



## Review

**Cite this article:** Cisek P, Pastor-Bernier A. 2014 On the challenges and mechanisms of embodied decisions. *Phil. Trans. R. Soc. B* **369**: 20130479. <http://dx.doi.org/10.1098/rstb.2013.0479>

One contribution of 18 to a Theme Issue 'The principles of goal-directed decision-making: from neural mechanisms to computation and robotics'.

### Subject Areas:

neuroscience, cognition

### Keywords:

decision-making, action selection, embodied behaviour, ecological psychology, affordances

### Author for correspondence:

Paul Cisek

e-mail: [paul.cisek@umontreal.ca](mailto:paul.cisek@umontreal.ca)

# On the challenges and mechanisms of embodied decisions

Paul Cisek<sup>1</sup> and Alexandre Pastor-Bernier<sup>2</sup>

<sup>1</sup>Groupe de Recherche sur le Système Nerveux Central (GRSNC), Département de Neurosciences, Université de Montréal, C.P. 6128 Succursale Centre-ville, Montréal, Québec, Canada H3C 3J7

<sup>2</sup>Department of Physiology, Development and Neuroscience (PDN), University of Cambridge, Cambridge, UK

Neurophysiological studies of decision-making have focused primarily on elucidating the mechanisms of classic economic decisions, for which the relevant variables are the values of expected outcomes and action is simply the means of reporting the selected choice. By contrast, here we focus on the particular challenges of embodied decision-making faced by animals interacting with their environment in real time. In such scenarios, the choices themselves as well as their relative costs and benefits are defined by the momentary geometry of the immediate environment and change continuously during ongoing activity. To deal with the demands of embodied activity, animals require an architecture in which the sensorimotor specification of potential actions, their valuation, selection and even execution can all take place in parallel. Here, we review behavioural and neurophysiological data supporting a proposed brain architecture for dealing with such scenarios, which we argue set the evolutionary foundation for the organization of the mammalian brain.

## 1. Introduction

Studies of human behaviour have historically been distributed across very diverse disciplines, whose practitioners seldom spoke with one another or exchanged ideas and insights. The earliest psychologists such as Wundt studied mental phenomena at a purely abstract level, undaunted by the sparseness of biological knowledge of their time. Likewise, the cognitive psychology of the late twentieth century was focused largely on mental function, often explicitly not concerned with brain structure (e.g. [1]). At the same time, studies of rational choice were largely the province of economic theory that focused primarily on understanding what the rational behaviour of *Economic Man* should be [2]. It could be said that this disciplinary division of labour reflected the three levels of analysis famously described by Marr [3]: economics studied the computational level of the problem, psychology studied the brain's algorithms and functional organization, and neurology was concerned with the details of implementation.

All that has changed in recent decades, and in almost all aspects of the behavioural sciences, there is now a trend to bring disciplines together and to build theories that at once address the phenomena in question at all three levels of analysis. While this is certainly a step in the right direction, marrying long estranged fields is a precarious business. Some caution must be taken in choosing which conceptual components are most valuable for guiding the fields towards a unified direction that is a fruitful one, rather than merely entrenching them all towards the same limiting assumptions.

Here, we take a cautionary look at the recent convergence of studies in behavioural economics and the neuroscience of decision-making, a field often called *neuroeconomics* [4,5]. As a new field, it is in great flux and difficult to characterize with a consistent set of agreed-upon axioms. Nevertheless, certain themes are repeatedly emphasized and at least lead to convergence on the questions that are being asked.

At the computational level, one persistent question inspired by economic theories is the degree to which human (and animal) behaviour is rational. For

example, it is obviously irrational (a violation of ‘transitivity’) to prefer A to B and B to C, but then choose C over A. VonNeumann & Morgenstern [6] formalized the concept of expected utility to describe how humans should make choices under conditions of risk, allowing for nonlinear relationships between objective quantities (amount of money) and their subjective utility to a decision-making agent. However, Simon [7] famously showed that humans do not behave in accordance with prescriptive economic theory, that they exhibit ‘bounded rationality’. This was echoed even more strongly by Kahneman & Tversky’s [8] seminal work on the heuristics and biases that drive human choices away from what is expected of purely rational thinking. While to some, the use of heuristics reveals flaws in human reasoning, others argue that heuristics actually do better than prescriptive statistical methods in real-world scenarios in which sample sizes are too small to ever usefully estimate the necessary probability distributions [9].

At the algorithmic level, a critical question concerns the representations and mechanisms that determine decisions. A classic model suggests that when presented with a choice, people compute the subjective utility of each option using a ‘common currency’ that integrates payoffs, costs, risks and any relevant variable [4,10–13]. The option with the highest value is selected as the goal, which is then ultimately transformed into an action plan. We may refer to the processes of calculating values (which may be prolonged in complex situations) as ‘deliberation’, and the process of goal selection as ‘commitment’. According to classic psychological theories, both of these processes take place within a central executive system [14,15] that presumably resides in the frontal lobes and is separate from sensorimotor control [16–19].

Recent studies seek to connect these principles to the level of neural implementation. For example, neural activity in the orbitofrontal cortex (OFC) and ventromedial prefrontal cortex behaves very much like an abstract representation of economic value [12,20–23]. In particular, activity in the OFC correlates with the value of an option independent of other options [24] and adjusts its gain to reflect the full range of values presented in a given block of trials [25]. The transformation of the chosen goal to a plan of action may take place in lateral prefrontal cortex [26], which projects to the premotor and motor regions that control movement [27]. Even departures from prescriptive economic theories can be studied using neuroscientific methods, which may explain differences of behaviour between individuals. For example, Tobler *et al.* [28] showed that the idiosyncratic distortions of probability estimates of individual subjects are correlated with activity patterns in the dorsolateral prefrontal cortex. Likewise, individual subjective policies of temporal discounting can be correlated with individual patterns of activity in ventral striatum, medial prefrontal cortex and posterior cingulate cortex [29]. In summary, questions posed by economic theory can be addressed with neuroscience, and some have argued that economic theory itself should be increasingly based on neuroscientific data [4].

## 2. Beyond economics

Despite the undeniable attraction of economic formalisms for interpreting neural data, all formalisms have a limited domain of applicability whose edges are important to

acknowledge. Economic theory is no exception, and here we focus on three of its limitations.

First, it has long been recognized that the most quantifiable formalisms of economic theory are designed to deal with only the most circumscribed situations, raising questions about their ‘external validity’ to the real world. For example, to quantify utility in the VonNeumann and Morgenstern sense, one needs to know the probabilities associated with the payoffs of different choices. Such probabilities may be known when presented with explicit gambles—what Savage [30] called ‘small worlds’—but that is rarely the case in the real world. Real choices are almost always made in ‘large world’ situations, where the number of samples is much too low to capture the variance along all of the potentially relevant dimensions. Statistical methods, Bayesian or otherwise, do not function well in such scenarios because their assumptions cannot be met [31], and simple heuristics often produce superior performance (see [9] for review). Thus, normative economic theories do not apply to all economic choice behaviour, and savvy marketers and advertisers often ignore them in practice.

Second, many theories of economic choice (particularly those currently studied in neuroscience) are aimed at the problem of how we decide between simultaneously presented offers, such as between two brands of peanut butter in a supermarket. In natural behaviour, however, simultaneously offered goods are relatively rare. Much more common are ‘sequential’ choices between either taking an opportunity that is encountered versus looking for other opportunities that might (or might not) be encountered within the current environment [32]. Because of this, the reinforced value of given options is strongly dependent on context—finding a small fruit during the dry season is more subjectively rewarding than finding a large fruit during times of plenty. Because animals presumably evolved to deal with the more common sequential choice scenarios, their learning mechanisms exploit the particularities of such situations. Consequently, when faced with the rare occurrence of a simultaneous choice, animals often make what appear as objectively irrational choices (i.e. choose the smaller fruit). An extensive series of studies by Kacelnik and co-workers [32–34] has shown that this is precisely the case, even when the offer value is unambiguously defined using temporal discounting.

Such apparently irrational behaviour is also exhibited by humans. A classic example is the phenomenon of preference reversals [35]: when humans are simultaneously presented with two gambles with similar expected values (e.g. 90% chance of winning \$8 versus 50% chance of winning \$16), they tend to prefer the one with a higher probability of payoff (e.g. the former). However, if the same gambles are presented sequentially and subjects are asked to assign to them a monetary value, they tend to put a higher price on the one they had previously rejected (e.g. the latter). This may occur because of inference errors about probability [36], overpricing of low-probability high-payoff bets [35] or because in both situations the subjects use completely different heuristics that are not congruent [37].

The above examples reveal well-known limitations of prescriptive economic theory as a general formalism for understanding all decision-making behaviour. Here, however, we wish to focus on a third limitation, which to our knowledge has not been widely acknowledged in neuroscience, and yet which may be the most relevant to linking behaviour with the functional neural architecture that implements it.

### 3. The challenges of embodied decisions

The brain is a product of evolution, which rests on two fundamental principles: natural selection; and descent with modification. The first of these is widely acknowledged in economic theory, motivating the search for mechanisms that confer a selective advantage. The second principle, however, is even more valuable for developments of theories that span across levels of analysis. This is because evolution is remarkably conservative, powerfully constrained by phylogenetic history. Evolution can only work with an existing ancestral system, and only make modifications that maintain the system's integrity. In the language of optimization theory, it does not perform an exhaustive search but is highly myopic and constrained to local regions whose basins may be separated from other classes of solutions by impassable crevasses. Major evolutionary change is primarily driven by changes in the environment, not by sudden creative leaps in the genome.

The consequences of phylogenetic constraints cannot be overemphasized. Contrary to persisting opinions among many psychologists and economists, brain evolution is not characterized by the addition of new structures or modules. The idea that the brain proceeds by superimposing a 'mammalian' cerebral cortex over a neopallium and 'reptilian' midbrain [38] has long ago been rejected by studies of comparative neuroanatomy [39,40]. Evolution cannot simply invent a new structure, complete with a full developmental schedule, and connect it up to an existing system. Instead, brain evolution involves the elaboration and specialization of existing structures and elongation of ancestral circuits whose neuroanatomical topology is remarkably conserved [41]. For example, circuits interconnecting the striatum, pallidum and cerebral cortex/neopallium are effectively the same in both mammals and birds [42], despite the fact that their respective lineages diverged about 310 Ma [43]. Indeed, such major features of the human telencephalon are shared with teleost fishes, which diverged from the lineage leading to humans about 450 Ma.

Because the neuroanatomical organization of these circuits was laid down so long ago, we obviously should not attempt to explain their functional roles from a perspective that focuses exclusively on human cognitive behaviour. Instead, we should first consider the kinds of behavioural challenges that were faced by our distant ancestors. This even pertains to the cerebral cortex, which is by no means a mammalian invention but a structure homologous to the dorsal pallium that is shared among nearly all terrestrial tetrapods [39,44]. Thus, even when we address the mechanisms of more modern cognitive functions, we can benefit by considering them as elaborations of ancient behavioural abilities. For example, as the mammalian brain evolved, the architecture of simple decisions may have been recapitulated towards progressively more anterior cortico-striatal circuits that serve progressively more abstract decision-making behaviour [45].

Acknowledging and addressing the phylogenetic constraints of any evolved system may appear as a burden for theorists, another challenge to what is already a challenging field of study. We are of precisely the opposite opinion. What theoretical neuroscience needs most are more constraints, not fewer, for they provide guidelines for effectively choosing which conceptual path is most promising. Even the most elegant and optimal theory is a dead-end if it is incompatible with the constraints of phylogenetic descent. By contrast, the

outline of an evolutionarily grounded theory is a promising topic for further study. Furthermore, we believe that an evolutionary perspective also helps us choose what kinds of questions to ask, and which to ask *first*.

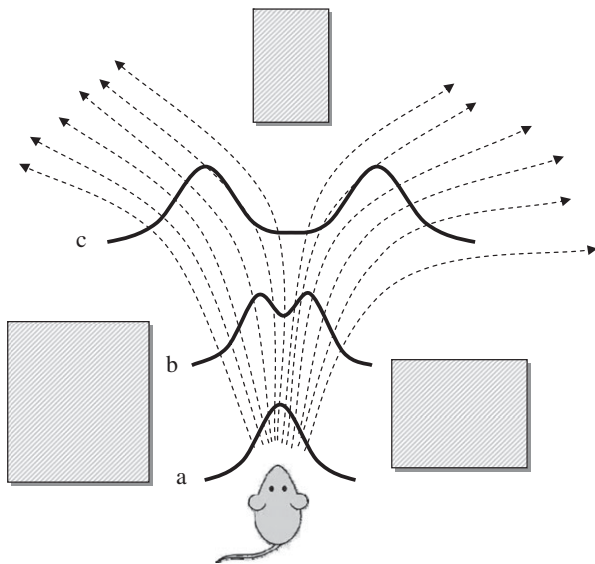
With respect to decision-making, the evolutionary perspective motivates us to build theories of decision-making that are fundamentally aimed at addressing the challenges of the kinds of decisions faced by our very distant ancestors, whose behaviour was primarily interactive and not deliberative. Here, we will take this approach and focus on what may be called 'embodied decisions'—decisions between actions during ongoing activity.

For example, an animal escaping from a predator is continuously making decisions about the direction to run, ways of avoiding obstacles [46], and even foot placement on uneven terrain [47]. Of course, humans also engage in such embodied decision-making during our daily lives, whether we are walking through a crowd or playing a sport. Importantly, embodied decisions have properties that are dramatically different from the economic choices that have dominated decision theories. First, the options themselves are potential action opportunities that are directly specified by the environment—what Gibson [48] called 'affordances'. The variables relevant to evaluating these options are overwhelmed by geometric and biomechanical contingencies and not merely related to offer values. Consequently, evaluation of the sensorimotor contingencies becomes the major challenge for the neural mechanism, whereas pure offer value estimation is computationally relatively trivial. Second, the options themselves are not categorical, like button presses in a psychology experiment. Instead, they are specified by spatio-temporal information, highly dependent on geometry, and even their identity is extended and blurry at the edges. Third, embodied decisions are perhaps the primary and archetypical kind of *simultaneous* decision. Animals encounter goods sequentially, but they are always surrounded by simultaneous action opportunities between which they must select.

Finally, embodied decisions are highly dynamic. As an animal moves through its world, available actions themselves are constantly changing, some are vanishing while others appear, and all the relevant variables (outcome values, success probability, action cost) are always in flux. This precludes any mechanism relying on careful deliberation about static quantities or estimation of probabilities from similar examples, because each embodied decision is a single-trial situation with unique settings. Consequently, the mechanisms that serve embodied decisions must process sensory information rapidly and continuously, specifying and re-specifying available actions in parallel while at the same time evaluating the options and deciding whether to persist in a given activity or switch to a new one. Thus, the temporal distinction between thinking about the choice and then implementing the response, so central to economic theory and laboratory experiments on decisions, simply does not apply to decisions made during interactive behaviour.

Consider the very simple example shown in figure 1. At the beginning (point a) environmental constraints present the mouse with a distribution of potential running directions that may be averaged into a single central forward path. As the mouse runs forward (point b) the obstacle in front begins to be relevant, but the choice between turning left versus right is not yet necessary and the distribution of potential directions can still be averaged into the forward path. As that obstacle is





**Figure 1.** Schematic embodied decision-making scenario. Dotted line arrows indicate possible paths for the mouse to move between obstacles depicted by shaded rectangles. Solid curves indicate the distribution of potential directions at three points in time. At point (a), the distribution can be averaged into a single central direction. At point (b) the distribution begins to separate, but averaging is still possible. At point (c), however, the average is no longer a viable direction and a decision must be made between directions to the right or left.

approached, however (point c), the forward path becomes increasingly undesirable, splitting the distribution of potential directions into two distinct choices, between which a decision must now be made. Note that in this scenario the processes of specifying potential actions and selecting between them have to take place in parallel during overt activity. While complete preplanning of running trajectories may seem acceptable in some situations, it will fail if the obstacles themselves can move unpredictably.

A formal description of such dynamic situations is difficult, but we believe it is both possible and necessary if we are to understand decision-making in the context within which it evolved. Some progress towards a description of embodied decision-making has been made in ecological psychology [49,50], in which decisions are seen as emerging within the dynamical system describing an agent's interaction with the world. In this view, the environment at a given moment defines a potential field through which the state of the agent may flow. Goals are attractors towards which the state tends to move and decision points are the ridges between the basins of attraction defined by different goals (mathematically, decisions are made on the unstable manifolds between different stable attractors of a dynamical system). Importantly, however, the flow field is not static but changes continuously as the agent moves and as the world changes around him/her. Considering again figure 1, the dynamical system at point b still contains a single attractor but as the mouse moves forward a bifurcation occurs and the new dynamical system contains two regions of attraction between which a decision must be made. Analysis of behaviour in sports shows that it exhibits properties characteristic of this kind of dynamical description [49].

While such work provides a formal language for describing embodied decision-making situations, it does not propose hypothetical mechanisms for how the brain may deal with them. What we desire is a theory that spans analysis levels

from descriptions of the problem all the way to neural data on the mechanisms involved in implementing the solutions. However, to date, most neurophysiological studies of the neural mechanisms of decisions use economic concepts, not dynamical system concepts, to define the experimental paradigm and interpret the data. Consequently, they do not even address the kinds of embodied decisions that, we would argue, defined the challenges faced by the vertebrate nervous system at the time, many millions of years ago, when its basic functional architecture was being established.

We have recently proposed that a great deal of neurophysiological data is in fact more naturally interpretable from an ecological and embodied perspective than from traditional cognitive perspectives [51,52]. In particular, the traditional distinctions between separate systems for perceiving the world, storing and retrieving memories, thinking about an abstract goal, and acting upon it do not find strong support in neurophysiological and neuroanatomical data [52–54]. Instead, more conducive to making sense of neural data may be a distinction between processes related to specifying potential opportunities for interaction with the world versus processes related to selecting between these. Here, we briefly review this proposal before describing some recent experiments relevant to its predictions and implications.

#### 4. A mechanism for embodied behaviour

The 'affordance competition hypothesis' [51] suggests that during interactive behaviour, the brain is continuously processing sensory information to specify a set of potential actions (affordances) currently available in the world, while at the same time collecting information to help select between these (see also [55–59]). For visually guided action, the specification takes place along the occipito-parietal 'dorsal' visual stream [60], while selection uses information from the occipito-temporal 'ventral' stream and its projections through the lateral prefrontal cortex [61]. From a computational perspective, action specification involves reciprocally interconnected sensorimotor maps in parietal and frontal cortex, each of which represents potential actions as peaks of activity in tuned neural populations and obstacles as valleys [62,63]. Within such maps, distinct potential actions compete against each other through mutual inhibition, through direct cortico-cortical connections or competing cortico-striatal loops [64], until one potential action suppresses the others and the decision is made. This distributed competition is continuously influenced by a variety of biasing inputs, including rule-based inputs from prefrontal regions [20,22,23,65–69], reward predictions from basal ganglia [70,71] or any variable pertinent to making a choice. While these diverse biases may contribute their 'votes' to different loci along the distributed fronto-parietal sensorimotor competition, their effects are shared across the network owing to its reciprocal connectivity. Consequently, the decision is not determined by any single central executive, but simply depends upon which regions are the first to commit to a given action strongly enough to pull the rest of the system into a 'distributed consensus' [72]. Simultaneous recordings across the cerebral cortex suggest that during simple decisions, such a consensus is reached in a diversity of brain regions at approximately 150 ms [73].

A role of cortical regions in action selection does not imply that subcortical regions are not also involved. In particular,

there is ample evidence that the basal ganglia play a significant role in action selection [64,74–76]. One possibility is that the basal ganglia played a major role in action selection in early vertebrates, which did not have an elaborate cortex, but over time the cortical part of the cortico-striatal system grew in complexity and took on a larger role in selection. A large topologically organized structure such as the cerebral cortex is particularly well-suited to performing selection between spatially distinct potential actions within a single effector system (e.g. selecting between reaching to different objects), while a centralized ‘hub’ like the basal ganglia may be most involved in selecting between different types of action (e.g. selecting whether to reach or to run, or to do nothing at all). Such a division of labour may provide a way to reconcile proposals that the basal ganglia are involved in selection with recent experiments showing that inactivation of the main output nuclei does not impair selection, but instead reduces movement vigour [77,78]. Perhaps disrupting the basal ganglia output to the reaching system did not interfere with a cortical selection of ‘where to reach’, but simply reduced a bias signalling ‘whether to reach’.

We refer the reader to earlier work reviewing the neurophysiological basis of the affordance competition hypothesis [51,52,72] as well as other theoretical frameworks that share similar features [55,59,79]. Here, we describe some of the empirical predictions made by such parallel frameworks, as well as some recent studies whose results bear upon these predictions.

## 5. Implications and predictions of an embodied brain architecture

The affordance competition hypothesis makes a number of testable predictions at the neurophysiological level. For example, it predicts that neural activity within the sensorimotor system will not only represent multiple potential actions, as previously shown [80–86], but that these representations will also be modulated by any factor relevant for making the choice. Decision-related modulation has of course been found throughout frontal and parietal regions of both the oculomotor and arm reaching systems [83,87–93], but the affordance competition hypothesis suggests more specific predictions. In particular, neural activity in the sensorimotor system should be modulated by all factors relevant for the animal’s choice, including offer values, action costs, temporal discounting or anything related to the subjective desirability of an action [87]. Furthermore, neural activity in regions closest to movement initiation should be strictly related to *relative* desirability. That is, neural activity related to a given potential action should be modulated by the desirability of that option relative to other options, and show no decision-related modulation at all when only a single option is present or when all are equally desirable. The reason for this is simple: the neural activity does not itself represent the desirability of options, but merely reflects the competition between options, biased by their desirability, that results in a choice. Finally, the strength of the competition between options should depend on their geometric relationship.

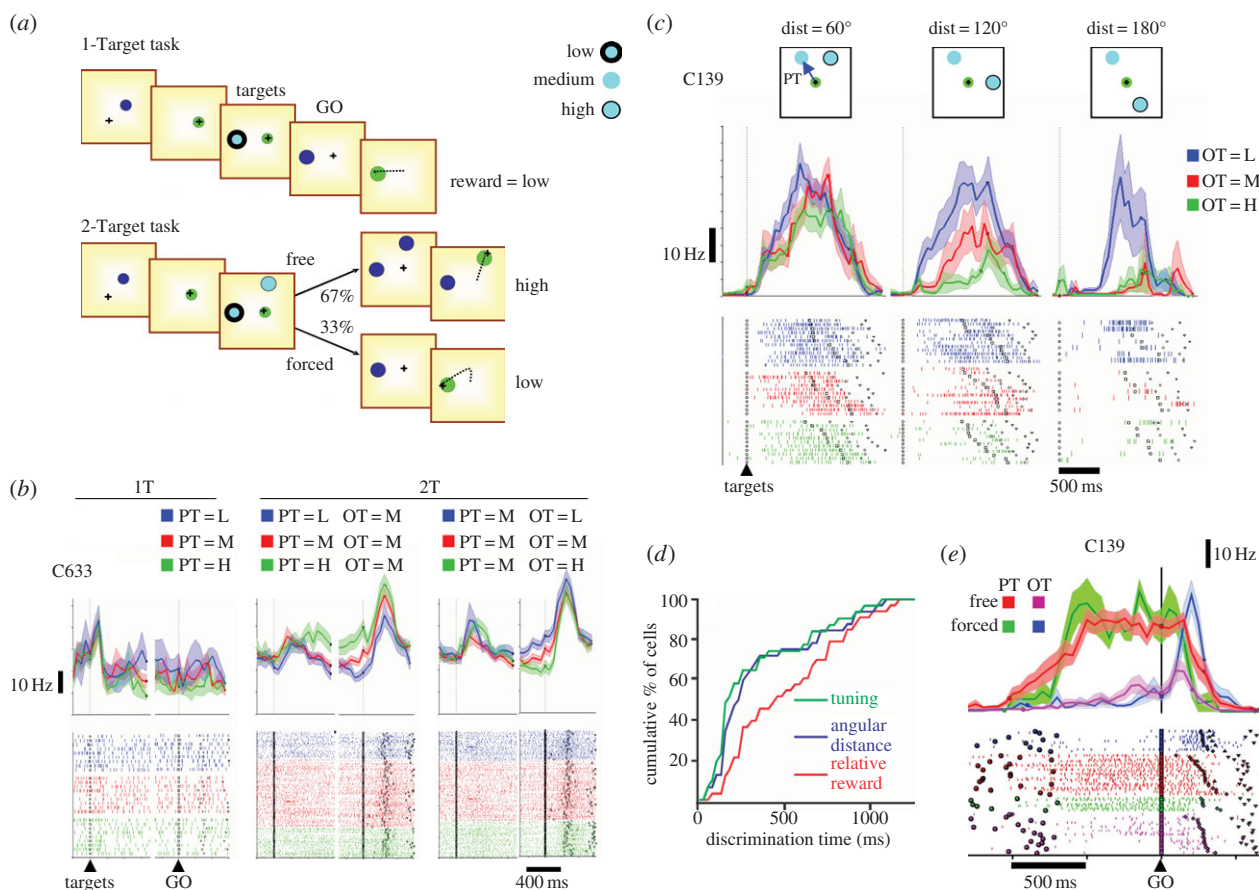
The last point deserves more explanation. The hypothesis that the competition between actions occurs within sensorimotor maps, which reflect the geometry of the environment, predicts that the mutual inhibition between two actions is stronger when they are further apart. The reason for this is

straightforward. When faced with two families of actions that are similar (e.g. turning by  $5^\circ$  to the right versus  $5^\circ$  to the left, as in point b in figure 1) the decision can be gradual, only slightly veering towards one option versus another. By contrast, when faced with very different response options (e.g. at point c in figure 1), a weighted-average solution is not acceptable anymore and an all-or-none decision has to be made. A simple way for a neural population to implement this sensorimotor contingency is for the inhibitory interaction between cells to depend on the similarity of the actions to which those cells are tuned [63]. Such an architecture makes the specific prediction that if one records from cells related to a given option while modulating the desirability of a different option, the gain of that modulation will be the strongest when that other option is most dissimilar to the one coded by the recorded cell.

We tested these predictions through neural recording experiments in dorsal premotor cortex (PMd) of monkeys trained to move a cursor to one of two potential targets with different reward values (figure 2a; [88]). When allowed a free choice the monkeys nearly always chose the target with a higher value, and we were interested in the patterns of neural activity in PMd prior to the GO signal. The results are shown in figure 2. To summarize, neural activity in PMd was indeed modulated by relative reward values when two targets were presented (figure 2b, centre and right), but these same cells never showed any reward-related modulation in trials with only one target (figure 2b, left). This is consistent with a purely relative code (or ‘full divisive normalization’), as predicted by the hypothesis of a recurrent competitive network that is biased by inputs carrying information about absolute values [63].

Most importantly, the gain of the interaction between targets was the strongest when they were furthest apart (figure 2c). This result is critical, in our opinion, because it specifically supports the hypothesis that embodied decisions are made within the space of a sensorimotor map [51], as opposed to being made in the space of outcome values [21]. The dependence of neural activity on the spatial separation between targets is not predicted by any theory based on economic principles, because the spatial separation has no bearing whatsoever on the benefits or the costs of either of the respective movements, or on their relative desirability. Receiving three drops of juice is better than receiving two drops to exactly the same extent regardless of the spatial separation between the targets that yield these rewards. Thus, spatial separation does not enter into any economic equation. And yet, the neural activity that appears to determine the choice of the target is strongly influenced by it (figure 2c). Admittedly, one could modify an economic model so as to incorporate notions of spatial separation between targets and to explain such data, but we believe that would be a post-hoc solution. What we really need from theories are motivated predictions.

If the competition between options takes place within a sensorimotor map, then the spatial relationships between the options should have behaviourally relevant effects. In particular, if the two competing actions are similar and their neural representations overlap, as proposed by Erlhagen & Schoner [95] and Cisek [63], then reaction times (RTs) should be slightly shorter and there should be more intermediate movements (especially with the shortest RTs) than when the candidate movements are further apart. Indeed, when we examined trials in which both targets were medium-valued, RTs were slightly but significant shorter when the targets were  $60^\circ$



**Figure 2.** An experiment on the neural mechanisms of spatial decisions. (a) Task design. In the 1-Target (1T) task, monkeys were presented with a cue whose border style indicated the expected size of reward (see legend at top right) and then given a GO signal to execute the movement. In the 2-Target (2T) task, the monkeys were presented with two possible targets. In 67% of trials (free), they could move to either of these after the GO signal, but in 33% of trials (forced) one disappeared, forcing them to move to the remaining target. (b) Comparison of activity of an example PMd neuron, as a function of reward size, during trials in which the cell's preferred target (PT) was present. In each pair of rasters and peri-event histograms, data are aligned on target onset and GO signal, with a break between them due to delay period variability. Black symbols in the rasters depict target onset, GO signal, and movement onset and offset. If the preferred target was the only one shown (1T task, right), neural activity showed no difference whether its value was low (blue), medium (red) or high (green). However, if two targets were present (2T task), then the same cell showed a strong effect of relative value, both when the value of the preferred target (PT) or the other target (OT) was varied (centre and right). (c) Analysis of the activity of a PMd cell when the PT value was medium and the OT value was varied, done separately depending on the angular distance between targets. Data are aligned on target onset. Note that the modulatory effect of OT value is much stronger when the distance between targets is 120° or 180°, than when they are only 60° apart. (d) Cumulative distributions of the latency with which PMd cells become tuned (green), reflect spatial interactions between targets (blue) and reflect relative value (red). (e) Activity of a PMd cell (aligned on the GO signal) comparing 2T trials in which the monkey was free to choose the higher valued target (free) that was either in the cell's preferred direction (red) or not (purple), and 2T trials in which the monkey was forced to go to the lower valued target (forced low) that was either in the cell's preferred direction (blue) or not (green). Note that activity prior to the GO signal, while both targets were still present, strongly indicates the monkey's preferred plan and activity after the GO signal switches abruptly in forced low trials. Panels (b–d) adapted from Pastor-Bernier & Cisek [88] and panel (e) from Pastor-Bernier *et al.* [94].

apart (mean 265 ms) than when they were 120° apart (273 ms). This difference was small because we used a delay period, but it was nevertheless highly significant (KS test,  $p < 10^{-10}$ ). When subjects are free to react right away, much larger effects of target separation can be seen. For example, Bock & Everaert [96] showed that RTs in target selection tasks depend much more on the angular range spanned by the distracters than by the number of distracters themselves. Several studies have shown that when subjects are forced to select quickly, they often launch their movements in-between targets that are close, but choose randomly between targets that are far apart [97–99].

The affordance competition hypothesis suggests that during visuo-motor decisions, the spatial information defining the options is processed along the dorsal visual stream, which is known to be remarkably fast [73,100]. Thus, if two options suddenly appear, there will a fast feedforward

sweep of information that defines their spatial parameters in fronto-parietal cortex. Because competition occurs within each region along this route, spatial interactions between targets should be evident as soon as cells become tuned. By contrast, because selection biases often take longer to compute, especially if they involve arbitrary mappings between stimulus features and reward size, then their influence on fronto-parietal activity should appear significantly later. Figure 2d illustrates this phenomenon in our PMd study. Of course, such a discrete sequence of events is only expected in the artificial situation of a laboratory experiment. During natural behaviour, spatial information specifying potential actions is constantly provided by the environment along with cues pertaining to their selection, and all of these processes must take place in parallel.

This brings us to the second critical property of embodied decisions: their dynamic nature. As the animal moves around



its environment, the available opportunities and demands for action are continuously changing, both in their spatio-temporal properties as well as their potential costs and payoffs. Consider again the situation faced by the mouse (figure 1). As it moves along a particular path, the relative cost of turning to the right or left is changing, and depends upon the animal's own momentum. Obviously, relevant objects in the world are also moving, redefining the options and motivating the re-evaluation of their relative desirability. This dynamic nature of embodied behaviour makes it challenging to study, especially through single-electrode neural recording experiments in which the data from individual trials is extremely noisy. Great progress is being made in developing methods to record activity from a much larger number of neurons [101] as well as methods to analyse that data [102], raising the possibility that in the future we may be able to conduct studies in which animals are faced with the kinds of dynamic, unique situations for which their nervous systems evolved. To date, however, we must be satisfied with more humble approaches for making sense of the dynamic nature of decisions.

Nevertheless, certain issues are clearly critical once we begin to examine decision-making in a dynamic world. One is the mechanism of how sensory information is treated during the process of deliberation. Most models suggest that sensory information pertinent to a choice is integrated over time until the accumulated information reaches a threshold [103–107]. This is usually called the 'drift diffusion' or 'bounded accumulation' model. Originally, it was conceived as a model of the deliberation that takes place within a purely cognitive system [104,106], but more recently it has also been suggested as a model of how action decisions are made within the sensorimotor system [105]. This relatively simple model explains the timing and accuracy of choices as a function of deliberation time and speed-accuracy trade-off conditions, and explains stimulus-dependent build-up of neural activity over time in a wide range of decision tasks [91,105,108,109]. Consequently, drift diffusion is currently one of the most influential models of decision-making.

However, if the world can suddenly change, then integration is highly sub-optimal. If you have collected evidence in favour of moving towards a particular goal (e.g. a tasty fruit) but suddenly the world changes (e.g. a threat appears near the fruit), then you should change your mind quickly and not have to 'integrate away' the previously collected evidence. Instead, you should always process information quickly, with time-constants appropriate for a given type of situation. This motivates the hypothesis that during dynamic decision-making the brain does not use a slow integration process to estimate the state of the world but instead uses a low-pass filter with a fairly short time-constant (short enough to respond quickly to changes on the relevant timescale, while long enough to filter out irrelevant high-frequency fluctuations in sensory input). The typical duration of fixations suggests that the time-constant of visual decisions is about 100 ms [110], predicting that neural estimates of sensory information should equilibrate within about 300 ms after stimulus onset, even in highly noisy tasks. But then what could be responsible for the neural activity build-up seen in so many decision-making tasks?

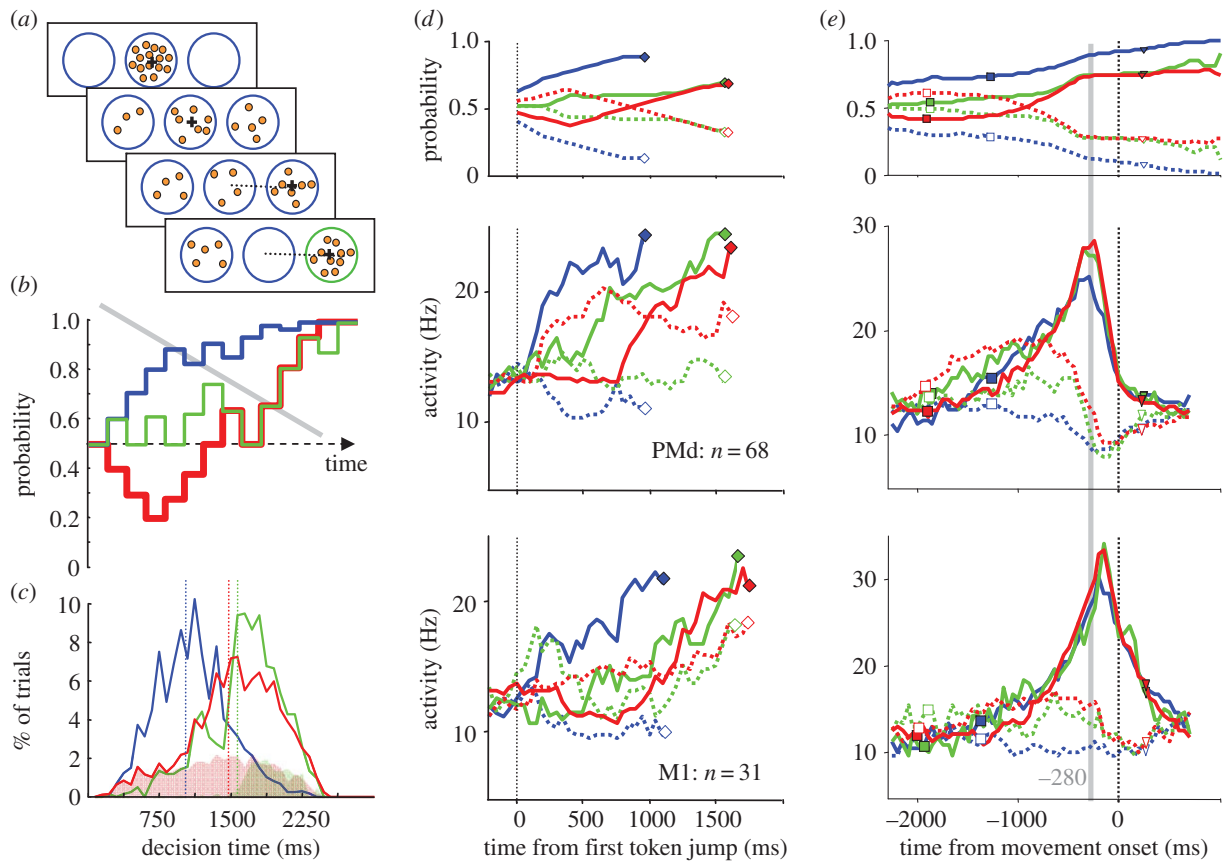
One possible answer comes from considering that during natural behaviour, what matters most is not the accuracy of choices but the reward rate [111,112]. Recent work has

shown, both numerically [113] and analytically [114], that to maximize rewards the brain should use a decision criterion that decreases over time within each trial, in a task-dependent manner. The intuition behind this is simple: if you quickly receive good information for making a decision then you should make the choice right away, but if you are not sure then it may be worthwhile to wait in case new information presents itself. However, if nothing happens and time is running out, you may need to just go with your current best guess. In other words, you should drop your standards over time.

Suppose now that we have a decision-making system in which potential actions compete against each other, biased by evidence in favour of each, and a decision is made when one of these gets strong enough to suppress the others. One simple way to implement a decreasing confidence criterion is to add to this system a non-specific 'urgency' signal that increases over time. When the urgency is low, only an option with very strong evidence will win the competition. As the urgency grows, however, the tension between the options increases (the basins of attraction become steeper) and a smaller difference in evidence may be sufficient to tip the scales [115]. Several neural experiments have suggested that such an urgency signal exists during decision-making in a number of brain regions, including frontal [116,117] and parietal cortex [118]. Indeed, it has been proposed that perhaps it is the urgency signal, not evidence accumulation, which may be responsible for the build-up of neural activity seen in most decision-making experiments [114,119–121].

To summarize, in dynamically changing situations the brain is motivated to process sensory information quickly and to combine it with an urgency signal that gradually increases over time. We call this the 'urgency-gating model', and some recent experiments in our laboratory set out to test its predictions. In these experiments, subjects were faced with decision tasks in which information about the correct choice changes continuously and the decision can be taken at any time [114,117,119]. In one series of studies, humans [119] or monkeys [117] have to guess which of two reaching targets will receive the majority of small circular tokens that jump into one or the other target every 200 ms (figure 3*a*). The subjects can decide at any time, and after they reach the selected target the remaining token jumps accelerate and reward is delivered if the correct guess was made. Thus, subjects are faced with a trade-off: either wait until there is enough information for choosing with confidence, or take an early guess and save some time (improving the reward rate). Importantly, the design of the task allows us to calculate the sensory evidence in favour of each target at each moment in time, and to characterize trials on the basis of the profile of this 'success probability' over time (figure 3*b*).

At the behavioural level, the results from both humans [119] and monkeys [117] are consistent with predictions of the urgency-gating model. In particular, during dynamic decision-making the sensory information appears to be processed quickly, with previous history of the state of evidence quickly 'forgotten' and updated. This was observed even in a task using noisy motion stimuli [114]. Furthermore, when we calculated the level of evidence available to subjects at the time of decision we found that it decreased as a function of decision duration, consistent with the hypothesis that an urgency signal pushes the system to commit with less and less sensory information as time is



**Figure 3.** An experiment on the neural mechanisms of decisions in dynamic situations. (a) Task design. Monkeys were presented with two targets on either side of a central target in which 15 tokens were randomly placed (first row). During each trial, these tokens jumped one-by-one every 200 ms into one or the other target, randomly (second row). The monkeys had to decide which target would receive the majority of the tokens but could make this guess at any time, and after the target was reached (third row), the remaining tokens accelerated and feedback was given when all 15 had jumped (fourth row). (b) The time-course of the probability that a target would be correct, during three different trials: in easy trials (blue) probability tended to quickly converge to a target; in ambiguous trials (green) it tended to remain close to 50% for most of the trial; in misleading trials (red), it tended towards the wrong target in the beginning. The grey line shows a schematic 'decreasing threshold' which explains behaviour in the task (see text). (c) Decision time distributions in easy (blue), ambiguous (green) and misleading trials (red). Shaded regions indicate error trials and dotted vertical lines indicate means. (d) The average activity of 68 PMd neurons (centre row) and 31 primary motor cortex (M1) neurons (bottom row) reflects the probability profile (top row) in easy, ambiguous and misleading trials when the monkey chooses each cell's preferred target (solid lines) or the opposite target (dotted lines). Neural activity also tends to build up over time in both regions. Data are aligned on the first token jump and, to prevent averaging artefacts, truncated at the estimated movement of decision. (e) Average activity of the same PMd and M1 cells aligned on movement onset. Note the prominent peak of activity in PMd cells tuned to the selected target, which occurs approximately 280 ms before movement onset regardless of the trial type (grey line). Note also that around the same time, there is a sharp suppression of activity in the M1 cells tuned to the unselected target, and a later peak of activity (140 ms before movement onset) in M1 cells tuned to the selected target. Adapted from Thura & Cisek [117].

passing—i.e. that subjects drop their standards over time (see the grey line in figure 3b).

At the neural level, these results were confirmed more explicitly [117]. As shown in figure 3d, neural activity in PMd reflected the temporal profile of the sensory information provided by the token movements, together with a trend for increasing activity over time. Similar results were found even in the primary motor cortex (M1), a structure normally seen as merely reflecting the output of decisions taken elsewhere. These observations are consistent with the 'continuous flow' model of Coles *et al.* [122] whereby the motor system is continuously receiving sensory information pertinent to actions, even long before movement onset. What we found in PMd and M1 appears to reflect the changing competition between potential action plans, playing out close to the edge of motor execution.

Importantly, the use of a dynamic decision task allows us to distinguish the process of deliberation from the moment of decision commitment. This lets us conclude that neural activity in the motor system (PMd and M1) reflects decision-making

processes *prior* to the moment of commitment and does not simply spill in from higher regions as proposed by serial models [21]. Furthermore, it allows us to examine the neural events that occur at the moment that, we estimate, the subjects commit to their choice.

As shown in figure 3e, if we align neural activity to the time of movement onset, we see a striking phenomenon occurring in PMd and M1. Approximately 280 ms before movement onset, the activity of PMd cells tuned to the selected target reaches a level of activity that is remarkably consistent across different trial types both in terms of timing and amplitude. At about the same time, the activity of M1 cells tuned to the unselected target is rapidly suppressed. This is exactly what one would predict from a biased competition taking place within the reciprocally connected PMd–M1 network: a sudden phase transition in the dynamical system that represents potential actions. Shortly after this putative 'moment of commitment', M1 cells related to movement production become active and presumably launch the selected action. It should be noted that the



phenomenon observed in PMd–M1 (figure 3e) cannot be simply related to movement execution, because it occurs much too early (280 ms is close to the monkeys' total RTs in a simple instructed task). Furthermore, it is virtually absent in a task in which both the choice of target and the timing of movement are externally instructed [117].

The results shown in figures 2 and 3 suggest that, consistent with the affordance competition hypothesis and the demands of embodied behaviour, when deciding between actions the brain represents the options within sensorimotor regions. These representations compete against each other through spatially dependent mutual inhibition (figure 2c), biased by a continuous flow of sensory information related to the choice as well as a growing signal related to the urge to act (figure 3d). Furthermore, these results suggest that what determines the choice that is made as well as the timing of commitment is the resolution of this competition within the sensorimotor system (figure 3e). This contradicts the proposal, made on the basis of classical economic theories that all decisions are made within a neural representation of outcome values [21].

Of course, not all decisions are the same. Obviously, nobody would suggest that a banker decides between two investments through a competition between the potential actions of signing on two different dotted lines! Here, we are explicitly talking about real-time, embodied decisions made between concrete action choices, such as choices between two directions to run from a predator. For such scenarios, it makes good sense that the resolution of competition between choices takes place in a space of actions. First, action costs must be taken into account [123–125], along with the geometrical relationship between options [88], and the motor system is best suited to combine these kinds of information. Second, the consequences of embodied decisions only begin to play out after an action is initiated, and so it makes good sense to commit to one action or another only in the latest stages of processing. Finally, even after the choice is taken and initiated, it makes sense to keep other options represented in case the world changes.

An analysis of neural responses to changes in the environment is illustrated in figure 2e. In some trials, the more valuable of two targets disappeared at the time of the GO signal, forcing the monkey to change its mind and reach for the remaining target. Importantly, the same PMd cells that appear to be involved in the initial choice between actions continue to be involved in changing the monkey's movement trajectory online (figure 2e). Similar results have been observed during 'target jump' experiments [126–128] and instructed switching tasks [129], suggesting that neurons in the motor system respond to changes of a target or goal within 150 ms of the relevant visual stimulus. The ability to rapidly respond to changes is of course of great survival value during natural behaviour in a dynamic environment.

Dramatic demonstrations of such rapid responses can be seen in recent studies by Scott and co-workers [130,131]. For example, Nashed *et al.* [131] compared electromyographic (EMG) responses to mechanical perturbations during movements made either with or without obstacles. They found that the mid-latency (45–75 ms) EMG response to perturbations was significantly stronger when the perturbation pushed the arm towards an obstacle, even though the obstacle was not in the path to the target. This means that when obstacles are present, the motor system is somehow

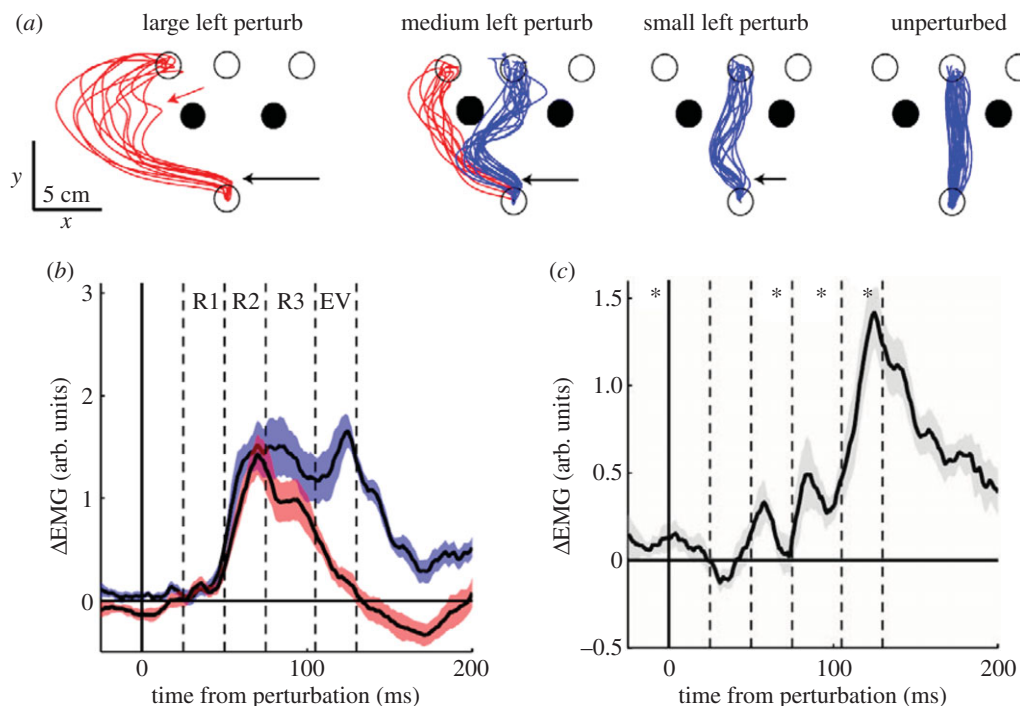
primed to resist perturbations that might cause collisions. Furthermore, a more recent study [130] shows how the motor system can select between different paths among obstacles, or even switch targets when necessary (figure 4a). With a small perturbation, the muscles quickly resisted and brought the trajectory back to its initial path (blue). By contrast, after a large perturbation, such resistance was absent and the limb was allowed to pass around the obstacle to a different target (red), which was an acceptable solution. Most interestingly, with medium-sized perturbations the muscle response was highly state-dependent, sometimes resisting and returning to the straight path, but sometimes 'giving in' and going around the obstacle and to a different target. What is critical to emphasize is that these task and state-dependent responses occurred about 45–75 ms after the perturbation (figure 4b,c). They were much too fast to allow any kind of re-planning or re-evaluation of decisions made in putatively cognitive regions and could involve only a single pass through the transcortical loop. Interestingly, with large perturbations there were two levels of corrective response evident in the data. In particular, in some trials when the arm was pushed around the obstacle to a new target (figure 4a, left), the trajectory first transiently returned towards the original central target (red arrow), as if the motor system selected a new path (around the obstacle) but had not yet selected a new target, and then only later a new target was selected.

In summary, the results of Nashed *et al.* [130] demonstrate that the mere presence of obstacles in the environment sets up the gains of the motor system's resistance to perturbations in a highly task-dependent manner. It appears that even the shortest transcortical loops are sensitive to the layout of objects in space, and are capable of quickly 'deciding' between alternative solutions (e.g. which way to go around an obstacle) and even to decide between different goals. From a dynamical systems perspective, such rapid online corrections could evolve over a potential flow field in which obstacles define unstable equilibria and decision points emerge as the edges between basins of attraction of desirable potential actions. Indeed, if we consider what human athletes (or wild monkeys) are capable of, then perhaps it should be no surprise that online sensorimotor control is this flexible and sophisticated.

## 6. Concluding remarks

In this review, we have argued that attempts to understand decision-making at both behavioural and neurophysiological levels would benefit by addressing the particular challenges of embodied decisions. Motivating this proposal are the assumptions that the brain evolved first and foremost to deal with such challenges, and that the resulting functional architecture has been highly conserved even to the present day.

The affordance competition hypothesis [51] suggests a candidate mechanism for dealing with the challenges of embodied behaviour. It proposes that sensory information is continuously used to specify the parameters of several potential actions in parallel, and these compete against each other under the influence of biases coming from a variety of sources, including the basal ganglia and prefrontal cortex. This differs from classical psychological models in several ways. Most importantly, classical models assume that all processes of deliberation and commitment occur within a prefrontal central executive [16,19,21] and their results are then relayed to motor



**Figure 4.** An experiment on rapid online correction to unexpected mechanical perturbations. (a) In one variant of the study, subjects made movements from a starting circle (at bottom) to any of three target circles (top) while avoiding obstacles (filled circles). In some trials, a perturbation was applied shortly after movement onset (black arrows). In many trials (blue), subjects moved between the obstacles to the central target, but in other trials (red), especially when the perturbation was medium or strong, they moved around the obstacle and reached the left target. In some trials (red arrow), there was a transient deflection towards the central target. (b) Difference in EMG activity of lateral triceps (an elbow extensor) between the unperturbed condition and the medium-perturbation condition in which subjects went between (blue) or around (red) the obstacle. Data are aligned to perturbation time. R1, R2 and R3 indicate the three components of stretch responses to perturbations, and EV indicates the early voluntary component. (c) Difference between the blue and red curves in (b). Note that there is a significant difference already in the R2 response component, 45–75 ms after perturbation, which is believed to involve the shortest latency transcortical loop. Adapted from Nashed *et al.* [130].

systems. By contrast, we propose that, at least in the case of embodied decisions, commitment takes place through a distributed consensus within sensorimotor systems. Deliberation involves both the dynamics of the sensorimotor competition as well as the computation of relevant biases in regions traditionally considered ‘cognitive’. Our hypothesis is quite compatible with recent proposals about activity in these regions. For example, computation of absolute economic value in OFC [22] would clearly be useful for biasing the competition between actions that yield different outcomes. A variety of stages may exist where relevant variables converge to form representations of stimulus value or action value, perhaps in OFC and anterior cingulate cortex, respectively [132,133]. However, a single ‘common currency’ is not strictly necessary because the biases can ultimately be integrated within the sensorimotor competition itself.

A major difference between our hypothesis and classical models is that there is no need for a value-to-action transformation [26]. Information about affordances is always available in sensory input and the dorsal stream and parieto-frontal systems are exquisitely good at extracting that information to specify potential actions. There is no need to convert these into abstract representations, make decisions to choose a goal, and then convert that goal back into an action plan. Instead, what is necessary is a way to map biases such as outcome value onto the (already specified) corresponding actions. While this may involve complex processes such as those ascribed to lateral prefrontal cortex [26,65,134], in many cases the mapping may be solved through very simple strategies. For example, creatures able

to move their gaze may use it to direct the processing of both the ventral and dorsal streams to the same spatial locus, allowing ventro-prefrontal mechanisms to simply instruct the dorsal sensorimotor systems to ‘grasp whatever is in the fovea’ [135].

Despite our emphasis on embodied decisions and the mechanisms that meet their particular challenges, we do not argue against the usefulness of studying the neural mechanisms of pure economic choices. Obviously, humans are capable of making decisions that have nothing to do with action, and understanding such abilities is of great scientific and clinical interest. In fact, it is quite possible that the distinction between different kinds of decisions, such as abstract versus embodied decisions, is paralleled by a distinction between different neural structures and circuits that subservise these scenarios. There is a wealth of theoretical work suggesting such distinctions, in psychology [136] as well as economics [4,137], and neural data suggest that different kinds of decisions involve different brain regions [45,132,133,138,139]. Indeed, phylogenetic continuity motivates us to consider *how* abstract decisions such as economic choice evolved within a system originally adapted for real-time embodied choices, and how the architectures subserving these abilities may be related.

In particular, the concept of affordances can generalize to higher levels of interaction. As long as the world contains reliable causal structure, the brain should be able to discover that structure and learn to detect opportunities for exploiting it [140]. For example, high-level causal structure exists in both social (if I show my teeth, you back off) and economic

domains (if I deposit money in a savings account, I can buy a house later), and it is possible that representations of potential high-level choices compete against each other under the influence of high-level biases. Along the primate lineage the expansion of the frontal lobes, together with the corresponding cortico-striatal and cortico-cerebellar systems, may correspond to a differentiation and specialization of the basic architecture of biased competition towards increasingly abstract domains of interaction [45,141]. Indeed, identifying and understanding the phylogenetic history of this process may be the single most valuable future research direction for understanding the fundamental architecture of human cognition.

## References

- Pinker S. 1997 *How the mind works*. New York, NY: Norton.
- Persky J. 1995 Retrospectives: the ethology of *Homo economicus*. *J. Econ. Perspect.* **9**, 221–231. (doi:10.1257/jep.9.2.221)
- Marr DC. 1982 *Vision*. San Francisco, CA: W. H. Freeman.
- Camerer C, Loewenstein G, Prelec D. 2005 Neuroeconomics: how neuroscience can inform economics. *J. Econ. Lit.* **XLIII**, 9–64. (doi:10.1257/0022051053737843)
- Glimcher PW. 2003 *Decisions, uncertainty, and the brain: the science of neuroeconomics*. Cambridge, MA: MIT Press.
- von Neumann J, Morgenstern O. 1944 *Theory of games and economic behaviour*. Princeton, NJ: Princeton University Press.
- Simon HA. 1957 *Models of man*. New York, NY: John Wiley.
- Kahneman D, Tversky A. 1979 Prospect theory: an analysis of decisions under risk. *Econometrica* **XLVII**, 263–291. (doi:10.2307/1914185)
- Gigerenzer G, Gaissmaier W. 2011 Heuristic decision making. *Annu. Rev. Psychol.* **62**, 451–482. (doi:10.1146/annurev-psych-120709-145346)
- Friedman M. 1953 *Essays in positive economics*. Chicago, IL: Chicago University Press.
- Fumagalli R. 2013 The futile search for true utility. *Econ. Phil.* **29**, 325–347. (doi:10.1017/S0266267113000291)
- Levy DJ, Glimcher PW. 2012 The root of all value: a neural common currency for choice. *Curr. Opin. Neurobiol.* **22**, 1027–1038. (doi:10.1016/j.conb.2012.06.001)
- Rangel A, Camerer C, Montague PR. 2008 A framework for studying the neurobiology of value-based decision making. *Nat. Rev. Neurosci.* **9**, 545–556. (doi:10.1038/nrn2357)
- Baddeley A, Hitch G. 1974 Working memory. In *The psychology of learning and motivation* (ed. IGA Bower), pp. 47–90. New York, NY: Academic Press.
- Norman DA, Shallice T. 1980 Attention to action: willed and automatic control of behavior. Center for Human Information Processing Technical Report no. 99.
- Baddeley A, Della Sala S. 1996 Working memory and executive control. *Phil. Trans. R. Soc. Lond. B* **351**, 1397–1404. (doi:10.1098/rstb.1996.0123)
- Fodor JA. 1983 *Modularity of mind: an essay on faculty psychology*. Cambridge, MA: MIT Press.
- Pylyshyn ZW. 1984 *Computation and cognition: toward a foundation for cognitive science*. Cambridge, MA: MIT Press.
- Shallice T. 1982 Specific impairments of planning. *Phil. Trans. R. Soc. Lond. B* **298**, 199–209. (doi:10.1098/rstb.1982.0082)
- Kennerley SW, Walton ME. 2011 Decision making and reward in frontal cortex: complementary evidence from neurophysiological and neuropsychological studies. *Behav. Neurosci.* **125**, 297–317. (doi:10.1037/a0023575)
- Padoa-Schioppa C. 2011 Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.* **34**, 333–359. (doi:10.1146/annurev-neuro-061010-113648)
- Padoa-Schioppa C, Assad JA. 2006 Neurons in the orbitofrontal cortex encode economic value. *Nature* **441**, 223–226. (doi:10.1038/nature04676)
- Wallis JD. 2007 Orbitofrontal cortex and its contribution to decision-making. *Annu. Rev. Neurosci.* **30**, 31–56. (doi:10.1146/annurev.neuro.30.051606.094334)
- Padoa-Schioppa C, Assad JA. 2008 The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat. Neurosci.* **11**, 95–102. (doi:10.1038/nn2020)
- Padoa-Schioppa C. 2009 Range-adapting representation of economic value in the orbitofrontal cortex. *J. Neurosci.* **29**, 14 004–14 014. (doi:10.1523/JNEUROSCI.3751-09.2009)
- Cai X, Padoa-Schioppa C. 2014 Contributions of orbitofrontal and lateral prefrontal cortices to economic choice and the good-to-action transformation. *Neuron* **81**, 1140–1151. (doi:10.1016/j.neuron.2014.01.008)
- Markov NT *et al.* 2014 A weighted and directed interareal connectivity matrix for macaque cerebral cortex. *Cereb. Cortex* **24**, 17–36. (doi:10.1093/cercor/bhs270)
- Tobler PN, Christopoulos GI, O'Doherty JP, Dolan RJ, Schultz W. 2008 Neuronal distortions of reward probability without choice. *J. Neurosci.* **28**, 11 703–11 711. (doi:10.1523/JNEUROSCI.2870-08.2008)
- Kable JW, Glimcher PW. 2007 The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* **10**, 1625–1633. (doi:10.1038/nn2007)
- Savage LJ. 1954 *Foundations of statistics*. New York, NY: Wiley.
- Volz KG, Gigerenzer G. 2012 Cognitive processes in decisions under risk are not the same as in decisions under uncertainty. *Front. Neurosci.* **6**, 105. (doi:10.3389/fnins.2012.00105)
- Freidin E, Kacelnik A. 2011 Rational choice, context dependence, and the value of information in European starlings (*Sturnus vulgaris*). *Science* **334**, 1000–1002. (doi:10.1126/science.1209626)
- Freidin E, Aw J, Kacelnik A. 2009 Sequential and simultaneous choices: testing the diet selection and sequential choice models. *Behav. Process.* **80**, 218–223. (doi:10.1016/j.beproc.2008.12.001)
- Vasconcelos M, Monteiro T, Kacelnik A. 2013 Context-dependent preferences in starlings: linking ecology, foraging and choice. *PLoS ONE* **8**, e64934. (doi:10.1371/journal.pone.0064934)
- Tversky A, Thaler RH. 1990 Anomalies: preference reversals. *J. Econ. Perspect.* **4**, 201–211. (doi:10.1257/jep.4.2.201)
- Bateman I, Day B, Loomes G, Sugden R. 2007 Can ranking techniques elicit robust values? *J. Risk Uncertainty* **34**, 49–66. (doi:10.1007/s11166-006-9003-4)
- Dudey T, Todd PM. 2001 Making good decisions with minimal information: simultaneous and sequential choice. *J. Bioecon.* **3**, 195–215. (doi:10.1023/A:1020542800376)
- MacLean PD. 1973 *A triune concept of the brain and behaviour*. Toronto, Canada: University of Toronto Press.
- Butler AB, Hodos W. 2005 *Comparative vertebrate neuroanatomy: evolution and adaptation*. New York, NY: Wiley-Liss.
- Deacon TW. 1990 Rethinking mammalian brain evolution. *Am. Zool.* **30**, 629–705.
- Karten HJ. 1969 The organization of the avian telencephalon and some speculations on the phylogeny of the amniote telencephalon. *Ann. NY*



- Acad. Sci.* **167**, 164–179. (doi:10.1111/j.1749-6632.1969.tb20442.x)
42. Jarvis ED *et al.* 2013 Global view of the functional molecular organization of the avian cerebrum: mirror images and functional columns. *J. Comp. Neurol.* **521**, 3614–3665. (doi:10.1002/cne.23404)
43. Kumar S, Hedges SB. 1998 A molecular timescale for vertebrate evolution. *Nature* **392**, 917–920. (doi:10.1038/31927)
44. Medina L, Reiner A. 2000 Do birds possess homologues of mammalian primary visual, somatosensory and motor cortices? *Trends Neurosci.* **23**, 1–12. (doi:10.1016/S0166-2236(99)01486-1)
45. Badre D, Kaysers AS, D'Esposito M. 2010 Frontal cortex and the discovery of abstract action rules. *Neuron* **66**, 315–326. (doi:10.1016/j.neuron.2010.03.025)
46. Fajen BR, Warren WH. 2003 Behavioral dynamics of steering, obstacle avoidance, and route selection. *J. Exp. Psychol. Hum. Percept. Perform.* **29**, 343–362. (doi:10.1037/0096-1523.29.2.343)
47. Lajoie K, Andujar JE, Pearson K, Drew T. 2010 Neurons in area 5 of the posterior parietal cortex in the cat contribute to interlimb coordination during visually guided locomotion: a role in working memory. *J. Neurophysiol.* **103**, 2234–2254. (doi:10.1152/jn.01100.2009)
48. Gibson JJ. 1979 *The ecological approach to visual perception*. Boston, MA: Houghton Mifflin.
49. Araujo D, Davids K, Hristovski R. 2006 The ecological dynamics of decision making in sport. *Psychol. Sport Exerc.* **7**, 653–676. (doi:10.1016/j.psychsport.2006.07.002)
50. Turvey MT, Shaw RE. 1999 Ecological foundations of cognition. I. Symmetry and specificity of animal-environment systems. *J. Conscious. Stud.* **6**, 95–110.
51. Cisek P. 2007 Cortical mechanisms of action selection: the affordance competition hypothesis. *Phil. Trans. R. Soc. B* **362**, 1585–1599. (doi:10.1098/rstb.2007.2054)
52. Cisek P, Kalaska JF. 2010 Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **33**, 269–298. (doi:10.1146/annurev.neuro.051508.135409)
53. Culham JC, Kanwisher NG. 2001 Neuroimaging of cognitive functions in human parietal cortex. *Curr. Opin. Neurobiol.* **11**, 157–163. (doi:10.1016/S0959-4388(00)00191-4)
54. Gaffan D. 2002 Against memory systems. *Phil. Trans. R. Soc. Lond. B* **357**, 1111–1121. (doi:10.1098/rstb.2002.1110)
55. Fagg AH, Arbib MA. 1998 Modeling parietal-premotor interactions in primate control of grasping. *Neural Netw.* **11**, 1277–1303. (doi:10.1016/S0893-6080(98)00047-1)
56. Kalaska JF, Sergio LE, Cisek P. 1998 Cortical control of whole-arm motor tasks. In *Sensory guidance of movement, Novartis Foundation Symposium no. 218* (ed. M Glickstein), pp. 176–201. Chichester, UK: John Wiley & Sons.
57. Kim J-N, Shadlen MN. 1999 Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* **2**, 176–185. (doi:10.1038/5739)
58. Passingham RE, Toni I. 2001 Contrasting the dorsal and ventral visual systems: guidance of movement versus decision making. *Neuroimage* **14**, S125–S131. (doi:10.1006/nimg.2001.0836)
59. Shadlen MN, Kiani R, Hanks TD, Churchland AK. 2008 Neurobiology of decision making: an intentional framework. In *Better than conscious? Decision making, the human mind, and implications for institutions* (eds C Engel, W Singer), pp. 71–101. Cambridge, MA: MIT Press.
60. Milner AD, Goodale MA. 1995 *The visual brain in action*. Oxford, UK: Oxford University Press.
61. Sakagami M, Pan X. 2007 Functional role of the ventrolateral prefrontal cortex in decision making. *Curr. Opin. Neurobiol.* **17**, 228–233. (doi:10.1016/j.conb.2007.02.008)
62. Chapman CS, Goodale MA. 2010 Obstacle avoidance during online corrections. *J. Vis.* **10**, 17. (doi:10.1167/10.11.17)
63. Cisek P. 2006 Integrated neural processes for defining potential actions and deciding between them: a computational model. *J. Neurosci.* **26**, 9761–9770. (doi:10.1523/JNEUROSCI.5605-05.2006)
64. Leblois A, Boraud T, Meissner W, Bergman H, Hansel D. 2006 Competition between feedback loops underlies normal and pathological dynamics in the basal ganglia. *J. Neurosci.* **26**, 3567–3583. (doi:10.1523/JNEUROSCI.5050-05.2006)
65. Miller EK. 2000 The prefrontal cortex and cognitive control. *Nat. Rev. Neurosci.* **1**, 59–65. (doi:10.1038/35036228)
66. Tanji J, Hoshi E. 2001 Behavioral planning in the prefrontal cortex. *Curr. Opin. Neurobiol.* **11**, 164–170. (doi:10.1016/S0959-4388(00)00192-6)
67. Tsujimoto S, Genovesio A, Wise SP. 2011 Comparison of strategy signals in the dorsolateral and orbital prefrontal cortex. *J. Neurosci.* **31**, 4583–4592. (doi:10.1523/JNEUROSCI.5816-10.2011)
68. Wallis JD, Kennerley SW. 2010 Heterogeneous reward signals in prefrontal cortex. *Curr. Opin. Neurobiol.* **20**, 191–198. (doi:10.1016/j.conb.2010.02.009)
69. Wallis JD, Miller EK. 2003 From rule to response: neuronal processes in the premotor and prefrontal cortex. *J. Neurophysiol.* **90**, 1790–1806. (doi:10.1152/jn.00086.2003)
70. Hollerman JR, Tremblay L, Schultz W. 2000 Involvement of basal ganglia and orbitofrontal cortex in goal-directed behavior. *Prog. Brain Res.* **126**, 193–215. (doi:10.1016/S0079-6123(00)26015-9)
71. Schultz W. 2004 Neural coding of basic reward terms of animal learning theory, game theory, microeconomics and behavioural ecology. *Curr. Opin. Neurobiol.* **14**, 139–147. (doi:10.1016/j.conb.2004.03.017)
72. Cisek P. 2012 Making decisions through a distributed consensus. *Curr. Opin. Neurobiol.* **22**, 927–936. (doi:10.1016/j.conb.2012.05.007)
73. Ledberg A, Bressler SL, Ding M, Coppola R, Nakamura R. 2007 Large-scale visuomotor integration in the cerebral cortex. *Cereb. Cortex* **17**, 44–62. (doi:10.1093/cercor/bhj123)
74. Brown JW, Bullock D, Grossberg S. 2004 How laminar frontal cortex and basal ganglia circuits interact to control planned and reactive saccades. *Neural Netw.* **17**, 471–510. (doi:10.1016/j.neunet.2003.08.006)
75. Mink JW. 1996 The basal ganglia: focused selection and inhibition of competing motor programs. *Prog. Neurobiol.* **50**, 381–425. (doi:10.1016/S0301-0082(96)00042-1)
76. Redgrave P, Prescott TJ, Gurney K. 1999 The basal ganglia: a vertebrate solution to the selection problem? *Neuroscience* **89**, 1009–1023. (doi:10.1016/S0306-4522(98)00319-4)
77. Desmurget M, Turner RS. 2008 Testing basal ganglia motor functions through reversible inactivations in the posterior internal globus pallidus. *J. Neurophysiol.* **99**, 1057–1076. (doi:10.1152/jn.01010.2007)
78. Desmurget M, Turner RS. 2010 Motor sequences and the basal ganglia: kinematics, not habits. *J. Neurosci.* **30**, 7685–7690. (doi:10.1523/JNEUROSCI.0163-10.2010)
79. Pezzulo G, Ognibene D. 2012 Proactive action preparation: seeing action preparation as a continuous and proactive process. *Motor Control* **16**, 386–424.
80. Baumann MA, Fluet MC, Scherberger H. 2009 Context-specific grasp movement representation in the macaque anterior intraparietal area. *J. Neurosci.* **29**, 6436–6448. (doi:10.1523/JNEUROSCI.5479-08.2009)
81. Cisek P, Kalaska JF. 2005 Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* **45**, 801–814. (doi:10.1016/j.neuron.2005.01.027)
82. Glimcher PW, Sparks DL. 1992 Movement selection in advance of action in the superior colliculus. *Nature* **355**, 542–545. (doi:10.1038/355542a0)
83. Klaes C, Westendorff S, Chakrabarti S, Gail A. 2011 Choosing goals, not rules: deciding among rule-based action plans. *Neuron* **70**, 536–548. (doi:10.1016/j.neuron.2011.02.053)
84. McPeck RM, Keller EL. 2002 Superior colliculus activity related to concurrent processing of saccade goals in a visual search task. *J. Neurophysiol.* **87**, 1805–1815.
85. Scherberger H, Andersen RA. 2007 Target selection signals for arm reaching in the posterior parietal cortex. *J. Neurosci.* **27**, 2001–2012. (doi:10.1523/JNEUROSCI.4274-06.2007)
86. Song JH, McPeck RM. 2008 Target selection for visually-guided reaching in the dorsal premotor area during a visual search task (Society for Neuroscience Abstracts). *J. Vis.* **8**, article 547. (doi:10.1167/8.6.547)
87. Dorris MC, Glimcher PW. 2004 Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* **44**, 365–378. (doi:10.1016/j.neuron.2004.09.009)
88. Pastor-Bernier A, Cisek P. 2011 Neural correlates of biased competition in premotor cortex. *J. Neurosci.*

- 31, 7083–7088. (doi:10.1523/JNEUROSCI.5681-10.2011)
89. Pesaran B, Nelson MJ, Andersen RA. 2008 Free choice activates a decision circuit between frontal and parietal cortex. *Nature* **453**, 406–409. (doi:10.1038/nature06849)
90. Platt ML, Glimcher PW. 1999 Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238. (doi:10.1038/22268)
91. Roitman JD, Shadlen MN. 2002 Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475–9489.
92. Shadlen MN, Newsome WT. 2001 Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *J. Neurophysiol.* **86**, 1916–1936.
93. Sugrue LP, Corrado GS, Newsome WT. 2005 Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat. Rev. Neurosci.* **6**, 363–375. (doi:10.1038/nrn1666)
94. Pastor-Bernier A, Tremblay E, Cisek P. 2012 Dorsal premotor cortex is involved in switching motor plans. *Front. Neuroeng.* **5**, 5.
95. Erlhagen W, Schöner G. 2002 Dynamic field theory of movement preparation. *Psychol. Rev.* **109**, 545–572. (doi:10.1037/0033-295X.109.3.545)
96. Bock O, Eversheim U. 2000 The mechanisms of movement preparation: a precuing study. *Behav. Brain Res.* **108**, 85–90. (doi:10.1016/S0166-4328(99)00134-5)
97. Chapman CS, Gallivan JP, Wood DK, Milne JL, Culham JC, Goodale MA. 2010 Reaching for the unknown: multiple target encoding and real-time decision-making in a rapid reach task. *Cognition* **116**, 168–176. (doi:10.1016/j.cognition.2010.04.008)
98. Favilla M. 1997 Reaching movements: concurrency of continuous and discrete programming. *Neuroreport* **8**, 3973–3977. (doi:10.1097/00001756-199712220-00025)
99. Ghez C, Favilla M, Ghilardi MF, Gordon J, Bermejo R, Pullman S. 1997 Discrete and continuous planning of hand movements and isometric force trajectories. *Exp. Brain Res.* **115**, 217–233. (doi:10.1007/PL00005692)
100. Schmolesky MT, Wang Y, Hanes DP, Thompson KG, Leutgeb S, Schall JD, Leventhal AG. 1998 Signal timing across the macaque visual system. *J. Neurophysiol.* **79**, 3272–3278.
101. Ahrens MB, Orger MB, Robson DN, Li JM, Keller PJ. 2013 Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nat. Methods* **10**, 413–420. (doi:10.1038/nmeth.2434)
102. Churchland MM, Yu BM, Sahani M, Shenoy KV. 2007 Techniques for extracting single-trial activity patterns from large-scale neural recordings. *Curr. Opin. Neurobiol.* **17**, 609–618. (doi:10.1016/j.conb.2007.11.001)
103. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. 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. (doi:10.1037/0033-295X.113.4.700)
104. Busemeyer JR, Townsend JT. 1993 Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* **100**, 432–459. (doi:10.1037/0033-295X.100.3.432)
105. Gold JJ, Shadlen MN. 2007 The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574. (doi:10.1146/annurev.neuro.29.051605.113038)
106. Ratcliff R. 1978 A theory of memory retrieval. *Psychol. Rev.* **83**, 59–108. (doi:10.1037/0033-295X.85.2.59)
107. Usher M, McClelland JL. 2001 The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* **108**, 550–592. (doi:10.1037/0033-295X.108.3.550)
108. Heekeren HR, Marrett S, Ungerleider LG. 2008 The neural systems that mediate human perceptual decision making. *Nat. Rev. Neurosci.* **9**, 467–479. (doi:10.1038/nrn2374)
109. Ratcliff R, Hasegawa YT, Hasegawa RP, Smith PL, Segraves MA. 2007 Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *J. Neurophysiol.* **97**, 1756–1774. (doi:10.1152/jn.00393.2006)
110. Uchida N, Kepecs A, Mainen ZF. 2006 Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nat. Rev. Neurosci.* **7**, 485–491. (doi:10.1038/nrn1933)
111. Balci F, Simen P, Niyogi R, Saxe A, Hughes JA, Holmes P, Cohen JD. 2011 Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Atten. Percept. Psychophys.* **73**, 640–657. (doi:10.3758/s13414-010-0049-7)
112. Chittka L, Skorupski P, Raine NE. 2009 Speed-accuracy tradeoffs in animal decision making. *Trends Ecol. Evol.* **24**, 400–407. (doi:10.1016/j.tree.2009.02.010)
113. Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, Pouget A. 2012 The cost of accumulating evidence in perceptual decision making. *J. Neurosci.* **32**, 3612–3628. (doi:10.1523/JNEUROSCI.4010-11.2012)
114. Thura D, Beauregard-Racine J, Fradet CW, Cisek P. 2012 Decision making by urgency gating: theory and experimental support. *J. Neurophysiol.* **108**, 2912–2930. (doi:10.1152/jn.01071.2011)
115. Standage D, You H, Wang DH, Dorris MC. 2011 Gain modulation by an urgency signal controls the speed-accuracy trade-off in a network model of a cortical decision circuit. *Front Comput. Neurosci.* **5**, 7. (doi:10.3389/fncom.2011.00007)
116. Gluth S, Rieskamp J, Buchel C. 2012 Deciding when to decide: time-variant sequential sampling models explain the emergence of value-based decisions in the human brain. *J. Neurosci.* **32**, 10 686–10 698. (doi:10.1523/JNEUROSCI.0727-12.2012)
117. Thura D, Cisek P. 2014 Deliberation and commitment in the premotor and primary motor cortex during dynamic decision-making. *Neuron* **81**, 1401–1416. (doi:10.1016/j.neuron.2014.01.031)
118. Churchland AK, Kiani R, Shadlen MN. 2008 Decision-making with multiple alternatives. *Nat. Neurosci.* **11**, 693–702. (doi:10.1038/nn.2123)
119. Cisek P, Puskas GA, El-Murr S. 2009 Decisions in changing conditions: the urgency-gating model. *J. Neurosci.* **29**, 11 560–11 571. (doi:10.1523/JNEUROSCI.1844-09.2009)
120. Ditterich J. 2006 Evidence for time-variant decision making. *Eur. J. Neurosci.* **24**, 3628–3641. (doi:10.1111/j.1460-9568.2006.05221.x)
121. Miller P, Katz DB. 2013 Accuracy and response-time distributions for decision-making: linear perfect integrators versus nonlinear attractor-based neural circuits. *J. Comput. Neurosci.* **35**, 261–294. (doi:10.1007/s10827-013-0452-x)
122. Coles MG, Gratton G, Bashore TR, Eriksen CW, Donchin E. 1985 A psychophysiological investigation of the continuous flow model of human information processing. *J. Exp. Psychol. Hum. Percept. Perform.* **11**, 529–553. (doi:10.1037/0096-1523.11.5.529)
123. Cos I, Belanger N, Cisek P. 2011 The influence of predicted arm biomechanics on decision-making. *J. Neurophysiol.* **105**, 3022–3033. (doi:10.1152/jn.00975.2010)
124. Cos I, Medleg F, Cisek P. 2012 The modulatory influence of end-point controllability on decisions between actions. *J. Neurophysiol.* **108**, 1764–1780. (doi:10.1152/jn.00081.2012)
125. Hosokawa T, Kennerley SW, Sloan J, Wallis JD. 2013 Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex. *J. Neurosci.* **33**, 17 385–17 397. (doi:10.1523/JNEUROSCI.2221-13.2013)
126. Archambault PS, Caminiti R, Battaglia-Mayer A. 2009 Cortical mechanisms for online control of hand movement trajectory: the role of the posterior parietal cortex. *Cereb. Cortex* **19**, 2848–2864. (doi:10.1093/cercor/bhp058)
127. Archambault PS, Ferrari-Toniolo S, Battaglia-Mayer A. 2011 Online control of hand trajectory and evolution of motor intention in the parietofrontal system. *J. Neurosci.* **31**, 742–752. (doi:10.1523/JNEUROSCI.2623-10.2011)
128. Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT. 1983 Interruption of motor cortical discharge subserving aimed arm movements. *Exp. Brain Res.* **49**, 327–340. (doi:10.1007/978-3-642-68915-4\_34)
129. Wise SP, Mauritz KH. 1985 Set-related neuronal activity in the premotor cortex of rhesus monkeys: effects of changes in motor set. *Proc. R. Soc. Lond. B* **223**, 331–354. (doi:10.1098/rspb.1985.0005)
130. Nashed JY, Crevecoeur F, Scott SH. 2014 Rapid online selection between multiple motor plans. *J. Neurosci.* **34**, 1769–1780. (doi:10.1523/JNEUROSCI.3063-13.2014)
131. Nashed JY, Crevecoeur F, Scott SH. 2012 Influence of the behavioral goal and environmental obstacles on rapid feedback responses. *J. Neurophysiol.* **108**, 999–1009. (doi:10.1152/jn.01089.2011)
132. Camille N, Tsuchida A, Fellows LK. 2011 Double dissociation of stimulus-value and action-value learning in humans with orbitofrontal or anterior cingulate cortex damage. *J. Neurosci.* **31**, 15 048–15 052. (doi:10.1523/JNEUROSCI.3164-11.2011)
133. Rudebeck PH, Behrens TE, Kennerley SW, Baxter MG, Buckley MJ, Walton ME, Rushworth MF. 2008 Frontal

- cortex subregions play distinct roles in choices between actions and stimuli. *J. Neurosci.* **28**, 13 775–13 785. (doi:10.1523/JNEUROSCI.3541-08.2008)
134. Quintana J, Fuster JM. 1999 From perceptions to actions: temporal integrative functions of prefrontal and parietal neurons. *Cereb. Cortex* **9**, 213–221. (doi:10.1093/cercor/9.3.213)
135. Cisek P, Turgeon M. 1999 'Binding through the fovea', a tale of perception in the service of action. *Psyche* **5**.
136. Toates F. 1998 The interaction of cognitive and stimulus-response processes in the control of behaviour. *Neurosci. Biobehav. Rev.* **22**, 59–83. (doi:10.1016/S0149-7634(97)00022-5)
137. Kahneman D. 2011 *Thinking, fast and slow*. Toronto, Canada: Doubleday.
138. van der Meer MA, Redish AD. 2010 Expectancies in decision making, reinforcement learning, and ventral striatum. *Front. Neurosci.* **4**, 6. (doi:10.3389/fnins.2010.00028)
139. Walton ME, Behrens TE, Buckley MJ, Rudebeck PH, Rushworth MF. 2010 Separable learning systems in the macaque brain and the role of orbitofrontal cortex in contingent learning. *Neuron* **65**, 927–939. (doi:10.1016/j.neuron.2010.02.027)
140. Cisek P. 1999 Beyond the computer metaphor: behaviour as interaction. *J. Conscious. Stud.* **6**, 125–142.
141. Cisek P, Kalaska JF. 2001 Common codes for situated interaction. *Behav. Brain Sci.* **24**, 883–884. (doi:10.1017/S0140525X0124391X)