Hierarchical Network Model Excitatory-Inhibitory Tone Shapes Alternative Strategies for Different Degrees of Uncertainty in Multi-Attribute Decisions

Warren Woodrich Pettine¹, Kenway Louie¹, John D Murray² and Xiao Jing Wang^{1,*}

¹Center for Neural Science, New York University, New York, NY 10003, USA

²Department of Psychiatry, Yale University School of Medicine, New Haven, Connecticut
06510

*Corresponding Author. E-mail: xjwang@nyu.edu

January 28, 2020

Abstract

We investigated two-attribute, two-alternative decision-making in a hierarchical neural network with three layers: an input layer encoding choice alternative attribute values; an intermediate layer of modules processing separate attributes; and a choice layer producing the decision. Depending on intermediate layer excitatory-inhibitory (E/I) tone, the network displays three distinct regimes characterized by linear (I), convex (II) or concave (III) choice indifference curves. In regimes I and II, each option's attribute information is additively integrated. To maximize reward at low environmental uncertainty, the system should operate in regime I. At high environmental uncertainty, reward maximization is achieved in regime III, with each attribute module selecting a favored alternative, and the ultimate decision based upon comparison between outputs of attribute processing modules. We then use these principles to examine multi-attribute decisions with autism-related deficits in E/I balance, leading to predictions of different choice patterns and overall performance between autism and neurotypicals.

1 Introduction

We are constantly faced with decisions between alternatives defined by multiple attributes. The true value of each attribute is at times clear, and other times uncertain. For example, on Friday one might choose between main courses at a restaurant where the flavor or healthiness attributes of all the dishes are familiar. The following Wednesday might be at a restaurant with an unknown cuisine, where one is highly uncertain as to different items' flavor or healthiness. To ensure the best meal, the brain must be able to optimize choice in both environments.

Systems neuroscientists have, for many years, been studying the specific circuits engaged in this kind of multi-attribute decision-making. Based on a robust set of electrophysiology and imaging findings (Xie and Padoa-Schioppa 2016; Raghuraman and Padoa-Schioppa 2014; Padoa-Schioppa and Assad 2006; O'Neill and Schultz 2018; Morrison and Salzman 2009; Conen and Padoa-Schioppa 2015; Chib et al. 2009; Pastor-Bernier, Stasiak, and Schultz 2019), many hold that all attribute signals are available in brain areas proximal to the final decision (Levy and Glimcher 2012; Padoa-Schioppa and Conen 2017). Indeed, when attribute values are clear, multi-attribute choice theoretically is simple: linearly weight and combine all attributes associated with a choice alternative, then select the one with the larger value. Though the subjective value of an attribute might be non-linearly related to the quantity offered, when the final choice is made in an environment without uncertainty, a weighted linear combination of attributes optimizes the choice between options (Nicholson and Snyder 2007).

However, an alternative perspective has recently arisen that takes into account the multiple brain areas implicated in decisions (Wunderlich, Rangel, and O'Doherty 2009; Peck, Lau, and Salzman 2013; Paton et al. 2006; Munuera, Rigotti, and Salzman 2018; Louie and Glimcher 2010; Chen and Stuphorn 2015; Steinmetz et al. 2019). Under this alternate view, a signal is transformed as it moves from hierarchically lower areas to those that are proximal to the final decision (Cisek and Kalaska 2010; Cisek 2012; Yoo and Hayden 2018). Such a process can allow for parallel computations, and produce clearer separation between the representation of choice alternatives where the decision is reached. Furthermore, these transformations can be highly non-linear, providing an additional layer of flexibility in the decision process.

When investigating the non-linear transformations that can occur in populations of neurons, it is imperative to consider excitatory and inhibitory (E/I) tone. In sensory areas, E/I tone can shape the stimulus tuning curve and responses timing (Mariño et al. 2005; Wehr and Zador 2003). When sensory information is used in perceptual decision-making, E/I tone can dictate the trade-off between speed and accuracy (Wong and Wang, 2006). In working memory tasks, E/I tone defines a network's ability to maintain a memory, and that memory's susceptibility to a distractor (Brunel and Wang 2001; Compte et al. 2000). When multiple

areas interact, the E/I tone can define their function in either memory maintenance or decision-making (Murray, Jaramillo, and Wang 2017). On a more theoretical level, E/I tone can shape a system's response as an input signal increases (Ahmadian and Miller 2019) and its transition to chaotic dynamics (Vreeswijk and Sompolinsky 1998). Yet, it remains to be described how E/I tone defines the functional interaction and transformation of multiple signals passing through a brain area, as in a hierarchical system.

Though a prior biophysical model of hierarchical neural networks engaged in multi-attribute decision-making captured key features of human reaction times and functional imaging (Hunt, Dolan, and Behrens 2014), how a hierarchical system shapes multi-attribute decisions is largely a matter of conjecture (Hunt and Hayden 2017). By their very nature, these transformations distort the signals available for the final decision. Distributed models have been shown to benefit motor control (Christopoulos, Bonaiuto, and Andersen 2015; Cisek 2006), but in what situations might such transformations improve a decision process?

To approach this question, we examine how E/I tone in hierarchical networks governs multi-attribute decision-making under varying degrees of environmental uncertainty. It is often the case that noise in the environment creates uncertainty as to the true value of attributes (Bach and Dolan 2012). Additionally, environmental uncertainty offers an entry-point to understanding neuropsychiatric conditions, such as autism spectrum disorder (ASD). A myriad of findings suggest that the diagnostic signs of behavioral rigidity and sensory abnormalities (American Psychiatric Association 2013) may be related to an intolerance of environmental uncertainty (Foss-Feig et al. 2017; Fujino et al. 2017; Boulter et al. 2014; Van de Cruys et al. 2014; Vasa et al. 2018). A mechanistic understanding of how hierarchical neural systems handle environmental uncertainty when making decisions can not only provide knowledge as to how the brain functions, but also may produce hypotheses as to the etiology of ASD.

We find that the E/I tone of the network creates distinct regimes defined by their choice indifference curve: a linear weighting of attribute values (regime I); a convex preference for balanced attributes (regime II); or a concave increased weighting of the larger attribute value (regime III). We then show that the degree of reward maximization by these regimes depends on the level of environmental uncertainty. When uncertainty is low, regimes I or II - where the choice area has access to all signal information - achieve the greatest reward. However, when uncertainty is high, regime III - with strong transforms along the hierarchy favoring the larger attribute value - on average, maximizes reward. After a detailed investigation of how these regimes arise and why they are optimal, we hypothesize a framework where the E/I tone can be modified according to the level of uncertainty. The framework is then used to create a model of neurotypicals with the full range of E/I tone, and a model of ASD populations defined by limited inhibitory tone. Based on the models, we predict that neurotypical and ASD subjects will adopt a linear weighting of attributes in environments with low uncertainty (regimes I/II). As environmental uncertainty increases, neurotypical subjects will fully move

into regime III, while ASD individuals will not. This will result in larger performances differences at high environmental uncertainty levels.

Using a hierarchical network parameterized by E/I tone, we are able to show that the the network can adopt multiple regimes (including linear) to maximize reward in various environments, and then we use these findings to provide an etiological hypothesis for behavior observed in ASD.

2 Results

2.1 Networks Perform Multi-Attribute Decision Task Under Uncertainty

To investigate the behavior of neural networks in environments with varying degrees of uncertainty, we used a multi-attribute decision task with varying degrees of environmental noise. On any given trial, the neural networks were presented with two choice alternatives (A or B). These choice alternatives were each composed of two attributes (1 or 2). For simplicity, attributes were assumed to have equal contributions to subjective value and actions were assumed to be symmetric. When indexing, the first value refers to the choice alternative, and the second to the attribute. Thus, choice alternative A was presented to the network using input $I_{A,1}$ and input $I_{A,2}$, while choice alternative B was presented to the network using input $I_{B,1}$ and input $I_{B,2}$. Inputs were translated from offer values to firing rates. As the final choice area competition was between alternatives, the signal at that stage is denoted by C_A or C_B , where the subscript indicates the choice alternative. On a given trial, the choice between A or B was determined by the first population in the choice area (C_A or C_B) to cross a firing rate threshold (35 Hz). To control the level of uncertainty in the environment, we added a noise term η_I to each input at the beginning of the trial, independently drawn from a normal distribution centered on 0, with a standard deviation of σ_{η_I} .

We created two basic network frameworks: the Linear Network and the Hierarchical Network. Areas in these networks were composed of mean-field approximations of population firing rates (see Methods). As the name suggests, the Linear Network computed an exact linear weighted sum of the attributes (weight = 0.5). Thus, with the Linear Network, the choice area received a signal that was linearly translated from the presented attributes of the choice alternatives (Figure 1A). This network has a similar architecture to that presented in (Rustichini and Padoa-Schioppa 2015), and represents the case where all attribute signals are fully available to a choice area.

In the Hierarchical Network (Figure 1B), inputs were first transformed by intermediate, attribute-specific areas. The intermediate layer transforming the input signal was a nonlinear dynamical system, composed of attractor states associated with each choice alternative (see Methods). To denote the transformation

performed by an intermediate layer area on the input signal, its output signal will be referred to as T , such that $I_{A,1}$ is related to $T_{A,1}$, etc. These continuous output signals were then fed into a choice area that determined the decision on a given trial. A single trial of the Hierarchical Network is shown in Figure 1C. This framework, similar in structure to that presented in (Hunt, Dolan, and Behrens 2014), represents the simplest form of a network with parallel processing streams, where transformations are performed on attribute signals passing through specialized areas prior to the final decision.

To first verify that the networks were able to perform the task, we computed simple psychometric functions of model choice behavior. This was done by varying the input $I_{A,1}$, while keeping inputs $I_{A,2}$, $I_{B,1}$, and $I_{B,2}$ fixed. We then measured the proportion of trials that networks chose A , and fit a sigmoid function to the resultant $P(A)$ (see Methods). We did this for environments where there was either no uncertainty ($\sigma_{\eta_I} = 0$ Hz), or a moderate amount of uncertainty ($\sigma_{\eta_I} = 0.75$ Hz).

Figure 2A shows the psychometric functions for the Linear Network in a certain environment (magenta) and an uncertain environment (green). The same is shown in Figure 2B for the Hierarchical Network with intermediate layer recurrent excitation ($J_+ = 0.33$ nA and cross inhibition ($J_- = -0.03$ nA. Both networks generated expected psychometric functions, with larger input value differences resulting in better performance. Furthermore, the slope of the functions was, as anticipated, determined by the level of environmental uncertainty, such that the greater uncertainty produced a less precise choice performance.

2.2 The Intermediate Layer Weights of Hierarchical Networks Dictate Decision Regimes

Having verified the networks were able to perform the task, we next examined how they utilize attribute information. While the Linear Network performs a simple linear sum of the attributes, the Hierarchical Network first passes the input signals through attribute-specific areas (Figure 1B and Figure 3A). This transform of inputs is defined by the recurrent excitation and cross-inhibition of areas in the intermediate layer.

To investigate how the transform of attribute inputs changed with weights, we used indifference curves. When generating an indifference curve for each weight-configuration, we varied the composition and magnitude of attributes in a single choice alternative ($I_{A,1}$ and $I_{A,2}$), while holding the composition and magnitude of the other choice alternative ($I_{B,1}$ and $I_{B,2}$) constant as a reference. For some values of $I_{A,1}$ and $I_{A,2}$, (for example when both are much less than their counterparts) $P(A)$ will be low. For other values of $I_{A,1}$ and $I_{A,2}$, (for example when both are much greater than their counterparts) $P(A)$ will be high. However, there will be a set of values of $I_{A,1}$ and $I_{A,2}$ where the subject is indifferent between the choices and $P(A) \approx 0.5$.

Figure 1: Network Schematics and Sample Trial. *I*: input; *C*: choice layer synaptic gating variable; *T*: intermediate transform layer synaptic gating variable; red: choice *A*; blue: choice *B*; black: attribute 1; white: 2; arrows: excitation; circles: inhibition. Both networks include an input layer and a final choice layer. (A) The Linear Network consists of two layers, with attribute signals from the input layer directly transmitted to the final choice layer. (B) The Hierarchical Network includes an additional intermediate layer that performs a functional transformation on the attribute signals prior to their passing to the choice area. (C) A sample trial of the Hierarchical Network, with intermediate layer weights $J_+ = 0.34$ nA and $J_- = -0.02$. For each area, the X-axis indicates time, while the Y-axis indicates the population firing rate. The first vertical line indicates the onset of the offer value signal, and the second vertical line indicates termination of the offer value signal.

Figure 2: Choice Performance Under Varying Environmental Uncertainty. The proportion of trials where the network chose option A, $P(A)$, as a function of the difference in offer inputs is shown, along with a fit sigmoid. Input $I_{A,1}$ was varied as all other inputs were held fixed. On the x-axis, inputs were normalized ($\hat{I}$) by the maximum choice alternative offer value's linear sum. (A) Performance of the Linear Network when attribute values are certain, with $\sigma_{\eta_I} = 0.0$ Hz (magenta), and when attribute values are uncertain, with $\sigma_{\eta_I} = 0.75$ Hz (green). (B) The same is given for the Hierarchical Network, with a intermediate layer weights of $J_+ = 0.33$ nA and $J_- = -0.03$ nA.

The values where $P(A) \approx 0.5$ were used to fit the indifference curve.

The Linear Network by definition computes a weighted sum of the attributes, so its indifference curve is trivially linear, following the line where $I_{A,1} + I_{A,2} = I_{B,1} + I_{B,2}$. For the Hierarchical Network, however, the indifference curve provides a tool for analyzing how the network weighs and combines attribute information (Figure 3). A linear indifference curve indicates that the network is performing a linear combination (similar to that of the Linear Network). A convex indifference curve indicates that the network places greater weight on attribute values composing a choice alternative that is balanced between the attributes, and less weight on attribute values composing an unbalanced choice alternative. A concave indifference curve indicates the opposite; specifically the network under-weights the smaller attribute value, and instead place preferentially greater weight on the larger of the attributes.

We parametrically varied the levels of excitatory (0.30 nA to 0.40 nA, spaced 0.05 nA) and inhibitory weights (0.0 nA to -0.10 nA, spaced 0.05 nA) within intermediate layer areas. The varied weights are colored yellow in Figure 3A. For each weight-configuration, we held the inputs $I_{B,1}$ and $I_{B,2}$ constant at 20 Hz, while independently varying $I_{A,1}$ and $I_{A,2}$ from 0 to 40 Hz. We did this for 1,000 trials at each combination and from those calculated $P(A)$.

We then fit a single-parameter exponential function to the points where $P(A) \approx 0.5$ (see Methods). That

Figure 3: Decision Regimes of Hierarchical Networks. I : input; $\hat{I}$: input normalized by the reference choice alternative values' linear sum; T : transform layer synaptic gating variable; $P(A)$: proportion of 1,000 trials where alternative A was chosen; J_+ : excitatory weight; J_- : inhibitory weight; yellow: varied weights; forest green: regime II; tan: regime I; sky blue: regime III. The first subscript indicates choice A or B, and the second subscript x indicates a generic attribute. Weights were systematically varied identically in all intermediate layer areas. (A) A diagram of a generic intermediate layer area and its inputs. (B) Examples of regime I (linear), regime II (convex) and regime III (concave) decision regimes. (C) The space of E/I weights is partitioned according to the decision regime of their indifference curves.

single parameter determined the curvature, such that a value of 1 indicates a linear indifference curve, a value less than one indicates convex, and a value greater than 1 indicates concave. We found that at different levels of excitation and inhibition, the network adopted distinct decision regimes, corresponding to linear (regime I), convex (regime II) and concave (regime III) indifference curves (Figure 3B). The appearance of these decision regimes formed distinct regions in the weight-space, such that we were able to use these regions to partition the weight-space (Figure 3C). This revealed that by configuring the levels of excitation and inhibition in the intermediate layer, the Hierarchical Network can adopt different decision-making regimes.

2.3 Optimal Intermediate Layer Weights Depend on Magnitude of Environmental Uncertainty

In what settings might distinct decision regimes in a hierarchical neural system be advantageous? Indeed, if the goal is to maximize reward in an environment without uncertainty, a linear combination of attribute values achieves optimality (Nicholson and Snyder 2007). Yet, many real-world decisions are made in environments where the true value of an attribute is uncertain. For example, when examining a menu to choose a meal at a restaurant, there may be dishes containing ingredients with which one is only vaguely familiar. We therefore investigated the performance of these networks in environments where varying degrees of environmental noise created uncertainty as to the true attribute values.

To do this, we sequentially presented networks with an array of 630 choice alternative combinations, which were translated to firing rates of 10 to 20 Hz, spaced by 1 Hz. The inputs were produced combinatorially, such that on one trial the A attributes might be, [11 Hz, 12 Hz] and B , [13 Hz, 18 Hz]. On the next trial, A could be composed of, [11 Hz, 13 Hz] and B , [12 Hz, 18 Hz], etc. These choices were presented in blocks, where each block had a fixed σ_{η_I} that ranged from 0 to 2 Hz.

For the Linear Network and each weight-configuration of the Hierarchical Network, we calculated the average amount of reward-per-trial. Heatmaps of the reward-per-trial for the Hierarchical Network are shown in Figure 4A, with the performance of the Linear Network indicated by the black dot on the heatmap colorbar. The lower-bound of the heatmap colorbar indicates the minimum possible reward-per-trial (set to 0), and the upper-bound indicates the maximum possible reward-per-trial (set to 1). Weight-configurations with regime I indifference curves are outlined in white. The region to the upper left of the white outlines corresponds to regime II, and to the right to regime III. On the left of Figure 4A is performance in a certain environment, and on the right performance in an environment with a high amount of environmental uncertainty. In both heatmaps, the optimally performing weight-configuration is indicated by a black “X.” The shift in the location of optimal weight-configuration is shown with more detail in Figure 4B, which plots

Figure 4: Hierarchical Network E/I Tone Governs Reward Performance Under Varying Uncertainty. (A) Heatmaps of the average reward-per-trial for different configurations of intermediate layer E/I weights. The X-axis is the intermediate layer inhibitory weight, and the Y-axis is the excitatory weight. On the heatmap colorbar, the 0 value is the minimum possible reward-per-trial, while 1 is the maximum possible reward-per-trial. The dot on the colorbar indicates the performance of the Linear Network. Weight-configurations in decision regime I are outlined in white on the heatmap. Regime II weight-configurations are to the upper left of those areas, and regime III are to the right. The weight-configuration achieving the most reward is indicated by the black “X.” The left heatmap of (A) shows performance in the absence of environmental uncertainty ($\sigma_{\eta_I} = 0$ Hz), while the right heatmap shows performance in the presence of high uncertainty ($\sigma_{\eta_I} = 1.5$ Hz). (B) The optimal E/I weights shift as uncertainty increases. The X-axis is environmental uncertainty (σ_{η_I}), the left Y-axis shows the optimal excitatory weight and the right Y-axis shows the optimal inhibitory weight. (C) The shift from regime I (tan) to regime III (sky blue, colors as in Figure 3C). The level of uncertainty is on the X-axis and the curvature of the fit to the indifference curve is on the Y-axis. (D) A measure of the change in reward-per-trial as a function of noise for a regime I ($J_+ = 0.34$, $J_- = -0.01$), regime II ($J_+ = 0.36$, $J_- = 0.0$) and regime III ($J_+ = 0.35$, $J_- = -0.10$) weight-configuration. The magnitude of the slope indicates the susceptibility to uncertainty. (E) The slopes of lines fit in (D) for all weight configurations. The slope for the Linear Network is again indicated by the black dot.

the optimal excitatory weight (magenta) and inhibitory weight (turquoise) across different magnitudes of environmental uncertainty. Note that the optimal excitatory weight fluctuates in a narrow range (0.32 to 0.34 nA), while the optimal inhibitory weight magnitude increases monotonically (-0.01 to -0.1 nA). The shape of the indifference curve produced by these weights is shown in Figure 4C, where the decision regime starts in regime I (linear), and then moves to regime III, becoming increasingly concave as the environment becomes more uncertain. Thus, the Hierarchical Network displays distinct decision regimes that maximize reward under a variety of environmental uncertainty levels. Furthermore, the shift in regimes with uncertainty is primarily achieved through increasing inhibitory tone.

We then measured how robust different weight-configurations of the Hierarchical Network were to changes in the level of uncertainty. For each weight-configuration, we fit a line to the proportion of maximum reward as the uncertainty level increased, as shown in Figure 4D. The magnitude of the slope was taken as the susceptibility, with a greater slope indicating that the weight-configuration was less robust to uncertainty changes. These slopes are shown in a heatmap on Figure 4E, with the slope of the Linear Network indicated again with a black dot on the colorbar. We found that the regime III areas, in addition to displaying greater performance with high uncertainty, were also more robust to changes in uncertainty levels.

2.4 Regime II Results from Nonlinear Summation, While Regime III Results from Max-like Operation

Having established that the configuration of E/I weights defines decision-making regimes, and that these decision-making regimes are optimal depending on the environment, we next examined the functional transformation taking place within an intermediate layer attribute-specific area, and we analytically investigated the conditions for which different regimes are optimal.

The phase planes of dynamical systems are a highly useful tool for analyzing the functional properties of decision networks (Wong and Wang 2006; Wong, Huk, et al. 2007). For an intermediate layer area, each combination of weight-configurations and input levels produced a unique phase plane. We calculated phase planes for all combinations of an intermediate layer area’s excitatory weights (0.30 nA to 0.40 nA, spaced 0.05 nA) and inhibitory weights (0.0 nA to -0.10 nA, spaced 0.05 nA), and all input combinations of 0 to 40 Hz, spaced 0.5 Hz. This produced 793,881 total phase planes. For each phase plane, we computed a noiseless trajectory starting near the origin. The endpoint of that trajectory was taken as the output of the specific combination of network structure and inputs. The transformation of an input to output for different weight-configurations is shown in Figure 5A, with the input on the left and the outputs of three different weight-configurations on the right. We used these fixed point trajectory endpoints to approximate $T_{A,1}$,

$T_{A,2}$, $T_{B,1}$, and $T_{B,2}$. The decision was determined by linearly summing these T values to obtain C_A and C_B , then choosing the maximum of these two choice values.

Though the intermediate layer outputs to the choice area in the full dynamical model often did not reach the fixed points (due to the decision threshold), the fixed points did a reasonable job approximating behavior. To measure this, for each weight-configuration we fit a sigmoid to the choice behavior of the full model as a function of the difference between choice alternative fixed points (see Methods). Supplementary Figure S2 shows the distribution of R^2 values. As over 90% of weight-configurations had an $R^2 > 0.70$, we felt confident in use of fixed points as an approximation for analysis of intermediate layer function.

Figure 5: Functional Properties of Hierarchical Network Decision Regimes. I : input; $\hat{I}$: input normalized by the maximum input value; T : the transformed output of an intermediate layer; J_+ : excitatory weight; J_- : inhibitory weight; red: choice alternative A; blue: choice alternative B; green: trajectory from a starting point of (0.06, 0.06); solid circles: stable fixed points; large solid circle: trajectory endpoint. (A) The input to an intermediate layer area is transformed to a fixed point output. The input values of $I_{A,x} = 15$ Hz and $I_{B,x} = 10$ Hz are presented graphically as coordinates on the left. These inputs were fed into intermediate layer areas (arrows) with unique weight-configurations. The resultant phase portraits of the transform layer are shown on the right and below for three weight-configurations. (B) The weight-space was partitioned according to the decision regimes arrived at through fixed point approximations. Indifference curves were computed for each weight-configuration by transforming offer input values using the endpoints of their phase-portrait trajectories, summing them to compute choice values C_A and C_B , then taking the larger of those values. The color scheme is the same as Figure 3C. (C) The input-output of a single attribute value as it is varied ($I_{A,x}$) while the other ($I_{B,x}$) is fixed for a regime II weight-configuration (top) and a regime III weight-configuration (bottom). The difference between the normalized fixed ($\hat{I}_{B,x}$) and the varied value ($\hat{I}_{A,x}$) is on the X-axis. The T value of the endpoints is on the Y-axis. Individual $T_{A,x}$ (red) and $T_{B,x}$ (blue) values are plotted, along with lines joining them to aid visualization.

We then used fixed points to compute indifference curves, similar to those shown in Figure 3B. In Figure 5B, we show the decision-making regimes as computed from the intermediate layer fixed points (as in Figure 3C). Its qualitatively similar form to that of Figure 3C further supports that the fixed points provide a reasonable approximation of the full network.

Confident in the representativeness of the intermediate layer fixed points, we used them to investigate

the input-output relationship of various intermediate layer weight-configurations. To do this, we first looked at how the output changes as one input is increased while all else is held constant. We then assessed the how multiple attributes interact as they pass through the intermediate layer.

To investigate how output changes as only one input is varied, we kept one of the attribute inputs ($I_{A,x}$) to an intermediate layer area fixed, while increasing the level of the other attribute input ($I_{B,x}$). The resultant $T_{A,x}$ and $T_{B,x}$ for a region II weight-configuration ($J_+ = 0.34$ nA, $J_- = -0.01$ nA), and a region III weight-configuration ($J_+ = 0.32$ nA, $J_- = -0.06$) are shown in Figure 5C. Note that in regime II (Figure 5C, top), the constant-input does not change as the magnitude of the variable input increases. Furthermore, the output level of the varied-attribute (outlined in gray) in regime II shows a decreasing rate of change as its magnitude grows. The same plot for regime III is shown in Figure 5C, on the bottom. In regime III, when the variable input is less than the fixed input, the output-value of the variable input is suppressed; however, when the variable input is greater than the fixed input, the output-value of the constant input is suppressed, while the output-value of the varied input grows approximately linearly. Thus, in region I and II, each input value passes through the intermediate layer whose degree of linearity is defined by the weight-configuration (Supplemental Figure S3A). In region III, however, a qualitatively different function is implemented such that the smaller input value is suppressed by a max-like operation (Supplemental Figure S3B).

Having established that different decision-making regimes maximize reward at different uncertainty levels, and then how the network produces the regimes, the natural question arises as to why these regimes are optimal in different environments. To answer this question, we turned to a mathematical analysis of simplified non-dynamical models. One model represented the extreme case of sequential max-operations (corresponding with regime III) and the other model that of a linear sum (regime I). We found that when the two choice alternatives are extremely dissimilar, sequential max operations can lead to sub-optimal choices. However, when the choice alternatives are similar, the max operations allow a network to improve the signal relative to noise (see Supplemental Information for proofs).

2.5 Choice Patterns of Those With Autism will Differ from Neurotypicals as Environment Uncertainty is Varied

Having shown that the Hierarchical Network can adopt multiple decision-making regimes, that the regimes are optimal in different environments, and how E/I tone defines regime-specific attribute transforms, we next hypothesized a framework where E/I tone can be modulated according to the degree of environmental uncertainty. We then applied this framework to investigate altered decision-making in ASD.

ASD is a developmental disorder with strong evidence for inhibitory dysfunction (Foss-Feig et al. 2017;

Figure 6: Task Simulation with Neurotypical and ASD Models. (A) The two-alternative, two-attribute task with varying levels of uncertainty. Alternatives, consisting of the attributes M&Ms and Jelly Beans, are simultaneously presented. The offer value of an attribute is indicated by the position of the horizontal white line on the vertical rectangle such that the higher the line, the greater the quantity of the attribute. Environmental uncertainty is controlled through a yellow bar obscuring the exact location of the white line (right trial). (B) A module was added (forest green) that modulated the intermediate layer E/I tone according to the uncertainty level. For the ASD model, inhibitory tone was restricted to a maximum magnitude of -0.01 nA, as is shown by the magenta dots in intermediate layer diagram. (C) A single session of the simulation, where the trial number is given on the X-axis and the proportion of trials where the larger option was chosen given on the Y-axis. The smoothed value for the neurotypical and ASD performance is given in black and magenta respectively. Blocks with no environmental uncertainty are in dark gray, medium uncertainty in light gray, and high uncertainty in white. (D) The reward-per-trial as a function of uncertainty for both the neurotypical and ASD models, with error bars drawn from sets of sessions. (E) The optimal indifference curvature as a function of uncertainty for the neurotypical and ASD models. To the right are example indifference curves for neurotypical (left) and ASD (right) at their maximum level of concavity. Input values ($\hat{I}$) are normalized by the reference choice alternative values' linear sum.

Lee, Lee, and Kim 2017; Zikopoulos and Barbas 2013). Although E/I alterations observed in animal models of ASD are highly diverse, one of the more consistent findings is alterations involving parvalbumin positive (PV) interneurons, which have a strong influence on the overall inhibitory tone of a circuit. Gogolla et al. found that PV interneurons were reduced across several areas of cortex and the hippocampus in a diverse set of ASD mouse models (Gogolla et al. 2009). This selective reduction in PV interneurons was similarly found in the cortex of human postmortem tissue, where only the proportion of PV interneurons was decreased relative to controls (Hashemi et al. 2017). Decreases in PV populations are behaviorally relevant for ASD. For example, mice specifically engineered for PV depletion display phenotypes strongly linked with ASD, such as abnormal reciprocal social interactions, impairments in communication, and repetitive and stereotyped patterns of behavior (Wöhr et al. 2015). These and other findings have provided strong support for the “E/I imbalance,” theory of autism, where inhibitory tone is restricted relative to excitatory tone (Foss-Feig et al. 2017). Thus, manipulations of E/I tone in the hierarchical model can provide an avenue for linking the biology of ASD with behavior.

In addition to restricted inhibitory tone in ASD, there is evidence for a general intolerance of uncertainty (Boulter et al. 2014; Fujino et al. 2017; Vasa et al. 2018), and altered performance with environmental uncertainty (Milne et al. 2002; Spencer et al. 2000). Furthermore, differences have been found in multi-attribute decision tasks (Foxy et al. 2015; Zaidel, Goin-Kochel, and Angelaki 2015). We therefore created a multi-attribute decision task with varying uncertainty as to the true value of attributes. The task, designed for pediatric subjects, involves choosing between baskets of food for “Harry the Hippo.” Harry’s favorite foods are M&Ms and Jelly Beans, which he enjoys equally. On a given trial, the network was presented with two baskets to feed Harry, each composed of M&Ms and Jelly Beans (Figure 6A). In the figure, quantities of the attributes are indicated by a horizontal lines on a vertical bars associated with each attribute. The goal over the course of the session is to feed Harry as much food as possible. Uncertainty is created by obscuring the exact position of the line with a yellow bar (right trial in Figure 6A); the wider the yellow bar, the more uncertain the attribute’s true value.

The E/I tone of a brain area is not necessarily fixed. Indeed, there are several mechanisms by which the tone can be temporarily modulated (Zucker 1989; Semyanov et al. 2004; Avery and Krichmar 2017). The cholinergic diffuse modulatory system, in particular, provides a plausible mechanism for the source of E/I tone shift due to expected environmental uncertainty (Yu and Dayan 2005).

To allow the network to adapt to uncertainty levels, we added a module that used a look-up table to identify the optimal intermediate layer weight-configuration for a given environmental uncertainty level, then shifted weights in the attribute-specific intermediate layers to that configuration (green in Figure 6B). The neurotypical model was able to explore the full range of possible weights. The ASD model was implemented

by limiting the possible inhibitory tone in intermediate attribute-specific areas to a maximum magnitude of -0.01 nA.

Sessions were composed of trial blocks with variable levels of uncertainty. A single session is presented in Figure 6C, with the smoothed proportion of the larger value chosen across trials shown for the neurotypical (black) and ASD (magenta) models. Note that during high-uncertainty blocks (white), one can already see a separation of performance levels. Analyzing food-per-trial as a function of uncertainty, we found that the ASD model did not differ from the neurotypical model when the uncertainty level was low, but began to diverge as uncertainty increased (Figure 6D). We also examined the change in the indifference curve as a function of uncertainty. The trend is shown in the left portion of Figure 6E, while specific instances of neurotypical and ASD indifference curves at the highest level are shown to the right. We found that there was a ceiling to the concavity of the ASD model’s decision behavior. Thus, the model makes two strong experimental predictions: 1) performance will be identical at low uncertainty, while subjects with ASD will show a greater fall-off in total reward as uncertainty increases; and 2) at high uncertainty, subjects with ASD will continue to show more linear-weighting of attribute values than neurotypicals, who shift to attending to the larger of the attribute values.

3 Discussion

We showed that a hierarchical network of areas defined by E/I tone is capable of performing a multi-attribute decision task, and can adopt behavioral regimes that maximize reward in a variety of environments. In environments where attribution values are certain, the network maximizes reward by adopting a regime that sums a linear translation of offer attribute values (regime I). However, as the environment becomes more uncertain, the network maximizes reward by implementing a max-like function to filter noise (regime III). We then proposed a framework where E/I tone can be modulated according to uncertainty. We used this framework to predict that, due to an imbalance of inhibition relative to excitation, subjects with ASD will diverge in behavior from neurotypicals as environmental uncertainty increases.

There is support in the literature for both the hierarchical and parallel structure of our model. Several prior computational models have proposed a similar parallel processing of attributes that are then joined in areas more proximal to the final decision (Roe, Busemeyer, and Townsend 2001; Balasubramani, Moreno-Bote, and Hayden 2018; Hunt, Dolan, and Behrens 2014; Cisek 2006; Christopoulos, Bonaiuto, and Andersen 2015). Furthermore, recent studies involving humans directly examined attribute-specific processing prior to integration in downstream areas. (Berker et al. 2019) looked at the representation of quantity and quality attributes in gift cards. They found that distinct areas were independently specific for each attribute, and

that a third area showed activity correlated with the combined signals from the earlier areas. The importance of later-areas for the integration of value signals was demonstrated by (Pelletier and Fellows 2019) in patients with damage to the OFC and vmPFC. They found that these patients were able to make decisions based on single attributes (the form of elements in an object versus the configuration of elements) no different from controls. However, when patients had to make decisions based on integration of the attribute signals, they showed deficits. These results are expected from a neural system with parallel structure that converges for integration of attribute information.

A down-stream area providing a common representation of attribute value information is a key component of our model. Several electrophysiology studies have specifically investigated the kind of two-alternative multi-attribute choices that we have examined, while recording from down-stream areas relevant for decisions such as the OFC or ACC. Reward magnitude, probability, information value, social hierarchy and juice type have all been found to be represented in these areas, and in a manner where all attribute signals were available for decision-making (Hunt, Malalasekera, et al. 2018; Pastor-Bernier, Plott, and Schultz 2017; Pastor-Bernier and Cisek 2011; Blanchard, Hayden, and Bromberg-Martin 2015; Cai and Padoa-Schioppa 2014; Raghuraman and Padoa-Schioppa 2014; Cai and Padoa-Schioppa 2012; Xie and Padoa-Schioppa 2016; Padoa-Schioppa and Assad 2006; Padoa-Schioppa 2009; Munuera, Rigotti, and Salzman 2018; Pastor-Bernier, Stasiak, and Schultz 2019). However, it is important to recognize that the stimuli used to signify attribute information in these studies were unambiguous, and that NHP subjects were over-trained on the use of attributes to optimize reward. Thus, these studies all take place in the “low uncertainty,” range of our results. A quasi-linear attribute value transmission is therefore expected. Our results suggest that if environmental uncertainty is introduced, one should observe an over-representation of the larger attribute value in areas such as the OFC or ACC, and that the activity will correlate with animal behavior. This is a distinct prediction, highly feasible for experimentation.

Though the biology of how E/I tone of specific areas can be tuned according to task demands remains to be fully characterized by experiments, diffuse modulatory neurotransmitter systems provide the most plausible mechanism. Dopamine, norepinephrine, serotonin, nitric oxide and acetylcholine (ACh) have all been shown to influence the E/I tone of the cells that they target (Avery and Krichmar 2017). These influences can be complex, with cell-type specific shifts in excitation or inhibition implemented via a variety of mechanisms (metabotropic receptors, ion channels, etc.). The ACh system in particular displays the area-specific targeting (Zaborszky et al. 2015), the modulation of inhibitory tone (Disney, Aoki, and Hawken 2012; Picciotto, Higley, and Mineur 2012; Herrero et al. 2008; Kang, Huppé-Gourgues, and Vaucher 2014; Phillips et al. 2000; Sarter, Lustig, Howe, et al. 2014; Sarter, Lustig, Berry, et al. 2016; Thomsen, Sørensen, and Dencker 2018; Disney and Aoki 2008; Disney, Domakonda, and Aoki 2006; Disney, Alasady, and Reynolds

2014; Disney and Reynolds 2014) and the activity associated with known uncertainty (Monosov, Leopold, and Hikosaka 2015; Voytko et al. 1994; Gill, Sarter, and Givens 2000; Dalley et al. 2004; Marshall et al. 2016) that is required by our model. However, it remains to be investigated precisely how the brain can identify an task-optimal E/I tone for a circuit and then shift the circuit to that value.

Our work makes additional strong predictions for decisions of ASD subjects performing a multi-attribute decision task as uncertainty is varied. There are a number of existing experimental results in the field of ASD research that suggest our predictions are reasonable and worth investigating. Several studies show that when attribute information is clear, performance between ASD and neurotypical subjects is identical, while it differs in the face of noise or uncertainty. Dot field motion tasks, in particular, have been used to study these questions in subjects with ASD. The general finding has been that when coherence of the moving dot field is high (uncertainty is low), subjects with ASD and neurotypicals show similar performance (Milne et al. 2002; Spencer et al. 2000; Zaidel, Goin-Kochel, and Angelaki 2015). Indeed, in detection of velocity, and some coherence studies, subjects with ASD outperform neurotypicals (Chen, Norton, et al. 2012). However, when coherence levels decrease (uncertainty increases), subjects with ASD consistently show significantly greater decreases in performance. Similar effects of environmental uncertainty have been shown in ASD multi-attribute studies (Zaidel, Goin-Kochel, and Angelaki 2015; Foxe et al. 2015). Our results suggest a plausible mechanism for the effects observed in prior studies, while also proposing a specific experiment with falsifiable findings.

In this paper, we used a hierarchical neural network to investigate how biophysical properties, such as parallel structure and E/I tone, shape multi-attribute, multi-alternative decisions. This investigation produced strong predictions for electrophysiology experiments, as well as experiments involving human behavior. Our work thus suggests a new avenue for research connecting neural circuits to multi-attribute economic choice and to deficits associated with autism.

4 Acknowledgments

This research was partly supported by the National Institutes of Health (NIH) grant 062349, and the Simons Collaboration on the Global Brain program grant 543057SPI To XJW.

5 Methods

5.1 Model Architecture

5.1.1 Cortical Areas

Areas in the networks were composed of mean-field attractors that have been used to model working memory, perceptual decision-making, as well as to model individual areas in multi-area models of cortex (Wong and Wang 2006; Wong, Huk, et al. 2007; Murray, Jaramillo, and Wang 2017; Mejias et al. 2016). Each area consisted of two populations, such that $c = A, B$. The dynamics of a given population is described by a single synaptic gating variable representing the fraction of activated N-methyl-D-aspartate receptor (NMDA) conductance and governed by the equation (Wong and Wang 2006),

$$\frac{dS}{dt} = -\frac{S}{\tau_{NMDA}} + \gamma(1 - S)r, \quad (1)$$

where the NMDA time constant $\tau_{NMDA} = 60$ ms, and the rate of saturation is controlled by $\gamma = 0.641$. The firing rate r was a function of the input current I , as defined by the curve relation first described in (Abbott and Chance 2005):

$$r = F(I) = \frac{aI - b}{1 - \exp[-d(aI - b)]}, \quad (2)$$

where $a = 270$ Hz/nA, $b = 108$ Hz, and $d = 0.154$ seconds. The total synaptic current consisted of recurrent (I_{rec}), noisy (I_{noise}), background ($I_o = 0.3297$ nA) and external components such that,

$$I = I_{rec} + I_{noise} + I_o + gI_{input}, \quad (3)$$

where the coupling constant $g = 0.0011$ nA/Hz converts the firing rate I_{input} to current. As described in section 5.1.2, specific values of I_{input} are notated as $I_{A,1}$, $I_{A,2}$, $I_{B,1}$, or $I_{B,2}$. For a given population c in area n of the the network, the recurrent current was given by the equation,

$$I_{rec,i}^n = \sum_{m,j} S_j^m J_{ij}^{(m \rightarrow n)}, \quad (4)$$

where $J_{ij}^{(m \rightarrow n)}$ is the strength of the connection from population j in area m to population i in area n . The use of this equation in the Linear Network and Hierarchical Network is described in the following sections.

The noise currents for each population were defined by an Ornstein-Uhlenbeck process:

$$\tau_{AMPA} \frac{dI_{noise}(t)}{dt} = -I_{noise}(t) + \eta_I(t) \sqrt{\tau_{AMPA} \sigma_{noise}^2}, \quad (5)$$

where the time constant $\tau_{AMPA} = 2$ ms, and η_I is a Gaussian white noise term with mean zero, and variance $\sigma_{noise}^2 = 0.003$ nA.

5.1.2 Inputs

External inputs were provided as firing rates, which could span the range of 0 to 40 Hz. Notationally, they are given by I_{input} , such that the input of attribute 1 for choice alternative A was written $I_{A,1}$, etc. An unspecified choice alternative is indicated with the population subscript “ c ” and and unspecified attribute by the subscript “ x .”

When varying environmental uncertainty, a value η_I was randomly drawn from $\mathcal{N}(0, \sigma_\eta^2)$ independently for each attribute on each trial, such that $\tilde{I}_{c,x} = I_{c,x} + \eta_I$. The standard deviation σ_{η_I} was used to control the amount of environmental uncertainty in a given block of trials.

In several figures, we refer to a normalized value of the input, $\hat{I}_{c,x}$. In those cases, the quantity is specified with respect to which the input is normalized.

5.1.3 Linear Network

The Linear Network consisted of two layers. The first layer provided inputs from the individual attributes, in the form of firing rates I_{input} . The attribute inputs to a given population were linearly summed, such that

$$I_c = w \sum_a I_{c,a}, \quad (6)$$

where $w = 0.5$. These values were then fed into the choice area described in section 5.1.1. As there was only one area, equation 4 simplified such that,

$$I_i = \sum_j S_j J_{ij}, \quad (7)$$

where for $i = j$, the connection strength was 0.3725 nA, and where $i \neq j$ the connection strength was -0.1137 nA (Wong and Wang 2006). The decision was determined when one of the choice area populations passed the firing rate threshold of 35 Hz.

5.1.4 Hierarchical Networks

The input layer to the Hierarchical Networks was also in Hz, and specified by $I_{c,a}$. In the Hierarchical Network, however, inputs were segregated by attribute and fed into an intermediate transform layer, consisting of a two-population area specific for each attribute. These areas were structured along the lines described in section 5.1.1.

In the intermediate layer areas of the Hierarchical Network, equation 4 was defined as follows. J_+ controlled the strength of the excitatory connection of a population to itself ($A \rightarrow A$, etc.), and J_- , the strength of the inhibitory connection from one population to another ($A \rightarrow B$, etc.). The strength of J between areas in the intermediate layer areas was set to 0. Connectivity from the intermediate layer (IL) areas to the choice layer (CL) was restricted to excitation between populations with the same selectivity, such that $J_{AA}^{(IL \rightarrow CL)} = 0.25$, $J_{AB}^{(IL \rightarrow CL)} = 0.0$, etc.

The excitatory and inhibitory weights in this intermediate layer (J_+ and J_-) were parametrically varied to define the specific Hierarchical Network weight-configurations described in the Results section. For simplicity, weights in all areas in the processing layer were symmetrically varied, such that when the recurrent excitation or cross-inhibition were changed, all areas in the processing layer assumed the new values. The outputs from this processing layer were dynamically fed into a choice area of the same type as the choice area of the Linear Network in section 5.1.1. To clarify notation, when an input r is passed through an intermediate transform layer area, its associated transformed output is referred to as T . The value T is that of the synaptic gating variable from equation 1.

5.2 Model Behavior

5.2.1 Psychometric Curves

For the psychometric curves shown in Figure 2, a single choice alternative attribute was varied while all others were held constant. 1,000 trials were run for the Linear Network, and for the Hierarchical Network. We then calculated the percent of those trials where the networks chose the option associated with the varied attribute. To these points, we fit a sigmoid of the form:

$$P(c) = \frac{1}{1 + \exp(-k(I_{c,x} - \mu))}, \quad (8)$$

where $P(c)$ represents the proportion of the time the varied choice alternative was chosen, k dictates the slope of the sigmoid, $I_{c,x}$ indicates the value of the varied attribute, and μ provides the centering of the sigmoid.

5.2.2 Indifference Curves

The indifference curves of Figure 3C and 5B were calculated by holding the values for choice alternative A ($I_{A,1}$ and $I_{A,2}$) constant as a reference, while independently varying $I_{B,1}$ and $I_{B,2}$. The reference attribute inputs were both 20 Hz, and the varied attribute inputs ranged from 0 Hz to 40 Hz. All possible combinations of the varied attributes were then simulated for 1,000 trials, where the networks chose between the varied alternative and the reference. This was done for the Linear Network, as well as for several processing-layer weight-configurations of the Hierarchical Network. The processing-layer weight-configurations were all possible combinations of the recurrent excitatory weights 0.30 to 0.40 nA, spaced 0.01 nA with the cross-inhibitory weights 0.0 to -0.1 nA, spaced 0.01 nA. Indifference values were then taken to be those where the network chose evenly between the alternatives.

To fit an indifference curve, we first normalized the indifference values such that the minimum in either direction was 0, and the maximum was 1. We then fit the equation,

$$I_{A,2} = (1 - I_{A,1}^a)^{1/a}, \quad (9)$$

where a was used to define the shape of the indifference curve. If $a < 1$, the curve was classified as convex. If $a = 1$, it was classified linear, and if $a > 1$ it was classified as concave. These classifications were then used in Figures 3C and 5B to partition the weight space into areas that produced convex, linear or concave indifference curves.

5.2.3 Reward-Per-Trial

To simulate a choice experiment, we computed the choice behavior of the network in response to $I_{c,x}$ values in all possible combinations of the inputs ranging from 10 to 20 Hz, with a step size of 1 Hz, producing 630 possible combinations. This was done for the Linear Network, as well as for the Hierarchical Network with all combinations of the recurrent excitatory weights ranging from 0.30 nA to 0.36 nA, spaced 0.01 nA, with the cross inhibitory weights ranging from 0.0 nA to -0.1 nA, spaced 0.01 nA. Each choice was run for 1,000 trials, from which the percentage chosen was calculated.

Reward was considered to be the total value of the attributes for the chosen alternative, regardless if it was the larger of the two. Thus, on each trial there was a maximum reward (the larger choice alternative) and a minimum reward (the smaller choice alternative). The reward-per-trial was computed by simply averaging the reward received across all trials. We did the same for the minimum and maximum reward. We then set 0 to be the minimum reward-per-trial and 1 to be the maximum reward-per-trial. Those values were used to bound the colorbar in Figure 4A.

5.2.4 Uncertainty Robustness

To compute the robustness to uncertainty, as shown in Figure 4E, we systematically stepped the σ_{η_I} as described in section 5.1.2 through the values 0 to 0.5 Hz with a step size of 0.1, as well as 0.75 to 1.0 Hz with a step size of 0.25, and 1.5 to 2 Hz with a step size of 0.5. For the Linear Network and each weight-configuration, the the reward-per-trial was calculated at each of these uncertainty levels. A linear function was then fit to the reward-per-trial as a function of uncertainty level of the form,

$$RPT = a\sigma_{\eta_I} + b \quad (10)$$

Where RPT is the reward-per-trial, a is the slope and b the intercept. The robustness to uncertainty was assessed using the magnitude of a , such that the greater the magnitude, the more susceptible.

5.3 Fixed Point Analysis

To compute the fixed points and phase portraits of a processing layer area, we used the python package pydstool (Clewley 2012), parameterized by the equations described in 5.1.1, with the internal noise was set to 0. We first defined an attribute input as a coordinate in euclidean space space whose positions are defined by the level of each input (Figure 5A, left). We then defined a phase plane by applying these inputs to each weight-configuration. A starting position in each phase plane was selected near the origin at (0.06, 0.06), and a noiseless trajectory was computed. The end point of that trajectory was taken as the output of the processing layer.

5.3.1 Calculation of Convexity

The trajectory endpoints were computed for all inputs and weight-configurations described in section 5.2.2. These endpoints were treated as the values T . They were then linearly summed to determine the value of a choice alternative, such that,

$$C_c = \sum_a T_{c,x}. \quad (11)$$

The indifference point was taken to be when the $C_A \approx C_B$. Convexity was then computed and the weight space was partitioned using the methods described in section 5.2.2.

5.3.2 Calculation of Functional I/O: Linearity

To determine the functional transformation of inputs as they passed through the processing layer, for each weight-configuration we kept $I_{B,x}$ constant at 20 Hz while increasing $I_{A,x}$ from 0 to 40 Hz. We then recorded the resultant $T_{A,x}$ and $T_{B,x}$ as the output of the functional transformation. We identified when the transition of $T_{A,x}$ stabilized and fit the function,

$$T_{A,x} = \alpha I_{A,x}^h + \beta, \quad (12)$$

where h was used to determine the linearity of the transformation. An $h \approx 1$ indicated that the network produced an output that was a linear function of the input. If $h < 1$ however, the network produced a sub-linear output. The value of h was plotted in the heatmap of Figure S3A.

5.3.3 Calculation of Functional I/O: Interaction

To calculate the full interaction of values, we computed $T_{A,x}$ as we varied both $I_{A,x}$ and $I_{B,x}$ from 0 to 40 Hz, with a step size of 0.5 Hz. We then plotted the values of $T_{A,x}$ shown in the heatmap of Figure S3B.

5.4 Neurotypical and ASD Models

We introduced a module that used a look-up table to assess the level of uncertainty in the environment, and then set the weight-configuration in the intermediate layer areas to a tone that optimized performance. The neurotypical model was able to utilize the full range of excitation (0.30 nA to 0.40 nA) and inhibition (0.0 nA to -0.1 nA). The ASD model was able to utilize the full range of excitation, but the inhibitory tone was limited to a maximum of -0.01 nA. We then presented the models with blocks of trials at different uncertainty levels, and calculated the average reward-per-trial (section 5.2.3) during each block.

Supplemental Information

1 Analysis of Linear Model and Sequential Max Model

The Hierarchical Network achieves maximum reward in an environment without uncertainty by adopting a weight-configuration that performed linear integration of attribute values (similar to the Linear Network), while in an environment with high uncertainty a max-like operation achieved maximum reward. In this section, we use simplified versions of the model to analytically investigate these observed phenomenon.

1.1 The Linear Model And the Sequential Max Model

We will consider and compare two models of multi-attribute decisions. In the linear model (LM), attribute input values (I) for each choice alternative were linearly combined to reach a choice value C_A or C_B , and then the maximum of those values was taken as the choice (as in Figure S1A). The sequential max model (SM) consisted of a series of max operations (as in Figure S1B-C). These two models are extreme cases of decision regimes adopted by the Hierarchical Network, as described in the main body of the paper. Though we use the same notation as the main body of the paper for clarity (I , C , etc.), the values in this section are unitless.

Figure S1: Static Decision Models. (A) The linear model, where attribute input values (I) are summed into a choice value (C_A or C_B). Choice is then determined with a max operation. (B) The sequential max model with a max operation at the attribute level and the choice level. The surviving values from the first max operation are summed to calculate C . Two examples of the sequential max model are shown, where the larger attribute input values are bolded. On the left, both attribute inputs associated with A are larger, and so these are used to calculate the C_A . Since the B attribute inputs are smaller, they do not survive the first stage, and so $C_B = 0$. On the right, an attribute associated with each choice alternative survives the first stage, and so both C values are nonzero.

Our analysis will focus on the most basic case, where there are two choices composed of two attributes. For notation, we will index the first choice as A and the second choice as B . Note that the number of choices could be expanded arbitrarily. We index the two attributes associated with each choice as 1 and 2. When the two attributes are interchangeable, it will be indicated with the subscript x . An attribute could be the

color of a choice alternative, the monetary value, or (if it is edible), its flavor. As an example, the input to the model of the monetary value (1) associated with choice alternative 1 (A) could be written as $I_{A,1}$. The flavor (2) associated with choice alternative 1 (A) could be written as $I_{A,2}$. We can then write the initial offers as $[I_{A,1}, I_{A,2}]$ and $[I_{B,1}, I_{B,2}]$.

1.1.1 The Linear Model: Equations

The LM was meant to capture the Linear Network, along with the case of the Hierarchical Network model performing a linear addition of the attribute values. The operations of the LM are shown in Figure S1A. In this model, the attributes associated with each choice alternative are first summed, such that,

$$C_A = I_{A,1} + I_{A,2}, \quad \text{and} \quad C_B = I_{B,1} + I_{B,2}. \quad (13)$$

The second decision stage then takes the maximum of these choice values as the choice of the network, consisting of the operation,

$$\max_{choice}(C_A, C_B). \quad (14)$$

For the linear model, we will often thus simplify the decision stage by writing,

$$\max_{choice}(I_{A,1} + I_{A,2}, I_{B,1} + I_{B,2}). \quad (15)$$

1.1.2 The Sequential Max Model: Equations

The SM is named for the series of max operations performed, first at the input level, and then with the final decision. Those operations are shown in Figure S1B-C. This model produces the concave decision behavior described in the main text. The first stage of SM implements operations on the attribute level,

$$\begin{aligned} \max_1(I_{A,1}, I_{B,1}), \\ \max_2(I_{A,2}, I_{B,2}). \end{aligned} \quad (16)$$

The key feature is that *only the maximum of that operation is passed onto the next level*. To gain intuition, we can walk through a couple of examples. First, consider the case where,

$$I_{A,1} > I_{A,2} \quad \text{and} \quad I_{A,1} > I_{B,2}.$$

The decision process is shown in on the left of Figure S1B. Since only the winners of the first max

operation (bolded in Figure S1) are passed to the next stage, the choice values will be,

$$C_A = I_{A,1} + I_{A,2}, \quad \text{and} \quad C_B = 0.$$

We can take another example to further illustrate the point (Figure S1C right). In this case, we are going to define the values such that,

$$I_{A,1} > I_{A,2} \quad \text{and} \quad I_{A,1} < I_{B,2}.$$

After these values pass through equation 16, the choice values are,

$$C_A = I_{B,1}, \quad \text{and} \quad C_B = I_{B,2}.$$

Because of this, the C values and composition will be specified in each examined case.

1.1.3 When Offers are Dissimilar, Max Produces Irrational Decisions

Here we will show that, in an edge case, the LM will choose the larger of the option, while the SM paradoxically chooses the lower.

We will consider the edge case where the value of A is less than the value of B ,

$$I_{A,1} + I_{A,2} < I_{B,1} + I_{B,2}.$$

We will further constrain the values so that,

$$I_{B,x} = I_{B,1} = I_{B,2}, \quad I_{A,1} < I_{B,x}, \quad \text{and} \quad I_{A,1} > I_{B,x}.$$

In the LM, by substituting the values into equation 15, we can specify the decision stage as the operation,

$$\max_{LM,choice}(I_{A,1} + I_{A,2}, 2I_{B,x}).$$

Given the constraints, the LM will choose B , which is the larger of the two offers.

In the SM, by use of equation 16, the first stage will consist of,

$$\max_1(I_{A,1}, I_{B,x}), \quad \text{and} \quad \max_2(I_{A,2}, I_{B,x})$$

.

Again, given the constraints, after the first stage we will have,

$$C_A = I_{A,2}, \quad \text{and} \quad C_B = I_{B,x}.$$

By use of equation 14, the decision stage will be,

$$\max_{SM,choice}(I_{B,2}, I_{B,x})$$

As $I_{A,2} > I_{B,x}$, the SM will choose A , even though the total value of B is greater.

Conclusion 1 When the offers are highly dissimilar, there will be cases where the SM will choose the offer with the smaller total value.

1.1.4 Uncertainty in the Linear Model and Sequential Max Model

In a biological system, uncertainty from noise can arise externally via the stimulus or value representation, or internally via transmission between areas, or internally during the decision-making operation. For comprehensiveness, this section will consider how, at all stages, noise effects the LM and SM frameworks. Here, we are agnostic as to the source of the noise (stimulus, representational, transmission, background, etc.), and simply assume a random process, such that the noise term η is drawn from a uniform distribution with the limits $[-\eta_{max}, \eta_{max}]$. By using a uniform distribution, we can more easily delineate the extreme cases. When there are two stages at which the noise term is added, we will simplify their expression notationally by writing,

$$\hat{\eta} = \eta_1 + \eta_2.$$

we will consider the two-attribute, two-alternative decision paradigm. To make the algebra clearer, we will introduce new variables D , E and scaling term γ such that,

$$D = I_{A,1}, \quad E = I_{B,1}, \quad \gamma E = I_{A,2}, \quad \gamma D = I_{B,2}.$$

We will also stipulate that,

$$D > E.$$

By maintaining the symmetry of the attributes across choice alternatives, we can finely control the magnitude differences, while keeping the choice alternatives similar to one another. This can be better understood by

examining the scaling term γ , which tunes the magnitude of the difference between the choice alternatives. For clarity, we will first write rewrite equation 14 using the values as,

$$\max_{LM,choice}(D + \gamma E, E + \gamma D).$$

It is easy then to see that the larger C will be dictated by the value of γ . If $\gamma > 1$, then B will be larger, if $\gamma < 1$, A will be larger, and if $\gamma = 1$ the total value of each choice alternative will be equal. Adding a combined noise term, we arrive at,

$$\max(D + \gamma E + \hat{\eta}_A, E + \gamma D + \hat{\eta}_B). \quad (17)$$

If we define γ as such that,

$$\gamma D > E,$$

$$D > \gamma E,$$

we can follow the algorithm, and arrive at the SM framework's final decision stage where,

$$\max_{SM,choice}(D + \hat{\eta}_A, \gamma D + \hat{\eta}_B). \quad (18)$$

The challenge then is to understand the relationship between scaling term γ (which controls the absolute difference between choice alternatives) and noise term η .

1.1.5 Signal to Noise in the Linear Model Scale with the Difference Between Attributes

We will show that the noise term will be irrelevant in decision making as long as,

$$\gamma_{LM} > 1 + \frac{\hat{\eta}_{max}}{|I_{A,x} - I_{B,x}|} \quad (19)$$

Rewriting equation 17 so that A is the winner we get,

$$E + \gamma_{LM} D + \hat{\eta}_A > D + \gamma_{LM} E + \hat{\eta}_B.$$

Giving alternative B the maximum benefit of noise we have,

$$E + \gamma_{LM} D - \eta_{max} > D + \gamma_{LM} E + \eta_{max}.$$

Solving for the γ_{LM} terms on one side of the equation we get,

$$\gamma_{LM}(D - E) > (D - E) + (\eta_{max} + \eta_{max}).$$

By defining

$$\hat{\eta}_{max} = 2\eta_{max},$$

and dividing through by $(D - E)$, and then back substituting the I values, we arrive at equation 19. It is important to see here that the minimum value of γ_{LM} depends primarily on the term,

$$\frac{\hat{\eta}_{max}}{|I_{A,x} - I_{B,x}|}.$$

The smaller this term, the smaller γ_{LM} can be.

Conclusion 2 The minimum possible value of scaling term γ_{LM} depends on the difference between $I_{A,x}$ and $I_{B,x}$. Thus, the larger the difference between the attribute values, the smaller the difference need be between the choice alternatives.

1.1.6 Signal to Noise in the Sequential Max Model Scales with Magnitude of the Larger Attribute

It can be shown that the value of γ_{SM} needed to render the noise irrelevant is given by,

$$\gamma_{SM} > 1 + \frac{\alpha \hat{\eta}_{max}}{\max(I_{A,x}, I_{B,x})} \quad (20)$$

Now, it is reasonable to consider the possibility that in a SM, the additional operation creates an additional source of internal noise. We will specify its scale using a variable α where,

$$\alpha > 1.$$

We will start by rewriting equation 18 so that A will be chosen over B . This can be specified by,

$$\gamma_{SM}D - \alpha\eta_{max} > D + \alpha\eta_{max}.$$

Solving for the γ_{SM} terms the equation becomes,

$$\gamma_{SM}D > D + \alpha(\eta_{max} + \eta_{max}).$$

By dividing through by D , making the $\hat{\eta}_{max}$ replacement, and back-substituting the I values, we arrive at equation 20. The important feature to note is that the minimum value of scaling term γ_{SM} depends on,

$$\frac{\alpha\hat{\eta}_{max}}{\max(I_{A,x}, I_{B,x})}.$$

Unlike the LM, the denominator has only the larger of the attribute values.

Conclusion 3 In the SM, the larger the largest attribute value relative to the noise, the less the magnitude difference between choice alternatives need be.

1.1.7 When Offers are Similar and in the Presence of Noise, Sequential Max Improves Signal

We will show that,

$$\gamma_{LM} - \gamma_{SM} > \frac{\hat{\eta}_{max}(I_{a,c_{larger}}(1 - \alpha) + \alpha I_{a,c_{smaller}})}{I_{a,c_{larger}}|I_{A,x} - I_{B,x}|}, \quad (21)$$

where,

$$I_{a,c_{larger}} = \max(I_{A,x}, I_{B,x}),$$

$$I_{a,c_{smaller}} = \min(I_{A,x}, I_{B,x}),$$

$$I_{a,c_{larger}} = \max(I_{A,x}, I_{B,x}),$$

$$I_{a,c_{smaller}} = \min(I_{A,x}, I_{B,x}).$$

We start by subtracting equation 20 from equation 19.

$$\gamma_{LM} - \gamma_{SM} > \left(1 + \frac{\hat{\eta}_{max}}{D - E}\right) - \left(1 + \frac{\alpha\hat{\eta}_{max}}{D}\right).$$

By expanding the right side of the inequality we get,

$$\gamma_{LM} - \gamma_{SM} > \frac{\hat{\eta}_{max}}{D - E} - \frac{\alpha\hat{\eta}_{max}}{D}.$$

Putting both terms on the right side over the same denominator produces equation 21. Because $D > E$

by definition, the denominator on the right will always be positive. However, for the inequality to hold, we need to consider what must be true of the values of D , E , a , γ_{LM} and γ_{SM} . By doing so, we will be able to draw some interesting conclusions regarding the two frameworks. First, let's consider the case where,

$$aE > D(1 - a).$$

This means that the numerator on the right side of the inequality will be positive. Thus for equation 21 to be true,

$$\gamma_{LM} - \gamma_{SM} > 0,$$

and therefore,

$$\gamma_{LM} > \gamma_{SM}.$$

Conclusion 4 When the two attributes are similar, it will require a greater difference between the choice alternatives (larger scaling term γ) for the LM model to recognize the more valuable choice alternative. If one compares equation 19 with equation 20, in equation 19, the noise is divided by $|I_{A,x} - I_{B,x}|$, while in equation 20 it is divided by $\max(I_{A,x}, I_{B,x})$. In the LM, the closer the attribute values, the less the noise is reduced. In the SM, the smaller attribute value is no longer present; thus the noise term is always reduced. Furthermore, if we consider the total noise in both models to be equal by setting,

$$a = 1,$$

then the numerator simplifies to,

$$\hat{\eta}_{\max} I_{\min(A,B),x},$$

which is always positive. In such a scenario, conclusion 2 will always be true.

1.1.8 If Internal Noise For the Sequential Max Model is Large, The Linear Model is Superior

Now, let us consider the extreme border-case where,

$$E = 0,$$

and thus,

$$0 > \frac{\hat{\eta}_{max}(1-a)}{\max(I_{A,x} - I_{B,x})},$$

By itself this would say nothing about the value of $\gamma_{LM} - \gamma_{SM}$ since it is possible that,

$$0 > \gamma_{LM} - \gamma_{SM} > \frac{\hat{\eta}_{max}(1-a)}{\max(I_{A,x} - I_{B,x})},$$

or,

$$\gamma_{LM} - \gamma_{SM} > 0 > \frac{\hat{\eta}_{max}(1-a)}{\max(I_{A,x} - I_{B,x})}.$$

Yet, as goal is the find the minimum possible value of the γ s such that the inequalities hold, we can look back to equations 19 and 20 for insight. In the border-case, equation 19 becomes,

$$\gamma_{LM} > 1 + \frac{\hat{\eta}_{max}}{\max(I_{A,x} - I_{B,x})}$$

If minimize the value of γ_{LM} that would keep the inequality true, and assume that a is sufficiently large, we get the inequality,

$$\gamma_{SM} > 1 + \frac{a\hat{\eta}_{max}}{\max(I_{A,x} - I_{B,x})} > \gamma_{LM} > 1 + \frac{\hat{\eta}_{max}}{\max(I_{A,x} - I_{B,x})}.$$

And thus,

$$\gamma_{SM} > \gamma_{LM}.$$

Conclusion 5 Due to the additional noise caused by multiple SM neural operations, if the difference between the largest attribute and smallest attribute in each alternative is sufficiently large, the LM framework will require less of a difference between the alternatives in order to choose the greater choice alternative.

2 Supplemental Figures

Figure S2: Intermediate Layer Area Fixed Points Fit to Full Network Behavior. Fixed point trajectory endpoints were used to approximate C_A and C_B . A sigmoid was then fit to the $P(A)$ of the network as a function of fixed point $C_A - C_B$ for each weight-configuration. The fit of a single weight-configuration is shown in A). An R^2 value was then computed for each fit. A histogram of those R^2 values is given in B), and in C) they are organized by weight-configuration.

Figure S3: Functional Transformations by Hierarchical Network Intermediate Layer Areas. (A) Polynomial functions of the form $T_{A,1} = \alpha I_{A,1}^h + \beta$ were fit to the constant-regions outlined in gray on Figure 5C. The value of the h exponent associated with each weight-configuration is indicated by the intensity of the heatmap. (E) The resultant $T_{A,x}$ as $I_{A,x}$ and $I_{B,x}$ are varied for a regime I (left) and a regime III (right) weight-configuration.

Bibliography

References

- Abbott, L. F. and Frances S. Chance (2005). “Drivers and modulators from push-pull and balanced synaptic input”. In: *Progress in Brain Research*. Vol. 149. Cortical Function: a View from the Thalamus. Elsevier, pp. 147–155. DOI: 10.1016/S0079-6123(05)49011-1. URL: <http://www.sciencedirect.com/science/article/pii/S0079612305490111> (visited on 07/11/2018).
- Ahmadian, Yashar and Kenneth D. Miller (2019). “What is the dynamical regime of cerebral cortex?” In: *arXiv:1908.10101 [q-bio]*. arXiv: 1908.10101. URL: <http://arxiv.org/abs/1908.10101> (visited on 09/04/2019).
- American Psychiatric Association, APA (2013). *The Diagnostic and Statistical Manual of Mental Disorders*. 5th. American Psychiatric Publishing, Inc.
- Avery, Michael C. and Jeffrey L. Krichmar (2017). “Neuromodulatory Systems and Their Interactions: A Review of Models, Theories, and Experiments”. In: *Frontiers in Neural Circuits* 11, e108. DOI: 10.3389/fncir.2017.00108. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5744617/> (visited on 11/11/2019).
- Bach, Dominik R. and Raymond J. Dolan (2012). “Knowing how much you don’t know: a neural organization of uncertainty estimates”. In: *Nature Reviews Neuroscience* 13 (8), pp. 572–586. DOI: 10.1038/nrn3289. URL: <https://www.nature.com/articles/nrn3289> (visited on 11/04/2019).
- Balasubramani, Pragathi P., Rubén Moreno-Bote, and Benjamin Y. Hayden (2018). “Using a Simple Neural Network to Delineate Some Principles of Distributed Economic Choice”. In: *Frontiers in Computational Neuroscience* 12, e22. DOI: 10.3389/fncom.2018.00022. URL: <https://www.frontiersin.org/articles/10.3389/fncom.2018.00022/full> (visited on 08/14/2018).
- Berker, Archy O. de et al. (2019). “Computing Value from Quality and Quantity in Human Decision-Making”. In: *Journal of Neuroscience* 39 (1), pp. 163–176. DOI: 10.1523/JNEUROSCI.0706-18.2018. URL: <http://www.jneurosci.org/content/39/1/163> (visited on 01/03/2019).
- Blanchard, Tommy C., Benjamin Y. Hayden, and Ethan S. Bromberg-Martin (2015). “Orbitofrontal Cortex Uses Distinct Codes for Different Choice Attributes in Decisions Motivated by Curiosity”. In: *Neuron* 85 (3), pp. 602–614. DOI: 10.1016/j.neuron.2014.12.050. URL: <http://www.sciencedirect.com/science/article/pii/S0896627314011593> (visited on 10/01/2018).
- Boulter, Christina et al. (2014). “Intolerance of Uncertainty as a Framework for Understanding Anxiety in Children and Adolescents with Autism Spectrum Disorders”. In: *Journal of Autism and Developmental*

- Disorders* 44 (6), pp. 1391–1402. DOI: 10.1007/s10803-013-2001-x. URL: <https://doi.org/10.1007/s10803-013-2001-x> (visited on 08/30/2019).
- Brunel, Nicolas and Xiao-Jing Wang (2001). “Effects of Neuromodulation in a Cortical Network Model of Object Working Memory Dominated by Recurrent Inhibition”. In: *Journal of Computational Neuroscience* 11 (1), pp. 63–85. DOI: 10.1023/A:1011204814320. URL: <https://link.springer.com/article/10.1023/A:1011204814320> (visited on 04/20/2018).
- Cai, Xinying and Camillo Padoa-Schioppa (2012). “Neuronal Encoding of Subjective Value in Dorsal and Ventral Anterior Cingulate Cortex”. In: *Journal of Neuroscience* 32 (11), pp. 3791–3808. DOI: 10.1523/JNEUROSCI.3864-11.2012. URL: <http://www.jneurosci.org/content/32/11/3791> (visited on 08/01/2018).
- (2014). “Contributions of Orbitofrontal and Lateral Prefrontal Cortices to Economic Choice and the Good-to-Action Transformation”. In: *Neuron* 81 (5), pp. 1140–1151. DOI: 10.1016/j.neuron.2014.01.008. URL: <http://www.sciencedirect.com/science/article/pii/S0896627314000130> (visited on 07/26/2018).
- Chen, Xiaomo and Veit Stuphorn (2015). “Sequential selection of economic good and action in medial frontal cortex of macaques during value-based decisions”. In: *eLife* 4. Ed. by Wolfram Schultz, e09418. DOI: 10.7554/eLife.09418. URL: <https://doi.org/10.7554/eLife.09418> (visited on 10/01/2018).
- Chen, Y., D. J. Norton, et al. (2012). “Enhanced local processing of dynamic visual information in autism: Evidence from speed discrimination”. In: *Neuropsychologia* 50 (5), pp. 733–739. DOI: 10.1016/j.neuropsychologia.2012.01.007. URL: <http://www.sciencedirect.com/science/article/pii/S0028393212000218> (visited on 09/04/2019).
- Chib, Vikram S. et al. (2009). “Evidence for a Common Representation of Decision Values for Dissimilar Goods in Human Ventromedial Prefrontal Cortex”. In: *Journal of Neuroscience* 29 (39), pp. 12315–12320. DOI: 10.1523/JNEUROSCI.2575-09.2009. URL: <http://www.jneurosci.org/content/29/39/12315> (visited on 07/25/2018).
- Christopoulos, Vassilios, James Bonaiuto, and Richard A. Andersen (2015). “A Biologically Plausible Computational Theory for Value Integration and Action Selection in Decisions with Competing Alternatives”. In: *PLOS Computational Biology* 11 (3), e1004104. DOI: 10.1371/journal.pcbi.1004104. URL: <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1004104> (visited on 10/04/2018).
- Cisek, Paul (2006). “Integrated Neural Processes for Defining Potential Actions and Deciding between Them: A Computational Model”. In: *Journal of Neuroscience* 26 (38), pp. 9761–9770. DOI: 10.1523/JNEUROSCI.5605-05.2006. URL: <http://www.jneurosci.org/content/26/38/9761> (visited on 10/01/2018).

- Cisek, Paul (2012). “Making decisions through a distributed consensus”. In: *Current Opinion in Neurobiology*. Decision making 22 (6), pp. 927–936. DOI: 10.1016/j.conb.2012.05.007. URL: <http://www.sciencedirect.com/science/article/pii/S0959438812000864> (visited on 08/02/2018).
- Cisek, Paul and John F. Kalaska (2010). “Neural Mechanisms for Interacting with a World Full of Action Choices”. In: *Annual Review of Neuroscience* 33 (1), pp. 269–298. DOI: 10.1146/annurev.neuro.051508.135409. URL: <https://doi.org/10.1146/annurev.neuro.051508.135409> (visited on 10/01/2018).
- Clewley, Robert (2012). “Hybrid Models and Biological Model Reduction with PyDSTool”. In: *PLOS Computational Biology* 8 (8), e1002628. DOI: 10.1371/journal.pcbi.1002628. URL: <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1002628> (visited on 12/18/2019).
- Compte, Albert et al. (2000). “Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model”. In: *Cerebral Cortex* 10 (9), pp. 910–923. DOI: 10.1093/cercor/10.9.910. URL: <https://academic.oup.com/cercor/article/10/9/910/289091> (visited on 10/27/2019).
- Conen, Katherine E. and Camillo Padoa-Schioppa (2015). “Neuronal variability in orbitofrontal cortex during economic decisions”. In: *Journal of Neurophysiology* 114 (3), pp. 1367–1381. DOI: 10.1152/jn.00231.2015. URL: <https://www-physiology-org.proxy.library.nyu.edu/doi/abs/10.1152/jn.00231.2015> (visited on 08/01/2018).
- Dalley, Jeffrey W. et al. (2004). “Cortical Cholinergic Function and Deficits in Visual Attentional Performance in Rats Following 192 IgG-Saporin-induced Lesions of the Medial Prefrontal Cortex”. In: *Cerebral Cortex* 14 (8), pp. 922–932. DOI: 10.1093/cercor/bhh052. URL: <https://academic.oup.com/cercor/article/14/8/922/268350> (visited on 11/21/2019).
- Disney, Anita A., Hussein A. Alasady, and John H. Reynolds (2014). “Muscarinic acetylcholine receptors are expressed by most parvalbumin-immunoreactive neurons in area MT of the macaque”. In: *Brain and Behavior* 4 (3), pp. 431–445. DOI: 10.1002/brb3.225. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/brb3.225> (visited on 11/20/2019).
- Disney, Anita A. and Chiye Aoki (2008). “Muscarinic acetylcholine receptors in macaque V1 are most frequently expressed by parvalbumin-immunoreactive neurons”. In: *Journal of Comparative Neurology* 507 (5), pp. 1748–1762. DOI: 10.1002/cne.21616. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cne.21616> (visited on 11/20/2019).
- Disney, Anita A., Chiye Aoki, and Michael J. Hawken (2012). “Cholinergic suppression of visual responses in primate V1 is mediated by GABAergic inhibition”. In: *Journal of Neurophysiology* 108 (7), pp. 1907–1923. DOI: 10.1152/jn.00188.2012. URL: <https://www.physiology.org/doi/full/10.1152/jn.00188.2012> (visited on 11/20/2019).

- Disney, Anita A., Kunal V. Domakonda, and Chiye Aoki (2006). “Differential expression of muscarinic acetylcholine receptors across excitatory and inhibitory cells in visual cortical areas V1 and V2 of the macaque monkey”. In: *Journal of Comparative Neurology* 499 (1), pp. 49–63. DOI: 10.1002/cne.21096. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cne.21096> (visited on 11/20/2019).
- Disney, Anita A. and John H. Reynolds (2014). “Expression of m1-type muscarinic acetylcholine receptors by parvalbumin-immunoreactive neurons in the primary visual cortex: A comparative study of rat, guinea pig, ferret, macaque, and human”. In: *Journal of Comparative Neurology* 522 (5), pp. 986–1003. DOI: 10.1002/cne.23456. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/cne.23456> (visited on 11/20/2019).
- Foss-Feig, Jennifer H. et al. (2017). “Searching for Cross-Diagnostic Convergence: Neural Mechanisms Governing Excitation and Inhibition Balance in Schizophrenia and Autism Spectrum Disorders”. In: *Biological Psychiatry*. Cortical Excitation-Inhibition Balance and Dysfunction in Psychiatric Disorders 81 (10), pp. 848–861. DOI: 10.1016/j.biopsych.2017.03.005. URL: <http://www.sciencedirect.com/science/article/pii/S0006322317313604> (visited on 08/30/2019).
- Foxe, John J. et al. (2015). “Severe Multisensory Speech Integration Deficits in High-Functioning School-Aged Children with Autism Spectrum Disorder (ASD) and Their Resolution During Early Adolescence”. In: *Cerebral Cortex* 25 (2), pp. 298–312. DOI: 10.1093/cercor/bht213. URL: <http://cercor.oxfordjournals.org.laneproxy.stanford.edu/content/25/2/298> (visited on 11/11/2016).
- Fujino, Junya et al. (2017). “Attitudes toward risk and ambiguity in patients with autism spectrum disorder”. In: *Molecular Autism* 8 (1), e45. DOI: 10.1186/s13229-017-0162-8. URL: <https://doi.org/10.1186/s13229-017-0162-8> (visited on 08/29/2019).
- Gill, T. Michael, Martin Sarter, and Ben Givens (2000). “Sustained Visual Attention Performance-Associated Prefrontal Neuronal Activity: Evidence for Cholinergic Modulation”. In: *Journal of Neuroscience* 20 (12), pp. 4745–4757. DOI: 10.1523/JNEUROSCI.20-12-04745.2000. URL: <https://www.jneurosci.org/content/20/12/4745> (visited on 11/21/2019).
- Gogolla, Nadine et al. (2009). “Common circuit defect of excitatory-inhibitory balance in mouse models of autism”. In: *Journal of Neurodevelopmental Disorders* 1 (2), pp. 172–181. DOI: 10.1007/s11689-009-9023-x. URL: <https://jneurodevdisorders.biomedcentral.com/articles/10.1007/s11689-009-9023-x> (visited on 09/21/2019).
- Hashemi, Ezzat et al. (2017). “The Number of Parvalbumin-Expressing Interneurons Is Decreased in the Prefrontal Cortex in Autism”. In: *Cerebral Cortex* 27 (3), pp. 1931–1943. DOI: 10.1093/cercor/bhw021. URL: <https://academic.oup.com/cercor/article/27/3/1931/3052664> (visited on 06/15/2018).

- Herrero, J. L. et al. (2008). “Acetylcholine contributes through muscarinic receptors to attentional modulation in V1”. In: *Nature* 454 (7208), pp. 1110–1114. DOI: 10.1038/nature07141. URL: <https://www.nature.com/articles/nature07141> (visited on 11/19/2019).
- Hunt, Laurence T., Raymond J. Dolan, and Timothy E. J. Behrens (2014). “Hierarchical competitions subserving multi-attribute choice”. In: *Nature Neuroscience* 17 (11), pp. 1613–1622. DOI: 10.1038/nn.3836. URL: <https://www.nature.com/articles/nn.3836> (visited on 11/16/2018).
- Hunt, Laurence T. and Benjamin Y. Hayden (2017). “A distributed, hierarchical and recurrent framework for reward-based choice”. In: *Nature Reviews Neuroscience* 18 (3), pp. 172–182. DOI: 10.1038/nrn.2017.7. URL: <https://www.nature.com/articles/nrn.2017.7> (visited on 10/01/2018).
- Hunt, Laurence T., W. M. Nishantha Malalasekera, et al. (2018). “Triple dissociation of attention and decision computations across prefrontal cortex”. In: *Nature Neuroscience* 21 (10), pp. 1471–1481. DOI: 10.1038/s41593-018-0239-5. URL: <https://www.nature.com/articles/s41593-018-0239-5> (visited on 10/16/2019).
- Kang, Jun Il, Frédéric Huppé-Gourgues, and Elvire Vaucher (2014). “Boosting visual cortex function and plasticity with acetylcholine to enhance visual perception”. In: *Frontiers in Systems Neuroscience* 8, p. 172. DOI: 10.3389/fnsys.2014.00172. URL: <https://www.frontiersin.org/articles/10.3389/fnsys.2014.00172/full> (visited on 11/19/2019).
- Lee, Eunee, Jiseok Lee, and Eunjoon Kim (2017). “Excitation/Inhibition Imbalance in Animal Models of Autism Spectrum Disorders”. In: *Biological Psychiatry. Cortical Excitation-Inhibition Balance and Dysfunction in Psychiatric Disorders* 81 (10), pp. 838–847. DOI: 10.1016/j.biopsych.2016.05.011. URL: <http://www.sciencedirect.com/science/article/pii/S0006322316323873> (visited on 04/29/2019).
- Levy, Dino J and Paul W Glimcher (2012). “The root of all value: a neural common currency for choice”. In: *Current Opinion in Neurobiology. Decision making* 22 (6), pp. 1027–1038. DOI: 10.1016/j.conb.2012.06.001. URL: <http://www.sciencedirect.com/science/article/pii/S0959438812001006> (visited on 07/18/2018).
- Louie, Kenway and Paul W. Glimcher (2010). “Separating Value from Choice: Delay Discounting Activity in the Lateral Intraparietal Area”. In: *Journal of Neuroscience* 30 (16), pp. 5498–5507. DOI: 10.1523/JNEUROSCI.5742-09.2010. URL: <http://www.jneurosci.org/content/30/16/5498> (visited on 07/25/2018).
- Mariño, Jorge et al. (2005). “Invariant computations in local cortical networks with balanced excitation and inhibition”. In: *Nature Neuroscience* 8 (2), pp. 194–201. DOI: 10.1038/nn1391. URL: <https://www.nature.com/articles/nn1391> (visited on 11/06/2019).

- Marshall, Louise et al. (2016). “Pharmacological Fingerprints of Contextual Uncertainty”. In: *PLOS Biology* 14 (11), e1002575. DOI: 10.1371/journal.pbio.1002575. URL: <https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.1002575> (visited on 11/21/2019).
- Mejias, Jorge F. et al. (2016). “Feedforward and feedback frequency-dependent interactions in a large-scale laminar network of the primate cortex”. In: *Science Advances* 2 (11), e1601335. DOI: 10.1126/sciadv.1601335. URL: <http://advances.sciencemag.org/content/2/11/e1601335> (visited on 06/11/2018).
- Milne, Elizabeth et al. (2002). “High motion coherence thresholds in children with autism”. In: *Journal of Child Psychology and Psychiatry* 43 (2), pp. 255–263. DOI: 10.1111/1469-7610.00018. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/1469-7610.00018> (visited on 09/03/2019).
- Monosov, Ilya E., David A. Leopold, and Okihide Hikosaka (2015). “Neurons in the Primate Medial Basal Forebrain Signal Combined Information about Reward Uncertainty, Value, and Punishment Anticipation”. In: *Journal of Neuroscience* 35 (19), pp. 7443–7459. DOI: 10.1523/JNEUROSCI.0051-15.2015. URL: <https://www.jneurosci.org/content/35/19/7443> (visited on 11/13/2019).
- Morrison, Sara E. and C. Daniel Salzman (2009). “The Convergence of Information about Rewarding and Aversive Stimuli in Single Neurons”. In: *Journal of Neuroscience* 29 (37), pp. 11471–11483. DOI: 10.1523/JNEUROSCI.1815-09.2009. URL: <http://www.jneurosci.org/content/29/37/11471> (visited on 07/25/2018).
- Munuera, Jérôme, Mattia Rigotti, and C. Daniel Salzman (2018). “Shared neural coding for social hierarchy and reward value in primate amygdala”. In: *Nature Neuroscience* 21 (3), pp. 415–423. DOI: 10.1038/s41593-018-0082-8. URL: <https://www.nature.com/articles/s41593-018-0082-8> (visited on 07/18/2018).
- Murray, John D., Jorge Jaramillo, and Xiao-Jing Wang (2017). “Working Memory and Decision-Making in a Frontoparietal Circuit Model”. In: *Journal of Neuroscience* 37 (50), pp. 12167–12186. DOI: 10.1523/JNEUROSCI.0343-17.2017. URL: <http://www.jneurosci.org/content/37/50/12167> (visited on 07/11/2018).
- Nicholson, Walter and Christopher Snyder (2007). *Microeconomic Theory: Basic Principles and Extensions*. 10th edition. Mason, OH: South-Western College Pub.
- O’Neill, Martin and Wolfram Schultz (2018). “Predictive coding of the statistical parameters of uncertain rewards by orbitofrontal neurons”. In: *Behavioural Brain Research*. SI: MCC 2016 355, pp. 90–94. DOI: 10.1016/j.bbr.2018.04.041. URL: <http://www.sciencedirect.com/science/article/pii/S0166432817307532> (visited on 10/02/2018).

- Padoa-Schioppa, Camillo (2009). “Range-Adapting Representation of Economic Value in the Orbitofrontal Cortex”. In: *Journal of Neuroscience* 29 (44), pp. 14004–14014. DOI: 10.1523/JNEUROSCI.3751-09.2009. URL: <http://www.jneurosci.org/content/29/44/14004> (visited on 09/07/2018).
- Padoa-Schioppa, Camillo and John A. Assad (2006). “Neurons in the orbitofrontal cortex encode economic value”. In: *Nature* 441 (7090), pp. 223–226. DOI: 10.1038/nature04676. URL: <https://www.nature.com/articles/nature04676> (visited on 07/25/2018).
- Padoa-Schioppa, Camillo and Katherine E. Conen (2017). “Orbitofrontal Cortex: A Neural Circuit for Economic Decisions”. In: *Neuron* 96 (4), pp. 736–754. DOI: 10.1016/j.neuron.2017.09.031. URL: <http://www.sciencedirect.com/science/article/pii/S0896627317309005> (visited on 08/01/2018).
- Pastor-Bernier, Alexandre and Paul Cisek (2011). “Neural Correlates of Biased Competition in Premotor Cortex”. In: *Journal of Neuroscience* 31 (19), pp. 7083–7088. DOI: 10.1523/JNEUROSCI.5681-10.2011. URL: <http://www.jneurosci.org/content/31/19/7083> (visited on 10/01/2018).
- Pastor-Bernier, Alexandre, Charles R. Plott, and Wolfram Schultz (2017). “Monkeys choose as if maximizing utility compatible with basic principles of revealed preference theory”. In: *Proceedings of the National Academy of Sciences* 114 (10), E1766–E1775. DOI: 10.1073/pnas.1612010114. URL: <http://www.pnas.org/content/114/10/E1766> (visited on 10/02/2018).
- Pastor-Bernier, Alexandre, Arkadiusz Stasiak, and Wolfram Schultz (2019). “Orbitofrontal signals for two-component choice options comply with indifference curves of Revealed Preference Theory”. In: *Nature Communications* 10 (1), pp. 1–19. DOI: 10.1038/s41467-019-12792-4. URL: <https://www.nature.com/articles/s41467-019-12792-4> (visited on 12/09/2019).
- Paton, Joseph J. et al. (2006). “The primate amygdala represents the positive and negative value of visual stimuli during learning”. In: *Nature* 439 (7078), pp. 865–870. DOI: 10.1038/nature04490. URL: <https://www.nature.com/articles/nature04490> (visited on 01/28/2018).
- Peck, Christopher J., Brian Lau, and C. Daniel Salzman (2013). “The primate amygdala combines information about space and value”. In: *Nature Neuroscience* 16 (3), pp. 340–348. DOI: 10.1038/nn.3328. URL: <https://www.nature.com/articles/nn.3328> (visited on 02/23/2020).
- Pelletier, Gabriel and Lesley K. Fellows (2019). “A critical role for human ventromedial frontal lobe in value comparison of complex objects based on attribute configuration”. In: *Journal of Neuroscience*, pp. 2969–18. DOI: 10.1523/JNEUROSCI.2969-18.2019. URL: <http://www.jneurosci.org/content/early/2019/03/12/JNEUROSCI.2969-18.2019> (visited on 03/15/2019).
- Phillips, Janice M. et al. (2000). “Cholinergic neurotransmission influences covert orientation of visuospatial attention in the rat”. In: *Psychopharmacology* 150 (1), pp. 112–116. DOI: 10.1007/s002130000437. URL: <https://doi.org/10.1007/s002130000437> (visited on 11/21/2019).

- Picciotto, Marina R., Michael J. Higley, and Yann S. Mineur (2012). “Acetylcholine as a Neuromodulator: Cholinergic Signaling Shapes Nervous System Function and Behavior”. In: *Neuron* 76 (1), pp. 116–129. DOI: 10.1016/j.neuron.2012.08.036. URL: <http://www.sciencedirect.com/science/article/pii/S0896627312008021> (visited on 11/19/2019).
- Raghuraman, Anantha P. and Camillo Padoa-Schioppa (2014). “Integration of Multiple Determinants in the Neuronal Computation of Economic Values”. In: *Journal of Neuroscience* 34 (35), pp. 11583–11603. DOI: 10.1523/JNEUROSCI.1235-14.2014. URL: <http://www.jneurosci.org/content/34/35/11583> (visited on 08/01/2018).
- Roe, Robert M., Jermon R. Busmeyer, and James T. Townsend (2001). “Multialternative decision field theory: A dynamic connectionist model of decision making.” In: *Psychological review* 108 (2), pp. 370–392. DOI: 10.1037/0033-295X.108.2.370. URL: <https://insights.ovid.com/psychological-review/plrev/2001/04/000/multialternative-decision-field-theory/4/00006832> (visited on 01/23/2020).
- Rustichini, Aldo and Camillo Padoa-Schioppa (2015). “A neuro-computational model of economic decisions”. In: *Journal of Neurophysiology* 114 (3), pp. 1382–1398. DOI: 10.1152/jn.00184.2015. URL: <https://www-physiology-org.proxy.library.nyu.edu/doi/abs/10.1152/jn.00184.2015> (visited on 08/01/2018).
- Sarter, Martin, Cindy Lustig, Anne S. Berry, et al. (2016). “What do phasic cholinergic signals do?” In: *Neurobiology of Learning and Memory* 130, pp. 135–141. DOI: 10.1016/j.nlm.2016.02.008. URL: <http://www.sciencedirect.com/science/article/pii/S1074742716000459> (visited on 11/20/2019).
- Sarter, Martin, Cindy Lustig, William M. Howe, et al. (2014). “Deterministic functions of cortical acetylcholine”. In: *European Journal of Neuroscience* 39 (11), pp. 1912–1920. DOI: 10.1111/ejn.12515. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/ejn.12515> (visited on 11/19/2019).
- Semyanov, Alexey et al. (2004). “Tonically active GABAA receptors: modulating gain and maintaining the tone”. In: *Trends in Neurosciences* 27 (5), pp. 262–269. DOI: 10.1016/j.tins.2004.03.005. URL: <http://www.sciencedirect.com/science/article/pii/S0166223604000906> (visited on 11/01/2019).
- Spencer, Janine et al. (2000). “Motion processing in autism: evidence for a dorsal stream deficiency”. In: *NeuroReport* 11 (12), pp. 2765–2767. DOI: 10.1097/00001756-200008210-00031. URL: https://journals.lww.com/neuroreport/Fulltext/2000/08210/Motion_processing_in_autism__evidence_for_a_dorsal.31.aspx (visited on 09/03/2019).
- Steinmetz, Nicholas A. et al. (2019). “Distributed coding of choice, action and engagement across the mouse brain”. In: *Nature*, pp. 1–8. DOI: 10.1038/s41586-019-1787-x. URL: <https://www.nature.com/articles/s41586-019-1787-x> (visited on 12/03/2019).

- Thomsen, Morgane, Gunnar Sørensen, and Ditte Dencker (2018). “Physiological roles of CNS muscarinic receptors gained from knockout mice”. In: *Neuropharmacology*. Neuropharmacology on Muscarinic Receptors 136, pp. 411–420. DOI: 10.1016/j.neuropharm.2017.09.011. URL: <http://www.sciencedirect.com/science/article/pii/S0028390817304288> (visited on 12/03/2019).
- Van de Cruys, Sander et al. (2014). “Precise minds in uncertain worlds: Predictive coding in autism”. In: *Psychological Review* 121 (4), pp. 649–675. DOI: 10.1037/a0037665. URL: <https://psycnet.apa.org/fulltext/2014-44506-004.html> (visited on 01/20/2020).
- Vasa, Roma A. et al. (2018). “Relationships between autism spectrum disorder and intolerance of uncertainty”. In: *Autism Research* 11 (4), pp. 636–644. DOI: 10.1002/aur.1916. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/aur.1916> (visited on 10/09/2019).
- Voytko, M. L. et al. (1994). “Basal forebrain lesions in monkeys disrupt attention but not learning and memory [published erratum appears in J Neurosci 1995 Mar;15(3): following table of contents]”. In: *Journal of Neuroscience* 14 (1), pp. 167–186. DOI: 10.1523/JNEUROSCI.14-01-00167.1994. URL: <https://www.jneurosci.org/content/14/1/167> (visited on 11/21/2019).
- Vreeswijk, C. van and H. Sompolinsky (1998). “Chaotic Balanced State in a Model of Cortical Circuits”. In: *Neural Computation* 10 (6), pp. 1321–1371. DOI: 10.1162/089976698300017214. URL: <https://doi.org/10.1162/089976698300017214> (visited on 09/04/2019).
- Wehr, Michael and Anthony M. Zador (2003). “Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex”. In: *Nature* 426 (6965), pp. 442–446. DOI: 10.1038/nature02116. URL: <https://www.nature.com/articles/nature02116> (visited on 11/06/2019).
- Wöhr, M. et al. (2015). “Lack of parvalbumin in mice leads to behavioral deficits relevant to all human autism core symptoms and related neural morphofunctional abnormalities”. In: *Translational Psychiatry* 5 (3), e525–e525. DOI: 10.1038/tp.2015.19. URL: <https://www.nature.com/articles/tp201519> (visited on 09/21/2019).
- Wong, Kong-Fatt, Alexander C. Huk, et al. (2007). “Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making”. In: *Frontiers in Computational Neuroscience* 1. DOI: 10.3389/neuro.10.006.2007. URL: <https://www.frontiersin.org/articles/10.3389/neuro.10.006.2007/full> (visited on 08/02/2018).
- Wong, Kong-Fatt and Xiao-Jing Wang (2006). “A Recurrent Network Mechanism of Time Integration in Perceptual Decisions”. In: *Journal of Neuroscience* 26 (4), pp. 1314–1328. DOI: 10.1523/JNEUROSCI.3733-05.2006. URL: <http://www.jneurosci.org/content/26/4/1314> (visited on 06/12/2018).
- Wunderlich, Klaus, Antonio Rangel, and John P. O’Doherty (2009). “Neural computations underlying action-based decision making in the human brain”. In: *Proceedings of the National Academy of Sciences* 106 (40),

- pp. 17199–17204. DOI: 10.1073/pnas.0901077106. URL: <http://www.pnas.org/content/106/40/17199> (visited on 10/01/2018).
- Xie, Jue and Camillo Padoa-Schioppa (2016). “Neuronal remapping and circuit persistence in economic decisions”. In: *Nature Neuroscience* 19 (6), pp. 855–861. DOI: 10.1038/nn.4300. URL: <https://www.nature.com/articles/nn.4300> (visited on 08/01/2018).
- Yoo, Seng Bum Michael and Benjamin Yost Hayden (2018). “Economic Choice as an Untangling of Options into Actions”. In: *Neuron* 99 (3), pp. 434–447. DOI: 10.1016/j.neuron.2018.06.038. URL: <http://www.sciencedirect.com/science/article/pii/S0896627318305348> (visited on 09/30/2018).
- Yu, Angela J. and Peter Dayan (2005). “Uncertainty, Neuromodulation, and Attention”. In: *Neuron* 46 (4), pp. 681–692. DOI: 10.1016/j.neuron.2005.04.026. URL: <http://www.sciencedirect.com/science/article/pii/S0896627305003624> (visited on 11/12/2019).
- Zaborszky, Laszlo et al. (2015). “Neurons in the Basal Forebrain Project to the Cortex in a Complex Topographic Organization that Reflects Corticocortical Connectivity Patterns: An Experimental Study Based on Retrograde Tracing and 3D Reconstruction”. In: *Cerebral Cortex* 25 (1), pp. 118–137. DOI: 10.1093/cercor/bht210. URL: <https://academic.oup.com/cercor/article/25/1/118/367269> (visited on 11/21/2019).
- Zaidel, Adam, Robin P. Goin-Kochel, and Dora E. Angelaki (2015). “Self-motion perception in autism is compromised by visual noise but integrated optimally across multiple senses”. In: *Proceedings of the National Academy of Sciences* 112 (20), pp. 6461–6466. DOI: 10.1073/pnas.1506582112. URL: <https://www.pnas.org/content/112/20/6461> (visited on 06/03/2019).
- Zikopoulos, Basilis and Helen Barbas (2013). “Altered neural connectivity in excitatory and inhibitory cortical circuits in autism”. In: *Frontiers in Human Neuroscience* 7. DOI: 10.3389/fnhum.2013.00609. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3784686/>.
- Zucker, R S (1989). “Short-Term Synaptic Plasticity”. In: *Annual Review of Neuroscience* 12 (1), pp. 13–31. DOI: 10.1146/annurev.ne.12.030189.000305. URL: <https://doi.org/10.1146/annurev.ne.12.030189.000305> (visited on 12/18/2019).