

Computational Psychiatry: A Primer

Ed: Peggy Seriès

46 Figures

4 Tables

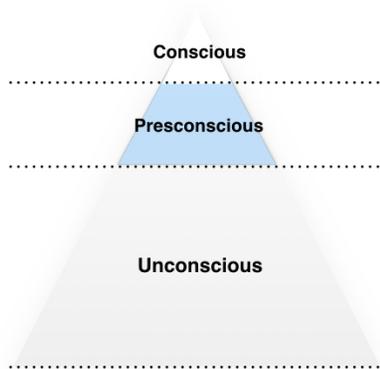


Figure 1.1: The three levels of consciousness. The conscious mind includes all thoughts, feelings and actions of which we are aware. The preconscious mind includes all mental activities that are not presently active, but stored and accessible when required. The unconscious mind includes mental activity of which we are unaware. According to Freud, some of the feelings, thoughts, urges and emotions that are actively buried into unconscious mind influence some of our unexplained behavior.

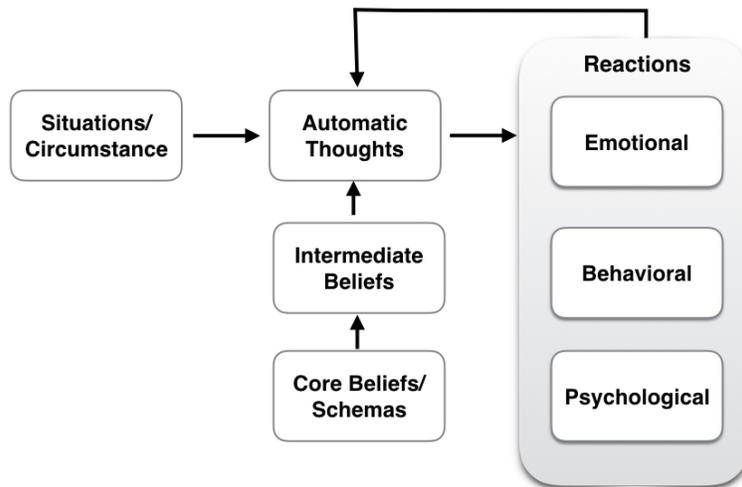


Figure 1.2: Beck's Cognitive model. The cognitive model developed by A. Beck explains individuals' emotional, physiological, and behavioral responses to circumstances and situations as mediated by their automatic thoughts (see e.g., Beck 1991). Automatic thoughts are influenced by underlying beliefs developed over time and through experience. Individuals' perceptions are often distorted and dysfunctional when they are distressed, leading to automatic negative thoughts. Through Cognitive Therapy, individuals can learn to identify, evaluate and correct their automatic negative thinking. When they do so, distress usually decreases and psychological function increases.

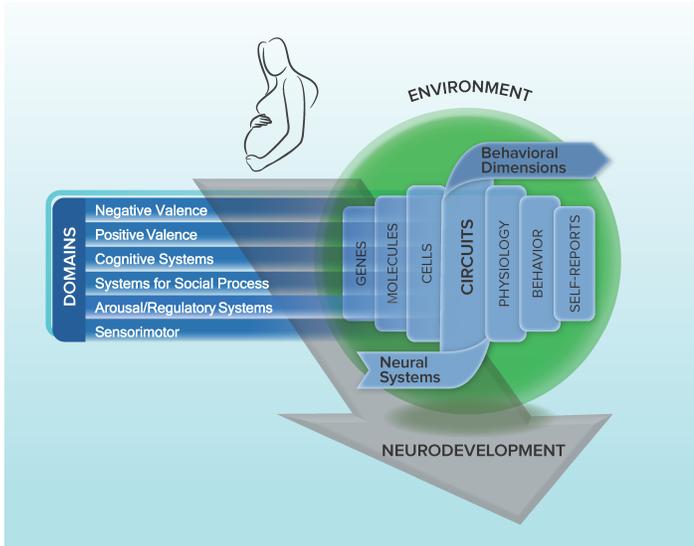


Figure 1.3: The RDoC matrix. RDoC is a research framework proposed by NIMH for new approaches to investigating mental disorders. It integrates many levels of information (from genomics and circuits to behavior and self-reports) in order to explore basic dimensions of functioning that span the full range of human behavior from normal to abnormal. RDoC is not meant to serve as a diagnostic guide, nor is it intended to replace current diagnostic systems. The goal is to understand the nature of mental health and illness in terms of varying degrees of dysfunctions in general psychological/biological systems.

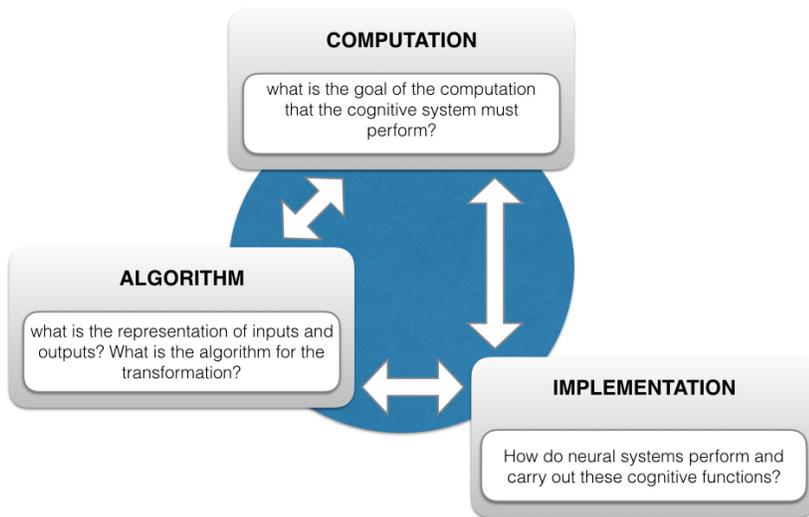


Figure 1.4: Marr's three levels of analyses. Originally introduced to advance the understanding of vision, Marr's approach postulated three distinct ways (computational, algorithmic, implementational) to consider information processing in the context of neuroscience. This distinction has since been used in research across other cognitive domains (*e.g.*, memory, attention, learning), though many have not strictly adopted the classically rigid framework of Marr's hierarchy. Instead the hierarchy is most commonly used as an organizing principle to highlight distinct conceptual questions at different levels of analysis. Image credit: Debbie Yee and Todd Braver.

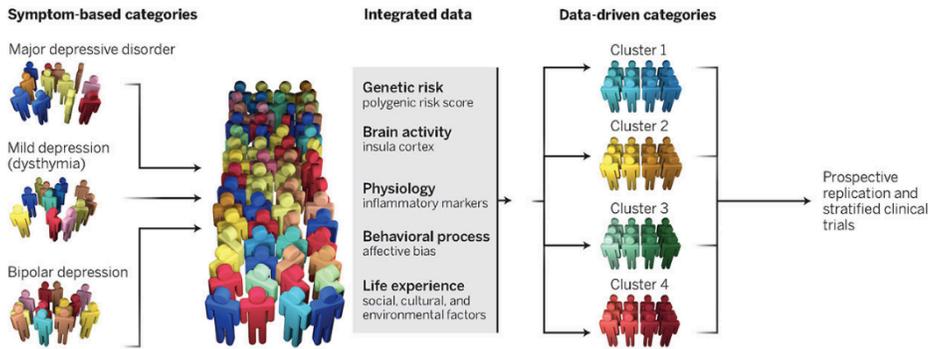


Figure 1.5: Data-driven computational psychiatry. A hypothetical example illustrates how a data-driven approach might lead to new descriptions and classifications, beyond traditional symptom-based categories of mental disorder. Consider a population of patients suffering from different types of mood disorders and for whom a variety of data has been collected (genetic, brain, physiological data etc.). New clusters might be found in the data that might connect more directly with mechanisms underlying their symptoms. Such clusters might form groups that are more homogeneous than the original classification, and possibly more relevant in terms of possible treatment. Reproduced from Insel and Cuthbert (2015) with permission.

	Sad	Ene	Con	Ins	Int	App	Bla	Wei	Agi	Ret	Sui	Hyp	Ref (%)	Profile Description
A													1.78	No symptoms
B	x	x	x	x	x	x	x	x	x	x			1.24	All but Sui and Hyp
C	x	x	x	x	x	x		x					1.19	Mixed profile
D	x	x	x	x	x	x	x	x					1.19	Mixed profile
E	x	x	x	x	x								1.13	Mixed profile
F	x	x	x	x	x		x						1.13	Mixed profile
G				x									1.08	Only Ins
H	x	x	x	x	x	x	x	x	x				1.00	All but Ret, Sui and Hyp
I	x	x	x	x									0.92	Mixed profile
J	x	x	x	x	x	x		x	x	x			0.89	All but Hyp, Bla and Sui

Table 1.1: Fried & Nesse (2014) examined DSM-5 depression symptom patterns in the “Sequenced Treatment Alternatives to Relieve Depression Study” (STAR*D). They found 1,030 unique symptom profiles in 3,703 patients all given the diagnosis of Major Depressive Disorder. The 10 most frequent symptom profiles are illustrated in this table. Cells with ‘x’ mark symptom presence. Abbreviations: Sad, sadness; Ene, energy loss; Con, concentration problems; Ins, insomnia; Int, interest loss; App, appetite problems; Bla, self-blame; Wei, weight problems; Agi, psychomotor agitation; Ret, psychomotor retardation; Sui, suicidal ideation; Hyp, hypersomnia; Freq, frequency of profiles. Reproduced from Fried and Nesse (2014).

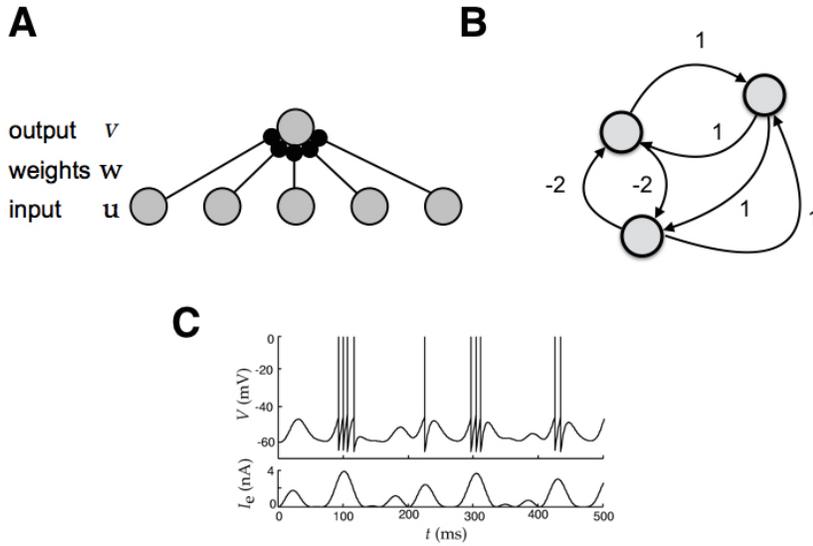


Figure 2.1: Example of neural network models. A. The perceptron is a feed-forward neural network. Here the output unit v receives inputs from all input units u , weighted by weight w . B. A Hopfield network is a network of binary units connected by recurrent connections. C. Integrate and fire neuron. This model accounts for the changes in voltage observed when a neuron received an input current I_e . When the voltage crosses a threshold, the voltage suddenly jumps to 0 and is reset. This models a spike, or action potential.

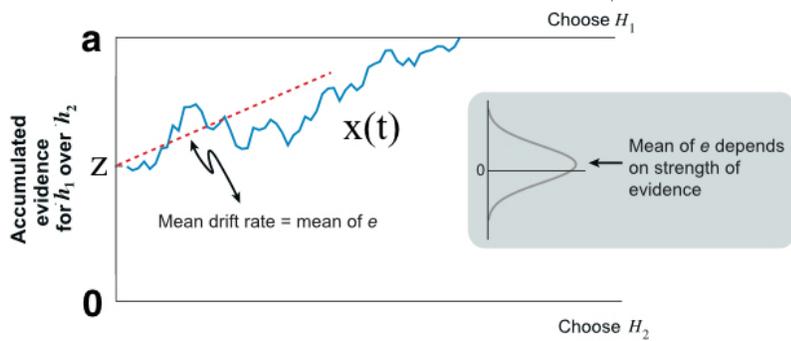


Figure 2.2: The Drift Diffusion Model (DDM). The decision variable is a noisy cumulative process (blue) composed of the evidence (e) with starting point in the middle of A and $-A$ (no starting bias). The evidence is sampled from a Gaussian distribution whose mean μ depends on the strength of the evidence. The bounds represent the stopping rule and their separation accounts for accuracy-speed trade-off. Reprinted from Gold and Shadlen (2007) with permission.

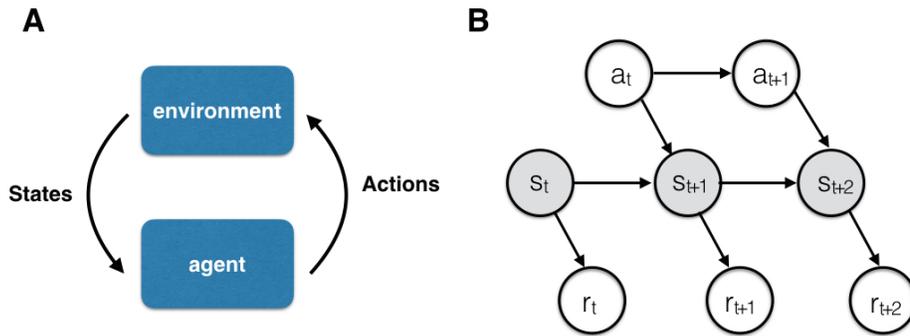


Figure 2.3: Markov Decision Problems. A) The setting: an agent interacts with an environment by choosing actions that in turn influence its current state. B) At each time step, the agent is in some state s and may choose any action available in that state a . This leads to him moving into a new state and giving the decision maker a corresponding reward r .

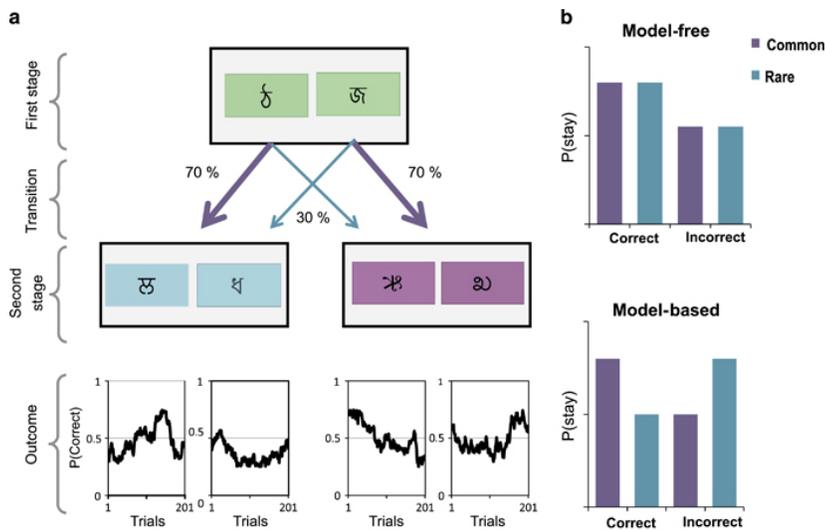


Figure 2.4: Two-steps decision-making task. Task: (a) On each trial, choosing between two stimuli leads with fixed probabilities (transition) to one of two pairs of stimuli in stage 2. Each of the four second-stage stimuli is associated with a probabilistic outcome (monetary reward). Those probabilities change slowly and independently across the trials. (b) Model-based and model-free strategies make different predictions about the influence of the outcome obtained after the second stage onto subsequent first-stage choices. They thus predict different choice patterns: in the model-free system, obtaining a reward increases the chance of choosing the same stimulus on the next trial independently of whether the type of transition was rare or common (upper row). In a model-based system, on the contrary, the choices of the stimuli on the next trial integrate the transition type (lower row). Reproduced from Worbe et al (2015) with permission.

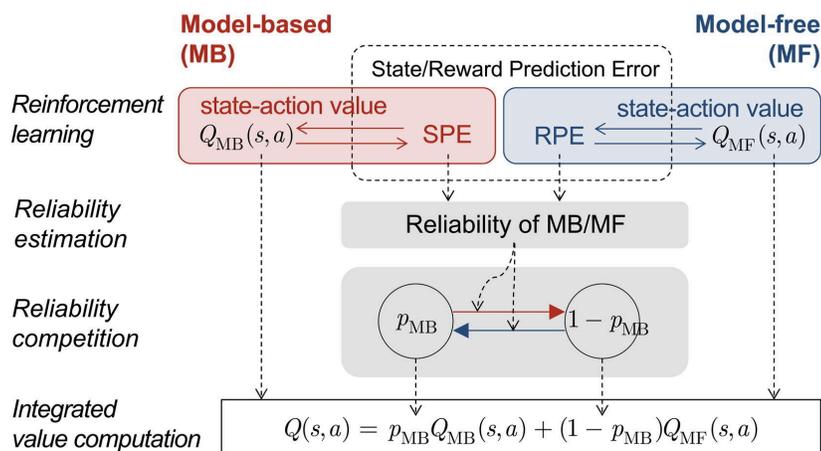


Figure 2.5: Example of a possible arbitration between Model-Based and Model-free Learning Strategies. Accumulating neural evidence support the existence of two distinct systems for guiding action selection, a deliberative “model-based” and a reflexive “model-free” system. However, little is known about how the brain determines which of these systems controls behavior at one moment in time. Lee et al (2014) propose an arbitration mechanism that allocates the degree of control over behavior by model-based and model-free systems as a function of the reliability of their respective predictions. Reliability is computed based on the state prediction error (SPE) in the model-based learning system and based on the reward prediction error (RPE) in the model-free learning system. The computed reliability functions as a transition rate for the two-state transition model, in which each state represents the probability of choosing the model-based learning strategy (p_{MB}) and the model-free ($1 - p_{MB}$), respectively. The state-action value regulating the actual choice behavior is given by the weighted average of values from the two reinforcement learning systems. Reproduced from Lee et al (2014) with permission.

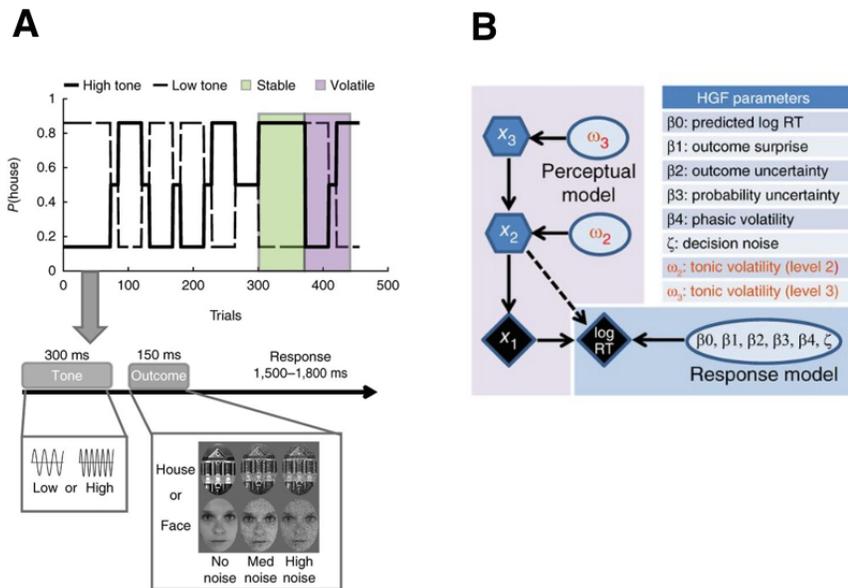


Figure 2.6: A. Example of the task used by Lawson, Mathys, and Rees (2017). Schematic of the task, showing the volatile environmental structure (top), for example, the probability of seeing a house (given the preceding high or low tone) across trials. The green area shows a 'stable' period of 72 trials in which the probabilities remained fixed, and the violet area shows a 'volatile' period of 72 trials in which the outcome probabilities switched three times. A single trial is also seen (bottom) showing example stimuli. **B. Hierarchical Gaussian Model (HGM) model.** Schematic depiction of the three-level HGF that was used to model this task. The perceptual model comprises three hierarchical states (x_1 , x_2 , and x_3). The lowest level variable x_1 describes the uncertainty about outcomes, i.e. the presence of a house or face, level 2 (x_2) addresses uncertainty about the cue-outcome contingencies, and level 3 (x_3) addresses uncertainty about environmental change, i.e. the volatility of the cue-outcome contingencies. Participant-specific free parameters (ovals) are estimated from individual reactions times (log RT) data.

Red parameters relate to the perceptual model, whereas black parameters relate to the response model. Diamonds, quantities that change over time (trials); hexagons quantities that change over time and that additionally depend on their previous state in time in a Markovian fashion. Reproduced from Lawson, Mathys, and Rees (2017) with permission.

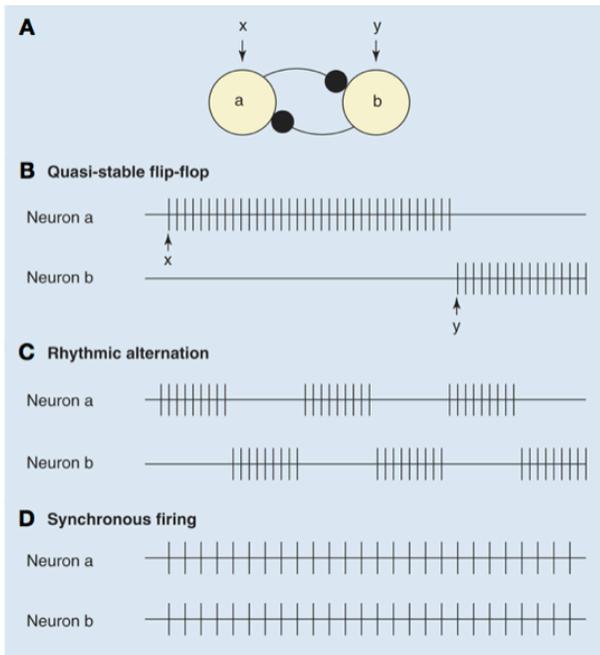


Figure 3.1: Distinct dynamical regimes in a circuit of two reciprocally connected inhibitory neurons. **(A)** Neurons a and b receive excitatory inputs x and y , respectively. **(B-D)** Under different conditions of neuronal and synaptic properties, the circuit can exhibit qualitatively distinct dynamical regimes, including quasi-stable flip-flop **(B)**, rhythmic alternation as a “half-center” oscillator **(C)**, and spike-by-spike synchrony. Adapted from Kristan and Katz (2006).

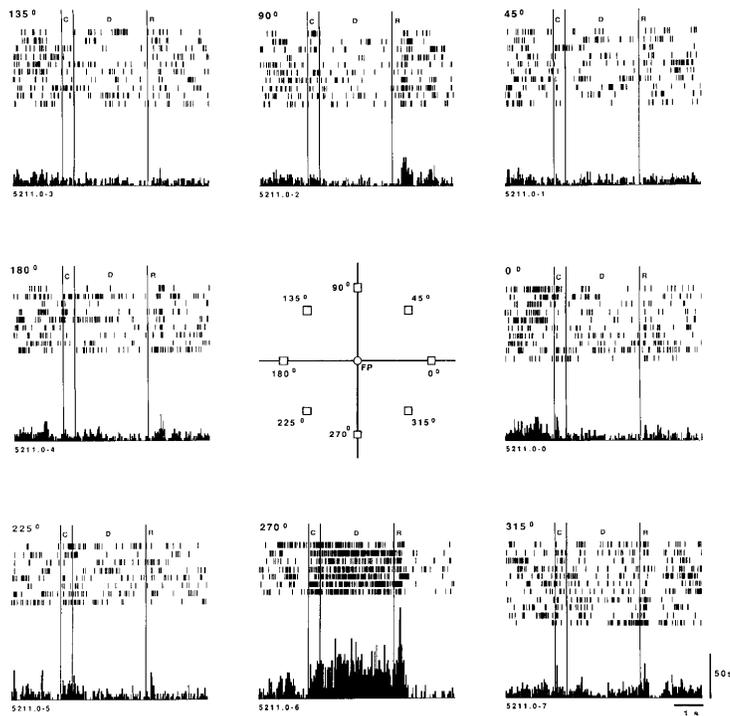


Figure 3.2. Example of directional delay period activity of a principle sulcus neuron during the oculomotor delayed-response task. In this seminal experiment described by Funahashi et al (1989), monkeys were trained to fixate a central spot during a brief presentation of a peripheral cue and throughout a subsequent delay period (3 sec) and then, upon the extinction of the fixation target, to make a saccadic eye movement to where the cue had been presented. Visual cues were randomly presented at one of the eight locations indicated in the center diagram. The neuron shown in this example had strongly directional delay period activity responding only when the cue had been presented at the bottom location. It was suppressed during the delay when the cue was presented in the upper visual field. Reproduced from Funahashi et al (1989) with permission.

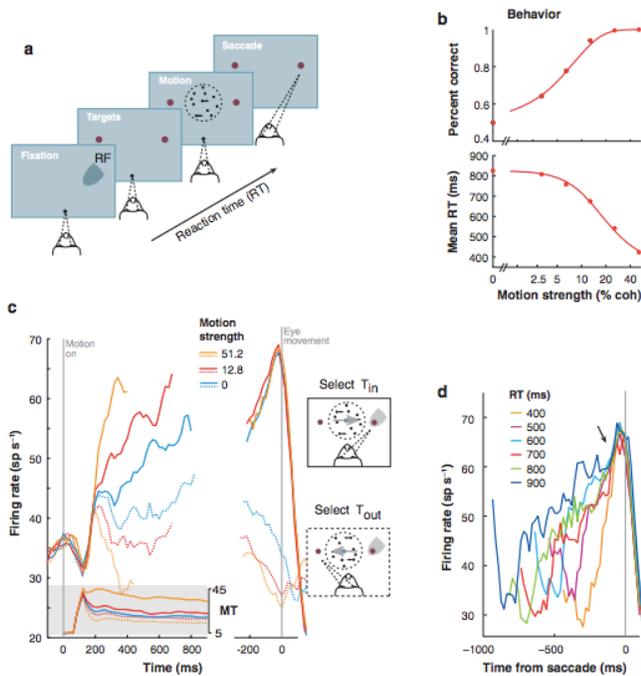


Figure 3.3: Neural mechanism of a decision about direction of motion. (a) The subject views a patch of dynamic random dots and is requested to indicate, whenever they are ready, which net direction they perceived for the motion (left or right). They need to indicate their decision by making an eye movement to a peripheral target. The gray patch shows the location of the response field (RF) of a LIP neuron. One of the choice targets (T_{in}) is in the response field (RF) of the LIP neuron; the other target (T_{out}), as well as the motion stimulus itself, lie outside the neuron's RF. (b) Effect of stimulus difficulty on accuracy and decision time. (c) Response of LIP neurons during decision formation. Average firing rate from 54 LIP neurons is shown for three levels of difficulty. Responses are grouped by motion strength and direction of choice, as indicated. **Left:** The responses are aligned to the onset of the random-dot motion. Shaded insert shows average responses from direction selective neurons in area MT to motion in the preferred and anti-preferred directions. After a transient, MT responds at a nearly constant rate. **Right:** The responses are aligned to the eye movement. The LIP firing rates ramp up or down, approximating the integral of a

difference in firing rate between MT neurons with opposite direction preferences. **(d)** Responses grouped by reaction time. Only T_{in} choices are shown. All trials reach a stereotyped firing rate ~ 70 ms before saccade initiation (arrow). Adapted with permission from Gold and Shadlen (2007).

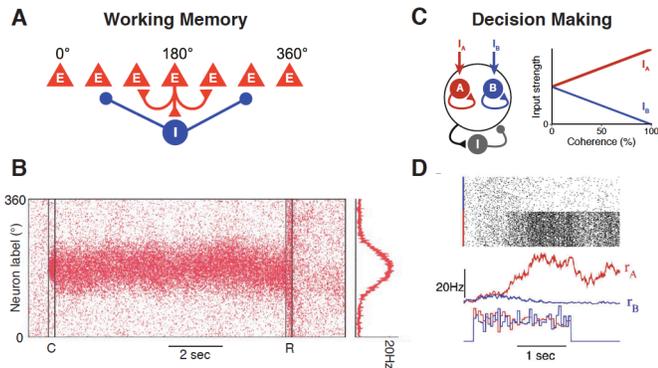
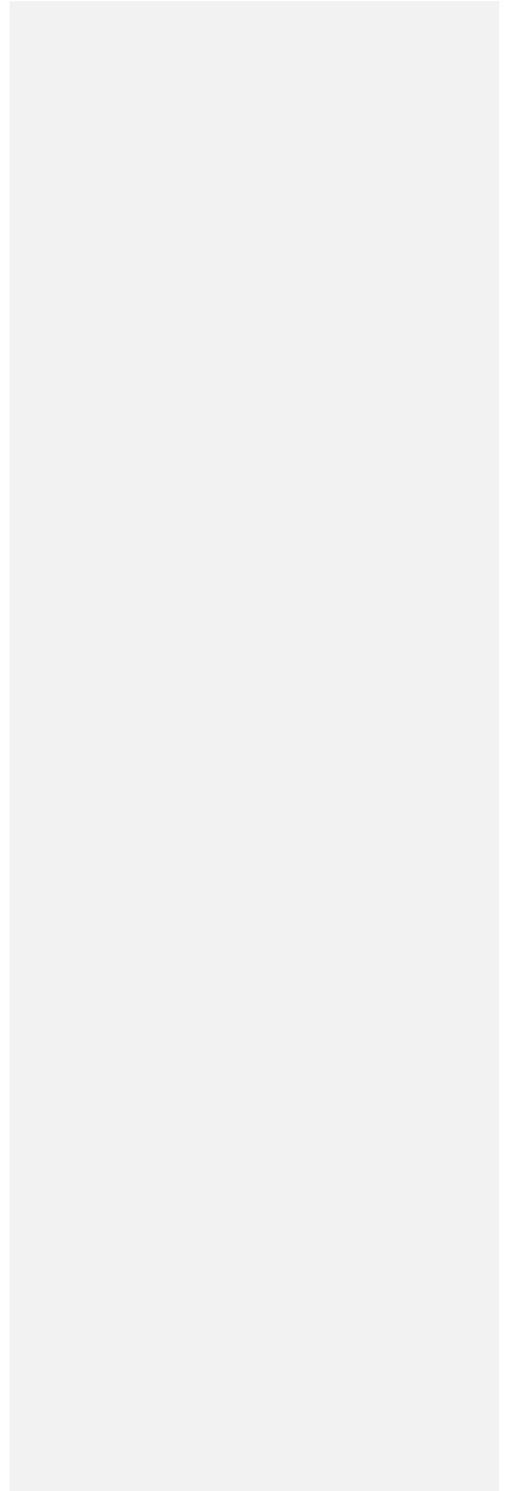


Figure 3.4: Biophysically based cortical circuit models of working memory and decision-making computations. (A) Schematic of the network architecture for a model of spatial working memory. The model consists of recurrently connected excitatory pyramidal cells (E) and inhibitory interneurons (I). Pyramidal cells are labeled by the angular location they encode (0–360°). Excitatory-to-excitatory connections are structured, such that neurons with similar preferred angles are more strongly connected. Connections between pyramidal cells and interneurons are unstructured and mediate feedback inhibition. (B) Spatiotemporal raster plot showing a bump attractor state in an example trial. A stimulus is presented at 180 deg during the brief cue epoch (denoted C) and during the subsequent delay, the stimulus location is encoded by persistent activity throughout the working memory delay until the response epoch (denoted R). On right is shown the firing rate profile of the working memory bump attractor state. (C) Schematic of the network architecture for a model of perceptual decision-making. The circuit contains two populations of pyramidal neurons which are each selective to one of the two stimuli (A and B). Within each pyramidal-neuron population there is strong recurrent excitation, and the two populations compete via feedback inhibition mediated by interneurons (I). Right: The selective populations receive sensory-related inputs determined by the stimulus coherence. (D) Example neuronal activity in a single trial for a zero-coherence stimulus. Top: Spatiotemporal raster plot for the two selective population. Middle: Population firing rates r_A and r_B . Bottom: Stochastic sensory-related inputs. During decision-making, the circuit exhibits an initial slow ramping, related to temporal integration of evidence, which leads to

categorical choice (for A in this trial). Panels (B) and (D) adapted from Compte et al (2000) and Wang (2002), respectively.



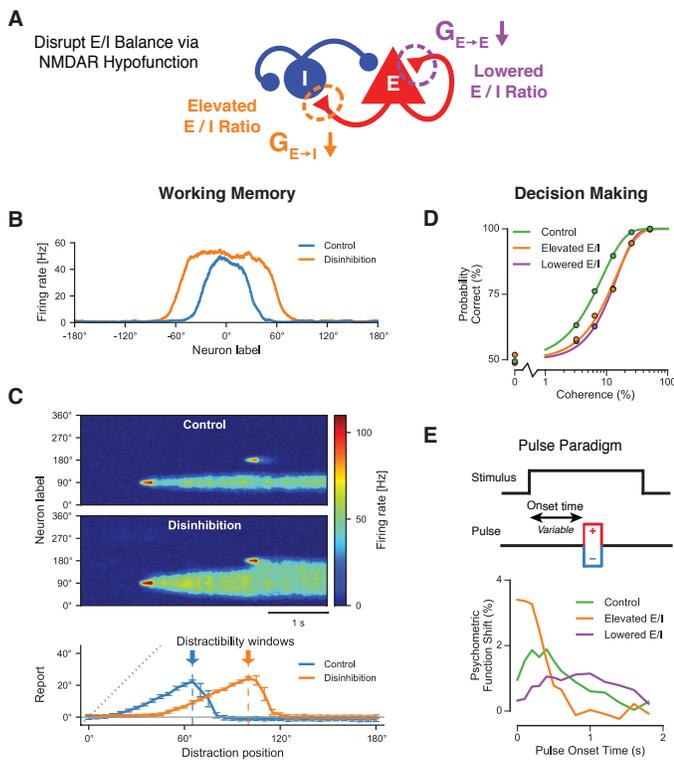


Figure 3.5: Effects of altered excitation-inhibition (E/I) balance in cortical circuit models of working memory and decision-making. (A) E/I ratio was perturbed bi-directionally via hypofunction of NMDA receptors at two recurrent synaptic sites: on inhibitory interneurons, which elevates E/I ratio via disinhibition; or on excitatory pyramidal neurons, which lowers E/I ratio. (B) For the working memory circuit, the firing rate profile of the “bump” attractor activity pattern during working memory maintenance. Elevated E/I ratio via disinhibition results in a broadened working memory representation. (C) Disinhibition impairs the network’s ability to filter out intervening distractors. Top: Spatiotemporal plot of network activity in response to a distractor presented during the delay at a distance of 90° from the target. Bottom: Deviation of the read-out report as a function of the angular distance between the distractor and the target. The “distractibility window” is widened by disinhibition. (D) In the decision-making

circuit, performance as quantified by the psychometric function, i.e., the proportion of correct choices as a function of stimulus coherence. Both perturbations, elevated and lowered E/I ratio, can comparably degrade performance relative to the control circuit. **(E)** A perceptual decision-making task paradigm that characterizes the time course of evidence accumulation can test dissociable behavioral predictions from elevated vs. lowered E/I ratio. Top: The pulse paradigm uses a brief pulse of additional perceptual evidence at different onset times. This pulse induces a shift the psychometric function, which quantifies the sensitivity of the choice on evidence presented at that time point. Bottom: Shift in the psychometric function as a function of pulse onset time, for the three E/I regimes. Relative to control, in the elevated-E/I circuit the pulse has a stronger impact at early onset times, but less impact at later onset times. The lowered-E/I circuit shows a flattened profile of the shift, with greater impact at late onset times.

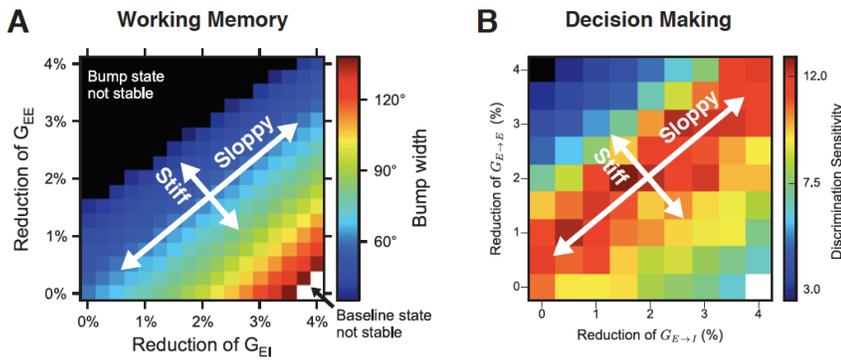


Figure 3.6: Dependence of circuit function on synaptic parameters: a critical role of excitation-inhibition (E/I) balance in both working memory and decision making. The plots illustrate a parameter space of reductions of two recurrent NMDAR conductance strengths from excitatory pyramidal neurons: onto inhibitory interneurons (G_{EI}) or onto excitatory pyramidal neurons (G_{EE}). This analysis characterizes the sensitivity of model function to joint perturbations of these two parameters. **(A)** For the working memory circuit, we measured the width of the working memory bump attractor state. Bump width affects mnemonic precision and distractibility during working memory maintenance. **(B)** For the decision-making circuit, we measured the discrimination sensitivity, which is defined as the inverse of the discrimination threshold (i.e., coherence which yields 81.6% correct). A higher sensitivity corresponds to better performance. For both working memory and decision-making circuits, within this range of perturbation, if G_{EI} and G_{EE} are reduced together in a certain proportion, circuit performance is essentially unaltered, because E/I balance is maintained. E/I balance defines a “sloppy” axis in parameter space along which the function is insensitive. In contrast, the function is highly sensitive to small orthogonal perturbations, along a “stiff” axis (Gutenkunst et al. 2007). Reduction of G_{EI} in greater proportion elevates E/I ratio and can degrade performance: for working memory, due to broadened mnemonic representations; for decision making, due to highly unstable integration leading to impulsive selection. In contrast, reduction of G_{EE} in

greater proportion lowers E/I ratio and can degrade performance: for working memory, due to loss of the bump attractor state; for decision making, due to indecisive selection. These findings indicate that E/I ratio is a crucial effective parameter for cognitive function in these circuits, with an “inverted-U” dependence of function on E/I ratio. Panels (A) and (B) adapted from Murray et al. (2014) and Lam et al. (2017), respectively, with permission.

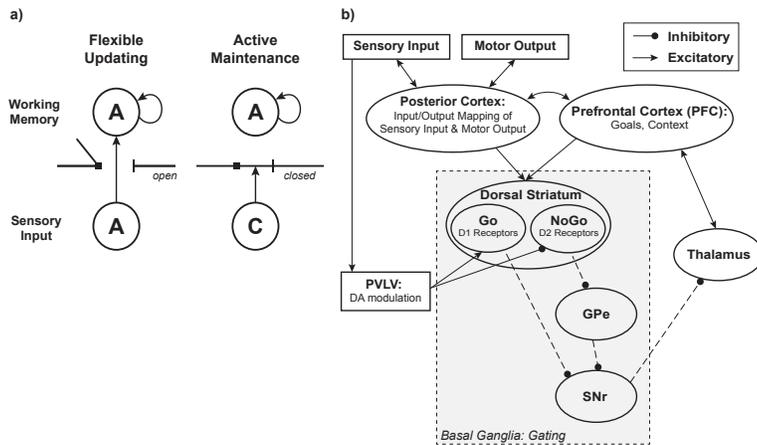


Figure 4.1: A) Gating Mechanism from Frank and O'Reilly's Prefrontal Basal Ganglia and Working Memory (PBWM) model (Frank, Loughry, and O'Reilly 2001). At the algorithmic level, this connectionist computational model features a gating function, which switches between active maintenance and flexible updating of working memory to incorporate task-relevant information, two core functions of cognitive control. **B) Neural Network Model Implementation of the PBWM.** Here, sensory inputs are mapped onto motor outputs via posterior ("hidden") layers. The prefrontal cortex (PFC) contextualizes this information and encodes relevant prior information and goals. The basal ganglia (BG) updates the PFC via dynamic gating, which is driven by dopaminergic (DA) modulation from a separate "PVLV" (primary value and learned value learning algorithm) system (O'Reilly et al. 2007). Specifically, DA is excitatory onto the Go neurons via D1 receptors and inhibitory onto NoGo neurons via D2 receptors. Thus, increased DA firing will inhibit SNr (substantia nigra pars reticulata) and disinhibit PFC to facilitate flexible updating of working memory representations in PFC. Decreased DA firing, on the other hand, counteracts this effect and facilitates active maintenance of current working memory representations in PFC.

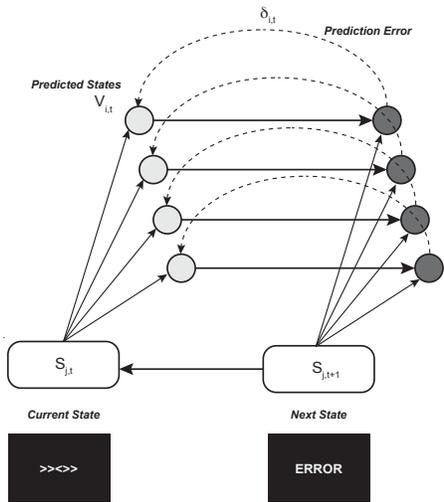


Figure 4.2: Schematic of the Prediction-Response Outcome (PRO) model by Alexander and Brown (2011, 2014). First, the PRO model learns predictions of multiple possible future outcomes of various chosen actions (indicated by $V_{i,t}$), using an error likelihood signal. Thus, activity in the PRO model reflects a temporally discounted prediction of such outcomes, which are proportionate to their likelihood of occurrence. Second, the PRO detects discrepancies between predicted and observed outcomes, and uses their prediction error signal (δ) to update and improve subsequent predictions. S refers to the representation of the stimulus (e.g., conflicting arrows from the Erikson flanker task) or task-related feedback (e.g., a screen indicating an error was made). Thus, the PRO model continually learns and updates associations between task-related cues and feedback in cognitive tasks.

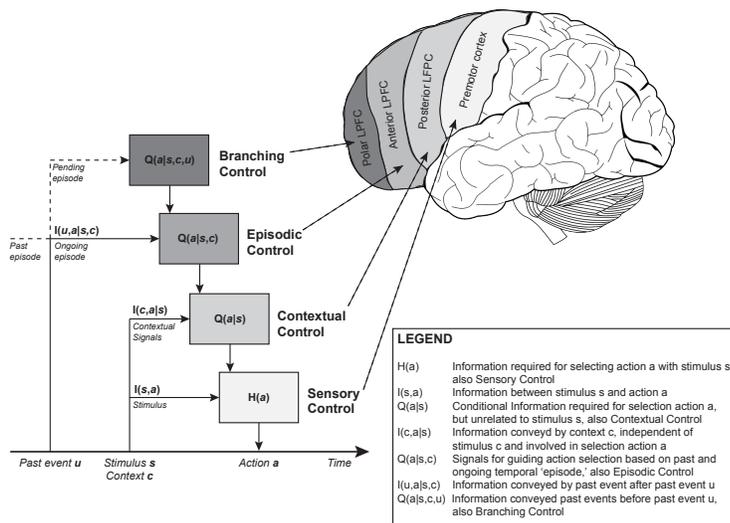


Figure 4.3: Model of Hierarchical Cognitive Control by Koehlin and colleagues (2003, 2007). This information-theoretic model posits that cognitive control operates according to three nested levels of control processes (branching, episodic, contextual), which are implemented as a cascade from anterior to posterior prefrontal regions. H(a) represents sensory control, the information required to select an action (a) to appropriate incoming stimuli, and is the sum of two control terms: bottom-up information conveyed by the stimulus (s) regarding the appropriate action [I(s,a)] and top-down information processed in the posterior lateral PFC [Q(a|s)]. The Q(a|s) term represents contextual control, the incoming signals congruent with the subject’s response, and is the sum of two control terms: bottom up information from the contextual (c) signals and stimulus [I(c,a|s), I(s,a)] and top-down information processed in anterior lateral PFC [Q(a|s,c)]. The Q(a|s,c) term represents episodic control, neural signals that guide actions based on information retrieved from past events stored in episodic memory (i.e., tonically maintained over a longer temporal interval), which is the sum of bottom-up information from past event u [I(u,a|s,c)] and top-down information processed in the polar lateral PFC [Q(a|s,c,u)]. The branching control term [Q(a|s,c,u)] relates to information conveyed by events prior to event u, and are maintained until the current episode or trial is complete. Thus, this computational model parses different levels of control based on

how much information must be internally represented and actively maintained in order to select and perform a correct action.

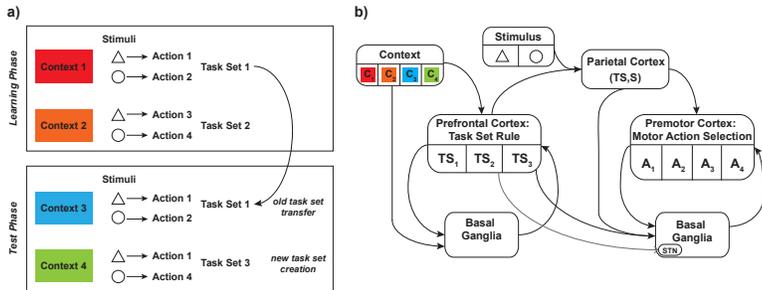


Figure 4.4: A) Context-Task Set (C-TS) Model by Collins and Frank (2013). This model solves the problem of how to learn hidden task-set rules (i.e., when in a given state and presented with sensory input, which action should be taken in order to maximize reward). The C-TS model posits that states are determined hierarchically; that is, an agent will consider some input dimension to act as a higher order context (C), which indicates a task-set (TS) and other dimensions to act as lower level stimuli (S), in determining which motor actions (A) to produce. Here, the color context determines a latent task-set that facilitates learning of shape stimulus-action associations in the learning phase (e.g., C1 is associated with TS1). In the test phase, C3 maps onto the same shape stimulus-action association as C1, so the C3 context is transferred to TS1, whereas C4 should be assigned to a new task set. Critically, the model predicts that it should be faster to transfer a task-set than learning a new task-set. **B) Schematic of two-loop corticostriatal gating neural network model.** These two loops are nested hierarchically, such that one loop learns to gate an abstract task-set (and will group together the contexts that are associated with the same task-sets), whereas the other loop learns to gate a motor action response conditioned on the task-set and perceptual stimulus. Here, color context (C) serves as the input for learning to select the correct task-set (TS) in the prefrontal cortex loop. This information is multiplexed with the shape stimulus in parietal cortex to modulate the motor loop and select the correct motor actions. Critically, these two loops accomplish two objectives: 1) constrain motor actions until a task-set is selected, and 2) allow conflict at the level of task-set selection to delay responding in the motor loop, preventing premature action selection

until a valid task-set is selected. Taken together, both algorithmic and neural network models similarly and accurately predict behavioral task performance. The synergism of different modeling levels provides an account of how humans engage cognitive control and learning to produce structured abstract representations that enable generalization in the long-term, even if it may be costly in the short-term.

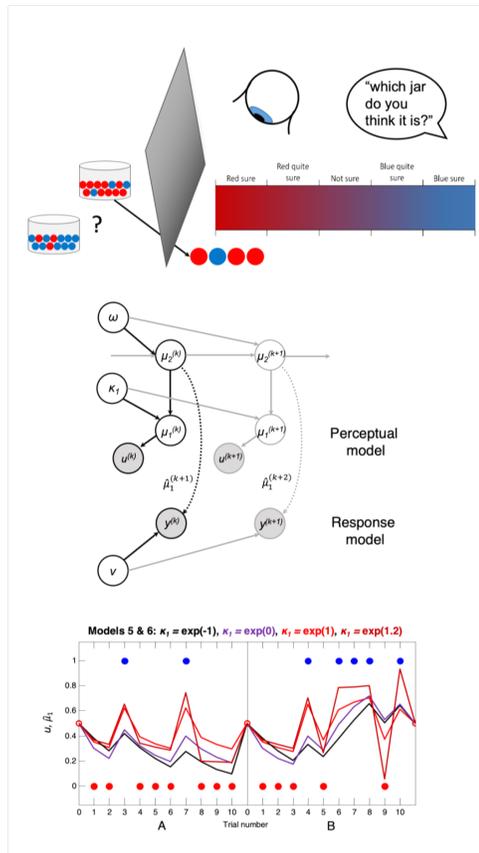


Figure 6.1: The ‘probability estimates’ version of the beads task and the winning model

Upper panel: This schematic illustrates the concept behind the beads task. A subject is shown two jars, each filled with opposite proportions of red and blue beads (e.g. 80:20 and 20:80 ratios), and the jars are then concealed from view. A sequence of beads is drawn, and the subject is asked to rate the probability that the beads are coming from one jar or another.

Middle panel: A schematic representation of the generative model in Model 5 and Model 6, the winning model. The black arrows denote the probabilistic network on trial k ; the grey arrows denote the network at other points in time. The perceptual model lies above the dotted arrows, and the response model below them. The shaded circles are known quantities, and the parameters and states in unshaded circles are estimated. The dotted line represents the result of an inferential process (the response model builds on a perceptual model inference); the solid lines are generative processes. μ_2 denotes the estimated tendency towards the blue or red jar and ω is a static source of variance at this level (greater variance means belief updates are larger). The bead seen by the subject, $u^{(k)}$, is generated by the estimated jar on trial k , $\mu_1^{(k)}$. The response model maps from $\mu_1^{(k+1)}$ – the predicted jar on the next trial, a sigmoid function s of $\mu_2^{(k)}$ – to $y^{(k)}$, the subject's indicated estimate of the probability the jar is blue. Variation in this mapping is modelled as the precision β of a beta distribution.

Lower panel: This figure illustrates the effects of κ_I (used in Models 5 and 6) on inference. It shows simulated perceptual model predictions; the second level μ_2 and simulated responses y have been omitted for clarity. The simulations use four different values of belief instability κ_I , which alters the sigmoid transformation: $\mu_1^{(k+1)} = s(\kappa_I \cdot \mu_2^{(k)})$. When $\kappa_I > \exp(0)$ updating is greater to unexpected evidence and lower to consistent evidence; when $\kappa_I < \exp(0)$ the reverse is true. The red and brown lines ($\kappa_I > \exp(0)$) illustrate the effects of increasingly unstable attractor networks, i.e. switching between states (jars) becomes more likely (see also Figure 2, upper panel). The black line ($\kappa_I = \exp(-1)$) illustrates slower updating around $\mu_1 = 0.5$, as was found in controls.

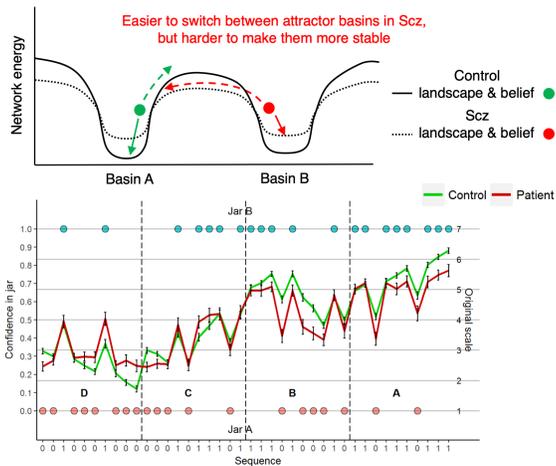


Figure 6.2: Attractors and belief updating, and dataset 2 group-averaged beliefs

Upper panel: This schematic illustrates the energy of a network with two attractors or fixed points – e.g. beliefs about blue and red jars – over a range of firing rates. The continuous line depicts a normal network whose ‘basins’ of attraction are relatively deep. The dotted line depicts the effect of NMDAR (or cortical dopamine 1 receptor) hypofunction on the attractor dynamics. The dots depict initial belief states, and the full and dashed arrows depict the effects of confirmatory and disconfirmatory evidence, respectively. A shallower basin of attraction means its corresponding belief state is harder to stabilize but easier to change. See also similar schematics elsewhere (Durstewitz and Seamans 2008; Rolls et al. 2008).

Lower panel: This panel shows the mean (\pm standard error) confidence ratings in the blue jar averaged across each group in dataset 2. These consist of four 10 bead sequences concatenated together (they were presented to the subjects separately during testing). The group with schizophrenia makes larger updates to unexpected beads but smaller updates to more consistent evidence at the end of sequences A and D.

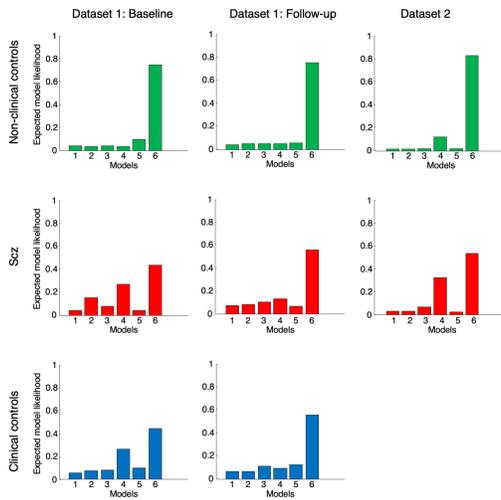


Figure 6.3: Expected model likelihoods for each group in each dataset.

This figure depicts the model likelihoods for the six models in each group in each dataset. The model likelihood is the probability of that model being the best for any randomly selected subject (Stephan et al. 2009). Model 6 wins in all groups in both datasets.

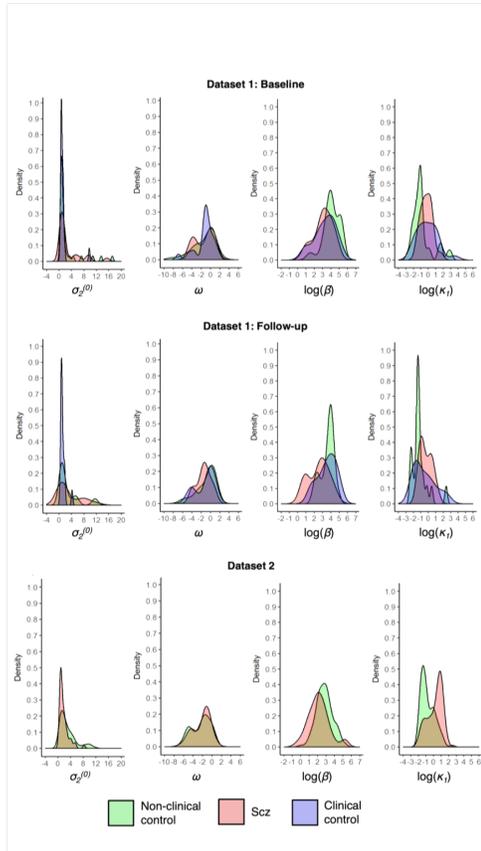


Figure 6.4: Probability density plots for the four Model 6 parameters in each dataset.

The upper and middle rows show the parameter estimates for dataset 1, at baseline ($n=81$) and follow-up ($n=53$) respectively, and the bottom row dataset 2 ($n=167$). There were only significant group differences in attractor instability κ_1 and β in each dataset (see text).

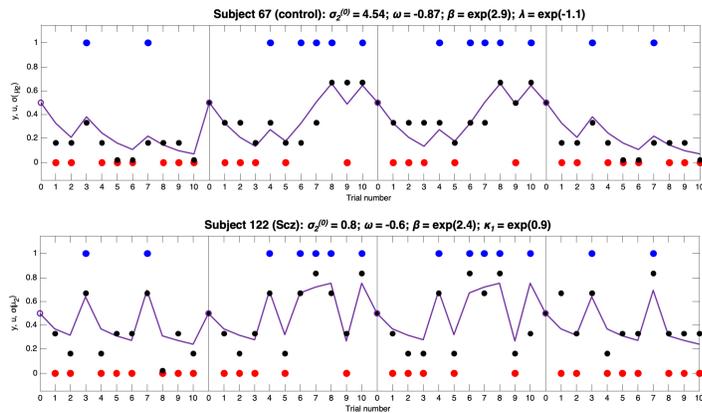


Figure 6.5: Model fits for two example subjects

These plots illustrate the beads seen by subjects (u , in blue and red), the ratings of the probability the jar is blue made by the subjects (y , in black) and the model fit line ($\sigma(\mu_2)$, in purple), for four 10-trial beads tasks concatenated together. The bead colors in two sequences have been swapped around for model estimation purposes. The upper plot is a control subject, the lower plot a subject with schizophrenia. The two subjects illustrate the effects of attractor instability κ_I – their values of κ_I (approximately $\exp(-1)$ and $\exp(1)$ respectively) are the modal values of κ_I in each group (see Figure 4, bottom row). The subject with schizophrenia makes much larger adjustments to changes in evidence (crossing the $p(\text{jar}=\text{blue})=0.5$ line repeatedly), explained by a higher κ_I . Note also that this patient’s responses are also more stochastic, explained by a lower β . These parameter differences may both be the result of attractor network instability in prefrontal cortex.

Model	Free parameters (prior mean, var). NB all models also contained ω (-2, 16) and β (exp(4.85), 1)	Description
1		Learning rate & response stochasticity only
2	$\sigma_2^{(0)}$ (exp(-5.1), 0.5)	Initial variance estimated
3	φ (0.1, 2)	Disconfirmatory bias
4	$\sigma_2^{(0)}$ (exp(-5.1), 0.5), φ (0.1, 2)	“ “ & initial var estimated
5	κ_l (0,1)	‘Attractor instability’
6	$\sigma_2^{(0)}$ (exp(-5.1), 0.5), κ_l (0,1)	“ “ & initial var estimated

Table 6.1: The six models tested on two beads task datasets. Each model contained a learning rate ω and response precision (similar to inverse temperature) β , along with the additional parameters listed. The brackets contain the prior mean and variance for each parameter used during model fitting.

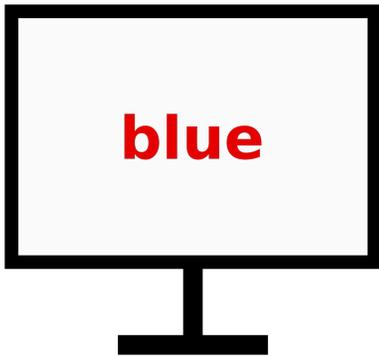


Figure 7.1: A sketch of the Stroop color-naming task, as used by Siegle, Steinhauer, and Thase (2004). Participants had to respond by indicating the color of the ink of the word (here *red*), while ignoring the written word (here *blue*).

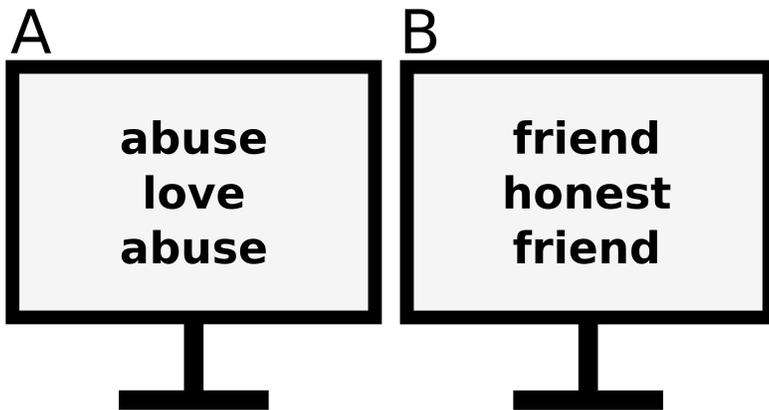


Figure 7.2: A sketch of two types of trials of the emotional flanker task, as used by Pe, Vandekerckhove, and Kuppens (2013). Participants had to classify the word in the center according to its valence. (A) An incongruent trial, in which the target word *love* and the flanking word *abuse* have differing valence. (B) A congruent trial, in which the valence of the flanking word is the same as the valence of the target. (Note that Dutch, four letters long, monosyllabic words were used by Pe, Vandekerckhove, and Kuppens (2013)).

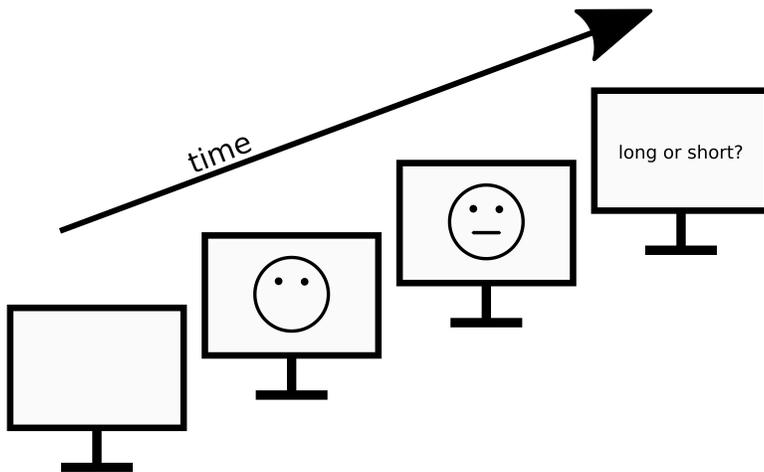


Figure 7.3: A sketch of the signal detection task (Huys et al. 2013). On each trial participants observe one of two possible cartoon faces which only differ slightly in the lengths of their mouths. They have to indicate which face they observed. The reward structure is asymmetrical with one of the stimuli being rewarded more frequently than the alternative.

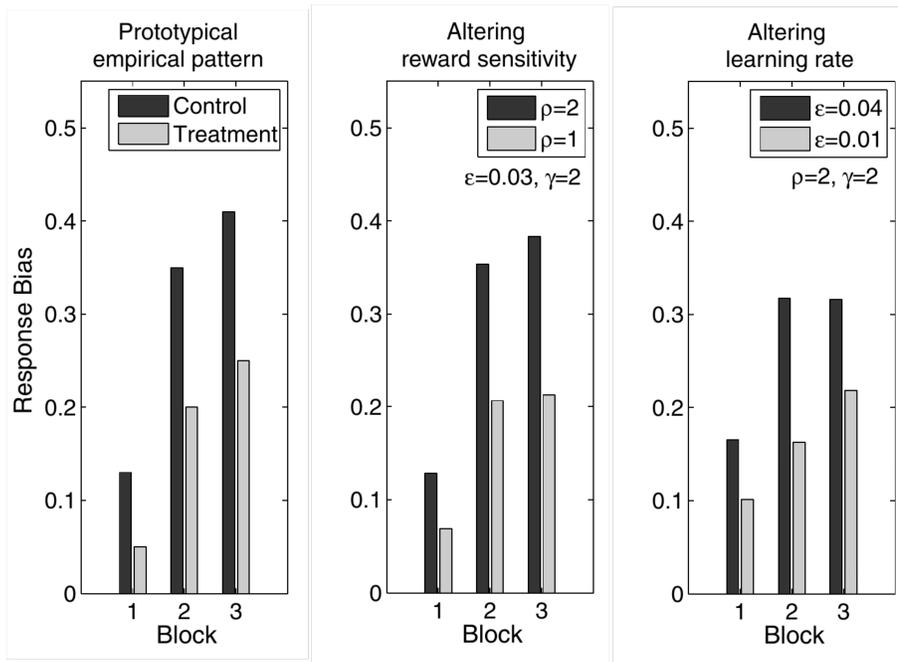


Figure 7.4: Modeling the signal detection task. Response bias on simulated data (adapted from Huys et al. (2013)). Three blocks of 100 trials were simulated and the development of the response bias is shown across these blocks in each bar chart. On the left, a typical pattern of group differences is shown, with controls developing a strong response bias over the three blocks, and patients showing a reduced bias. The middle chart shows how a reduced reward sensitivity (ρ) could lead to these observed differences. The right chart shows how a reduced learning rate could also lead to similar differences.

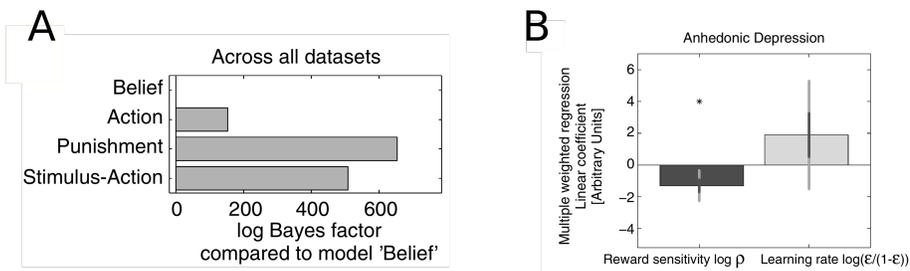


Figure 7.5: Results of the signal detection task (adapted from Huys et al. (2013)). (A) Results of the model comparison. Compared to the three alternative models, the model 'Belief' was shown to be the most parsimonious explanation for the data. (B) Linear (correlation) coefficients between anhedonic depression and reward sensitivity (left; significant at $p < .05$) and learning rate (right; not significant) parameters. (See Huys et al. (2013) for details on this hierarchical regression analysis.)

Model name	Prediction error	W update
Stimulus-action	$\delta_t = \rho r_t - Q_t(a_t, s_t)$	$W_t(a_t, s_t) = \gamma l(a_t, s_t) + Q_t(a_t, s_t)$
Belief	$\delta_t = \rho r_t - Q_t(a_t, s_t)$	$W_t(a_t, s_t) = \gamma l(a_t, s_t) + \zeta Q_t(a_t, s_t) + (1 - \zeta) Q_t(a_t, \bar{s}_t)$
Punishment	$-(1 - r_t) - Q_t(a_t, s_t)$ $\delta_t = \rho r_t + \rho$	$W_t(a_t, s_t) = \gamma l(a_t, s_t) + \zeta Q_t(a_t, s_t) + (1 - \zeta) Q_t(a_t, \bar{s}_t)$
Action	$\delta_t = \rho r_t - Q_t(a_t, s_t)$	$W_t(a_t, s_t) = \gamma l(a_t, s_t) + 0.5 Q_t(a_t, s_t) + 0.5 Q_t(a_t, \bar{s}_t)$

Table 7.1: Summary of models. The choice probability is always $p(a_t \parallel s_t) = \sigma(W_t(a_t, s_t) - W_t(\bar{a}_t, s_t))$ and the Q update is always $Q_{t+1}(a_t, s_t) = Q_t(a_t, s_t) + \varepsilon \times \delta_t$.

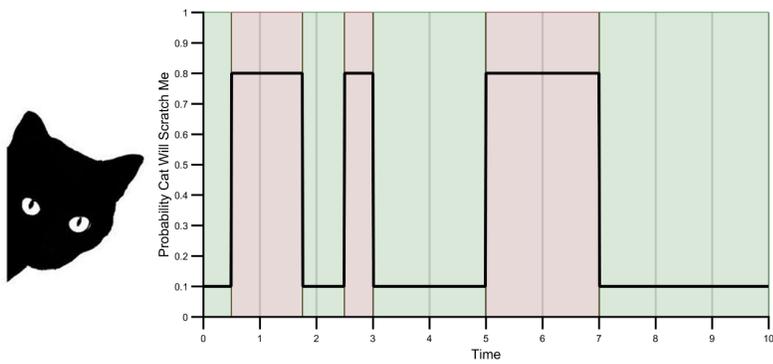


Figure 8.1: Ambiguity in Fear Learning (Pulcu and Browning 2019). Your cat scratches you 10% of the times you stroke it when it is in a good mood (green areas) and 80% of the time you stroke it when it is in a bad mood (red areas). You can't observe the cat's mood; all you can observe is whether it has scratched you when you previously stroked it. In order to learn what the cat's mood is and therefore how likely it is to scratch you the next time you stroke it, you need to account for two sources of uncertainty 1) even if you know exactly what the cat's mood is, you can't be certain what its behavior will be (e.g. even when it is in a good mood it will scratch you 10% of the times you stroked it) and 2) the cat's mood changes over time.

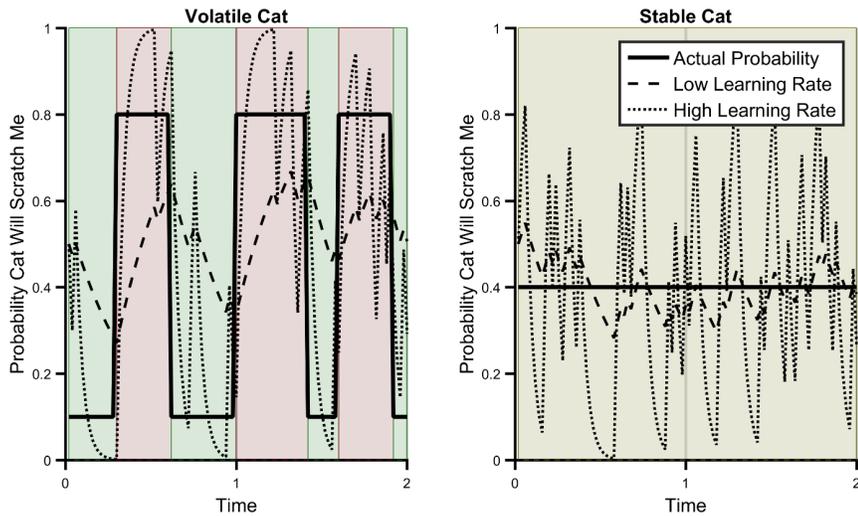


Figure 8.2: The optimal learning rate to use depends on the volatility of the association being learned (Pulcu and Browning 2019). Two Rescorla-Wagner models try and learn how likely two cats are to scratch them. The first cat (left panel) is volatile with periods of better mood (scratch probability 10%) and worse mood (scratch probability 80%). The second cat (right panel) is stable, with one mood (scratch probability 40%). One model uses a low learning rate (dashed line), the other a high learner rate (dotted line), the solid black line is the underlying truth the models are trying to learn. For the volatile cat (left panel), the model with the high learning rate captures the cat’s behavior more accurately than the low learning rate model. This is because the high learning rate model puts more weight on new events than previous events, and the volatility of the cat’s behavior reduces how informative previous events are. For the stable cat (right panel), the low learning rate model more accurately captures its behaviour because previous events are more informative.

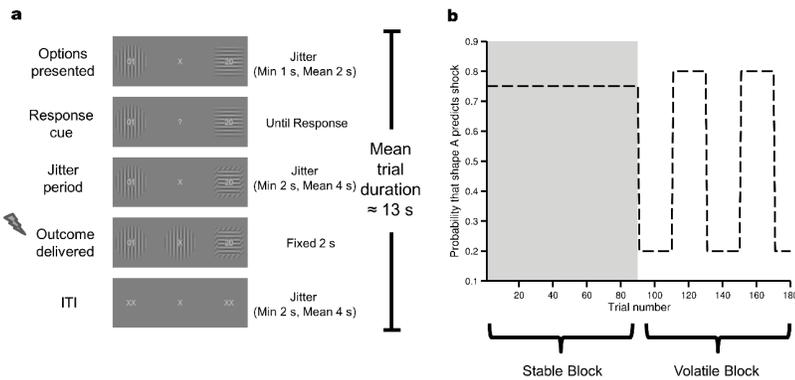


Figure 8.3: Task used in Browning et al. study. a) Example trial from the task. Participants were presented with 2 shapes, each with a number in the center. They chose one shape, if the chosen shape was associated with the electric shock for that trial; they then received the shock after a short delay. The magnitude of the delivered shock was reported by the number in the center of each shape (shock intensity had been calibrated for each participant before the task). Participants therefore had to learn which shape was most likely to be associated with the shock and combine this information with the displayed intensities of the shock for each shape when deciding which shape to choose. b) Structure of the task. Y-axis reports the probability that the shock was associated with ‘shape a’ (the probability that the shock was associated with the other shape is 1 minus these values). The task consisted of 180 trials divided into 2 blocks. In one block, the shock was stably associated with one shape on 75% of trials. In the other block, the association reversed every 20 trials from 20 to 80% and back again. Block order (stable, volatile) was counterbalanced across participants.

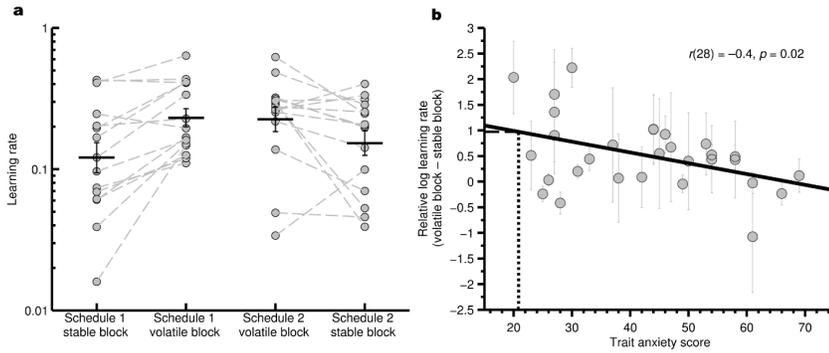


Figure 8.4: Estimated learning rates from the study. a) Individual learning rates from the stable and volatile blocks for all participants. As expected, learning rates were significantly higher in the volatile than the stable blocks. b) Relationship between trait anxiety (x axis) and the degree to which an individual adjusted their learning rate between the volatile and stable blocks (y axis). A significant negative correlation was seen, the more anxious a participant the less they adjusted their learning rate between blocks. The dashed line shows the behavior of a normative Bayesian learner (Behrens et al. 2007) which performs the task optimally, the dotted line illustrates that the behavior of the learner is similar to those participants with low levels of anxiety. The behavior of participants with low anxiety is more like this learner than participants with high anxiety. Anxiety was not associated with differences in any of the other model parameters.

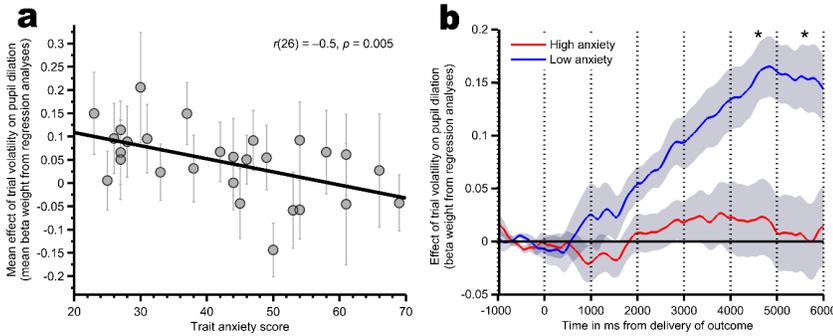


Figure 8.5: Relationship between pupil dilation and trait anxiety. The degree to which pupil dilation during the outcome of a trial was influenced by trial volatility was estimated using regression analyses performed separately for each participant, the larger the beta weights from these analyses the more that participant's pupil dilated during volatile relative to stable trials. Panels a) and b) illustrate the same results. Panel a) demonstrates that the mean effect of volatility on a participant's pupil dilation across the entire 6 second outcome period was negatively correlated with trait anxiety. Panel b) uses a median split on participant anxiety to illustrate the mean (shaded area is SEM) time course of the volatility effect in participants with high and low trait anxiety. These results demonstrate that the pupils of participants with high anxiety show less differentiation between volatile and stable trials than those of low anxiety participants. *=significant difference between the groups after Bonferroni correction for number of time bins.

Disorder	Brief Description
Obsessive-Compulsive Disorder *	Patients experience unpleasant, intrusive thoughts (“obsessions”) and/or the need to perform associated actions (“compulsions”)
PTSD *	Patients experience anxious symptoms following a threatening experience (e.g. being in a car crash, being attacked)
Specific Phobia	Patients experience anxiety when faced with specific objects, places, animals etc. (e.g. spider phobia)
Social Anxiety Disorder	Patients experience anxiety particularly when in social situations or when having to perform (e.g. give a talk)
Panic Disorder	Patients experience recurrent “panic attacks” which include both catastrophic thoughts (e.g. feeling like they are going to die) and physical symptoms such as chest pain or shortness of breath. Often associated with agoraphobia
Agoraphobia	Patients experience anxiety particularly when in situations in which it would be difficult to escape from (e.g. on airplanes) or where help may not be available (e.g. outside the home)
Generalized Anxiety Disorder	Patients worry about a range of activities and/or events. This is commonly associated with depression

Table 8.1: Anxiety Disorder Diagnoses and brief descriptions. Separate diagnoses classified as anxiety disorders in the DSM-IV manual are included. * these disorders were moved to different categories in DSM-V. PTSD= post-traumatic stress disorder

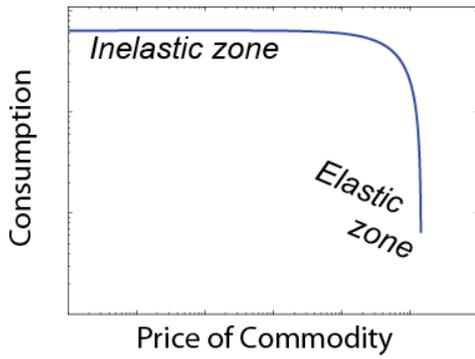


Figure 9.1: Demand curves. The shape of a typical demand curve. As the price of the commodity increases, the number of samples consumed decreases. There is typically an inelastic zone, where large ratio changes have little effect, and an elastic zone where large ratio changes have a larger effect. Note that both axes are logarithmic. Compare, for example, to real demand curves as seen in Bruner and Johnson (2013) where subjects were asked how much cocaine they would buy at a hypothetical given price.

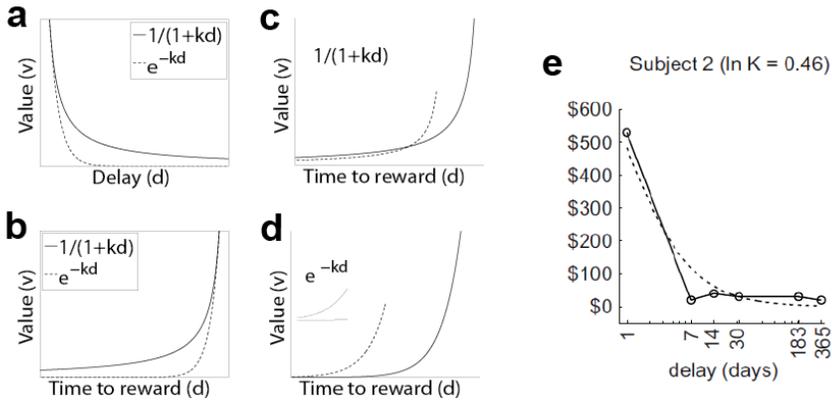


Figure 9.2: Delay discounting. (a) Delay discounting entails a loss of value as a function of delay to an event. Two discounting functions are typically used, hyperbolic [$V = r/(1 + kd)$] and exponential [e^{-kd}], where d is the delay to the event and k is a parameterization factor. (b) Logically, this can be understood in terms of value of an expected event as one approaches the event in time. (c) Hyperbolic discounting functions can create a preference reversal where one prefers one option (the larger later, solid line) to another (the smaller sooner, dashed line) that reverses as one approaches the options in time. (d) Exponential discounting, however, does not reverse, even when both options are far away (see inset) showing an expansion of the far-left edge of the graph. (e) A real discounting curve from an individual, reprinted from Kurth-Nelson and Redish (2010).

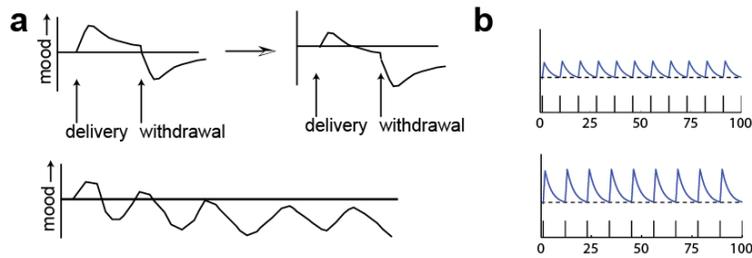


Figure 9.3: Homeostatic /Allostatic processes. (a) Drug delivery produces a positive reaction state, which then adapts and collapses to a negative state when the drug is removed. Over time, the user is hypothesized to adapt to the positive state, producing a shift in the allostatic set-point towards the negative state. Redrawn after Koob (2013). (b) Tsibulsky and Norman (1999) and Keramati et al (2017) modeled self-administration as an attempt to maintain the total level of drug at a given set-point. As the drug was processed internally and reduced beyond the set-point, the animal was hypothesized to seek the drug through lever pressing. This model explains the different rates of lever pressing as a function of the drug dose. Redrawn after Tsibulsky and Norman (1999) and Keramati et al (2017).

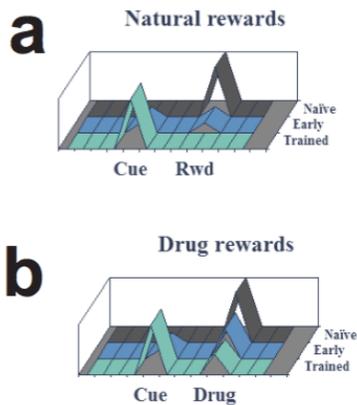


Figure 9.4: The delta signal – dopamine and delta. Diagram of delta (vertical axis) by time (horizontal axis) over three conditions: naïve (untrained), early (with limited training), and trained. **(a)** With normal rewards, the delta signal shifts from appearing at the unexpected reward to the unexpected cue-that-predicts-reward. **(b)** In the Redish (2004) model, there are two components in the delta signal, a reward-related component that shifts and a pharmacological component that remains at the reward time. Compare the classic data from Schultz (1998). When the expected reward is not delivered, dopamine cells pause their firing. Aragona et al. (2009) tested the double bump hypothesis and found that the cue-related signal occurred in accumbens core, while the pharmacological component occurred in shell.

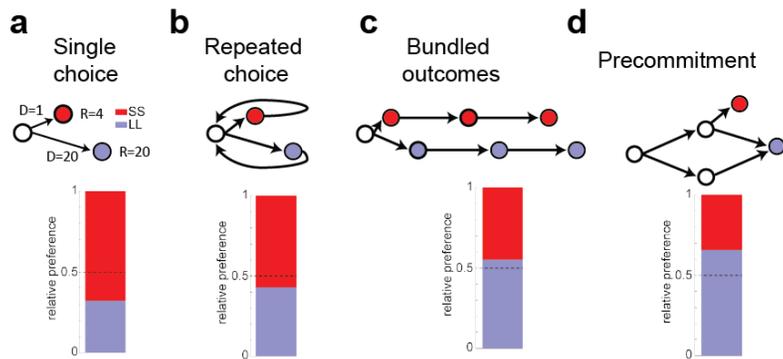


Figure 9.5: Changing state spaces. (a) Imagine a single choice between a smaller reward ($R=4$) delivered sooner (after 1s), compared to a larger reward ($R=20$) delivered later (after 20s). A typical agent might prefer the smaller-sooner over the larger-later reward. (b) If the agent realizes that this is going to be a repeated choice, then it is possible to drive the relative preference to 50/50 with a long look-ahead, but it is impossible to change the actual preference. An agent that prefers the smaller-sooner option in (a) will still prefer it in (b). (c) Bundling creates new options such that there are consequences to one's decision. An agent can switch preferences by bundling. (d) Precommitment adds a new option to skip the choice. An agent making a decision at the earlier option can prefer the larger-later and learn to skip the choice in the right conditions. After models in Kurth-Nelson and Redish (2012).

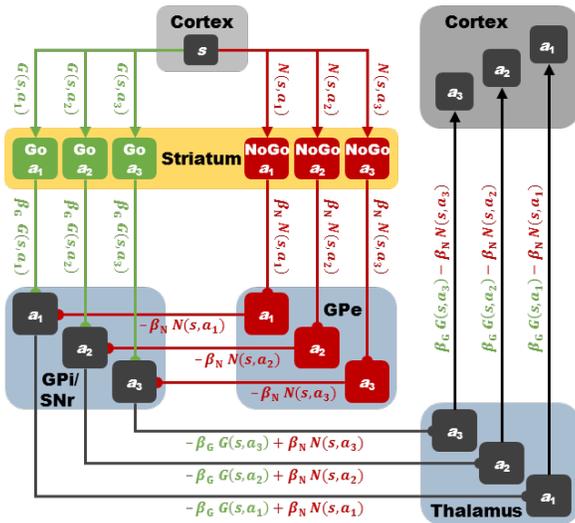
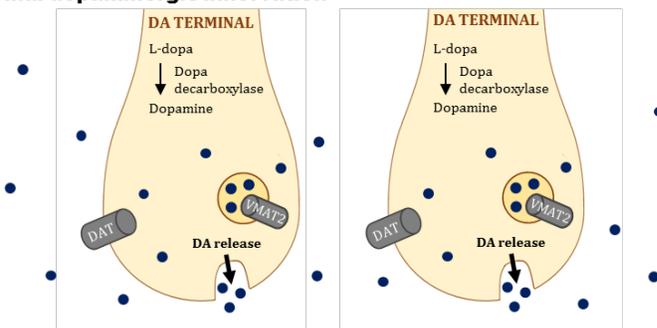


Figure 10.1: Computational roles of the direct (or Go) and indirect (or NoGo) motor pathways in action learning and selection. The situation or state, s , is a potentially rich, multidimensional representation encompassing representations of the external world (perception), interoception, motivational and emotional states, etc., potentially represented in a wide range of cortical (and some subcortical) regions—which likely explains why many regions project to the striatum (Choi, Yeo, and Buckner 2012; Postuma and Dagher 2006). The figure represents three possible actions (a_1 , a_2 , and a_3); their representation is kept separate in each anatomical region as required to have action specificity. The striatum contains two populations of neurons: D₁-expressing (Go) and D₂-expressing (NoGo) medium spiny neurons (MSNs), which are represented in green and red, respectively. Each of these populations contains a representation of the three actions. Go (G) and NoGo (N) values that represent the learned associations between state s and each of the actions are represented in corticostriatal synapses onto Go and NoGo MSNs, respectively. The gains of Go and NoGo MSNs (β_G and β_N , respectively) are modulated by striatal dopamine levels. The activation, and therefore the output, of Go and NoGo MSNs is thus given by $\beta_G G(s, a_i)$ and $\beta_N N(s, a_i)$, respectively. Mathematically, inhibitory projections (GABAergic projections, represented by circles) flip the sign of the information (provided that there is intrinsic activity

in the target structures, as is the case here). The anatomy of the basal ganglia seems therefore precisely suited to represent the difference between Go and NoGo activations, $\beta_G G(s, a_i) - \beta_N N(s, a_i)$, in the thalamus. The latter, in turn, helps to select actions in cortex in proportion to this difference. Internal variables, structures, and projections related to the direct and indirect pathways are coded in green and red, respectively. Glutamatergic (excitatory) projections are represented by arrowheads. GPe: external segment of the globus pallidus; GPi: internal segment of the globus pallidus; SNr: substantia nigra *pars reticulata*. Figure and caption adapted, with permission, from Maia and Conceição (2017).

A. Normal dopaminergic innervation



B. Dopaminergic hyperinnervation

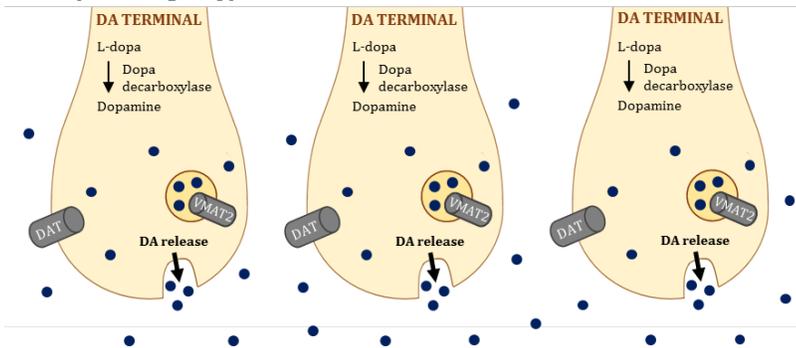


Figure 10.2: The dopaminergic-hyperinnervation hypothesis of Tourette syndrome (TS). (A) Normal dopaminergic innervation. Multiple characteristics of dopamine terminals can be investigated using molecular imaging *in vivo*: the dopamine transporter (DAT), the vesicular monoamine transporter 2 (VMAT2), the activity of dopa decarboxylase (by measuring F-dopa uptake), and the extent of amphetamine-induced dopamine (DA) release. (B) The hypothesis that TS involves dopaminergic hyperinnervation—that is, an increased number of DA terminals—explains why all of these markers (DAT and VMAT2 binding, F-dopa uptake, and amphetamine-induced DA release) seem to be increased in TS (section 10.3.1; Maia and Conceição 2018).

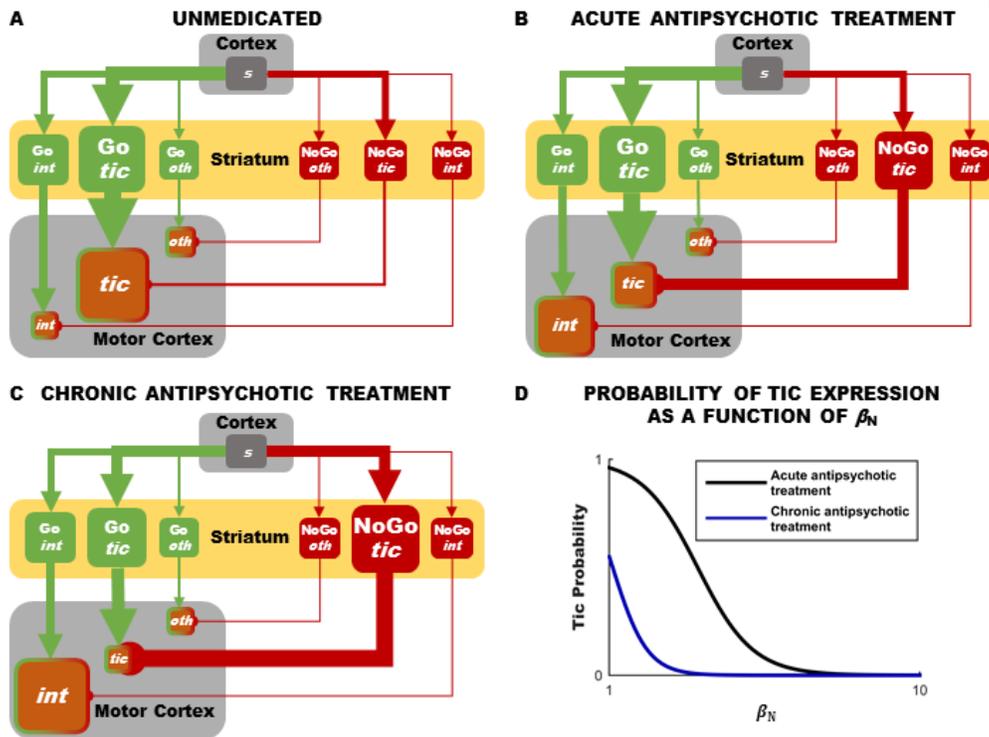


Figure 10.3. Effects of striatal dopamine levels and antipsychotic treatment on Tourette syndrome (TS). Panels A–C represent the Go and NoGo pathways (green and red, respectively) for three actions: an intended action (*int*), a tic (*tic*), and some other weakly supported action (*oth*). The pathways are represented schematically, through their ultimate effects on motor cortex, by omitting the globus pallidus and thalamus. Go and NoGo MSNs are therefore shown directly stimulating (or, more precisely, disinhibiting) and inhibiting the motor cortex, respectively, because activation in motor cortex ultimately reflects a subtraction of the outputs of the Go and NoGo pathways (section 10.3.2; Figure 10.1). (In the figure, arrowheads represent excitation or disinhibition; circles represent inhibition.) The size of each square and the width of the arrow that departs from it represent the level of neuronal activity. Panels A–C differ in terms of patient medication status. (A) In unmedicated patients with TS, tics may have strong learned Go values, stored in corticostriatal synapses onto Go MSNs (see the thick green arrow from cortex

Commented [VC1]: Corrected section and Figure citation

to “Go tic” in the striatum); these Go values may be learned through ill-timed, exaggerated phasic dopamine responses or through negative reinforcement due to the temporary relief from the preceding premonitory urge (as mechanistically explained in **section 10.3.3**). Tics may also have relatively strong NoGo values, stored in corticostriatal synapses onto NoGo MSNs (see the relatively thick red arrow from cortex to “NoGo tic” in the striatum); these NoGo values may be learned through negative life experiences with tics (e.g., being embarrassed because of tics, feeling sore because of a tic, etc.). In unmedicated patients, however, the expression of these NoGo values is likely suppressed by the high striatal dopamine—predicted to occur under dopaminergic hyperinnervation (**section 10.3.1**; Hienert et al. 2018; Maia and Conceição 2018). Therefore, in an unmedicated patient with TS, the Go activity for tics overcomes the NoGo activity, making tic expression likely. **(B)** As soon as a patient begins antipsychotic treatment, or as soon as the antipsychotic reaches a sufficiently high dose, the antipsychotic blocks D_2 receptors in the NoGo pathway, disinhibiting that pathway, which then becomes stronger and better able to counteract the activity in the Go pathway. The tic thereby becomes less likely to be expressed. This very early effect of the antipsychotic may act mostly through this effect on excitability; the corticostriatal synapses representing Go and NoGo values may not yet be changed. **(C)** In addition to the effect on excitability, chronic antipsychotic treatment also increases NoGo values (by increasing the weight of corticostriatal synapses onto NoGo MSNs), decreases Go values (by decreasing the weight of corticostriatal synapses onto Go MSNs), or both (**section 10.3.2**; Maia and Conceição 2017). Thus, with the same level of D_2 occupancy as acute antipsychotic treatment, the tic becomes even less likely to be expressed (compare panels B and C). (For illustrative purposes, panel C shows the case of both increased NoGo values and decreased Go values, but the same effect would be obtained with changes in just one or the other.) **(D)** Due to the fact that only chronic antipsychotic treatment is likely to strongly affect NoGo and/or Go values (i.e., to change synaptic weights rather than just excitability), acute and chronic antipsychotic treatment likely provide different levels of symptomatic control, as quantified by the probability of executing a given tic in each case. The x axis in the figure represents β_N as a proxy to the effect of the antipsychotic on excitability; larger values of β_N correspond to larger antipsychotic doses. The black and blue lines represent (qualitatively) the probabilities of tic execution following acute and chronic treatment, respectively. For a given dose of antipsychotic (i.e., for a given value of β_N), the probability of executing a tic is lower following chronic administration than following acute administration. This explains why antipsychotics may have a gradual cumulative effect and why, during chronic treatment, the dose may sometimes be reduced gradually without loss of efficacy (Maia and

Conceição 2017). Still, if the dose is reduced too drastically or medication is completely stopped, tics that were completely absent may return. In the plot, this corresponds to moving left along the blue line, from a point in which the probability is nearly 0 to a point in which it becomes more substantial. Still, tics may be less severe than before treatment started (at least temporarily, until relearning occurs). In fact, tics may also be less severe right after stopping chronic treatment than they would be after a single acute dose of an antipsychotic wears off (compare the intercept of the blue vs. the black line). Figure and caption adapted, with permission, from Maia and Conceição (2017).

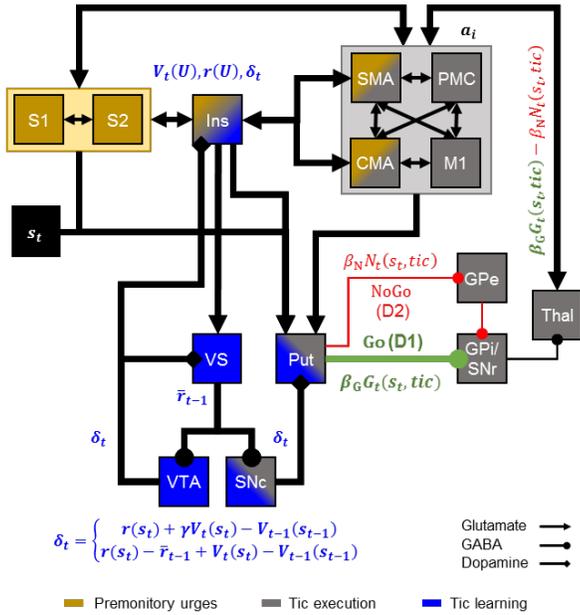


Figure 10.4: Regions and computations involved in premonitory urges, tic execution, and tic learning, in Tourette syndrome (TS). The figure depicts the main regions involved in premonitory urges (gold), tic execution (gray), and tic learning (blue), and the most important connections between them, according to the framework that we have previously proposed (section 10.3.3; Conceição et al. 2017). This framework addresses both the computational roles of dopamine in motor-loop-mediated tic learning and execution (section 10.3.2; Maia and Conceição 2017) and the neural substrates and mechanistic roles of premonitory urges in tic learning and execution (Conceição et al. 2017). *Tic execution:* Cortical motor areas represent candidate actions (a_i), including tics, being considered for gating. Other cortical (and some subcortical) areas represent the current state or situation (s_t , where the subscript t denotes time). The putamen contains Go (direct-pathway, green) and NoGo (indirect-pathway, red) medium spiny neurons (MSNs; Figure 10.1). Striatal hyperdopaminergia, due to dopaminergic hyperinnervation in TS (section 10.3.1; Hienert et al. 2018; Maia and Conceição 2018), increases the activation of the Go relative to the NoGo pathway (by increasing the value of β_G relative to β_N ; Figure 10.3). Such increase is

disproportionately larger for actions with large Go (G) values, because β_G is a multiplicative gain parameter. Consequently, as tics become strongly learned behaviors (see below and **sections 10.3.2–3.3**), striatal hyperdopaminergia will make $\beta_G G_t(s_t, tic) \gg \beta_N N_t(s_t, tic)$ (compare the width of the green and red arrows leaving from the putamen; **Figures 10.1 and 10.3**). As a consequence, there is strong inhibition of the tic representation in the basal ganglia output nuclei, the globus pallidus internal segment (GPi) and substantia nigra *pars reticulata* (SNr), by the direct pathway [$\beta_G G_t(s_t, tic)$], with weak disinhibition by the NoGo pathway [with a value of $\beta_N N_t(s_t, tic)$]. Given that $\beta_G G_t(s_t, tic) \gg \beta_N N_t(s_t, tic)$, tic execution is promoted (**Figure 10.3**; Maia & Conceição, 2017). Tic execution may also possibly be driven directly by cortico-cortical projections from somatosensory regions and the insula to motor cortices (Conceição et al. 2017). *Tic learning*: Tic execution commonly terminates a preceding premonitory urge, yielding a positive prediction error, δ_t , that promotes tic learning (Conceição et al. 2017). Like δ_t , $r(U)$ is a relevant variable to explain tic learning in the two reinforcement-learning (RL) accounts that we have used to explain premonitory-urge-driven tic learning: standard RL and average-reward RL (Conceição et al. 2017). In contrast, $V_t(U)$ is necessary for the account using standard RL and optional for the account using average-reward RL, and r_{t-1} is only necessary for the account using average-reward RL (Conceição et al. 2017). In this figure, those variables are depicted near the regions and/or connections that we have hypothesized to subserve them (Conceição et al. 2017). The prediction errors represented in the insula, however, may be mostly aversive. *Additional figure details*: Some anatomical projections are omitted for simplicity. For clarity, both somatosensory cortical areas [primary (S1) and secondary (S2) somatosensory cortices] and motor cortical areas [cingulate motor area (CMA), supplementary motor area (SMA), premotor cortex (PMC), and primary motor cortex (M1)] are grouped. The distinct ways of calculating δ_t according to each of the proposed computational accounts are indicated in curly braces: $\delta_t = r(s_t) + \gamma V_t(s_t) - V_{t-1}(s_{t-1})$ (standard RL) and $\delta_t = r(s_t) - r_{t-1} + V_t(s_t) - V_{t-1}(s_{t-1})$ (average-reward RL; section 3.3). *Additional abbreviations*: Ins: insula; Put: putamen; SNc: substantia nigra *pars compacta*; Thal: thalamus; VS: ventral striatum; VTA: ventral tegmental area. Figure and caption adapted, with permission, from Conceição et al. (2017).

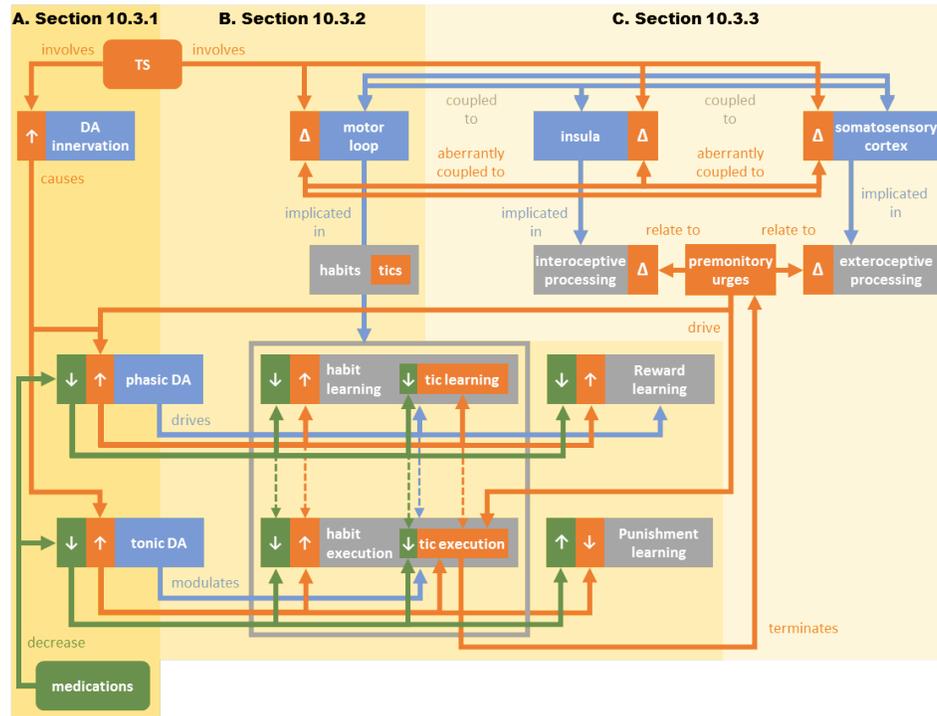


Figure 10.5: Concept map summarizing the main theoretical ideas in the chapter. The key brain regions that we discuss are shown in blue, the key cognitive processes are shown in gray, the key alterations in Tourette syndrome (TS) are shown in orange, and the key medication effects are shown in green. Upward and downward pointing arrows represent increases and decreases, respectively (due to TS or the medication if they are in orange or green boxes, respectively). Triangles show alterations. Each region shaded in a different color corresponds to a specific chapter section. **(A) Section 10.3.1.** As discussed in **section 10.3.1** and depicted in the area shaded in gold, patients with TS likely suffer from dopaminergic hyperinnervation (Hienert et al. 2018; Maia and Conceição 2018), which causes increases in both phasic and tonic dopamine. Those dopaminergic increases help to explain why all medications with well-established efficacy for TS reduce phasic and/or tonic dopaminergic neurotransmission (Maia and Conceição 2018). **(B) Section 10.3.2.** As discussed in **section 10.3.2** and depicted in the area shaded in medium-light gold, TS involves structural and functional disturbances in the motor loop (Maia and

Conceição 2017; Worbe, Lehericy, and Hartmann 2015; Worbe et al. 2015), which is implicated in both habit learning and execution (Delorme et al. 2016; Horga et al. 2015; Yin and Knowlton 2006). Habit learning and habit execution are strongly mediated by phasic and tonic dopamine, respectively (Collins and Frank 2014; Maia and Conceição 2017), with some novel evidence also implicating phasic dopamine in habit execution (hence, the blue dashed arrow; da Silva et al. 2018). Tics seem to be pathological motor habits (hence, their depiction in the orange rectangle inside the blue “habits” rectangle), which explains why, in TS, phasic- and tonic-dopamine levels may play pathological roles in tic learning and execution that parallel the roles normally played by normal phasic- and tonic-dopamine levels in habit learning and execution (note the similarities between the respective orange and blue arrows; Maia and Conceição 2017). The latter associations help to explain mechanistically how medications may reduce tic severity in TS: by downregulating phasic and/or tonic dopaminergic neurotransmission, medication should reduce tic learning and/or execution (Maia and Conceição 2017). The role of phasic dopamine in signaling positive prediction errors explains why unmedicated patients with TS, with their likely increase in phasic dopamine, exhibit increased learning from rewards (Palminteri et al. 2009; 2011). The role of tonic dopamine in possibly blunting the signaling of negative prediction errors by phasic decreases in dopamine explains why TS involves reduced learning from punishments (Palminteri et al. 2009). In turn, the previously mentioned effects of TS medications on dopaminergic transmission—possibly along with other, more complex effects of the medication (Maia and Conceição 2017)—explain why patients with TS under antipsychotics (other than aripiprazole) are impaired at reward learning (Palminteri et al. 2009; 2011; Worbe et al. 2011). **(C) Section 10.3.3.** As discussed in section 10.3.3 and depicted in the area with the lightest gold shading, TS involves structural and/or functional abnormalities in the somatosensory cortices and insula, in addition to the motor loop. TS, moreover, involves structural and functional abnormalities in the connectivity between all those regions (Conceição et al. 2017; Neuner, Schneider, and Shah 2013; Sigurdsson et al. 2018). The somatosensory cortices are implicated in exteroception (Cox, Seri, and Cavanna 2018) and the insula is implicated in both (natural) urges (Naqvi and Bechara 2010; Jackson, Parkinson, Kim, et al. 2011) and interoception (Quadt, Critchley, and Garfinkel 2018; Cox, Seri, and Cavanna 2018). Premonitory urges, in turn, seem to be driven by the interplay between pathological exteroceptive and/or interoceptive processing (Cox, Seri, and Cavanna 2018), which explains the implication of the somatosensory cortices and insula in premonitory urges (**Figure 10.4**; Conceição et al. 2017). Premonitory urges likely play a key role in tic learning, as the termination of premonitory urges, via tic execution, may elicit positive prediction errors, signaled by phasic dopamine, that reinforce tics

(Conceição et al. 2017). Premonitory urges may also directly drive tic execution through the connections from somatosensory regions and the insula to motor cortices (Conceição et al. 2017).