Flexible control of representational dynamics in a disinhibition-based model of decision making

Bo Shen[1], Kenway Louie[2,1], Paul Glimcher[1,2]

[1]Neuroscience Institute, New York University Grossman School of Medicine, New York, NY, United States, 10016

[2]Center for Neural Science, New York University, New York, NY, United States, 10003

Manuscript Information:

150 words in abstract, 5482 words in main text, 8 figures, and supplementary information

Keywords:

Decision-making circuit, disinhibition, divisive normalization, winner-take-all choice, persistent activity

Correspondence:

Bo Shen
Neuroscience Institute
Grossman School of Medicine
New York University
435 E 30th, Room 1260
New York, NY 10016
Tel: +1.212.263.8167
Fax: +1.212.263.8170
Email: bs3667@nyu.edu

Abstract

Current models utilize two primary circuit motifs to replicate neurobiological decision making. Recurrent gain control implements normalization-driven relative value coding, while recurrent excitation and competitive inhibition together implement winner-take-all (WTA) dynamics. Despite evidence for concurrent valuation and selection computations in individual brain regions, existing models focus selectively on either normalization or WTA dynamics and how both arise in a single circuit architecture is unknown. Here we show that a novel hybrid motif unifies both normalized representation and WTA competition, with dynamic control of circuit state governed by local disinhibition. In addition to capturing empirical psychometric and chronometric data, the model produces persistent activity consistent with working memory. Furthermore, disinhibition provides a simple mechanism for flexible top-down control of network states, enabling the circuit to capture diverse task-dependent neural dynamics. These results suggest a new biologically plausible mechanism for decision making and emphasize the importance of local disinhibition in neural processing.

Introduction

Two fundamental processes in decision making are value coding and option selection. In the early stage of a decision, value representations serve as integrated decision variables that combine outcome information such as expected gain and probability of realization. Such representations are central to formal decision theories in ecology, economics, and psychology[1]. Consistent with behavioral theories, neural firing rates vary with option values in numerous decision-related brain areas, including the frontal[2–10] and parietal[10–21] cortices and basal ganglia[22–25]. Recent research shows that neural value coding is contextual in nature, with the value of a given option represented relative to the value of available alternatives[2,9,15,19,26–30]. Furthermore, this relative value coding employs divisive normalization[9,27,31,32], a canonical computation prevalent in sensory processing and thought to implement efficient coding principles[33–39] and temporal adaptation[28,32,36,40–42].

As the decision process progresses, a common and powerful neural mechanism for categorical choice is winner-take-all (WTA) competition[43,44]. WTA dynamics are widely observed in multiple brain regions: the neural firing rate representing the chosen option or action target increases in concert with selection (often reaching a common activity threshold at choice), while firing rates representing the other unchosen alternatives are suppressed, resulting in a single categorical choice[14–16,19,20,45–50]. The wide prevalence of WTA dynamics in decision-related neural activities suggests that it is a general feature of biological choice.

Computational decision models have identified core circuit motifs that produce either normalized value representation or WTA selection (**Fig. 1**). For normalized value representation, dynamic circuit-based models emphasize a crucial role for both lateral and feedback inhibition[29,51]. In the dynamic normalization model (DNM), paired excitatory and inhibitory neurons represent each choice option (**Fig. 1A**); feedforward excitation inputs value information, lateral connectivity mediates contextual interactions, and feedback inhibition drives divisive scaling. This simple differential equation model emphasizes a crucial role of lateral connectivity and feedback inhibition in driving empirically-observed divisive scaling and contextual interactions (**Fig. 1B**).

For WTA selection, the predominant class of decision models (recurrent network models, hereafter RNM) propose a central role for recurrent connectivity[52–58] and non-selective feedback inhibition[43,44] (**Fig. 1C**). RNMs replicate psychophysical and neurophysiological results in perceptual[49,50,59,60] and economic[31,61–63] choices, capturing the complex nonlinear dynamics of empirical neurons (**Fig. 1D**). Furthermore, the competitive nature of the RNM generates attractor states which maintain continued activity even in the absence of stimuli, consistent with persistent spiking activity associated with working memory during delay intervals[48–50,59,64–72].

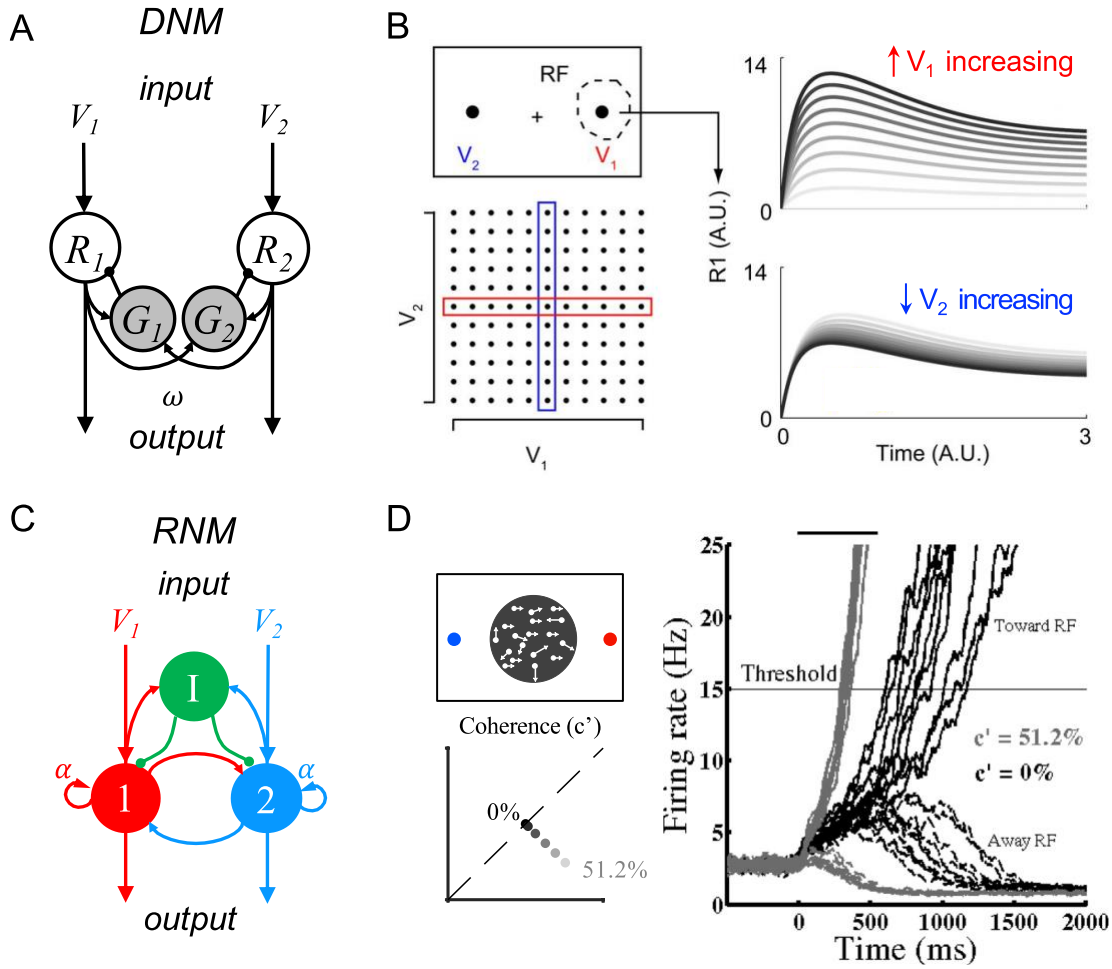

Figure 1. Standard circuit motifs and neural dynamics in existing decision-making models. A. Dynamic normalization model (DNM). Each pair of excitatory (*R*) and inhibitory (*G*) units corresponds to an option in the choice set, with *R* receiving value-dependent input *V* and providing output. Lateral interactions implement a cross-option gain control that produces normalized value coding. **B.** DNM predicted dynamics replicate empirical contextual value coding. Example task involves orthogonal manipulation of both option values. R_1 activity increases with the direct input value V_1 (array framed in red) but is suppressed by the contextual input V_2 (array framed in blue), consistent with value normalization. **C.** Recurrent network model (RNM). The network consists of excitatory pools with self-excitation (1 and 2) and a common pool of inhibitory neurons (I). **D.** RNM predicted dynamics generate winner-take-all (WTA) selection. Example task involves motion discrimination of the main direction of a random dot motion stimulus with varying coherence (c') levels (left). Model activity (right) under two different levels of input coherence (0% and 51.2%) predicts different ramping speeds to the decision threshold, and generates a selection even with equal inputs. Graphics A and B adapted from Louie et al. (2014); graphics C and D adapted from Wong & Wang (2006), Copyright 2006 Society for Neuroscience.

Despite the existence of prominent models for normalized valuation and WTA selection, no current model integrates both properties within a single unified circuit. The current DNM cannot capture late-stage choice dynamics because it lacks a mechanism for WTA competition and categorical choice. Similarly, due to the lack of structured lateral inhibition, RNMs typically neither exhibit contextual value coding nor predict contextual choice patterns[73]. Thus, existing DNM and RNM models highlight a need for an integrated decision model that explains both the normalized value coding and option selection seen in the same single brain regions.

To unify these key features of decision-making, we develop and characterize a novel biological circuit consisting of three neuronal types including local disinhibition. At its core, this model hybridizes the architectural features of divisive gain control and recurrent self-excitation used in existing models, but utilizes disinhibition rather than the commonly-assumed pooled inhibition to implement competition for two reasons. First, in contrast to the non-selective or broadly tuned inhibition seen in visual cortex during stimulus representation[74–80], increasing evidence shows that inhibitory neurons are selective in coding decision variables in frontal cortex[81], parietal cortex[81,82], and striatum[83], suggesting structured inhibition instead pooled inhibition may play a role in decision circuits. Second, GABAergic inhibitory interneurons exhibit a diversity of subtypes with distinct information processing roles in local circuits[84–88], including a prominent role for local disinhibition[89–94]. In disinhibitory architectures, vasoactive intestinal peptide (VIP)-expressing interneurons inhibit the neighboring interneurons expressing somatostatin (SST) or parvalbumin (PV) that inhibit dendritic or perisomatic areas in pyramidal neurons, respectively[91–100] (**Fig. 2C**). Local circuit inputs to VIP neurons suggest that disinhibition may be a key mechanism for generating the mutual competition necessary in option selection and driving

selective inhibition. In addition, given the existence of long-range inputs[85,91,93,94,101] and neuromodulatory inputs[87,88,93,96,102,103] to VIP neurons, local disinhibition has been proposed to play a particular role in dynamic gating of circuit activity; such gating may be essential in decision circuits underlying flexible behavior, mediating top-down control of network function[91,92,94,96,101,104,105]. We find that a novel hybrid model incorporating recurrent excitation, lateral inhibition, and local disinhibition unifies multiple characteristics of decision activity including normalized value coding, WTA choice, and working memory. These findings suggest that local disinhibition is crucial for a biologically plausible integration of normalization and WTA selection in a single circuit architecture. Moreover, a top-down gating signal operating via this disinhibition enables the model to reproduce decision activity in a range of experimental paradigms with diverse task timing and activity dynamics.

Results

Local disinhibition decision model (LDDM)

To develop an integrated circuit model of decision making, we systematically tested a series of models incorporating the core elements of existing models, namely divisive gain control, recurrent excitation, and mutual competition (supplementary **Fig. S1**). This analysis identified *local disinhibition* as the crucial component that can integrate mutual competition and value normalization within the existing circuit architecture of DNM. In the rest of this paper, we focus on the novel disinhibitory hybrid model (local disinhibition decision model, hereafter LDDM).

In the LDDM (**Fig. 2A**), an option-specific disinhibitory *D* unit receives input from its associated excitatory *R* unit and inhibits the inhibitory *G* unit in the local sub-circuit. Biased disinhibition – via different value inputs to option-selective *R* units – can thus introduce an effective asymmetry to global inhibition, generating an unbalanced gain control between local and opponent circuits and leading to WTA competition. In this model, the network shifts from value coding to WTA competition regimes in response to an onset of disinhibition (controlled by the coupling strength between *R* and *D*). With zero or weak *R-D* coupling, the circuit preserves normalized value coding consistent with the DNM; with strong *R-D* coupling, the circuit switches to a state of WTA selection (**Fig. 2B**). Inhibitory units, as a result, dynamically switch from a non-selective response pattern to inputs to a selective response pattern (*G* and *D* units in **Fig. 2B**). This flexible onset of disinhibition is modeled after biological findings, which show that activation of disinhibition in cortical circuits arises from exogenous, long-distance projections[91,94,96,104,105] (**Fig.**

2C). This form of top-down control allows for flexibility in the relative timing of the valuation and selection processes, consistent with neural and behavioral data in different task paradigms (see *Gated disinhibition provides top-down control of choice dynamics*).

Activity dynamics of the LDDM are described by a set of differential equations:

$$\tau_R \frac{dR_i}{dt} = -R_i + \frac{V_i + \alpha R_i + B}{1 + G_i}, \tag{1}$$

$$\tau_G \frac{dG_i}{dt} = -G_i + \sum_{j=1}^{N} \omega_{ij} R_j - D_i, \tag{2}$$

$$\tau_I \frac{dD_i}{dt} = -D_i + \beta R_i. \tag{3}$$

where i = 1, …, N designates choice alternatives, each of which receive selective input V_i and non-selective baseline input B. τ_R, τ_G, and τ_D are the time constants for the R, G, and D units. The weights ω_{ij} represent the coupling strength between excitatory units R_j and gain control units G_i, with each inhibitory unit driven by a weighted sum of its excitatory inputs; the parameter α controls the strength of recurrent self-excitation on R units. Finally, β weights the coupling strength between the excitatory R_i and the disinhibitory D_i units and is presumed to be under external (task-triggered) control.

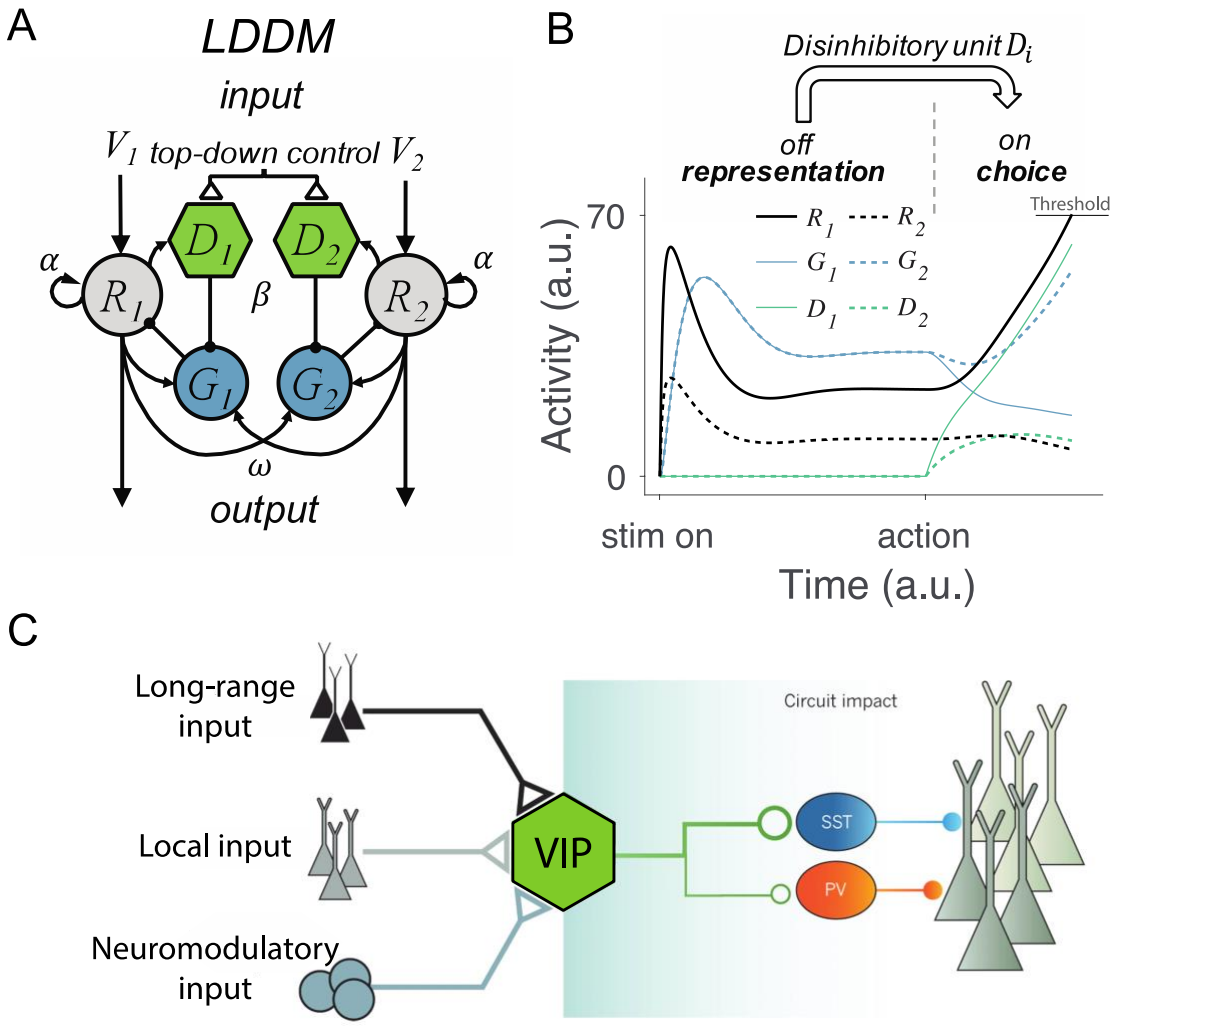

Figure 2. Local disinhibition decision model (LDDM) and its biological plausibility. **A**. LDDM hybridizes DNM and RNM by incorporating a disinhibitory *D* unit to mediate the local disinhibition of the associated excitatory *R* unit; strength of disinhibition is controlled by the parameter β presumed via an external top-down control. V_i, α, and ω indicate the corresponding input value to each option, self-excitation of *R* unit, and the coupling weights from *R* to *G* unit, respectively. **B**. The network phase transition between representation and choice under gated disinhibition. With the disinhibitory module silent, the network performs dynamic divisive normalization on *R* units and predicts non-selective inhibition via *G* units; after the disinhibitory module is triggered via an external top-down control signal, the network switches to a winner-take-all competition dynamic and predicts selective inhibition via *G* and *D* units. **C**. Biological basis of disinhibition. Disinhibition provides a mechanism for dynamic gating of circuit states. Vasoactive intestinal peptide (VIP)-expressing interneurons typically inhibit somatostatin (SST) and parvalbumin (PV)-positive interneurons, resulting in a disinhibition of pyramidal neurons. VIP neurons receive local, long-range, and neuromodulatory input, providing different potential mechanisms to modulate local circuit dynamics. Adapted with permission from Kepecs & Fishell (2014).

149

Dynamic divisive normalization preserved in the LDDM

We first examine whether the LDDM retains the dynamics of divisively normalized value coding seen in the DNM[29,51]. As discussed above, during initial option evaluation the disinhibitory units are silent ($\beta = 0$) and the sole difference between the LDDM and the DNM is recurrent excitation (controlled by α). Example activity traces in **Fig. 3B** show that the LDDM preserves characteristic early-stage dynamics and contextual modulation seen in both empirical data (**Fig. 3C**) and the original DNM[27,29,51]. Immediately after stimulus onset, R_1 activities replicate the transient peak observed in a wealth of studies[11,13,15,17,19,27,29,106]. Further, the network settles to equilibrium displaying relative value coding: R_1 activities increase with V_1 and decrease with V_2 (**Fig. 3B,** R_1 activities across V_1 inputs (upper panels) and V_2 inputs (bottom panels)), reflecting a contextual representation of value (see Methods for details of parameters used in visualization).

Taking advantage of its simplified mathematical form, we analytically analyzed the LDDM and found that it represents each set of input values $(V_1, \dots, V_N)$ as one unique and stable equilibrium point in its output space $(R_1, \dots, R_N)$ when $\beta = 0$. **Fig. 3D** shows that the nullclines of R_1 (solid) and R_2 (dashed), which represent the equilibrium state of each individual unit, intersect at a unique and stable equilibrium point regardless of equal or unequal inputs (see *Equilibria and stability analysis of the LDDM* in **Supplementary Results** for the mathematical proof). The steady state of neural activity (R_i^*) at the equilibrium reflects divisive normalization (Eq. 4), as in the original DNM[29,51]. The only difference at the equilibrium is the negative constant ($-\alpha$) in the denominator introduced by recurrent excitation; this change rescales the activity magnitudes but preserves normalized value coding.

$$R_i^* = \frac{V_i + B}{1 - \alpha + \sum_{j=1}^{N} \omega_{ij} R_j^*} \tag{4}$$

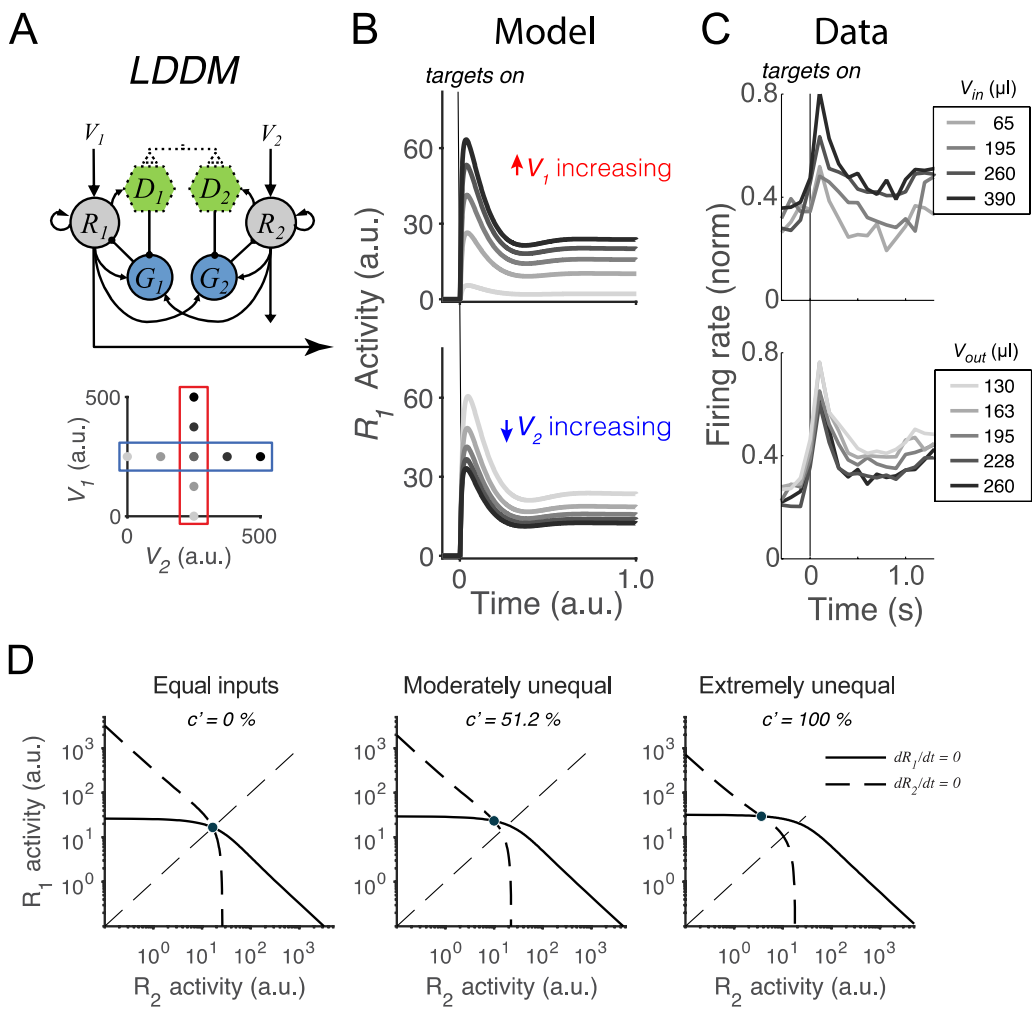

Figure 2. Normalized value coding in the LDDM. **A**. In this example, the LDDM receives a set of binary input values with varying V_1 (framed in red) and V_2 (framed in blue). **B**. Example LDDM dynamics show relative value coding. R_1 activity shows a transient peak before a sustained period of coding. Increasing V_1 increases R_1 activity but increasing V_2 decreases R_1 activity. **C**. Value coding dynamics recorded in monkey parietal cortex. The model prediction we showed is consistent with the empirical observation. Reprinted from Louie et al. (2011). **D**. Phase-plane analysis of the system under equal (left), weakly unequal (middle), and extremely unequal (right) inputs. The nullclines of R_1 (solid) and R_2 (dashed) indicating the equilibrium state of the individual units intersect at a unique and stable equilibrium point with divisively normalized coding.

We next verified that the normalized value coding produced by the LDDM cannot be implemented by standard recurrent RNM models. **Fig. 4A** compares the activity of R_1^* as a function of binary inputs (V_1 and V_2) in the LDDM (left panel), the original DNM (middle panel), and the RNM (right panel). Both the LDDM and the DNM exhibit R_1^* activities (indicated by color) that monotonically increase with input V_1 but decrease with V_2, with a slightly steeper V_2 dependence in the LDDM versus the DNM model depending on the rescaling of α. In contrast, strong WTA dynamics in the RNM implement categorical (choice) coding rather than relative value representation, with high or low binary coding of input values (right panel).

To quantitatively test value normalization, we fit the models to observed firing rates of monkey lateral intraparietal (LIP) neurons under varying reward conditions[27]. In the empirical data (**Fig. 4B**, dots), LIP activity increases with the reward (juice quantity) associated with the target inside the neuronal response field (V_{in}) and decreases with the summed rewards of targets outside the response field (V_{out}). The fitting results show that the DNM precisely captures the rescaled firing rates with only one free parameter (B = 70.6 in Eq. 1) (middle panel in **Fig. 4B**, R^2 = .9640). Importantly, value normalization is equivalently replicated by the LDDM (with one additional parameter α; left panel in **Fig. 4B**, R^2 = .9640, B = 71.2, $\alpha = 0$). Note that the LDDM reduces to the DNM in this case with $\alpha = 0$; however, a non-zero self-excitation parameter may be important in other scenarios (see *A novel form of persistent activity*).

In contrast to the LDDM (and DNM), fitting the standard RNM with four parameters (see supplementary **Fig. S4** for details) does not capture very well the magnitudes of neural activities as a function of V_{in} (right panel in **Fig. 4B**) (R^2 = .8920). The curvature of neural activities as a

function of V_{out} shows a linear type of lateral inhibition, contrast to the concave curvature predicted by divisive normalization in LDDM (and DNM). Furthermore, fitting the RNM to the data results in a parameter regime that can longer generate WTA competition; instead, the model predicts mean firing rates in a low-activity regime with maximum value ~ 3.5 Hz (supplementary **Fig. S4**).

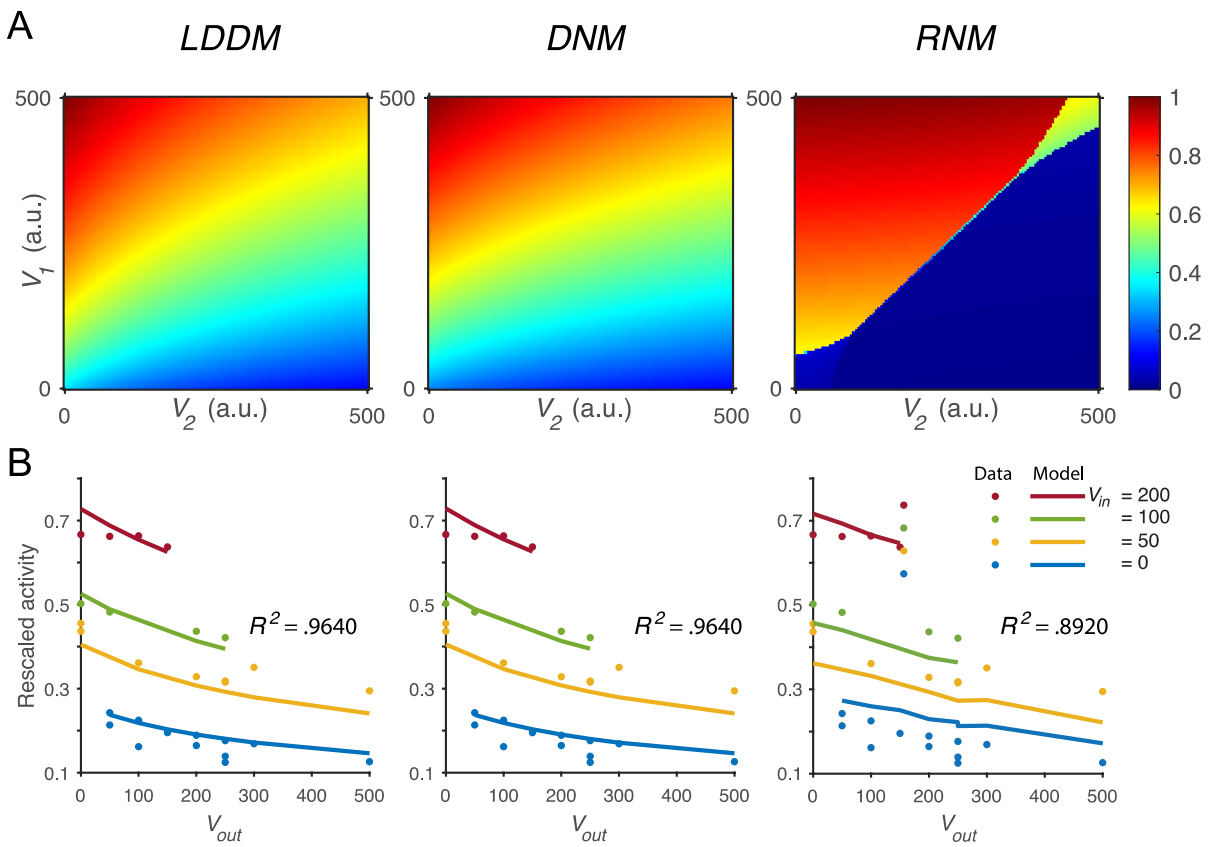

Figure 4. Quantitative comparison of contextual value coding across the LDDM, DNM, and RNM models. **A**. Comparison between the LDDM (left), the DNM (middle), and the RNM (right) in value coding. The LDDM and the DNM show normalized value coding. The neural activity of R_1 (indicated by color) increase with the direct input V_1 but decrease with the contextual input V_2. The LDDM shows slightly stronger contextual modulation than the DNM but qualitatively replicated normalized value coding. The RNM shows a qualitatively different pattern consistent with WTA competition. Within the regime of WTA competition (V_1 and V_2 within a reasonable scale), R_1 activity is high when $V_1 > V_2$ and low when $V_1 < V_2$. **B**. Fitting the models to a trinary choice dataset shows that the LDDM (left panel) and the DNM (middle panel) precisely capture the neural activities responding to values inside (V_{in}) and outside (V_{out}) of the receptive field. Fitting the RNM to the dataset does not capture the neural activities as well as the LDDM (and DNM) (right panel); furthermore, in this parameter regime the RNM no longer implements WTA selection.

Local disinhibition drives winner-take-all competition

A key question is whether the LDDM also produces WTA competition. Given the architecture of the LDDM, local disinhibition is hypothesized to break the symmetry between option-specific *R*-*G* sub-circuits, enabling a competitive interaction between sub-circuits. To examine whether this competition produces WTA selection, we simulated model activity in a reaction-time version of a motion discrimination task, a standard perceptual decision-making paradigm in non-human primates[14,19]. The task contains two stages of processing: the pre-motion stage with only the choice targets presented and the motion stage presenting a random-dot motion stimulus. Animals are allowed to select an option, indicating their percept of the main direction of motion, at any time following motion stimulus onset (see timeline, **Fig. 5A**). During the pre-motion stage, we simulated equal value inputs, given the equal prior probability of either target being correct. The simulated pre-motion dynamics replicate the characteristic transient peak observed in both perceptual and economic decision-making tasks[11,15,19,27]. At motion stimulus onset, inputs to the two *R* units are changed according to the task design; because the animals could make their decision at any time in this reaction-time task, disinhibition is assumed to increase when motion inputs begin.

We find that the LDDM replicates neural and behavioral aspects of WTA competition. In **Fig. 5A**, we show example model activity for five input strengths corresponding to different motion coherence levels. Consistent with electrophysiological recordings in posterior parietal cortex[14,16,19], model *R* unit activities bifurcate based on the input strengths, with the unit receiving stronger input ramping-up to an (arbitrary) decision threshold while the activity of the

opponent unit is suppressed. The speed of bifurcation depends on the contrast between the inputs, a variable equivalent to motion coherence in the experimental literature[14,16]. When we simulated the LDDM with noise (additive Ornstein-Uhlenbeck noise; see Methods for details), the model reproduces the characteristic psychometric and chronometric functions that relate choice accuracy and reaction time (RT) to choice difficulty (**Fig. 5B**). Thus, the LDDM exhibits mutual competition that generates WTA selection, and reproduces empirical decision behavior previously described by RNM models (**Fig. 5A**, inset).

What features of the LDDM are essential to generate WTA competition? As shown in the phase-plane analysis (**Fig. 5C**), the network in the choice regime ($\beta_{on} = .9$ in this example) shows a different configuration of nullcline intersections than the network in the value representation regime ($\beta_{off} = 0$; **Fig. 3B**). Given equal inputs, the nullclines of R_1 and R_2 intersect at three equilibrium points (left panel in **Fig. 5C**), with the central point unstable and the two peripheral points stable. Thus, given an initial configuration of R_1-R_2 activities (in the presence of noise), the system will converge to the closer peripheral attractor (see example activity traces in blue and red thin lines) and implement WTA competition. Given moderately unequal inputs, the basin of attraction is biased towards the side with higher input, resulting in a higher probability falling into the side with higher input (middle panel in **Fig. 5C**). When inputs are extremely unequal, the stable equilibrium in the middle of the basin and the unstable equilibrium point associated with weaker input no longer exist, leaving only the attractor associated with stronger input (**Fig. 5C**, right). Thus, across varying degrees of input coherences, disinhibition drives the LDDM towards a selection of one of the potential choices. This can be seen in **Fig. 5D** by viewing the output ratio ($\frac{R_1^*}{R_1^*+R_2^*}$) of the preferred attractor as a function of input ratio ($\frac{V_1}{V_1+V_2}$): under active

disinhibition we observe categorical coding (green line), in contrast to inactive disinhibition where the output ratio faithfully preserves the original ratio of inputs (dark line).

To understand the operating regimes of the LDDM, we quantified model behavior across the full parameter space defined by recurrent excitation weight (α) and local disinhibition weight (β), both of which are critical in determining the properties of the system (see *Equilibria and stability analysis of the LDDM* in **Supplementary Results** for the mathematical proof). Decisions with equivalent inputs are a critical test of WTA behavior, since WTA systems should select an option (stochastically) even in these symmetric scenarios[48–50,59]; we therefore analyzed system behavior under equal value inputs. As shown in **Fig. 5E**, this analysis revealed two distinct territories corresponding to value representation and WTA operating regimes. The value representation regime generates a unique attractor for normalized value representation but no WTA attractors; in contrast, the WTA regime (induced by a change in β) generates no normalization attractor but instead R_1 and R_2 always diverge into high-contrast attractors (see supplementary **Fig. S2** for a full description of regime parcellation). While the WTA regime asymptotically approaches zero disinhibition when recurrent excitation is extremely strong, local disinhibition is always required to generate WTA choice. Models with a wide range of recurrent excitation can transit from value representation to WTA choice with an increase in local disinhibition strength (for example, red arrow in **Fig. 5E**). These findings emphasize the importance of local disinhibition to WTA choice, and highlight a particular role for a dynamic gating signal in controlling the transition from value coding to option selection.

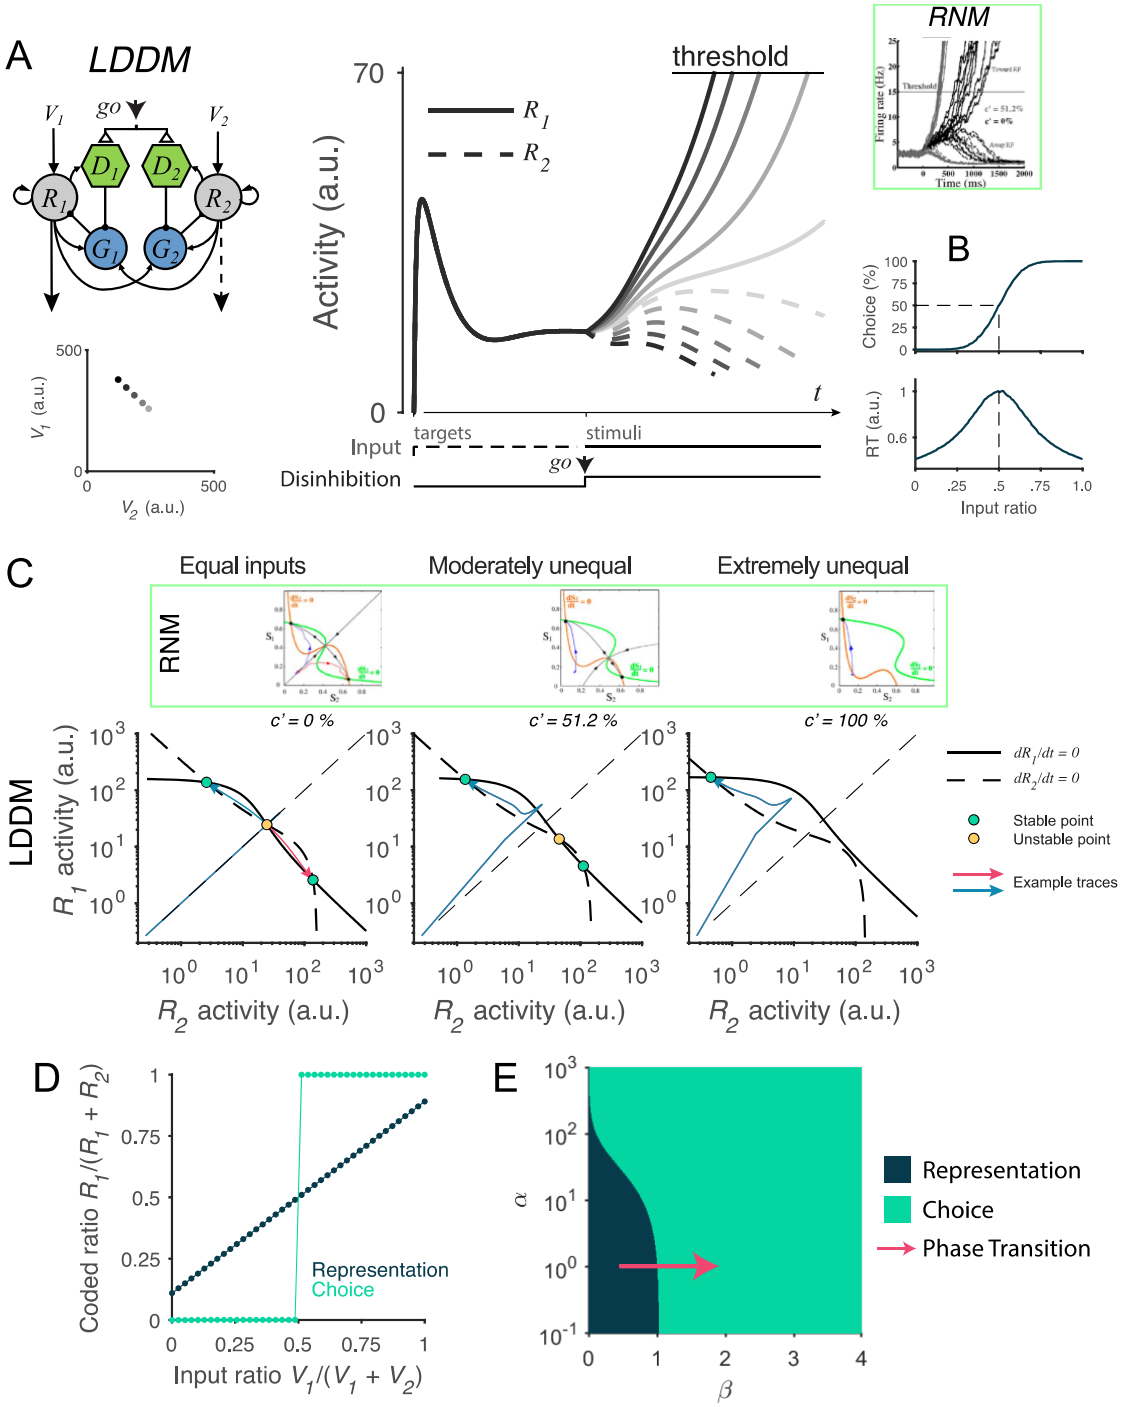

Figure 5. RNM-like WTA selection dynamics in the LDDM. **A**. Example R_1 (solid) and R_2 (dashed) dynamics in a classic reaction-time motion discrimination task. The model predicts phasic stimulus onset dynamics during the pre-stimulus stage and WTA dynamics during the stimulus stage when receiving different input values (left inset). Consistent with RNM dynamics (upper right inset), the *R* unit receiving stronger input ramps up to reach the decision threshold while the opponent *R* unit activity is suppressed; the speed of bifurcation depends on the input strength. **B**. Choice accuracy (top) and RT (bottom) predicted by the LDDM with noise. **C**. Phase-plane analysis of the LDDM (lower) compared with the original RNM (upper inset) shows basis for WTA dynamics under equal (left), moderately unequal (middle), and extremely unequal (right) inputs. Both models show similar features across input values: Under equal inputs, the nullclines of R_1 and R_2 intersect on three equilibrium points, with one unstable point (yellow) and two stable attractors (green) (left). Under unequal inputs, the basin of nullclines is biased to the side with stronger input (middle). When the inputs are strongly biased, only the attractor associated with stronger input retains (right). Red and blue lines show example traces of R_1 and R_2 activities. **D**. Comparison of the coded ratio between the representation (black) and WTA competition (green) regimes. While the LDDM preserves the input ratios during value representation, it shifts to a categorical coding of choice during WTA selection. **E**. Distinct normalized value coding (dark) and WTA competition (green) regimes in the parameter space defined by α and β. Acoss a wide range of α, the transition between valuation and selection regimes can be implemented by an increase in β (pointed by red arrow).

The LDDM captures empirical choice behavior and neural activity

While the preceding analyses show that the LDDM can generate value normalization and WTA selection, a critical question is whether this circuit architecture accurately captures behavioral and neural aspects of empirical decision making. Here, we take advantage of the limited number of parameters in this differential equation-based LDDM (compared to more complicated conductance-based biophysical models[49,50,70,72]), which allows model fitting to empirical data. Specifically, we fit LDDM parameters to nonhuman primate behavior from the reaction-time version of the motion discrimination task described above. These choice and RT data from monkeys align with a reduced form model of decision making (the drift diffusion model)[107], and the activity of posterior parietal neurons recorded during this task display characteristic decision-related features (motion-dependent ramping, a common decision threshold, and WTA activity).

To fit the LDDM to behaviorally observed RTs, we employed the standard quantile maximum likelihood method (QMLE) to the RT distributions across input coherence levels (0 – 51.2%), with correct and error trials dissociated[108–110]. After defining connection weights ω_{ij} as 1 for simplicity, the model is reduced to seven parameters: recurrent excitation weight α, local disinhibition weight β, noise parameter σ, input value scaling parameter S, and time constants τ_R, τ_G, and τ_D (see Methods for model-fitting details). Predictions of the best fitting model are shown in **Fig. 6A** (best fitting parameters: $\alpha = 0$, $\beta = 1.434$, $\sigma = 25.36$, $S = 3251$, $\tau_R = .1853$, $\tau_G = .2244$, and $\tau_D = .3231$). Model-predicted RT distributions (lines) closely follow the empirical distributions (bars) for both correct (blue) and error (red) trials across different levels of input coherence. The aggregated mean choice accuracy and RT data are shown in **Fig. 6C**.

Model choice accuracy (line) captures the average empirical psychometric function (crosses); model RT captures coherence-dependent changes in the chronometric function, including longer RTs in error trials (dashed line and empty dots) compared to correct trials (solid line and dots). Beyond mean RT data, the LDDM accurately captured aspects of the empirical RT distributions, as evident in the quantile-quantile (Q-Q) plot of RT quantiles as functions of chosen ratio (**Fig. 6B**).

We compared the performance of the LDDM in fitting this classical dataset with the reduced form of the RNM, which can be reduced to eight free parameters (see supplementary **Fig. S3** for details)[50]. The negative loglikelihood (nLL) and AIC values of the two models are close, with nLL_{LDDM} = 16546, nLL_{RNM} = 16573, AIC_{LDDM} = 33109, AIC_{RNM} = 33165, suggesting that the LDDM performs as well as the original RNM in fitting behavior in the RT task.

Importantly, the LDDM – fit only to behavior – generates predictions about the underlying neural dynamics that can be compared to electrophysiological findings. We examined *R* unit activity in the best-fitting model, with predicted activity aggregated across trials and aligned to the onset of stimuli and the time of decision as in the original study[14]. Aligned to the onset of stimuli **(Fig. 6D**, left), neural responses are aggregated by coherence level and eventual choice, and truncated at median RT. These data show clear evidence of WTA competition: chosen (solid) and unchosen (dashed) activity traces diverge over time. Moreover, neural activity is stimulus-dependent: the dynamics of both chosen and unchosen units ramp at different, coherence-dependent speeds, consistent with empirical findings consistent with an accumulation process. More quantitatively, we examined the relationship between activity and coherence at the specific

time point (arrow points **a** and **b**) reported in the original work (**Fig. 6E**). Model predictions align well with empirical observations: chosen activity shows a significant increase with input coherence (**a**, 18.56 spikes/second/100% coherence), while unchosen activity shows a decrease (**b**, -30.18 spikes/second/100% coherence).

Aligned to the onset of decision (**Fig. 6D,** right), model *R* unit activity near the time of choice shows further evidence of WTA competition observed in real neurons: the initial divergence between chosen and unchosen activity traces extends into a categorical coding of choice. The relationship between activity and coherence quantitatively replicates the empirical pattern immediately preceding the decision time[14]: chosen activity (indicated by arrow **c** in **Fig. 6D** and plotted in **Fig. 6F**) no longer shows much difference across coherence conditions (-.009 spikes /second/100% coherence), while unchosen activity (indicated by **d** in **Fig. 6D** and plotted in **Fig. 6F**) retains a decrease (**d**, -47.49 spikes/second/100% coherence). Thus, *R* unit activity – in a model with parameters fit only to behavior – replicates the recorded activity of parietal neurons during both initial decision processing and eventual choice selection.

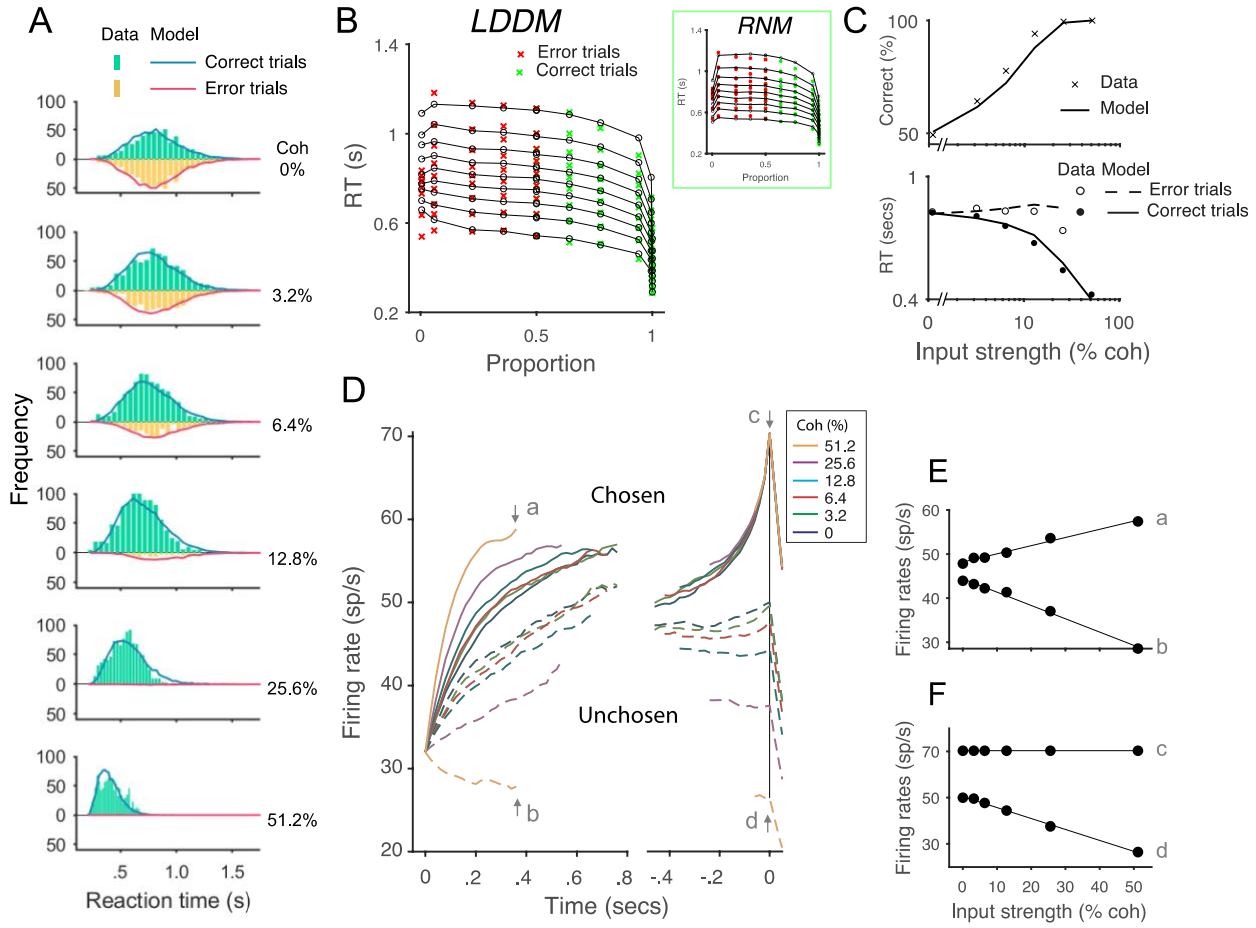

Figure 6. The LDDM performs as well as the RNM in capturing empirical behavior and neurophysiological data during perceptual decision-making. **A**. Model predicted RT distributions fit to behavioral data. Predicted RT distribution (lines) match the histogram of empirical RT distribution (bars), with correct and error trials separately (indicated by color) across levels of input strength (% coherence). **B**. The fitting results of the LDDM and the original RNM (upper-right inset) visualized in Q-Q plots. Nine quantiles of RT under each condition are stacked on the x-axis indicating the correct choice proportion under each input coherence (0 - .5 are error trials, shown in red cross; .5 - 1 are correct trials, shown in green cross). LDDM Model predicts the choice proportion and the shape of RT distribution as well as the original RNM. **C**. Model predicted psychometric function (upper) and chronometric function (lower). Choice accuracy aggregated by input strength (lines) fit well to the empirical data (crosses). The predicted RT aggregated by input strength for correct (solid line) and error (dashed line) trials capture well the RT for correct (filled dots) and error (empty dots) trials in empirical data. **D**. The best-fitting model replicates the neural dynamic features of the recorded neural activity. *R* unit activities aligned to the onset of stimulus inputs (left) and aligned to the time of model decision (right) replicate the stereotyped ramping dynamics of units associated with the chosen side (solid lines) and suppression of units associated with the unchosen side (dashed lines) under different levels of input strength. **E.** Effect of input strength on early stage model activity shortly after stimulus onset, at times indicated by arrows **a** (chosen side) and **b** (unchosen side). Consistent with empirical findings, the activity associated with chosen and unchosen options linearly increase and decrease with input strength. **F.** Effect of input strength on late stage model activity shortly before model choice, at times indicated by arrows **c** and **d**. Consistent with empirical findings, chosen unit activity reaches a common decision bound and show little input dependence, while unchosen activity remains suppressed by input strength. Empirical and behavioral and neural dataset from Roitman & Shadlen (2002).

A novel form of persistent activity

We next examine implications of the local disinhibition architecture for another characteristic of decision-related neural firing: persistent activity. In cortical areas such as parietal[10,14,16,21,111], prefrontal[112–115], and premotor cortices[2], neurons show elevated firing in the absence of stimulus-driven input over intervals of seconds; such persistent activity is thought to underlie working memory and enable decisions based on internally maintained information. In RNMs, recurrent excitation and feedback inhibition preserve categorical choice information after input withdrawal because of point-attractor dynamics[49,50,59]. Here, we answer two questions: does the LDDM generate persistent activity, and how does this persistent activity differ from that in standard RNMs?

We found that the LDDM can generate two distinct forms of persistent activity, controlled by the state of disinhibition. **Fig. 7A** shows example dynamics of two R units before and after withdrawal of inputs while disinhibition is silent. Following input withdrawal, network activity decreases but still preserves elevated firing rates, governed by the self-excitation parameter α (the network loses elevated activity when $\alpha \leq 1$). The persistent activity ratio between R_1 and R_2 preserves the ratio between the input values V_1 and V_2, in contrast to RNMs which only preserve categorical information about the largest value (see supplementary **Fig. S5** for proof). Phase-plane analysis suggests that relative value coding in persistent activity arises from a line-attractor dynamic in the network during the inactivation of disinhibition, unlike point-attractor dynamics in the RNM (**Fig. 7B**). Like other line-attractor models of persistent activity that store continuous-valued information[65,116–118], an unbiased coding of the input ratio requires perfectly

balanced gain control weights from *G* to *R*. Unbalanced weights will result in distorted coding of the input ratio and graded coding of the inputs will decay over time (supplementary **Fig. S5**). For perfectly balanced weights, the line attractor state is vulnerable to noise perturbation. A small perturbation can easily drive the activity to drift on the line of attractors, with the summed value of R_1 and R_2 as a constant ($\frac{\alpha-1}{\omega}$). The preserved ratio between R_1 and R_2 drifts stochastically over time, similar to the prediction of other line-attractor circuits and consistent with behavioral and neural variability related to working memory[117,119].

However, a line attractor is not the only state that the LDDM predicts. If disinhibition is activated during the delay interval, the network switches to a point attractor dynamic (see supplementary **Fig. S6** for proof). **Fig. 7D** shows example dynamics of two *R* units before and after withdrawal of binary inputs. Disinhibition drives a competition between the two *R* units, resulting in a switch between graded coding of the input ratio to a categorical coding of the largest value ($\beta = .4$ in visualization). Interestingly, a transition of coded information from input values to categorical information has been widely observed in firing rates in decision related regions, such as LIP and superior colliculus, during the delay period of decision making[15,16,120]. The point attractor predicted by the circuit under disinhibition (**Fig. 7E**) is highly tolerable to perturbations compared to the line attractor, and choice performance over long delays may require a switch from the value coding to the categorical regimes. As a plausible biological mechanism for mediating top-down control, disinhibition may gate such a transition without changing the network architecture.

The LDDM can be easily expanded to multiple alternatives. Here we show an example of a 5-alternative case, with 5 sets of option-specific *R-G-D* units. A line attractor network with silent disinhibition (**Fig. 7C**, right) is able to retain input value information of the 5 items simultaneously in the network. Due to normalization, the neural activity representing each alternative decreases with the total number of alternatives, with the summed value as a constant ($\sum_1^N R_i = \frac{\alpha-1}{\omega}$), leading to a lower signal-to-noise ratio when coding more items; this set-size effect may be related to WM memory span constraints[121–125]. When disinhibition is active, the LDDM exhibits a point attractor (**Fig. 7F**, right), and the network only holds the information of the largest item as a categorical code during persistent activity.

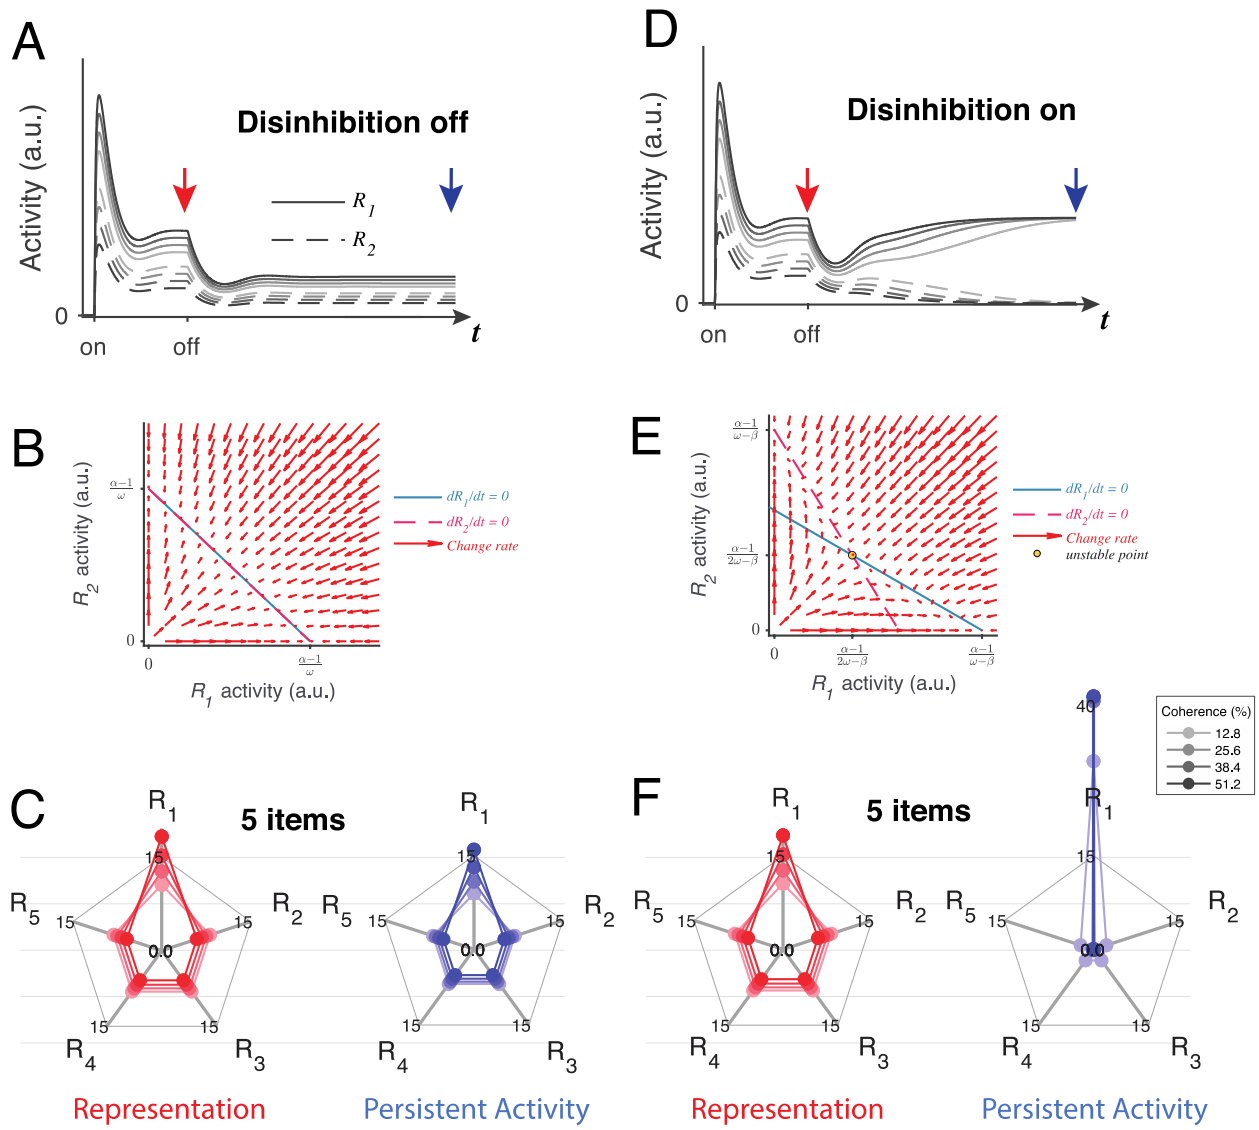

Figure 7. LDDM disinhibition controls flexible implementation of either line attractor or point attractor dynamics in persistent activity**. A-C**. LDDM under silent disinhibition preserves the input ratio information during persistent activity. **A**. Example R_1 (solid) and R_2 (dashed) activities before and after withdrawal of stimuli under different levels of inputs. Neural activity decreases after withdrawal but reaches a new steady that preserves graded coding of the inputs. **B.** Phase-plane analysis of persistent activity exhibits a line attractor under inactivated disinhibition. The nullclines of R_1 (blue) and R_2 (red) intersect on the line of attractors, on which the summed value of R activities is a constant ($\frac{\alpha-1}{\omega}$). Red arrows indicate the instantaneous change rate of R_1 and R_2 at given initial values, following the direction that preserves the R_1-R_2 ratio. **D-F**. Persistent activity under active disinhibition preserves only the largest item as categorical information. **D**. Example R_1 and R_2 activities before and after withdrawal of stimuli. Disinhibition activates at the same time as the offset of stimuli. During the delay period, the activity dynamic gradually switches from a graded coding of the inputs to a WTA type of categorical coding, preserves only the larger item. **E**. Delay period phase-plane analysis exhibits a point-attractor state under activated disinhibition. The nullclines of R_1 (blue) and R_2 (red) intersect on an unstable point. Red arrows indicating the instantaneous change rate of R_1 and R_2 bifurcate from the middle to the side corners, resulting in a high-contrast categorical coding. **C** and **F**. Expansion of the LDDM from a 2-item circuit to a 5-item circuit, under inactivation and activation of disinhibition. Each axis on the radar plot indicates the activity of one R unit. Dots connected with a line indicate the R activities under the same input conditions. The input values change according to coherence level (c') as $S^*[1+c']$ for R_1 and S^* [1-c'] for R_2-R_5. Representation before withdrawal of inputs (left panels in **C** and **F**) and persistent activity without disinhibition (right panel in **C**) preserve the information about the input values. While persistent activity under disinhibition (right panel in **F**) only preserves the item received the largest input, with activities of the other items suppressed.

Gated disinhibition provides top-down control of choice dynamics

In addition to its crucial role in generating WTA competition, local disinhibition provides an intrinsic mechanism for top-down control of choice dynamics. Decision circuits show remarkable flexibility in timing, with similar neurophysiological data recorded in a variety of task paradigms: in addition to reaction-time tasks, in which subjects can choose at any time immediately after onset of stimulus, decision-related neural activity has been widely studied in fixed-duration and delayed-response tasks. In fixed-duration tasks, subjects are required to withhold selection of an action until an instruction signal. Neural activity prior to the instruction signal reflects value information, for example about reward characteristics[13,17,18,27,126] or accumulating perceptual evidence[3,10,14–16,21,111]; however, this activity never fully diverges or reaches the decision threshold until after the instruction cue, suggesting a gating of the selection process. In delayed-response (working memory) tasks, subjects must postpone selection for an interval that includes both stimulus presentation and an additional subsequent interval after the stimulus is withdrawn. As in fixed-duration tasks, neural activity in delayed-response tasks typically carries decision–related information (across both the stimulus and delay periods) but WTA selection – and behavioral choice – is withheld until the instruction cue is given[3,10,14,16,21,111]. Thus, biological decision circuits are able to evaluate choice options while selectively initiating the WTA selection process with variable context-dependent timing.

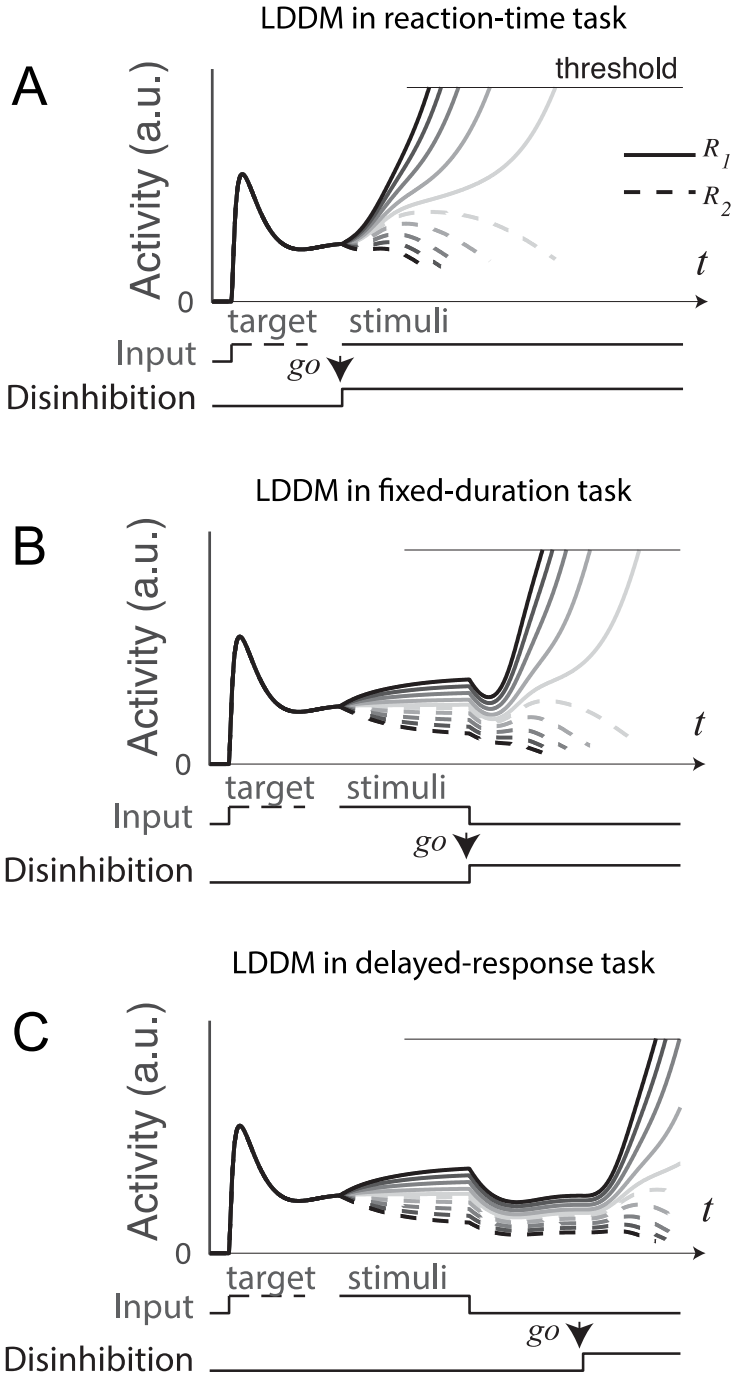

Figure 8. Gated disinhibition flexibly adapts the dynamics of the circuit to various types of tasks. All of the tasks consist of a pre-stimulus stage with equal inputs to R_1 (solid) and R_2 (dashed) and a stimulus stage with input values determined by the stimuli (indicated by grayscale, the same value matrix as used in Figure 5A). **A**. Reaction-time task. Subjects are free to respond at any time following stimulus onset, and model disinhibition is activated with the onset of stimuli. Model dynamics show WTA competition right after the onset of stimuli. **B**. Fixed duration task. Subjects are required to wait for a fixed duration of stimulus viewing before choice, and model disinhibition is turned on only at the onset of the instruction cue (usually indicated in experiments by fixation point offset). Model dynamics show normalized value coding before the instruction cue and a transition to WTA choice afterward. **C**. Working memory (delayed response) task. Subject choice occurs after an interval of stimulus presentation and a subsequent delay interval without stimuli, and model disinhibition is turned on at the end of the delay period. Model dynamics exhibit normalized value coding during stimulus input, preserved relative value information during the delay period, and a transition to WTA choice dynamic after the instruction cue.

How neural circuits implement dynamic control of selection – and a temporal separation of evaluation and WTA choice – is largely unaddressed in current decision models. In RNM models, neural activity is driven by attractor dynamics; option evaluation and the selection process cannot be disambiguated, and WTA competition is not under top-down control. Here, we examine how the timing of a dynamic top-down control signal – modulating the strength of disinhibition via long-range inputs and neuromodulation – allows the LDDM to capture neural activity in different task paradigms. In these simulations, disinhibition is activated when the instruction cue to choose is presented. **Fig. 8A** shows LDDM activity in a reaction-time task, a standard paradigm in perceptual decision-making[14,19]. As in prior analyses (**Figs. 5** and **6**), LDDM *R* units show simultaneous evaluation (coherence-dependent ramping) and WTA selection (rise to threshold) processes, driven by an immediate activation of disinhibition at motion stimulus onset.

In a fixed-duration task (**Fig. 8B**), disinhibition is activated after a required interval of stimulus presentation. Compared to the reaction-time task, LDDM activity here shows distinct, temporally separated patterns during stimuli viewing and option selection; this temporal segregation is driven by the activation of disinhibition (a step function on β in this example), which promotes a transition between value representation and WTA choice.

A further demonstration of this temporal flexibility arises from considering delayed-response tasks (**Fig. 8C**), which include an interval between stimuli offset and onset of the instruction cue. Consistent with its ability to maintain persistent activity (**Fig. 7**), the LDDM shows value coding across the delay interval and implements WTA selection until instruction and accompanying

activation of disinhibition. These results show that the LDDM – via modulation in the timing of disinhibition activation - can temporally separate the value representation and selection processes, enabling it to capture the diversity of neural dynamics seen in reaction-time, fixed-duration, and delayed-response tasks.

Discussion

Normalized value coding, WTA choice, and persistent activity are three important and characteristic features of decision making. We find that a hybrid model that unifies these key features can be implemented with a specific biologically plausible architecture – local disinhibition. In the LDDM, gated disinhibition separates the value normalization and WTA choice computations, enabling them to be generated in the same circuit architecture. When fit to empirical behavioral observation, this hybrid model accurately captures choice and RT patterns, driven by underlying model dynamics that reproduce the neural dynamics of empirical neurophysiological findings. Furthermore, via disinhibitory control, the LDDM can exhibit both line attractor and point attractor forms of persistent activity without a change in the circuit structure. Finally, gated disinhibition in the LDDM provides a mechanism for top-down control of decision dynamics. Controlling the timing of disinhibition paces the decision process and replicates neural dynamics from various choice task variants.

While normalized value coding and WTA selection have largely been modeled separately, the LDDM offers a biologically-plausible circuit architecture that integrates the two features. Existing neurophysiological evidence show that WTA dynamics and normalized coding co-exist in the same brain regions. On the one hand, neural activities show relative value coding in the early stage of decision-making, reflecting a context-dependent modulation consistent with the canonical divisive normalization computation[2,9,15,19,26,27,30]. On the other hand, WTA choice dynamics are widely observed during decision making across multiple brain regions of non-human primates[2–25], including many of the brain regions that show normalized value coding. In

addition, a transition from graded coding to WTA choice has been widely documented in the decision relevant regions mentioned above. Neural firing rates shows a graded coding of perceptual evidence and reward during the early stage of decision-making and gradually transition to a categorical coding for choice in the late period of decision-making[13–19,45,120,127]. However, the evidence for one alternative is typically inversely related to the evidence for the other alternative, making it difficult to dissociate the dynamic effects of evidence integration and contextual information about other alternatives.

In the LDDM, disinhibition modulates the dynamics of the circuit without requiring changes in circuit structure. Existing models capture activity dynamics only in specific temporal intervals during decision-making tasks, or across trials in specific task paradigms[29,31,49,50,67], and thus typically do not generalize across tasks. In contrast, gated disinhibition in the LDDM – driven by the external action instruction cue - controls the timing of valuation-to-WTA regime transition, enabling the LDDM to replicate neural dynamics in diverse task paradigms with different stimulus and action timing schedules[14–16,21]. Recent research on neuromodulatory control of disinhibition offers biologically plausible mechanisms for such top-down control of circuit dynamics. In addition to evidence that VIP neurons are recruited by long-range projections from distanced regions[91,105], VIP neurons are recruited by neuromodulatory projections such as acetylcholine[96] from the basal forebrain and pedunculopontine nuclei and serotonin from the red nucleus. With ionotropic acetylcholine receptor (nAChR) and serotonin receptors ($5HT_{3a}R$ and $5HT_2R$), VIP neurons depolarize to acetylcholine and serotonin[87,88,93,102]. The spiking mode of a major type of VIP neurons in layer II/III of the cortex switches from an input-insensitive burst-quiescent mode to an input-sensitive tonic mode under cholinergic and serotonin modulation[103].

Such a mode-switching feature allows the disinhibitory neurons to receive excitatory projections with different gain under different level of neuromodulation, providing a mechanism to modulate network dynamics via disinhibition without a change in network structure. *In vivo* studies show that disinhibition mediated by cholinergic activation is triggered in a surprisingly fast time scale of tens of milliseconds[92,102], supporting a fast modulation mechanism of disinhibition and network plasticity.

An interesting feature of the LDDM is that it can produce both point attractor[128–131] and continuous/line attractor[116,119,132] dynamics in persistent activity, a balance controlled by the level of disinhibition. Given ambiguous empirical evidence, it remains controversial whether persistent activity in neural circuits exhibit point attractor[135–138] or continuous/line attractor[116,119,139] dynamics, and existing circuit models of persistent activity exclusively predict either a point attractor[64,72,133,134] or line attractor[65,116–118,135]. The LDDM can generate both line attractor and point attractor states, suggesting that attractor dynamics might not be a fixed property of a network; rather, it may be adaptive and controllable by a top-down signal operating via gated disinhibition.

In conclusion, here we introduce a novel, biologically-plausible architecture for decision making based on local disinhibition, unifying the characteristic decision-making features of normalized value coding, WTA competition, and persistent activity into a single circuit. The LDDM captures psychometric and chronometric aspects of behavioral choice and predicts realistic neural dynamics in standard decision-making tasks. Local disinhibition provides a mechanism for top-down control of local decision circuit dynamics, enabling the LDDM to replicate variable task-

dependent timing in diverse decision-making paradigms and implement speed-accuracy tradeoffs. These results suggest a new circuit mechanism for decision making, and emphasize the importance of incorporating interneuron diversity, local circuit architecture, and top-down control into models of the decision process.

Methods

Numerical simulations

To quantify neural dynamics and behavioral performance (choice/RT), time varying activity was represented by a system of differential equations (see Results) which was solved numerically using the Runge-Kutta method implemented in MATLAB (MathWorks) at time step of 1 ms. Evaluations using smaller time steps (0.1 ms) were examined and produced similar results. At each time step, the model unit activities were updated based on their values at the previous step according to the differential equations. Considering the biological reality that spike rates cannot be negative, the activities were constrained to be non-negative. For the simulations including noise, we assumed an additive noise term for each unit, which evolved independently based on an Ornstein-Uhlenbeck process (Eq. 5),

$$\tau_{noise}\frac{dNoise(t)}{dt} = -Noise(t) + \eta(t)\sqrt{\sigma^2 dt} \qquad (5)$$

where σ^2 is the variance of the noise, η is a Gaussian white noise with zero mean and unit variance, and τ_{noise} is the time constant for the noise fluctuation process. The time constant for the noise process (τ_{noise}) was set to 2 ms aligned with previous studies[49,50].

All parameters used for visualization were set as following unless specified elsewhere or fitted as free parameters in **Fig. 6**: τ_R, τ_G, and τ_D were set as the same value of 100 ms; the gain control weight ω_{ij} was set as unit value of 1 for simplicity; the self-excitation weight α was set as 15; the disinhibition weight β was assumed as zero in representation (i.e., $\beta_{off} = 0$) and set as 1.1 in WTA competition; the input values V_1 and V_2 were set as S*(1+c') and S*(1-c'), where c'

indicates the motion coherence of the stimulus, with varied values [0, 3.2, 6.4, 12.8, 25.6, 38.4, 51.2]%) and *S* indicates the scale of input (set as 250). Baseline input *B* was set as 70 in **Figs. 3** and **5** and set as 0 in **Fig. 7**. Ornstein-Uhlenbeck noise was set as zero in most of the simulation but $\sigma = 10$ in **Fig. 5B**. In **Fig. 7**, parameters are adjusted between 2-item and 5-item cases in order to get comparable scale of activities in visualization: for 2-item case, $S = 250$, $\alpha = 15$, and $\beta_{on} = .4$; for 5-item case, $S = 50$, $\alpha = 37.5$, and $\beta_{on} = .1$.

Fit the LDDM to empirical behavioral data

The LDDM with seven free parameters (the weights of self-excitation (α) and disinhibition (β), the variance of Gaussian white noise in the Ornstein-Uhlenbeck process (σ^2), the scaling parameter of input (*S*), and time constants for three types of units τ_R, τ_G, and τ_D) was fit to choice behavior (RT and choice accuracy) in a classic perceptual decision-making dataset[14]. We employed the commonly used quantile maximum likelihood estimation (QMLE) method[109,136]. The rationale of QMLE is to minimize the differences between the predicted data and the empirical data on the proportion of number of trials located in each RT bin. Choice accuracy was implicitly estimated because the algorithm accounts for the proportion of number of trials between correct and error trials. Nine quantiles (from .1 to .9 with .1 of step size) were used resulting in ten RT bins, with correct and error trials accounted for separately at each coherence level. Because the LDDM has no closed-form analytic expression for the RT distribution, we evaluated the prediction by Monte Carlo simulations (10240 repetitions for each input coherence). In each simulated trial, the initial values of R_1 and R_2 activities were set as 32 Hz to be comparable to empirical data[14,19]. Visual stimulus (motion) inputs were defined as S*(1+c')

and S*(1-c') for V_1 and V_2, where the free parameter S models input scaling and the coherence c' replicated values in the original experiment (0, 3.2, 6.4, 12.8, 25.6, and 51.2 %)[14]. At visual stimulus onset, a gap period (90 ms) was implemented in order to capture the commonly observed initial dip in empirical firing rates[14]. Gated disinhibition was activated along with inputs after the gap. A decision was reached when either of the *R* unit activities reached a decision threshold of 70 Hz, the biological threshold observed in the empirical data[14]. 30 ms was added to the RT of threshold hitting to capture the delay in the down-streaming motor execution. After the decision, the input values, self-excitation, and disinhibition were reset to zero. The negative loglikelihood (nLL) of QMLE was minimized using Bayesian adaptive direct search algorithm in Matlab[137]. The estimation was conducted using GPU (NVIDIA Tesla V100) parallel computation on a high-performance cluster (NYU Langone), with 160 chains of random initial parameter values to prevent local minima. The chain with the smallest nLL in its fitting result was selected as the best fitting result.

The visualization of the predicted RT distribution (**Fig. 6A**) was calculated based on 60 evenly distributed RT bins, with correct and error trials calculated separately under each coherence. The predicted neural dynamics (**Fig. 6D**) were generated using the model best fit to behavior. *R* unit activities were aggregated across correct trials, segregated by units associated with the chosen side and unchosen side. As in the original experiment data visualization[14], activity early in trials was aligned to stimulus onset and data within 100 ms of boundary crossing were omitted to reduce the impact of decision dynamics on visualizing early stage ramping dynamics. Early activity traces were cut off at the median value of RT for each coherence level to ensure that the

average trace was based on at least half of the trials. Activity late in trials was aligned to the time of decision, and data within 200 ms of stimulus onset was omitted.

Fit the LDDM and the DNM to the neural firing rates of normalized value coding

In order to quantify the performance of the LDDM in fitting to the neural dynamics of normalized value coding and compare with the original DNM, we fit the equilibrium values of the LDDM and DNM to the dataset of normalized value coding (**Fig. 4** in Louie et al., 2011). In this task, monkeys are asked to represent the reward targets (1, 2, or 3) on the corresponding location of the screen. The neural activity in the response field receiving direct input V_1 is recorded. Different combinations of V_1, V_2, and V_3 are provided to the monkeys based on the associated volumes of juice in the presented targets (varying from 50, 100, 200, and 250 µl or omitted target marked as 0; see details in Louie et al., 2011), resulting in 28 data points. The LDDM and DNM are expanded to a trinary-choice circuit, with R_1, R_2, and R_3 receiving input from V_1, V_2, and V_3, respectively. One free parameter is fitted for the DNM (Setting $\alpha = 0$ reduces the form of the LDDM in Eq. 1 to DNM; Baseline input B in Eq. 1 is fitted). And two free parameters are fitted for LDDM (α and B in Eq. 1). The fitted neural firing rates were further rescaled by a scaling parameter S to match the rescaled empirical data points (**Fig. 4B**). Other parameters in Eq. 1 – 3 are set as following: The dimensionless values of V_1, V_2, and V_3 set equals to the volumes of juice reported in the original paper, $\beta = 0$, $\omega_{ij} = 1$, $\tau_R = \tau_G = \tau_D = 100\,ms$, and Ornstein-Uhlenbeck noise $\sigma = 0$. Bayesian adaptive direct search algorithm

(BADS[137]) is implemented to minimize the ordinary squared error between the steady state of the predicted neural firing rates on R_1 and the empirical data.

References

1. Glimcher, P. W., Camerer, C. F., Fehr, E. & Poldrack, R. A. Chapter 1 - Introduction: A Brief History of Neuroeconomics. in *Neuroeconomics* (eds. Glimcher, P. W., Camerer, C. F., Fehr, E. & Poldrack, R. A.) 1–12 (Academic Press, 2009). doi:10.1016/B978-0-12-374176-9.00001-4.
2. Pastor-Bernier, A. & Cisek, P. Neural Correlates of Biased Competition in Premotor Cortex. *J. Neurosci.* **31**, 7083–7088 (2011).
3. Kim, J.-N. & Shadlen, M. N. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* **2**, 176–185 (1999).
4. Padoa-Schioppa, C. Neuronal Origins of Choice Variability in Economic Decisions. *Neuron* **80**, 1322–1336 (2013).
5. Padoa-Schioppa, C. & Conen, K. E. Orbitofrontal Cortex: A Neural Circuit for Economic Decisions. *Neuron* **96**, 736–754 (2017).
6. Roesch, M. R. & Olson, C. R. Impact of Expected Reward on Neuronal Activity in Prefrontal Cortex, Frontal and Supplementary Eye Fields and Premotor Cortex. *J. Neurophysiol.* **90**, 1766–1789 (2003).
7. Thura, D. & Cisek, P. Deliberation and Commitment in the Premotor and Primary Motor Cortex during Dynamic Decision Making. *Neuron* **81**, 1401–1416 (2014).
8. Thura, D. & Cisek, P. Modulation of premotor and primary motor cortical activity during volitional adjustments of speed-accuracy trade-offs. *J. Neurosci.* **36**, 938–956 (2016).
9. Yamada, H., Louie, K., Tymula, A. & Glimcher, P. W. Free choice shapes normalized value signals in medial orbitofrontal cortex. *Nat. Commun.* **9**, 162 (2018).

10. Kiani, R., Cueva, C. J., Reppas, J. B. & Newsome, W. T. Dynamics of Neural Population Responses in Prefrontal Cortex Indicate Changes of Mind on Single Trials. *Curr. Biol.* **24**, 1542–1547 (2014).

11. Andersen, R. A. & Buneo, C. A. Intentional Maps in Posterior Parietal Cortex. *Annu. Rev. Neurosci.* **25**, 189–220 (2002).

12. Louie, K. & Glimcher, P. W. Separating Value from Choice: Delay Discounting Activity in the Lateral Intraparietal Area. *J. Neurosci.* **30**, 5498–5507 (2010).

13. Platt, M. L. & Glimcher, P. W. Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238 (1999).

14. Roitman, J. D. & Shadlen, M. N. Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task. *J. Neurosci.* **22**, 9475–9489 (2002).

15. Rorie, A. E., Gao, J., McClelland, J. L. & Newsome, W. T. Integration of Sensory and Reward Information during Perceptual Decision-Making in Lateral Intraparietal Cortex (LIP) of the Macaque Monkey. *PLOS ONE* **5**, e9308 (2010).

16. Shadlen, M. N. & Newsome, W. T. Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. *J. Neurophysiol.* **86**, 1916–1936 (2001).

17. Sugrue, L. P., Corrado, G. S. & Newsome, W. T. Matching Behavior and the Representation of Value in the Parietal Cortex. *Science* **304**, 1782–1787 (2004).

18. Dorris, M. C. & Glimcher, P. W. Activity in Posterior Parietal Cortex Is Correlated with the Relative Subjective Desirability of Action. *Neuron* **44**, 365–378 (2004).

19. Churchland, A. K., Kiani, R. & Shadlen, M. N. Decision-making with multiple alternatives. *Nat. Neurosci.* **11**, 693–702 (2008).

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

21. Kiani, R., Hanks, T. D. & Shadlen, M. N. Bounded Integration in Parietal Cortex Underlies Decisions Even When Viewing Duration Is Dictated by the Environment. *J. Neurosci.* **28**, 3017–3029 (2008).

22. Ding, L. & Gold, J. I. Caudate Encodes Multiple Computations for Perceptual Decisions. *J. Neurosci.* **30**, 15747–15759 (2010).

23. Ding, L. & Gold, J. I. Neural Correlates of Perceptual Decision Making before, during, and after Decision Commitment in Monkey Frontal Eye Field. *Cereb. Cortex* **22**, 1052–1067 (2012).

24. Ding, L. & Gold, J. I. The Basal Ganglia's Contributions to Perceptual Decision Making. *Neuron* **79**, 640–649 (2013).

25. Thura, D. & Cisek, P. The Basal Ganglia Do Not Select Reach Targets but Control the Urgency of Commitment. *Neuron* **95**, 1160-1170.e5 (2017).

26. Kira, S., Yang, T. & Shadlen, M. N. A Neural Implementation of Wald's Sequential Probability Ratio Test. *Neuron* **85**, 861–873 (2015).

27. Louie, K., Grattan, L. E. & Glimcher, P. W. Reward Value-Based Gain Control: Divisive Normalization in Parietal Cortex. *J. Neurosci.* **31**, 10627–10639 (2011).

28. Louie, K., Khaw, M. W. & Glimcher, P. W. Normalization is a general neural mechanism for context-dependent decision making. *Proc. Natl. Acad. Sci.* **110**, 6139–6144 (2013).

29. Louie, K., LoFaro, T., Webb, R. & Glimcher, P. W. Dynamic Divisive Normalization Predicts Time-Varying Value Coding in Decision-Related Circuits. *J. Neurosci.* **34**, 16046–16057 (2014).

30. Strait, C. E., Blanchard, T. C. & Hayden, B. Y. Reward Value Comparison via Mutual Inhibition in Ventromedial Prefrontal Cortex. *Neuron* **82**, 1357–1366 (2014).

31. Hunt, L. T. *et al.* Mechanisms underlying cortical activity during value-guided choice. *Nat. Neurosci.* **15**, 470–476 (2012).

32. Louie, K., Glimcher, P. W. & Webb, R. Adaptive neural coding: from biological to behavioral decision-making. *Curr. Opin. Behav. Sci.* **5**, 91–99 (2015).

33. Carandini, M., Heeger, D. J. & Anthony Movshon, J. Linearity and Gain Control in V1 Simple Cells. in *Models of Cortical Circuits* (eds. Ulinski, P. S., Jones, E. G. & Peters, A.) 401–443 (Springer US, 1999). doi:10.1007/978-1-4615-4903-1_7.

34. Carandini, M. & Heeger, D. J. Summation and division by neurons in primate visual cortex. *Science* **264**, 1333–1336 (1994).

35. Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 (2012).

36. Heeger, D. J. Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* **9**, 181–197 (1992).

37. Heeger, D. J. Modeling simple-cell direction selectivity with normalized, half-squared, linear operators. *J. Neurophysiol.* **70**, 1885–1898 (1993).

38. Schwartz, O. & Simoncelli, E. P. Natural signal statistics and sensory gain control. *Nat. Neurosci.* **4**, 819–825 (2001).

39. Silver, R. A. Neuronal arithmetic. *Nat. Rev. Neurosci.* **11**, 474–489 (2010).

40. Chau, B. K., Law, C.-K., Lopez-Persem, A., Klein-Flügge, M. C. & Rushworth, M. F. Consistent patterns of distractor effects during decision making. *eLife* **9**, e53850 (2020).

41. Steverson, K., Brandenburger, A. & Glimcher, P. Choice-theoretic foundations of the divisive normalization model. *J. Econ. Behav. Organ.* **164**, 148–165 (2019).

42. Webb, R., Glimcher, P. W. & Louie, K. Rationalizing Context-Dependent Preferences: Divisive Normalization and Neurobiological Constraints on Choice. *SSRN Electron. J.* (2014) doi:10.2139/ssrn.2462895.

43. Wickens, J. R., Arbuthnott, G. W. & Shindou, T. Simulation of GABA function in the basal ganglia: computational models of GABAergic mechanisms in basal ganglia function. in *Progress in Brain Research* (eds. Tepper, J. M., Abercrombie, E. D. & Bolam, J. P.) vol. 160 313–329 (Elsevier, 2007).

44. Wilson, C. J. GABAergic inhibition in the neostriatum. in *Progress in Brain Research* (eds. Tepper, J. M., Abercrombie, E. D. & Bolam, J. P.) vol. 160 91–110 (Elsevier, 2007).

45. Gold, J. I. & Shadlen, M. N. The Neural Basis of Decision Making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).

46. Hanes, D. P. & Schall, J. D. Neural Control of Voluntary Movement Initiation. *Science* **274**, 427–430 (1996).

47. Lo, C.-C., Wang, C.-T. & Wang, X.-J. Speed-accuracy tradeoff by a control signal with balanced excitation and inhibition. *J. Neurophysiol.* **114**, 650–661 (2015).

48. 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).

49. Wang, X.-J. Probabilistic Decision Making by Slow Reverberation in Cortical Circuits. *Neuron* **36**, 955–968 (2002).

50. Wong, K.-F. & Wang, X.-J. A Recurrent Network Mechanism of Time Integration in Perceptual Decisions. *J. Neurosci.* **26**, 1314–1328 (2006).

51. LoFaro, T., Louie, K., Webb, R. & Glimcher, P. W. The Temporal Dynamics of Cortical Normalization Models of Decision-making. *Lett. Biomath.* 12 (2014).

52. Houck, B. D. & Person, A. L. Cerebellar Loops: A Review of the Nucleocortical Pathway. *The Cerebellum* **13**, 378–385 (2014).

53. Ito, M. Historical Review of the Significance of the Cerebellum and the Role of Purkinje Cells in Motor Learning. *Ann. N. Y. Acad. Sci.* **978**, 273–288 (2002).

54. Ito, M. Cerebellar circuitry as a neuronal machine. *Prog. Neurobiol.* **78**, 272–303 (2006).

55. Ito, M. Control of mental activities by internal models in the cerebellum. *Nat. Rev. Neurosci.* **9**, 304–313 (2008).

56. Llinás, R. R. THE CORTEX OF THE CEREBELLUM. *Sci. Am.* **232**, 56–71 (1975).

57. Sathyanesan, A. *et al.* Emerging connections between cerebellar development, behaviour and complex brain disorders. *Nat. Rev. Neurosci.* **20**, 298–313 (2019).

58. Sillitoe, R. V. & Joyner, A. L. Morphology, Molecular Codes, and Circuitry Produce the Three-Dimensional Complexity of the Cerebellum. *Annu. Rev. Cell Dev. Biol.* **23**, 549–577 (2007).

59. Furman, M. & Wang, X.-J. Similarity Effect and Optimal Control of Multiple-Choice Decision Making. *Neuron* **60**, 1153–1168 (2008).

60. Wong, K.-F., Huk, A. C., Shadlen, M. N. & Wang, X.-J. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front. Comput. Neurosci.* **1**, (2007).

61. Jocham, G., Hunt, L. T., Near, J. & Behrens, T. E. J. A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex. *Nat. Neurosci.* **15**, 960–961 (2012).

62. Rustichini, A. & Padoa-Schioppa, C. A neuro-computational model of economic decisions. *J. Neurophysiol.* **114**, 1382–1398 (2015).

63. Soltani, A. A Biophysically Based Neural Model of Matching Law Behavior: Melioration by Stochastic Synapses. *J. Neurosci.* **26**, 3731–3744 (2006).

64. Brunel, N. & Wang, X.-J. Effects of Neuromodulation in a Cortical Network Model of Object Working Memory Dominated by Recurrent Inhibition. *J. Comput. Neurosci.* **11**, 63–85 (2001).

65. Compte, A. Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model. *Cereb. Cortex* **10**, 910–923 (2000).

66. Constantinidis, C. *et al.* Persistent Spiking Activity Underlies Working Memory. *J. Neurosci.* **38**, 7020–7028 (2018).

67. Hart, E. & Huk, A. C. Recurrent circuit dynamics underlie persistent activity in the macaque frontoparietal network. *eLife* **9**, e52460 (2020).

68. Macoveanu, J., Klingberg, T. & Tegnér, J. A biophysical model of multiple-item working memory: A computational and neuroimaging study. *Neuroscience* **141**, 1611–1618 (2006).

69. Murray, J. D., Jaramillo, J. & Wang, X.-J. Working Memory and Decision-Making in a Frontoparietal Circuit Model. *J. Neurosci.* **37**, 12167–12186 (2017).

70. Tegnér, J., Compte, A. & Wang, X.-J. The dynamical stability of reverberatory neural circuits. *Biol. Cybern.* **87**, 471–481 (2002).

71. Wang, M. *et al.* NMDA Receptors Subserve Persistent Neuronal Firing during Working Memory in Dorsolateral Prefrontal Cortex. *Neuron* **77**, 736–749 (2013).

72. Wang, X.-J. Synaptic Basis of Cortical Persistent Activity: the Importance of NMDA Receptors to Working Memory. *J. Neurosci.* **19**, 9587–9603 (1999).

73. Wang, X.-J. Neural dynamics and circuit mechanisms of decision-making. *Curr. Opin. Neurobiol.* **22**, 1039–1046 (2012).

74. Bock, D. D. *et al.* Network anatomy and in vivo physiology of visual cortical neurons. *Nature* **471**, 177–182 (2011).

75. Chen, T.-W. *et al.* Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300 (2013).

76. Hofer, S. B. *et al.* Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nat. Neurosci.* **14**, 1045–1052 (2011).

77. Kerlin, A. M., Andermann, M. L., Berezovskii, V. K. & Reid, R. C. Broadly Tuned Response Properties of Diverse Inhibitory Neuron Subtypes in Mouse Visual Cortex. *Neuron* **67**, 858–871 (2010).

78. Liu, B. *et al.* Visual Receptive Field Structure of Cortical Inhibitory Neurons Revealed by Two-Photon Imaging Guided Recording. *J. Neurosci.* **29**, 10520–10532 (2009).

79. Niell, C. M. & Stryker, M. P. Highly Selective Receptive Fields in Mouse Visual Cortex. *J. Neurosci.* **28**, 7520–7536 (2008).

80. Sohya, K., Kameyama, K., Yanagawa, Y., Obata, K. & Tsumoto, T. GABAergic Neurons Are Less Selective to Stimulus Orientation than Excitatory Neurons in Layer II/III of Visual Cortex, as Revealed by In Vivo Functional Ca2+ Imaging in Transgenic Mice. *J. Neurosci.* **27**, 2145–2149 (2007).

81. Allen, W. E. *et al.* Global Representations of Goal-Directed Behavior in Distinct Cell Types of Mouse Neocortex. *Neuron* **94**, 891-907.e6 (2017).

82. Najafi, F. *et al.* Excitatory and Inhibitory Subnetworks Are Equally Selective during Decision-Making and Emerge Simultaneously during Learning. *Neuron* **105**, 165-179.e8 (2020).

83. Gage, G. J., Stoetzner, C. R., Wiltschko, A. B. & Berke, J. D. Selective Activation of Striatal Fast-Spiking Interneurons during Choice Execution. *Neuron* **67**, 466–479 (2010).

84. Jiang, X. *et al.* Principles of connectivity among morphologically defined cell types in adult neocortex. *Science* **350**, (2015).

85. Kepecs, A. & Fishell, G. Interneuron cell types are fit to function. *Nature* **505**, 318–326 (2014).

86. Markram, H. *et al.* Interneurons of the neocortical inhibitory system. *Nat. Rev. Neurosci.* **5**, 793–807 (2004).

87. Rudy, B., Fishell, G., Lee, S. & Hjerling-Leffler, J. Three groups of interneurons account for nearly 100% of neocortical GABAergic neurons. *Dev. Neurobiol.* **71**, 45–61 (2011).

88. Tremblay, R., Lee, S. & Rudy, B. GABAergic Interneurons in the Neocortex: From Cellular Properties to Circuits. *Neuron* **91**, 260–292 (2016).

89. Acsády, L., Görcs, T. J. & Freund, T. F. Different populations of vasoactive intestinal polypeptide-immunoreactive interneurons are specialized to control pyramidal cells or interneurons in the hippocampus. *Neuroscience* **73**, 317–334 (1996).

90. Hájos, N., Acsády, L. & Freund, T. F. Target Selectivity and Neurochemical Characteristics of VIP-immunoreactive Interneurons in the Rat Dentate Gyrus. *Eur. J. Neurosci.* **8**, 1415–1431 (1996).

91. Lee, S., Kruglikov, I., Huang, Z. J., Fishell, G. & Rudy, B. A disinhibitory circuit mediates motor integration in the somatosensory cortex. *Nat. Neurosci.* **16**, 1662–1670 (2013).

92. Letzkus, J. J. *et al.* A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature* **480**, 331–335 (2011).

93. Pfeffer, C. K., Xue, M., He, M., Huang, Z. J. & Scanziani, M. Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nat. Neurosci.* **16**, 1068–1076 (2013).

94. Pi, H.-J. *et al.* Cortical interneurons that specialize in disinhibitory control. *Nature* **503**, 521–524 (2013).

95. Chiu, C. Q. *et al.* Compartmentalization of GABAergic Inhibition by Dendritic Spines. *Science* **340**, 759–762 (2013).

96. Fu, Y. *et al.* A Cortical Circuit for Gain Control by Behavioral State. *Cell* **156**, 1139–1152 (2014).

97. Urban-Ciecko, J. & Barth, A. L. Somatostatin-expressing neurons in cortical networks. *Nat. Rev. Neurosci.* **17**, 401–409 (2016).

98. Karnani, M. M. *et al.* Opening Holes in the Blanket of Inhibition: Localized Lateral Disinhibition by VIP Interneurons. *J. Neurosci.* **36**, 3471–3480 (2016).

99. Fino, E. & Yuste, R. Dense inhibitory connectivity in neocortex. *Neuron* **69**, 1188–1203 (2011).

100. Karnani, M. M., Agetsuma, M. & Yuste, R. A blanket of inhibition: functional inferences from dense inhibitory connectivity. *Curr. Opin. Neurobiol.* **26**, 96–102 (2014).

101. Schuman, B., Dellal, S., Prönneke, A., Machold, R. & Rudy, B. Neocortical Layer 1: An Elegant Solution to Top-Down and Bottom-Up Integration. *Annu. Rev. Neurosci.* **44**, null (2021).

102. Alitto, H. J. & Dan, Y. Cell-type-specific modulation of neocortical activity by basal forebrain input. *Front. Syst. Neurosci.* **6**, (2013).

103. Prönneke, A., Witte, M., Möck, M. & Staiger, J. F. Neuromodulation Leads to a Burst-Tonic Switch in a Subset of VIP Neurons in Mouse Primary Somatosensory (Barrel) Cortex. *Cereb. Cortex* 488–504 (2020) doi:10.1093/cercor/bhz102.

104. Kamigaki, T. Dissecting executive control circuits with neuron types. *Neurosci. Res.* **141**, 13–22 (2019).

105. Zhang, S. *et al.* Long-range and local circuits for top-down modulation of visual cortex processing. *Science* **345**, 660–665 (2014).

106. Gnadt, J. W. & Andersen, R. A. Memory related motor planning activity in posterior parietal cortex of macaque. *Exp. Brain Res.* **70**, 216–220 (1988).

107. Ratcliff, R. & McKoon, G. The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Comput.* **20**, 873–922 (2008).

108. Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R. & Brown, S. D. Revisiting the Evidence for Collapsing Boundaries and Urgency Signals in Perceptual Decision-Making. *J. Neurosci.* **35**, 2476–2484 (2015).

109. Heathcote, A., Brown, S. & Mewhort, D. J. K. Quantile maximum likelihood estimation of response time distributions. *Psychon. Bull. Rev.* **9**, 394–401 (2002).

110. Ratcliff, R. & Tuerlinckx, F. Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychon. Bull. Rev.* **9**, 438–481 (2002).

111. Kiani, R. & Shadlen, M. N. Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. *Science* **324**, 759–764 (2009).

112. Funahashi, S., Bruce, C. J. & Goldman-Rakic, P. S. Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J. Neurophysiol.* **61**, 331–349 (1989).

113. Fuster, J. M. & Alexander, G. E. Neuron Activity Related to Short-Term Memory. *Science* **173**, 652–654 (1971).

114. Goldman-Rakic, P. S. Cellular basis of working memory. *Neuron* **14**, 477–485 (1995).

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

116. Ganguli, S., Huh, D. & Sompolinsky, H. Memory traces in dynamical systems. *Proc. Natl. Acad. Sci.* **105**, 18970–18975 (2008).

117. Seung, H. S. How the brain keeps the eyes still. *Proc. Natl. Acad. Sci.* **93**, 13339–13344 (1996).

118. Burak, Y. & Fiete, I. R. Accurate Path Integration in Continuous Attractor Network Models of Grid Cells. *PLOS Comput. Biol.* **5**, e1000291 (2009).

119. Wimmer, K., Nykamp, D. Q., Constantinidis, C. & Compte, A. Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nat. Neurosci.* **17**, 431–439 (2014).

120. Zhang, B. *et al.* Transforming absolute value to categorical choice in primate superior colliculus during value-based decision making. *Nat. Commun.* **12**, 3410 (2021).

121. Cowan, N. The Magical Mystery Four: How Is Working Memory Capacity Limited, and Why? *Curr. Dir. Psychol. Sci.* **19**, 51–57 (2010).

122. Cowan, N. *Working Memory Capacity: Classic Edition*. (Psychology Press, 2016).

123. Engle, R. W. What is working memory capacity? in *The nature of remembering: Essays in honor of Robert G. Crowder.* (eds. Roediger, H. L., Nairne, J. S., Neath, I. & Surprenant, A. M.) 297–314 (American Psychological Association, 2001). doi:10.1037/10394-016.

124. Engle, R. W. Working Memory Capacity as Executive Attention. *Curr. Dir. Psychol. Sci.* **11**, 19–23 (2002).

125. Oberauer, K., Farrell, S., Jarrold, C. & Lewandowsky, S. What limits working memory capacity? *Psychol. Bull.* **142**, 758–799 (2016).

126. Watanabe, M. Reward expectancy in primate prefrontal neurons. **382**, 4 (1996).

127. Shadlen, M. N. & Newsome, W. T. Motion perception: seeing and deciding. *Proc. Natl. Acad. Sci.* **93**, 628–633 (1996).

128. Bathellier, B., Ushakova, L. & Rumpel, S. Discrete Neocortical Dynamics Predict Behavioral Categorization of Sounds. *Neuron* **76**, 435–449 (2012).

129. Niessing, J. & Friedrich, R. W. Olfactory pattern classification by discrete neuronal network states. *Nature* **465**, 47–52 (2010).

130. Wills, T. J., Lever, C., Cacucci, F., Burgess, N. & O'Keefe, J. Attractor Dynamics in the Hippocampal Representation of the Local Environment. *Science* **308**, 873–876 (2005).

131. Kopec, C. D., Erlich, J. C., Brunton, B. W., Deisseroth, K. & Brody, C. D. Cortical and Subcortical Contributions to Short-Term Memory for Orienting Movements. *Neuron* **88**, 367–377 (2015).

132. Yoon, K. *et al.* Specific evidence of low-dimensional continuous attractor dynamics in grid cells. *Nat. Neurosci.* **16**, 1077–1084 (2013).

133. Hopfield, J. J. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci.* **79**, 2554–2558 (1982).

134. Amit, D. J. & Brunel, N. Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. *Cereb. Cortex* **7**, 237–252 (1997).

135. Amari, S. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol. Cybern.* **27**, 77–87 (1977).

136. Ratcliff, R. & McKoon, G. The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Comput.* **20**, 873–922 (2007).

137. Acerbi, L. & Ma, W. J. Practical Bayesian Optimization for Model Fitting with Bayesian Adaptive Direct Search. *ArXiv170504405 Q-Bio Stat* (2017).