



HHS Public Access

Author manuscript

Behav Neurosci. Author manuscript; available in PMC 2022 April 01.

Published in final edited form as:

Behav Neurosci. 2021 April ; 135(2): 245–254. doi:10.1037/bne0000457.

Heterogeneous value coding in orbitofrontal populations

Pierre Enel¹, Aster Q. Perkins¹, Erin L. Rich¹

¹Nash Family Department of Neuroscience and Friedman Brain Institute, Icahn School of Medicine at Mount Sinai, New York, New York, USA

Abstract

Value signals in the brain are important for learning, decision-making, and orienting behavior toward relevant goals. Although they can play different roles in behavior and cognition, value representations are often considered to be uniform and static signals. Nonetheless, contextual and mixed representations of value have been widely reported. Here, we review the evidence for heterogeneity in value coding and dynamics in the orbitofrontal cortex. We argue that this diversity plays a key role in the representation of value itself and allows neurons to integrate value with other behaviorally relevant information. We also discuss modeling approaches that can dissociate potential functions of heterogeneous value codes and provide further insight into its importance in behavior and cognition.

Value

In psychology and neuroscience, value as a concept is often loosely defined. Something of value may be simple, like a chocolate donut, or complex like a job promotion; it may be concrete, like a bank note, or abstract like an idea; it may meet biological needs, like food and drink, or not, like art and music. Value can be something sought or something experienced. It can drive learning or be learned, and can alter attention, memory, and motivation. The field of economics quantifies “utility”, a concept that is similar to value but in many cases is unable to capture the full spectrum of behavior and psychology related to value itself (Camerer, 2014). Despite this intangibility, there is a common understanding that value exists on a continuum from bad to good, and for the purposes of this review we will largely rely on this intuitive definition. Departing from the traditional perspective, however, we propose that value representations in the brain should not be uniform, idealized signals that vary in direct relation to this continuum. Instead, our discussion will focus on the heterogeneity of value signals in neural circuits, and how this might reflect the myriad ways that representing value is critical for adaptive behavior and cognition. We also examine how this diversity of coding might emerge and what roles it might play in network computations. However, we start from the premise that value is a poorly constrained construct, compared, for instance, to concrete sensory or motor events, and encompasses broadly different ethological functions, types of information, and cognitive processes.

Corresponding Author: Erin L. Rich, 1 Gustave Levy Place, New York, NY 10029, Tel: (212) 824-9304, erin.rich@mssm.edu.

Perhaps because value has varied qualities and roles, it is richly represented throughout the brain (Vickery et al., 2011). Across species, value signals can be resolved at the level of single neurons, local field potentials, population codes and large-scale network activity (Padoa-Schioppa & Assad, 2006; Rich & Wallis, 2017; Vickery et al., 2011; Wallis & Miller, 2003). Subcortical regions including the amygdala and areas of the ventral and medial frontal cortex are most strongly associated with signaling value; however, different areas are believed to play distinct roles in learning and value-mediated behavior (Wallis & Kennerley, 2010). For example, since the early demonstrations of intracranial self-stimulation, the ventral striatum (VS) and related circuits have been associated with positive reinforcement (Olds & Milner, 1954), and converging evidence indicates that value signals here underlie hedonic effects of reward (Berridge et al., 2009). In contrast, most frontal cortical regions do not support self-stimulation (Mora et al., 1980), or do so in a context-dependent manner (Mora et al., 1979), suggesting that value-related activity in these regions plays a different role.

Further insights into functional specificity across frontal cortex can similarly be gained through analysis of value-related activity in different regions. For instance, although value is encoded throughout the frontal cortex (e.g. (Matsumoto et al., 2003; Roesch & Olson, 2003; Wallis & Miller, 2003; Watanabe, 1996)), the orbitofrontal cortex (OFC) and anterior cingulate cortex (ACC) have been most strongly implicated in value-related behavior (Amiez et al., 2005; Knutson et al., 2003; Knutson et al., 2005; Sallet et al., 2007). Both regions encode value signals, but differences in response profiles have suggested distinct functions (Wallis & Kennerley, 2010). For instance, OFC neurons signal both pre- and post-decision variables, including the value of available options (Padoa-Schioppa & Assad, 2006; Rich & Wallis, 2016), consistent with the idea that they causally contribute to decision formation (Padoa-Schioppa & Conen, 2017). On the other hand, ACC encodes the value of a reward chosen by the subject but not the value of each individual option, suggesting that these areas are involved in processing that occurs after a choice is made (Cai & Padoa-Schioppa, 2012). ACC value coding also tends to lag that of OFC, relates actions to their associated outcomes, and encodes prediction error signals (Kennerley et al., 2011; Luk & Wallis, 2013; Wallis & Kennerley, 2010; Wallis & Rich, 2011), so that in comparison to OFC, ACC value signals appear to be downstream and therefore more proximal to the behavioral execution of a choice. As such, the ACC has been linked to a variety of evaluative functions including updating of beliefs and models of the environment, committing to a course of action, determining the benefits and costs of an action, and exploring or searching for alternatives (Amemori & Graybiel, 2012; Heilbronner & Hayden, 2016; Kolling et al., 2016), whereas the OFC is believed to be critical for learning, expectation, and decision-making functions, particularly when they rely on unobservable variables and task structures (Schoenbaum et al., 2011; Wilson et al., 2014).

Comparisons such as these illustrate how heterogeneity in value signals can underlie distinct functions of different brain regions. In the remainder, we expand on this perspective to suggest that value signals within a brain area are not uniform, but dynamic and heterogeneous, and this complexity plays important roles in network function. We take a particular focus on the OFC and first discuss recent experimental results demonstrating heterogeneity of OFC value signals and the roles these may play in cognitive processes. We

then suggest ways in which artificial network models that mimic these patterns could be used to provide additional insights into the origin and importance of diversity in value signaling.

Heterogeneous value signals in OFC

A wealth of evidence supports the notion that OFC function is critical for learning and updating value expectations (Rudebeck et al., 2013; Schoenbaum et al., 1998; Schoenbaum & Roesch, 2005), and this is reflected in the predominance of value coding neurons observed in OFC populations. Across neurons, however, value codes are not uniform. Now-classic studies have shown that, in a simple decision task involving different amounts and types of reward, most responses in OFC can be classified as encoding the subjective value of the chosen reward (often referred to as “chosen value” or “goal value”), the type of reward that was chosen, or the value of one type of reward that was offered (“offer value”) (Padoa-Schioppa, 2013; Padoa-Schioppa & Assad, 2006). Here, chosen value is associated with a subject’s choice in anticipation of receiving a reward, and is distinct from “outcome value,” or the value computed upon receiving the reward itself (Peters & Buchel, 2010). As task demands change, however, the types of value signals found in OFC become more complex. Human and rodent studies support the notion that many value signals in OFC are sensory-specific, and therefore represent the combination of outcome identity and value (Howard et al., 2015; Stalnaker et al., 2014). For instance, even when valued the same, different food items like cupcakes or potato chips elicit distinguishable value responses in human OFC (Howard et al., 2015), and these change when the appeal of the food is altered by feeding the participant to satiety (Howard & Kahnt, 2017). We have also found sensory-specific value signals encoded in monkey OFC neurons, though they are relatively weak compared to nonspecific values (Rich & Wallis, 2016, 2017). Value information can also be mixed with other decision variables like choice, expected outcome and recent history (Kimmel et al., 2020; Saez et al., 2018).

The neural signals that reflect the evaluation of different items or aspects of items can also exhibit nonrandom trial-by-trial variability, suggesting that they are more deeply entwined in cognitive processing than simply signaling value. For instance, chosen value neurons have slightly stronger signals when the competing offer is more valuable (Padoa-Schioppa, 2013). Although one possibility is that this stronger signal originates from variability in subjective evaluation that biases choices, the fact that chosen values are defined by the subject’s decision suggests that the variability may be more important to post-decision processing. Nonetheless, value signals such as these, that integrate aspects of the specific decision or task, may be important for specifying the particular goals of motivated behavior, while nonspecific value signals may serve more general roles in reward-based learning.

Beyond task-specific information, OFC value signals are also modulated by attention. For instance, OFC neurons that encode the value of a visual cue become more selective when visual fixations focus on the cue (McGinty et al., 2016). When there are two cues, visual perturbations that attract attention to one or the other influence the firing rate of value-coding neurons (Xie et al., 2018). Moreover, value signals can interact in an attention-dependent manner: in OFC they vary with the difference between the attended and

unattended options (Lim et al., 2011) indicating that attention influences which values are encoded, how these values relate to one another, and the magnitude of the value signal itself. Together, the impact of factors such as choice type, choice context, task structure, and attention on OFC value signals demonstrate that value representations are heterogeneous and dynamically modulated. In addition, in some cases this variability can be traced back to features of the task and behavior that are unrelated to value itself.

Non-canonical value coding

The canonical single neuron value representation described so far is a monotonic and approximately linear relationship between value and firing rate (Padoa-Schioppa & Assad, 2006). Although this is the most prevalent form of value representation in OFC, we have also found that a non-negligible proportion of neurons robustly encode value with nonlinear and non-monotonic activities (Enel et al., 2020). In the task, monkeys were presented one of 8 possible stimuli distributed among 4 discrete values, and received the corresponding reward after a delay. Approximately 13% of the recorded units were not linearly related to value, but significantly encoded value as measured with an ANOVA. Some neurons had a clear preference for a specific value, whereas others had highest firing rates for the extreme values, and others had profiles that could not be readily interpreted (Figure 1). This nonlinear encoding was not exclusive to the timing of stimulus presentation, and occurred at various times during the delay between cue presentation and outcome delivery. Because linear value coding was most prevalent and decoding value with linear neurons only did not decrease accuracy, the role of these nonlinear activities remains unclear. However, the presence of nonlinear and non-monotonic value representations is consistent with reported representations of other magnitude variables like duration, distance and numbers, all of which are represented with both monotonic and non-monotonic activities (Dehaene & Brannon, 2011; Eichenbaum, 2014; Funahashi, 2013; Moser & Moser, 2008; Nieder & Dehaene, 2009; Wittmann, 2013).

One possibility is that nonlinear value neurons reflect an extreme discretization of value information in an overlearned task. Monkeys were overtrained in this experiment, and this was reflected in their highly accurate and efficient behavioral performance. As a consequence, they were very familiar with the stimuli and the 4 different levels of value. It is therefore possible that the value scale of the task came to be represented more categorically than ordinally, and a subset of neurons developed unique responses to value categories. In the task, monkeys were also presented with choice trials in which two reward-predicting cues appeared, and the animal made a choice among the cues with a saccade. One potential benefit of discretizing value is that comparing simultaneously presented stimuli is more computationally straightforward when individual values have their own representations. For instance, when the number of options is small and consistent, subjects quickly learn which stimuli correspond to the highest or lowest possible rewards. This can make for an easy choice using a simple heuristic: choose the highest valued stimulus if available, and avoid the lowest valued. This hypothesis could explain the existence of OFC neurons encoding the lowest and highest values available in this task. Similarly, a small number of stimuli means that only a limited number of combinations of two-value choices exist. By representing individual values, another heuristic could potentially emerge with overtraining, in the form

of a direct mapping between value combinations and choice, bypassing a potentially costly and time-consuming value comparison.

Another perspective is that nonlinear coding increases the dimensionality of value signals in a neural population, and this could offer computational advantages to a value processing network. Nonlinear mixing of task variables has been shown to serve a similar function in other regions of prefrontal cortex (Fusi et al., 2016; Rigotti et al., 2013). In this case, prefrontal neurons were found to combine multiple task variables in nonlinear ways that at first appeared to be random, but actually served to increase the dimensionality of the task representations the network could generate, which facilitates efficient and flexible read-outs with simple linear methods. The proportion of neurons with nonlinear mixed selectivity also increases with learning, supporting the view that this is an efficient means of coding information (Dang et al., 2020). Similarly, nonlinear coding should facilitate value decoding from a population of neurons, and future studies could investigate whether such coding also increases with learning.

To demonstrate that nonlinear coding improves decodability, we created synthetic neural data sets with defined properties (Figure 2). Our goal was to hold all aspects of the simulated populations constant while systematically varying the mapping between mean firing rate and value to create populations that vary monotonically or non-monotonically. Thus, each “neuron” was characterized by a firing rate on each of 1200 trials, with trials randomly assigned an ordinal value (1 to 4). One group of neurons had average firing rates that increased or decreased monotonically with value. These were simulated by defining a mean firing rate between 0 and 20 Hz for each neuron responding to each value, and drawing the neuron’s response on a given trial from a Gaussian distribution centered on that value. Two sources of noise were added to these firing rates. The first was neuron-specific noise that defined the standard deviation of the Gaussians from which the trial-wise firing rates were drawn. This varied by neuron between 15 and 30. The second was shared noise that varied by trial but was the same across all neurons in our simulated populations. Shared noise was drawn from a Gaussian with 0 mean and standard deviation ranging between 1.5 and 3. Non-monotonic neurons were identical to monotonic value neurons, except that the order of the value-to-mean firing rate assignments were shuffled pseudorandomly, with the constraint that the mapping could not be constantly increasing or decreasing. A third group of nonselective neurons was created by following the same procedure, except the response on each trial was drawn from any of the four value distributions. Examples of each type of neuron are shown in Figure 2A. To ensure each group had the characteristics we aimed for, we used a general linear model to predict each neuron’s firing rate from trial-wise value, and showed that monotonic neurons were linearly related to value, with either positive or negative regression coefficients. Non-monotonic neurons were less linearly related to value with regression coefficients widely distributed around zero, and nonselective neurons had regression coefficients tightly grouped around zero.

With these simulated neurons, we next tested our ability to decode trial-wise value from different populations using linear discriminant analysis with 10-fold cross validation. To do this, we created populations of 100 neurons, where 50% were non-selective and 50% were value-coding. We compared these populations when all value-coding neurons were

monotonic, and as the proportion of non-monotonic neurons increased one neuron at a time, until there were equal proportions of each type. We found that decoding accuracy increased asymptotically as more non-monotonic neurons were added to the population (Figure 2B). This demonstrates that, although the non-monotonic codes may not be intuitively related to values, when present in a population they increase decodability at a network level, effectively making the value signal more robust. Similar to the demonstration that nonlinear mixed selectivity offers computational advantages to a network encoding multiple types of information (Fusi et al., 2016; Rigotti et al., 2013), nonlinear coding of value increases dimensionality of neural population activity, which can enhance algorithmic readout by simple linear classifiers.

Dynamic value signals in OFC

So far, we have described multiple facets of heterogeneity in OFC value signals, including those reflecting non-value information relevant to the decision process. Incorporating such information into value-related signals may be a critical mechanism for coping with changing cognitive or behavioral demands that arise across the course of a decision. In many cases, there are gaps of time between a predictive event or cue and the valuable outcome it signals, and unrelated cognitive demands may intervene between the two. In order for value expectations to be updated effectively, there must be a neural mechanism that either maintains these expectations across time or retrieves them at the time an outcome is experienced. Behavioral evidence indicates that value expectations are likely maintained across these delays when they are relatively short. In the time following a predictive cue, attention is oriented to rewarding stimuli, and higher value stimuli more effectively attract attention (McGinty et al., 2016). Similarly, motor responses are faster, memory performance is more accurate, and task completion is more consistent when subjects expect a higher value outcome (Kennerley & Wallis, 2009; Rich & Wallis, 2016), all suggesting the presence of value information during cue-outcome delays. Using value representations held in mind to organize processing like attention and motor preparation is evolutionarily advantageous, because it increases the likelihood that a potential reward will be obtained. Thus, beyond neural signals that could serve the role of a “value eligibility trace” by holding an expected value online to facilitate subsequent updating, value expectations maintained across delays can also influence behavior to prepare the animal to obtain a reward.

In other cognitive tasks, holding information across delays is a key function of working memory (WM), which is known to depend on the dorsolateral prefrontal cortex (dlPFC) in primates (Barbey et al., 2013; Butters & Pandya, 1969; Levy & Goldman-Rakic, 1999). In dlPFC, different neural dynamics have been found to represent information stored in working memory, and these proposed mechanisms of WM can provide a framework for investigating how value representations are maintained across delays. On one hand, trial-averaged data shows persistent spiking, supporting the view that there is a stable neural representation of the information held in WM (Constantinidis et al., 2018). However, individual neurons can exhibit great variability in firing rate across cell groups, within individual cells, and across the duration of the trial (Brody et al., 2003; Shafi et al., 2007), suggesting a more dynamic signal and raising the question of how stable representations might be stored in changing neural activity. Recent work has found stable subspaces in

dIPFC activity, such that WM signals are maintained in an unchanging format across delays (Murray et al., 2017; Spaak et al., 2017). These appear to be embedded in time varying coding of task variables, in which selectivity is transiently present within individual neurons. In this case, stable coding may permit consistent readout of WM information, while dynamic coding provides time-relevant information (Lee et al., 2020). Given these perspectives on maintenance of cognitive information in WM, we reasoned that similar insights into the cause and function of heterogeneous value signals in OFC might be gained by examining their dynamics across delays.

To investigate the dynamics of value representations, we recorded populations of neurons from OFC while monkeys performed a value-based decision-making task in which a visual cue predicted a delayed reward (Enel et al., 2020). Because stable but not dynamic codes should generalize across time, we used a cross-temporal decoding method (Meyers et al., 2008) to assess the similarity of value representations between two time points during the delay (Figure 3). A stable signal would result in equally high decoding accuracy at all pairs of time points (i.e., a decoder trained on one time point can generalize to distant time points), whereas high accuracy restricted to pairs of data points that are close in time is an indication that the value signal evolves across the delay, and is therefore dynamic.

Our results did not strictly conform to either stable or dynamic pictures. Instead, OFC neural populations showed an intermediate level of stability. The most reliable value signals were extracted when the same time windows were used to train and test the decoder, consistent with a dynamic signal that evolved over time. However, we could also decode weaker but significant value signals when the decoder was trained and tested on more distant time points, suggesting the existence of a stable signal in the same neural population. Using further analysis methods developed to isolate each type of signal, we confirmed that stable and dynamic value representations were dissociable. To do this, we combined recent population methods to select an ensemble of neurons with stable value coding and then extract a stable value subspace from this population activity. In this way, we could identify a very stable representation of value across the delay, consistent with what is typically described as a persistent representation (Constantinidis et al., 2018). Conversely, temporally local representations of value could be isolated with an ensemble selection method targeting dynamic value encoding. This yielded a representation that evolved throughout the delay via reliably sequenced encoding within the population (Figure 3B). Interestingly, the ensembles that yielded the most stable or local dynamics comprised on average only 13% and 28% of the full OFC population respectively. The population in dynamic ensembles was larger, most likely because covering the delay with successive local representations required more units than a stable representation, which can rely on the same units covering longer portions of the delay.

Thus, we demonstrated that opposite dynamics coexist within the same population of OFC neurons coding expected values. These two dynamical regimes have unique benefits in terms of information processing for specific behavioral demands, and are both able to represent value information over a delay. On one hand, stable value representations are relevant for robust and time-independent maintenance of value information. Such stability allows any local or downstream neurons to access the expected value information in a time independent

manner, making this encoding more robust because a single set of connections is able to extract value information in a reliable format across the duration of the delay. In other words, the downstream network does not have to expect value information at a particular time, as it is always available with a single set of readout connections.

On the other hand, dynamic representations are likely to participate in the temporal organization of behavior. This is because temporal information can be extracted in addition to value itself, and since most experiences and cognitive tasks are organized in time, this allows a network to be prepared to process incoming stimuli and generate behaviors. As noted above, expected values can influence attentional, motor, and sensory processes, and dynamic value representations may allow these influences to be organized in time. This organization may be key for both optimal processing (a network receiving an unexpected input is not as likely to process it optimally or even correctly) and faster behavioral output. Although temporal and value information could have separate and largely independent representations, i.e. a stable value representation and a context independent representation of time, this would require that a downstream network perform the extra step of integrating time and value. Conversely, a mixed representation means that time dependent value signals are readily available. Our results show that these mixed representations are robust across trials (Enel et al., 2020), suggesting they have been strengthened through learning, likely due to their relevance to performing the task. These results have notable similarity to the dynamic and stable WM representations reported in the dlPFC, and similar interpretations have been drawn regarding their relevance to behavior (Meyers, 2018). The OFC recordings were taken from a broad region of central OFC, including targets in Walker's areas 11 and 13. These were grouped into one population based on previous results that found only minor differences in value encoding among neurons within the recorded field (Rich & Wallis, 2017). However other data have suggested subregional differences that could manifest in different value dynamics. Specifically, inactivations of the more posterior area 13 disrupted value updating in a devaluation paradigm but left subsequent decisions intact, while inactivation of anterior area 11 showed it is necessary for selections based on a devalued state but not during the updating process (Murray et al., 2015). Further work is needed to investigate whether and how different dynamics might support these proposed roles in behavior. Regardless, in the case of WM or evaluation, stable representations may provide a robust signal that can be read out by a single set of synapses, whereas dynamic representations could enable the concurrent encoding of time to support temporally-organized information processing.

Insights from artificial networks

The example of artificial neurons above suggests ways that nonlinear coding might be computationally advantageous. Similarly, simulations of interconnected neurons that form artificial neural networks (ANN) can help us gain insights and test hypotheses about populations of neurons in general, and about value coding dynamics specifically. As simulations of biological networks, ANNs have many parameters like training schedules and network architectures that can be systematically manipulated in isolation to test effects on coding, dynamics and behavior (for example, as in (Barak et al., 2013)). A key benefit to this approach is that it allows experiments that are not possible *in vivo*, such as precisely and

repeatedly manipulating inputs or selectively lesioning neurons with particular coding properties. By permitting us to “edit” and examine different aspects, artificial networks allow us to identify and test principles that may underlie neural systems at large. Successful ANN models also point to potential directions of investigation in biological systems, and in particular, reveal ways to test theories developed with models.

Given our results that show mixed dynamics in OFC value representations, we hypothesized that stable attractor-like representations and temporally local ones associated with dynamical activities each have their own potential benefits. One of the surprising features of our study was the concurrent existence of both of these dynamical regimes in the same neural population. Traditional models of working memory are associated with attractors, whether they are sets of fixed points of the neural dynamics that represent discrete items (Compte et al., 2000) or one-dimensional manifolds in the form of linear/ring attractors that represent the continuous/cyclical dimension of a stimulus feature (Wang, 2001; Wimmer et al., 2014). In these models, the entire population of units is involved in the attractor, as if to freeze the network’s dynamics in a state that is specific to the information that must be remembered until it can be used. As a consequence, these models cannot include dynamical activities. Although a small proportion of persistently firing neurons recorded in prefrontal regions do indicate the presence of an attracting dynamic of this sort (Constantinidis et al., 2018; Fuster & Alexander, 1971; Goldman-Rakic, 1995), many more neurons display dynamic or transient patterns of activity (Barak et al., 2010; Goldman-Rakic, 1995; Lundqvist et al., 2016). Our study indicates that, in OFC, the representation of value follows both of these patterns as well. Consistent with a recent studies (Murray et al., 2017; Parthasarathy et al., 2019) our results hint at the possibility that attracting and dynamical patterns of activity representing task-relevant information are concurrently present throughout the PFC. Although pure attractor networks are very efficient for holding information “online”, they are not biologically plausible. That is, attracting dynamics cannot involve the entire population activity of a brain region, as this would leave the network unresponsive to incoming inputs. Mixing stable and dynamic regimes may offer the advantage of allowing the circuit to process new inputs while maintaining a stable representation. A future challenge will be to assess whether these stable representations arise from attractors in neural populations or stable read-outs of attractor-less dynamics.

In addition, concurrent attracting and dynamic activities might have a larger role, as suggested by other modeling studies that successfully combined attractors and dynamic processing. Artificial neural networks with different architectures can display mixed dynamical regimes similar to those described here. These are most commonly random recurrent networks of the reservoir type, in which units are connected with randomly generated weights that remain fixed. Trained readout units feed back to the recurrent layer to allow for the emergence of attractors (Enel et al., 2016; Maass et al., 2007; Pascanu & Jaeger, 2011). In these networks, attractors involve only a fraction of the total dynamics of the population, which leaves room for concurrent dynamic activity. The PFC might use similar mechanisms to represent different types of information across time, such that the network represents memoranda over long periods while still processing incoming inputs. In the case of expected values, an attractor could stably maintain the information, while the rest of the network processes incoming stimuli in a way that is biased by the value being

represented in the attractor. These dynamics could provide a mechanism by which attention is enhanced by sustained value representation, resulting in lower rates of errors and shorter reaction times with higher predicted rewards (Anderson et al., 2011; Roesch & Olson, 2007). In this type of mixed dynamical network, stimulus processing and behavioral output may both be influenced by the expected value represented by an attractor state.

In addition to the existence of attractors as part of network dynamics, the degree of their involvement in the representation of value may directly reflect the amount of training on the task that includes these values (Barak et al., 2010). Upon first exposure, the association between a stimulus and its outcome is expected to be weak, so representations of stimulus values are unlikely to be stable and carried by persistent activity. Consequently, initial representations are likely to be less organized and more dynamic. With training, more efficient schemes of value encoding might emerge and tend towards the stable and persistent activities. To illustrate this, three models of working memory, trained with increasingly stable dynamics, were developed in the context of a delayed discrimination task in which monkeys compared the frequency of two tactile stimulations separated by a delay (Barak et al., 2010). A highly dynamic random network with fixed recurrent connections (known as reservoir computing, described above (Jaeger, 2001; Maass et al., 2002)) was the most dynamic network, and a stable, hand tuned line attractor network (sometimes referred to as a bump attractor, (Amit, 1995; Machens et al., 2005; Wang, 2001) was the most stable network. An intermediate network was also created, with recurrent connections that were initially random but modified through training and exhibited dynamics that were a mix of the first two models. This intermediate network best reproduced several aspects of the electrophysiological recordings of macaque monkeys performing this task, suggesting that cortical dynamics may result from partial training of somewhat random neuronal connections. The authors suggest that the three models potentially represent the successive dynamical states visited by the PFC through learning, from unstructured and dynamic to more structured and stable. Future experiments could explore this challenging question with neural recordings obtained during the course of learning.

The degree of attracting dynamics may also depend on the task demands. A recent study showed that among models trained on different WM tasks, those exposed to tasks requiring manipulation of the remembered information were more likely to develop persistent activities with training (Masse et al., 2019). Similarly, neural networks trained to perform different WM tasks demonstrated that tasks with variable delays elicited more persistent activity, while those with greater temporal complexity generated more dynamic activities (Orhan & Ma, 2019). Supporting this theory, existing literature suggests that the structure of WM tasks influences representation dynamics in the PFC (Cavanagh et al., 2018; Parthasarathy et al., 2017; Spaak et al., 2017). This further supports the notion that stable and dynamic value signals arise to serve specific roles in task performance, and the existence of both types of signals allows the network to flexibly support a variety of cognitive demands.

Summary

Here we have outlined the many ways that OFC exhibits heterogeneity in value coding. Value signals are frequently mixed with task-relevant information and modified by ongoing attentional or cognitive processes; they also take on different dynamics and non-canonical coding patterns. Through converging examples drawn from modeling and experimental work, we have proposed that this variety may be a hallmark of value representations that are adaptively integrated into cognitive processes. For instance, different dynamics could underlie different roles for value signals in integrated networks (Enel et al., 2020), and mixing of value with other task-relevant information could serve to increase the efficiency of neural coding (Fusi et al., 2016).

We close with a final observation that flexibility in value scales could explain the generality of value as a concept, which we alluded to in the beginning of this paper. Because the dynamic range of neural firing rates is limited, OFC neurons rescale their responses to accommodate current contexts, for example reflecting the range of values of recently available options (Kobayashi et al., 2010; Padoa-Schioppa, 2009). Similar scaling effects have been observed on EEG recordings from humans when they encounter numbers in blocks with different ranges (Sheahan et al., 2020). Using ANNs to model humans' ability to grade numbers as lower or higher than the previous number in a sequence, numbers of the same relative magnitude within a block became aligned across different contexts. It was found that divisive normalization ensured that number scales had the same size in the neural representation space across contexts, and subtractive normalization registered these scales to align them.

The effects of these normalization steps on number scales are strikingly similar to the relative representation of value in OFC, suggesting an underlying mechanism for range adaptation. This simple explanation drawn from network modeling could have broad impacts on our understanding of value. It provides a potential mechanism by which disparate concepts and experiences can be placed on a generalizable scale from bad to good. In effect, this may give traction for understanding how we can have an intuitive, but poorly constrained concept of value. Insights such as these illustrate the importance of embracing variability in data as apparently simple as value codes, and the promise of combining experimental and theoretical approaches to interpret them.

Acknowledgements

The authors thank M.G. Baxter for comments and suggestions on the manuscript.

References

- Amemori K, & Graybiel AM (2012). Localized microstimulation of primate pregenual cingulate cortex induces negative decision-making. *Nat Neurosci*, 15(5), 776–785. 10.1038/nn.3088 [PubMed: 22484571]
- Amiez C, Joseph J-P, & Procyk E (2005). Anterior cingulate error-related activity is modulated by predicted reward. *European Journal of Neuroscience*, 21(12), 3447–3452.
- Amit DJ (1995). The Hebbian paradigm reintegrated: Local reverberations as internal representations. *Behav Brain Sci*, 18(4), 617–657.

- Anderson BA, Laurent PA, & Yantis S (2011). Value-driven attentional capture. *Proc Natl Acad Sci U S A*, 108(25), 10367–10371. 10.1073/pnas.1104047108 [PubMed: 21646524]
- Barak O, Sussillo D, Romo R, Tsodyks M, & Abbott LF (2013). From fixed points to chaos: three models of delayed discrimination. *Prog Neurobiol*, 103, 214–222. 10.1016/j.pneurobio.2013.02.002 [PubMed: 23438479]
- Barak O, Tsodyks M, & Romo R (2010). Neuronal population coding of parametric working memory. *J Neurosci*, 30(28), 9424–9430. 10.1523/JNEUROSCI.1875-10.2010 [PubMed: 20631171]
- Barbey A, Koenigs M, & Grafman J (2013). Dorsolateral prefrontal contributions to human working memory. *Cortex*, 49(5), 1195–1205. [PubMed: 22789779]
- Berridge KC, Robinson TE, & Aldridge JW (2009). Dissecting components of reward: ‘liking’, ‘wanting’, and learning. *Curr Opin Pharmacol*, 9(1), 65–73. 10.1016/j.coph.2008.12.014 [PubMed: 19162544]
- Brody CD, Hernández A, Zainos A, & Romo R (2003). Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb Cortex*, 13(11), 1196–1207. 10.1093/cercor/bhg100 [PubMed: 14576211]
- Butters N, & Pandya D (1969). Retention of delayed-alternation: effect of selective lesions of sulcus principalis. *Science*, 165(3899), 1271–1273. [PubMed: 4979528]
- Cai X, & Padoa-Schioppa C (2012). Neuronal encoding of subjective value in dorsal and ventral anterior cingulate cortex. *J Neurosci*, 32(11), 3791–3808. 10.1523/JNEUROSCI.3864-11.2012 [PubMed: 22423100]
- Camerer CF (2014). Behavioral economics. *Curr Biol*, 24(18), R867–R871. 10.1016/j.cub.2014.07.040 [PubMed: 25247364]
- Cavanagh SE, Towers JP, Wallis JD, Hunt LT, & Kennerley SW (2018). Reconciling persistent and dynamic hypotheses of working memory coding in prefrontal cortex. *Nat Commun*, 9(1), 3498. 10.1038/s41467-018-05873-3 [PubMed: 30158519]
- Compte A, Brunel N, Goldman-Rakic PS, & Wang XJ (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb Cortex*, 10(9), 910–923. [PubMed: 10982751]
- Constantinidis C, Funahashi S, Lee D, Murray JD, Qi XL, Wang M, & Arnsten AFT (2018). Persistent Spiking Activity Underlies Working Memory. *J Neurosci*, 38(32), 7020–7028. 10.1523/JNEUROSCI.2486-17.2018 [PubMed: 30089641]
- Dang W, Jaffe RJ, Qi X-L, & Constantinidis C (2020). Emergence of non-linear mixed selectivity in prefrontal cortex after training. *bioRxiv*, doi: 10.1101/2020.1108.1102.233247.
- Dehaene S, & Brannon EM (2011). *Space, time and number in the brain : searching for the foundations of mathematical thought : an attention and performance series volume (1st ed.)*. Elsevier Academic Press.
- Eichenbaum H (2014). Time cells in the hippocampus: a new dimension for mapping memories. *Nat Rev Neurosci*, 15(11), 732–744. 10.1038/nrn3827 [PubMed: 25269553]
- Enel P, Procyk E, Quilodran R, & Dominey PF (2016). Reservoir Computing Properties of Neural Dynamics in Prefrontal Cortex. *PLoS Comput Biol*, 12(6), e1004967. 10.1371/journal.pcbi.1004967 [PubMed: 27286251]
- Enel P, Wallis JD, & Rich EL (2020). Stable and dynamic representations of value in the prefrontal cortex. *Elife*, 9. 10.7554/eLife.54313
- Funahashi S (2013). Space representation in the prefrontal cortex. *Progress in Neurobiology*, 103, 131–155. [PubMed: 22521602]
- Fusi S, Miller EK, & Rigotti M (2016). Why neurons mix: high dimensionality for higher cognition. *Curr Opin Neurobiol*, 37, 66–74. 10.1016/j.conb.2016.01.010 [PubMed: 26851755]
- Fuster JM, & Alexander GE (1971). Neuron activity related to short-term memory. *Science*, 173(3997), 652–654. [PubMed: 4998337]
- Goldman-Rakic PS (1995). Cellular basis of working memory. *Neuron*, 14(3), 477–485. [PubMed: 7695894]
- Heilbronner SR, & Hayden BY (2016). Dorsal Anterior Cingulate Cortex: A Bottom-Up View. *Annu Rev Neurosci*, 39, 149–170. 10.1146/annurev-neuro-070815-013952 [PubMed: 27090954]

- Howard JD, Gottfried JA, Tobler PN, & Kahnt T (2015). Identity-specific coding of future rewards in the human orbitofrontal cortex. *Proc Natl Acad Sci U S A*, 112(16), 5195–5200. 10.1073/pnas.1503550112 [PubMed: 25848032]
- Howard JD, & Kahnt T (2017). Identity-Specific Reward Representations in Orbitofrontal Cortex Are Modulated by Selective Devaluation. *J Neurosci*, 37(10), 2627–2638. 10.1523/JNEUROSCI.3473-16.2017 [PubMed: 28159906]
- Jaeger H (2001). The “Echo State” Approach to Analysing and Training Recurrent Neural Networks (German National Research Center for Information Technology GMD Technical Report, Issue.
- Kennerley SW, Behrens TE, & Wallis JD (2011). Double dissociation of value computations in orbitofrontal and anterior cingulate neurons. *Nat Neurosci*, 14(12), 1581–1589. 10.1038/nn.2961 [PubMed: 22037498]
- Kennerley SW, & Wallis JD (2009). Reward-Dependent Modulation of Working Memory in Lateral Prefrontal Cortex. *Journal of Neuroscience*, 29(10), 3259–3270. 10.1523/jneurosci.5353-08.2009 [PubMed: 19279263]
- Kimmel DL, Elsayed GF, Cunningham JP, & Newsome WT (2020). Value and choice as separable and stable representations in orbitofrontal cortex. *Nat Commun*, 11(1), 3466. 10.1038/s41467-020-17058-y [PubMed: 32651373]
- Knutson B, Fong GW, Bennett SM, Adams CM, & Hommer D (2003). A region of mesial prefrontal cortex tracks monetarily rewarding outcomes: characterization with rapid event-related fMRI. *NeuroImage*, 18(2), 263–272. [PubMed: 12595181]
- Knutson B, Taylor J, Kaufman M, Peterson R, & Glover G (2005). Distributed neural representation of expected value. *The Journal of Neuroscience*, 25(19), 4806–4812. [PubMed: 15888656]
- Kobayashi S, Pinto de Carvalho O, & Schultz W (2010). Adaptation of reward sensitivity in orbitofrontal neurons. *J Neurosci*, 30(2), 534–544. 10.1523/JNEUROSCI.4009-09.2010 [PubMed: 20071516]
- Kolling N, Behrens T, Wittmann MK, & Rushworth M (2016). Multiple signals in anterior cingulate cortex. *Curr Opin Neurobiol*, 37, 36–43. 10.1016/j.conb.2015.12.007 [PubMed: 26774693]
- Lee H, Choi W, Park Y, & Paik S-B (2020). Distinct role of flexible and stable encodings in sequential working memory. *Neural Networks*, 121, 419–429. [PubMed: 31606611]
- Levy R, & Goldman-Rakic P (1999). Association of storage and processing functions in the dorsolateral prefrontal cortex of the nonhuman primate. *The Journal of Neuroscience*, 19(12), 5149–5158. [PubMed: 10366648]
- Lim SL, O’Doherty JP, & Rangel A (2011). The decision value computations in the vmPFC and striatum use a relative value code that is guided by visual attention. *J Neurosci*, 31(37), 13214–13223. 10.1523/JNEUROSCI.1246-11.2011 [PubMed: 21917804]
- Luk CH, & Wallis JD (2013). Choice coding in frontal cortex during stimulus-guided or action-guided decision-making. *J Neurosci*, 33(5), 1864–1871. 10.1523/JNEUROSCI.4920-12.2013 [PubMed: 23365226]
- Lundqvist M, Rose J, Herman P, Brincat SL, Buschman TJ, & Miller EK (2016). Gamma and Beta Bursts Underlie Working Memory. *Neuron*, 90(1), 152–164. 10.1016/j.neuron.2016.02.028 [PubMed: 26996084]
- Maass W, Joshi P, & Sontag ED (2007). Computational aspects of feedback in neural circuits. *PLoS Comput Biol*, 3(1), e165. 10.1371/journal.pcbi.0020165
- Maass W, Natschläger T, & Markram H (2002). Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput*, 14(11), 2531–2560. 10.1162/089976602760407955 [PubMed: 12433288]
- Machens CK, Romo R, & Brody CD (2005). Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science*, 307(5712), 1121–1124. 10.1126/science.1104171 [PubMed: 15718474]
- Masse NY, Yang GR, Song HF, Wang XJ, & Freedman DJ (2019). Circuit mechanisms for the maintenance and manipulation of information in working memory. *Nat Neurosci*, 22(7), 1159–1167. 10.1038/s41593-019-0414-3 [PubMed: 31182866]
- Matsumoto K, Suzuki W, & Tanaka K (2003). Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science*, 301(5630), 229–232. 10.1126/science.1084204 [PubMed: 12855813]

- McGinty VB, Rangel A, & Newsome WT (2016). Orbitofrontal Cortex Value Signals Depend on Fixation Location during Free Viewing. *Neuron*, 90(6), 1299–1311. 10.1016/j.neuron.2016.04.045 [PubMed: 27263972]
- Meyers EM (2018). Dynamic population coding and its relationship to working memory. *J Neurophysiol*, 120(5), 2260–2268. 10.1152/jn.00225.2018 [PubMed: 30207866]
- Meyers EM, Freedman DJ, Kreiman G, Miller EK, & Poggio T (2008). Dynamic population coding of category information in inferior temporal and prefrontal cortex. *J Neurophysiol*, 100(3), 1407–1419. 10.1152/jn.90248.2008 [PubMed: 18562555]
- Mora F, Avrith DB, Phillips AG, & Rolls ET (1979). Effects of satiety on self-stimulation of the orbitofrontal cortex in the rhesus monkey. *Neurosci Lett*, 13(2), 141–145. [PubMed: 119182]
- Mora F, Avrith DB, & Rolls ET (1980). An electrophysiological and behavioural study of self-stimulation in the orbitofrontal cortex of the rhesus monkey. *Brain Res Bull*, 5(2), 111–115. [PubMed: 6769550]
- Moser EI, & Moser MB (2008). A metric for space. *Hippocampus*, 18(12), 1142–1156. 10.1002/hipo.20483 [PubMed: 19021254]
- Murray EA, Moylan EJ, Saleem KS, Basile BM, & Turchi J (2015). Specialized areas for value updating and goal selection in the primate orbitofrontal cortex. *Elife*, 4. 10.7554/eLife.11695
- Murray JD, Bernacchia A, Roy NA, Constantinidis C, Romo R, & Wang XJ (2017). Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proc Natl Acad Sci U S A*, 114(2), 394–399. 10.1073/pnas.1619449114 [PubMed: 28028221]
- Nieder A, & Dehaene S (2009). Representation of number in the brain. *Annu Rev Neurosci*, 32, 185–208. 10.1146/annurev.neuro.051508.135550 [PubMed: 19400715]
- Olds J, & Milner P (1954). Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain. *J Comp Physiol Psychol*, 47(6), 419–427. 10.1037/h0058775 [PubMed: 13233369]
- Orhan AE, & Ma WJ (2019). A diverse range of factors affect the nature of neural representations underlying short-term memory. *Nat Neurosci*, 22(2), 275–283. 10.1038/s41593-018-0314-y [PubMed: 30664767]
- Padoa-Schioppa C (2009). Range-adapting representation of economic value in the orbitofrontal cortex. *J Neurosci*, 29(44), 14004–14014. 10.1523/JNEUROSCI.3751-09.2009 [PubMed: 19890010]
- Padoa-Schioppa C (2013). Neuronal origins of choice variability in economic decisions. *Neuron*, 80(5), 1322–1336. 10.1016/j.neuron.2013.09.013 [PubMed: 24314733]
- Padoa-Schioppa C, & Assad JA (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature*, 441(7090), 223–226. 10.1038/nature04676 [PubMed: 16633341]
- Padoa-Schioppa C, & Conen KE (2017). Orbitofrontal Cortex: A Neural Circuit for Economic Decisions. *Neuron*, 96(4), 736–754. 10.1016/j.neuron.2017.09.031 [PubMed: 29144973]
- Parthasarathy A, Herikstad R, Bong JH, Medina FS, Libedinsky C, & Yen SC (2017). Mixed selectivity morphs population codes in prefrontal cortex. *Nat Neurosci*, 20(12), 1770–1779. 10.1038/s41593-017-0003-2 [PubMed: 29184197]
- Parthasarathy A, Tang C, Herikstad R, Cheong LF, Yen SC, & Libedinsky C (2019). Time-invariant working memory representations in the presence of code-morphing in the lateral prefrontal cortex. *Nat Commun*, 10(1), 4995. 10.1038/s41467-019-12841-y [PubMed: 31676790]
- Pascanu R, & Jaeger H (2011). A neurodynamical model for working memory. *Neural Netw*, 24(2), 199–207. 10.1016/j.neunet.2010.10.003 [PubMed: 21036537]
- Peters J, & Buchel C (2010). Neural representations of subjective reward value. *Behavioural Brain Research*, 213, 135–141. [PubMed: 20420859]
- Rich EL, & Wallis JD (2016). Decoding subjective decisions from orbitofrontal cortex. *Nat Neurosci*, 19(7), 973–980. 10.1038/nn.4320 [PubMed: 27273768]
- Rich EL, & Wallis JD (2017). Spatiotemporal dynamics of information encoding revealed in orbitofrontal high-gamma. *Nat Commun*, 8(1), 1139. 10.1038/s41467-017-01253-5 [PubMed: 29074960]

- Rigotti M, Barak O, Warden MR, Wang XJ, Daw ND, Miller EK, & Fusi S (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451), 585–590. 10.1038/nature12160 [PubMed: 23685452]
- Roesch MR, & Olson CR (2003). Impact of expected reward on neuronal activity in prefrontal cortex, frontal and supplementary eye fields and premotor cortex. *Journal of Neurophysiology*, 90(3), 1766–1789. 10.1152/jn.00019.2003 [PubMed: 12801905]
- Roesch MR, & Olson CR (2007). Neuronal activity related to anticipated reward in frontal cortex - Does it represent value or reflect motivation? *Linking Affect to Action: Critical Contributions of the Orbitofrontal Cortex*, 1121, 431–446. 10.1196/annals.1401.004
- Rudebeck PH, Saunders RC, Prescott AT, Chau LS, & Murray EA (2013). Prefrontal mechanisms of behavioral flexibility, emotion regulation and value updating. *Nat Neurosci*, 16(8), 1140–1145. 10.1038/nn.3440 [PubMed: 23792944]
- Saez I, Lin J, Stolk A, Chang E, Parvizi J, Schalk G, Knight RT, & Hsu M (2018). Encoding of Multiple Reward-Related Computations in Transient and Sustained High-Frequency Activity in Human OFC. *Curr Biol*, 28(18), 2889–2899.e2883. 10.1016/j.cub.2018.07.045 [PubMed: 30220499]
- Sallet J, Quilodran R, Rothe M, Vezoli J, Joseph J-P, & Procyk E (2007). Expectations, gains, and losses in the anterior cingulate cortex. *Cognitive, Affective, & Behavioral Neuroscience*, 7, 327–336.
- Schoenbaum G, Chiba AA, & Gallagher M (1998). Orbitofrontal cortex and basolateral amygdala encode expected outcomes during learning. *Nat Neurosci*, 1(2), 155–159. 10.1038/407 [PubMed: 10195132]
- Schoenbaum G, & Roesch M (2005). Orbitofrontal cortex, associative learning, and expectancies. *Neuron*, 47(5), 633–636. 10.1016/j.neuron.2005.07.018 [PubMed: 16129393]
- Schoenbaum G, Takahashi Y, Liu TL, & McDannald MA (2011). Does the orbitofrontal cortex signal value? *Ann N Y Acad Sci*, 1239, 87–99. 10.1111/j.1749-6632.2011.06210.x [PubMed: 22145878]
- Shafi M, Zhou Y, Quintana J, Chow C, Fuster J, & Bodner M (2007). Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience*, 146(3), 1082–1108. 10.1016/j.neuroscience.2006.12.072 [PubMed: 17418956]
- Sheahan H, Luyckx F, Nelli, Teupe C, & Summerfield C (2020). Neural state space alignment for magnitude generalization in humans and recurrent networks. *bioArxiv*, doi: 10.1101/2020.1107.1122.215541.
- Spaak E, Watanabe K, Funahashi S, & Stokes MG (2017). Stable and Dynamic Coding for Working Memory in Primate Prefrontal Cortex. *J Neurosci*, 37(27), 6503–6516. 10.1523/JNEUROSCI.3364-16.2017 [PubMed: 28559375]
- Stalnaker TA, Cooch NK, McDannald MA, Liu TL, Wied H, & Schoenbaum G (2014). Orbitofrontal neurons infer the value and identity of predicted outcomes. *Nat Commun*, 5, 3926. 10.1038/ncomms4926 [PubMed: 24894805]
- Vickery TJ, Chun MM, & Lee D (2011). Ubiquity and specificity of reinforcement signals throughout the human brain. *Neuron*, 72(1), 166–177. 10.1016/j.neuron.2011.08.011 [PubMed: 21982377]
- Wallis JD, & Kennerley SW (2010). Heterogeneous reward signals in prefrontal cortex. *Curr Opin Neurobiol*, 20(2), 191–198. 10.1016/j.conb.2010.02.009 [PubMed: 20303739]
- Wallis JD, & Miller EK (2003). Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur J Neurosci*, 18(7), 2069–2081. [PubMed: 14622240]
- Wallis JD, & Rich EL (2011). Challenges of Interpreting Frontal Neurons during Value-Based Decision-Making. *Front Neurosci*, 5, 124. 10.3389/fnins.2011.00124 [PubMed: 22125508]
- Wang XJ (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci*, 24(8), 455–463. 10.1016/s0166-2236(00)01868-3 [PubMed: 11476885]
- Watanabe M (1996). Reward expectancy in primate prefrontal neurons. *Nature*, 382(6592), 629–632. 10.1038/382629a0 [PubMed: 8757133]
- Wilson RC, Takahashi YK, Schoenbaum G, & Niv Y (2014). Orbitofrontal cortex as a cognitive map of task space. *Neuron*, 81(2), 267–279. 10.1016/j.neuron.2013.11.005 [PubMed: 24462094]

- Wimmer K, Nykamp DQ, Constantinidis C, & Compte A (2014). Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nat Neurosci*, 17(3), 431–439. 10.1038/nn.3645 [PubMed: 24487232]
- Wittmann M (2013). The inner sense of time: how the brain creates a representation of duration. *Nat Rev Neurosci*, 14(3), 217–223. 10.1038/nrn3452 [PubMed: 23403747]
- Xie Y, Nie C, & Yang T (2018). Covert shift of attention modulates the value encoding in the orbitofrontal cortex. *Elife*, 7. 10.7554/eLife.31507

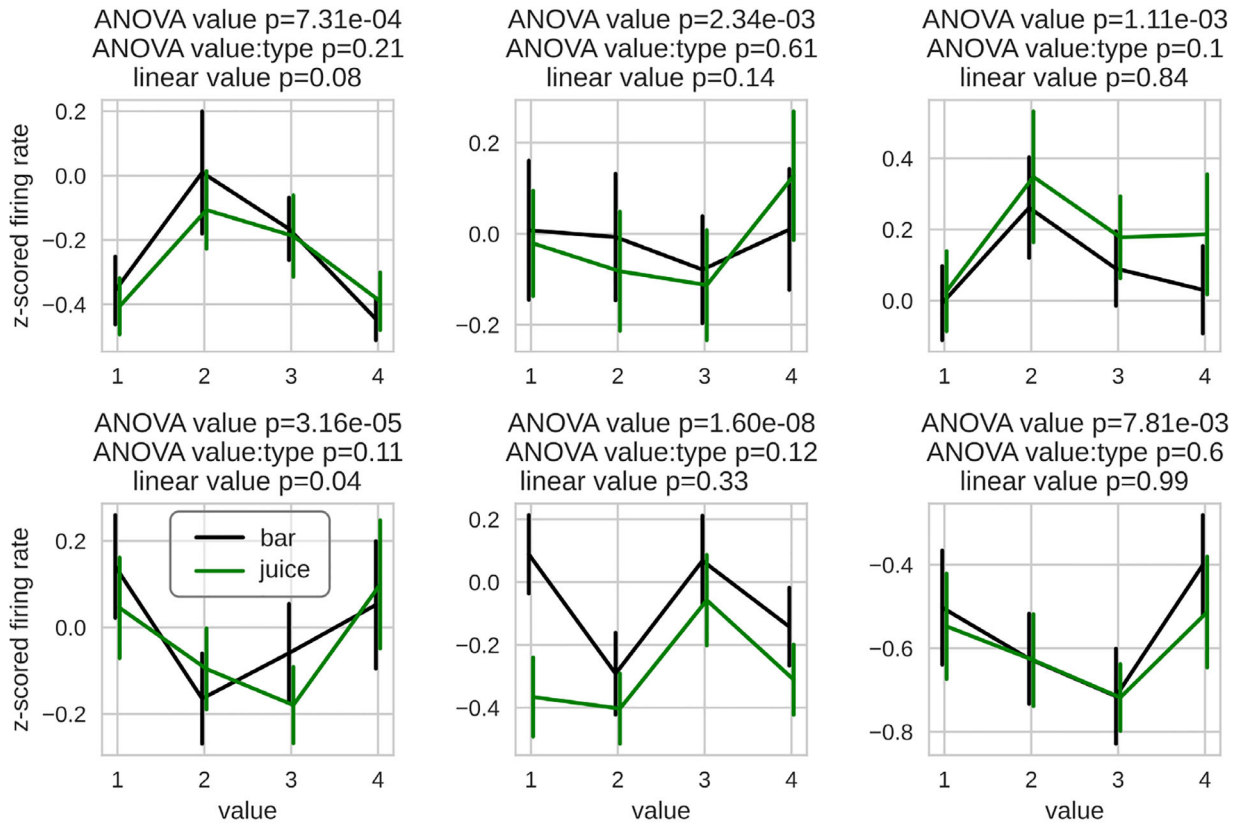


Figure 1. Nonlinear value coding in OFC neurons. Z-scored firing rates of six example neurons recorded from monkeys performing a reward expectation task. Green and black lines show neurons’ responses to cues that predict different types of rewards. Amounts of each reward were titrated so that outcomes assigned the same ordinal value (1 to 4) were chosen with similar probabilities. To be labeled as nonlinear, (1) the coefficient for value in a linear regression must not be significant, (2) the interaction between value and type in an ANOVA is not significant (to avoid cue confound) (3) and the value coefficient is significant in the same ANOVA. Reproduced from Enel et al. (Enel et al., 2020).

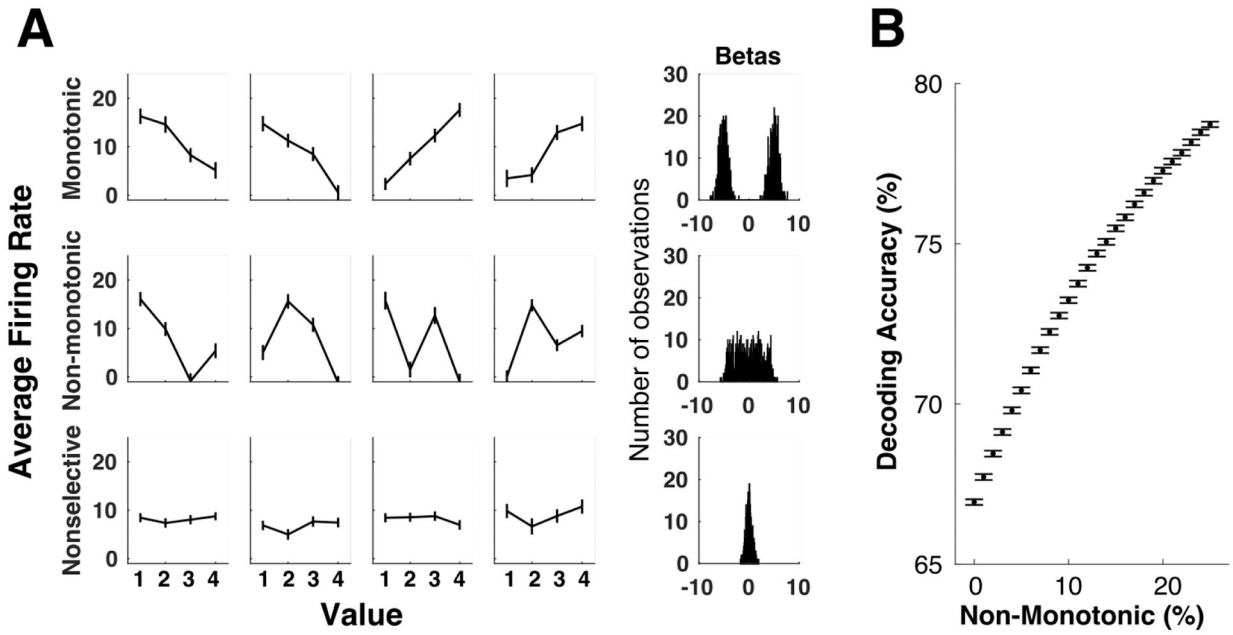
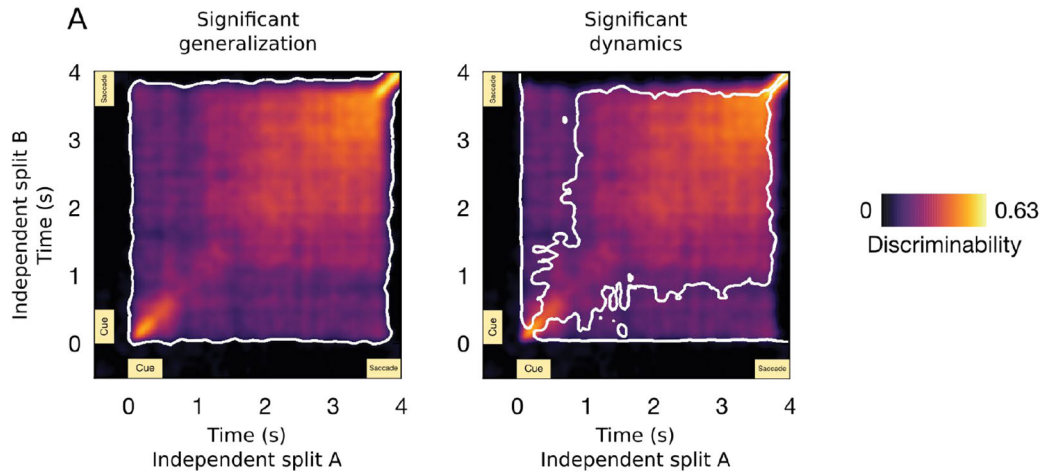
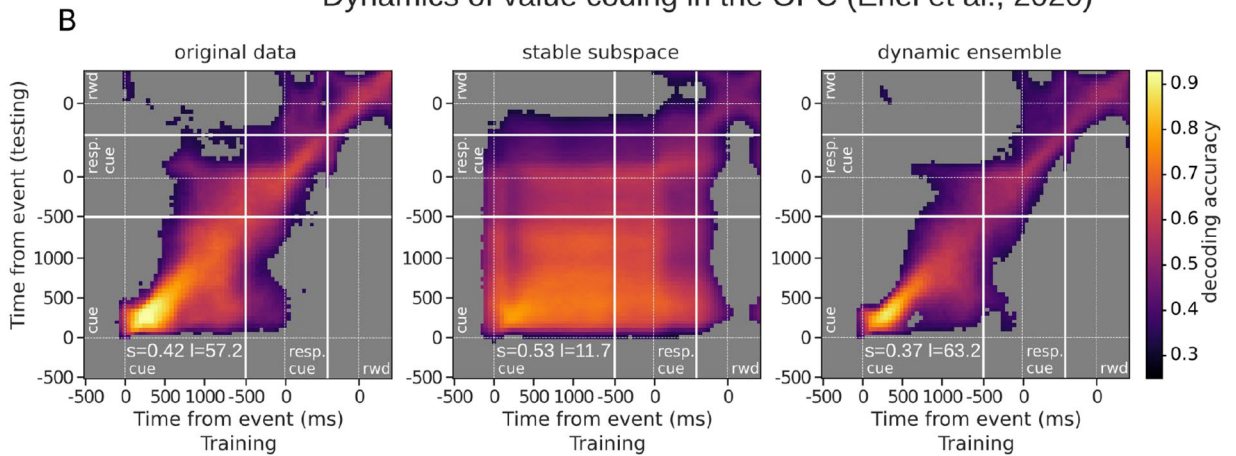


Figure 2. Non-monotonic codes improve value decoding from artificial neuron populations. **A.** Left panels show examples of artificial units designed to monotonically or non-monotonically encode value, as well as nonselective units with the same noise statistics. Each plot shows the mean \pm SEM response across trials of the same ordinal value (1 to 4). Right panels show histograms of regression coefficients for value from 1000 generated units, corresponding to each unit type to the left. **B.** Value decoding increases asymptotically with the percent of non-monotonic value coding neurons in the population. Each point is decoding accuracy (with 10-fold cross validation), averaged across 500 different synthetic populations. Chance decoding of the four simulated values is 25%. Error bars = SEM.

Dynamics of spatial WM in the dIPFC (Spaak et al., 2017)



Dynamics of value coding in the OFC (Enel et al., 2020)

**Figure 3.**

Similarities between spatial and value information maintained across delays. In two studies, cross-temporal analyses compared representations of remembered information across time to determine the extent of generalization. Each pixel is a score resulting from the comparison of representations at two time points, such that the diagonal compares a time point with itself, and higher scores (warmer colors) indicate greater similarity. **A.** Data from Spaak et al. (Spaak et al., 2017), showing dIPFC neurons recorded from monkeys performing a spatial working memory task. Colors represent discriminability scores, which are pairwise correlations of condition differences between time points. Panels show the same scores. White contours on the left indicate significant cross-temporal correlations, and those on the right outline significant reductions in generalizations (i.e. more dynamic signals), suggesting the presence of both stable and dynamic representations in the same dIPFC population. Black = non-significant correlations. **B.** Data from Enel et al. (Enel et al., 2020), showing OFC neurons recorded from monkeys performing a value-based choice task. Colors represent decoding accuracies, where the decoding algorithm was trained and tested on all pairwise combinations of time points. The first panel shows decoding from the original population of firing rates that suggests the presence of both stable and dynamic

representations in the population. The second shows decoding from the population found to have the most stable encoding subspace, and the last shows decoding from the population with the most dynamic encoding. White lines indicate task events. Gray = non-significant decoding.