Perceptual Decision Making in Rodents, Monkeys, and Humans

Timothy D. Hanks^{1,2} and Christopher Summerfield^{3,*}

¹Center for Neuroscience, University of California, Davis, 1554 Newton Court, Davis, CA 95618, USA

²Department of Neurology, University of California, Davis, 4860 Y Street, Suite 3700, Sacramento, CA 95817, USA

³Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD, UK

*Correspondence: christopher.summerfield@psy.ox.ac.uk

http://dx.doi.org/10.1016/j.neuron.2016.12.003

Perceptual decision making is the process by which animals detect, discriminate, and categorize information from the senses. Over the past two decades, understanding how perceptual decisions are made has become a central theme in the neurosciences. Exceptional progress has been made by recording from single neurons in the cortex of the macaque monkey and using computational models from mathematical psychology to relate these neural data to behavior. More recently, however, the range of available techniques and paradigms has dramatically broadened, and researchers have begun to harness new approaches to explore how rodents and humans make perceptual decisions. The results have illustrated some striking convergences with findings from the monkey, but also raised new guestions and provided new theoretical insights. In this review, we summarize key findings, and highlight open challenges, for understanding perceptual decision making in rodents, monkeys, and humans.

Introduction

The study of perceptual decision making within the cognitive and neural sciences seeks to understand how animals detect, discriminate, and categorize information from the senses. Over the past quarter of a century, a canonical theory has emerged of how perceptual decisions are made in the mammalian brain. Inspired by a marriage of quantitative modeling and neural recordings in non-human primates, the proposal states that neurons in sensorimotor areas, prominently including the parietal and dorsal prefrontal cortex, contribute to perceptual decisions by optimizing input signals through repeated sequential sampling and linear integration to a fixed decision threshold (Gold and Shadlen, 2007; Schall, 2003). This work has brought the study of perceptual decisions to the fore within neuroscience and psychology, and has exemplified the benefits of convergent mathematical and biological approaches to understanding brain function.

However, the last 5 years have seen a tremendous diversification of the theories and methods that are available to study perceptual decision making, and have thrown wide open a number of central questions concerning both the computation of decision variables and their expression in neural circuits. Building on the foundational work using monkeys, there has been an expansion toward studies employing rodents, which allow access to a greater range of experimental methods to measure and manipulate neural activity, and studies involving humans, which permit the investigation of a broader range of complex cognitive behaviors. In the first half of this review (section 1), we describe new methodological approaches using monkeys, rodents, and humans, and discuss how the resulting theoretical insights have begun to reshape the field. We separate this into a set of three sections organized around each model system. In the second half of the review (section 2), we discuss three areas that offer great promise for future research linking model sys-

tems. Each of these sections highlights existing work that cuts across species while noting directions where stronger connections may provide insights not possible in any model system alone.

Canonical Perspectives

Perceptual decisions involve the conversion of noisy sensory signals to a discrete motor act. Psychophysical tasks allow researchers to control the nature and quality of sensory input variables, and to reward the animal for specific sensorimotor behaviors. For example, a macaque monkey viewing a random dot kinetogram (RDK; a field of randomly moving dots) might receive a liquid reward for producing a saccade to a spatial target that it has learned to associate with a given motion direction (Newsome and Paré, 1988) (Figure 1A). Where information quality is low (for example, when most dots move randomly, and only some in a coherent direction), decisions can be optimized by repeatedly sampling sensory information and integrating (i.e., summing) the resulting direction estimates over time. Accordingly, a long tradition in mathematical psychology has argued that perceptual decisions are initiated when cumulative estimates of noisy sensory variables reach a criterion response threshold (Wald and Wolfowitz, 1949). In one canonical version of this model, with its roots in decision-theoretic accounts of binary choices, decisions and their latencies are principally controlled by the (drift) rate at which relative information in favor of each of two choices is acquired, the amount of stochasticity (noise) in the representation of this signal, and the level of cumulative information at which choices occur (threshold). This "drift-diffusion" model (DDM) successfully accounts for the empirically measured function that relates decision speed and accuracy, and elaborations thereupon can account for the shape of the observed distribution of response times for both correct and error trials under manipulations of





signal quality, time pressure, and the relative probability or reward value of either response (Bogacz et al., 2006; Ratcliff and Rouder, 1998; Ratcliff et al., 2016).

Building on this algorithmic framework for understanding decision making, the neural mechanisms underlying perceptual choices have been studied in the macaque monkey using extracellular recordings from single neurons. Research has focused on cortical neurons that encode task-relevant sensory signals, such as motion direction in visual area MT (Britten et al., 1993) and vibrational frequency in somatosensory area S1 (Hernández et al., 2000), and those that fire in advance of a choice of the rele-

Figure 1. Challenges for Understanding Perceptual Decision Making in Monkeys

(A) The random dot motion discrimination task that has served as a foundational task for understanding neural contributions to perceptual decision making. A monkey maintains eye fixation while presented with an RDK, consisting of a field of dots that is refreshed at regular intervals. On each refresh, a set fraction of dots are replotted with coherent motion toward a choice target and the rest are replotted randomly. The monkey is rewarded for making a saccade to the target corresponding to the direction of the coherent motion. In the reaction time version of this task, the monkey freely responds when it is ready.

(B) Two models for single-trial dynamics that yield ramping trajectories when averaged across trials. The left panel shows the canonical view that individual trial spike rates meander in the fashion of a drift-diffusion process with a terminating bound at the upper extreme. The right panel shows an alternative model where individual trial spike rates step up or down with the step direction and step time determined stochastically. Reprinted from Latimer et al. (2015).

(C) Unilateral reversible inactivation of LIP with muscimol causes negligible effects on choices in a pulse-based motion discrimination task. A lack of effect is also seen following unilateral LIP inactivation when both choice targets are placed in the hemifield contralateral to the inactivated hemisphere. Reprinted from Katz et al. (2016).

vant option, such as the lateral intraparietal cortex (LIP) and frontal eye fields (FEFs) for saccadic choices (Hanes and Schall, 1996; Shadlen and Newsome, 1996) and premotor cortex for manual choices (Cisek and Kalaska, 2005; Romo et al., 2004). In LIP and FEF, a number of findings support the view that neurons mediate decisions between rival saccadic responses. These findings have been reviewed extensively elsewhere (Gold and Shadlen, 2007; Huk and Meister, 2012), but we summarize them briefly here. First, after a stereotyped dip in activity locked to stimulus onset, average firing rates in LIP and FEF increase steadily when the sensory evidence favors a saccade toward a target in the neuron's response field, and decrease steadily in advance of responses

to the opposing target. Second, the buildup rate depends on the quality of the sensory information, e.g., the level of motion coherence in the stimulus, with steeper slopes for stronger evidence. Third, when response times are controlled by the monkey, firing rates reach a common level prior to responding in the neurons' preferred directions irrespective of signal quality, as if a criterion threshold or bound had been breached (Roitman and Shadlen, 2002). In other words, there is evidence that at least a subset of oculomotor neurons encode cumulative tallies of sensory information that favor a saccade toward a given spatial target, in close correspondence to the drift rate parameter in the DDM

(Gold and Shadlen, 2007). The relative effects of microstimulating MT and LIP on decision accuracy and latencies are also consistent with simulations that increase the rate of arrival of sensory information and boost the cumulative information tally, respectively (Ditterich et al., 2003; Hanks et al., 2006). This work has prompted a compelling argument that LIP and FEF subserve linear integration of sensory inputs during oculomotor choices, providing a neural implementation for approximating the optimal principles of the sequential sampling algorithm proposed by Bayesian decision theory.

Section 1: Insights from New Approaches 1.1 New Challenges for Understanding Perceptual Decision Making in Monkeys

Recently, however, challenges have arisen to the canonical perspective described above. New work with non-human primates has emphasized the heterogeneity of neural responses in putative decision areas, furnished new statistical approaches for unravelling single-trial neural response dynamics, and questioned the causal role of LIP in decision formation (Churchland and Kiani, 2016). However, we argue that these challenges have strengthened, rather than weakened, our foundational understanding of how the brain makes perceptual decisions, by offering new opportunities to disclose the neural mechanisms at the microcircuit level (Murakami and Mainen, 2015).

The link between LIP responding and an integration-to-bound computation was initially forged by pooling average responses over selected subsets of neurons under stereotyped task conditions, where one choice target was positioned centrally in the response field of the recorded neuron and the other target was placed in the opposite hemifield. While there were good reasons to do this, pooling across neurons may obscure response variability (Churchland and Kiani, 2016), and stereotyped experimentation can paint an oversimplified picture of neural encoding (Huk and Meister, 2012; Machens et al., 2010; Murakami and Mainen, 2015). For example, LIP neurons exhibit diverse response motifs that are strongly modulated by sensory and motor-related variables when these are teased apart (Bennur and Gold, 2011; Park et al., 2014) or through experiments that used a flexible association decision task to separate sensory-driven from motor processes (Bennur and Gold, 2011). The diversity of responses also includes transient bursts of activity that scale with the salience of the cues that signal response contingencies (Leathers and Olson, 2012) and responses altered by other decision-irrelevant factors such as saccadic target duration (Bisley et al., 2004; Meister et al., 2013). In other words, neural signals in LIP may only resemble an integration-to-bound signal on aggregate, when heterogenous neuronal responses are averaged together (Bennur and Gold, 2011; Meister et al., 2013; Park et al., 2014). Beyond LIP, the FEF and dorsal striatum have also been shown to have trial-average correlates of an integration-to-bound process in conjunction with other decisionrelated signals (Ding and Gold, 2010, 2012). Interacting sensory and decision factors (such as the input modality and response location) are also encoded during decision tasks in rodent parietal cortex (Raposo et al., 2014). Together, these findings have led to the suggestion that mixed selectivity is a general principle

of neural coding across brain regions and species (Fusi et al., 2016).

At the computational level, one advantage of mixed selectivity is that it allows stimulus or task variables to be mapped onto a wide range of complex responses using simple linear readout operations alone (Fusi et al., 2016; Rigotti et al., 2013). Heterogeneous neural encoding may also allow information to be maintained over multiple distinct timescales (Barak et al., 2013; Bernacchia et al., 2011), potentially providing a circuit mechanism for sustained firing during information integration (Tegnér et al., 2002). However, it remains unclear how the brain meets the computational challenge of finding the correct axis along which to unmix (or "decode") multiplexed information from a neural population. One possibility is that a substantial component of observed heterogeneity arises because recording experiments unwittingly sample neurons with distinct anatomical or neurochemical properties. Knowledge of the cell types, cortical layers, and projection patterns of recorded neurons may be key to understanding the coding of decision information at the level of neural microcircuits, a view we expand on below. Indeed, some accounts have conversely argued that coding of decision information in LIP is remarkably low dimensional, with homogeneous population firing rates sufficient to distinguish among perceptual categories (Fitzgerald et al., 2013; Ganguli et al., 2008). New statistical methods, including multivariate decoding models (Park et al., 2014) and dimensionality reduction techniques (Cunningham and Yu, 2014; Kobak et al., 2016), as well as new multi-electrode recording methods that allow simultaneous data acquisition from multiple neurons (Kiani et al., 2014b), are all likely to be helpful for understanding the pressing issue of how decision information is encoded in neural populations.

The finding that average responses of individual LIP neurons show gradual evidence-dependent firing rate increases does not necessarily imply that these dynamics are present on single trials. For example, sudden step-like changes in activity levels that occur with differing latency will resemble a gradual buildup when averaged together over trials (Figure 1B). Innovative statistical methods have begun to ask whether LIP activity increases are step-like or more gradual. One approach segregates two sources of variance in neural responding: variance proportional to the spike count and a residual variance component that fluctuates between trials. Measured in this way, levels of variability in neural data are insufficient to support a step-like model, but have been argued to favor gradual evidence accumulation during decision making (Churchland et al., 2011; Ding, 2015), although this view has been challenged (Latimer et al., 2015). Another approach uses model comparison to arbitrate among step-like and gradual accounts of single-trial spike train generation. Not surprisingly, the results depend critically on the functional form of the models considered, and in particular on the assumed distribution of step latencies and types of steps allowed. When a uniform distribution is assumed and the step direction is yoked to the choice, dynamics again favor a gradual process for the majority of LIP neurons (Bollimunta and Ditterich, 2012), but models allowing a more flexibly parameterized non-uniform latency distribution with probabilistic correspondence of step direction and choice instead favor step-like dynamics for the majority of LIP neurons (Latimer et al., 2015) (Figure 1B). However, where sensory evidence is presented for a fixed latency before a response is allowed, it becomes challenging to distinguish neural signals that precede an implicit categorical decision from those that follow, and it is unclear how other multiplexed signals described above may affect these sorts of analyses. These contradictions have prompted a lively debate, and it will fall to future work to offer a definitive arbitration among these accounts. One important caveat is that at the population level, a group of neurons that step up and down at different times could have an equivalent impact on downstream neurons to a ramping process, and so either conclusion is potentially compatible with LIP involvement in implementing an evidence accumulation process.

Another challenge to the role of LIP in implementing evidence accumulation comes from recent work examining its necessity for perceptual decision-making tasks. While microstimulation of LIP biases choices and reaction times in an oculomotor decision-making task (Hanks et al., 2006), unilateral pharmacological inactivation that eliminates spiking activity in this region has negligible effects on behavior (Katz et al., 2016) (Figure 1C). This is the case even when both choice options are contralateral to the side of inactivation and thus both contained primarily within response fields of neurons that are inactivated. Unilateral LIP inactivation does, however, robustly bias free saccadic choices away from the contralesional side (Balan and Gottlieb, 2009; Katz et al., 2016; Wardak et al., 2002; Wilke et al., 2012). Together, these findings suggest that LIP does not play an obligatory role in evidence accumulation, but it leaves open the possibility that LIP participates alongside other brain regions. During the formation of the oculomotor decisions used in these experiments, deviations in artificially induced saccades depend on accumulated evidence (Gold and Shadlen, 2000), so one possibility is that other parts of the oculomotor system with which LIP is heavily interconnected may afford behavioral compensation during evidence accumulation. This question can, in theory, be addressed by inactivation methods that are faster than the timescales of compensation and/or through simultaneous perturbation of multiple brain regions. In the next section, we describe new approaches for studying perceptual decision making in rodents that will very likely facilitate these sorts of experiments. 1.2 New Approaches for Studying Perceptual Decision

Making in Rodents

Recently, many researchers have turned to rodent models to investigate perceptual decision making (Carandini and Churchland, 2013). Remarkably, the rodent cognitive repertoire seems well suited to the study of many key elements of perceptual decision making, including evidence accumulation (Brunton et al., 2013; Hanks et al., 2015), perception-to-action remapping (Duan et al., 2015), and even decision confidence (Kepecs et al., 2008). The rodent model has promoted a diversification of sensory modalities under consideration, to include auditory (Brunton et al., 2013; Znamenskiy and Zador, 2013), tactile (Guo et al., 2014), olfactory (Uchida et al., 2006), and multisensory (Raposo et al., 2014), as well as the visual studies that dominate among non-human primate work. While more evolutionarily distant from humans, rodents confer a number of advantages over monkeys that have led to new insights into the circuit mechanisms of perceptual decision making. First, rodents are most readily amenable to the application of cuttingedge technologies for neural measurement and manipulation, such as optogenetics (Grosenick et al., 2015; Luo et al., 2008), calcium imaging (Guo et al., 2014), and cell-type- and projection-specific tagging of neurons (Znamenskiy and Zador, 2013). Second, rodents increase the feasibility of large-scale data collection, including multi-area mapping of the effects of the neural perturbations (Guo et al., 2014) and semi- or fully automated, high-throughput training (Brunton et al., 2013; Poddar et al., 2013). A third, and perhaps underappreciated, factor is that lower investment costs incurred by rodent research encourage a shift toward higher risk, higher reward experiments, offering an opportunity to advance the frontiers of decision research more rapidly.

Gradual, signal-dependent neural buildup has now been observed in multiple rodent brain areas during accumulation of decision evidence in the auditory domain. These include the rat posterior parietal cortex (PPC) and frontal orienting fields (FOFs), putative homologs of monkey PPC and FEF (Brody and Hanks, 2016). One key line of research has relied on a new "Poisson click" task that involves discriminating whether auditory pulses occurred more frequently in a stream coming from the left or right (Figure 2A). Like the "Weather Prediction" task previously used in humans and monkeys (Kira et al., 2015; Knowlton and Squire, 1993; Yang and Shadlen, 2007), this "discretepulse" approach has the advantage that the stream of sensory stimulation is fully known on each trial, allowing researchers to estimate precisely the state of the accumulator at each time point in the trial. Linking this estimate to neural signals permits a characterization of neural tuning curves that describe how single-cell responses map onto cumulative decision variables. For example, the integration-to-bound framework implies that each pulse should provoke a sustained increment in the firing rate of the neuron, a prediction that has been validated in PPC (Hanks et al., 2015) and is consistent with similar analyses in monkey LIP in a visual decision task (Huk and Shadlen, 2005). In trying to link these findings to the canonical perspectives provided by monkey work, it is important to note the differences in sensory modality and neuronal selection. In contrast to the monkey work that typically involved visual targets with one centered in the response field of a neuron pre-screened for selectivity, the rodent work involved many neurons recorded simultaneously with a fixed geometry for the sources of auditory information and screening for selectivity performed afterward.

Although monkey studies have emphasized the commonalities among decision signals in LIP and FEF, rodent work has revealed key differences in parietal and frontal tuning curves for accumulated evidence. Whereas firing rates display an approximately linear relationship with the accumulator value in PPC, in FOF they change more abruptly as the accumulator switches sign (i.e., as accumulated evidence comes to favor a leftward, rather than rightward, choice) (Figure 2B). In other words, prefrontal neurons encode a more discretized, categorical signal than parietal neurons during evidence accumulation (Hanks et al., 2015). This presumably makes prefrontal neural signals more robust to rapidly fluctuating noise in sensory signals, and suggests that they may exhibit dynamics that are even more



Figure 2. New Approaches for Studying Perceptual Decision Making in Rodents

(A) The "Poisson clicks" task, a rodent evidence accumulation task inspired by the motion discrimination task used in monkeys. In this auditory discrimination task, rats fixate their nose while presented with two competing streams of auditory clicks, one from the left and one from the right. These clicks are generated by Poisson processes with different underlying rates. Rats are rewarded for making orienting movements when cued to the nose port corresponding to the side that played the greater number of clicks.

(B) Neural "tuning curves" for accumulated evidence for pre-movement side-selective PPC and FOF neurons during the Poisson clicks task. This shows the relationship between the accumulator value from a continuous-time model of the decision process and neural firing rate. PPC has a more linear encoding of accumulated evidence compared to the more step-like encoding of FOF. The latter seems more tightly linked to the choice supported by the accumulated evidence. Reprinted from Hanks et al. (2015).

(C) Bilateral inactivation of rat PPC has negligible effects on choices in the Poisson clicks task. Endpoints on each side show control trials with a constantly illuminated LED cuing the correct side that therefore did not require evidence accumulation. Reprinted from Erlich et al. (2015).

(D) Method to selectively stimulate neurons in primary auditory cortex that project to the striatum. First, the striatum is injected with herpes simplex virus-1 (HSV) that expresses Cre recombinase. HSV is transported in retrograde fashion up axons to cell bodies projecting to the injected region, some of which reside in primary auditory cortex. This region is then injected with a Cre-dependent adeno-associated virus (AAV) driving expression of channelrhodopsin-2 (ChR2) in neurons co-infected with Cre—that is, those projecting to the striatum. ChR2 is a light-sensitive ion channel that allows the artificial stimulation of neurons that express it. Thus, this allows for the selective stimulation of primary auditory cortex neurons that project to the striatum without stimulating those that do not project to the striatum. Reprinted from Znamenskiy and Zador (2013).

step-like than PPC. We suggest that analyses such as these that estimate the accumulator value on individual trials may provide more traction for resolving whether neural responses on individual trials follow step-like or ramp-like dynamics (Latimer et al., 2015), one of the challenges raised in the previous section.

Robustness of prefrontal representations may depend on time to recover, as demonstrated by temporally precise, optogenetically mediated inactivation methods available in rodent models. For example, inactivation of a related frontal region in the mouse, the anterior lateral motor cortex (ALM), only biases licking choices in a tactile decision task near the time of decision report, but not earlier (Li et al., 2016). Neural encoding of choice-related signals recovered rapidly following this brief window, unless inactivation was bilateral or interhemispheric communication was blocked with callosal resection. Similarly, in the rat Poisson click task, a bias to the ipsilateral stream was provoked by unilateral FOF inactivation that occurred late, but not early, in the accumulation epoch (Hanks et al., 2015). This evidence for a robust, choice-related signal in the prefrontal cortex of both rats and mice suggests that it may correspond to a broader property of macrocircuitry that could be observed in other species, including monkeys and humans.

Interestingly, pharmacological inactivation of PPC had little or no effect on perceptual decisions in the rat (Figure 2C), but did impair free choices that did not depend on evidence accumulation (Erlich et al., 2015), which was mirrored by the monkey work highlighted above. This is also consistent with a further rodent study employing both pharmacological inactivation and optogenetic perturbation methods, which found no effect of PPC inactivation or perturbation on decisions about auditory pulse frequency (Raposo et al., 2014). That same study did find that inactivation reduced discrimination sensitivity in a visual pulse task, but the pattern of results was more consistent with PPC playing an auxiliary role, rather than being directly involved in evidence accumulation (Licata et al., 2016). Together with the monkey work described above, these studies furnish an emerging picture in which the contributions of the parietal cortex to perceptual decisions are not causally related to evidence accumulation per se. Rather, it may relate to auxiliary processes that can usefully employ an evidence accumulation signal and contribute to decision making in conjunction with other brain regions.

Another one of the challenges alluded to above is how to make sense of the heterogenous coding properties of putative decision neurons. Recent work has begun to shed light on this question by stimulating distinct classes of neurons on the basis of their axonal projection targets. This is currently possible in the rodent, while still in nascent stages in the monkey, through the use of retrogradely transported viruses that control of the expression of optogenetic constructs in a projection-specific manner (Figure 2D). For example, when rats were trained to report whether a stream of auditory tones was composed of predominantly higher or lower frequencies, stimulation of tonotopic regions of primary auditory cortex only biased choices toward the corresponding frequency when the targeted neurons had striatal projections (Znamenskiy and Zador, 2013). Relatedly, response heterogeneity in motor preparatory activity during a whisker-based object location discrimination task can be partly

explained by patterns of axonal projections (Li et al., 2015). Specifically, while layer 5 neurons in mouse ALM that project within the cortex have mixed selectivity for ipsilateral and contralateral movements, those that project to the brainstem are biased toward contralateral movements. In other words, the projection targets of relevant neurons may be a key factor in understanding response heterogeneity in perceptual decision tasks. We also note that this study combined optogenetics and cell-type-specific imaging in a total of 52 mice, a prohibitive cohort size in monkey research. These findings again allude to the importance of understanding how local circuit and long-range connections contribute to decision making, a theme that may guide future work in other model systems.

1.3 New Paradigms and Neural Decision Signals in Humans

Unlike monkeys and rodents, humans can readily perform complex decision tasks with minimal instruction. Moreover, single-cell studies typically focus on just a handful of recording sites, whereas macroscopic imaging techniques such as fMRI offer the promise of revealing the wider brain networks involved in making perceptual decisions. Thus, building on work in experimental animals, over the past decade cognitive neuroscientists have sought to identify canonical signatures of perceptual decision making in humans and to characterize the neural systems involved at a whole-brain level.

One major challenge is that fMRI measures neural activity indirectly and aggregates over the activity of many millions of neurons within a single voxel. This makes it difficult to directly probe for neural selectivity to sensory or decision variables (e.g., motion direction) in psychophysical tasks such as the RDK paradigm. One solution to the former problem has been to devise discrimination tasks involving visual stimuli that preferentially activate contiguous clusters of voxels within the extrastriate cortex, such as faces and buildings (Heekeren et al., 2008). Like firing rates in sensory regions (such as MT or the auditory cortex), blood oxygenation level-dependent (BOLD) amplitude in these regions scales with the relative strength of evidence for each category (e.g., visibility of face > house in the fusiform region). One approach has searched for brain regions that correlate with the difference in relative BOLD signal amplitude in these extrastriate areas, in line with the assumption that the decision variable is a cumulative differential of sensory inputs. Like recent rodent work, this approach has identified lateral portions of the prefrontal cortex as a candidate structure (Heekeren et al., 2004), and indeed, temporary inactivation of this region with transcranial magnetic stimulation (TMS) dampens the estimated drift rate in human perceptual decision tasks (Philiastides et al., 2011).

However, due to the torpid nature of the BOLD signal, fMRI is poorly suited to measuring neural dynamics on a millisecond scale, making it difficult, for example, to distinguish neural signals that occur before or after a decision. Accordingly, a consensus has yet to emerge about how perceptual evidence accumulation is expressed in fMRI signals (Mulder et al., 2014). In rodents and monkeys, firing rates in posterior parietal neurons grow as information is integrated toward a saccade or manual action. By contrast, BOLD signals in homologous parietal regions of the human tend to vary inversely with the level of evi-

Neuron Review

dence in a perceptual stimulus, scaling instead with levels of decision uncertainty or conflict. A similar pattern is observed in other cortical regions, prominently including the anterior cingulate cortex and anterior insular cortex (Ho et al., 2009; Liu and Pleskac, 2011; Wheeler et al., 2015). One explanation for this apparent contradiction is that during speeded decisions, the highest aggregate spiking activity in target-selective neurons will be observed when evidence is weak or ambiguous because on those trials, responses are slower and the buildup in firing (although shallower) is more prolonged. Under the framework of the DDM, the BOLD signal in parietal, cingulate, and insular cortex is successfully predicted by the cumulative excursion from zero of the best-fitting decision variable for a given trial or condition (Basten et al., 2010; Ho et al., 2009). However, this theory fails to explain why BOLD signals continue to scale with decision uncertainty even in fixed-response settings, where integration latencies are presumably constant, or why parietal BOLD signals are stronger under regimes that favor speed over accuracy (van Veen et al., 2008) (see below). An alternative account appeals to the intuition that in LIP, the numbers of neurons that are selective for the target response (e.g., a saccade to the required target) are typically outnumbered by those coding for competing responses. In the presence of other computational mechanisms, such as divisive normalization, many neurons may be silenced when sensory signals are most reliable. Even within target-selective neurons, stronger sensory signals elicit steeper firing rate slopes but tend to build up from a lower level on these trials, potentially reducing aggregate firing (Meister et al., 2013). Together, these factors further complicate the relationship between sensory evidence accumulation and the BOLD signal. Thus, for now predictions about how fMRI signals vary during rapid information integration necessarily rest on untested assumptions about the link between the BOLD signal and underlying neural dynamics.

Unlike fMRI, magneto/electroencephalography (M/EEG) allows the neural consequences of decision formation to be directly charted with millisecond resolution (Kelly and O'Connell, 2015). Once again, inventive experiments have been conducted that dissociate relevant task variables at the macroscopic level. For example, when opposing perceptual decisions involve lateralized hand responses, the amplitude of high-frequency (gamma and beta band) MEG activity over motor regions diverges steadily between hemispheres that are contralateral and ipsilateral to the response effector (Donner et al., 2009). These signals are attenuated on error trials, implying that they are not merely response related, and like signals in LIP, they begin earlier when an advance cue signals the likely motion direction (de Lange et al., 2013) (see below). More recently, EEG studies have identified a positive potential recorded over midline parietal electrodes that grows in a signal-dependent fashion as sensory evidence accumulates (Kelly and O'Connell, 2013; O'Connell et al., 2012). This potential (known as the CPP) terminates at a fixed plateau, mirroring the firing rate acceleration observed in LIP neurons (Figure 3A). Intriguingly, the same dynamics are observed during detection of deviant stimuli that typically elicit the classic P300 potential, with which the CPP shares a scalp topography, suggesting that although they are differently named, these two potentials may



Figure 3. Neural Data and Models for Human Perceptual Decision Making

(A) The centro-parietal positivity (CPP) is a scalp EEG signal recorded over medial parietal electrodes (inset; red shows the maximum) that shows build-tothreshold dynamics during the random dot motion task (inset). The buildup varies with coherence level (left panel), showing the steepest ramp for high coherence events. The buildup terminates at a common plateau, irrespective of coherence left (right panel). Reprinted from Kelly and O'Connell (2013).

(B) Neural signals during accumulation of information from discrete samples. Right panel: a parietal potential co-varies negatively (at 250 ms) and then positively (at 500 ms) with decision information. Center: the strength of neural encoding at 500 ms on each sample k positively predicts the behavioral weight given to sample k, but negatively predicts the weight attributed to samples k + 1 and k - 1. Right: this is indicative of a rhythmic gain control mechanism (shown schematically for a tilt averaging task, right panel). Reprinted from Wyart et al. (2012).

(C) Left: selective integration model. Noisy sensory inputs are accumulated in parallel streams. A central "bottleneck" stage mediates competition among samples, processing the winning sample with relatively higher gain, controlled by a selective gating parameter. "Late" noise also arises during integration. Right: under this model, when late noise is absent, perfect integration yields greatest accuracy (cyan line). However, as late noise grows, simulated observers with higher late noise perform better. The same is true for human participants. Reprinted from Tsetsos et al. (2016).

be common manifestations of a dynamically growing decision signal (Twomey et al., 2015).

An important next step for cognitive neuroscientists is to move beyond the search for human signatures of sensory evidence accumulation that mirror those found in the monkey, toward a computational account of how those signals underpin perceptual decisions. One promising approach, related to the pulsatile stimulation approach favored by rodent researchers (Brunton et al., 2013; Raposo et al., 2014), has been to develop tasks that involve categorizing the information in a stream of discrete events (or "samples"), each of which conveys partial evidence about the correct response. For example, participants might be asked to average the tilt in a sequential stream of visual gratings (Cheadle et al., 2014; Wyart et al., 2012), the most frequent direction among a series of arrows (de Lange et al., 2010), or the average size in a succession of objects (Gorea et al., 2014; Hubert-Wallander and Boynton, 2015). This allows researchers to characterize the perceptual information (the disparity between one sample and the next), decision information (the momentary information conveyed by each sample), and response information (the cumulative evidence for one choice over another) that accompany each sample. In conjunction with EEG recordings, these quantities can be regressed against neural signals to reveal their respective encoding over the occipital, parietal, and motor cortices, respectively (Wyart et al., 2012). Going a step further, it is then possible to link the single-trial residuals from this analysis to choices, disclosing how encoding at each neural stage predicts behavior, an approach related to the calculation of neural "choice probabilities" in single-cell recordings (Nienborg and Cumming, 2009). This class of analysis has revealed that the strength of parietal encoding of momentary decision information predicts the multiplicative weight (or influence) that each sample carries in the final choice (Figure 3B), whereas the premotor signal scales with an additive bias to respond with the left or right hand (Wyart et al., 2012). In other words, sensory information may be transformed in distinct, sequential multiplicative and additive stages in the parietal and premotor cortex, respectively.

A key open question concerns the sources of noise or loss that corrupt perceptual decisions (Brunton et al., 2013; Hunt, 2014; Renart and Machens, 2014; Scott et al., 2015; Wyart and Koechlin, 2016). While some have suggested that human decisions are limited mainly by noise arising during sensory encoding (Körding, 2007), a traditional view from cognitive science has emphasized that human information processing is capacity limited, and that sensory information must pass through a central "bottleneck" before influencing decisions (Broadbent, 1958). One emerging proposal argues that parietal cortex may filter information in time, acting as just such a bottleneck during human perceptual decisions. Over the parietal cortex, neuroelectric activity measured with EEG fluctuates slowly, and the influence that each sample of evidence wields over the decision depends on its timing with respect to this rhythm, with samples falling at the preferred phase of parietal oscillations carrying more weight, and those falling in the anti-preferred phase being relatively overlooked (Spitzer et al., 2016; Wyart et al., 2012) (Figure 3B). This suggests that the central bottleneck may occur because a fluctuating gain control mechanism resolves competition among temporally proximal samples of information, a form of "active sensing" that has also been observed in monkey sensory cortices (Lakatos et al., 2008). This variable-gain mechanism might explain why behavioral data in decision tasks are best explained by models that incorporate between-trial drift rate variability in addition to within-trial noise (Ratcliff and Rouder, 1998). A similar approach has also been used to understand how attention modulates perceptual decisions, with the observation that dividing attention does not dampen multiplicative weighting of sensory inputs over the parietal cortex, but incurs a loss at a later stage, as if information "leaks" away during additive integration in motor circuits (Wyart et al., 2015). Further work has explored competition arising between sequential samples within a trial, suggesting that undue weight is given to information that is consistent with the current cumulative decision variable, and that these consistent samples are encoded more strongly over parietal cortex, a form of confirmation bias in human perceptual choices (Cheadle et al., 2014).

Cognitive studies of human judgment have traditionally emphasized that decisions are limited not just by sensory noise but also by processing capacity, with an additional source of information loss incurred as perceptual signals are discretized for maintenance in central circuits. Limitations in central processing

Neuron Review

may be one reason why humans show irrational biases and reversals of preference that fail to maximize financial outcomes (Kahneman, 2011). To provide a process-level explanation for these economic suboptimalities, recent work has extended the sequential sampling framework that underpins perceptual choices (Busemeyer and Townsend, 1993; Krajbich et al., 2010; Summerfield and Tsetsos, 2015). One key motivating intuition is that each discrete sample (e.g., a grating) conveys unique information about the probability that an alternative will be rewarded, just as each attribute of an economic prospect partially signals its value (e.g., when purchasing a car, one might consider the price, reliability, and fuel economy). A new sequential sampling model, known as "selective integration," argues that when choosing among two simultaneous streams of discrete samples, simultaneously occurring events compete for limited neural resources, with humans attributing greater multiplicative weight to the sample (e.g., a symbolic number) with higher value, as if they attend preferentially to salient choice attributes (Tsetsos et al., 2012) (Figure 3C). This theory offers a process-level account of human violations of axiomatic rationality, including framing effects, intransitivity, and non-independence from irrelevant alternatives (Tsetsos et al., 2012, 2016). Interestingly, although selective integration is suboptimal for an ideal observer, it can be shown to maximize economic outcomes under the assumption that decisions are limited not just by early sensory uncertainty, but also an additional, "late" source of noise that arises when multiple attributes or features are combined to make a decision (Scott et al., 2015; Tsetsos et al., 2016; Wyart and Koechlin, 2016) (Figure 3C). These paradigms provide a new opportunity to harness the sequential sampling framework to offer a normative account of human economic biases, uniting the study of perceptual and value-guided choices under a common framework (Summerfield and Tsetsos, 2012). However, more work is needed to understand the neural underpinnings of these effects, work that will likely benefit from insights garnered from multiple model systems. In the following sections, we highlight three topics within the study of perceptual decisions where connections have already begun to be formed between findings that cut across species.

Section 2: Opportunities to Bridge across Model Systems

2.1 Deciding When to Decide: The Speed and Accuracy of Perceptual Decisions

For all animal species, perceptual decisions are made in the context of ongoing motivational states. For example, a thirsty monkey will be motivated by the receipt of liquid reward that follows correct trials, or a human participant might wish to complete an onerous psychophysical experiment in the shortest possible time. In order to satisfy ongoing goals, observers must decide both what to decide (e.g., left versus right) and when to decide (now versus later). Studies have begun to address how this is achieved at the neural level in both human and non-human subjects.

Mathematical models of the decision process (such as the DDM) at the algorithmic level propose that responses are initiated when cumulative decision information achieves a fixed criterion value, i.e., when a flat decision "bound" is reached. An



CellPress

Figure 4. Decision Urgency

Evidence for a stimulus-independent "urgency' signal during speeded decisions. Three studies (Hanks et al., 2015; Heitz and Schall, 2012; Thura and Cisek, 2016) show evidence that when monkeys make perceptual decisions under a reward regime that encourages speed over accuracy, decisions build up to threshold faster. The left panel (A) shows shows LIP neurons recorded in a motion discrimination, saccadic decision task. The middle panel (B) shows movement neurons recorded in FEF in a visual search, saccadic decision task. The right panel (C) shows PMd neurons recorded during discrete visual accumulation, reaching decision task. Interestingly, none of these studies report evidence for reduced terminal firing rates (i.e., decision threshold) under speed compared to accuracy emphasis, inconsistent with the predictions provided by the standard DDM.

important part of the decision policy is thus to specify the height of the bound. Because the precision of noisy decision signals gradually increases during sequential sampling, a high bound will lead to slow but more accurate decisions, whereas a low bound will reduce deliberation times at the expense of performance. Participants can thus trade off speed and accuracy by setting the height of the bound, effectively "deciding when to decide" (Bogacz et al., 2010). Extant behavioral data suggest that they do so approximately optimally (Simen et al., 2009).

The claim that firing rates in LIP are involved in the implementation of a DDM-like algorithm by encoding the value of a cumulative decision variable is bolstered by the observation that neural signals reach a common, signal-independent level prior to a saccade, as if a criterion firing rate had been achieved (Roitman and Shadlen, 2002). However, a prediction of this account is that the threshold value should vary systematically under regimes that differentially emphasize speed and accuracy. In the monkey, a number of recent studies have varied incentives to encourage more cautious or incautious responding, while recording from FEF, LIP, premotor cortex (PMd), and primary motor cortex (M1) (Hanks et al., 2014; Heitz and Schall, 2012; Thura and Cisek, 2016). While thought-provoking differences are apparent between these studies, they nevertheless provide an emerging picture of how neural correlates of decision variables change when speed and accuracy are traded off. In all three studies, firing rates under speed emphasis started higher and/or ramped up more quickly than under accuracy emphasis, a feature not predicted by the standard DDM (Figure 4). Further evidence for heightened gain under speed pressure emerges from analysis of local field potentials in FEF (Heitz and Schall, 2013). In addition, none of the studies showed reductions in the neural threshold preceding speeded decisions, the firstorder prediction based on the standard DDM. Even more surprisingly, one of the studies reported a reduced neural threshold under the conditions of accuracy emphasis compared to speed emphasis (Heitz and Schall, 2012). The authors postulated an additional stage of leaky downstream integration to account for this difference, but that suggestion has been more recently challenged (Cassey et al., 2014). Nonetheless, all three studies are in agreement that the standard bounded accumulation account provides an incomplete picture of how neural responses contribute to the timing of perceptual decisions. Interestingly,

both increased pre-decision response levels and evidence-independent ramping for the speeded condition predict overall stronger neural responses under speed pressure, and are thus consistent with the finding that BOLD signals in parietal and prefrontal regions are higher when speed is emphasized (van Veen et al., 2008). These findings illustrate how neural data can help to inform psychological models of perceptual decisions and link human findings to neural mechanisms.

In psychophysical studies, experimenters can choose to manipulate the reliability of sensory signals (e.g., level of motion coherence) over either blocks or trials. In the former case, reward rates will be maximized by the application of a fixed bound, i.e., one with a height that remains constant across each trial within a block (Wald and Wolfowitz, 1949). The optimal height of the bound depends on the observer's belief about the signal reliability, which in the latter case varies over the course of the trial. In particular, after prolonged deliberation, it is more likely that the current trial is one with low signal quality, where the information obtained from further sampling may be limited. Optimal models thus predict that under unknown sensory reliability, the height of the bound should "collapse" over time (Deneve, 2012; Drugowitsch et al., 2012; Frazier and Yu, 2008). However, the empirical question of whether decisions about signals with unknown reliability respect a collapsing bound algorithmically, and how this might be implemented neurally, represents a new frontier for research in perceptual decision making.

One emerging view is that decisions may be driven to the bound by a strong, evidence-independent quantity referred to as an "urgency" signal, which effectively implements a collapsing bound by inflating later accumulator states away from zero. One line of evidence for such a signal can be observed by analyzing trials on which sensory evidence is entirely ambiguous (e.g., 0% coherence), where firing rates in LIP nevertheless grow toward the bound associated with the eventual response (Churchland et al., 2008). Indeed, LIP neurons show growing firing rates up to a response when monkeys estimate the time elapsed across an interval, consistent with an evidence-independent component to their responses (Janssen and Shadlen, 2005). A heightened urgency signal may also explain the generalized increase in neural gain-described above-observed under speed relative to accuracy pressure (Hanks et al., 2014), and complementary findings regarding post-error slowing (Purcell and Kiani,

2016). However, studies that have asked whether fixed or collapsing bounds provide the best explanation of human and monkey behavior in perceptual decision tasks have yielded mixed results. One large-scale analysis of five datasets found that humans performing a limited number of trials were best described by a fixed bound, whereas a dynamic bound captured better the performance of highly trained monkeys (Boehm et al., 2016; Hawkins et al., 2015). Another, more controversial model eschews the sequential sampling approach entirely, proposing that decisions are initiated by the joint influence of an urgency signal and momentary excursion of decision evidence from zero (Thura et al., 2012), but this intriguing account has been challenged on numerous grounds (Winkel et al., 2014). While the interesting notion that a stimulus-independent neural signal controls the timing of decisions is likely to have a strong influence over coming years, its existence remains controversial in humans and monkeys, and it has yet to be systematically explored in rodent model systems. For rodents, this will require the development of tasks complementary to those described here for humans and monkeys where "deciding when to decide" becomes more important, as opposed to most of the rodent work that we described above (where the timing of the decision report was cued). Techniques available in rodents hold the promise of helping to determine the circuit mechanisms responsible for urgency signals, which will help to explain their source, and allow more refined tests of their role in decision making.

2.2 Modulation of Perceptual Decisions by Probability and Value

Another topic that has begun to see complementary approaches using multiple model systems is the study of how probability and value modulate perceptual decisions. In natural environments, sensory stimuli occur with differing frequency and are associated with differing reward and punishments. One long-standing concern in mathematical psychology is how noisy sensory evidence is combined with contextual information encoding the probability or value of choice alternatives, in order to produce an optimal decision. Recently, studies have begun to explore how humans and monkeys achieve this at both the behavioral and neural levels.

Where one response (e.g., leftward target) is associated with more positive or less negative outcomes, reward rates will be maximized by selecting it more often. Similarly, stimuli (e.g., leftward motion) that have higher base rates of presentation should elicit responses more readily than those that occur infrequently. Where signal quality is known, this can be achieved under the sequential sampling framework via a simple additive offset to the starting point of evidence accumulation toward the expected bound (Bogacz et al., 2006). However, when signals are of unknown sensory reliability, for example, where coherence levels are randomly intermixed between trials, an additional sensory time-varying component is required for optimal responding (Moran, 2015). This is because prior information about stimulus probability or value can be deployed most effectively when sensory signals are weak or ambiguous, which is usually only evident later in the deliberation process. Multiple studies focused on LIP have asked how perceptual decisions are modulated by stimulus value or probability, and these have consistently observed an early increase in firing rates when the

Neuron Review

saccadic target associated with the more probable or valuable response falls within the neuron's receptive field. This occurs irrespective of whether responses are speeded (Hanks et al., 2011) or where integration latencies are controlled by the experimenter (Gold et al., 2008; Rao et al., 2012; Rorie et al., 2010), and is consistent with LIP recordings conducted while the monkey makes decisions about unambiguous sensory signals of varying probability or value, confirming the view that the LIP signal incorporates a range of contextual variables relevant to saccadic choice (Platt and Glimcher, 1999). However, evidence for a modulation that occurs during evidence accumulation has only emerged in conjunction with reaction time decisions, in the form of heightened stimulus-independent drift toward the bound for the expected choice. This signal is akin to a biased version of the urgency signal described above (Hanks et al., 2011), suggesting a possible shared neural circuit mechanism for both. We expect newer techniques available in rodents to help tease apart contributions from distinct brain sources or microcircuit components in testing this hypothesis.

Several human studies have also attempted to address this question, typically using a combination of computational modeling and functional neuroimaging (Summerfield and de Lange, 2014). In humans, an additive offset in the origin of evidence accumulation typically provides the best-fit behavioral data (Mulder et al., 2012), and fMRI studies have reported correlates of a bias signal in sensory cortex (Kok et al., 2014), as well as parietal and prefrontal regions (Chen et al., 2015; Mulder et al., 2012; Summerfield and Koechlin, 2010), during perceptual decisions about visual stimuli. However, it is hard to know whether visual bias signals in these studies are a consequence, or a cause, of any effects observed in higher brain areas. More generally, a popular framework (known as "predictive coding") suggests that priors and sensory evidence are integrated through reciprocal interactions between decision (e.g., parietal) and sensory (e.g., visual) regions, and this framework has been used to interpret a wide range of human imaging data in relation to this topic. The theory makes specific predictions about various types of neural response that would be observed in the cortical microcircuitry (Bastos et al., 2012), but single-cell researchers are only just beginning to engage with this hypothesis (Bell et al., 2016). In other work, how human decisions are influenced by the potentially time-varying economic value of stimuli and actions is a topic that has received considerable recent attention (Rangel and Hare, 2010; Rushworth et al., 2012), but an overview of this literature is beyond the scope of the current review.

2.3 Decision Confidence

For humans, perceptual decisions are often accompanied by a strong subjective sense of certainty or uncertainty, which we are often able to quantify when asked to report our "confidence." Decision confidence has been studied in humans for more than a century, but definitions have proven controversial, and no consensus has yet emerged about what function, if any, confidence may have across different species (Fetsch et al., 2014; Meyniel et al., 2015; Pouget et al., 2016; Shea et al., 2014). Nevertheless, rapid progress has been made in recent years thanks to convergent, cross-species research, and evidence for the neural implementation of decision confidence



Figure 5. Decision Confidence

(A and B) Two tasks used to measure decision confidence in experimental animals. In the opt-out task (A), the animal (here, a monkey) first makes a perceptual discrimination, and then, on a subset of trials, is offered a third option to "decline" the choice, which, when chosen, leads to a small but certain reward. Lower confidence should result in a higher proportion of "sure bet" choices. In the waiting task (B), an animal (here, a rat) first makes a perceptual discrimination, and then experiences a variable delay before reward on most correct trials. The rat either chooses to wait for the reward, or to initiate a new trial. The resulting waiting times scale with accuracy and are putative estimates of the animal's confidence. (A) was reprinted from Fetsch et al. (2014). (B) was reprinted from Lak et al. (2014).

(C) The left panel shows how (inverse) decision confidence should vary with signal strength (here, odor mixture) on correct and incorrect trials, under a signal detection model in which confidence is proportional to the absolute decision value. The same characteristic X-like pattern is observed in the firing rates of OFC neurons performing the task. Reprinted from Kepecs et al. (2008).

(D) The relationship between time, decision value, and the (log odds) probability of a correct response in the DDM, when signal reliability is unknown. As time passes, the variation in the decision value has minimal impact on the probability of a correct choice because if decisions are prolonged, it is likely that signal reliability is low. This motivates the foreclosure of deliberation with an urgency signal, and provides the basis for a model of decision confidence. Reprinted from Kiani and Shadlen (2009).

is beginning to emerge in rodents, monkeys, and humans (Kepecs and Mainen, 2012).

Humans are able to report confidence explicitly on an interval scale ("I'm 70% sure I'm right"). This behavior seems to require reflection upon one's own internal states, a class of computation that psychologists define as "metacognitive" (Yeung and Summerfield, 2012). Although it is not possible to elicit overt subjective reports of decision confidence in experimental animals, techniques allow the measurement of whether animals can use internal estimates of decision certainty to optimize their behavior. One creative approach adapts the two-choice discrimination paradigm by offering, after a subset of choices, a third option that "declines" the choice for a small but certain reward (Hampton, 2001) (opt-out task; Figure 5A). Subjects select this option more often when signal reliability is low, and performance is higher on those trials where the choice was voluntarily accepted than when the decline option was not presented (Kiani and Shadlen, 2009; Komura et al., 2013). Another elegant paradigm offers a graded estimate of decision confidence in animals, by imposing a variable delay to reward after most correct responses (and no feedback for errors) but allowing the animal to restart the next trial at will (waiting task; Figure 5B). The time that an animal was willing to wait for a reward is proportional to its accuracy (Lak et al., 2014). These new paradigms circumvent many of the traditional criticisms associated with the study of confidence in animals, and the findings suggest that rodents and monkeys, as well as humans, have access to internal estimates of decision certainty and use them to maximize outcomes (Kepecs and Mainen, 2012).

In order to understand how confidence is encoded in neural circuits, it is first necessary to come up with a formal definition of how it is computed from sensory or decision information (Pouget et al., 2016). A traditional view states that confidence reflects the divergence of a noisy decision variable from an indifference point or decision criterion (Galvin et al., 2003). This simple theory makes a key prediction: confidence should grow with signal

reliability on correct trials, but fall with signal reliability on error trials (Figure 5C, left panel; showing here the inverse pattern). To understand why, consider the two Gaussian distributions encoding sensory evidence for each category. The degree of separation between the distributions is greatest for stronger signals, and so errors will only occur when a decision variable falls close to the criterion (low confidence), whereas the converse is true for weaker signals. Indeed, in the waiting task, confidence estimated in this way successfully predicts how long a rat will postpone initiation of the next trial (Kepecs et al., 2008; Lak et al., 2014). Thus, one way of understanding confidence is that decision information is read out from the frame of reference of choice (e.g., probability of left versus right) and recoded in the frame of reference of accuracy (probability of et al., 2010).

Evidence from rodents has highlighted the orbitofrontal cortex (OFC) as a candidate site for this computation. In rats performing an olfactory mixture categorization task, about 25% of neurons show neural responses that exhibit an "X-like" pattern of positive and negative correlations with signal strength on correct and error trials, respectively (or the converse), consistent with predicted estimates of decision confidence on these trials (Kepecs et al., 2008) (Figure 5C, right panel). Moreover, ablation of rat OFC disrupts the otherwise monotonic relationship between waiting times and accuracy (Lak et al., 2014). Despite difficulties with identifying homologies among brain regions in rodents and humans, it is salutary that functional neuroimaging studies have also identified nearby polar regions of the frontal cortex as contributing to human confidence judgments about both perceptual and economic judgments (De Martino et al., 2013; Fleming et al., 2010, 2012). These findings have led to a converging perspective that the most anterior regions of the prefrontal cortex may participate in the computation of "metacognitive" signals that allow an animal to reflect on its performance and optimize behavior accordingly. However, other brain regions may also be involved. For example, neurons in the pulvinar nucleus of the macaque thalamus show a similar X-like pattern of responding during an opt-out task, and silencing of these responses increased the tendency to decline without impairing discrimination performance (Komura et al., 2013).

This perspective sees "confidence" and "certainty" as related but separate quantities that are potentially computed in distinct neural circuits (Pouget et al., 2016). An alternative is that confidence is equivalent to certainty, and is an intrinsic property of first-order decision signals themselves, encoding a graded belief about some state of the world. For example, uncertain sensory information (e.g., a low-coherence RDK) might elicit a broader neural population tuning curve, which is decoded into a weaker decision signal (Ma et al., 2006). This class of encoding scheme has the virtue of translating principles of probabilistic computation directly to neural circuits, and thus ensuring that more reliable signals wield greater influence over choices. It therefore provides an elegant mechanism for understanding how humans and monkeys optimally combine noisy sensory information across modalities during multisensory perception (Ernst and Banks, 2002; Ma and Jazayeri, 2014). In support of this view, in the opt-out task, single-cell activity in LIP both during and after stimulus viewing predicts a monkey's later decision to decline

Neuron Review

when available, with mean firing rates on opt-out trials falling intermediate between (but not simply reflecting a mixture of) those where a saccade is made to the target or distracter. In other words, decision information and choice certainty are not encoded separately, but rather in a common neural population (Kiani and Shadlen, 2009).

This latter perspective raises the question of how confidence is used to guide behavior, such as the decision to select an "opt-out" option. One long-standing difficulty with explaining confidence judgments in terms of the decision variable proposed by the DDM is that if decision bounds are fixed at a single value, then terminal certainty (and reported confidence) should be identical on every trial. One solution to this problem draws upon a different sequential sampling model, in which cumulative tallies of evidence for each option "race" toward a single threshold. This allows confidence to be computed as the "balance of evidence" between the winning trace and its nearest competitor at the time of the decision (Vickers, 1979). A different solution to this problem invokes a notion discussed above, namely that where sensory reliability is unknown, the time elapsed in a trial is a good proxy for whether a decision will be accurate or not (Figure 5D). Indeed, a model that uses both signal strength and time elapsed to compute the probability of a correct response can account for monkey behavior and concomitant neural data during the opt-out task (Kiani and Shadlen, 2009). Interestingly, unlike the models that overlook elapsed time as a predictor of confidence, this account correctly predicts that when choice and confidence are signaled with a single, ballistic movement (precluding post-decision information or "changes of mind" from polluting certainty estimates), then confidence will increase with signal strength on both correct and error trials. This finding has been supported by recent data (Kiani et al., 2014a). Thus, under this perspective confidence is a product of both signal strength and elapsed time, and first- and second-order decision signals share a common neural substrate in LIP, obviating the need for a separate evaluative or "metacognitive" system encoding decision confidence. However, these issues remain controversial, and it will fall to future research to unpick the complex empirical and theoretical arguments surrounding decision confidence (Insabato et al., 2016).

Summary and Perspective

In this review, we have discussed central questions in the study of perceptual decision making concerning both the computation of decision variables and their expression in neural circuits. The first section describes insights derived from new approaches using monkeys, rodents, and humans that have begun to reshape the field. In section 1.1, we discussed new challenges to the canonical view that neurons in area LIP of the macaque monkey implement a gradual, build-to-threshold decision process. In particular, we note that a new emphasis is being placed on the heterogenous coding properties (mixed selectivity) of neurons in this area, and questions have been raised about the nature of the dynamics of integration. We also discuss the key recent finding that inactivating LIP seems to have little or no impact on perceptual decisions, calling into question its causal role in this function. In section 1.2, we survey exciting new techniques that have been developed to study perceptual decisions

in association with a rodent model. New methods offer the possibility of targeting specific cell types for both recording and perturbation. Although this field is new, it has already pointed to hitherto overlooked dissociations among brain regions previously implicated in perceptual choice, including the parietal and prefrontal cortices. Early work in this area has hinted that characterizing the circuit mechanisms underlying perceptual decisions, for example, by classifying neurons according to their cell type, selectivity, or projection targets, may be a key to unravelling the computational mechanisms of perceptual decision making and pinpointing their neural implementation. In section 1.3, we focus on new work that has studied perceptual decision making in humans, using modeling in conjunction with whole-brain functional neuroimaging techniques. This work has identified new candidate brain signals that could reflect a build-to-threshold decision process in humans, as well as extending the sequential sampling framework to account for suboptimal biases that occur in human decisions about both sensory stimuli and economic prospects.

The second section expands the discussion to topics that we believe offer great promise for future research linking model systems. In section 2.1, we discuss the timing of perceptual decisions, with a focus on both monkeys and humans. New single-cell recordings have alerted researchers to the idea that the timing of many decisions may be driven by a dynamically growing, evidence-independent "urgency" signal not predicted by standard computational models that draws decisions to a close when information is weak or ambiguous. However, behavioral and neural evidence for this signal remain preliminary. In section 2.2, we review neural studies that have revisited the long-standing question of how perceptual decisions are biased by contextual signals encoding the probability or value of responses. Once again, we find some evidence that in addition to an early additive component, a dynamically growing signal may bias decisions toward a preferred bound, especially under free-response conditions. Finally, in section 2.3, we discuss intriguing new debates surrounding the computation of decision confidence, and its neural implementation. We compare rival views suggesting that confidence may be a "second-order" or "metacognitive" signal, and those proposing that decisions and confidence rely on shared computations and neural circuits. We discuss the view that elapsed decision time is a key quantity determinant of decision confidence.

We would like to close by bringing together a number of common themes that have emerged from the diverse questions we have considered, which suggest some promising avenues for research over the coming years.

The elegant, statistically optimal framework provided by the sequential sampling framework has been of great benefit to the field, providing a strong computational basis for interpreting neural data. However, we argue that the next step will be to move to understand perceptual decisions at a lower, more biologically plausible level of description that refers to both cell types and neural dynamics at the level of microcircuits. Both the computational framework and the recording tools are becoming available to tackle this challenge, and we urge researchers to build on important first steps to bring them together. We expect that studies focused on unpicking the contribution of distinct

neuronal classes within the microcircuit will shed light on the perplexing heterogeneity of neural responding observed during perceptual decisions in both rodents and monkeys, and help further characterize the dynamics that underlie sequential integration of decision information.

Relatedly, we argue that current understanding of how distinct regions coordinate perceptual decision making across the brain is very limited. This is in part because until now, most research groups have focused exclusively on a limited number of recording sites, such as LIP and FEF. The advent of highthroughput recording opens new doors to multi-site recordings that can explore the relative latencies of, and interactions between, multiple brain regions (Siegel et al., 2015). Moreover, future research promises to reveal key contributions to perceptual decision making from currently under-explored brain regions, such as the striatum, as well as clarifying the relative contributions of structures such as the parietal and orbitofrontal cortices. Characterizing the nature of the decision signals at diverse cortical and subcortical sites also offers the opportunity to link animal work more closely with human functional neuroimaging studies, which have emphasized the contribution of regions such as the anterior cingulate cortex and insular cortex that are relatively unexplored in studies of rodent and monkey perceptual decisions. An important step toward realizing this goal will be more work that explicitly compares neural responses across species, in an attempt to identify functional as well as structural homologies (Narayanan et al., 2013).

In parallel, we hope that new insights from recording and interference studies across species will also help resolve debates over computational mechanisms of decision making. The DDM has been a tremendously useful and reliable workhorse for understanding extant neural data, but a full mechanistic description of perceptual decisions will need to elaborate this modeling framework to incorporate additional terms and parameters. One example that has been highlighted in the current review is the notion that decisions are not merely driven by accumulation of noisy sensory evidence, but by time-varying bias signals that help curtail deliberation in the face of ambiguous information. This "urgency" signal seems not only to help control the tradeoff between speed and accuracy, but may also contribute to the biasing of decisions by probability, and even play a role in the computation of decision confidence. We argue that characterizing this signal at the neural level is an important future goal for neurophysiologists. Other examples are emerging from human work, where subjects can more easily engage in more sophisticated tasks and are less overtrained, placing suboptimal decision policies more clearly in view and offering the opportunity to understand their computational substrates. New neural data will undoubtedly help further constrain and guide our modeling framework, and other computational details concerning how information is transformed and integrated will no doubt emerge over coming years. As ever, the pace of progress will be accelerated if researchers keep an open mind about the classes of model that may best describe perceptual decisions.

In summary, we think that the field is at an important juncture. New techniques have shown that advanced optogenetic and imaging methods are feasible in the rodent, and new paradigms have opened doors to understanding the cognitive



underpinnings of perceptual decisions in humans. The resultant data have already inspired new theoretical insights concerning the neural computations that support decision formation. The next steps are to connect these insights across species, furnishing general principles for perceptual decision making in rodents, monkeys, and humans.

ACKNOWLEDGMENTS

T.D.H. and C.S. contributed equally to this manuscript.

REFERENCES

Balan, P.F., and Gottlieb, J. (2009). Functional significance of nonspatial information in monkey lateral intraparietal area. J. Neurosci. *29*, 8166–8176.

Barak, O., Sussillo, D., Romo, R., Tsodyks, M., and Abbott, L.F. (2013). From fixed points to chaos: three models of delayed discrimination. Prog. Neurobiol. *103*, 214–222.

Basten, U., Biele, G., Heekeren, H.R., and Fiebach, C.J. (2010). How the brain integrates costs and benefits during decision making. Proc. Natl. Acad. Sci. USA *107*, 21767–21772.

Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., and Friston, K.J. (2012). Canonical microcircuits for predictive coding. Neuron *76*, 695–711.

Bell, A.H., Summerfield, C., Morin, E.L., Malecek, N.J., and Ungerleider, L.G. (2016). Encoding of stimulus probability in macaque inferior temporal cortex. Curr. Biol. *26*, 2280–2290.

Bennur, S., and Gold, J.I. (2011). Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. J. Neurosci. *31*, 913–921.

Bernacchia, A., Seo, H., Lee, D., and Wang, X.J. (2011). A reservoir of time constants for memory traces in cortical neurons. Nat. Neurosci. 14, 366–372.

Bisley, J.W., Krishna, B.S., and Goldberg, M.E. (2004). A rapid and precise onresponse in posterior parietal cortex. J. Neurosci. 24, 1833–1838.

Boehm, U., Hawkins, G.E., Brown, S., van Rijn, H., and Wagenmakers, E.J. (2016). Of monkeys and men: Impatience in perceptual decision-making. Psychon. Bull. Rev. 23, 738–749.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. Psychol. Rev. *113*, 700–765.

Bogacz, R., Wagenmakers, E.J., Forstmann, B.U., and Nieuwenhuis, S. (2010). The neural basis of the speed-accuracy tradeoff. Trends Neurosci. 33, 10–16.

Bollimunta, A., and Ditterich, J. (2012). Local computation of decision-relevant net sensory evidence in parietal cortex. Cereb. Cortex 22, 903–917.

Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1993). Responses of neurons in macaque MT to stochastic motion signals. Vis. Neurosci. 10, 1157–1169.

Broadbent, D.E. (1958). Perception and Communication (Pergamon Press).

Brody, C.D., and Hanks, T.D. (2016). Neural underpinnings of the evidence accumulator. Curr. Opin. Neurobiol. 37, 149–157.

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95–98.

Busemeyer, J.R., and Townsend, J.T. (1993). Decision field theory: a dynamiccognitive approach to decision making in an uncertain environment. Psychol. Rev. *100*, 432–459.

Carandini, M., and Churchland, A.K. (2013). Probing perceptual decisions in rodents. Nat. Neurosci. *16*, 824–831.

Cassey, P., Heathcote, A., and Brown, S.D. (2014). Brain and behavior in decision-making. PLoS Comput. Biol. *10*, e1003700. Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., de Gardelle, V., Herce Castañón, S., and Summerfield, C. (2014). Adaptive gain control during human perceptual choice. Neuron *81*, 1429–1441.

Chen, M.Y., Jimura, K., White, C.N., Maddox, W.T., and Poldrack, R.A. (2015). Multiple brain networks contribute to the acquisition of bias in perceptual decision-making. Front. Neurosci. 9, 63.

Churchland, A.K., and Kiani, R. (2016). Three challenges for connecting model to mechanism in decision-making. Curr Opin Behav Sci *11*, 74–80.

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

Churchland, A.K., Kiani, R., Chaudhuri, R., Wang, X.J., Pouget, A., and Shadlen, M.N. (2011). Variance as a signature of neural computations during decision making. Neuron 69, 818–831.

Cisek, P., and Kalaska, J.F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. Neuron 45, 801–814.

Cunningham, J.P., and Yu, B.M. (2014). Dimensionality reduction for largescale neural recordings. Nat. Neurosci. 17, 1500–1509.

de Lange, F.P., Jensen, O., and Dehaene, S. (2010). Accumulation of evidence during sequential decision making: the importance of top-down factors. J. Neurosci. *30*, 731–738.

de Lange, F.P., Rahnev, D.A., Donner, T.H., and Lau, H. (2013). Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. J. Neurosci. *33*, 1400–1410.

De Martino, B., Fleming, S.M., Garrett, N., and Dolan, R.J. (2013). Confidence in value-based choice. Nat. Neurosci. *16*, 105–110.

Deneve, S. (2012). Making decisions with unknown sensory reliability. Front. Neurosci. 6, 75.

Ding, L. (2015). Distinct dynamics of ramping activity in the frontal cortex and caudate nucleus in monkeys. J. Neurophysiol. *114*, 1850–1861.

Ding, L., and Gold, J.I. (2010). Caudate encodes multiple computations for perceptual decisions. J. Neurosci. *30*, 15747–15759.

Ding, L., and Gold, J.I. (2012). Separate, causal roles of the caudate in saccadic choice and execution in a perceptual decision task. Neuron 75, 865–874.

Ditterich, J., Mazurek, M.E., and Shadlen, M.N. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. Nat. Neurosci. *6*, 891–898.

Donner, T.H., Siegel, M., Fries, P., and Engel, A.K. (2009). Buildup of choicepredictive activity in human motor cortex during perceptual decision making. Curr. Biol. *19*, 1581–1585.

Drugowitsch, J., Moreno-Bote, R., Churchland, A.K., Shadlen, M.N., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. J. Neurosci. *32*, 3612–3628.

Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Requirement of prefrontal and midbrain regions for rapid executive control of behavior in the rat. Neuron 86, 1491–1503.

Erlich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. eLife *4*, http://dx.doi.org/10.7554/eLife.05457.

Ernst, M.O., and Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. Nature *415*, 429–433.

Fetsch, C.R., Kiani, R., and Shadlen, M.N. (2014). Predicting the accuracy of a decision: a neural mechanism of confidence. Cold Spring Harb. Symp. Quant. Biol. *79*, 185–197.

Fitzgerald, J.K., Freedman, D.J., Fanini, A., Bennur, S., Gold, J.I., and Assad, J.A. (2013). Biased associative representations in parietal cortex. Neuron 77, 180–191.

Fleming, S.M., Weil, R.S., Nagy, Z., Dolan, R.J., and Rees, G. (2010). Relating introspective accuracy to individual differences in brain structure. Science 329, 1541–1543.

Fleming, S.M., Huijgen, J., and Dolan, R.J. (2012). Prefrontal contributions to metacognition in perceptual decision making. J. Neurosci. *32*, 6117–6125.

Frazier, P., and Yu, A.J. (2008). Sequential hypothesis testing under stochastic deadlines. Adv. Neural Inf. Process. Syst. 20, 465–472.

Fusi, S., Miller, E.K., and Rigotti, M. (2016). Why neurons mix: high dimensionality for higher cognition. Curr. Opin. Neurobiol. 37, 66–74.

Galvin, S.J., Podd, J.V., Drga, V., and Whitmore, J. (2003). Type 2 tasks in the theory of signal detectability: discrimination between correct and incorrect decisions. Psychon. Bull. Rev. *10*, 843–876.

Ganguli, S., Bisley, J.W., Roitman, J.D., Shadlen, M.N., Goldberg, M.E., and Miller, K.D. (2008). One-dimensional dynamics of attention and decision making in LIP. Neuron 58, 15–25.

Gold, J.I., and Shadlen, M.N. (2000). Representation of a perceptual decision in developing oculomotor commands. Nature 404, 390–394.

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

Gold, J.I., Law, C.T., Connolly, P., and Bennur, S. (2008). The relative influences of priors and sensory evidence on an oculomotor decision variable during perceptual learning. J. Neurophysiol. *100*, 2653–2668.

Gorea, A., Belkoura, S., and Solomon, J.A. (2014). Summary statistics for size over space and time. J. Vis. 14, 22.

Grosenick, L., Marshel, J.H., and Deisseroth, K. (2015). Closed-loop and activity-guided optogenetic control. Neuron *86*, 106–139.

Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., and Svoboda, K. (2014). Flow of cortical activity underlying a tactile decision in mice. Neuron *81*, 179–194.

Hampton, R.R. (2001). Rhesus monkeys know when they remember. Proc. Natl. Acad. Sci. USA 98, 5359–5362.

Hanes, D.P., and Schall, J.D. (1996). Neural control of voluntary movement initiation. Science 274, 427–430.

Hanks, T.D., Ditterich, J., and Shadlen, M.N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. Nat. Neurosci. 9, 682–689.

Hanks, T.D., Mazurek, M.E., Kiani, R., Hopp, E., and Shadlen, M.N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. J. Neurosci. *31*, 6339–6352.

Hanks, T., Kiani, R., and Shadlen, M.N. (2014). A neural mechanism of speedaccuracy tradeoff in macaque area LIP. eLife 3, http://dx.doi.org/10.7554/ eLife.02260.

Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220–223.

Hawkins, G.E., Forstmann, B.U., Wagenmakers, E.J., Ratcliff, R., and Brown, S.D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. J. Neurosci. 35, 2476–2484.

Heekeren, H.R., Marrett, S., Bandettini, P.A., and Ungerleider, L.G. (2004). A general mechanism for perceptual decision-making in the human brain. Nature *431*, 859–862.

Heekeren, H.R., Marrett, S., and Ungerleider, L.G. (2008). The neural systems that mediate human perceptual decision making. Nat. Rev. Neurosci. *9*, 467–479.

Heitz, R.P., and Schall, J.D. (2012). Neural mechanisms of speed-accuracy tradeoff. Neuron 76, 616–628.

Heitz, R.P., and Schall, J.D. (2013). Neural chronometry and coherency across speed-accuracy demands reveal lack of homomorphism between computational and neural mechanisms of evidence accumulation. Philos. Trans. R. Soc. Lond. B Biol. Sci. 368, 20130071.

Hernández, A., Zainos, A., and Romo, R. (2000). Neuronal correlates of sensory discrimination in the somatosensory cortex. Proc. Natl. Acad. Sci. USA 97, 6191–6196.

Ho, T.C., Brown, S., and Serences, J.T. (2009). Domain general mechanisms of perceptual decision making in human cortex. J. Neurosci. *29*, 8675–8687.

Hubert-Wallander, B., and Boynton, G.M. (2015). Not all summary statistics are made equal: Evidence from extracting summaries across time. J. Vis. *15*, 5.

Huk, A.C., and Meister, M.L. (2012). Neural correlates and neural computations in posterior parietal cortex during perceptual decision-making. Front. Integr. Nuerosci. *6*, 86.

Huk, A.C., and Shadlen, M.N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J. Neurosci. *25*, 10420–10436.

Hunt, L.T. (2014). What are the neural origins of choice variability? Trends Cogn. Sci. 18, 222–224.

Insabato, A., Pannunzi, M., Rolls, E.T., and Deco, G. (2010). Confidencerelated decision making. J. Neurophysiol. *104*, 539–547.

Insabato, A., Pannunzi, M., and Deco, G. (2016). Neural correlates of metacognition: A critical perspective on current tasks. Neurosci. Biobehav. Rev. 71, 167–175.

Janssen, P., and Shadlen, M.N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. Nat. Neurosci. *8*, 234–241.

Kahneman, D. (2011). Thinking, Fast and Slow (Farrar, Straus and Giroux).

Katz, L.N., Yates, J.L., Pillow, J.W., and Huk, A.C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. Nature 535, 285–288.

Kelly, S.P., and O'Connell, R.G. (2013). Internal and external influences on the rate of sensory evidence accumulation in the human brain. J. Neurosci. *33*, 19434–19441.

Kelly, S.P., and O'Connell, R.G. (2015). The neural processes underlying perceptual decision making in humans: recent progress and future directions. J. Physiol. Paris *109*, 27–37.

Kepecs, A., and Mainen, Z.F. (2012). A computational framework for the study of confidence in humans and animals. Philos. Trans. R. Soc. Lond. B Biol. Sci. 367, 1322–1337.

Kepecs, A., Uchida, N., Zariwala, H.A., and Mainen, Z.F. (2008). Neural correlates, computation and behavioural impact of decision confidence. Nature *455*, 227–231.

Kiani, R., and Shadlen, M.N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. Science 324, 759–764.

Kiani, R., Corthell, L., and Shadlen, M.N. (2014a). Choice certainty is informed by both evidence and decision time. Neuron *84*, 1329–1342.

Kiani, R., Cueva, C.J., Reppas, J.B., and Newsome, W.T. (2014b). Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. Curr. Biol. *24*, 1542–1547.

Kira, S., Yang, T., and Shadlen, M.N. (2015). A neural implementation of Wald's sequential probability ratio test. Neuron *85*, 861–873.

Knowlton, B.J., and Squire, L.R. (1993). The learning of categories: parallel brain systems for item memory and category knowledge. Science *262*, 1747–1749.

Kobak, D., Brendel, W., Constantinidis, C., Feierstein, C.E., Kepecs, A., Mainen, Z.F., Qi, X.L., Romo, R., Uchida, N., and Machens, C.K. (2016). Demixed principal component analysis of neural population data. eLife 5, e10989.

Kok, P., Failing, M.F., and de Lange, F.P. (2014). Prior expectations evoke stimulus templates in the primary visual cortex. J. Cogn. Neurosci. *26*, 1546–1554.

Komura, Y., Nikkuni, A., Hirashima, N., Uetake, T., and Miyamoto, A. (2013). Responses of pulvinar neurons reflect a subject's confidence in visual categorization. Nat. Neurosci. *16*, 749–755.



Körding, K. (2007). Decision theory: what "should" the nervous system do? Science 318, 606–610.

Krajbich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. Nat. Neurosci. *13*, 1292– 1298.

Lak, A., Costa, G.M., Romberg, E., Koulakov, A.A., Mainen, Z.F., and Kepecs, A. (2014). Orbitofrontal cortex is required for optimal waiting based on decision confidence. Neuron *84*, 190–201.

Lakatos, P., Karmos, G., Mehta, A.D., Ulbert, I., and Schroeder, C.E. (2008). Entrainment of neuronal oscillations as a mechanism of attentional selection. Science 320, 110–113.

Latimer, K.W., Yates, J.L., Meister, M.L., Huk, A.C., and Pillow, J.W. (2015). NEURONAL MODELING. Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. Science *349*, 184–187.

Leathers, M.L., and Olson, C.R. (2012). In monkeys making value-based decisions, LIP neurons encode cue salience and not action value. Science *338*, 132–135.

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

Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. Nature 532, 459–464.

Licata, A.M., Kaufman, M.T., Raposo, D., Ryan, M.B., Sheppard, J.P., and Churchland, A.K. (2016). Posterior parietal cortex guides visual decisions in rats. bioRxiv. http://dx.doi.org/10.1101/066639.

Liu, T., and Pleskac, T.J. (2011). Neural correlates of evidence accumulation in a perceptual decision task. J. Neurophysiol. *106*, 2383–2398.

Luo, L., Callaway, E.M., and Svoboda, K. (2008). Genetic dissection of neural circuits. Neuron 57, 634–660.

Ma, W.J., and Jazayeri, M. (2014). Neural coding of uncertainty and probability. Annu. Rev. Neurosci. 37, 205–220.

Ma, W.J., Beck, J.M., Latham, P.E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. Nat. Neurosci. 9, 1432–1438.

Machens, C.K., Romo, R., and Brody, C.D. (2010). Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex. J. Neurosci. 30, 350–360.

Meister, M.L., Hennig, J.A., and Huk, A.C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. J. Neurosci. 33, 2254–2267.

Meyniel, F., Sigman, M., and Mainen, Z.F. (2015). Confidence as Bayesian probability: from neural origins to behavior. Neuron *88*, 78–92.

Moran, R. (2015). Optimal decision making in heterogeneous and biased environments. Psychon. Bull. Rev. 22, 38–53.

Mulder, M.J., Wagenmakers, E.J., Ratcliff, R., Boekel, W., and Forstmann, B.U. (2012). Bias in the brain: a diffusion model analysis of prior probability and potential payoff. J. Neurosci. *32*, 2335–2343.

Mulder, M.J., van Maanen, L., and Forstmann, B.U. (2014). Perceptual decision neurosciences - a model-based review. Neuroscience 277, 872–884.

Murakami, M., and Mainen, Z.F. (2015). Preparing and selecting actions with neural populations: toward cortical circuit mechanisms. Curr. Opin. Neurobiol. 33, 40–46.

Narayanan, N.S., Cavanagh, J.F., Frank, M.J., and Laubach, M. (2013). Common medial frontal mechanisms of adaptive control in humans and rodents. Nat. Neurosci. *16*, 1888–1895.

Newsome, W.T., and Paré, E.B. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). J. Neurosci. 8, 2201–2211.

Nienborg, H., and Cumming, B.G. (2009). Decision-related activity in sensory neurons reflects more than a neuron's causal effect. Nature *459*, 89–92.

O'Connell, R.G., Dockree, P.M., and Kelly, S.P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. Nat. Neurosci. *15*, 1729–1735.

Park, I.M., Meister, M.L., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. Nat. Neurosci. *17*, 1395–1403.

Philiastides, M.G., Auksztulewicz, R., Heekeren, H.R., and Blankenburg, F. (2011). Causal role of dorsolateral prefrontal cortex in human perceptual decision making. Curr. Biol. *21*, 980–983.

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

Poddar, R., Kawai, R., and Ölveczky, B.P. (2013). A fully automated high-throughput training system for rodents. PLoS ONE *8*, e83171.

Pouget, A., Drugowitsch, J., and Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. Nat. Neurosci. 19, 366–374.

Purcell, B.A., and Kiani, R. (2016). Neural mechanisms of post-error adjustments of decision policy in parietal cortex. Neuron 89, 658–671.

Rangel, A., and Hare, T. (2010). Neural computations associated with goaldirected choice. Curr. Opin. Neurobiol. 20, 262–270.

Rao, V., DeAngelis, G.C., and Snyder, L.H. (2012). Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas. J. Neurosci. *32*, 10063–10074.

Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. Nat. Neurosci. *17*, 1784–1792.

Ratcliff, R., and Rouder, J.N. (1998). Modeling response times for two-choice decisions. Psychol. Sci. 9, 347–356.

Ratcliff, R., Smith, P.L., Brown, S.D., and McKoon, G. (2016). Diffusion decision model: current issues and history. Trends Cogn. Sci. 20, 260–281.

Renart, A., and Machens, C.K. (2014). Variability in neural activity and behavior. Curr. Opin. Neurobiol. 25, 211–220.

Rigotti, M., Barak, O., Warden, M.R., Wang, X.J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature *497*, 585–590.

Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J. Neurosci. *22*, 9475–9489.

Romo, R., Hernández, A., and Zainos, A. (2004). Neuronal correlates of a perceptual decision in ventral premotor cortex. Neuron *41*, 165–173.

Rorie, A.E., Gao, J., McClelland, J.L., and Newsome, W.T. (2010). Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. PLoS ONE *5*, e9308.

Rushworth, M.F., Kolling, N., Sallet, J., and Mars, R.B. (2012). Valuation and decision-making in frontal cortex: one or many serial or parallel systems? Curr. Opin. Neurobiol. *22*, 946–955.

Schall, J.D. (2003). Neural correlates of decision processes: neural and mental chronometry. Curr. Opin. Neurobiol. *13*, 182–186.

Scott, B.B., Constantinople, C.M., Erlich, J.C., Tank, D.W., and Brody, C.D. (2015). Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. eLife *4*, e11308.

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

Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C., and Frith, C.D. (2014). Supra-personal cognitive control and metacognition. Trends Cogn. Sci. 18, 186–193.

Siegel, M., Buschman, T.J., and Miller, E.K. (2015). Cortical information flow during flexible sensorimotor decisions. Science 348, 1352–1355.

Simen, P., Contreras, D., Buck, C., Hu, P., Holmes, P., and Cohen, J.D. (2009). Reward rate optimization in two-alternative decision making: empirical tests of

theoretical predictions. J. Exp. Psychol. Hum. Percept. Perform. 35, 1865–1897.

Spitzer, B., Blankenburg, F., and Summerfield, C. (2016). Rhythmic gain control during supramodal integration of approximate number. Neuroimage *129*, 470–479.

Summerfield, C., and de Lange, F.P. (2014). Expectation in perceptual decision making: neural and computational mechanisms. Nat. Rev. Neurosci. *15*, 745–756.

Summerfield, C., and Koechlin, E. (2010). Economic value biases uncertain perceptual choices in the parietal and prefrontal cortices. Front. Hum. Neurosci. *4*, 208.

Summerfield, C., and Tsetsos, K. (2012). Building bridges between perceptual and economic decision-making: neural and computational mechanisms. Front. Neurosci. *6*, 70.

Summerfield, C., and Tsetsos, K. (2015). Do humans make good decisions? Trends Cogn. Sci. 19, 27–34.

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

Thura, D., and Cisek, P. (2016). Modulation of premotor and primary motor cortical activity during volitional adjustments of speed-accuracy trade-offs. J. Neurosci. *36*, 938–956.

Thura, D., Beauregard-Racine, J., Fradet, C.W., and Cisek, P. (2012). Decision making by urgency gating: theory and experimental support. J. Neurophysiol. *108*, 2912–2930.

Tsetsos, K., Chater, N., and Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. Proc. Natl. Acad. Sci. USA *109*, 9659–9664.

Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., and Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. Proc. Natl. Acad. Sci. USA *113*, 3102–3107.

Twomey, D.M., Murphy, P.R., Kelly, S.P., and O'Connell, R.G. (2015). The classic P300 encodes a build-to-threshold decision variable. Eur. J. Neurosci. *42*, 1636–1643.

Uchida, N., Kepecs, A., and Mainen, Z.F. (2006). Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. Nat. Rev. Neurosci. 7, 485–491.

van Veen, V., Krug, M.K., and Carter, C.S. (2008). The neural and computational basis of controlled speed-accuracy tradeoff during task performance. J. Cogn. Neurosci. *20*, 1952–1965.

Vickers, D. (1979). Decision Processes in Visual Perception (Academic Press).

Wald, A., and Wolfowitz, J. (1949). Bayes solutions of sequential decision problems. Proc. Natl. Acad. Sci. USA *35*, 99–102.

Wardak, C., Olivier, E., and Duhamel, J.R. (2002). Saccadic target selection deficits after lateral intraparietal area inactivation in monkeys. J. Neurosci. *22*, 9877–9884.

Wheeler, M.E., Woo, S.G., Ansel, T., Tremel, J.J., Collier, A.L., Velanova, K., Ploran, E.J., and Yang, T. (2015). The strength of gradually accruing probabilistic evidence modulates brain activity during a categorical decision. J. Cogn. Neurosci. *27*, 705–719.

Wilke, M., Kagan, I., and Andersen, R.A. (2012). Functional imaging reveals rapid reorganization of cortical activity after parietal inactivation in monkeys. Proc. Natl. Acad. Sci. USA *109*, 8274–8279.

Winkel, J., Keuken, M.C., van Maanen, L., Wagenmakers, E.J., and Forstmann, B.U. (2014). Early evidence affects later decisions: why evidence accumulation is required to explain response time data. Psychon. Bull. Rev. 21, 777–784.

Wyart, V., and Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. Current Opinion in Behavioral Sciences *11*, 109–115.

Wyart, V., de Gardelle, V., Scholl, J., and Summerfield, C. (2012). Rhythmic fluctuations in evidence accumulation during decision making in the human brain. Neuron 76, 847–858.

Wyart, V., Myers, N.E., and Summerfield, C. (2015). Neural mechanisms of human perceptual choice under focused and divided attention. J. Neurosci. *35*, 3485–3498.

Yang, T., and Shadlen, M.N. (2007). Probabilistic reasoning by neurons. Nature 447, 1075–1080.

Yeung, N., and Summerfield, C. (2012). Metacognition in human decisionmaking: confidence and error monitoring. Philos. Trans. R. Soc. Lond. B Biol. Sci. 367, 1310–1321.

Znamenskiy, P., and Zador, A.M. (2013). Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. Nature 497, 482–485.