Context-specific configuration of orthogonal integrator dynamics for flexible foraging decisions

Lyle Kingsbury[1,2,*], Grace Zhang[1,2], Juan Ignacio Sanguinetti-Scheck[1,2,3], Naoshige Uchida[1,2,*]

1. Dept. of Molecular and Cellular Biology, Harvard University, Cambridge, MA, USA
2. Center for Brain Science, Harvard University, Cambridge, MA, USA
3. Current address: Dept. of Psychology, University of Pennsylvania, Philadelphia, PA, USA

* Correspondence: lylekingsbury@fas.harvard.edu, uchida@mcb.harvard.edu

Abstract

The capacity to adapt behavior across different contexts is a fundamental feature of intelligence and is crucial for animal life, yet the mechanisms by which neural circuits are contextually reconfigured to alter their function are poorly understood. As a universal, ecological behavior, foraging provides a natural window into the neural basis of decision making and cognitive flexibility. We developed complementary free-moving and virtual tasks in which mice make foraging decisions and switch between context-specific decision strategies by toggling distinct configurations of integrator dynamics for evidence accumulation. Neural recordings reveal brain-wide activity states recruited by foraging environments which organize divergent dynamics subspaces and orthogonal coding of context-specific decision variables in separate neural subpopulations. In the dorsal frontal cortex, this coding structure is preserved across experimental settings, supports rapid dynamics reconfiguration for volitional mode switching, and is functionally necessary for decision flexibility.

Introduction

The capacity to adapt cognition and behavior across different contexts is a core feature of natural intelligence, and for many animals is crucial for surviving in dynamic natural environments. A fundamental goal of neuroscience is to understand how such cognitive processes arise from cellular and circuit mechanisms in the brain. In mammals, the frontal cortex is critical for cognitive flexibility, which is thought to involve regulation of network modes for context-specific information processing and behavioral control.[1–5] Theory and modeling work give insight into how network modes, such as attractor states, may be organized by properties of cell physiology, connectivity, and circuit architecture to perform context-specific computations.[6–10]

Despite these advances, experimental progress in linking population dynamics-level models with their cellular underpinnings has been limited. The cornerstone of investigations in animal models has been rule-based perceptual decision tasks, which seek to elucidate how context signals configure decision processes.[11–15] For these problems, many studies have identified mechanisms that locate context effects upstream of decision dynamics, incorporating context contingencies by attentional shifting,[13,14,16–18] sensorimotor mapping,[12] or altering how sensory inputs perturb decision activity.[11,15,19] Population coding theories and other studies also raise the possibility that decision activity itself could be parsed into different functional modes (e.g. encoded in orthogonal subspaces) to enable flexible selection across contexts by control signals.[7,10,20–25]

Neural solutions for flexible computation, and their cellular-level instantiations, are in fact likely to be diverse, task-dependent, and even idiosyncratic[8,19]. For example, while context-specific sensory tuning can readily solve perceptual decision tasks, the picture of how varied mechanisms are used across problem domains, including for social or economic decision problems, remains muddled. Arguably, this knowledge gap persists because of a narrow focus on reduced task designs that, while in line with classical approaches, can be difficult for animals to learn and are often peripheral to their ethological repertoire.[26,27] Indeed, it is unclear to what extent neural mechanisms tuned for highly trained choice tasks may generalize, or whether (and what) computational and neural solutions may be at play in more complex or naturalistic settings.

In response to such critiques, recent years have seen a conceptual shift in neuroscience toward natural behaviors that preserve the native regime of brain functions shaped by evolution to solve ecological problems.[26,27] Foraging – how to source and harvest resources from the environment – is a challenge faced across the animal kingdom.[28,29] Animals routinely make decisions about where to forage and when to leave depleting resource locations ("patches") in search of alternatives (the "patch leaving problem"),[30,31] and can adjust behavioral strategies across different environments or situations. In contrast to alternative-choice tasks, patch foraging is immediately intuitive for many animal subjects and requires relatively little laboratory training. Inspired by this, recent studies have begun to uncover the neural basis of foraging behaviors;[32,33] so far, these indicate that patch leaving decisions can solved through a process akin to bounded evidence accumulation, linking natural cognition to classical models of decision making.[34–38] The flexibility and ecological grounding of foraging thus open a new window onto the algorithmic logic and neural basis of context-dependent decision making that meets animal subjects closer to "where they are."

We developed a foraging task in which mice flexibly toggle patch leaving decision strategies across environment contexts using distinct configurations of a common evidence

accumulation process (i.e. a neural integrator). Using an equivalent virtual task for large-scale neural recordings, we find that foraging environments recruit distinct, brain-wide activity states. These states reconfigure neural dynamics in a subset of brain areas, organizing context-specific integrator activity for foraging decisions into orthogonal coding dimensions realized in separate neuronal subpopulations. In the dorsal frontal cortex (dFC), this coding structure is shaped by experience and preserved across task settings, even when unconstrained mice engage in repeated volitional context switches. Inactivation experiments demonstrate causal necessity of dFC for context-specific configuration of decision modes. Together, our results reveal a cellular-level neural basis for context-dependent decision making in orthogonal integrator dynamics, isolate a functional bottleneck for circuit configurability in dFC, and establish foraging as an ecologically grounded, tractable alternative to classical approaches for investigating complex cognition.

Results

A patch foraging task for context-dependent decision making in mice

In patch foraging,[30,31] animals encounter resource patches that deplete over time and must decide when to leave in search of alternatives (**Figure 1A**). We hypothesized that mice foraging in environments with different resource distributions could learn and deploy context-specific patch leaving decision strategies. We developed a task in which mice freely navigate a split chamber for water rewards from multiple ports ("patches," **Figures 1B-C**). Mice nose-poke a port to initiate a trial and immediately (t = 0 sec) receive a 2μL water droplet, after which they can stay for more rewards. While patches in the "Stochastic" (Stc, **Figure 1B**, top) environment yield probabilistic reward each second with decaying probability (**Figure S1A**), patches in the "Deterministic" (Det, **Figure 1B**, bottom) environment always (100% probability) yield three rewards in a fixed sequence (t = 0, 1, 2 sec). 15% of trials in the Stc environment were programmed with this Det reward sequence ("Stc:012" trials) to create reward-matched trials across contexts for analysis.

Mice first foraged in the Det environment/context to learn a decision strategy for the Det distribution. In Phase 1 (**Figure 1C**, left), mice quickly learned to harvest water, performing an average of 64 trials on session 1 with the first within three minutes (**Figures S1B-C**). Though initially impatient, mice rapidly learned to collect all rewards by the end of the first session (80% of trials, **Figure S1D**). Examination of patch residence time (PRT, how long mice stay at a patch) distributions across Phase 1 revealed calibration of the decision process over time (**Figures S1E-G**) to the Det reward return.

Mice then proceeded to Phase 2 (**Figure 1C**, middle), in which they explored both Det and Stc contexts once per session in blocks, allowing them to learn context-specific decision strategies. In Phase 3, the door separating the chambers was left open and mice could roam freely and switch contexts volitionally (**Figure 1C**, right). Overall, mice collected more rewards per trial on Det patches and preferred them (**Figure S1H-I**), possibly reflecting risk aversion or small value difference. Still, mice explored both contexts extensively, switching 2-3 times per minute and sampled all patches multiple times in every session (**Figure S1J-K**).

Mice exhibit context-specific decision strategies for different foraging environments

Figure 1

A flexible foraging task for context-dependent decision making

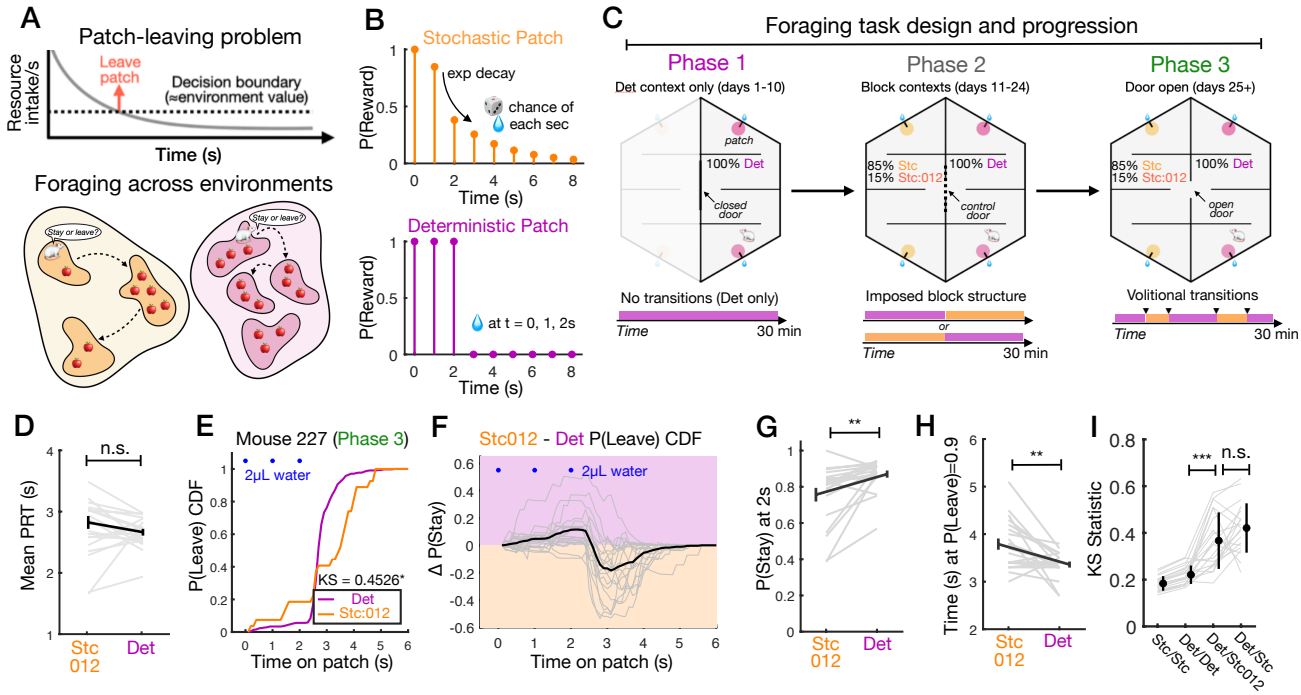

A configurable integrator explains context-specific foraging decisions

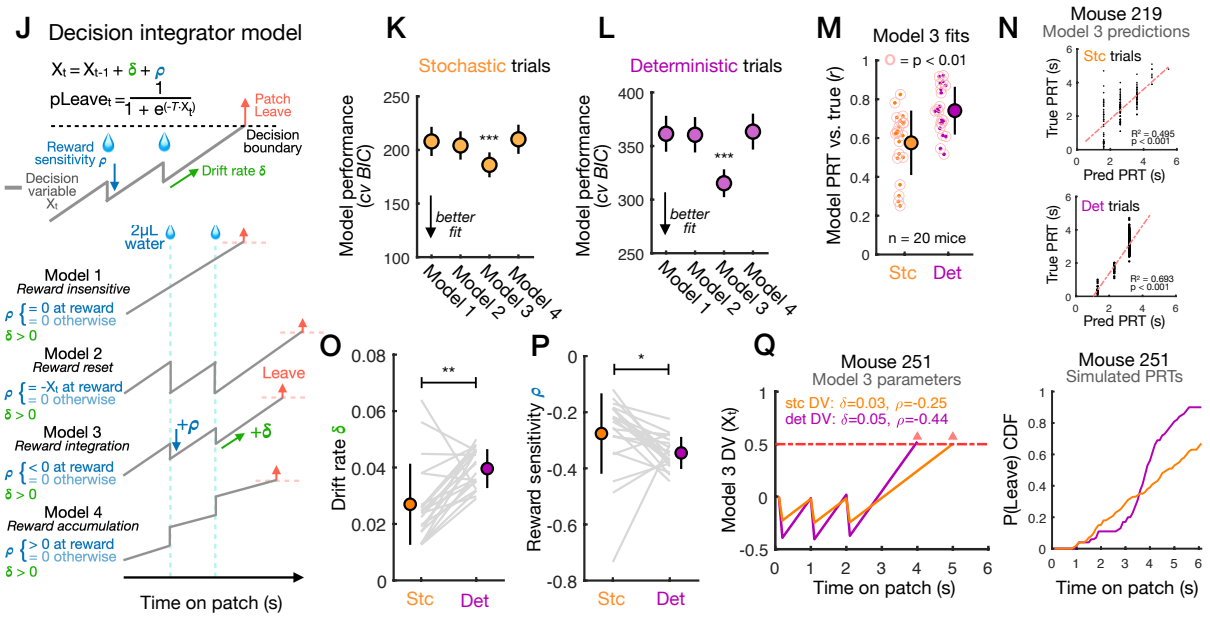

Figure 1: Mice use rapidly flexible foraging strategies across different environments.

A) Left, Illustration of the classical patch-leaving problem (when to leave a patch with depleting resources); animals must make patch leaving decisions across two environment contexts with distinct resource distributions (Deterministic "Det" or Stochastic "Stc").
B) Reward returns (2μL water drops with probability each second) for Stc- (orange) and Det-environment (purple) patches.
C) Progression through flexible foraging task, from Det environment (Phase 1, 7-10 days) to block structure (Phase 2, 14 days) and unconstrained volitional transitions (door open, Phase 3).
D) Mean patch residence times (PRT, mean ± s.d.) for Det and Stc:012 patch trials (SR test, p = 0.218, n = 20 mice).
E) CDF of leaving probabilities (pLeave) over time for Det and Stc:012 (Det reward return in the Stc context) patches from example Mouse 227 (KS test, p=3.25e-5).
F) Difference in stay probability (pStay) on Det vs. Stc:012 patches for each mouse (gray lines) and average across mice (black line, n = 20 mice).
G) Comparison of pStay at t = 2 seconds for Det vs. Stc:012 trials (mean ± s.d., SR test, p = 0.0013, n = 20 mice).
H) Comparison of average time on patch when pLeave = 0.9 for Det vs. Stc:012 trials (mean ± s.d., SR test, p = 0.0038, n = 20 mice).
I) Comparison of KS statistics comparing different distributions of trial type PRTs (SR tests, p = 2.19e-4, Det/Det vs. Det/Stc:012, p = 0.0522, Det/Stc:012 vs. Det/Stc, n = 20 mice).
J) Schematic of drift-diffusion integrator model for patch-leaving decisions (top). Decision variable (DV) X_t is incremented each time step by drift rate δ and perturbed by an amount ρ upon reward receipt, accumulating to a decision boundary to trigger a leaving decision. $pLeave_t$ is generated by a logistic function transform of X_t with temperature parameter *T*. Parameters δ and control ρ the integrator process implementing different computations for leaving decisions (Models 1-4).
K) Performance (cross-validated Bayesian information criterion, cv-BIC) of integrator models 1-4 fit to mouse PRT behavior and tested on held-out data for Stc patch trials (SR test between models 1v3, 2v3, 3v4, p=0.857e-5 for each, n = 20 mice, see Methods for details).
L) Model performance as in **(K)** for Det trials (SR test between models 1v3, 2v3, 3v4; p=0.857e-5 for each, n = 20 mice).
M) Performance of "reward integration" Model 3 (correlation *r* between predicted and true PRTs) for Stc and Det patch trials. Each mouse is one dot; red circles indicate independent statistical significance (p<0.01) of regression on predicted vs. true PRTs for that animal.
N) Model 3 prediction vs. true PRTs for example mouse 219; red line indicates linear regression.
O) Comparison of fit drift rate parameter δ to Model 3 for Stc vs. Det trial behavior (mean ± s.d., SR test, p = 0.0036, n = 20 mice).
P) Comparison of fit reward sensitivity parameter ρ to Model 3 for Stc vs. Det trial behavior (mean ± s.d., SR test, p = 0.023, n = 20 mice).
Q) Left, Rollout of Model 3 DV for integrator parameterizations fit to Stc or Det trial behavior for example mouse 251. Right, example of pLeave CDFs generated from simulated PRTs using mouse 251 parameters (see Methods), recapitulating the typical pattern of decision behavior observed across animals.

Standard deviation (s.d.); Standard error of mean (SEM); Kolmogorov-Smirnov test (KS test); Wilcoxon signed-rank test (SR test); ***p<0.001; **p<0.01; *p<0.05; n.s. p>0.05.

To characterize context-specific decision strategies, we compared PRTs across different trials. Stc trials included four hidden types with different reward probabilities (**Figure S1A)** which correlated positively with PRTs, confirming that mice use reward information to make leaving decisions (**Figure S2A**). Examining Det vs. Stc:012 trials, we observed no difference in average PRTs (**Figure 1D**). However, examination of PRT distributions by plotting the CDF of leaving probabilities (pLeave) revealed clear differences in decision behavior across contexts. For example, while Mouse 227 typically waited to collect all Det rewards and left shortly after, it did not exhibit this "step-like" pattern on Stc:012 patches, often leaving before the final reward or staying after (**Figure 1E**). This pattern was consistent across mice (**Figure 1F**), and indeed, mice were significantly more likely to collect all rewards on Det vs. Stc:012 trials (**Figure 1G**) and faster to leave on Det trials after collecting the final reward (**Figure 1H**).

To quantify this decision flexibility, we measured the difference between Det and Stc:012 PRT distributions using the Kolmogorov-Smirnov (KS) statistic, which measures the vertical difference in pLeave CDFs at their maximum difference (**Figures 1E, S2B**). We found that PRT distributions for Det vs. Stc:012 trials were significantly more different than subsampled Det trials, despite identical reward structure (**Figure 1I**). Divergence of PRT distributions was detectable early in Phase 2 and increased over time, indicating continued shaping of decision strategies that are maintained with volitional transitions in Phase 3 (**Figures S2C, S1G**). While this suggests that mice learn distinct strategies, another possibility is that they rely on recent within-context outcomes to continually update their decision process. To test this, we analyzed PRTs on trial sequences in which mice performed a Stc:012 followed by a Det trial (**Figure S2D**). We observed a similar difference in PRT distributions (**Figure S2E**), even for paired trial comparisons (**Figure S2F).** This shows that mice immediately adapt to context transitions and exhibit context-specific behavior without needing to sample recent outcomes. Taken together, these results demonstrate that mice calibrate decision behavior to environment reward structure and can flexibly and rapidly deploy context-specific patch leaving decisions.

Configurations of a common integrator process explain decision flexibility

Patch leaving decisions in mice are well described by a bounded accumulator process instantiated by neural dynamics that integrate time and reward as evidence for, or against, patch leaving (**Figure 1J**).[34,37] We hypothesized that an integrator process may explain decision behavior in the present task and account for contextual flexibility. To test this, we modeled leaving decisions as a drift-diffusion decision variable (DV) which increments each time step (by drift rate $\boldsymbol{\delta}$) and is offset on reward receipt (by reward sensitivity $\boldsymbol{\rho}$), generating a leaving probability pLeave_t (**Figure 1j, top).** In this model, the parameters $\boldsymbol{\delta}$ and $\boldsymbol{\rho}$ govern decision behavior and, as in classical models of perceptual decisions,[35,36] can be thought of as weights of evidence for or against the leaving decision. Different configurations produce different algorithms for solving the patch leaving problem (**Figure 1J, bottom**). For example, while setting $\boldsymbol{\rho} = 0$ yields a reward-insensitive integrator that tracks time on the patch (equivalent to the "giving up time" heuristic in foraging theories),[30] setting $\boldsymbol{\rho} > 0$ yields a counting strategy that accumulates reward quantity.

In principle, mice could use different algorithms for Stc vs. Det leaving decisions. To test which algorithm best explained behavior, we fit each model to empirical PRTs and measured predictive performance on held-out trials. Consistent with our previous report,[37] we found that the

"reward integration" algorithm ("Model 3") outperformed others (**Figures 1K-L**) and was the best model for all mice for both patch types. This suggests that mice do not use simple heuristics but instead integrate time and reward as decision evidence. Despite distinct patterns of behavior (**Figures 1E-F**), the reward integration model captured PRT variability for both patch types consistently across mice (**Figures 1M-N**). As these results suggest a common integration algorithm, this raises the possibility that context-specific decision strategies could be driven by different parameter configurations (i.e. $\boldsymbol{\delta}$ and $\boldsymbol{\rho}$) for the integrator process.

To test this, we compared $\boldsymbol{\delta}$ and $\boldsymbol{\rho}$ for models fit to Stc vs. Det patches and found consistent differences in parameterization: while Det trials had significantly larger drift rates $\boldsymbol{\delta}$ (**Figure 1O**), they also had more negative values for reward sensitivity $\boldsymbol{\rho}$ (**Figure 1P**). To see how these configurations could produce the observed behavior, we plotted DVs following a Det reward return from example mice (**Figure 1Q**, left). These produced distinct DV rollouts, and sampling simulated PRTs from them qualitatively replicated the pattern of behavior in most mice (**Figure 1Q,** right). Taken together, these results suggest that context-specific behavior may be driven by different configurations of a common integrator process underlying patch leaving decisions.

An equivalent head-fixed virtual task for wide sampling of neural activity

We next aimed to investigate the neural mechanisms underlying contextual configuration of the decision process. While naturalistic, freely moving (FM) tasks are indispensable for understanding complex behavior,[26,27] constraint with head-fixation can minimize confounds such as unmeasured sensory inputs and movements (**Figure 2A**) and enables wide sampling with electrophysiology.[39] Therefore, we developed a structural equivalent virtual-reality[37] (VR) version of the flexible foraging task (**Figure 2B**, left). Mice run on a treadmill to navigate a virtual corridor and stop at patches – landmarks indicated with visual and odor cues – to collect 2μL water rewards. While stopped, they can receive water each second after an initial reward. As in the FM task, Det patches always yield three rewards in a fixed sequence (at t = 0, 2, 4 sec, **Figure 2B**, right), and 10% of Stc-context trials had this same return ("Stc:024" trials). Otherwise, Stc patches gave rewards each second with decaying probability (**Figure 2B**, **S3A**). Stc and Det patches were contained in distinct virtual environments indicated with visuals on the corridor walls. After habituation (**Figure 2C**), animals foraged in the Det context for 7-10 days (Phase 1), then proceeded to Phase 2, where they foraged in Det and Stc contexts in 20-25-minute blocks once per session with one context transition.

While all mice learned to forage in VR, we observed striking individual variability in task engagement, waiting behavior, and decision strategies, consistent with our previous report.[37] On average, mice performed 70.9 trials/session with a wide range of engagement rates (**Figure S3B**). While mice invested an average of 9.5 sec/trial, PRTs also varied widely, indicating individual differences in patience levels (**Figure S3C**). PRTs on Stc patches were again correlated with reward frequency (**Figure 2D**) and were similar between Det vs. Stc:024 trials (**Figure 2E**).

Despite this variability, like in the FM task, pLeave CDFs revealed a tendency for mice to stay more at early times on Det patches and to over-stay on Stc:024 patches (**Figure 2F**). Analysis of PRT distributions confirmed that decision behavior was more distinct across Det vs. Stc:024 patches than among Det patches (**Figure 2G**), and many individuals (e.g. Mouse 124, **Figure 2H**) showed dramatic contextual shifts in behavior consistent with this group trend.

Figure 2

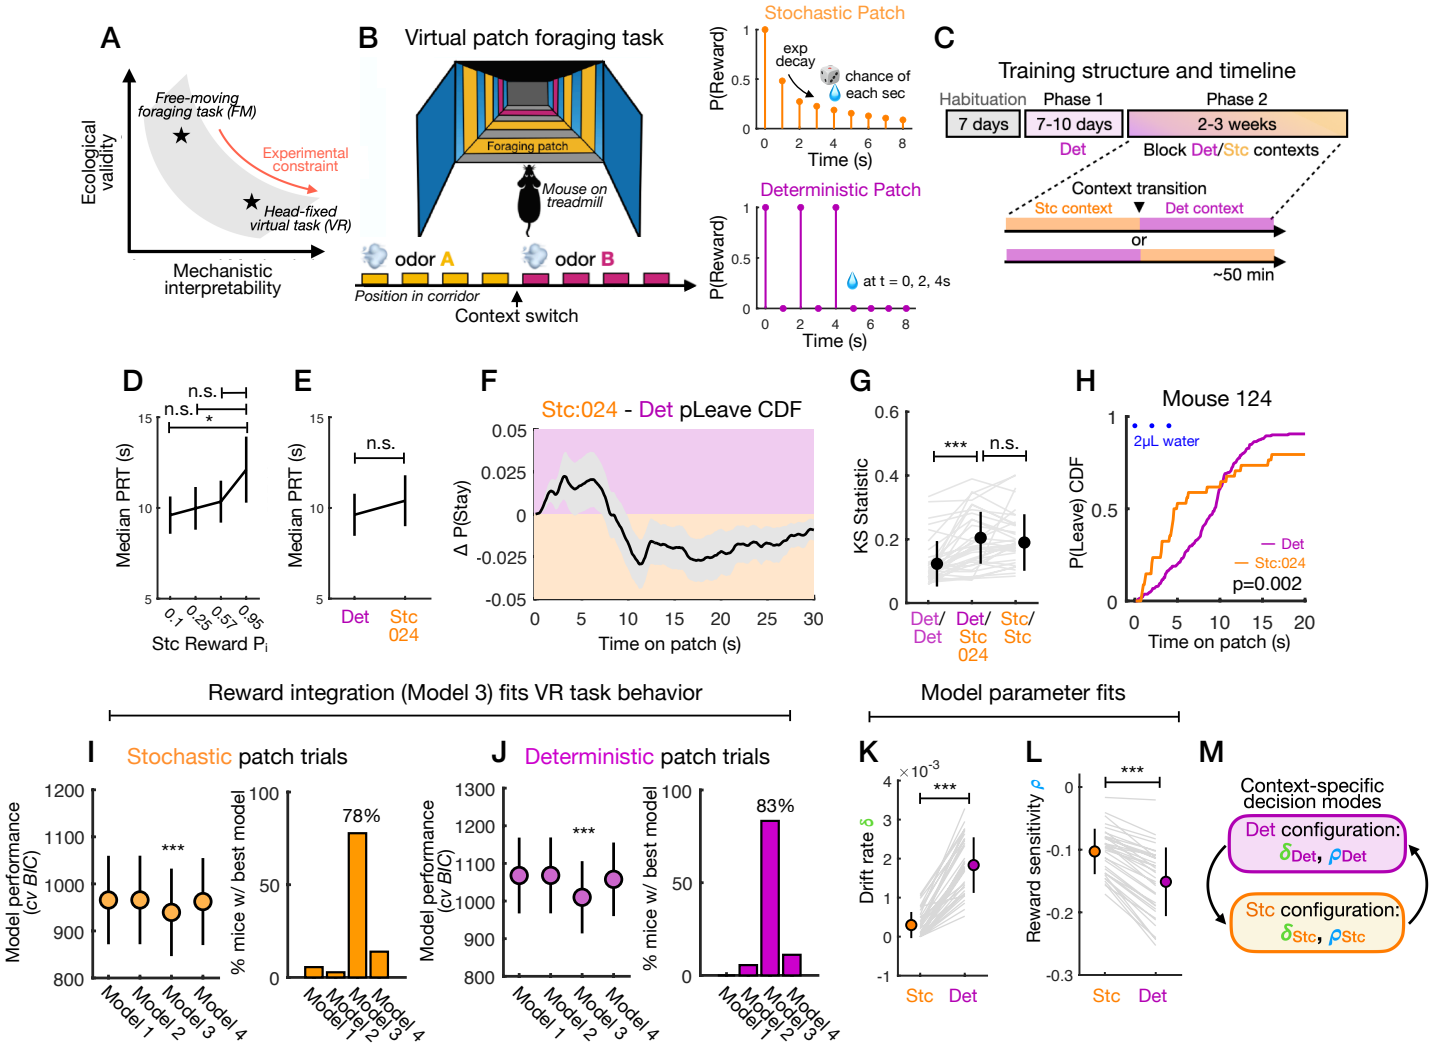
A
Ecological validity
Free-moving foraging task (FM)
Experimental constraint
Head-fixed virtual task (VR)
Mechanistic interpretability
B
Virtual patch foraging task
Foraging patch
Mouse on treadmill
odor A
odor B
Position in corridor
Context switch
Stochastic Patch
P(Reward)
exp decay
chance of each sec
Time (s)
Deterministic Patch
at t = 0, 2, 4s
C
Training structure and timeline
Habituation
Phase 1
Phase 2
7 days
7-10 days
2-3 weeks
Det
Block Det/Stc contexts
Context transition
Stc context
Det context
or
~50 min
D
Median PRT (s)
n.s.
*
Stc Reward Pi
E
n.s.
Det
Stc 024
F
Stc:024 - Det pLeave CDF
Δ P(Stay)
Time on patch (s)
G
KS Statistic
***
n.s.
Det/ Det
Det/ Stc 024
Stc/ Stc
H
Mouse 124
2μL water
P(Leave) CDF
Det
Stc:024
p=0.002
Time on patch (s)
Reward integration (Model 3) fits VR task behavior
Model parameter fits
I
Stochastic patch trials
Model performance (cv BIC)
% mice w/ best model
78%
Model 1
Model 2
Model 3
Model 4
J
Deterministic patch trials
83%
K
Drift rate δ
Stc
Det
L
Reward sensitivity ρ
M
Context-specific decision modes
Det configuration: δDet, ρDet
Stc configuration: δStc, ρStc

Figure 2. Mice exhibit decision flexibility in an equivalent virtual-reality task.

A) Cartoon of "experiment space" tradeoff illustrating how experimental constraint can improve interpretability and mechanistic access at the expense of ecological validity.
B) Left, illustration of structural equivalent virtual-reality (VR) task in which mice do patch foraging for water rewards and encounter stochastic (Stc) and deterministic (Det) patches in different environment contexts cued with visuals and odors. Right, reward returns for Stc and Det patches (also see **Figure S3A**).
C) Training timeline in which mice are habituated to head fixation and running wheel, exposed to the Det context (Phase 1), and forage in block-structured contexts (Phase 2).
D) Mean patch residence times (PRT, mean ± SEM) for Stc patches with different initial reward probability levels (P_i = 0.1, 0.25, 0.57, 0.95, SR tests, p = 0.0315 [0.1 vs. 0.95], p = 0.01746 [0.25 vs. 0.95], p = 0.0996 [0.57 vs. 0.95], n = 40 mice).
E) Comparison of PRTs (mean ± SEM) for Det vs. Stc:024 (Det reward return in Stc context) trials (SR test, p=0.6653, n = 40 mice).
F) Average difference in stay probability on Det vs. Stc:024 patches (mean ± SEM, n = 36 mice).
G) KS statistics for differences in PRT distributions for trial type comparisons (SR tests, p = 2.1661e-7 [Det/Det vs. Det/Stc:024], p = 0.3379 [Det/Stc:024 vs. Det/Stc], n = 36 mice).
H) PRT distributions (CDF of pLeave) over time for Det and Stc:024 patches from from example Mouse 124 (KS test, p=0.002).
I) Performance (cross-validated Bayesian information criterion, cv-BIC) of integrator models fit to mouse PRT behavior and tested on held-out data for Stc patch trials (mean ± SEM, SR test between models 1v3, 2v3, 3v4, p = 2.06e-5, p = 1.92e-5, p = 5.05e-5, n = 36 mice).
J) Model performance and best models as in **(I)** for Det trials (mean ± SEM, SR test between models 1v3, 2v3, 3v4, p = 2.84e-6, p = 5.8543e-7, p = 7.4660e-7, n = 36 mice).
K) Comparison of Model 3 drift rate parameter δ fit to Stc vs. Det trial behavior (mean ± s.d., SR test, p = 1.68e-7, n = 36 mice).
L) Comparison of Model 3 reward sensitivity parameter ρ fit to Stc vs. Det trial behavior (mean ± s.d., SR test, p = 1.68e-7, n = 36 mice).
M) Illustration of toggling between context-specific decision modes. Different configurations of the integrator process captured with parameters δ and ρ produce computations for patch leaving tuned to environment reward returns.

Standard deviation (s.d.); Standard error of mean (SEM); Kolmogorov-Smirnov test (KS test); Wilcoxon signed-rank test (SR test); ***p<0.001; **p<0.01; *p<0.05; n.s. p>0.05.

We further investigated whether decision flexibility could similarly be captured by different configurations of a common integrator process. As before, we fit four versions of the integrator model to Stc and Det PRT behavior. Model comparisons showed that time/reward integration ("Model 3") was the best fit algorithm, both in terms of predictive performance (**Figures 2I-J**, left) and as the best fit for each mouse (**Figures 2I-J**, right). We also observed similar differences in fits of $\boldsymbol{\delta}$ and $\boldsymbol{\rho}$ (**Figures 2K-M**) to Det vs. Stc behavior as in the FM task (**Figures 1O-P**). Male and female mice showed no differences in behavior in either task (**Figure S4**). Taken together, these results establish that mice use similar strategies for flexible foraging across experimental conditions, opening a window to explore the structure of underlying neural dynamics in a more controlled setting.

Foraging contexts recruit distinct brain-wide neural activity states

To generate flexible decisions, context information from the environment must configure neural processes that instantiate or feed into decision computations (**Figure 3A**).[8,9,11] To investigate how context shapes neural activity, we performed high-density *in vivo* electrophysiological recordings using Neuropixels[40] probes targeting a range of brain regions likely involved in coding sensory context and decision dynamics (**Figures 3B-C, S5A-D**). Following our previous study,[37] we densely sampled the frontal cortex, as well as the basal ganglia, subcortical and thalamic nuclei, and sensory cortex spanning olfactory, somatosensory, visual and associative modalities (**Figure 3D**). After clustering and manual curation of spiking data (see Methods), we included a total of 7804 units recorded from 47 sessions across 10 mice for analysis (**Figure S6A**).

We first explored large-scale activity patterns by splitting the dataset into seven "macro-region" categories: dorsal frontal cortex (dFC), ventral frontal cortex (vFC), sensory cortex (SC), hippocampal (HC), basal ganglia (BG), thalamic (THL), and subcortical (SUB) areas (**Figure 3D, Table S1**). We visualized single-neuron activity in a heat map, centering a 20-minute window at the time of the context switch in each session (**Figure 3E**). Global coding of sensory context (i.e. identity of the foraging environment) in single neurons was immediately evident in every region, with changes in activity aligned to the context switch and persisting for at least 10 minutes. To analyze the time-course of this state transition, we measured population vector distances between timepoints following the switch and a preceding time window (**Figure 3F**). Population activity exhibited a transient deviation which decayed over 1-2 minutes and stabilized away from the pre-transition state. While this response was largest in dFC, we observed similar dynamics in all regions, suggesting that environments recruit distinct activity states represented across the brain.

To examine more local dynamics structure, we plotted the correlation matrix of population vectors for individual sessions (**Figure 3G**), which a revealed clear "block" pattern aligned to the context switch, indicating context-specific activity patterns. We also projected population activity onto principal components (**Figure 3H**), which confirmed a dynamical trajectory spanning two distinct states, like the pattern observed across the full dataset (**Figure 3F**). To rule out the possibility that sensory context coding could be spuriously driven by slow-timescale fluctuations or drift artifacts, we quantified the population state change using the angle θ between population vectors sampled across contexts and compared these to within-context samples. Population vectors were nearly orthogonal ($\theta = 97°$) across contexts, and significantly more aligned within contexts ($\theta = 66°, 62°$), indicating more distinct activity states for Det and Stc environments (**Figure 3I**).

Figure 3

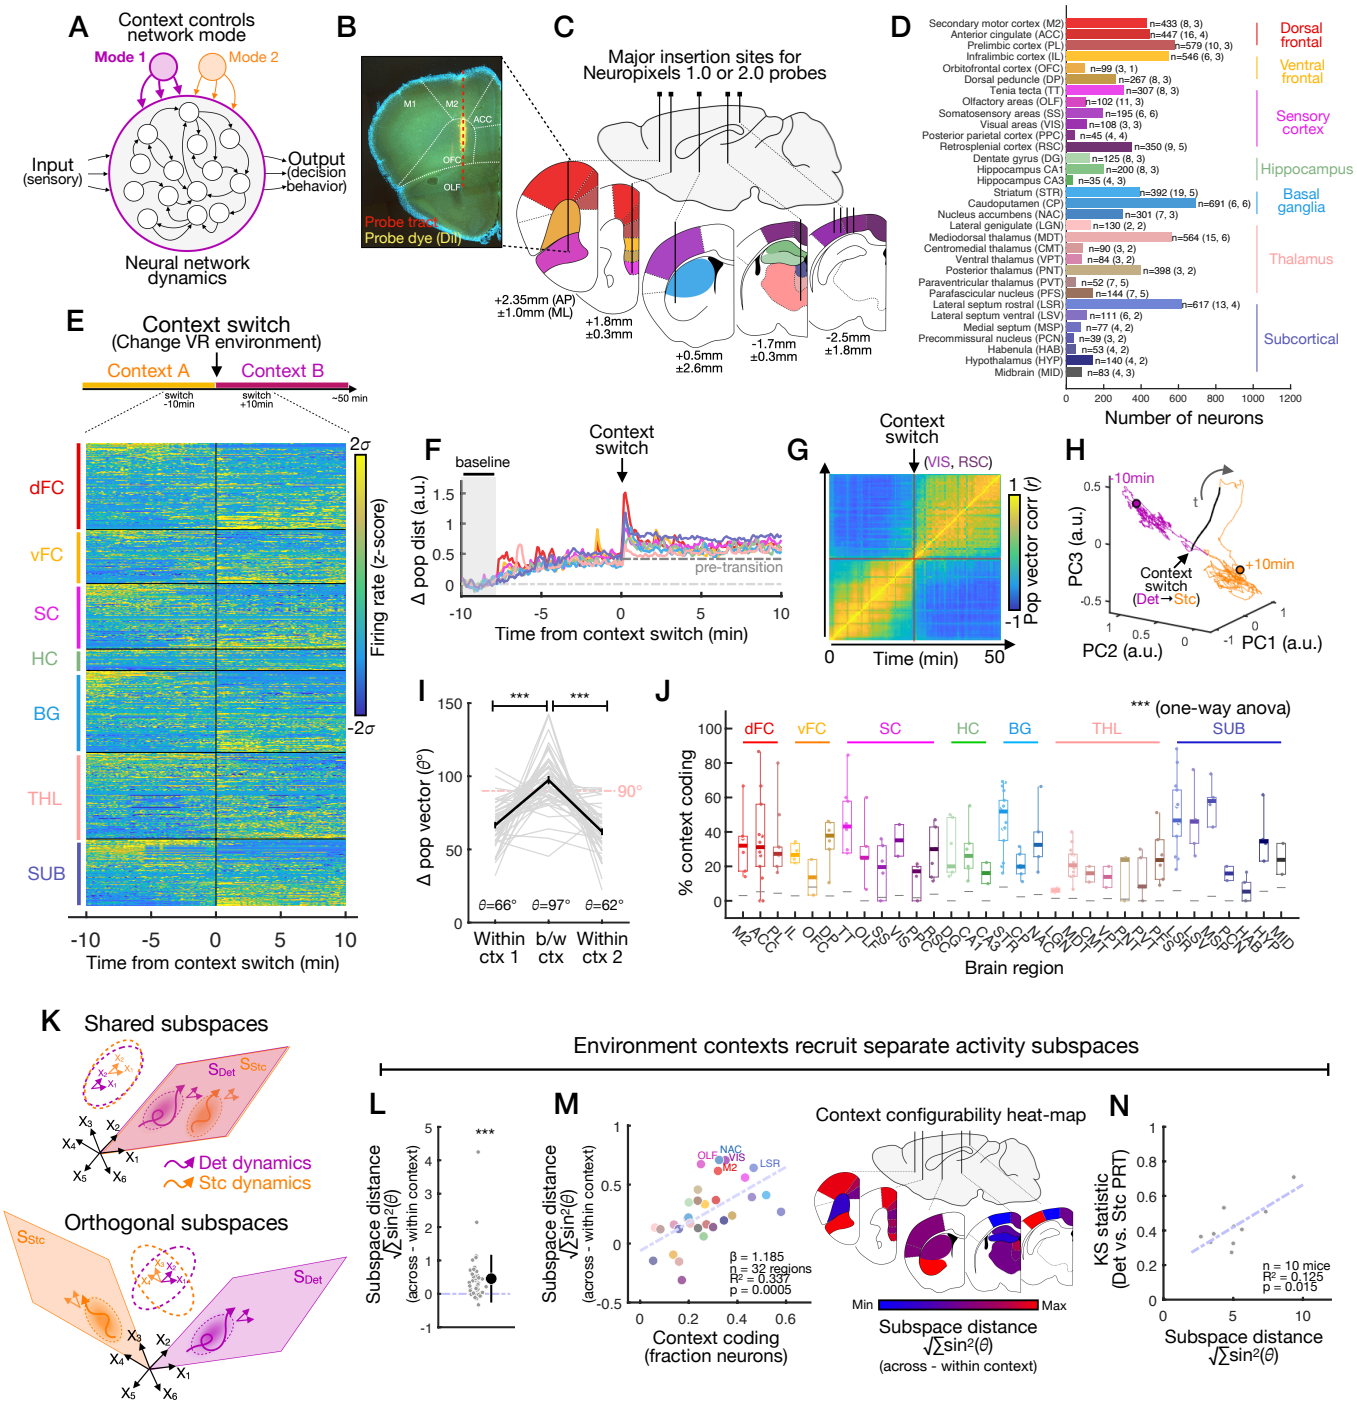
A
Context controls network mode
Mode 1
Mode 2
Input (sensory)
Output (decision behavior)
Neural network dynamics
B
M1
M2
ACC
OFC
OLF
Probe track
Probe dye (DiI)
C
Major insertion sites for Neuropixels 1.0 or 2.0 probes
+2.35mm (AP) ±1.0mm (ML)
+1.8mm ±0.3mm
+0.5mm ±2.6mm
-1.7mm ±0.3mm
-2.5mm ±1.8mm
D
Dorsal frontal
Ventral frontal
Sensory cortex
Hippocampus
Basal ganglia
Thalamus
Subcortical
Number of neurons
E
Context switch (Change VR environment)
Context A
Context B
switch -10min
switch +10min
~50 min
dFC
vFC
SC
HC
BG
THL
SUB
2σ
-2σ
Firing rate (z-score)
Time from context switch (min)
F
baseline
Context switch
Δ pop dist (a.u.)
pre-transition
Time from context switch (min)
G
Context switch
(VIS, RSC)
Pop vector corr (r)
Time (min)
H
10min
+10min
Context switch (Det→Stc)
PC3 (a.u.)
PC2 (a.u.)
PC1 (a.u.)
I
Δ pop vector (θ°)
90°
θ=66°
θ=97°
θ=62°
Within ctx 1
b/w ctx
Within ctx 2
J
*** (one-way anova)
% context coding
Brain region
K
Shared subspaces
Det dynamics
Stc dynamics
Orthogonal subspaces
Environment contexts recruit separate activity subspaces
L
Subspace distance √∑sin²(θ) (across - within context)
M
β = 1.185
n = 32 regions
R² = 0.337
p = 0.0005
Context coding (fraction neurons)
Context configurability heat-map
Min
Max
Subspace distance √∑sin²(θ) (across - within context)
N
KS statistic (Det vs. Stc PRT)
n = 10 mice
R² = 0.125
p = 0.015
Subspace distance √∑sin²(θ)

Figure 3. Foraging environments recruit distinct brain-wide activity states.

A) Cartoon illustrating contextual configuration of neural network modes for context-dependent computation (adapted from Mante et al. 2013).
B) Example histological section showing location of a Neuropixels probe insertion through frontal cortex (dotted red line, probe tract, yellow, probe dye; see also **Figure S5** for examples).
C) Schematic showing major insertion locations with coordinates for acute recording experiments.
D) Summary of recording dataset; for each region, bar indicates the number of recorded neurons, and number of (sessions, mice) to the right (see **Figure S6** for full recording summaries).
E) Heat map showing activity of all recorded neurons (n = 7804, rows) over a 20-minute window centered around the context transition (change of virtual environment, t = 0).
F) Time course of population vector distances (euclidean norm) from pre-transition time window (-10 to -8 minutes) for pooled pseudo-populations of neurons from each region.
G) Correlation matrix of population vectors from an example recording in visual/retrosplenial cortex; block structure indicates within-context activity patterns that are reorganized at the time of context transition with greater within- than across-context similarity.
H) Principal component (PC) projections of activity from session in **(G)** showing transient deviation followed by persistent change in population activity at context switch (marked with arrow; gray arrow indicates direction of time).
I) Angle measurements (θ in degrees) comparing within- vs. between-context population vectors (mean ± SEM, SR tests, p = 9.09e-8 [within context 1 vs. between], p = 3.13e-8 [between vs. within context 2], n = 47 sessions).
J) Single neuron context coding across sampled regions (one-way anova, p = 5.07e-4, n = 187 samples).
K) Schematic illustrating population activity subspace divergence and measurement of subspace alignment using canonical correlation analysis (CCA, see **Figure S7** and Methods).
L) Normalized (between context - within context) chordal distance ($\sqrt{\sum \sin^2(\theta)}$ of angles between canonical vectors/directions) for Det- vs. Stc-context subspaces (mean ± s.d., t-test, p = 7.23e-5, n = 47 sessions).
M) Left, correlation across regions between single-neuron context coding and subspace difference (chordal distance, linear regression, p = 4.95e-4, n = 32 regions); right, heat-map of normalized chordal distance across brain regions (scaled min-max in blue-red colormap).
N) Correlation across mice between subspace difference (chordal distance) and flexible decision behavior, measured using the KS statistic of Stc vs. Det PRT distributions (linear regression, p = 0.0112, n = 10 mice).

Standard deviation (s.d.); Standard error of mean (SEM); Wilcoxon signed-rank test (SR test); ***p<0.001; **p<0.01; *p<0.05; n.s. p>0.05.

Finally, we characterized the finer regional distribution of sensory context coding by identifying single neurons that exhibit significant and persistent changes in activity at context switch ("context coding" neurons). Overall, 29.6% of cells were context coding, and although non-uniform, context information was widespread across the brain. In fact, while sensory areas (e.g. tenia tecta) exhibited substantial context coding, likely attributable to the sensory cues that signal environment identity, subcortical (e.g. septal nuclei) and most frontal areas were also modulated by context. This suggests that context signals organize distributed activity states which could structure context-specific dynamics for decision computations (**Figure 3A**).

Foraging contexts organize dynamics into separate activity subspaces

Defined in this way, "context coding" indicates sensitivity to sensory information that informs environment identity but does not obligately imply that the structure of neural dynamics within each context is distinct, as might be expected for different configurations of a decision process. For example, neural dynamics for each context could lie within shared activity subspaces (**Figure 3K**). Alternatively, context could organize activity into partially separate or orthogonal subspaces, creating context-specific functional modes.[7,22,41,42] To test whether dynamics are reconfigured by context, we quantified the alignment of Det vs. Stc activity subspaces. We used principal components analysis (PCA) to find orthonormal bases (subspaces $\mathbf{S_{Det}}$ and $\mathbf{S_{Stc}}$) capturing within-context activity, and then applied canonical correlation analysis[43,44] (**Figure S7A**, see Methods) to obtain vector pairs, and their canonical correlations, which describe the alignment between $\mathbf{S_{Det}}$ and $\mathbf{S_{Stc}}$. This method can be used to quantify the number of orthogonal dimensions between $\mathbf{S_{Det}}$ and $\mathbf{S_{Stc}}$ or to measure overall subspace distance (e.g. the "chordal distance" $\sqrt{\sum \sin^2(\theta)}$ using angles θ from the canonical correlations, **Figure S7A-E**). We quantified these for each session, subtracting within-context values to control for baseline variability.

We found that subspace distances (**Figure 3L**) and the number of orthogonal vectors (**Figure S7H**) were significantly higher than within-context comparisons, suggesting that environment contexts not only drive linear changes in firing rate ("context coding") but also structure context-specific population activity. Interestingly however, we observed substantial variability in subspace distances across regions, which was itself correlated with context coding (**Figure 3M**) and much greater than expected from effects of context coding by itself (**Figure S7F-I**). This indicates that regions which are more sensitive to sensory context also tend to exhibit more context-specific dynamics, potentially supporting distinct functional modes for decision making. Consistent with this idea, subspace differences were also correlated, across mice and sessions, with decision flexibility in the task (**Figures 3N, S7J**), suggesting a connection between dynamics modes and behavior. Thus, we next focused our investigation specifically on how environment contexts shape decision activity for foraging.

Context-specific decision variables are encoded in separate neuronal subpopulations

Previous studies of context-dependent decision making have mostly evidenced context effects that are upstream of decision activity.[11–14,19] Broadly, these can be thought of as "**input configuration**" mechanisms (**Figure 4A**, left), which alter decision behavior by affecting the inputs to some context-invariant decision process. Alternatively, context could parse decision

activity into distinct modes that occupy different subspaces ("**subspace configuration**", **Figure 4A**, right), a different neural structure which could offer computational and learning advantages (see Discussion).[7,8,22,24] These abstract models describe the population geometry of decision dynamics and, as instantiated at the cellular level, relate to the correlation structure of coding in single neurons (**Figure 4B**). That is, while correlated decision coding produces aligned coding vectors (**Figure 4B**, right), less or uncorrelated coding produces divergent (potentially orthogonal) coding vectors for dynamics supporting context-specific decision computation (**Figure 4B**, left).

To test which model best describes our data, we first quantified the correlation between single-neuron activity and model-fit DVs against shuffled data to identify "DV-coding" neurons (see Methods). Many neurons (38.9% of all) exhibited trial activity that was highly correlated with the DV on either Stc patches, Det patches, or both (**Figure 4C**, top). Interestingly, while a small subset (5.7% of all) encoded DVs across both contexts, most cells were specific to either Stc (12.9%) or Det (20.3%) patches, reflecting the context-specific DV with a mixture of positive and negative coding (**Figure 4C**, bottom). From this, we could identify neurons with trial-averaged activity that exhibited ramping and reward deflections in the preferred context yet were nearly unresponsive in the non-preferred context (**Figures 4D-E**, top). Averaging across these categories yielded context-specific coding of the Stc or Det DV in completely non-overlapping subpopulations of neurons (**Figures 4D-E**, bottom).

This single-neuron analysis suggests that context signals organize decision activity into separate subpopulations, which in our framework represents a particular realization of subspace configuration (**Figure 4B**, left) and supports that model. However, this does not necessitate that DV-coding neurons track distinct configurations of the decision process. Based on our modeling results, we reasoned that neurons encoding configured DVs should exhibit distinct activity patterns during patch leaving decisions (**Figure 4F**). Specifically, we hypothesized that: 1) Det-DV neurons exhibit faster ramping activity than Stc-DV neurons (reflecting a larger drift rate $\boldsymbol{\delta}$), and 2) Det-DV neurons exhibit larger reward responses (reflecting a larger reward sensitivity $\boldsymbol{\rho}$).

To test these predictions, we examined activity of Stc- and Det-DV neurons during patch trials. At trial onset (**Figure 4G**), we observed negative responses in both populations, with a larger negative response in Stc-DV neurons reflecting the trial onset and trial-initial reward. Following this, both populations exhibited ramping, with Det-DV neurons ramping faster (**Figure 4G**, right). As this was partially confounded with trial initiation, we also examined responses around additional patch rewards. As predicted, rewards elicited stronger negative responses in Det- vs. Stc-DV neurons as well as faster ramping after the reward (**Figure 4H**).

Finally, to test more precisely whether ramping is faster in Det-DV cells, we examined activity leading up to the leaving decision (**Figure 4I**) in trials with no preceding rewards, eliminating the potential confound of reward response dynamics (see Methods). Det-DV neurons also exhibited faster ramping before leaving decisions. Taken together, this shows that neurons encoding context-specific DVs exhibit distinct dynamics reflecting parameter configurations predicted by the integrator model: Det-DV neurons exhibit both larger reward responses ($\boldsymbol{\rho}$) and faster ramping ($\boldsymbol{\delta}$) during patch leaving decisions.

Figure 4

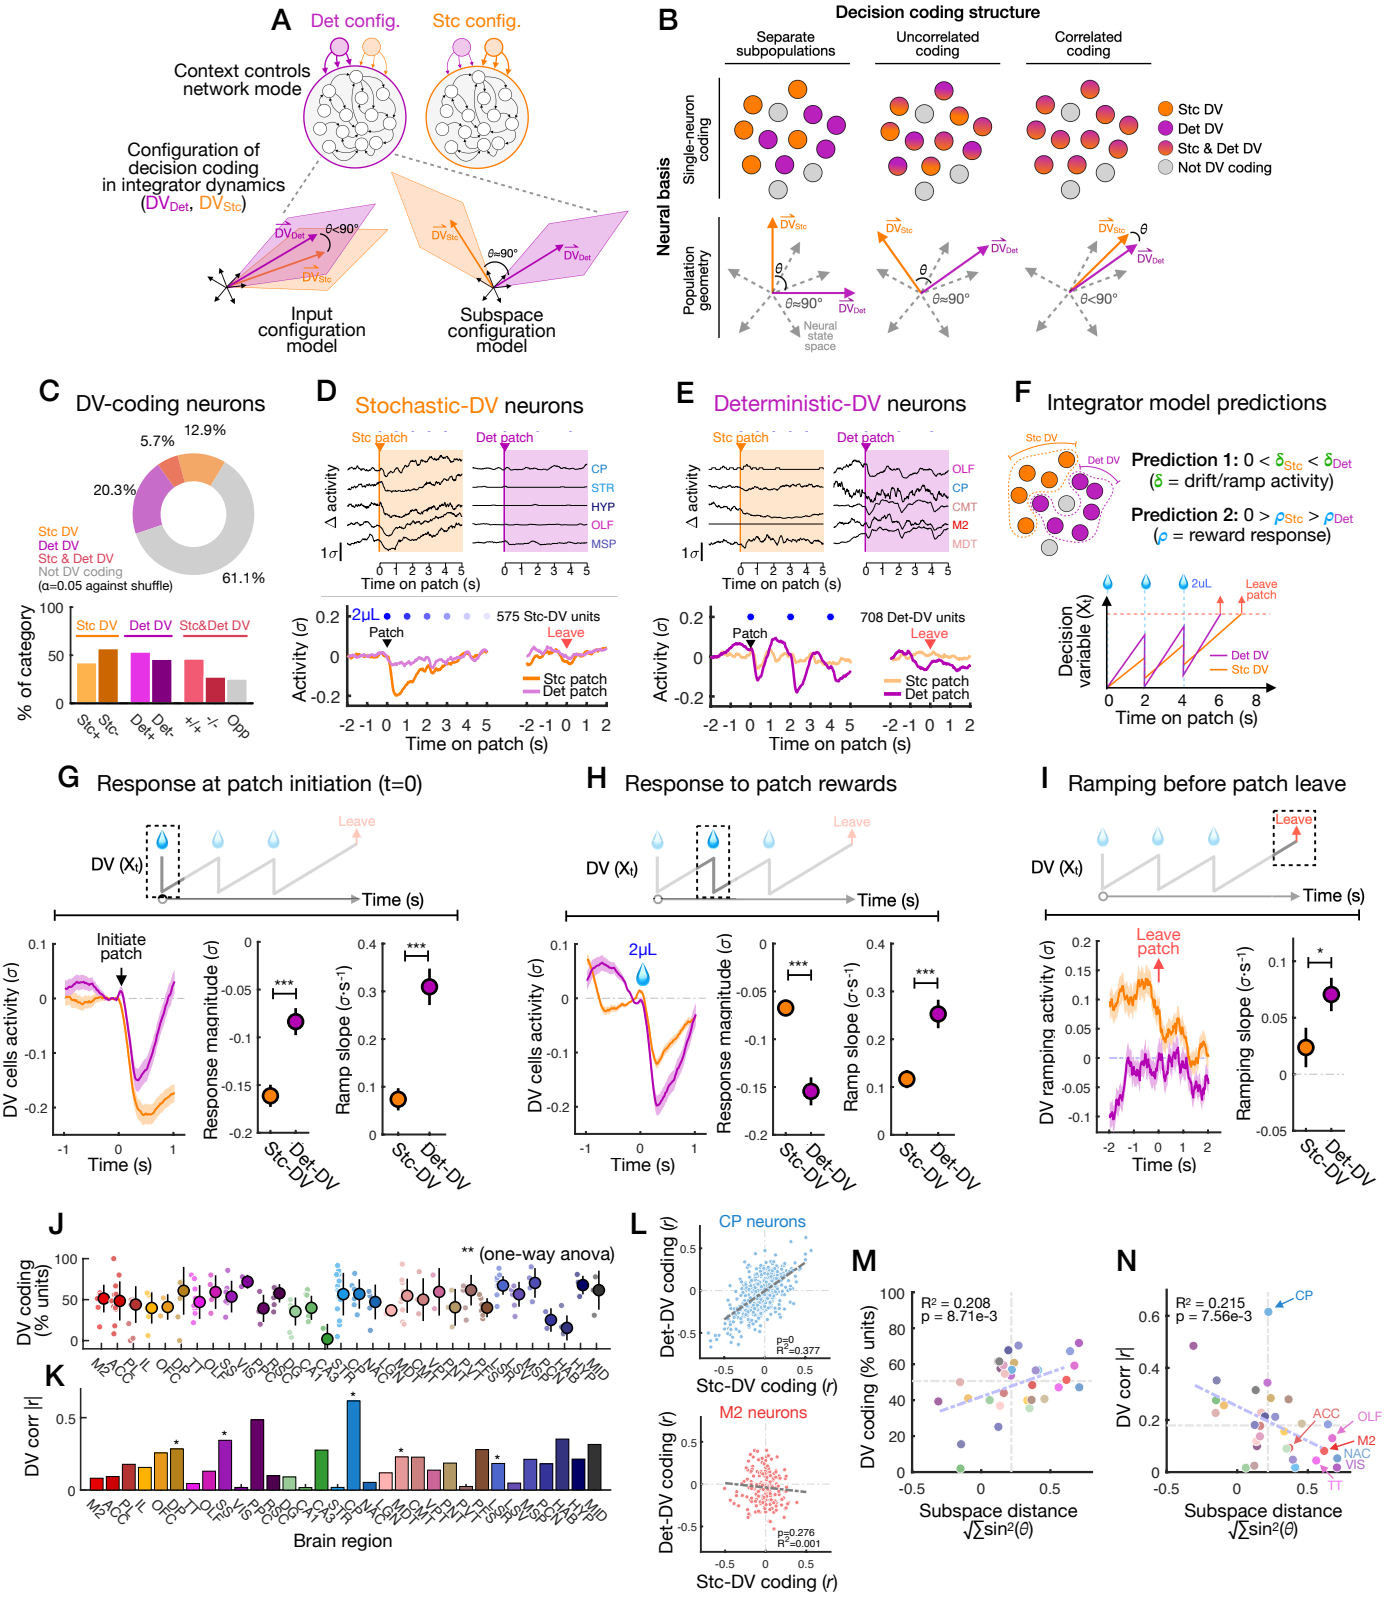
A
Det config.
Stc config.
Context controls network mode
Configuration of decision coding in integrator dynamics (DV_{Det}, DV_{Stc})
Input configuration model
Subspace configuration model
B
Decision coding structure
Separate subpopulations
Uncorrelated coding
Correlated coding
Neural basis
Single-neuron coding
Population geometry
Stc DV
Det DV
Stc & Det DV
Not DV coding
Neural state space
C
DV-coding neurons
5.7%
12.9%
20.3%
61.1%
Stc DV
Det DV
Stc & Det DV
Not DV coding
(α=0.05 against shuffle)
% of category
Stc&Det DV
D
Stochastic-DV neurons
Stc patch
Det patch
Δ activity
CP
STR
HYP
OLF
MSP
Time on patch (s)
2μL
Patch
575 Stc-DV units
Leave
Activity (σ)
E
Deterministic-DV neurons
OLF
CP
CMT
M2
MDT
708 Det-DV units
F
Integrator model predictions
Prediction 1: 0 < δ_{Stc} < δ_{Det}
(δ = drift/ramp activity)
Prediction 2: 0 > ρ_{Stc} > ρ_{Det}
(ρ = reward response)
Decision variable (X_t)
Leave patch
G
Response at patch initiation (t=0)
DV (X_t)
Time (s)
Initiate patch
DV cells activity (σ)
Response magnitude (σ)
Ramp slope (σ·s⁻¹)
Stc-DV
Det-DV
H
Response to patch rewards
I
Ramping before patch leave
DV ramping activity (σ)
Ramping slope (σ·s⁻¹)
J
** (one-way anova)
DV coding (% units)
K
DV corr |r|
Brain region
L
CP neurons
M2 neurons
Det-DV coding (r)
Stc-DV coding (r)
M
DV coding (% units)
Subspace distance
$\sqrt{\sum \sin^2(\theta)}$
N
DV corr |r|
CP
ACC
OLF
M2
NAC
VIS
TT

Figure 4. Subspace dynamics structure context-specific configurations of decision activity.

A) Illustration of subspace (left) vs. input (right) configuration models for how context can shape decision activity. Subspace configuration exhibits orthogonal Det vs. Stc decision coding vectors.
B) Cartoon of different possibilities for cellular-level coding structure underlying population geometry of decision activity (linearly captured on coding vectors). Orthogonal coding vectors correspond to uncorrelated or non-overlapping single-neuron decision coding.
C) Top, percentage breakdown of Stc-DV, Det-DV, both, and non-coding neurons from among all recorded; bottom, breakdown of positive and negative coding cells within each category. For cells encoding both Stc and Det DVs, "opp" indicates opposite-sign coding for Det vs. Stc.
D) Traces of trial-average activity for example Stc-DV cells (top) and population average on Stc or Det patch trials (bottom).
E) Traces of trial-average activity for example Det-DV cells (top) and population average (bottom).
F) Predictions of distinct activity profiles for Stc-DV and Det-DV cells based on parameter fits from integrator models (**Figures 2I-L**).
G) Left, average responses of Stc-DV (orange) and Det-DV (purple) cells at patch initiation (mean ± SEM); middle, comparison of response magnitude (t-test, p = 4.08e-5); right, comparison of ramping slope following response (t-test, p = 6.87e-7).
H) Left, average responses of Stc- and Det-DV cells at additional patch rewards (mean ± SEM); middle, comparison of response magnitude (t-test, p = 3.37e-7); right, comparison of ramping slope following response (t-test, p = 2.59e-4).
I) Left, average responses of Stc- and Det-DV cells at patch leave (mean ± SEM); right, comparison of ramping slope before patch leave (t-test, p = 0.0392).
J) Breakdown of single-neuron DV coding (fraction significant for either or both category at α=0.05 - shuffle) across regions (ones-way anova, p = 0.0006).
K) DV coding similarity (|correlation r| between Stc-DV and Det-DV coding measures) for neurons from sampled regions (* indicates independent correlation for that region at α=0.05).
L) Correlation (*r*) of Stc- and Det-DV coding across single neurons from two example regions (CP, Caudoputamen; M2, secondary motor cortex).
M) Correlation between subspace difference (from **Figure 3N**) and DV coding as in **(J)** for sampled regions; grey lines indicate average of these metrics across all regions (linear regression, R^2 = 0.215, p = 7.56e-3).
N) Correlation between subspace difference (from **Figure 3N**) and coding similarity as in **(I)** for sampled regions; grey lines indicate average metrics across all regions (linear regression, R^2 = 0.208, p = 8.71e-3).

Standard error of mean (SEM); ***p<0.001; **p<0.01; *p<0.05; n.s. p>0.05.

Activity subspaces structure context-specific decision coding

We next examined how coding of configured DVs was represented across the brain. Quantifying the fraction of neurons encoding either or both DVs, we found that most regions contained substantial single-neuron DV coding (**Figure 4J**), consistent with recent reports of nearly brain-wide integrator dynamics during perceptual decisions.[45,46] Regions also exhibited broadly similar coding strength (**Figure S8A**).

However, it is possible that despite widespread integrator dynamics, context-specific DV coding may be more localized or graded.[8] While the dynamics configuration model predicts correlated single-neuron DV coding (**Figure 4B**), subspace configuration predicts uncorrelated coding. Thus, for each region, we quantified DV similarity as the correlation between single-neuron Stc- and Det-DV coding strengths. Intriguingly, we found substantial regional heterogeneity in coding similarity (**Figure 4K**). Indeed, while some regions exhibited highly correlated DV coding (e.g. Cautoputamen, **Figure 4L**, top), others exhibited completely uncorrelated, context-specific DV coding (e.g. secondary motor cortex, **Figure 4L**, bottom).

How might this heterogeneity arise? One possibility is that contextual configuration of activity subspaces decorrelates Stc- and Det-DV coding into orthogonal dimensions. To test this idea, we plotted the subspace distance for each region ("chordal distance" from **Figure 3N**) against measures of overall DV-coding (**Figure 4M**) or coding similarity (**Figure 4N**). Strikingly, both were positively correlated with subspace distance, suggesting that subspace configurations, recruited by environment context, parse decision activity into uncorrelated modes. A subset of interconnected regions, including sensory areas (olfactory and visual cortex), dFC regions, and some subcortical areas exhibited strong subspace divergence and context-specific DV coding. To further investigate this coding structure and its role in behavior, we thus narrowed our focus to the dFC, which has been functionally linked to self-timed movement and foraging.[37,38,47–49]

dFC neurons encode orthogonalized context-specific decision variables

To examine decision coding in dFC (including M2, ACC, and PL), we first used a linear decoder analysis to confirm that the Stc- and Det-DV could be read out from population activity (**Figure 5A**). Decision coding within each context was stable enough that we could visualize single-trial DVs by projecting activity onto either Det- or Stc-DV coding vectors (**Figure 5B**). If dFC decision coding is orthogonal, as in the subspace configuration model, these coding dimensions should be: 1) geometrically perpendicular (i.e. angle $\theta \approx 90°$), and 2) mutually uninformative of the DV. We measured the angle θ between DV coding vectors and found these to be close to 90° and greater than within-context comparisons (**Figure 5C**), indicating orthogonality. Further, while within-context decoder performance was above chance for dFC neurons (**Figure 5A**), decoders could not predict DVs from the non-preferred context (**Figure 5D**), indicating that DV coding is context-specific and mutually uninformative.

As dFC receives direct input from sensory areas, we wondered whether orthogonal DV coding was shaped by task demands or was inherited from upstream signals encoding sensory context. We reasoned that if context specificity is independent of cognitive demand, then it should also be present in task-naïve mice in the same sensory conditions. To test this, we recorded dFC neurons in a separate cohort of mice on the first 3 days of the task (**Figures 5E, S6B**) and, first,

Figure 5

dFC neurons learn orthogonal representations of the integrator decision variable

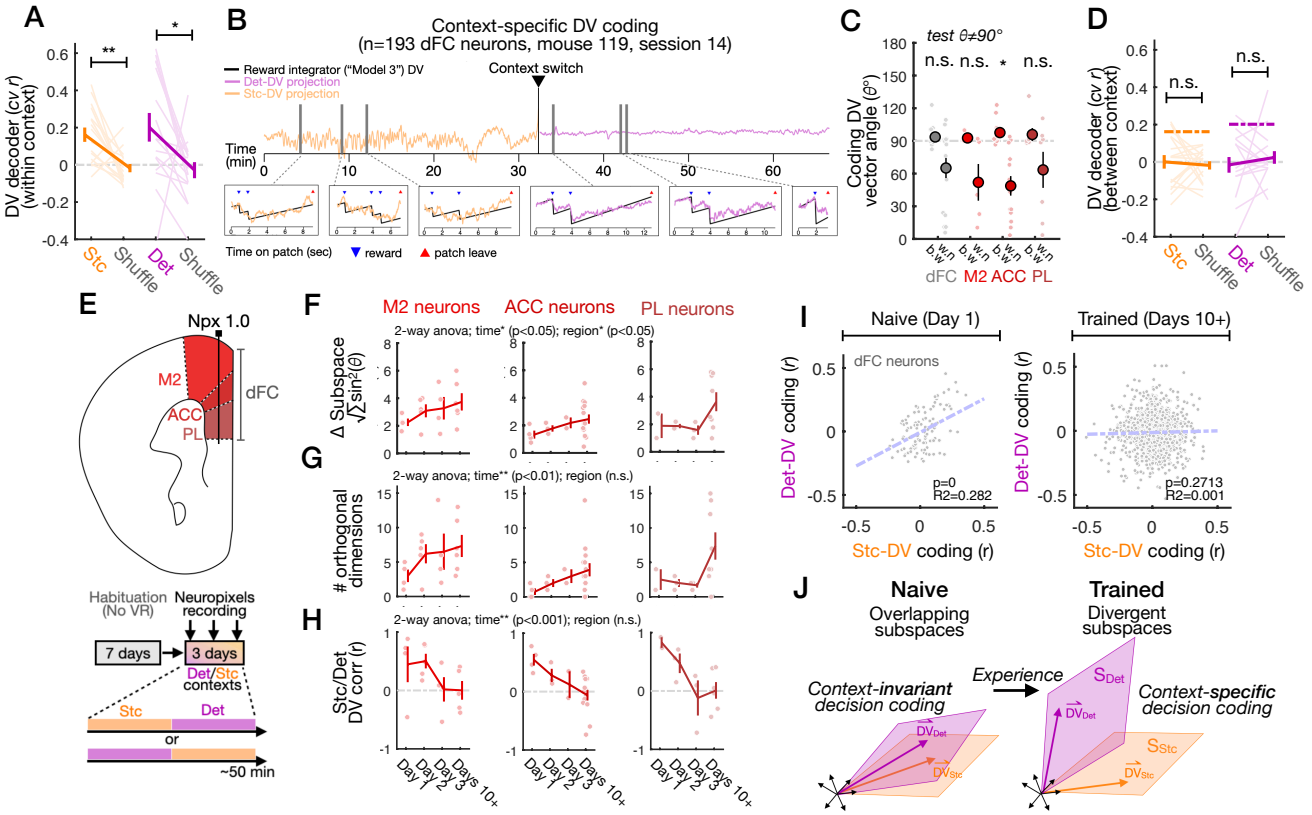

dFC activity causally contributes to decision flexibility

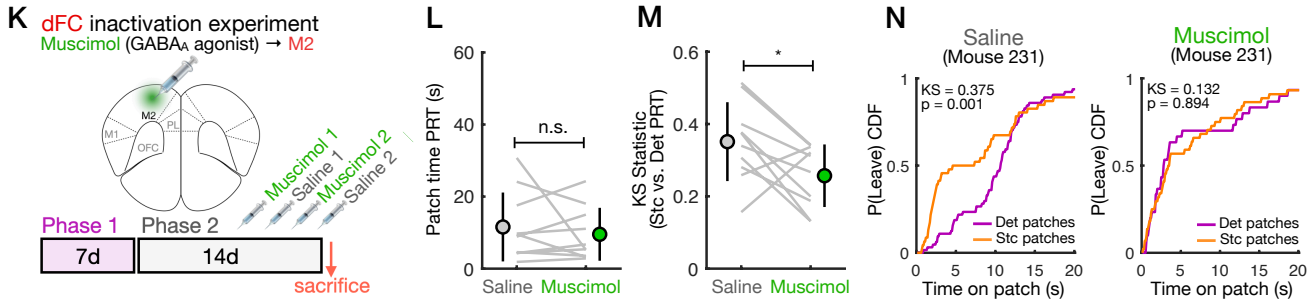

Figure 5. Orthogonalized integrator dynamics in dFC support decision flexibility.

A) Performance of Linear decoders for Stc or Det decision variables (DVs, cross-validated correlation (cv-r) between predicted and model-fit DVs on held out trials) from dFC activity (SR tests, $p = 0.0084$, Stc, $p = 0.0151$, Det, $n = 16$ sessions).
B) dFC activity projected onto Det-DV or Stc-DV coding vectors (defined using coefficients from regression used for decoders, see Methods) from an example session (Mouse 116, session 10), showing DV coding on individual Stc and Det trials.
C) Angle (θ in degrees) between Det-DV and Stc-DV coding vectors for dFC neurons (including secondary motor cortex M2, anterior cingulate ACC, and prelimbic cortex PL) or for each region separately (mean ± s.d., t-tests against $\theta = 90°$, $p = 0.1444$ [dFC], $p = 0.2244$ [M2], $p = 0.0276$ [ACC], $p = 0.3598$ [PL]; 2-way anova, $p = 0$ (within vs. between), $p = 0.7308$, region, d.f. = 49).
D) Performance of Stc-DV and Det-DV decoders to predict DVs from other context (dashed lines indicate within-context performance from (**A**), SR tests, $p = 0.8767$, Stc, $p = 0.7174$, Det, $n = 16$ sessions).
E) Schematic of experiment to measure neural activity in task-naive mice; After habituation to running wheel and head-fixation, acute Neuropixels recordings targeting dFC were performed on the first 3 days of Phase 2 of the VR foraging task (Phase 1 was not included).
F) Subspace difference (chordal distance $\sqrt{\sum \sin^2(\theta)}$) in dFC neurons from task-naive (days 1-3) and trained sessions (mean ± SEM, 2-way anova, $p = 0.0083$, region, $p = 0.0124$, time, d.f. = 57).
G) Number of orthogonal vectors ($\theta > 80°$) in dFC neurons from task-naive (days 1-3) and trained sessions (mean ± SEM, 2-way anova, $p = 0.0083$, region, $p = 0.0137$, time, d.f. = 57).
H) DV coding similarity (correlation between Stc-DV and Det-DV coding metric r) across days (mean ± SEM, 2-way anova, $p = 0.8174$, region, $p = 0$, time, d.f. = 53).
I) Left, correlation across dFC neurons between Stc-DV and Det-DV coding in day 1 sessions (linear regression, $R^2 = 0.282$, $p = 2.01e\text{-}9$, $n = 111$ cells); right, correlation between Stc-DV and Det-DV coding in experienced sessions (days 10+, $R^2=8.83e\text{-}4$, $p=0.271$, $n = 1372$ cells).
J) Cartoon illustration of progressive subspace divergence in dFC across days and de-correlation of context-specific decision coding as shown quantitatively in **(I)**.
K) Experimental approach for dFC inactivation experiments. After 2 weeks of training in the VR task, mice were injected unilaterally in M2 region of dFC with saline or Muscimol (alternating across four days) and then tested in the task each day (see Methods for details).
L) Comparison of overall patch residence time (median PRT for each mouse) behavior under Muscimol vs. saline control sessions (mean ± s.d., SR test, $p = 0.9219$, $n = 10$ mice).
M) Comparison of difference between Det and Stc PRT distributions (KS statistic) under Muscimol vs. saline control (mean ± s.d., SR test, $p = 0.0195$, $n = 10$ mice).
N) Det and Stc PRT distributions (CDF of pLeave over time) for Mouse 231 under saline (left) and Muscimol (right) showing collapse of context-specific behavior with inactivation.

Standard deviation (s.d.); Standard error of mean (SEM; Wilcoxon signed-rank test (SR test).
***$p<0.001$; **$p<0.01$; *$p<0.05$; n.s. $p>0.05$.

repeated the subspace alignment analysis on early (days 1-3) and trained (days 10+) sessions. We observed a significant increase in subspace distances and orthogonal vectors over days (**Figures 5F-G**), suggesting that dFC dynamics are organized into context-specific subspaces with experience. We also quantified similarity in Stc vs. Det DVs with the correlation of single-neuron coding. Notably, while DV coding was initially correlated on day 1 (**Figure 5H**) across dFC, it was successively decorrelated across days and completely uncorrelated in trained sessions (**Figure 5I**), consistent with the successive divergence of activity subspaces (**Figure 5J**). Together, these results show that while task variables including context and decision coding are generally present in naïve animals, task experience organizes different activity subspaces and orthogonalizes decision dynamics into context-specific modes.

dFC activity causally contributes to context-specific foraging decisions

This coding structure positions dFC as a potentially key node for contextual configuration of foraging decisions. Thus, we targeted dFC to test its causal contribution to decision flexibility in our task (**Figure 5K**). After training mice, we injected either saline or the $GABA_A$ receptor agonist Muscimol (on alternating days) into the M2 region of dFC to locally silence neural activity. While dFC inactivation had no effect on average PRTs (**Figure 5L**) or task engagement (**Figure S9A),** suggesting that foraging behavior remained intact, we observed a significant decrease in decision flexibility under Muscimol (**Figures 5M, S9B**). In some mice, Stc and Det PRT distributions were indistinguishable under Muscimol (Mouse 231, **Figure 5N**), indicating complete abolition of context specificity. This demonstrates that dFC activity, while apparently redundant for the core behavior, contributes to contextual configuration of the foraging decisions.

dFC coding structure is preserved across virtual and naturalistic settings

The VR task enabled us to sample neural activity widely, revealing distributed activity states that organize context-specific, orthogonal decision coding, with functional significance in dFC. We noted, however, that despite similar tendencies, behavior in the VR task was more variable than in the FM task (**Figure S3**), speculatively due to the artificiality of head fixation, virtual reality, and/or impoverished proprioceptive feedback. Thus, we next sought to explore how this coding structure, identified under experimental constraint, might generalize to the naturalistic setting of free foraging (**Figure 6A**).[26,27,29] For this, we returned to the FM task we initially developed (**Figures 1C, 6B**). We implanted chronic Neuropixels probes[50] in four mice targeting dFC, recorded across daily sessions in the task, and retained 822 dFC neurons from 16 sessions (Phase 2) for analysis (**Figures 6C, S6C**).

We first tested whether dFC neurons encoded context (i.e. chamber identity) in the FM task, which lacks explicit cues. Binary decoders could accurately predict Stc vs. Det identity from population activity on both patch trials and inter-trial intervals (ITI, **Figure 6D**), indicating stable context coding. Visualizing neural activity in a heatmap centered around the "door open" time, when mice transition between contexts (**Figure 6E**), we observed large-scale reorganization of population activity, with many neurons exhibiting sustained changes in firing rate. We confirmed this coding structure in individual sessions by examining population vector correlation matrices (**Figure 6F**) and trajectory dynamics (**Figure 6G**), which exhibited similar structure and state

Figure 6

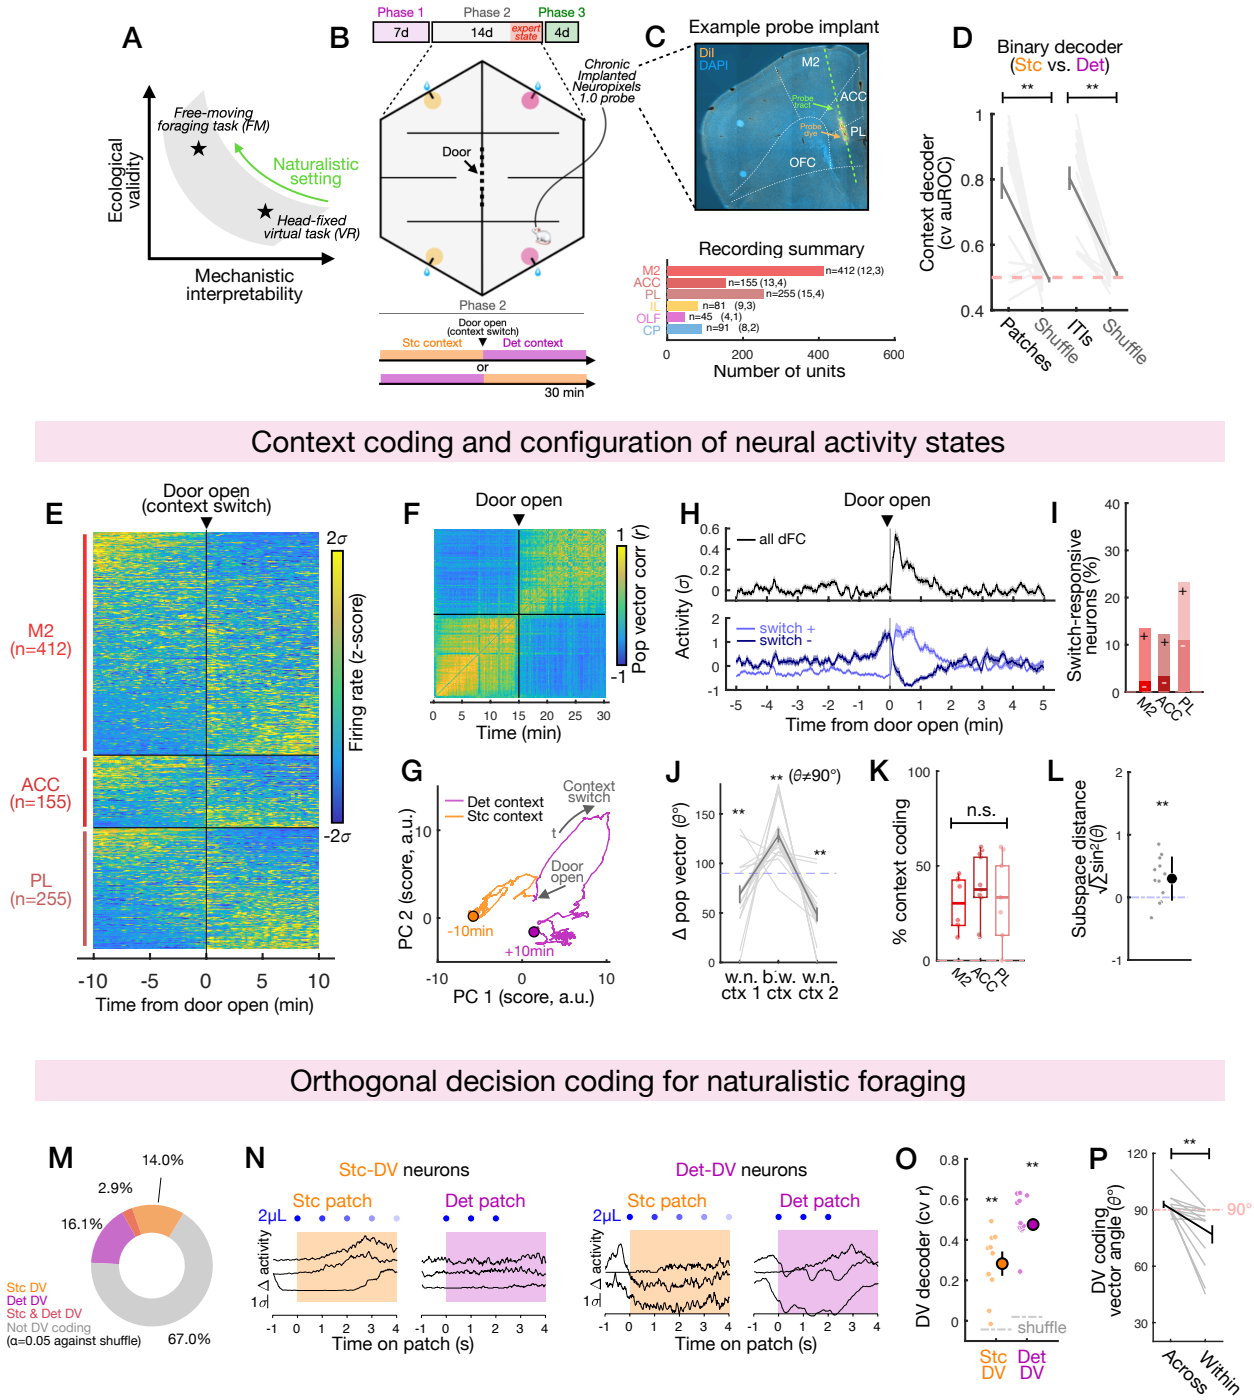
A
Ecological validity
Mechanistic interpretability
Free-moving foraging task (FM)
Naturalistic setting
Head-fixed virtual task (VR)
B
Phase 1
Phase 2
Phase 3
7d
14d
4d
exper state
Door
Chronic Implanted Neuropixels 1.0 probe
Phase 2
Door open (context switch)
Stc context
Det context
or
30 min
C
Example probe implant
DiI
DAPI
M2
ACC
PL
OFC
Recording summary
M2 n=412 (12,3)
ACC n=155 (13,4)
PL n=255 (15,4)
IL n=81 (9,3)
OLF n=45 (4,1)
CP n=91 (8,2)
Number of units
D
Binary decoder (Stc vs. Det)
Context decoder (cv auROC)
Patches
Shuffle
ITIs
Shuffle
Context coding and configuration of neural activity states
E
Door open (context switch)
M2 (n=412)
ACC (n=155)
PL (n=255)
Firing rate (z-score)
Time from door open (min)
F
Door open
Pop vector corr (r)
Time (min)
H
Door open
all dFC
switch +
switch -
Activity (σ)
Time from door open (min)
I
Switch-responsive neurons (%)
M2
ACC
PL
G
Det context
Stc context
Context switch
Door open
-10min
+10min
PC 2 (score, a.u.)
PC 1 (score, a.u.)
J
(θ≠90°)
Δ pop vector (θ°)
w.n. ctx 1
b.w. ctx
w.n. ctx 2
K
n.s.
% context coding
M2
ACC
PL
L
Subspace distance √Σsin²(θ)
Orthogonal decision coding for naturalistic foraging
M
14.0%
2.9%
16.1%
67.0%
Stc DV
Det DV
Stc & Det DV
Not DV coding
(α=0.05 against shuffle)
N
Stc-DV neurons
Det-DV neurons
Stc patch
Det patch
2μL
Δ activity
1σ
Time on patch (s)
O
DV decoder (cv r)
shuffle
Stc DV
Det DV
P
DV coding vector angle (θ°)
90°
Across
Within

Figure 6. Structure of dFC dynamics is preserved across experimental settings.

A) Cartoon illustration (as in **Figure 2A**) of tradeoffs in experimental approach. The free-moving (FM) volitional task design can test generalization of neural circuit models with less structural constraint and increasing ecological validity (green arrow).
B) Schematic of FM foraging task phase 2, in which mice forage in Det and Stc contexts across two 15-minute blocks with one transition in the middle of the session through a controllable door.
C) Top, example histology showing location of chronic Neuropixels probe implant targeting the dorsal frontal cortex (dFC); bottom, summary of single neuron recordings from Phase 2. A total of n = 822 neurons were recorded from secondary motor cortex (M2), anterior cingulate (ACC), and prelimbic cortex (PL) across 16 sessions in 4 mice.
D) Performance of decoders (using Fisher's linear discriminant) to predict Det vs. Stc context identity from dFC population activity during patch trials (left) or inter-trial intervals (ITI, right), quantified using the area under ROC (auROC) curve for predicted vs. true labels on held-out trials (SR test, p = 0.0023, trials, p = 4.3778e-4, ITIs, n = 12 sessions).
E) Heat map of single-neuron activity centered around the time of door opening/context transition.
F) Correlation matrix of dFC population vectors sampled across an example session. Block structure indicates within-context similarity reorganized at the context switch (door open time).
G) Principal component (PC) projection of activity from example session in **(F)** showing persistent change in population activity upon context transition.
H) Top, transient response in mean population activity of dFC neurons centered around the door open time; bottom, averages of subpopulations that are excited (light blue, n = 93) or inhibited (dark blue, n = 41) transiently at context transition ("switch-responsive" neurons, permutations tests with $\alpha = 0.01$, see Methods).
I) Breakdown of switch-responsive neurons across dFC regions (light, positive; dark, negative responses).
J) Angles (θ in degrees) comparing within- vs. between-context population vectors (t-tests against $\theta = 90°$, p = 0.0227 [within context 1], p = 6.2025e-7 [between], p = 9.2935e-9 [within context 2], n = 12 sessions).
K) Breakdown of context-coding neurons across dFC regions (anova, p=0.0362, n = 12 sessions).
L) Normalized (between context - within context) chordal distance $\sqrt{\sum \sin^2(\theta)}$ for Det- vs. Stc-context subspaces (mean ± s.d., t-test, p = 0.0132, n = 12 sessions).
M) Breakdown of Stc-DV (orange), Det-DV (purple), Stc and Det DV (red), and non-DV coding neurons across dFC (n = 822 neurons, permutation test at $\alpha = 0.05$).
N) Trial-averaged Stc and Det patch activity for example Stc-DV (left) and Det-DV (right) coding neurons.
O) Performance of linear decoders to predict Stc or Det DVs from dFC activity on held-out trials (SR tests against shuffle, p = 9.76e-4 [Stc trials], p = 4.88e-4 [Det trials], n = 12 sessions).
P) Angle θ between DV-coding vectors for across-context DVs vs. within-context control comparisons (SR test, p = 4.88e-4, n = 12 sessions).

Standard error of mean (SEM); Wilcoxon signed-rank test (SR test). ***p<0.001; **p<0.01; *p<0.05; n.s. p>0.05.

transition patterns as in the VR task (**Figures 3G-H**). We also observed a transient excitation of population activity immediately on context transition (**Figure 6H**, top), which reflected a sum over bidirectional, single-neuron responses (**Figures 6H-6I**). We ruled out contributions from drift artifacts by comparing population vector angles within vs. between contexts (**Figure 6J**). Again, consistent with the VR task, these were significantly larger between contexts, indicating that activity states are more separable than expected from context-independent fluctuations. Context identity was encoded in single neurons (**Figure 6K**) and drove separation of activity subspaces in dFC neurons (**Figure 6L**).

To examine the structure of decision coding in the FM task, we repeated our previous analysis to identify neurons that encoded model-fit DVs. Overall, 33% encoded the DV (**Figure 6M),** with the majority of those (91%) doing so only in one context but not the other. Stc- and Det-DV neurons exhibited ramping activity and reward deflections specifically on Stc or Det patches (**Figure 6N**). While 9% of DV-coding neurons were context-invariant, the dominant pattern of context-specific DV-coding suggests an orthogonal coding structure, again supporting the "subspace configuration" model. Consistent with this, DVs could be independently decoded from dFC activity (**Figure 6O**), and the angle between Stc-DV and Det-DV coding vectors was close to 90° (**Figure 6P**), indicating orthogonal geometry. Altogether, these results demonstrate that contextual configuration of orthogonal DV coding is a general neural solution for context-dependent decision making for foraging, preserved across experimental settings with different levels of constraint and ecological validity.

Decision dynamics are reconfigured immediately on context transitions

As most studies impose block structure or rule cues to indicate context, it remains unclear how neural circuitry is configured to enable flexible behavior when animals are free to adjust cognitive modes of their own volition, as would be the case in natural settings. We took advantage of our task design to address this, continuing to record dFC neurons into Phase 3 of the task ("volitional transitions", **Figures 7A-B**) and retaining 712 neurons across 16 sessions from the same mice (**Figure 7B**). We again confirmed encoding of context identity (**Figure 7C**) and context-specific DVs (**Figure 7D**) from dFC activity, indicating that these variables are stably represented even when block structure is removed and mice are unconstrained. Despite now switching contexts repeatedly each session (2-3 times/min, **Figure S1K**), orthogonal DV coding was preserved (**Figure 7E**), suggesting that mice "toggle" decision modes on context transitions and do not require blocks of repeated trial types to stabilize coding structure.

To explore this more precisely, we analyzed neural population trajectories (visualized with PC projections) during patch trials within and across contexts. Taking the Euclidean distance between trajectories as a difference metric (ΔTraj, **Figure 7F**), we found that average Stc vs. Det trajectories were more distinct than same-context trajectories (**Figure 7G**), consistent with the orthogonal DV coding structure (**Figures 5C, 6P, 7E**). Then, to home in on the immediate effect of transitions, we examined trajectories associated with sequential within-context (e.g. Det→Det) or across-context (e.g. Det→Stc) trial pairs. Visualization in example sessions suggested that single-trial trajectories were more distinct across transitions (**Figure 7H**), and quantitatively, ΔTraj for across-context trials sequences was greater than within-context comparisons (**Figure 7I**). This indicates that context-specific neural dynamics for foraging decisions are immediately

Figure 7

dFC dynamics exhibit rapid reconfiguration for decision flexibility

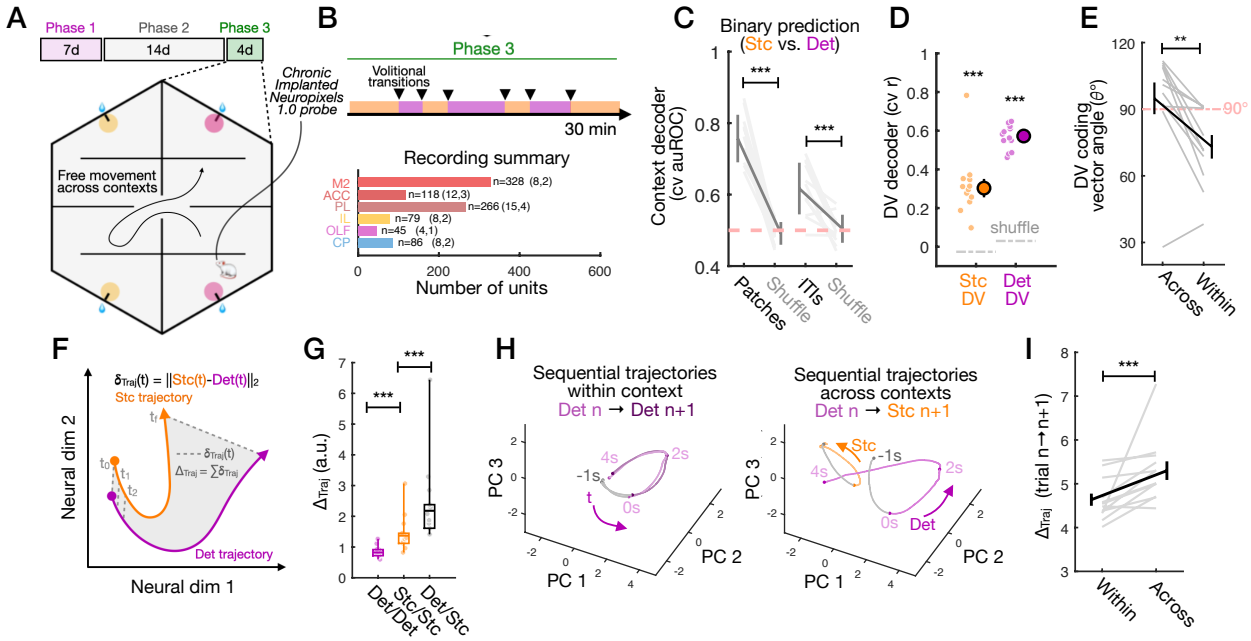

dFC activity is necessary for integrator configurations under volitional mode switching

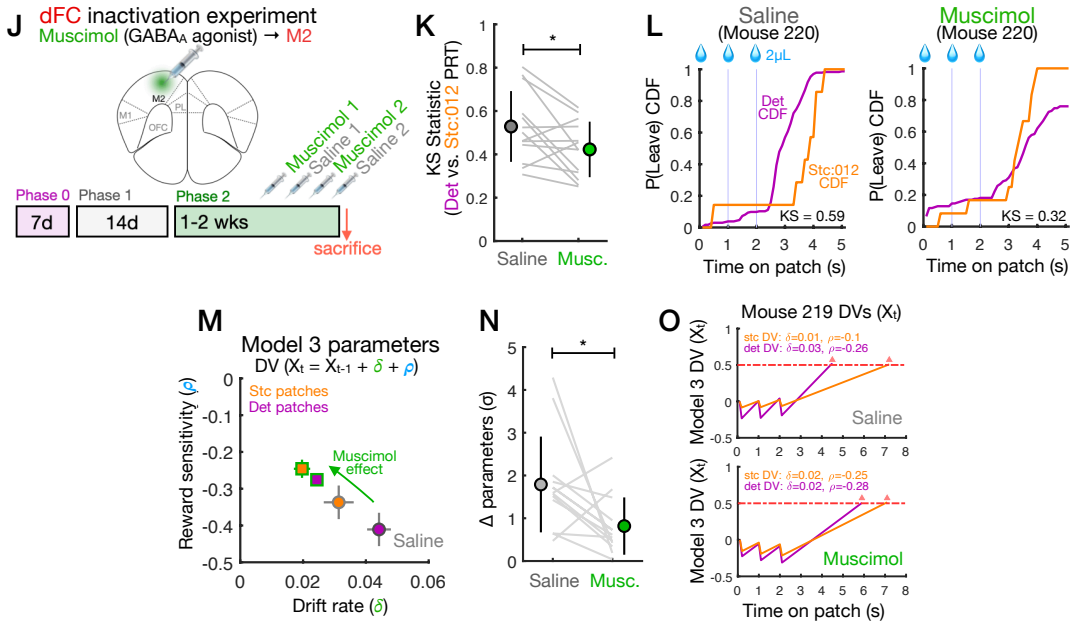

Figure 7. dFC supports volitional mode switching and rapid configuration of foraging decisions.

A) Schematic of chronic Neuropixels recording in Phase 3 of the free-moving foraging task when mice are free to switch environments volitionally.
B) Cartoon of volitional transitions in Phase 3 and recording summary (n = 712 dFC neurons, 16 sessions, 4 mice).
C) Performance (cross-validated *auROC*) of binary decoders for context identity on patch trials (left) and inter-trial interval (ITI) periods (right) (mean ± SEM, SR tests against shuffle, p = 4.88e-4 [trials], p = 0.0015 [ITIs], n = 12 sessions).
D) Performance (cross-validated *r*) of Stc-DV and Det-DV decoders compared to shuffle controls (mean ± SEM, SR tests against shuffle, p = 4.88e-4 [Stc DV], p =4 .88e-4 [Det DV], n = 12 sessions).
E) Angle θ between DV-coding vectors for across-context DVs vs. within-context control comparisons (SR test, p = 0.0049, n = 11 sessions).
F) Schematic illustrating neural trajectory analysis and quantification of ΔTraj metric (the sum of time point-matched euclidean distances between patch trial trajectories, see Methods).
G) Comparison of ΔTraj across within- and between-context trial-average trajectories (U tests, p = 4.77e-4 [within Det v.s. within Stc], p = 0.0035 [within Stc v.s. across contexts]).
H) Average of trajectories for Det→Det (left) and Det→Stc (right) sequential trial pairs from one example session (Mouse 251, session 17); arrow indicates direction of time evolution.
I) Comparison of session average ΔTraj for within- vs. between-context sequence pairs as visualized in **(H)** (median ± SEM, SR test, p = 9.76e-4, n = 12 sessions).
J) Schematic of dFC Muscimol inactivation experiments and timeline.
K) Comparison of Det vs. Stc:012 PRT distribution KS statistics for saline vs. Muscimol sessions (SR test, p=0.0398, n = 13 mice).
L) Det and Stc:012 patch trial PRTs (CDF of leaving probability over time) from example Mouse 220 under saline (left) and Muscimol (right) sessions, showing collapsed PRT distributions and reduced KS statistic under dFC inactivation.
M) Scatters of average parameter fits (***ρ*** and ***δ***) for Stc and Det patches under saline (grey) and Muscimol (green) (mean ± SEM).
N) Comparison of parameter differences (euclidean distance of z-scored values for ***ρ*** and ***δ***) for saline vs. muscimol sessions (SR test, p=0.0398, n = 13 mice).
O) DV rollouts for a Det patch reward sequences (t = 0, 1, 2) using parameter fits from example Mouse 219 for saline (top) or Muscimol (bottom) sessions.

Standard error of mean (SEM); Mann-Whitney U test (U test); Wilcoxon signed-rank test (SR test).
***p<0.001; **p<0.01; *p<0.05; n.s. p>0.05.

reconfigured upon volitional context switches and that, once learned, these decision modes are flexibly and rapidly deployable.

dFC is necessary for integrator configurations and volitional transitions

The striking consistency of neural coding across experiment designs led us to hypothesize that dFC is generally required for decision flexibility, even without imposed block structure or context cues. To test this, we repeated the Muscimol inactivation experiment to silence dFC neurons during Phase 3 of the FM task (**Figure 7J**). Like the previous result (**Figure 5M**), decision flexibility was significantly reduced under Muscimol compared to saline (**Figures 7K-L**). In the integrator model, this behavioral pattern should reflect a collapse of parameter configurations which underlie to distinct decision modes. Indeed, we found that $\boldsymbol{\delta}$ and $\boldsymbol{\rho}$ were reduced in magnitude under dFC inactivation (**Figure 7M**), and that the difference in parameter fits across Stc vs. Det patches was significantly smaller (**Figure 7N**). Plots of example DV rollouts confirmed the interpretation that Muscimol disrupted context-specific DV configurations (**Figure 7O**).

Finally, we observed some additional effects of dFC inactivation in the FM task. While mice still performed well, they were significantly less likely to make volitional port or context transitions under Muscimol (**Figures S9C-D**) and spent significantly more time on patch trials (**Figure S9E**). While this may indicate perseverative behavior, it is also consistent with the established role for dFC in tracking task performance and regulating explore-exploit modes,[18,51–53] which may be relevant for context transitions in our foraging task. Longer PRTs under Muscimol are also broadly consistent with frontal cortex involvement in biasing waiting behavior and sustaining persistent ramping dynamics underlying decision computations.[38,45,47,48,54–56]

Discussion

The question of how neural processes are configured by context to alter functionality is fundamental to our understanding of cognition and flexible behavior.[1,4] Here, we investigated how the brain organizes neural dynamics to support context-dependent decision making in the ecological setting of patch foraging. Our approach stands in contrast to classical designs, which often require animals to learn arbitrary sensory mapping rules for alternative-choice perceptual decisions, and typically locate context effects upstream of decision dynamics, operating to regulate attention or gate sensory inputs.[11–13,17,19] Instead, we found evidence of a strikingly different neural solution for circuit configuration: different foraging environments recruit brain-wide activity states, which organize decision dynamics into orthogonal modes instantiated in separate subpopulations of neurons. In our task, these neurons encoded two distinct configurations of an integrator decision variable (DV) that performs evidence accumulation for patch foraging in different environments. In the dorsal frontal cortex (dFC), these decision modes are rapidly toggled as animals transition between contexts, and activity in dFC is required for expression of context-specific decision strategies.

Configurations of a common integrator process explain context-specific decision behavior

Recent efforts to re-interpret classical foraging problems through the lens of decision theory have led to fresh insights and proposals for biologically realistic computational models and underlying neural mechanisms.[34,35] In line with this, we recently demonstrated that foraging decisions in mice are well described by an integrator model that weighs time and rewards as decision evidence (**Figure 1J**).[37] Ours and other studies have also identified neural dynamics encoding integrator DVs in the frontal cortex of rodents,[45,57,58] as well as parallel encoding of DVs that match different algorithms for foraging problems.[38,47,48] Taken together, this burgeoning literature is beginning to clarify the algorithmic logic and neural implementation of foraging decisions.[28,29,59] But beyond this, it raises new questions about how the brain may adapt these dynamics to implement different cognitive strategies or heuristics for ecological problems such as foraging. For example, different configurations of an integrator process can track time investment, time since last reward, decision evidence, and reward quantity (**Figure 1J**, Models 1-4).[34,60]

How, then, do animals contextually adapt such decision processes? Here, we addressed this by studying mice foraging for water rewards across two environment contexts: a Stochastic (Stc) context, which yields probabilistic rewards, and a Deterministic (Det) context, which yields a fixed number of rewards (**Figure S1A-C**). *A priori*, an intuitive hypothesis is that mice may switch between a "reward rate" computation (i.e. "Model 3") for Stc patches, and a "reward accumulation" computation (i.e. "Model 4") for Det patches. Instead, we found that mice solve this problem using the *same computational strategy* (**Figures S1K-L, 2I-J**) but with *different parameter configurations* (**Figures S1O-P, 2K-L).** Model fits to Det patches had larger values for both the drift rate $\boldsymbol{\delta}$ and reward sensitivity $\boldsymbol{\rho}$, and these were clearly reflected in the neural correlates of context-specific DVs (**Figure 4G-I**). In our task, mice apparently adapt this process to produce leaving behavior for Det patches that matches a "counting" strategy (leaving immediately after the final reward), but at the neural level, DVs do not exhibit features of reward accumulation (i.e. Model 4). This suggests that reconfiguration of a common evidence accumulation computation may be a general solution to calibrate foraging decisions to different reward distributions and contingencies.[34,60] It has been shown that rodents can exhibit rudimentary "counting" behavior.[38,61,62] Future studies could clarify whether, in the context of foraging, counting algorithms are instantiated by integrator configurations which could be flexibly deployed in settings where they are effective.

What circuit mechanisms underlie changes in parameter configuration for integrator dynamics? The drift rate $\boldsymbol{\delta}$ in our model effectively captures the weight of passing time as evidence *for* patch leaving, and is analogous to the same parameter in drift-diffusion models for perceptual decisions, where it can be modulated by the quality of sensory evidence, confidence, and value.[36,37,63–66] Mechanistically, this may reflect influence of "top-down" cortical signals which encode drift rates, organization of recurrent interactions (e.g. by thalamocortical inputs) which drive persistent activity, and/or neuromodulator signals. For the patch foraging problem, some promising candidate mechanisms include frontal or thalamic inputs, which bias waiting times for reward seeking[48,54] and modulate recurrent dynamics,[14,67] and 5-HT signaling, which has been linked to behavioral "patience" for uncertain rewards and could potentially tune drift rates.[68–70]

The reward sensitivity term $\boldsymbol{\rho}$ can be thought of as a weight of pulsatile sensory evidence *against* patch leaving decisions. As rewards have rapid influence on integrator dynamics (**Figure 4H**), this signal could be mediated, at least in part, by detection of water on the tongue/mouth through somatosensory relays. The gain (or sign) of this signal onto integrator activity could be modulated by top-down sensory gating mechanisms, local interneuron interactions, or calibrated

network dynamics. In general, what mechanisms control the "parameters" for integrator dynamics, and how these are learned from reward feedback (e.g. through reinforcement learning systems), are important and ripe domains for continued investigation.

Orthogonal coding of context-specific integrator dynamics in separate neuron populations

Another major aim of our study was to understand how neural activity is structured to support context-dependent decision making through configuration of integrator dynamics. To express context-specific behavior, context signals must organize different functional modes which alter how sensory information affects the decision process (**Figures 3A, 4A**).[8,9] Theoretically, this could happen by adjusting how sensory signals (e.g. rewards or sensory stimuli) are processed or how they perturb decision activity.[11,12,16,17,19] This latter mechanism, described as a "selection vector" produced by tuned network dynamics, has been proposed for context-dependent perceptual decisions in monkeys and rats.[11] Together, mechanisms which locate context effects "upstream" of a decision process can be thought of as "input configuration" solutions (**Figure 4A**). These have recently been unified in a theory that explains individual variability in flexibility as idiosyncratic expression of circuit-level configuration mechanisms.[19]

However, the input configuration model is, in general, agnostic as to how decision activity itself is structured. Alternatively (or additionally), the brain could use a "subspace configuration" strategy to structure different instantiations of a decision process for different contexts.[7,20,22,71] If context-specific DVs are represented in orthogonal coding vectors, then they can be tuned without interference, read out independently by downstream circuits, and flexibly "toggled" by control signals.[7,24,41,72] Indeed, subspace coding has been widely explored and evidenced in other domains, including to decorrelate sensory representations,[22,73] for cortical communication,[20] and to stay premotor activity from immediate expression.[45,74]

Most studies of context-dependent perceptual decisions posit various input configuration mechanisms coupled to an invariant decision process.[11–13,19] In contrast, we find widespread subspace configuration and orthogonal coding of context-specific DVs. Furthermore, orthogonal DV coding is instantiated at the cellular level in separate neuron subpopulations. In principle, this simple structure allows the brain to calibrate decision dynamics by activating subsets of neurons in different contexts, exposing them selectively to teaching signals and plasticity. Here, we only included two environments, but it remains to be clarified how scalable this neural solution is. For example, what is the "carrying capacity" for different configurations of decision dynamics? What are the constraints on what can be learned? And how are populations selected to implement context-specific computations?

More fundamentally, why do we observe orthogonal, context-specific decision coding when many other studies have not? One possibility is that neural solutions for flexible computation are highly task dependent.[8] For perceptual decision tasks, filtering or gating sensory signals onto common decision dynamics is a straightforward way to embed context contingencies – this may be especially so if such signals are sensory stimuli (e.g. lights or tones) that appear transiently or must be inferred from performance.[11–13,18,19] In our foraging task, however, context is not only continuously given through stable cues, but the notion of context is also less abstract, and perhaps more ecologically intuitive (for a mouse) as a spatially extended environment.[26] It is tempting to speculate that, when "context" is synonymous with "environment," which carries intrinsic

motivational meaning (as in a place to forage), the brain simply cares much more about it. Indeed, we found that environment identity (i.e. sensory context) was not only encoded in sensory areas but widely throughout the brain, including in frontal, subcortical, and thalamic structures.[8,73,75,76] Beyond linear context coding, foraging environments also recruit separate activity subspaces in many regions, which organize orthogonal integrator dynamics (**Figure 4M**). These observations are consistent with a picture of context-sensitivity as distributed rather than localized in the brain, yet graded such that some circuits are more "configurable" than others, perhaps operating as functional bottlenecks for flexible cognition and behavior.[8]

Placement of dFC in a network of context-sensitive configurable circuitry

In this study, we sampled widely across the brain to characterize how context shapes neural dynamics and integrator DVs (**Figures 3-4**), and we focused our investigation on the dFC to analyze the development of orthogonal coding structure (**Figure 5**), test causality, and study consistency across experimental settings (**Figures 6-7**). We observed encoding of context-specific DVs in dFC regions (including M2, ACC, and PL) and found that inactivation of M2 disrupted context sensitivity in both tasks. Interestingly, however, M2 inactivation did not disrupt foraging in the VR task, and only minorly affected engagement in the FM task, suggesting that M2 may be redundant for the core decision process. This is perhaps not surprising given the distributed nature of integrator dynamics[45,46,77] (**Figure 4J**), which could be sustained through large-scale recurrent interactions. However, it may suggest a more specific role for M2 in *selecting functional modes* based on context recognition. Indeed, some previous studies support a role for M2 in cognitive flexibility and in shaping ongoing behavior and action policies based on sensory contingencies.[78–80] The anatomy also supports this picture, as M2 receives substantial direct input from sensory areas and projects broadly to motor control circuitry to regulate movement initiation.[47,49,57,77]

However, M2 and dFC do not exist in isolation, and their contributions to behavior likely depend on interactions with a wide network of interconnected circuits. Indeed, although we focused on the dFC, our subspace alignment and coding analyses identified several other brain regions that exhibit configurable dynamics and context-specific DVs (**Figure 4M**), including sensory (olfactory areas, tenia tecta, visual cortex) and subcortical (the septum and striatum) areas which are connected with dFC. This raises questions about how these regions interact, how dynamics across the larger network are structured, and whether M2 or other dFC regions play an outsized role in contextually shaping network activity. Leveraging high-density recording technologies[40,50] to sample neural populations in these areas simultaneously, including with targeted perturbations in specific circuits or neural populations, will be illuminating.

Methodological consideration of ecological validity and complementary task designs

Finally, our approach takes seriously the limitations of both highly controlled and unstructured task designs (**Figures 2A, 6A**), as well as the importance of ecological relevance to drive interpretable behavior.[26,27,39] Motivated to address a fundamental question – how the brain adapts cognitive function to context – we took inspiration from the real-world challenges faced by animal subjects and designed a task that encourages cognitive flexibility for the ecological goal of foraging.[28,29,39] Rather than taking weeks or months to learn the task, animals naturally calibrated

their own flexible decision strategies within days. Practically, this allowed us to characterize behavior across dozens of animals and record widely across the brain, leading to insights (e.g. reconfiguration of a common decision process and brain-wide activity states) that may have eluded studies using more classical task designs with fewer animal subjects.

In the same vein, our study demonstrates how complementary task designs can be used synergistically to elucidate and validate models of neural circuit function. We began with a naturalistic foraging task and closely studied the behavior to understand, computationally, how mice contextually adapt their foraging strategies. We then moved to a "reduced preparation" using a virtual, head-fixed task to widely sample neural activity and characterize coding structures that support contextual configuration of integrator dynamics. Finding clear evidence for subspace configuration, we returned to the free-moving setting to study how this neural solution is used under conditions that better match real-world experience. Thus, the dual-task approach not only gave us multiple perspectives on the same problem but also allowed us to test consistency of the neural mechanisms underlying flexible computation under different experimental settings. Taken together, our work reveals a general neural basis for context-dependent foraging decisions in orthogonal integrator dynamics and paves the way for future investigation into mechanisms of complex cognition through the lens of ecologically grounded behavior.

Figure S1

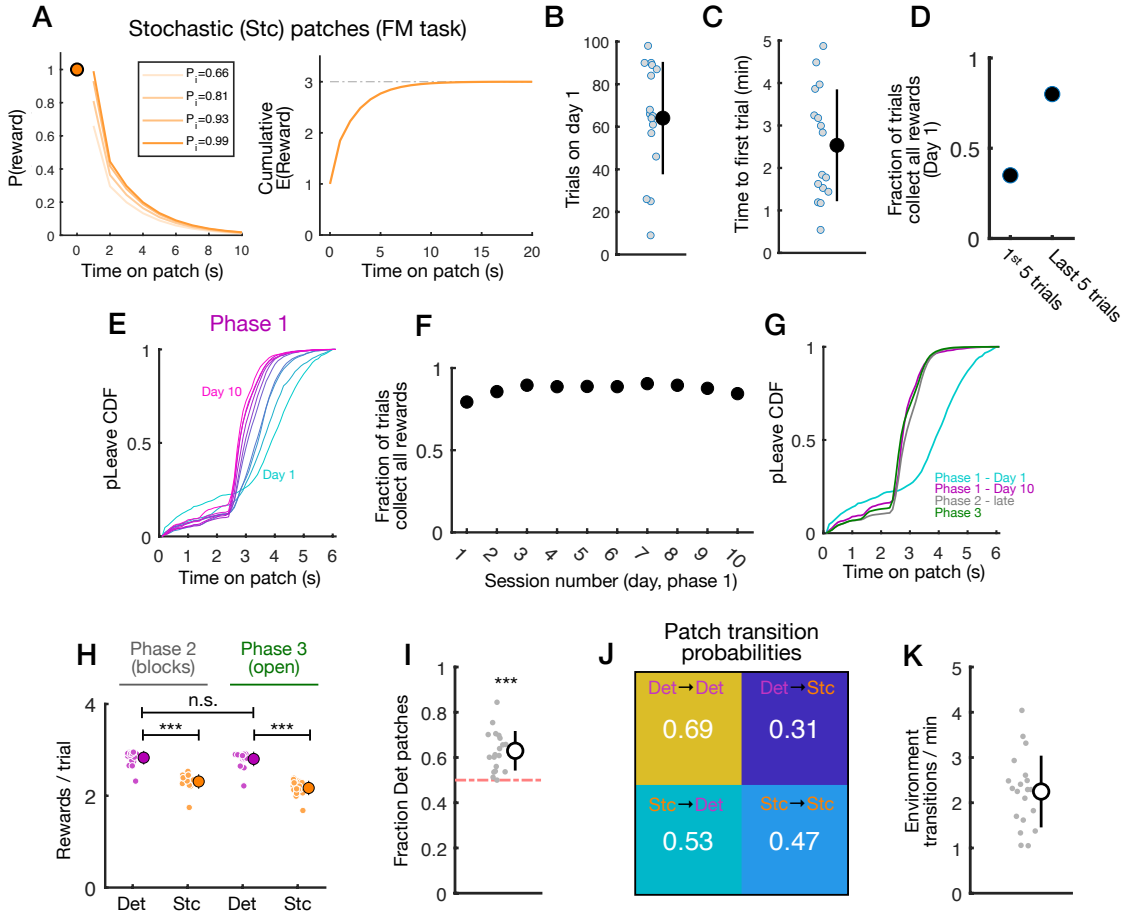
A
Stochastic (Stc) patches (FM task)
P(reward)
Time on patch (s)
Cumulative E(Reward)
Time on patch (s)
B
Trials on day 1
C
Time to first trial (min)
D
Fraction of trials collect all rewards (Day 1)
1st 5 trials
Last 5 trials
E
Phase 1
pLeave CDF
Day 10
Day 1
Time on patch (s)
F
Fraction of trials collect all rewards
Session number (day, phase 1)
G
pLeave CDF
Phase 1 - Day 1
Phase 1 - Day 10
Phase 2 - late
Phase 3
Time on patch (s)
H
Phase 2 (blocks)
Phase 3 (open)
n.s.
***
***
Rewards / trial
Det
Stc
Det
Stc
I
***
Fraction Det patches
J
Patch transition probabilities
Det→Det
0.69
Det→Stc
0.31
Stc→Det
0.53
Stc→Stc
0.47
K
Environment transitions / min

Figure S1. Additional analysis of behavior through foraging task progression.

A) Left, reward probability over time for stochastic patches. Each trial yields one reward at t=0s, and then one each second after with decaying probability starting with random P_i = 0.66, 0.81, 0.93, or 0.99. Right, cumulative expected reward over time for stochastic patches.
B) Number of patch trials performed by naive mice in the first session on Phase 1 of the foraging task (each dot one mouse, mean ± s.d., n = 16 mice).
C) Latency (minutes) to the first patch trial in the first session of the task (each dot one mouse, mean ± s.d., n = 16 mice).
D) Fraction of trials among the first 5 in the first session (left) or the last 5 in the first session (right) in which mice collected all three deterministic rewards (data pooled across mice).
E) CDF of leaving probabilities for trials on each session (days 1-10) in Phase 1 (data pooled across mice).
F) Fraction of trials in which mice collected all three rewards for each session (days 1-10) in Phase 1.
G) CDF of leaving probabilities for deterministic trials on day 1 in Phase 1 (cyan), day 10 in Phase 1 (purple), late (last 4 days) in Phase 2 (grey), and Phase 3 (green) showing similar PRT distributions across all phases.
H) Average number of rewards collected per trial for Det and Stc trials in Phase 2 and Phase 3 (each dot one mouse, mean ± s.d., SR test, p = 8.86e-5 (P2 Stc vs. Det), p = 8.86e-5 (P3 Stc vs. Det), p = 0.0522 (P2 Det vs. P3 Det), n = 20 mice).
I) Choice probability for Det patches in Phase 3 of the task (each dot one mouse, mean ± s.d., t-test > 0.5, p = 2.4e-6, n = 20 mice).
J) Patch type transition matrix for Det and Stc patches in Phase 3 of the task (data pooled across mice).
K) Volitional context transitions per minute in Phase 3 (each dot one mouse, mean ± s.d.).

Wilcoxon signed rank test (SR test); ***p<0.001; **p<0.01; *p<0.05; n.s. p>0.05.

Figure S2

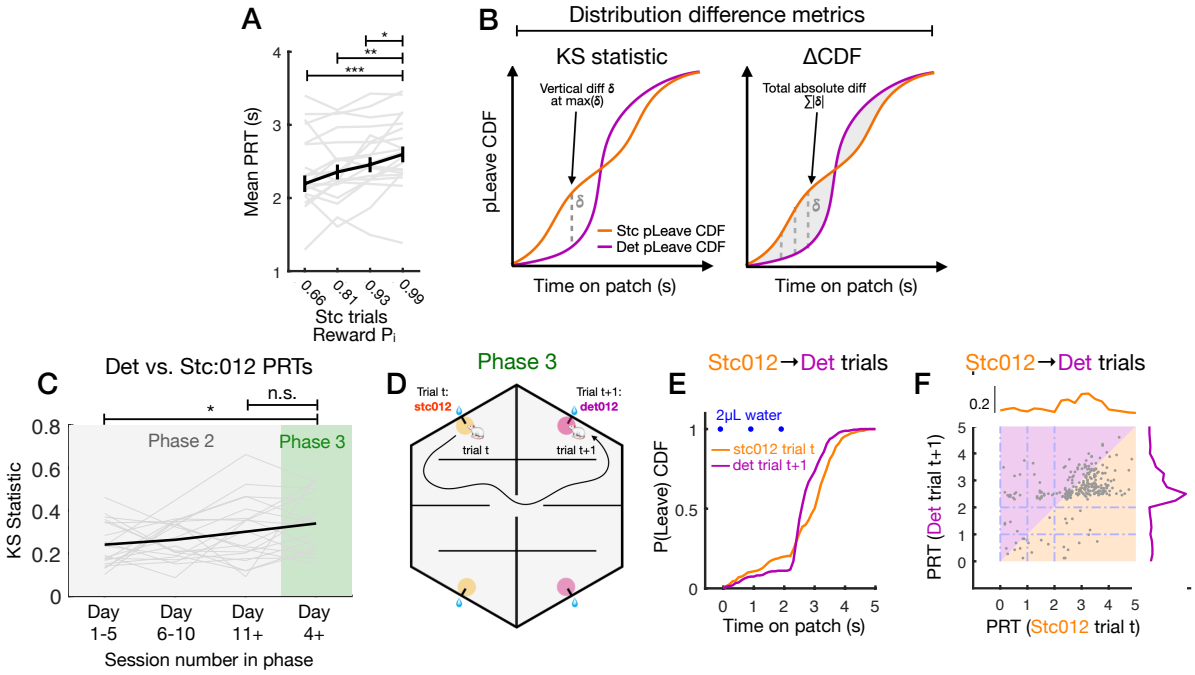
A
Mean PRT (s)
Stc trials
Reward Pi
0.66
0.81
0.93
0.99
B
Distribution difference metrics
KS statistic
ΔCDF
Vertical diff δ at max(δ)
Total absolute diff ∑|δ|
pLeave CDF
Stc pLeave CDF
Det pLeave CDF
Time on patch (s)
C
Det vs. Stc:012 PRTs
n.s.
Phase 2
Phase 3
KS Statistic
Day 1-5
Day 6-10
Day 11+
Day 4+
Session number in phase
D
Phase 3
Trial t: stc012
Trial t+1: det012
trial t
trial t+1
E
Stc012→Det trials
2μL water
stc012 trial t
det trial t+1
P(Leave) CDF
Time on patch (s)
F
Stc012→Det trials
PRT (Det trial t+1)
PRT (Stc012 trial t)

Figure S2. Additional analysis of decision behavior and context effects.

A) Mean patch residence times (PRT, seconds) for hidden Stc trial types with different reward probabilities P_i = 0.66, 0.81, 0.93, or 0.99 (**Figure S1A**, mean ± s.d.; SR test, p = 8.85e-5 (0.66 vs. 0.99), p = 0.004, (0.81 vs. 0.99), p = 0.0479 (0.93 vs. 0.99), n = 20 mice).
B) Diagram showing difference metrics for comparing PRT distributions. The Kolmogorov-Smirnov (KS) statistic measures the maximum distance δ between CDFs (left), and ΔCDF measures the total absolute difference between CDFs (right).
C) KS statistic comparing Det vs. Stc:012 PRT distributions across early (days 1-5), middle (days 6-10), and late (days 11+) Phase 2 progression and Phase 3 (SR test, p = 0.0438 (early P2 vs. P3), p =0 .3703 (late P2 vs. P3), n = 20 mice).
D) Cartoon showing example of a Stc:012 (trial t) followed by Det (trial t+1) sequence in Phase 3.
E) pLeave CDF comparing PRT distributions for Stc:012→Det trial sequences illustrated in **(d)** (KS test, p = 8.83e-8, n = 271 trials).
F) Scatter plot showing PRTs for Stc:012→Det trial sequences in Phase 3 and histograms of PRT distributions.

Kolmogorov-Smirnov test (KS test); Wilcoxon signed-rank test (SR test); ***p<0.001; **p<0.01; *p<0.05; n.s. p>0.05.

Figure S3

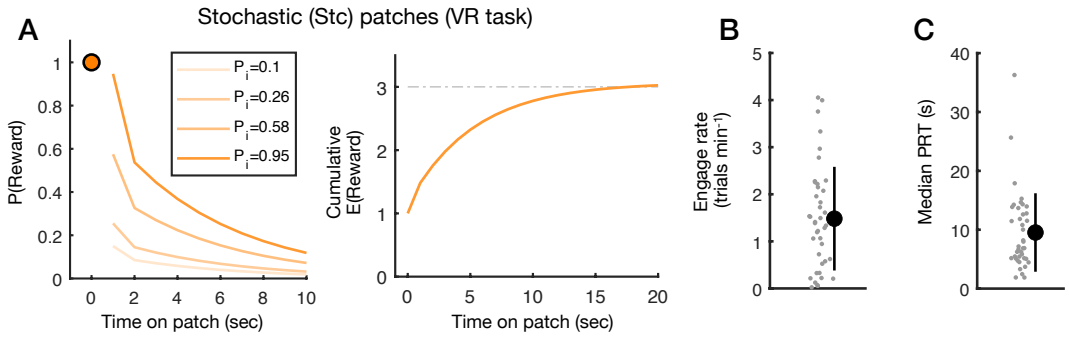
A
Stochastic (Stc) patches (VR task)
P(Reward)
Time on patch (sec)
P_i=0.1
P_i=0.26
P_i=0.58
P_i=0.95
Cumulative E(Reward)
Time on patch (sec)
B
Engage rate (trials min⁻¹)
C
Median PRT (s)

Figure S3. Additional details of virtual foraging task and behavior.

A) Left, reward probability over time for stochastic patches. Each trial yields one reward at t=0s, and then one reward each second with a decaying probability starting at t=1s with random P_i = 0.15, 0.255, 0.575, 0.947. Right, cumulative expected reward over time for stochastic patches.
B) Patch engagement rate (trials per minute) for virtual (VR) foraging task (each dot one mouse, mean ± s.d.).
C) Median patch residence time (PRT) for virtual foraging task (each dot one mouse, mean ± s.d.).

Figure S4

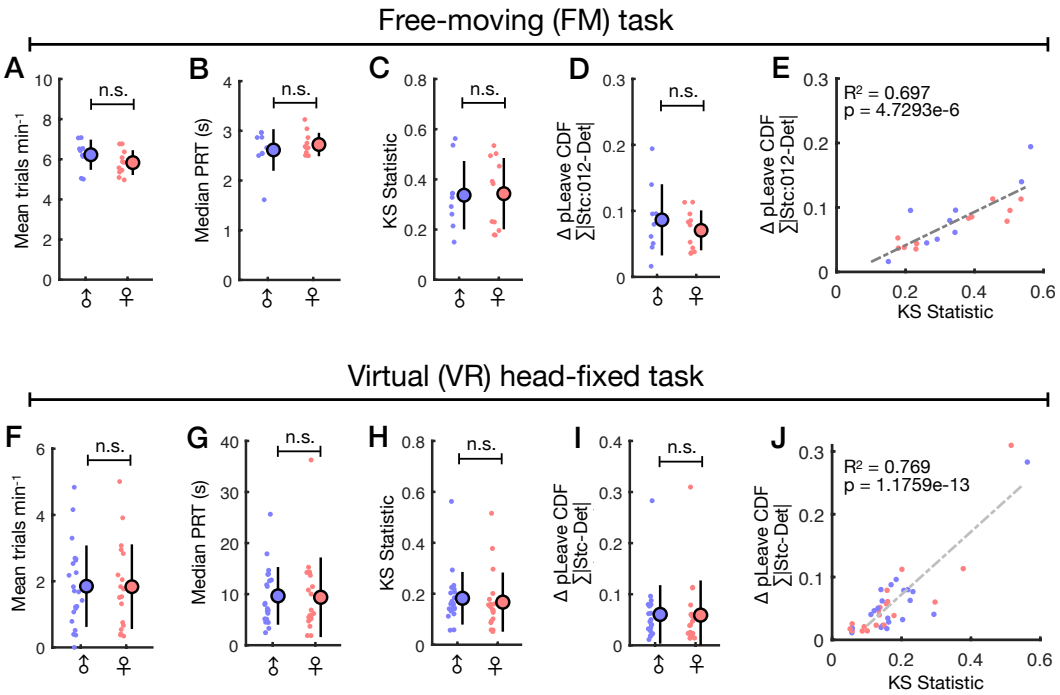
Free-moving (FM) task
A
B
C
D
E
n.s.
Mean trials min-1
Median PRT (s)
KS Statistic
Δ pLeave CDF
∑|Stc:012-Det|
R² = 0.697
p = 4.7293e-6
KS Statistic
Virtual (VR) head-fixed task
F
G
H
I
J
Δ pLeave CDF
∑|Stc-Det|
R² = 0.769
p = 1.1759e-13

Figure S4. Lack of sex differences in foraging behavior and decision flexibility.

A) Comparison of patch engagement rate (trials per minute) between male vs. female mice in the free-moving (FM) foraging task (mean ± s.d., U test, $p = 0.1965$, $n = 9$ males, $n = 11$ females).
B) Comparison of median patch residence time (PRT, seconds) between male vs. female mice in the FM task (mean ± s.d., U test, $p = 0.9697$).
C) Comparison of KS statistic for Stc:012 vs. Det patch PRTs between male vs. female mice in the FM task (mean ± s.d., U test, $p = 1$).
D) Comparison of ΔCDF for Stc:012 vs. Det patch PRTs between male vs. female mice in the FM task (mean ± s.d., U test, $p = 0.4941$).
E) Correlation between KS **(C)** statistic and ΔCDF **(D)** for all mice in the FM task (linear regression, $R^2 = 0.69$, $p = 4.73e-6$, $n = 20$ mice).
F) Comparison of patch engagement rate between male vs. female mice in the virtual (VR) foraging task (mean ± s.d., U test, $p = 0.9296$, $n = 21$ males, $n = 19$ females).
G) Comparison of median patch residence time between male vs. female mice in the VR task (mean ± s.d., U test, $p = 0.5513$).
H) Comparison of KS statistic for Stc vs. Det patch PRTs between male vs. female mice in the VR task (mean ± s.d., U test, $p = 0.203$).
I) Comparison of ΔCDF for Stc vs. Det patch PRTs between male vs. female mice in the VR task (mean ± s.d., U test, $p = 0.5157$).
J) Correlation between KS statistic **(H)** and ΔCDF **(I)** for all mice in the VR task (linear regression, $R^2 = 0.77$; $p = 1.16e-13$, $n = 40$ mice).

For all plots, each dot is one mouse (blue, males; red, females). Mann-Whitney U test (U test); ***$p<0.001$; **$p<0.01$; *$p<0.05$; n.s. $p>0.05$.

Figure S5

A

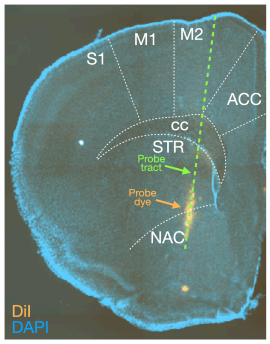

Mouse 160; Bregma +1.35 AP

Sampled regions:
Secondary motor cortex (M2)
Anterior cingulate (ACC)
Striatum (STR)
Nucleus Accumbens (NAC)

B

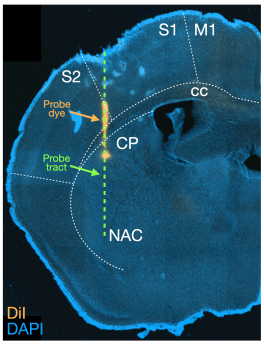

Mouse 257; Bregma -0.3 AP

Sampled regions:
Somatosensory areas (SS)
Caudoputamen (CP)
Nucleus Accumbens (NAC)

C

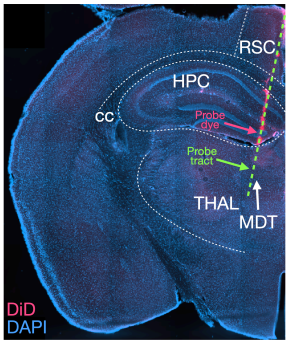

Mouse 283; Bregma -1.85 AP

Sampled regions:
Retrosplenial cortex (RSC)
Hippocampus CA1
Dentate Gyrus (DG)
Habenula (HAB)
Mediodorsal thalamus (MDT)
Centromedial thalamus (CMT)

D

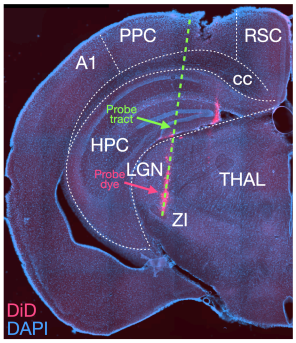

Mouse 283; Bregma -2.65 AP

Sampled regions:
Posterior parietal cortex (PPC)
Hippocampus CA1
Dentate Gyrus (DG)
Lateral geniculate (LGN)

Figure S5. Examples of Neuropixels probe insertions.

A) Example histology (mouse 160) showing a 1-shank Neuropixels probe insertion targeting frontal and basal ganglia regions (M2, ACC, STR, NAC).
B) Example histology (mouse 257) showing a probe insertion targeting sensory cortex and basal ganglia regions (SS, CP, NAC).
C) Example histology (mouse 283) showing a probe insertion targeting sensory cortex, hippocampal, subcortical and anterior thalamus regions (RSC, CA1, DG, HAB, MDT, CMT).
D) Example histology (mouse 283) showing a probe insertion targeting sensory cortex, hippocampal, and posterior thalamus regions (PPC, CA1, DG, LGN).

For all images, blue = DAPI, orange/red = probe dye (DiI/DiD), green dashes = probe tract)

Figure S6

A — Acute recordings in virtual (VR) task (n = 10 mice, 47 sessions)

			dFC			vFC			SC						HC			BG			THL							SUB						
Mouse	Exp #	Probe	M2	ACC	PL	IL	OFC	DP	TT	OLF	SS	VIS	PPC	RSC	DG	CA1	CA3	STR	CP	NAC	LGN	MDT	CMT	VPT	PNT	PVT	PFS	LSR	LSV	MSP	PCN	HAB	HYP	MID
83	16	1 shank	.	31	.	.	.	.	.	.	.	.	.	.	.	.	.	31	.	78	.	.	.	.	.	.	.	51	4	.	.	.	.	.
83	17	1 shank	24	123	.	.	.	.	.	.	.	.	.	.	.	.	.	6	.	.	.	.	.	.	.	.	.	24	.	.	.	.	.	.
83	18	1 shank	3	15	.	.	.	.	.	.	.	.	.	.	.	.	.	15	.	.	.	.	.	.	.	.	.	44	23	.	.	.	26	.
83	19	1 shank	9	15	.	.	.	.	.	.	.	.	.	.	.	.	.	3	.	.	.	.	.	.	.	.	.	51	17	.	.	.	.	.
88	14	1 shank	.	.	1	.	.	3	10	2	.	.	.	.	.	.	.	13	.	.	.	.	.	.	.	.	.	17	.	34	.	.	.	.
88	15	1 shank	.	12	4	9	.	46	41	1	.	.	.	.	.	.	.	26	.	.	.	.	.	.	.	.	.	59	.	10	.	.	.	.
88	16	1 shank	.	4	25	31	.	30	27	4	.	.	.	.	.	.	.	31	.	.	.	.	.	.	.	.	.	41	.	14	.	.	.	.
88	17	1 shank	.	9	.	.	.	.	.	.	.	.	.	.	.	.	.	27	.	.	.	.	.	.	.	.	.	60	21	.	.	.	59	.
88	18	1 shank	.	13	.	.	.	.	.	.	.	.	.	.	.	.	.	20	.	.	.	.	.	.	.	.	.	24	37	.	.	.	13	.
88	19	1 shank	.	53	.	.	.	.	.	.	.	.	.	.	.	.	.	12	.	.	.	.	.	.	.	.	.	58	9	.	.	.	42	.
116	10	1 shank	96	2	14	.	25	.	.	19	.	.	.	.	.	.	.	12	.	110	.	.	.	.	.	.	.	.	.	.	.	.	.	.
116	11	1 shank	57	8	22	.	44	.	.	24	.	.	.	.	.	.	.	32	.	80	.	.	.	.	.	.	.	.	.	.	.	.	.	.
116	12	1 shank	.	12	5	.	.	3	13	1	.	.	.	.	.	.	.	48	.	.	.	.	.	.	.	.	.	129	.	19	.	.	.	.
116	14	1 shank	.	53	95	33	.	55	40	4	.	.	.	.	.	.	.	31	.	2	.	.	.	.	.	.	.	44	.	.	.	.	.	.
116	15	1 shank	203	.	.	.	30	.	.	30	.	.	.	.	.	.	.	20	.	24	.	.	.	.	.	.	.	.	.	.	.	.	.	.
119	13	1 shank	.	5	140	181	.	38	54	8	.	.	.	.	.	.	.	26	.	5	.	.	.	.	.	.	.	.	.	.	.	.	.	.
119	14	1 shank	2	50	141	163	.	26	64	4	.	.	.	.	.	.	.	25	.	2	.	.	.	.	.	.	.	.	.	.	.	.	.	.
119	15	1 shank	39	42	132	129	.	66	58	5	.	.	.	.	.	.	.	11	.	.	.	.	.	.	.	.	.	15	.	.	.	.	.	.
257	19	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	59	.	.	.	6	31	.	.	.	2	.	.	.
257	20	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	106	.	.	.	8	.	.	.	.	20	7	.	42
257	21	4 shank	.	.	.	.	.	.	.	.	.	79	7	73	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
257	22	1 shank	.	.	.	.	.	.	.	.	13	.	.	.	.	.	.	.	182	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
258	18	4 shank	.	.	.	.	.	.	.	.	.	.	.	40	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
258	19	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	15	.	.	.	.	.	.	.	.	.	.	.	1
258	20	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	30	.	.	.	13	16	.	.	.	17	.	.	.
258	22	1 shank	.	.	.	.	.	.	.	.	25	.	.	.	.	.	.	.	48	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
261	17	4 shank	.	.	.	.	.	.	.	.	.	.	.	145	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
261	19	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	71	.	.	.	12	11	.	.	.	.	.	.	13
261	20	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	32	.	2	.	.	15	.	.	.	.	.	.	27
261	21	4 shank	.	.	.	.	.	.	.	.	.	27	.	21	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
261	22	1 shank	.	.	.	.	.	.	.	.	7	.	.	.	.	.	.	3	111	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
262	18	4 shank	.	.	.	.	.	.	.	.	.	.	.	23	1	1	4	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
262	19	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	77	10	.	.	2	1	.	.	.	.	.	.	.
262	20	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	34	.	.	.	.	29	.	.	.	.	.	.	.
262	21	4 shank	.	.	.	.	.	.	.	.	.	2	5	17	.	10	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
262	22	1 shank	.	.	.	.	.	.	.	.	25	.	.	.	.	.	.	.	156	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
283	17	4 shank	.	.	.	.	.	.	.	.	7	.	.	.	.	.	.	.	143	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
283	18	4 shank	.	.	.	.	.	.	.	.	.	.	28	.	24	26	20	.	.	.	103	3	.	.	.	.	.	.	.	.	.	.	.	.
283	19	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	12	74	.	.	.	.	.	.	.	.	203	.	.	.	.	.	.	.	.	.
283	21	4 shank	.	.	.	.	.	.	.	.	.	.	.	26	.	.	.	.	.	.	.	42	25	.	.	1	.	.	.	.	.	.	.	.
283	22	1 shank	.	.	.	.	.	.	.	.	.	.	.	4	20	.	9	.	.	.	.	5	55	.	.	.	.	.	.	.	.	12	.	.
283	23	4 shank	.	.	.	.	.	.	.	.	.	.	.	1	.	.	.	.	.	.	.	47	.	.	.	10	.	.	.	.	.	6	.	.
283	24	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	7	23	.	.	.	.	.	36	.	.	.	.	.	.	.	.	.	28	.	.
283	25	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	31	27	.	.	.	.	.	4	.	.	.	.	41	.	.	.	.	.	.	.
284	17	4 shank	.	.	.	.	.	.	.	.	118	.	.	.	.	.	.	.	51	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
284	18	4 shank	.	.	.	.	.	.	.	.	.	.	.	.	12	10	2	.	.	.	27	3	.	77	18	.	.	.	.	.	.	.	.	.
284	19	4 shank	.	.	.	.	.	.	.	.	.	.	5	.	18	29	.	.	.	.	.	.	.	5	177	.	.	.	.	.	.	.	.	.

B — Acute recordings in VR task (n = 6 naive nice, 17 sessions)

			dFC			vFC			SC						HC			BG			THL							SUB						
Mouse	Exp #	Probe	M2	ACC	PL	IL	OFC	DP	TT	OLF	SS	VIS	PPC	RSC	DG	CA1	CA3	STR	CP	NAC	LGN	MDT	CMT	VPT	PNT	PVT	PFS	LSR	LSV	MSP	PCN	HAB	HYP	MID
166	2	1 shank	65	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	107	21	.	.	.	.	.	.	.	.	.	.	.	.	.	.
166	3	1 shank	87	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	87	11	.	.	.	.	.	.	.	.	.	.	.	.	.	.
167	1	1 shank	.	7	7	.	.	.	.	.	.	.	.	.	.	.	.	.	102	12	.	.	.	.	.	.	.	.	.	.	.	.	.	.
167	2	1 shank	1	.	11	.	.	.	.	.	.	.	.	.	.	.	.	.	74	5	.	.	.	.	.	.	.	.	.	.	.	.	.	.
167	3	1 shank	.	11	19	.	.	.	.	.	.	.	.	.	.	.	.	.	102	18	.	.	.	.	.	.	.	.	.	.	.	.	.	.
168	1	1 shank	11	17	29	.	.	.	.	.	.	.	.	.	.	.	.	.	90	32	.	.	.	.	.	.	.	.	.	.	.	.	.	.
168	2	1 shank	8	18	19	.	.	.	.	.	.	.	.	.	.	.	.	.	76	37	.	.	.	.	.	.	.	.	.	.	.	.	.	.
168	3	1 shank	7	2	4	.	.	.	.	.	.	.	.	.	.	.	.	.	57	13	.	.	.	.	.	.	.	.	.	.	.	.	.	.
169	1	1 shank	23	7	1	.	.	.	.	.	.	.	.	.	.	.	.	.	5	1	.	.	.	.	.	.	.	.	.	.	.	.	.	.
169	2	1 shank	37	12	2	.	.	.	.	.	.	.	.	.	.	.	.	.	31	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
169	3	1 shank	28	22	1	.	.	.	.	.	.	.	.	.	.	.	.	.	51	5	.	.	.	.	.	.	.	.	.	.	.	.	.	.
170	1	1 shank	25	6	.	.	.	.	.	.	.	.	.	.	.	.	.	.	40	55	.	.	.	.	.	.	.	.	.	.	.	.	.	.
170	2	1 shank	32	8	13	.	.	.	.	.	.	.	.	.	.	.	.	.	31	7	.	.	.	.	.	.	.	.	.	.	.	.	.	.
170	3	1 shank	2	1	7	.	.	.	.	.	.	.	.	.	.	.	.	.	35	9	.	.	.	.	.	.	.	.	.	.	.	.	.	.
171	1	1 shank	40	.	.	.	66	.	.	21	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
171	2	1 shank	48	.	2	.	112	.	.	23	.	.	.	.	.	.	.	.	.	18	.	.	.	.	.	.	.	.	.	.	.	.	.	.
171	3	1 shank	60	.	5	.	109	.	.	35	.	.	.	.	.	.	.	.	.	33	.	.	.	.	.	.	.	.	.	.	.	.	.	.

C — Chronic recordings in free-moving (FM) task (n = 4 animals, 32 sessions)

			dFC			vFC			SC						HC			BG			THL							SUB						
Mouse	Exp #	Probe	M2	ACC	PL	IL	OFC	DP	TT	OLF	SS	VIS	PPC	RSC	DG	CA1	CA3	STR	CP	NAC	LGN	MDT	CMT	VPT	PNT	PVT	PFS	LSR	LSV	MSP	PCN	HAB	HYP	MID
248	11	1 chronic	40	11	12	.	.	.	.	.	.	.	.	.	.	.	.	.	11	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
248	12	1 chronic	37	12	10	.	.	.	.	.	.	.	.	.	.	.	.	.	11	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
248	13	1 chronic	46	9	10	.	.	.	.	.	.	.	.	.	.	.	.	.	13	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
248	14	1 chronic	41	12	9	.	.	.	.	.	.	.	.	.	.	.	.	.	10	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
248	15	1 chronic	46	11	13	.	.	.	.	.	.	.	.	.	.	.	.	.	15	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
248	16	1 chronic	37	11	9	.	.	.	.	.	.	.	.	.	.	.	.	.	10	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
248	17	1 chronic	41	13	9	.	.	.	.	.	.	.	.	.	.	.	.	.	10	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
248	18	1 chronic	31	17	13	.	.	.	.	.	.	.	.	.	.	.	.	.	7	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
249	11	1 chronic	1	.	.	2	.	.	.	12	.	.	.	.	.	.	.	.	15	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
249	12	1 chronic	1	1	2	5	.	.	.	12	.	.	.	.	.	.	.	.	9	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
249	13	1 chronic	1	.	2	4	.	.	.	11	.	.	.	.	.	.	.	.	13	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
249	14	1 chronic	2	.	2	5	.	.	.	10	.	.	.	.	.	.	.	.	9	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
249	15	1 chronic	.	.	1	5	.	.	.	9	.	.	.	.	.	.	.	.	11	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
249	16	1 chronic	.	.	.	4	.	.	.	11	.	.	.	.	.	.	.	.	12	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
249	17	1 chronic	.	.	2	3	.	.	.	13	.	.	.	.	.	.	.	.	10	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
249	18	1 chronic	.	.	2	5	.	.	.	12	.	.	.	.	.	.	.	.	11	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
250	11	1 chronic	.	22	45	11	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
250	12	1 chronic	.	24	54	18	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
250	13	1 chronic	.	26	55	18	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
250	14	1 chronic	.	16	46	17	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
250	15	1 chronic	.	13	62	12	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
250	16	1 chronic	.	14	42	6	.	.	.	.	.	.	.	.	.	.	.	1	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
250	17	1 chronic	.	14	50	22	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
250	18	1 chronic	.	7	55	22	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
251	11	1 chronic	61	5	2	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
251	12	1 chronic	60	5	2	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
251	13	1 chronic	68	8	2	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
251	14	1 chronic	54	4	2	1	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
251	15	1 chronic	47	6	1	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
251	16	1 chronic	34	4	4	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
251	17	1 chronic	45	4	1	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.
251	18	1 chronic	47	4	2	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.

Figure S6. Additional information about Neuropixels recordings.

A) Extended summary table for acute Neuropixels recordings in trained mice performing the virtual (VR) foraging task, showing mouse, session number, probe type (1-shank/4-shanks), and number of recorded neurons from each region (n = 47 recording sessions from 10 mice).
B) Extended summary table for acute Neuropixels recordings in task-naive (first 3 days) mice performing the VR task as in **(A)** (n = 17 recording sessions from 6 mice).
C) Extended summary table for acute Neuropixels recordings in mice performing the free-moving (FM) foraging task showing details as in **(A-B)** (n = 32 recording sessions from 4 mice).

Figure S7

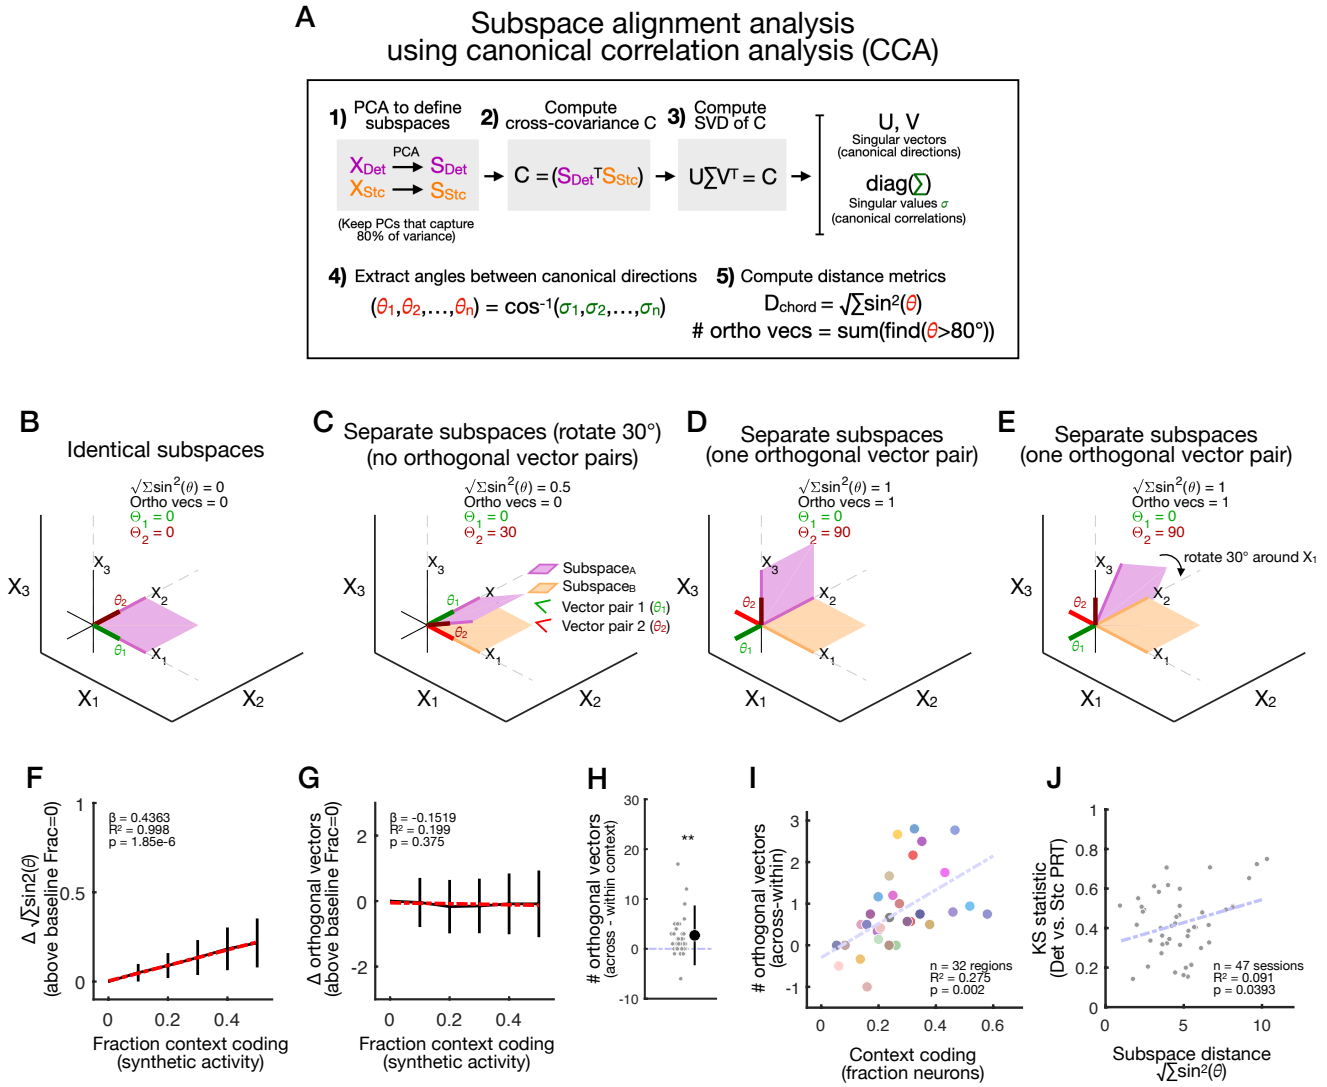
A
Subspace alignment analysis
using canonical correlation analysis (CCA)
1) PCA to define subspaces
2) Compute cross-covariance C
3) Compute SVD of C
PCA
(Keep PCs that capture 80% of variance)
U, V
Singular vectors (canonical directions)
diag(∑)
Singular values σ (canonical correlations)
4) Extract angles between canonical directions
5) Compute distance metrics
ortho vecs = sum(find(θ>80°))
B
Identical subspaces
C
Separate subspaces (rotate 30°)
(no orthogonal vector pairs)
D
Separate subspaces
(one orthogonal vector pair)
E
Separate subspaces
(one orthogonal vector pair)
rotate 30° around X1
Subspace A
Subspace B
Vector pair 1 (θ1)
Vector pair 2 (θ2)
F
β = 0.4363
R² = 0.998
p = 1.85e-6
Δ √∑sin2(θ) (above baseline Frac=0)
Fraction context coding (synthetic activity)
G
β = -0.1519
R² = 0.199
p = 0.375
Δ orthogonal vectors (above baseline Frac=0)
Fraction context coding (synthetic activity)
H
orthogonal vectors (across - within context)
I
# orthogonal vectors (across-within)
n = 32 regions
R² = 0.275
p = 0.002
Context coding (fraction neurons)
J
KS statistic (Det vs. Stc PRT)
n = 47 sessions
R² = 0.091
p = 0.0393
Subspace distance √∑sin²(θ)

Figure S7. Description of subspace alignment analysis and additional controls.

A) Graphical description of subspace alignment analysis pipeline. For each session, principal component analysis (PCA) is used on neuron spiking data within each context to determine subspaces S_{Det} and S_{Stc}, and canonical correlation analysis (CCA) is used to find canonical correlations and compute subspace distance metrics.
B) Example visualization of subspace alignment analysis applied to identical 2D subspaces (green = 1st pair of canonical directions; red = 2nd pair of canonical directions; θ = angle between vector pairs).
C) Visualization of analysis applied to separate 2D subspaces with 30° angle between them and one shared canonical direction.
D) Visualization of analysis applied to separate 2D subspaces with one orthogonal pair of canonical directions and one shared canonical direction.
E) Visualization after rotating (30°) one subspace in **(D)**, maintaining the same subspace distance.
F) Relationship between chordal distance $\sqrt{\sum \sin^2(\theta)}$ and addition of synthetic context coding activity in single neurons (fraction of context-coding neurons, linear regression, $\beta = 0.44$, $R^2 = 0.99$, $p = 1.85e\text{-}6$).
G) Relationship between orthogonal vectors and addition of synthetic context coding activity in single neurons (linear regression, $\beta = -0.1519$, $R^2 = 0.199$, $p = 0.375$).
H) Normalized (between - within) number of orthogonal vectors for Det- vs. Stc-context subspaces (mean ± s.d., t-test, $p = 0.0034$, n = 47 sessions).
I) Correlation across regions between context-coding (fraction of neurons) and number of orthogonal vectors between context subspaces ($R^2 = 0.275$, $p = 0.002$, n = 32 regions).
J) Correlation across sessions between subspace distance $\sqrt{\sum \sin^2(\theta)}$ and decision flexibility, indexed as the KS statistic between Det vs. Stc PRT distributions ($R^2 = 0.091$, $p = 0.039$, n = 47 sessions).

***$p<0.001$; **$p<0.01$; *$p<0.05$; n.s. $p>0.05$.

Figure S8

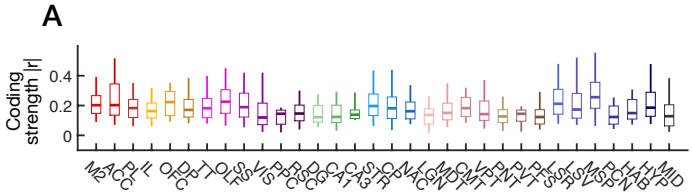
A
Coding strength |r|
0.4
0.2
0
M2
ACC
PIL
OFC
DP
TT
OLF
SS
VIS
PPC
RSC
DG
CA1
CA3
STR
CP
NAC
LGN
MD
CM
VPT
PNT
PVT
PFT
FS
FSR
MSV
PCN
HAB
HPF
MID

Figure S8. Additional plots and neural data for virtual foraging task.

A) Comparison of decision variable (DV) coding strength in DV-coding neurons across sampled brain regions, measuring (for each neuron) the absolute-value correlation |r| between its firing rate and the model-fit DV (median, IQR, range of session samples for each region).

Figure S9

Virtual reality (VR) patch foraging task
dFC inactivation (Muscimol → M2)

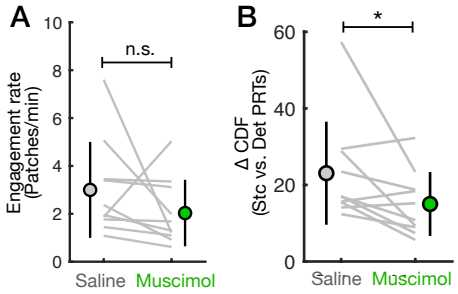

Free-moving (FM) patch foraging task
dFC inactivation (Muscimol → M2)

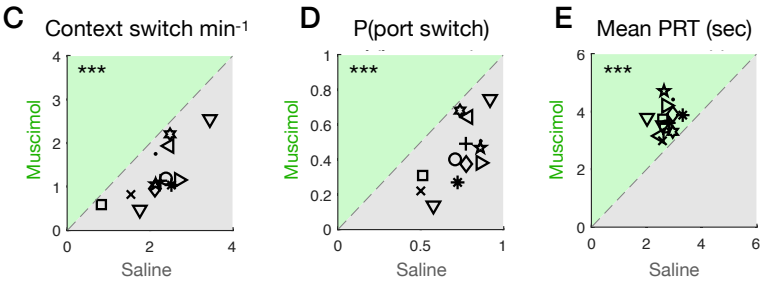

Figure S9. Additional data for dFC inactivation experiments.

A) Comparison of patch engagement (trials/min) for mice in the virtual (VR) foraging task under muscimol vs. saline (SR test, $p = 0.0645$, $n = 10$ mice).
B) Comparison of dCDF of Det vs. Stc PRTs for mice in the VR task under muscimol vs. saline (SR test, $p = 0.0273$, $n = 10$ mice, similar to KS statistic, **Figure S2B**).
C) Comparison of context transition rate (switches/min) for mice in the free-moving (FM) foraging task under muscimol vs. saline (SR test, $p = 2.44e-4$, $n = 10$ mice).
D) Comparison of port switch rate (average probability of port transition) for mice in the FM task under muscimol vs. saline (SR test, $p = 2.44e-4$, $n = 10$ mice).
E) Comparison of average patch residence times (PRT) for mice in the FM task under muscimol vs. saline (SR test, $p = 2.44e-4$, $n = 10$ mice).

For all plots, each dot or line is one mouse; Wilcoxon signed rank test (SR test); ***$p<0.001$; **$p<0.01$; *$p<0.05$; n.s. $p>0.05$.

Table S1. Brain region categories and Allen atlas correspondence

Region name	Tag	Macro-region category	Allen atlas labels
Secondary motor cortex	M2	Dorsal frontal cortex (dFC)	'Secondary motor area'
Anterior Cingulate cortex	ACC	Dorsal frontal cortex (dFC)	'Anterior cingulate area'
Prelimbic cortex	PL	Dorsal frontal cortex (dFC)	'Prelimbic'
Infralimbic cortex	IL	Ventral frontal cortex (vFC)	'Infralimbic'
Orbitofrontal cortex	OFC	Ventral frontal cortex (vFC)	'Orbital'
Dorsal peduncle	DP	Ventral frontal cortex (vFC)	'Dorsal peduncular'
Tenia tecta	TT	Sensory cortex (SC)	'Tenia tecta'
Olfactory areas	OLF	Sensory cortex (SC)	'Olfactory areas'
Somatosensory areas	SS	Sensory cortex (SC)	'somatosensory'
Visual areas	VIS	Sensory cortex (SC)	'visual areas'
Posterior parietal cortex	PPC	Sensory cortex (SC)	'anterior area'
Retrosplenial cortex	RSC	Sensory cortex (SC)	'Retrosplenial'
Dentate gyrus	DG	Hippocampal (HC)	'Dentate gyrus'
Hippocampus CA1	CA1	Hippocampal (HC)	'Field CA1'
Hippocampus CA3	CA3	Hippocampal (HC)	'Field CA3'
Striatum	STR	Basal ganglia (BG)	'Striatum'
Caudoputamen	CP	Basal ganglia (BG)	'Caudoputamen'
Nucleus accumbens	NAC	Basal ganglia (BG)	'accumbens'
Lateral geniculate nucleus	LGN	Thalamic (THL)	'lateral geniculate'
Mediodorsal thalamus	MDT	Thalamic (THL)	'Mediodorsal nucleus of the thalamus'
Centromedial thalamus	CMT	Thalamic (THL)	'Central medial nucleus of the thalamus' 'Rhomboid nucleus'
Ventral thalamus	VPT	Thalamic (THL)	'Ventral posteromedial nucleus of the thalamus' 'Ventral posterolateral nucleus of the thalamus'
Posterior thalamus	PNT	Thalamic (THL)	'posterior nucleus of the thalamus' 'Posterior complex of the thalamus' 'Posterior triangular thalamic nucleus' 'Suprageniculate nucelus' 'Ethmoid nucleus of the thalamus'
Paraventricular thalamus	PVT	Thalamic (THL)	'Paraventricular nucleus of the thalamus'
Parafascicular nucleus	PFS	Thalamic (THL)	'Subparafascicular area'
Lateral septum rostral	LSR	Subcortical (SUB)	'Lateral septal nucleus rostral'
Lateral septum ventral	LSV	Subcortical (SUB)	'Lateral septal nucleus ventral'
Medial septum	MSP	Subcortical (SUB)	'Medial septal nucleus'
Precommissural nucleus	PSN	Subcortical (SUB)	'Precommissural nucleus'
Hebnula	HAB	Subcortical (SUB)	'habenula'
Hypothalamus	HYP	Subcortical (SUB)	'Hypothalamus' 'Preoptic area'
Midbrain	MID	Subcortical (SUB)	'Midbrain'

Methods

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Mice

A total of 67 adult mice (C57/BL6j, 33 male, 34 female, age 2-6 months) were used in the experiments. Mice were housed on a 12h dark/12h light cycle (dark from 07:00 to 19:00) and performed behavior tasks in the dark period. All procedures were performed in accordance with the National Institute of Health Guide for the Care and Use of Laboratory Animals and approved by the Harvard Animal Care and Use Committee.

METHOD DETAILS

Surgical procedures

For virtual-reality (VR) Muscimol inactivation experiments, mice were anaesthetized using isoflurane (4% induction, 1-2% maintenance). For experiments requiring head-fixation, a custom titanium head-plate was attached to the skull with Metabond (Parkell). For both acute and chronic Neuropixels recordings, fiducial marks were made at the target sites for probe insertion using a fine-tipped pen, and a ground pin was inserted into the skull above contralateral cortex or cerebellum and attached with Metabond. For Muscimol inactivation experiments, a craniotomy was placed above the M2 region of the dorsal frontal cortex (Bregma +2.35 AP, +1.0 ML) and Muscimol or saline injected 1.0 ventral to the skull surface.

For chronic Neuropixels implants, mice were anaesthetized using Ketamine/medetomidine (Ketamine 10 mg/mL, Dexdomitor 0.04 mg/mL, 6 mL/kg). We used a custom adapted version of the Apollo implant described in Bimbard et al.,[50] which uses separable "docking" and "payload" modules to secure and cement the probe to the skull after insertion and recover the probe after completion of experiments. Neuropixels 1.0 single-shank probes were used for chronic recordings targeted in dorsal frontal cortex, and we tracked electrophysiology signals online using SpikeGLX to guide probe placement. After implanting and securing the probe, we injected Antisedan (Atipamezole hydrochloride, 5mg/mL) to reverse the effects of Ketamine and allowed animals to fully recover for 1 week before behavior experiments.

Foraging behavior experiments

For all experiments, mice were water restricted to maintain body weight > 85% of their baseline weight before training to ensure task motivation and maintain health. Additional water supplementation was provided following training sessions with a total of task and post-task water ranging from 1.0-2.4mL per day, determined individually per mouse. Mice that showed task motivation at higher weights were given greater amounts of total water supplementation accordingly to minimize stress from water restriction. Body weights ranged from 85-105% of baseline.

Male and female mice were used in approximately equal proportions for all behavior and recording experiments. We observed no significant differences in task engagement, decision behavior, or contextual flexibility in foraging decisions across sexes (**Figure S4**).

Virtual reality setup for VR foraging task

Virtual reality setups were identical to those used in Bukwich et al.[37] Three monitors (width 53cm, height 30cm) were placed in front and on either side of the animal. Virtual reality scenes were generated using VirMEn software in MATLAB on a workstation computer (DELL Precision 5810). Mice were head-fixed on a cylindrical Styrofoam treadmill (diameter 20.3cm, width 10.4cm). The rotational velocity of the treadmill was recorded using a rotary encoder. The output pulses of the encoder were converted into continuous velocity signals using custom Arduino code running on a microprocessor (Teensy 3.2). Velocity was integrated within VirMEn to compute position. Water was delivered to the mouse from a spout placed in front of the mouse's mouth, and licks were monitored using an infrared sensor (OPB819Z, TT electronics). Voltage signals from the rotary encoder were digitized and recorded on the virtual reality computer using a data acquisition system (PCIe-6323, National Instruments). Water timing and amount was controlled using a solenoid valve (LHDA 1221111H, The Lee Company) and a switch (2N7000, On Semiconductor), with TTL pulses generated by the virtual reality computer via the PCIe-6323 data acquisition card.

Virtual reality (VR) patch foraging task

Habituation: Mice were habituated to the treadmill prior to initial head-fixation with delivery of occasional water droplets from the reward spout while mice were free to explore on top of the treadmill, which was held in place to prevent rotation. Once mice showed consistent interest in consuming water droplets (usually 2 days, 15-minute exposure), head fixation began on the following session. Initial head-fixation habituation sessions were run for durations of 15 minutes, with gradual increases in session duration across consecutive sessions reaching a maximum duration of 50-60 minutes. During these sessions, water rewards (2μL) were delivered in 10 second intervals with additional rewards given during running bouts to reinforce running behavior. Once starting head-fixation, mice performed daily sessions with 1 day off per week (6 days on, 1 day off) for the remainder of the task. Once mice were consistently running for most of the session (usually 4-6 days), they proceeded on the next session to the first day of the foraging task.

Phase 1: Once mice demonstrated consistent running behavior during habituation (above), they began training on the linear virtual track. Mice navigated forward down the virtual corridor (track) and could stop at visually indicated patch locations to receive water rewards (they never received water between patches). The VR gain was set so that the travel distance between patches was ~40cm. On patch entry, mice received an odor puff indicating that they were on the patch, and if they stopped moving for 0.1s, could receive an initial water reward (2μL). Following this reward, they could stay on the patch to receive additional rewards following patch-specific statistics. For the first 7-10 days, mice were only exposed to the Deterministic ("Det") environment, which was indicated with a unique visual pattern on the corridor walls, unique visuals for patch locations, and a unique patch odor (isoamyl acetate or ethyl butyrate). Det patches always yield the patch-initial reward (at t=0s) followed by rewards at t=2s and t=4s for a total of three guaranteed rewards. Mice

could leave the patch at any time by moving forward on the track. The immobile time required to initiate the patch trial (starting at 0.1s) was gradually increased by 0.05s each day to a maximum of 0.5s. Occasionally, this was reduced for animals that showed faltering engagement rates and then increased again to 0.5s as engagement behavior recovered.

Phase 2: Mice then proceeded to Phase 2 of the task, where they foraged in the Det environment and the Stochastic (Stc) environment once per session in two blocks (20-25 minutes each) with one context transition in between. The Stc environment was indicated with distinct visuals for the corridor walls and patches, and a distinct odor (isoamyl acetate or ethyl butyrate) on patch entry. On Stc patches, after an initial reward at t = 0s, mice could receive additional rewards each second with probably P(Reward) that decayed over time (**Figure S3A**). Initial P(Reward) P_i at t = 1s for Stc patches was drawn randomly (uniformly) from P_i = [0.1, 0.26, 0.58, 0.95] and followed exponential decay with $\boldsymbol{\tau}$ = 0.189s (i.e. $\mathbf{P_t = P_i * e^{(-\tau * t)}}$), with P = 0 after 30 seconds. 10% of the patches in the Stc environment were programmed with the Det patch return ("Stc:024 patches" with rewards at t = 0, 2, 4s) but were otherwise indistinguishable from other Stc patches. Initial context identity alternated each day, and the context transition in the middle of each session was only trigged between patches and included a 10 second blank screen with no visual graphics before initiating the second environment. All rewards were 2µL water droplets.

Arena setup for free-moving foraging task

Mice navigated a hexagonal arena (inner radius 26.35 cm) with waterspouts placed at four locations in the center of the arena walls (see **Figure** 1C) and a wall with a controllable door separating the arena into two chambers (the "Stc" and "Det" environments/contexts, each with two waterspouts). Mice were detected at waterspouts using IR beam breaks (3mm LED IR beam sensors, Adafruit), and water timing and amount was controlled using solenoid valves (LHDA 1221111H, The Lee Company). Task events (water delivery and door) were controlled and recorded using custom Arduino code running on a microprocessor (Teensy 4.1). Transparent acrylic barriers were placed between the patches to increase and equalize the effective path length between patches (approximately 30 inches, 76.2 cm) and encourage greater time investment for patch trials. Mice could approach and nose-poke a reward port to initiate a trial, upon which they would immediately receive a 2µL water droplet (marking t = 0s in the patch trial). If they remained with their nose inserted, they could continue to receive water rewards each second with decreasing probability based on patch statistics (see below) and were free to leave the patch at any time by removing their nose. After leaving a patch, that reward port became "locked out" and would not initiate another trial for 6 seconds. All rewards were 2µL water droplets.

Free-moving (FM) patch foraging task

Habituation: Mice began water restriction one day prior to the task and were handled for 15 minutes prior to starting experiments. Mice were not habituated to the experiment chamber before exposure on day 1.

Phase 1: For the first 10 sessions (30 minutes), mice were exposed only to the right chamber (the Det environment) which included two reward ports yielding Deterministic ("Det") reward returns (the door separating the chambers remained closed). After the initial reward at t = 0s, mice

could receive additional rewards at t = 1 and t = 2s (for a total of three guaranteed rewards). Mice shaped their patch leaving decision strategies to this deterministic structure over the course of training in Phase 1 (**Figure S1**).

Phase 2: After Phase 1, mice proceeded to Phase 2, in which they were exposed to the Det environment (the right chamber) and the Stochastic ("Stc") environment (the left chamber) once per session in 15-minute blocks, alternating the starting environment each day. After 15 minutes in the session, the door separating the chambers opened, and once the mouse explored the other chamber and completed a patch trial, the door was automatically shut, imposing a block structure. On Stc patches, after an initial reward at t = 0s, mice could receive additional rewards each second with decaying probability. Initial P(Reward) P_i at t = 1s for Stc patches was drawn randomly (uniformly) from P_i = [0.66, 0.81, 0.93, 0.99] and followed exponential decay with $\boldsymbol{\tau}$ = 0.325s (i.e. $\mathbf{P_t = P_i * e^{(-\tau * t)}}$), with P = 0 after 20 seconds. 15% of the patches in the Stc environment were programmed with the Det patch return ("Stc:012 patches" with rewards at t = 0, 1, 2s) but were otherwise indistinguishable from other Stc patches. Mice experienced the Phase 2 task structure for 2 weeks, typically with 1 day off per week (6 days on, 1 day off).

Phase 3: After Phase 2, mice transitioned to Phase 3 of the task, in which the controllable door separating the environments/chambers was always left open. Mice were therefore free to engage context transitions and forage completely under their own volition. Although mice gained on average more rewards on Det compared to Stc patches (**Figure S1H**) and preferred to exploit patches in the Det context (**Figure 1I**), they still explored both contexts extensively in every session, typically making 2-3 context transitions each minute. Despite the lack of block structure or specific context cues, mice exhibited the same differences in decision behavior across contexts as observed in Phase 2 (**Figures S2C-F**).

Muscimol inactivation experiments

To test the functional role of dorsal frontal cortex (dFC) activity in context-specific decisions, we inactivated dFC neurons using local infusion of the $GABA_A$ receptor agonist Muscimol (Sigma-Aldrich M1523) targeting the secondary motor cortex (M2) subregion of dFC. For the VR task, mice were trained as described above for two weeks in Phase 2. Then, once per day about 2 hours before experiments, they were briefly (~15 minutes) anesthetized and injected with either Muscimol (1ng/nL, 150 nL) or saline (150 nL) over 5 minutes targeting the right side M2 using the coordinates (Bregma +2.35 AP, +1.0 ML, -1.0 DV from the skull surface). Mice were completely recovered from anesthesia before experiments and performed the task as usual. We injected Muscimol or saline for 2 sessions each on alternating days. Mice showed no deficits in task engagement (**Figure S9A**).

For inactivation experiments in the FM task, mice were trained as described above through Phase 1 and Phase 2, and for four days in Phase 3. We then injected either Muscimol or saline on alternating days as we did for the VR task and performed behavior experiments (FM task Phase 3) about 2 hours after injection. Mice exhibited several behavioral changes under Muscimol (**Figures 7J-O, S9C-E**) but did not show obvious deficits in motor behavior and still completed 10s-100s of trials per session.

Electrophysiological recording and spike sorting

Acute recordings in the VR task: Spiking data were collected using Neuropixels 1.0 single shank or Neuropixels 2.0 4-shank probes and acquired using SpikeGLX (https://billkarsh.github.io/SpikeGLX/). Craniotomies were performed 12-24 hours before the recording using a dental drill and covered with Kwik-cast (World precision Instruments), and mice recovered fully before experiments. On the day of recording, a 3D-printed shield was placed above the animal's head to block light and sight of the experimenter. Prior to insertion, probes were dyed with DiI, DiD, or DiO (Thermofisher) to label probe tracts from each recording session in tissue. Neuropixels probes were lowered using a Thorlabs micromanipulator (PT1-Z8) at 12-14 µm/sec. After reaching the target depth, the probe was allowed to settle in the tissue for 20 minutes prior to the start of the recording. Recordings typically lasted ~50-60 minutes.

Chronic recordings in the FM task: Spiking data were collected using Neuropixels 1.0 single shank probes targeting the dorsal frontal cortex (+2.25 AP, +0.8 ML) and recorded with Bonsai software (OpenEphys.Onix1 software library) using the ONIX PCIe acquisition system (open Ephys). Craniotomies and probe implantation were performed at least one week before behavior and recording sessions began. Artificial dura (Dura-Gel, Cambridge Neurotech) was applied during probe implantation to protect the site of the craniotomy. Probes were mounted with dental cement using a custom adaptation of the "Apollo" implant system to enable probe recovery and were dyed using DiI to mark the tissue trajectory of the probe.[50]

For all neural recording experiments, Neuropixels data were spike sorted offline using Kilosort 4 (https://github.com/MouseLand/Kilosort) followed by manual curation in Phy (https://github.com/cortex-lab/phy) to ensure quality spike waveforms and recording stability for isolated units and to split doublet clusters. Multi-unit and LFP signals were not used for any analysis. See **Figure S6** for full recording summaries.

Histology and probe tract analysis

After the last recording session, mice were transcardially perfused with phosphate buffered saline (PBS) followed by 4% paraformaldehyde in PBS. The brains were sliced in 100 µm coronal sections using a vibratome (Leica). Brain sections were mounted on glass slides and stained with 4', 6-diamidino-2-phenylindole (DAPI, Vectashield) in mounting medium. Slides were then imaged with a fluorescence microscope (Zeiss, Axio Scan.Z1). Probe tracts (labeled with fluorescence from DiI, DiD, and DiO) were identified in histology images (processed using Zen 3.3, Carl Zeiss Microscopy) and aligned using histological and electrophysiological landmarks to a reference atlas (Allen brain atlas CCF) using MATLAB software written by Dr. Andrew Peters (https://github.com/petersaj/AP_histology). Single units were accordingly sorted into 32 region and 7 macro-region categories using region identities from the reference atlas. Single units localized in white-matter structures within 50mm (vertical) of an identified region were assigned to it. Units localized in white-matter structures outside of 50mm from other regions, or those assigned to regions outside the 32 categories, were not used.

Mouse and trial inclusion criteria

For basic behavior and recording experiments in the VR head-fixed foraging task, 41 mice (22 males and 19 females) were trained and used for analysis, of which 10 were recorded from. 47 recording sessions were used for neural analysis. For analysis of decision behavior, trials from individual sessions were only included if at least 10 trials were performed in the session (70548/70827, 99.61% of all trials). Mice occasionally disengaged from the task for brief periods of time. Therefore, to prevent contamination of decision behavior by the disengaged state, we did not include trials with PRT (patch residence time) greater than 60 seconds for behavior analysis (2243/70827, 3.17% of all trials dropped).

For task-naïve recordings in the VR task, 6 mice (4 females, 2 males) were habituated to head fixation as described above and then proceeded directly to Phase 2 of the foraging task for 3 days, during which acute Neuropixels recordings were performed. 17 recording sessions were used for neural analysis.

For Muscimol inactivation experiments in the VR task, mice were required to engage at least 10 trials of each type (Det and Stc:024) during both Muscimol and saline sessions to be used for analysis. 10/14 mice trained on the task passed this inclusion criterion and were used for behavior analysis.

For basic behavior and recording experiments in the FM foraging task, 20 mice (9 males and 11 females) were trained and used for analysis, of which 4 were recorded from. 32 recording sessions (16 from Phase 2 and 16 from Phase 3) were used for neural analysis. All sessions were used for behavior analysis. We did not include trials with PRT greater than 6 seconds for behavior analysis.

For Muscimol inactivation experiments in the FM task, mice were required to engage at least 10 trials of each type (Det and Stc:012) during both Muscimol and saline sessions to be used for analysis. 13/13 mice trained on the task passed this inclusion criterion and were used for behavioral analysis.

Analysis of patch residence time (PRT) decision behavior

We analyzed decision behavior for patch foraging trials using the patch residence time (PRT) for different trial types. To evaluate whether mice made patch leaving decisions in a value-sensitive manner, we compared PRTs across four hidden trial types within Stc patches which had different initial reward probabilities (P_i). P_i was not cued to animals but was drawn randomly on trial initiation (see task descriptions above). In both the VR and FM tasks, average PRTs scaled positively with P_i (**Figures S2A, 2D**).

We compared distributions of PRTs across patch types as a measure of contextual flexibility in decision behavior, visualizing them using the cumulative density function (CDF) of leaving probabilities over time (as in **Figure 1E**). We quantified their divergence using the Kolmogorov-Smirnov statistic (KS statistic), which measures the difference in CDFs at their point of maximum vertical difference. We compared the KS statistic across Det and reward-matched Stc:012 or Stc:024 trial types, or across subsets of Det trials (**Figures 1I, 2G**), and found that PRT distributions across contexts were more distinct than within-context Det comparisons. However,

the KS statistic can be sensitive to large, local deviations in the shape of a distribution and can be skewed by small samples. Thus, for some analyses which required comparing distributions with very different sample sizes (e.g. comparing PRTs on Det vs. Stc trials for Muscimol experiments), we estimated the KS statistic by down-sampling trials to equalize sample sizes and taking the average resampled KS statistic across 1000 draws. We additionally used an alternative measure which is less sensitive to localized differences, the total absolute difference in CDFs (ΔCDF, see **Figure S2B**), to validate behavior results (**Figure S9B**). ΔCDF was highly correlated with the KS statistic in both tasks (**Figures S4E, S4J**), indicating that they are similar measures of PRT distribution difference.

Integrator models of patch leaving decisions

Following our previous report, we modeled patch leaving decisions in both the VR and FM foraging tasks as a drift-diffusion decision variable (DV) that integrates time on patch ($\boldsymbol{\delta}$) and reward receipt ($\boldsymbol{\rho}$) toward a boundary as evidence for or against the leaving decision.[34,37] In the model (schematized in **Figure 1J**), a decision variable $\mathbf{X_t}$ is initialized to 0 at t = 0 (for each patch trial), incremented by a "drift rate" $\boldsymbol{\delta}$ each time step, and is perturbed upon reward receipt by $\boldsymbol{\rho}$ (the "reward sensitivity"). At each time step, $\mathbf{X_t}$ is converted into a leaving probability ($\mathbf{0 <= P_{Leave} <= 1}$) using a sigmoid function centered at 0.5 with temperature $\boldsymbol{T}$ ($\mathbf{P_{Leave} = 1/(1+e^{(-T*(Xt - 0.5))})}$). Thus, the model contains three parameters ($\boldsymbol{\delta}$, $\boldsymbol{\rho}$, $\boldsymbol{T}$), and different versions of the decision computation ("Models 1-4") are produced by setting different constraints and definitions for $\boldsymbol{\delta}$ and $\boldsymbol{\rho}$. Specifically, the "reward insensitive" model 1 has ($\boldsymbol{\delta > 0, \rho = 0}$); the "reward reset" model 2 has ($\boldsymbol{\delta > 0, \rho = -X_t}$); the "reward integration" model 3 has ($\boldsymbol{\delta > 0, \rho < 0}$); and the "reward accumulation" model 4 has ($\boldsymbol{\delta > 0, \rho > 0}$). Integrator models were fit by using empirical reward returns and mouse PRT behavior to generate DV rollouts, and then estimating the parameters ($\boldsymbol{\delta}$, $\boldsymbol{\rho}$, $\boldsymbol{T}$) using the fmincon function in MATLAB (model fitting toolbox) to maximize the log-likelihood between $\mathbf{P_{Leave}}$ and the empirical patch leave times (PRTs). For parameter comparisons across patch types (**Figures 1O-P, 2K-L**), full models were used to estimate $\boldsymbol{\delta}$ and $\boldsymbol{\rho}$ using trials with 0.5s <= PRT < 30s (for VR behavior) and 0.5s <= PRT < 10s (for FM behavior). For comparison and quantification of model performance, we fit models on 90% of the trials, computed the Bayesian Information Criterion (BIC) between predicted $\mathbf{P_{Leave}}$ and ground truth leave times for held-out trials, and compared average (across folds) cross-validated BIC (cv-BIC) for Models 1-4 (**Figures 1K-L, 2I-J**).

We reported in a previous study (Bukwich et al.)[37] that, in a similar VR patch foraging task, patch leaving decisions were best described using an integrator model that is scaled by a latent "patience" state $\boldsymbol{\phi}$, which captures individual variability in waiting behavior as well as within-session slow fluctuations in PRTs that are not explained by patch rewards. For behavior in the VR task, we used a similar scaling parameter in the models to improve fits to behavior while preserving interpretability of the evidence weight parameters $\boldsymbol{\delta}$ and $\boldsymbol{\rho}$. For each patch trial, we computed the current value of a "patience" variable $\boldsymbol{\phi}$ by taking the average of gaussian smoothed PRTs of the two adjacent trials and then scaled these to the average $\boldsymbol{\phi}$ across all trials used for model fitting ($\boldsymbol{\phi = \phi/\mathrm{mean}(\phi)}$). Then, to compute the DV $\mathbf{X_t}$ in the model fitting procedure, we used $\boldsymbol{\delta/\phi}$ as the "drift rate." As smaller drift rates yield longer PRTs, this captures the effect of the patience state on leaving decisions by scaling the drift rate to bias trial PRTs in the direction of those for neighboring trials. As we did not observe substantial variability in FM task behavior, we did not

include this "patience" scaling by ϕ for fits of the integrator model for integrator model of decisions in the FM task.

Generating predicted PRTs from model outputs

To simulate PRTs from specific DV rollouts fit for example mice (**Figure 1Q**), we took the deterministic DVs $\mathbf{X_t}$ generated from model fits (**Figure 1Q**, left), added gaussian noise (**σ=0.1**), computed $\mathbf{P_{Leave}}$ as described above, took the first timepoint where $P_{Leave} > 0.5$ as one simulated PRT, and repeated this procedure 100 times. We used this distribution of simulated PRTs to compute $\mathbf{P_{Leave}}$ CDFs for different parameterizations of the DV (fit to behavior from Stc and Det patch trials) and plotted these.

Normalization of single-neuron firing rate traces

Firing rate traces for each unit were generated by binning spikes into 50 ms time bins and then z-scoring across the entire time series. Throughout (e.g. in **Figures 3E and 4D-E**), σ indicates 1 standard deviation of firing rate activity. For some analyses, local single-neuron activity was further normalized by subtracting activity within a pre-event (e.g. such as pre-patch trial) window and reported as "Δ activity".

Smoothing and de-noising for principal component projections

Where principal component (PC) analysis was used to visualize low-dimensional projections of population activity, single-unit traces were further smoothed or averaged to reduce noise contributions to PC identification[42]. For PC projections of activity across long timescales capturing the context transition (**Figures 3H and 6G**), single-neuron traces were smoothed using a 10-second left moving-mean filter (i.e. only using previous, and not future, timepoints). For PC projections of single-trial activity, single-neuron activity was averaged within trial types to reduce noise from trial variability and then smoothed with a 1-second left moving-mean filter.

Analysis of single-neuron and population coding of sensory context

We identified individual neurons as encoding context identity by measuring consistent changes in firing rate corresponding to environment identity ("Stc" or "Det") that occur around the time of context transition. For each neuron, we sampled activity in 60-second bins across the entire session and performed a 2-sample t-test at $\alpha = 0.05$ comparing the mean activity across samples from context 1 and context 2. Then, we measured the change in mean activity "Δ activity" across 5-minute windows before and after the context transition and obtained a p-value (testing at $\alpha = 0.01$) by comparison to a null distribution generated from sampling Δ activity around random within-context timepoints. Neurons that independently passed both of these statistical tests and exhibited the same direction of activity change across both (i.e. consistent global and transient changes in firing rate) were considered "context coding." For regional comparisons of single-neuron context coding (**Figure 3J**), only session/region samples with at least 5 neurons were used for analysis.

For quantification of population-level context discriminability (**Figure 4I**), we split the single-neuron traces into rough quarter bins (for each session) corresponding to the first and second half of each context block, and then took the average activity of each neuron in each of these bins, forming population vectors for context 1 (half 1 and half 2) and for context 2 (half 1 and half 2). We then computed three angles comparing these population vectors by first normalizing their lengths and then taking the inverse cosine of their dot product ($\boldsymbol{\theta} = \mathbf{cos^{-1}(V_1 \cdot V_2)}$). For within-context comparisons, we computed the angle between the two halves within each context, and for across context comparison, we took the angle between the second half of context 1 and the first half of context 2 (**Figures 3I and 6J**).

Analysis of subspace divergence

To analyze how context reconfigures population dynamics, we quantified the divergence of activity subspaces that separately capture population activity for Det and Stc contexts. To do this, we used canonical correlation analysis (CCA), which finds directions of maximum covariance between two sets of data and can be used to measure their subspace overlap (see **Figure S7** for graphical description and examples).[43] For each session or session/region sample of population activity, we first separately applied PCA to epochs corresponding to the Stc and Det context to define directions in the state space (the principal components) which capture within-context dynamics, and retained the top $\mathbf{N_{PC}}$ components for each, where $\mathbf{N_{PC}}$ is the number of PCs computed for the full dataset that capture 80% of total variance. These sets of PCs ($\mathbf{S_{Stc}}$ and $\mathbf{S_{Det}}$) are orthonormal bases that define subspaces of the full state space.[22,42] To measure their alignment, we first computed their cross-covariance matrix ($\mathbf{C = S_{Stc}^T S_{Det}}$) and then performed a singular value decomposition of C ($\mathbf{U \sum V^T = C}$). This yields two sets of singular vectors (or canonical directions) **U** and **V** which are successively maximally correlated with each other under the constraint that **U** and **V** are themselves orthogonal[43,44]. The diagonal of $\sum$ contains the corresponding singular values (or canonical correlations) for each vector pair, which are related to their alignment and the angle between them. These angles ($\boldsymbol{\theta}$) can be used to quantify subspace alignment; completely overlapping subspaces yield all angles $\boldsymbol{\theta} = \mathbf{0°}$, while completely orthogonal subspaces yield all $\boldsymbol{\theta} = \mathbf{90°}$ (see **Figures S7B-E** for illustrations). We measured overall subspace divergence as $\sqrt{\sum \mathbf{sin^2(\boldsymbol{\theta})}}$, the "chordal distance" between the subspaces, as well as the number of singular vector pairs with $\boldsymbol{\theta} > \mathbf{80°}$ ("orthogonal vectors"), which can be interpreted as the dimensional size of the components of the subspaces that are orthogonal to one another.

To evaluate whether subspace divergence was greater than expected from within-context neural variability (**Figures 3L-M**), we performed this analysis across epochs corresponding to the second half of context 1 and the first half of context 2 (across-context) and compared these by subtraction) to the first and second halves of context 2 (within-context) as direct within-session controls. We measured overall subspace difference metrics for each brain region as the median of all region/session samples with at least 5 neurons.

Because we observed a strong correlation between subspace difference and context coding across regions (**Figures 3M, S7I**), we analyzed how much these alignment metrics were obligately correlated with linear firing rate changes along one dimension (e.g. along a "context coding" vector), as this could yield a spurious association. To test this, we repeated the subspace alignment analysis using data from each session taken from two halves of one context and synthetically added

activity with gaussian noise (μ = 0-2 s.d., σ = 0-0.1 s.d.) to different sized subpopulations of neurons to replicate the presence of linear "context coding." Although adding synthetic context coding to real neural activity mildly increased chordal distance measures (**Figure S7F**), the relationship (measured using linear regression) was far weaker than that observed in the real data ($\beta = 0.436$ synthetic, $\beta = 1.185$ real). Additionally, adding synthetic context coding did not increase the number of orthogonal vectors between subspaces (**Figure S7G**), while real context coding was significantly correlated with orthogonal vectors (**Figure S7I**).

Analysis of single-neuron decision variable (DV) coding

To identify single-neuron correlates of the foraging decision variable (DV), we used a similar approach as in Bukwich et al.[37] We first fit the cognitive DV to decision behavior for each mouse using the integrator model as described, extracted the DV for trials corresponding to each trial type (e.g. Det trials or Stc trials), and concatenated these separately. For each neuron, we then extracted the (0.5s mean smoothed) firing rate corresponding to the same patch trials, concatenated these, and computed the correlation (***r***) between the unit activity and the DV (separately for Stc and Det trials), which we took as a measure of "coding strength." To test statistical significance, we repeated this process 1000 times, shuffling the trial order before concatenating, and computed a null distribution of ***r*** to test the observed value against. Neurons that exceeded either 0.5% tails ($\alpha = 0.01$) of this distribution were considered "DV coding" for either the Stc or Det DV.

Analysis of reward responses and ramping activity in DV coding neurons

To assess whether DV coding neurons represented differently parameterized decision variables as predicted from fits of the integrator models, we measured local reward responses and ramping activity in neurons encoding the Stochastic DV (Stc-DV) or Deterministic DV (Det-DV). We included all neurons (pooled across all sampled regions) which passed our bootstrap statistical test at $\alpha = 0.01$ for only one (Det or Stc) DV. We measured average activity in these Stc-DV and Det-DV neurons around three trial events: the reward marking trial initiation (t = 0s), additional patch rewards (at t >= 1s), and patch leave (the time the animal moves off the patch location). For the trial initiation response, we included activity from all trials with PRT >= 1s, measured the reward response as the local change in mean activity across surrounding 0.5s windows, and measured ramping activity as the slope (β) for regression of activity against time in the interval [0.25 to 1s] after trial initiation (reported as change in activity/sec, or $\sigma \cdot s^{-1}$. Patch reward responses were analyzed the same way, except centering the window to rewards that occur after t = 0s. For ramping activity before patch leave, we included trials with PRT > =7s and no rewards at t > 4s (so that neural activity was not contaminated by recent reward responses). We measured ramping activity as the slope β for regression of activity against time in the interval [-2 to 0.5s], reported in activity/sec, or $\sigma \cdot s^{-1}$.

Analysis of population DV coding and DV coding vectors

To define DV coding vectors for measurements of population coding geometry, we concatenated Stc or Det DVs as described above for the single-neuron analysis, extracted corresponding single-

neuron activity for all neurons in the population (within sessions), and performed linear regression of the neuron traces (**X**) against the DV (i.e. $\mathbf{DV} = \mathbf{X}^{\mathbf{T}}\boldsymbol{\beta}$). The normalized coefficients $\boldsymbol{\beta}/\mathbf{norm}(\boldsymbol{\beta})$ can be interpreted as a direction in the population activity space that best captures how neural activity covaries with the DV and thus defines a coding vector (or coding dimension) $\boldsymbol{\beta}_{\mathbf{Stc}}$ (Stc-DV) or $\boldsymbol{\beta}_{\mathbf{Det}}$ (Det-DV).[11,42] The angle between these coding vectors can be measured as the inverse cosine of their dot product ($\boldsymbol{\theta} = \mathbf{cos}^{-1}(\boldsymbol{\beta}_{\mathbf{Det}} \cdot \boldsymbol{\beta}_{\mathbf{Stc}})$).

Within-context comparisons of the angle between DV coding vectors were computed in the same way but instead using trial data corresponding to two subsampled halves of trial data within each context rather than across contexts.

Population decoder analysis of decision variables

For linear decoding of DVs, we fit DVs to mouse behavior and concatenated trial DVs and corresponding neural activity (for region/session populations) as described above for the coding vector analysis. However, instead of concatenating all trials, we performed leave-one-out cross-validation by iterating through trials, leaving each trial out once, computing the DV coding vector using the remaining trial data, projecting the neural data (**X**) for the held-out trial onto the coding vector using $\mathbf{DV}_{\mathbf{pred}} = \mathbf{X}^{\mathbf{T}}\boldsymbol{\beta}$, and computing the correlation $\boldsymbol{r}$ between the ground truth DV and $\mathbf{DV}_{\mathbf{pred}}$. For each session/region sample, we took the average $\boldsymbol{r}$ across all held-out trial predictions as a summary measure of decoder performance (***cv-r***). For cross-context decoder analyses, we projected trial data from one context onto the coding vector for the other context and computed ***cv-r***. For shuffle comparisons, we first randomly circularly permuted each neuron activity trace and then repeated the analysis as described. Trial data (DVs and corresponding activity) were all time-warped to common length so that, once concatenated, they would contribute equally to the regression regardless of variable trial lengths. For DV decoders for the VR task, trials with 0.5s < PRT < 30s were used, and for the FM task, trials with 0.5s < PRT < 10s were used.

Population decoder analysis of context identity

For linear binary decoding of context identity, we took the average activity for each patch trial (trial decoder) or each inter-trial-interval (ITI) period in the session as population vector samples, then randomly up-sampled the group (i.e. Stc or Det trials/ITIs) with fewer samples (to equalize sample sizes and remove bias from the decoder). Iterating 10 times, we then held out 10% of the samples (non-overlapping across iterations), trained a Fisher's Linear Discriminant (FLD) model on the remaining 90% to discriminate Stc vs. Det identity (fitcdiscr in MATLAB), and tested predicted labels against ground truth, using the area under the receiver operating characteristic (auROC) curve as a performance metric. The average across auROCs for the 10 folds (*cv-auROC*) was used as an overall summary measure. For shuffle comparisons, class labels were randomly shuffled before model fitting.

QUANTIFICATION AND STATISTICAL ANALYSIS

Unless otherwise noted, data were analyzed with custom code written in MATLAB (version 2019b, Mathworks). All box and whisker plots represent the median, interquartile range (box), and full range (whiskers) of the underlying distribution. For all statistical tests throughout, normality of data and equal variance of groups was not assumed, and non-parametric (Wilcoxon rank-sum and signed-rank) tests were used for unpaired and paired group comparisons, respectively. Test for normality was used before applying parametric tests (t-tests, Anova), and for all significance testing we used two-tailed tests at $\alpha = 0.05$ or $\alpha = 0.01$ (noted in figure legends). For statistical comparisons of PRT distributions, two-sample Kolmogorov-Smirnov tests were used. Resampling methods based on temporally-permuted neuron activity traces or shuffled time points were used to assess significance of single-neuron coding (described in detail above). Statistical significance of population decoder analyses was assessed by comparison to null distributions of performance statistics generated from shuffled trial sequences or class labels. The sizes of mouse groups were not pre-specified and approximated those of previous studies.

ACKNOWLEDGEMENTS

We thank Mike Bukwich and Malcolm Campbell for help setting up the VR foraging task and early discussions on design of experiments; Edward Soucy and Yuwei Li of the Harvard Neurotechnology Core for engineering assistance, the Harvard Center for Biological Imaging for support with microscopy; Mitsuko Watabe-Uchida and Uchida lab members for critical input and feedback on the manuscript. This work was supported by grants from the National Institutes of Health (NIH) (5U19NS113201 to N.U. and 5R01DA05975 to N.U. and M.W.-U.) and the Simons Collaboration on Global Brain (to N.U.). L.K. is supported by a Helen Hay Whitney Fellowship.

AUHOR CONTRIBUTIONS

L.K. designed the foraging tasks, basing the virtual-reality (VR) task on a previous design reported in Bukwich et al.[37] L.K. adapted and built behavioral and recording setups. L.K. performed all behavior and recording experiments for the VR task, L.K. performed all behavior experiments for the free-moving (FM) task, and L.K. and N.S.S. set up and performed chronic Neuropixels recordings for the FM task. L.K. performed Muscimol inactivation experiments with assistance from G.Z. L.K. analyzed all behavioral and neural data with feedback from N.U. L.K. wrote the manuscript with feedback from N.U.

DECLARATION OF INTERESTS

The authors declare no competing interests.

DATA AND CODE AVAILABILITY

All code and data will be uploaded to a public repository upon publication

References

1. Miller, E. K. & Cohen, J. D. An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* **24**, 167–202 (2001).

2. Diamond, A. Executive functions. *Annu. Rev. Psychol.* **64**, 135–168 (2013).

3. Passingham, R. *Understanding the Prefrontal Cortex*. (Oxford University Press (OUP), 2021).

4. Stokes, M. G. *et al.* Dynamic Coding for Cognitive Control in Prefrontal Cortex. *Neuron* **78**, 364 (2013).

5. Cole, M. W. *et al.* Multi-task connectivity reveals flexible hubs for adaptive task control. *Nat. Neurosci.* **16**, 1348 (2013).

6. Rigotti, M. *et al.* The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).

7. Liu, Y. & Wang, X. J. Flexible gating between subspaces in a neural network model of internally guided task switching. *Nat. Commun.* **15**, 1–20 (2024).

8. Okazawa, G. & Kiani, R. Neural Mechanisms That Make Perceptual Decisions Flexible. *Annu. Rev. Physiol.* **85**, 191–215 (2023).

9. Nakajima, M. Neuronal identity and cognitive control dynamics in the PFC. *Semin. Cell Dev. Biol.* **129**, 14–21 (2022).

10. Rigotti, M., Rubin, D. B. D., Wang, X. J. & Fusi, S. Internal representation of task rules by recurrent dynamics: the importance of the diversity of neural responses. *Front. Comput. Neurosci.* **4**, (2010).

11. Mante, V., Sussillo, D., Shenoy, K. V & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).

12. Wu, Z. *et al.* Context-Dependent Decision Making in a Premotor Circuit. *Neuron* **106**, 316-328.e6 (2020).

13. Rikhye, R. V., Gilra, A. & Halassa, M. M. Thalamic regulation of switching between cortical representations enables cognitive flexibility. *Nat. Neurosci.* **21**, 1753–1763 (2018).

14. Schmitt, L. I. *et al.* Thalamic amplification of cortical connectivity sustains attentional control. *Nature* **545**, 219–223 (2017).

15. Kumano, H., Suda, Y. & Uka, T. Context-Dependent Accumulation of Sensory Evidence in the Parietal Cortex Underlies Flexible Task Switching. *J. Neurosci.* **36**, 12192–12202 (2016).

16. Wimmer, R. D. *et al.* Thalamic control of sensory selection in divided attention. *Nature* **526**, 705–709 (2015).

17. Nakajima, M., Schmitt, L. I. & Halassa, M. M. Prefrontal Cortex Regulates Sensory

Filtering through a Basal Ganglia-to-Thalamus Pathway. *Neuron* **103**, 445-458.e10 (2019).

18. Lam, N. H. *et al.* Prefrontal transthalamic uncertainty processing drives flexible switching. *Nature* **637**, 127–136 (2025).

19. Pagan, M. *et al.* Individual variability of neural computations underlying flexible decisions. *Nature* **639**, 421–429 (2025).

20. MacDowell, C. J. *et al.* Multiplexed subspaces route neural activity across brain-wide networks. *Nat. Commun.* **16**, 1–21 (2025).

21. Ebitz, R. B. & Hayden, B. Y. The population doctrine in cognitive neuroscience. *Neuron* **109**, 3055–3068 (2021).

22. Langdon, C., Genkin, M. & Engel, T. A. A unifying perspective on neural manifolds and circuits for cognition. *Nat. Rev. Neurosci. 2023* 1–15 (2023). doi:10.1038/s41583-023-00693-x

23. Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T. & Wang, X. J. Task representations in neural networks trained to perform many cognitive tasks. *Nat. Neurosci.* **22**, 297 (2019).

24. Flesch, T., Juechems, K., Dumbalska, T., Saxe, A. & Summerfield, C. Orthogonal representations for robust context-dependent task performance in brains and neural networks. *Neuron* **110**, 1258 (2022).

25. Okazawa, G., Hatch, C. E., Mancoo, A., Machens, C. K. & Kiani, R. Representational geometry of perceptual decisions in the monkey parietal cortex. *Cell* **184**, 3748-3761.e18 (2021).

26. Ulanovsky, N. Natural neuroscience : toward a systems neuroscience of natural behaviors. (2025).

27. Miller, C. T. *et al.* Natural behavior is the language of the brain. *Curr. Biol.* **32**, R482–R493 (2022).

28. Mobbs, D., Trimmer, P. C., Blumstein, D. T. & Dayan, P. Foraging for foundations in decision neuroscience: insights from ethology. *Nat. Rev. Neurosci.* **19**, 419 (2018).

29. Pearson, J. M., Watson, K. K. & Platt, M. L. Decision making: the neuroethological turn. *Neuron* **82**, 950–965 (2014).

30. Stephens, D. & Krebs, J. *Foraging Theory*. (Princeton University Press, 1987).

31. Charnov, E. L. Optimal foraging, the marginal value theorem. *Theor. Popul. Biol.* **9**, 129–136 (1976).

32. Hall-McMaster, S. & Luyckx, F. Revisiting foraging approaches in neuroscience. *Cogn. Affect. Behav. Neurosci.* **19**, 225 (2019).

33. Hayden, B. Y., Pearson, J. M. & Platt, M. L. Neuronal basis of sequential foraging decisions in a patchy environment. *Nat. Neurosci. 2011 147* **14**, 933–939 (2011).

34. Davidson, J. D. & El Hady, A. Foraging as an evidence accumulation process. *PLoS Comput. Biol.* **15**, (2019).

35. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).

36. Shadlen, M. N. & Kiani, R. Decision making as a window on cognition. *Neuron* **80**, 791–806 (2013).

37. Bukwich, M. *et al.* Competitive integration of time and reward explains value-sensitive foraging decisions and frontal cortex ramping dynamics. *Neuron* **113**, 3458-3475.e12 (2025).

38. Cazettes, F. *et al.* A reservoir of foraging decision variables in the mouse brain. *Nat. Neurosci. 2023* 1–10 (2023). doi:10.1038/s41593-023-01305-8

39. Fetsch, C. R. The importance of task design and behavioral control for understanding the neural basis of cognitive functions. *Curr. Opin. Neurobiol.* **37**, 16 (2016).

40. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).

41. Maggi, S. & Humphries, M. D. Activity Subspaces in Medial Prefrontal Cortex Distinguish States of the World. *J. Neurosci.* **42**, 4131–4146 (2022).

42. Cunningham, J. P. & Yu, B. M. *Dimensionality reduction for large-scale neural recordings*. *Nature Neuroscience* **17**, 1500–1509 (Nature Publishing Group, 2014).

43. Wang, H. T. *et al.* Finding the needle in a high-dimensional haystack: Canonical correlation analysis for neuroscientists. *Neuroimage* **216**, 116745 (2020).

44. Zhang, X. *et al.* Inter-brain neural dynamics in biological and artificial intelligence systems. *Nature* **645**, 991–1001 (2025).

45. Khilkevich, A. *et al.* Brain-wide dynamics linking sensation to action during decision-making. *Nature* **634**, 890 (2024).

46. Zhang, Y. *et al.* A brain-wide map of neural activity during complex behaviour. *Nature* **645**, 177–191 (2025).

47. Murakami, M., Vicente, M. I., Costa, G. M. & Mainen, Z. F. Neural antecedents of self-initiated actions in secondary motor cortex. *Nat. Neurosci.* **17**, 1574–1582 (2014).

48. Murakami, M., Shteingart, H., Loewenstein, Y. & Mainen, Z. F. Distinct Sources of Deterministic and Stochastic Components of Action Timing Decisions in Rodent Frontal Cortex. *Neuron* **94**, 908-919.e7 (2017).

49. Adam, E. M., Johns, T. & Sur, M. Dynamic control of visually guided locomotion through corticosubthalamic projections. *Cell Rep.* **40**, (2022).

50. Bimbard, C. *et al.* An adaptable, reusable, and light implant for chronic Neuropixels probes. *Elife* **13**, (2025).

51. Karlsson, M. P., Tervo, D. G. R. & Karpova, A. Y. Network resets in medial prefrontal

cortex mark the onset of behavioral uncertainty. *Science* **338**, 135–139 (2012).

52. Akam, T. *et al.* The Anterior Cingulate Cortex Predicts Future States to Mediate Model-Based Action Selection. *Neuron* **109**, 149-163.e7 (2021).

53. Durstewitz, D., Vittoz, N. M., Floresco, S. B. & Seamans, J. K. Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron* **66**, 438–448 (2010).

54. Lak, A. *et al.* Orbitofrontal Cortex Is Required for Optimal Waiting Based on Decision Confidence. *Neuron* **84**, 190–201 (2014).

55. Yang, W., Tipparaju, S. L., Chen, G. & Li, N. Thalamus-driven functional populations in frontal cortex support decision-making. *Nat. Neurosci. 2022 2510* **25**, 1339–1352 (2022).

56. Finkelstein, A. *et al.* Attractor dynamics gate cortical information flow during decision-making. *Nat. Neurosci.* **24**, 843–850 (2021).

57. Li, N., Chen, T. W., Guo, Z. V., Gerfen, C. R. & Svoboda, K. A motor cortex circuit for motor planning and movement. *Nat. 2015 5197541* **519**, 51–56 (2015).

58. Inagaki, H. K., Fontolan, L., Romani, S. & Svoboda, K. Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* **566**, 212–217 (2019).

59. Rudebeck, P. H. & Izquierdo, A. Foraging with the frontal cortex: A cross-species evaluation of reward-guided behavior. *Neuropsychopharmacol. 2021* 1–13 (2021). doi:10.1038/s41386-021-01140-0

60. Kilpatrick, Z. P., Davidson, J. D. & El Hady, A. Uncertainty drives deviations in normative foraging decision strategies. *J. R. Soc. Interface* **18**, (2021).

61. Taniuchi, T., Sugihara, J., Wakashima, M. & Kamijo, M. Abstract numerical discrimination learning in rats. *Learn. Behav.* **44**, 122–136 (2016).

62. Toptaş, P., Gür, E. & Balcı, F. Count-based decision-making in mice: numerosity vs. stimulus control. *Anim. Cogn.* **25**, 1621–1630 (2022).

63. Hanks, T. D., Kiani, R. & Shadlen, M. N. A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *Elife* **3**, e02260 (2014).

64. Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E. & Shadlen, M. N. Elapsed Decision Time Affects the Weighting of Prior Probability in a Perceptual Decision Task. *J. Neurosci.* **31**, 6339 (2011).

65. Mochol, G., Kiani, R. & Moreno-Bote, R. Prefrontal cortex represents heuristics that shape choice bias and its integration into future behavior. *Curr. Biol.* **31**, 1234 (2021).

66. Lo, C. C. & Wang, X. J. Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. Neurosci.* **9**, 956–963 (2006).

67. Scott, D. N., Mukherjee, A., Nassar, M. R. & Halassa, M. M. Thalamocortical architectures for flexible cognition and efficient learning. *Trends Cogn. Sci.* **28**, 739–756 (2024).

68. Luo, M., Zhou, J. & Liu, Z. Reward processing by the dorsal raphe nucleus: 5-HT and beyond. *Learn. Mem.* **22**, 452–60 (2015).

69. Miyazaki, K., Miyazaki, K. W. & Doya, K. The role of serotonin in the regulation of patience and impulsivity. *Mol. Neurobiol.* **45**, 213–224 (2012).

70. Miyazaki, K., Miyazaki, K. W. & Doya, K. Activation of dorsal raphe serotonin neurons underlies waiting for delayed rewards. *J. Neurosci.* **31**, 469–479 (2011).

71. Miller, E. K. The prefontral cortex and cognitive control. *Nat. Rev. Neurosci.* **1**, 59–65 (2000).

72. Weber, J. *et al.* Subspace partitioning in the human prefrontal cortex resolves cognitive interference. *Proc. Natl. Acad. Sci. U. S. A.* **120**, (2023).

73. Avitan, L. & Stringer, C. Not so spontaneous: Multi-dimensional representations of behaviors and context in sensory areas. *Neuron* **110**, 3064–3075 (2022).

74. Churchland, M. M. & Shenoy, K. V. Preparatory activity and the expansive null-space. *Nat. Rev. Neurosci.* **25**, 213–236 (2024).

75. Hyman, J. M., Ma, L., Balaguer-Ballester, E., Durstewitz, D. & Seamans, J. K. Contextual encoding by ensembles of medial prefrontal cortex neurons. *Proc. Natl. Acad. Sci.* **109**, 5086–5091 (2012).

76. Siegel, M., Buschman, T. J. & Miller, E. K. Cortical Information Flow During Flexible Sensorimotor Decisions. *Science* **348**, 1352 (2015).

77. Inagaki, H. K. *et al.* A midbrain-thalamus-cortex circuit reorganizes cortical dynamics to initiate movement. *Cell* **185**, 1065-1081.e23 (2022).

78. Siniscalchi, M. J., Phoumthipphavong, V., Ali, F., Lozano, M. & Kwan, A. C. Fast and slow transitions in frontal ensemble activity during flexible sensorimotor behavior. *Nat. Neurosci.* **19**, 1234–1242 (2016).

79. Barthas, F. & Kwan, A. C. Secondary Motor Cortex: Where 'Sensory' Meets 'Motor' in the Rodent Frontal Cortex. *Trends Neurosci.* **40**, 181–193 (2017).

80. Kim, J. H., Daie, K. & Li, N. A combinatorial neural code for long-term motor memory. *Nature* **637**, 663–672 (2025).