibility is that aging alters transcriptional permissiveness. Even if dopamine transients retain sufficient amplitude to activate receptors, transcriptional or synaptic mechanisms required for reinforcement-dependent updating may become less responsive. Variability in the capacity for stimuli-responsive transcription has previously been demonstrated with predator odor stress paradigms in mice (Long et al., 2025). Under this model, reinforcement signals remain present but the ability of striatal circuits to translate those signals into stable behavioral updates may be diminished. Further work will be required to determine how aging affects the coupling between dopamine release, intracellular signaling, and synaptic plasticity.

## Behavioral experience may preserve reinforcement-related circuit function across aging

Notably, in an ongoing longitudinal two-armed bandit task behavioral experiment, behavioral experience appears to preserve reinforcement sensitivity in large subsets of animals repeatedly exposed to the task across their lifespan. These animals maintain stable behavioral performance and model parameters into advanced age as compared to mice of the same age executing the task with less prior exposure, suggesting that sustained engagement in the paradigm across life may stabilize underlying neural or molecular processes. Repeated behavioral experience could preserve dopaminergic circuit function through sustained activity-dependent plasticity, preserve transcriptional responsiveness, or prevent maladaptive molecular remodeling associated with aging. These observations raise the possibility that behavioral engagement itself may influence the trajectory of reinforcement-related neural function across aging, and suggest potential avenues for behavioral interventions aimed at preserving cognitive flexibility.

## Conclusions and implications

Taken together, these findings characterize parallel age-dependent changes in adaptive behavior and dopaminergic physiology and identify tonic dopamine baseline as a physiological feature that tracks these changes. While causal relationships remain to be established, these results provide a mechanistic framework linking dopaminergic signaling dynamics to age-dependent differences in reinforcement sensitivity. By establishing a quantitative link between behavioral, physiological, and molecular levels of analysis, this work provides a foundation for future

studies aimed at identifying the causal mechanisms that regulate tonic dopamine tone and determining how those mechanisms may be leveraged to promote cognitive resilience during aging.

**Limitations of the study**

*Limitation 1.* Behavioral measurements and fluorescence lifetime recordings were obtained in separate cohorts. As a result, it was not possible to directly test how tonic dopamine baseline relates to reinforcement sensitivity within the same animal. Although both behavioral reinforcement sensitivity and tonic dopamine baseline decline with age, this cohort separation prevents direct assessment of whether individual differences in dopamine baseline predict reinforcement learning parameters or behavioral updating dynamics. Future experiments combining chronic fluorescence lifetime recordings with 2-ABT behavior in the same implanted animals will be essential for establishing whether tonic dopamine baseline directly tracks reinforcement sensitivity at the individual level. Such experiments would further enable trial-by-trial coupling between dopamine dynamics and behavioral updating to be quantified, providing a more direct test of the proposed mechanistic relationship.

*Limitation 2.* Present findings establish a physiological correlate of altered reinforcement sensitivity but do not demonstrate causality. Direct manipulation of tonic dopamine levels will be required to determine whether restoring baseline dopamine tone is sufficient to rescue reinforcement-dependent updating in aged animals. A chemogenetic “tickle” approach has previously been successfully demonstrated in the Sabatini lab.

*Limitation 3.* rDlight sensor has not been fully characterized and mapped to absolute dopamine concentration *in vivo*. While lifetime signals reliably reflect changes in dopamine binding to the

sensor, additional approaches such as electrochemical recordings, complementary optical sensors, or direct measurements of dopamine synthesis and clearance mechanisms will be necessary to fully characterize the biochemical and cellular basis of the observed baseline shift.

*Limitation 4.* Dopamine measurements were restricted to the NAcc and it remains unclear whether similar tonic baseline shifts occur across other dopaminergic projection targets, including dorsal striatum or cortical regions, which are known to influence 2-ABT gameplay. Because different dopaminergic pathways support distinct aspects of learning and behavior, determining the anatomical specificity of tonic dopamine changes will be important for understanding how aging affects different functionally distinct aspects of the behavior.

# V. References


1. Oberman L, Pascual-Leone A (2013) Changes in plasticity across the lifespan: Cause of disease and target for intervention. Prog Brain Res 207:91–120.
2. Voglewede RL, Vandemark KM, Davidson AM, DeWitt AR, Heffler MD, Trimmer EH, Mostany R (2019) Reduced sensory-evoked structural plasticity in the aging barrel cortex. Neurobiol Aging 81:222–233.
3. Marzola P, Melzer T, Pavesi E, Gil-Mohapel J, Brocardo PS (2023) Exploring the role of neuroplasticity in development, aging, and neurodegeneration. Brain Sci 13:1610.
4. Citri A, Malenka RC (2008) Synaptic plasticity: Multiple forms, functions, and mechanisms. Neuropsychopharmacol 33:18–41.
5. Pan Y, Monje M (2020) Activity shapes neural circuit form and function: A historical perspective. J Neurosci 40:944–954.
6. Bozorgmehr T, Ardiel EL, McEwan AH, Rankin CH (2013) Mechanisms of plasticity in a *Caenorhabditis elegans* mechanosensory circuit. Front Physiol 4:88.
7. Sasakura H, Mori I (2013) Behavioral plasticity, learning, and memory in *C. elegans*. Curr Opin Neurobiol 23:92–99.
8. Favila N, Capece Marsico J, Pacheco CM, et al. (2025) Heterogeneous plasticity of amygdala interneurons in associative learning and extinction. Nat Commun 16:9926.

9. Hille M, Kühn S, Kempermann G, Bonhoeffer T, Lindenberger U (2024) From animal models to human individuality: Integrative approaches to the study of brain plasticity. Neuron 112:3522–3541.
10. Herzberg MP, Nielsen AN, Luby J, et al. (2025) Measuring neuroplasticity in human development: The potential to inform the type and timing of mental health interventions. Neuropsychopharmacol 50:124–136.
11. Whyte AJ, Kietzman HW, Swanson AM, Butkovich LM, Barbee BR, Bassell GJ, Gross C, Gourley SL (2019) Reward-related expectations trigger dendritic spine plasticity in the mouse ventrolateral orbitofrontal cortex. J Neurosci 39:4595–4605.
12. Breton YA, Seeland KD, Redish AD (2015) Aging impairs deliberation and behavioral flexibility in inter-temporal choice. Front Aging Neurosci 7:41.
13. Schoenbaum G, Nugent S, Saddoris MP, Gallagher M (2002) Teaching old rats new tricks: Age-related impairments in olfactory reversal learning. Neurobiol Aging 23:555–564.
14. Zhang J, Liu D, Fu P, Liu Z-Q, Lai C, Yang C-Q, Chen K, Bao W-D, Hu F, Du H-Y, Yang W, Wang J, Man H-Y, Lu Y, Zhu L-Q (2022) Social isolation reinforces aging-related behavioral inflexibility by promoting neuronal necroptosis in basolateral amygdala. Mol Psychiatry 27:4050–4063.
15. Tait DS, Chase EA, Brown VJ (2013) Tacrine improves reversal learning in older rats. Neuropharmacology 73:284–289.
16. Norris CM, Korol DL, Foster TC (1996) Increased susceptibility to induction of long-term depression and long-term potentiation reversal during aging. J Neurosci 16:5382–5392.
17. Foster TC, Kumar A (2007) Susceptibility to induction of long-term depression is associated with impaired memory in aged Fischer 344 rats. Neurobiol Learn Mem 87:522–535.
18. Ducote AL, Voglewede RL, Mostany R (2024) Dendritic spines of layer 5 pyramidal neurons of the aging somatosensory cortex exhibit reduced volumetric remodeling. J Neurosci 44:e1378242024.
19. Boric K, Muñoz P, Gallagher M, Kirkwood A (2008) Potential adaptive function for altered long-term potentiation mechanisms in aging hippocampus. J Neurosci 28:8034–8039.
20. Wang Y (2008) Differential effect of aging on synaptic plasticity in the ventral and dorsal striatum. Neurobiol Learn Mem 89:70–75.
21. Chepkova AN, Schönfeld S, Sergeeva OA (2015) Age-related alterations in the expression of genes and synaptic plasticity associated with nitric oxide signaling in the mouse dorsal striatum. Neural Plast 2015:458123.
22. Pauwels L, Chalavi S, Swinnen SP (2018) Aging and brain plasticity. Aging (Albany NY) 10:1789–1790.
23. James CE, Müller DM, Müller CAH, Van De Looij Y, Altenmüller E, Kliegel M, Van De Ville D, Marie D (2024) Randomized controlled trials of non-pharmacological interventions for healthy seniors: Effects on cognitive decline, brain plasticity and activities of daily living—A 23-year scoping review. Heliyon 10:e26674.
24. Speranza L, di Porzio U, Viggiano D, de Donato A, Volpicelli F (2021) Dopamine: The neuromodulator of long-term synaptic plasticity, reward, and movement control. Cells 10:735.

25. Pignatelli M, Bonci A (2015) Role of dopamine neurons in reward and aversion: A synaptic plasticity perspective. Curr Opin Neurobiol 35:1145–1157.
26. Vidal-Gadea AG, Pierce-Shimomura JT (2012) Conserved role of dopamine in the modulation of behavior. Commun Integr Biol 5:440–447.
27. Menegas W, Bergan JF, Ogawa SK, Isogai Y, Venkataraju KU, Osten P, Uchida N, Watabe-Uchida M (2015) Dopamine neurons projecting to the posterior striatum form an anatomically distinct subclass. eLife 4:e10032.
28. Trutti AC, Mulder MJ, Hommel B, Forstmann BU (2019) Functional neuroanatomical review of the ventral tegmental area. NeuroImage 191:258–268.
29. Haber SN (2011) Neuroanatomy of reward: A view from the ventral striatum. In: *Neurobiology of Sensation and Reward* (Gottfried JA, ed). Boca Raton, FL: CRC Press/Taylor & Francis.
30. Cai J, Tong Q (2022) Anatomy and function of ventral tegmental area glutamate neurons. Front Neural Circuits 16:867053.
31. Zhuang Y, Xu P, Mao C, Wang L, Krumm B, Zhou XE, Huang S, Liu H, Cheng X, Huang X-P, Shen D-D, Xu T, Liu Y-F, Wang Y, Guo J, Jiang Y, Jiang H, Melcher K, Roth BL, Zhang Y, Zhang C, Xu HE (2021) Structural insights into the human D1 and D2 dopamine receptor signaling complexes. Cell 184:931–942.e18.
32. Teng X, Chen S, Wang Q, Chen Z, Wang X, Huang N, Zheng S (2022) Structural insights into G protein activation by D1 dopamine receptor. Sci Adv 8:eabo4158.
33. Gerfen CR, Engber TM, Mahan LC, Susel Z, Chase TN, Monsma FJ Jr, Sibley DR (1990) D1 and D2 dopamine receptor-regulated gene expression of striatonigral and striatopallidal neurons. Science 250:1429–1432.
34. Skeberdis VA, Chevaleyre V, Lau CG, Hemar A, Pertwee RG, Bliss TVP, Malenka RC (2006) Protein kinase A regulates calcium permeability of NMDA receptors. Nat Neurosci 9:501–510.
35. Iino Y, Sawada T, Yamaguchi K, et al. (2020) Dopamine D2 receptors in discrimination learning and spine enlargement. Nature 579:555–560.
36. Lau GC, Saha S, Faris R, Russek SJ (2004) Up-regulation of NMDAR1 subunit gene expression in cortical neurons via a PKA-dependent pathway. J Neurochem 88:564–575.
37. Nayak A, Zastrow D, Lickteig R, Zahniser NR, Browning MD (1998) Maintenance of late-phase LTP is accompanied by PKA-dependent increase in AMPA receptor synthesis. Nature 394:680–683.
38. Vaughn AR, Davis MJ, Sivamani RK, Isseroff RR (2018) A concise review of the conflicting roles of dopamine-1 versus dopamine-2 receptors in wound healing. Molecules 23:50.
39. Schultz W (2016) Dopamine reward prediction-error signalling: A two-component response. Nat Rev Neurosci 17:183–195.
40. Rescorla RA, Wagner AR (1972) Classical conditioning: Current research and theory. New York: Appleton-Century-Crofts.
41. Dayan P, Daw ND (2008) Decision theory, reinforcement learning, and the brain. Cogn Affect Behav Neurosci 8:429–453.
42. Katahira K (2015) The relation between reinforcement learning parameters and the influence of reinforcement history on choice behavior. J Math Psychol 66:59–69.
43. Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science 275:1593–1599.

44. Eshel N, Tian J, Bukwich M, Uchida N (2016) Dopamine neurons share common response function for reward prediction error. Nat Neurosci 19:479–486.
45. Steinberg EE, Keiflin R, Boivin JR, Witten IB, Deisseroth K, Janak PH (2013) A causal link between prediction errors, dopamine neurons, and learning. Nat Neurosci 16:966–973.
46. Waelti P, Dickinson A, Schultz W (2001) Dopamine responses comply with basic assumptions of formal learning theory. Nature 412:43–48.
47. Yamaguchi K, Maeda Y, Sawada T, Iino Y, Tajiri M, Nakazato R, Ishii S, Kasai H, Yagishita S (2022) A behavioural correlate of the synaptic eligibility trace in the nucleus accumbens. Sci Rep 12:1921.
48. Yagishita S, Hayashi-Takagi A, Ellis-Davies GCR, Urakubo H, Ishii S, Kasai H (2014) A critical time window for dopamine actions on the structural plasticity of dendritic spines. Science 345:1616–1620.
49. Shouval HZ, Kirkwood A (2025) Eligibility traces as a synaptic substrate for learning. Curr Opin Neurobiol 91:102978.
50. Floresco SB, West AR, Ash B, Moore H, Grace AA (2003) Afferent modulation of dopamine neuron firing differentially regulates tonic and phasic dopamine transmission. Nat Neurosci 6:968–973.
51. Grace, A. A. (1991). Phasic versus tonic dopamine release and the modulation of dopamine system responsivity: A hypothesis for the etiology of schizophrenia. *Neuroscience* 41, 1–24.
52. Wang Y, Toyoshima O, Kunimatsu J, Yamada H, Matsumoto M (2021) Tonic firing mode of midbrain dopamine neurons continuously tracks reward values changing moment-by-moment. eLife 10:e63166.
53. Mikhael JG, Bogacz R (2016) Learning reward uncertainty in the basal ganglia. PLoS Comput Biol 12:e1005062.
54. Romero Pinto S, Uchida N (2025) Tonic dopamine and biases in value learning linked through a biologically inspired reinforcement learning model. Nat Commun 16:7529.
55. Xu Y, Lin Y, Yu M, Zhou K (2024) The nucleus accumbens in reward and aversion processing: Insights and implications. Front Behav Neurosci 18:1420028.
56. Klawonn AM, Malenka RC (2019) Nucleus accumbens modulation in reward and aversion. Cold Spring Harb Symp Quant Biol 83:119–129.
57. Yin HH, Ostlund SB, Balleine BW (2008) Reward-guided learning beyond dopamine in the nucleus accumbens: The integrative functions of cortico-basal ganglia networks. Eur J Neurosci 28:1437–1448.
58. Hikida T, Yao S, Macpherson T, et al. (2020) Nucleus accumbens pathways control cell-specific gene expression in the medial prefrontal cortex. Sci Rep 10:1838.
59. Yawata S, Yamaguchi T, Danjo T, Hikida T, Nakanishi S (2012) Pathway-specific control of reward learning and its flexibility via selective dopamine receptors in the nucleus accumbens. Proc Natl Acad Sci USA 109:12764–12769.
60. Lai C-W, Shih C-W, Chang C-H (2022) Analysis of collateral projections from the lateral orbitofrontal cortex to nucleus accumbens and basolateral amygdala in rats. J Neurophysiol 127:1535–1546.
61. Xu Y, Lin Y, Yu M, Zhou K (2024) The nucleus accumbens in reward and aversion processing: Insights and implications. Front Behav Neurosci 18:1420028.

62. Ibrahim KM, Massaly N, Yoon HJ, et al. (2024) Dorsal hippocampus to nucleus accumbens projections drive reinforcement via activation of accumbal dynorphin neurons. Nat Commun 15:750.
63. Castro DC, Bruchas MR (2019) A motivational and neuropeptidergic hub: Anatomical and functional diversity within the nucleus accumbens shell. Neuropsychopharmacology 102:529–552.
64. Kravitz AV, Kreitzer AC (2012) Striatal mechanisms underlying movement, reinforcement, and punishment. Physiology 27:167–177.
65. Shen W, Flajolet M, Greengard P, Surmeier DJ (2008) Dichotomous dopaminergic control of striatal synaptic plasticity. Science 321:848–851.
66. Gerfen CR, Engber TM, Mahan LC, Susel Z, Chase TN, Monsma FJ Jr, Sibley DR (1990) D1 and D2 dopamine receptor-regulated gene expression of striatonigral and striatopallidal neurons. Science 250:1429–1432.
67. Macpherson T, Morita M, Hikida T (2014) Striatal direct and indirect pathways control decision-making behavior. Front Psychol 5:1301.
68. Lee SJ, Lodder B, Chen Y, et al. (2021) Cell-type-specific asynchronous modulation of PKA by dopamine in learning. Nature 590:451–456.
69. Steinberg EE, Boivin JR, Saunders BT, Witten IB, Deisseroth K, Janak PH (2014) Positive reinforcement mediated by midbrain dopamine neurons requires D1 and D2 receptor activation in the nucleus accumbens. PLoS One 9:e94771.
70. Morrison JH, Hof PR (1997) Life and death of neurons in the aging brain. Science 278:412–419.
71. Goldman-Rakic PS, Brown RM (1981) Regional changes of monoamines in cerebral cortex and subcortical structures of aging rhesus monkeys. Neuroscience 6:177–187.
72. Kish SJ, Shannak K, Rajput A, Deck JHN, Hornykiewicz O (1992) Aging produces a specific pattern of striatal dopamine loss: Implications for the etiology of idiopathic Parkinson's disease. J Neurochem 58:642–648.
73. Karrer TM, Josef AK, Mata R, Morris ED, Samanez-Larkin GR (2017) Reduced dopamine receptors and transporters but not synthesis capacity in normal aging adults: A meta-analysis. Neurobiol Aging 57:36–46.
74. Dang LC, Castrellon JJ, Perkins SF, Le NT, Cowan RL, Zald DH, Samanez-Larkin GR (2017) Reduced effects of age on dopamine D2 receptor levels in physically active adults. NeuroImage 148:123–129.
75. Fearnley JM, Lees AJ (1991) Ageing and Parkinson's disease: Substantia nigra regional selectivity. Brain 114:2283–2301.
76. Rodriguez M, Morales I, Rodriguez-Sabate C, Sanchez A, Castro R, Brito JM, Sabate M (2014) The degeneration and replacement of dopamine cells in Parkinson's disease: The role of aging. Front Neuroanat 8:80.
77. Berry AS, Jagust WJ, Hsu M (2019) Age-related variability in decision-making: Insights from neurochemistry. Cogn Affect Behav Neurosci 19:415–434.
78. Eppinger B, Schuck NW, Nystrom LE, Cohen JD (2013) Reduced striatal responses to reward prediction errors in older compared with younger adults. J Neurosci 33:9905–9912.
79. Dreher J-C, Meyer-Lindenberg A, Kohn P, Berman KF (2008) Age-related changes in midbrain dopaminergic regulation of the human reward system. Proc Natl Acad Sci USA 105:15106–15111.

80. Tryon VL, Baker PM, Long JM, Rapp PR, Mizumori SJM (2020) Loss of sensitivity to rewards by dopamine neurons may underlie age-related increased probability discounting. Front Aging Neurosci 12:49.
81. Lei HC, Parker KE, Kuo C-C, Yuede CM, McCall JG, Imai S-I (2025) Aging reduces motivation through decreased Bdnf expression in the ventral tegmental area. Mol Psychiatry.
82. Cools R, Frank MJ, Gibbs SE, Miyakawa A, Jagust W, D'Esposito M (2009) Striatal dopamine predicts outcome-specific reversal learning and its sensitivity to dopaminergic drug administration. J Neurosci 29:1538–1543.
83. Rutledge RB, Lazzaro SC, Lau B, Myers CE, Gluck MA, Glimcher PW (2009) Dopaminergic drugs modulate learning rates and perseveration in Parkinson's patients in a dynamic foraging task. J Neurosci 29:15104–15114.
84. Chowdhury R, Guitart-Masip M, Lambert C, Dayan P, Huys Q, Düzel E, Dolan RJ (2013) Dopamine restores reward prediction errors in old age. Nat Neurosci 16:648–653.
85. Moustafa AA (2011) Levodopa enhances reward learning but impairs reversal learning in Parkinson's disease patients. Front Hum Neurosci 4:240.
86. Lodder B, Kamath T, Savenco E, et al. (2025) Absolute measurement of fast and slow neuronal signals with fluorescence lifetime photometry at high temporal resolution. Neuron 113:3554–3566.e7.
87. Chen Y, Saulnier JL, Yellen G, Sabatini BL (2014) A PKA activity sensor for quantitative analysis of endogenous GPCR signaling via 2-photon FRET-FLIM imaging. Front Pharmacol 5:56.
88. Mastro K, Lee W-C, Wang W, Meng S, Stevens B, Sabatini B (2024) Developmental maturation of frontal cortical circuits impacts decision-making. *bioRxiv*. doi:10.1101/2024.05.24.595609.
89. Beron CC, Neufeld SQ, Linderman SW, Sabatini BL (2022) Mice exhibit stochastic and efficient action switching during probabilistic decision making. Proc Natl Acad Sci USA 119:e2113961119.
90. Linker KE, Duran-Laforet V, Ollivier M, Yu X, Schafer DP, Khakh BS (2025) Aging in mice alters regionally enriched striatal astrocytes. Nat Commun 16:8496.
91. Long KLP, Muroy SE, Sorooshyari SK, Ko MJ, Jaques Y, Mohan K, Sudmant P, Kaufer D (2025) Transcriptomic profiles of susceptibility and resilience to stress in the amygdala and hippocampus of male rats. Neurobiol Stress 38:100754.
92. Kilfeather P, Khoo JH, Wagner K, Liang H, Caiazza MC, An Y, Zhang X, Chen X, Connor-Robson N, Shang Z, Wade-Martins R (2024) Single-cell spatial transcriptomic and translatomic profiling of dopaminergic neurons in health, aging, and disease. Cell Rep 43:113784.

# VI. Supplementary

## A. Statistical Supplementary

***All statistical tests, comparisons, and their presentation status in the main text***

**Key**

| Symbol / Shading | Meaning |
|---|---|
| ✓ (green highlight) | Comparison selected for presentation in main text or figure |
| n.s. | Not significant — reported explicitly as null result |
| No highlight | Significant but not the primary comparison highlighted in the paper |
| Not presented | Test run for completeness; not reported in main text or figure |

### Figure 1. Age-dependent differences in acquisition of deterministic task structure

***n = 8, 27, 9, 17, 8, 11 mice for age groups 0–2, 2–4, 4–6, 6–12, 12–18, 18–24 months. All pairwise tests: Tukey HSD post-hoc following one-way ANOVA unless stated.***

#### Fig 1b. Acquisition speed (sessions to criterion)

**One-way ANOVA: F(5,74) = 17.06, p < 0.0001**

| Comparison | Mean Diff (sessions) | p-adj (Tukey) | Presented in Paper |
|---|---|---|---|
| 0–2 vs 12–18 | 7.88 | < 0.001 | |
| 0–2 vs 18–24 | 9.03 | < 0.001 | |
| 0–2 vs 4–6 | 4.26 | 0.010 | ✓ (selected) |
| 0–2 vs 6–12 | 6.89 | < 0.001 | |
| 18–24 vs 2–4 | -6.03 | < 0.001 | |
| 18–24 vs 4–6 | -4.76 | 0.001 | |

| | | | |
|---|---|---|---|
| 12–18 vs 2–4 | -4.88 | < 0.001 | |
| 12–18 vs 4–6 | -3.61 | 0.045 | ✓ (selected) |
| 2–4 vs 6–12 | 3.89 | < 0.001 | |

**Fig 1c — Rewarded and unrewarded trials across learning phases**

**Early phase. Unrewarded trials [ANOVA: F(5,74) = 6.37, p < 0.001]**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 0–2 vs 18–24 | 114.5 | 0.001 | |
| 12–18 vs 18–24 | 95.3 | 0.007 | |
| 18–24 vs 2–4 | -103.7 | < 0.001 | ✓ |
| 18–24 vs 4–6 | -114.0 | < 0.001 | |
| 18–24 vs 6–12 | -96.7 | 0.001 | |

**Early phase. Rewarded trials [ANOVA: F(5,74) = 7.80, p < 0.001]**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 0–2 vs 18–24 | 214.05 | 0.0004 | |
| 18–24 vs 2–4 | -199.2167 | < 0.001 | ✓ |
| 18–24 vs 4–6 | -229.1636 | < 0.001 | |

| | | | |
|---|---|---|---|
| 18–24 vs 6–12 | -148.0667 | 0.0078 | |

**Middle phase. Unrewarded trials [ANOVA: F(5,74) = 3.62, p = 0.006]**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 0–2 vs 18–24 | 121.9 | 0.042 | ✓ |
| 18–24 vs 4–6 | -148.8 | 0.002 | ✓ |

**Middle phase. Rewarded trials [ANOVA: F(5,74) = 1.14, p = 0.345]**

*Not significant. Not presented.*

**Late phase. Rewarded and Unrewarded trials**

*Late rewarded: ANOVA F(5,70) = 0.84, p = 0.526 — not significant. Late unrewarded: ANOVA F(5,70) = 3.29, p = 0.010 — significant overall but no individual Tukey comparisons reached significance. Consistent with convergence of age groups in later learning. Not highlighted.*

**Fig 1d. Reward rate across learning phases**

**Early phase [ANOVA: F(5,132) = 7.48, p < 0.001]**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 0–2 vs 18–24 | -0.262 | < 0.001 | |
| 12–18 vs 18–24 | -0.194 | 0.005 | |
| 18–24 vs 2–4 | 0.213 | < 0.001 | ✓ |
| 18–24 vs 4–6 | 0.262 | < 0.001 | |
| 18–24 vs 6–12 | 0.194 | 0.001 | |

**Middle phase [ANOVA: F(5,284) = 1.82, p = 0.110]**

*Not significant. Not presented.*

**Late phase [ANOVA: F(5,221) = 1.16, p = 0.330]**

*Not significant. Not presented.*

### Fig 1e. Switch probability across learning phases

*Early: F(5,132) = 0.97, p = 0.439 | Middle: F(5,284) = 0.75, p = 0.585 | Late: F(5,221) = 0.86, p = 0.510. No significant differences at any phase. Reported in text as null result confirming preserved exploratory switching.*

### Fig 1f. Trial duration across learning phases

**Early phase [ANOVA: F(5,132) = 4.74, p < 0.001]**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 12–18 vs 2–4 | -0.211 | 0.022 | ✓ |
| 2–4 vs 4–6 | 0.232 | 0.004 | |
| 4–6 vs 6–12 | -0.223 | 0.013 | |

**Middle phase [ANOVA: F(5,284) = 9.46, p < 0.001]**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 0–2 vs 12–18 | 0.109 | 0.007 | |
| 0–2 vs 18–24 | 0.111 | 0.004 | |
| 0–2 vs 4–6 | 0.088 | 0.036 | |
| 12–18 vs 2–4 | -0.106 | < 0.001 | ✓ |
| 18–24 vs 2–4 | -0.108 | < 0.001 | ✓ |
| 18–24 vs 6–12 | -0.065 | 0.048 | |
| 2–4 vs 4–6 | 0.085 | < 0.001 | |

**Late phase [ANOVA: F(5,221) = 24.95, p < 0.001]**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 0–2 vs 12–18 | 0.110 | < 0.001 | ✓ |
| 0–2 vs 18–24 | 0.087 | < 0.001 | ✓ |
| 0–2 vs 4–6 | 0.075 | < 0.001 | |
| 0–2 vs 6–12 | 0.072 | < 0.001 | |
| 12–18 vs 2–4 | -0.105 | < 0.001 | |
| 18–24 vs 2–4 | -0.083 | < 0.001 | |
| 2–4 vs 4–6 | 0.070 | < 0.001 | ✓ |
| 2–4 vs 6–12 | 0.067 | < 0.001 | |

***Trial duration differences in figure highlight that duration increases early but stabilises across older age groups. Key interpretation: duration does not track error rates, indicating preserved task engagement.***

## Figure 2. Aging impairs reinforcement-dependent updating during probabilistic decision-making

***n = 8, 24, 11, 15, 8, 9 mice for age groups 0–2, 2–4, 4–6, 6–12, 12–18, 18–24 months.***

### Fig 2a. Rewarded trials

**One-way ANOVA: F(5,69) = 3.76, p = 0.005 | Kruskal-Wallis: H = 15.07, p = 0.010**

***Significant differences driven by 0–2 month group having more rewarded trials than all others.***

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 0–2 vs 2–4 | -59.5944 | 0.0047 | ✓ |

### Fig 2a. Unrewarded trials

**One-way ANOVA: F(5,69) = 12.01, p < 0.001 | Kruskal-Wallis: H = 29.60, p < 0.001**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|

| | | | |
|---|---|---|---|
| **0–2 vs 12–18** | **-51.6** | **0.002** | |
| **0–2 vs 2–4** | **-59.0** | **< 0.001** | ✓ |
| **0–2 vs 4–6** | **-74.2** | **< 0.001** | |
| **0–2 vs 6–12** | **-62.5** | **< 0.001** | |
| **18–24 vs 2–4** | **-37.5** | **0.004** | |
| **18–24 vs 4–6** | **-52.6** | **< 0.001** | ✓ |
| **18–24 vs 6–12** | **-41.0** | **0.003** | ✓ |

### Fig 2b. Trial duration

**One-way ANOVA: F(5,69) = 3.99, p = 0.003 | Kruskal-Wallis: H = 32.19, p < 0.001**

**Mixed linear model confirms age-related differences (all groups vs 0–2 months significant). Key finding: duration stabilises from 4 months onward — not increasing with error rates.**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| **0–2 vs 12–18** | **0.051** | **< 0.001** | |
| **0–2 vs 18–24** | **0.049** | **< 0.001** | |
| **0–2 vs 4–6** | **0.054** | **< 0.001** | |
| **0–2 vs 6–12** | **0.052** | **< 0.001** | |

| | | | |
|---|---|---|---|
| **12–18 vs 2–4** | **-0.029** | **0.026** | |
| **18–24 vs 2–4** | **-0.028** | **0.028** | |
| **2–4 vs 4–6** | **0.032** | **0.003** | ✓ |
| **2–4 vs 6–12** | **0.031** | **0.001** | |

### Fig 2c. High-port selection probability

**One-way ANOVA: F(5,69) = 3.99, p = 0.003**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| **18–24 vs 2–4** | **-0.057** | **0.009** | ✓ |
| **18–24 vs 4–6** | **-0.076** | **0.001** | |
| **18–24 vs 6–12** | **-0.047** | **0.082** | **n.s.** |

### Fig 2d, 2e, 2f, 2g. Non-significant measures

***The following measures did not yield significant results. Directional trends are discussed in the main text as consistent with the overall aging narrative, with the caveat that these may reflect limited statistical power given current sample sizes.***

| Measure | ANOVA statistic | p-value | Status |
|---|---|---|---|
| **Pre-switch pHigh** | **F(5,69) = 2.14** | **0.071** | **Not presented** |
| **Post-reversal recovery** | **F(5,68) = 0.70** | **0.623** | **Not presented** |
| **Post-reversal switch prob.** | **F(5,69) = 1.20** | **0.319** | **Not presented** |
| **Lose-switch probability** | **F(5,69) = 1.58** | **0.177** | **Not presented** |
| **Win-stay probability** | **F(5,69) = 0.86** | **0.512** | **Not presented** |

## Figure 3. Model reveals age-dependent reduction in reinforcement sensitivity and prolonged persistence of prior evidence

***n = 8, 24, 11, 15, 8, 10 mice for age groups 0–2, 2–4, 4–6, 6–12, 12–18, 18–24 months.***

**Summary of RFLR model parameter statistics**

| Parameter | ANOVA statistic | p-value | Presentation status |
|---|---|---|---|
| Alpha (persistence) | F(5,70) = 1.43 | 0.224 | Not significant, not highlighted |
| Beta (reinf. sensitivity) | F(5,70) = 2.68 | 0.029 | ✓ Significant, highlighted |
| Tau (integration timescale) | F(5,70) = 3.10 | 0.014 | ✓ Significant, highlighted |
| Model fit quality (LLS) | F(5,70) = 0.87 | 0.504 | Not significant, reported as validation |

**Fig 3d. Beta (reinforcement sensitivity) post-hoc comparisons**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 18–24 vs 4–6 | 0.704 | 0.023 | ✓ |

**Fig 3e. Tau (integration timescale) post-hoc comparisons**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 18–24 vs 2–4 | -2.110 | 0.007 | ✓ |
| 18–24 vs 4–6 | -2.238 | 0.017 | ✓ |

**Fig 3f. Model fit quality**

**One-way ANOVA: F(5,70) = 0.87, p = 0.504**

**Spearman correlation with age: ρ = −0.037, p = 0.753**

***Reported as validation that the model provides a comparably good description of behavior across all age groups. No significant differences, this is the desired result.***

## Figure 4. Dopamine transients encode reward history in deterministic and probabilistic versions of the task

***No formal inferential statistics presented for Figure 4. Values reported as mean ± SEM. Differences between age groups and across sequence positions are described qualitatively in the main text.***

**Peak z-score values (mean ± SEM)**

**100/0 contingency, Aligned to consumption**

| Group | Sequence 1 | Sequence 2 | Sequence 3 |
|---|---|---|---|
| P120 | 3.82 ± 0.05 | 3.57 ± 0.05 | 3.34 ± 0.06 |
| P60 | 3.36 ± 0.06 | 2.84 ± 0.07 | 2.74 ± 0.07 |

**100/0 contingency,  Aligned to cue**

| Group | Sequence 1 | Sequence 2 | Sequence 3 |
|---|---|---|---|

| | | | |
|---|---|---|---|
| **P120** | **1.82 ± 0.05** | **2.39 ± 0.05** | **2.34 ± 0.05** |
| **P60** | **1.91 ± 0.05** | **2.28 ± 0.05** | **2.22 ± 0.05** |

**100/0 contingency, Loss sequence aligned to consumption**

| Group | Sequence 1 | Sequence 2 | Sequence 3 |
|---|---|---|---|
| **P120** | **−1.27 ± 0.02** | **−1.31 ± 0.03** | **−1.26 ± 0.03** |
| **P60** | **−1.47 ± 0.03** | **−1.41 ± 0.03** | **−1.36 ± 0.04** |

**80/20 contingency, Aligned to consumption**

| Group | Sequence 1 | Sequence 2 | Sequence 3 |
|---|---|---|---|
| **P120** | **4.05 ± 0.05** | **3.28 ± 0.04** | **3.11 ± 0.05** |
| **P60** | **3.60 ± 0.05** | **2.50 ± 0.05** | **2.36 ± 0.06** |

## Figure 5. Aging produces a downward shift in tonic dopamine baseline while preserving relative phasic signalling

***n = 10, 6, 14 recordings from 4, 3, and 5 mice at 3–6, 18–21, and 24+ months respectively. Several animals contributed recordings at multiple timepoints (longitudinal design).***

### Fig 5d. Baseline lifetime (modal signal)

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| **18–21 mo vs 24+ mo** | **-0.201** | **0.001** | ✓ |
| **24+ mo vs 3–6 mo** | **0.1288** | **0.013** | ✓ |
| **18–21 mo vs 3–6 mo** | **-0.0718** | **0.269** | **n.s.** |

### Fig 5e. Peak lifetime

**One-way ANOVA: F(2,21) = 9.47, p = 0.001**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| **18–21 mo vs 24+ mo** | **-0.161** | **0.001** | ✓ |
| **24+ mo vs 3–6 mo** | **0.107** | **0.013** | ✓ |
| **18–21 mo vs 3–6 mo** | **-0.054** | **0.3287** | **n.s.** |

### Fig 5f. Phasic amplitude (peak minus baseline)

***No significant differences between age groups (all p > 0.05).***

**One-way ANOVA: F(2,21) = 2.97, p = 0.073**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 18–21 mo vs 24+ mo | 0.0393 | 0.063 | n.s. |
| 24+ mo vs 3–6 mo | -0.0217 | 0.3054 | n.s. |
| 18–21 mo vs 3–6 mo | 0.0176 | 0.5098 | n.s. |

### Fig 5g. Autofluorescence control

***Fluorescence intensity scales linearly with excitation power across all animals ($R^2$ = 0.977–0.998 for representative 24 month animals). Confirms sensor-dependent emission.***

## Figure 6. Tonic dopamine baseline does not predict motivational effort or hunger state

***Fig 6b–c: n = 34 recordings from 13 mice. Fig 6d: n = 8 animals. Fig 6e: see per-group session counts below.***

### Fig 6b. Breakpoint vs body weight

**Mixed effects model: β = −6.56, SE = 4.41, p = 0.137 (not significant)**

**One-way ANOVA by weight bin: F(3,30) = 4.65, p = 0.009 | Kruskal-Wallis: H = 12.90, p = 0.005**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 80–85% vs 90–95% | -151.2 | 0.009 | ✓ |
| 80–85% vs 85–90% | -67.3 | 0.249 | n.s. |

### Fig 6c. Breakpoint across age groups

**One-way ANOVA: F(2,30) = 1.38, p = 0.268 | Kruskal-Wallis: H = 4.72, p = 0.194**

| Comparison | Mean Diff (trials) | p-adj (Tukey) | Presented |
|---|---|---|---|
| 3 mo vs 12 mo | -81.5 | 0.234 | n.s. |
| 3 mo vs 18 mo | -24.5 | 0.969 | n.s. |
| 3 mo vs 24 mo | -59.0 | 0.627 | n.s. |

### Fig 6d. Dopamine lifetime vs PR breakpoint

| Measure | Correlation | p-value | Conclusion |
|---|---|---|---|
| Baseline DA vs Breakpoint | r = 0.005 | 0.991 | Not significant |
| Peak DA vs Breakpoint | r = -0.183 | 0.665 | Not significant |

### Fig 6e. Baseline dopamine lifetime: sated vs hungry

***3–6 months: n = 10 sated, 12 hungry and p = 0.176 (Mann–Whitney U test)***

***18–21 months: n = 10 sated, 7 hungry and p = 0.364 (Mann–Whitney U test)***

*24+ months: n = 12 sated, 3 hungry and p = 0.136 (Mann–Whitney U test)*

*No significant difference between hunger states in any age group. Reported as null result confirming that tonic dopamine baseline is not acutely modulated by hunger.*

## Figure 7. Aging drives transcriptional changes in striatal interneurons without altering canonical marker expression

*F(*

### Fig 7e. Astrocyte contamination marker expression across age

**Mann-Whitney U test, two-sided. Young n = 436, Aged n = 105.**

| Gene | p-value | Sig |
|---|---|---|
| Slc1a2 | 2.26e-12 | *** |
| Vim | 1.93e-02 | * |
| S100b | 2.43e-02 | * |
| Aldoc | — | n.s. |

### Fig 7g. Interneuron subtype marker expression across age (Striatal Interneuron cluster)

**Mann-Whitney U test, two-sided. Young n = 438, Aged n = 726.**

| Gene | p-value | Significance |
|---|---|---|
| Sst | 3.72e-05 | *** |
| Npy | 4.59e-06 | *** |
| Kcnq3 | 4.58e-02 | * |
| Cnr1 | 3.27e-02 | * |
| Chat | 1.0000 | n.s. |
| Pvalb | 6.02e-01 | n.s. |
| Hcn1, Gabra1, Th | — | n.s. |

**Sst and Npy show significant age-related reductions in the striatal interneuron population. Canonical cholinergic (Chat), fast-spiking (Pvalb), and other interneuron markers (Hcn1, Gabra1, Th) are not significantly altered by age. Astrocyte markers Slc1a2, Vim, and S100b show elevated expression in the aged subcohort; Aldoc does not differ significantly.**