10

Cognitive Control as Cost-Benefit Decision Making

Wouter Kool,[1] Amitai Shenhav,[1]
and Matthew M. Botvinick

Many everyday situations afford us a set of default behaviours and cognitive processes that could play out automatically in response to stimuli in our environment. Cognitive control enables us to modify our thoughts and actions away from those defaults in a variety of ways, allowing us as a species to perform great intellectual feats such as planning (D. A. Simon & Daw, 2011), reasoning (Christoff et al., 2001), inhibition (Aron, 2011), and working memory maintenance (Goldman-Rakic, 1987). But what is it that determines when we exert control, how much we do so, and what form(s) this control takes? In other words, by what computational and neural mechanisms is the controller itself controlled (Botvinick & Cohen, 2015; Dayan, 2012)?

In this chapter, we address this question by framing it as a reward-based decision-making problem. This approach views the exertion of cognitive effort as being determined by the output of a decision that considers both the costs and benefits of mobilising cognitive control at a given moment. We begin by enumerating a set of factors that weigh in favour of the exertion of control and those that oppose it. We then present a theoretical framework that specifies how these costs and benefits are integrated together to form a decision about whether and how control should be deployed. Finally, we describe the neural underpinnings of this decision process, with a particular focus on the role of the dorsal anterior cingulate cortex (dACC) in determining how best to allocate control.

Why Exert Control? The Demands for (and Benefits of) Control

As mentioned above, some behavioural situations demand control, whereas others permit a more automatic approach to response selection. A key question, central to research on cognitive control, is how the brain detects this difference. How, in other words, is a demand for control detected? Research has pointed to a number of signals, both internal and external, that appear to serve as cues or signals of demand.

---

[1] The first two authors contributed equally to this work.

Cues to Control Demands

One straightforward way of determining how much control to allocate would be to rely entirely on automatic processes (e.g., habits) until the individual discovers that she has performed the wrong action, for instance because she is told that this is the case or it is obvious as soon as the response is made (e.g., in a situation where time pressure encourages premature responses). After encountering an error, one type of control she can exert is to exercise caution in her subsequent actions (Laming, 1968; Rabbitt, 1966), which can be operationalised as an adjustment in her threshold for responding so as to allow for additional evidence accumulation (and consequently greater accuracy) at the cost of slower responding. This type of sequential adjustment in response to previous-trial errors is commonly observed in the behavioural literature (see Danielmeier & Ullsperger, 2011 for a review, including discussion of inconsistencies in observations of post-error improvements in accuracy). Other kinds of post-error adjustments include reallocating attention away from a stimulus property that is distracting and/or towards a task-relevant property (Danielmeier & Ullsperger, 2011; see Chapter 17 by Ullsperger in this volume).

Rather than relying exclusively on explicit error feedback, which is often unavailable in our daily activities, or expecting errors to otherwise be relatively unambiguous, one can instead simply look out for situations where outcomes are worse than expected (i.e., negative prediction errors) as indicators that greater control is needed. Such a mechanism avoids relying on explicit instruction and provides a continuous rather than binary measure on which to base performance adjustments (i.e., indicating not only that outcomes were worse than expected, but also how much worse). But valuable information can also be garnered from monitoring how different outcomes were from one's predictions, *irrespective* of the direction. Rather than informing specific adjustments in behaviour/control, these *unsigned* prediction errors (or surprise signals) can signal a demand for attending to the elicitor of surprise (Pearce & Hall, 1980), and more generally can act as indicators of volatility/unpredictability in one's environment (Alexander & Brown, 2011; Behrens, Woolrich, Walton, & Rushworth, 2007; Cavanagh & Frank, 2014). Control systems can improve performance in response to these changes in volatility by, for instance, adjusting the rate at which one's world model is adjusted to reflect the changing outcomes, and by extension, the rate at which behaviours are adjusted on the basis of recent feedback (Behrens et al., 2007).

Still, monitoring primarily for prediction errors has the drawback that task feedback can occur too late to adjust control when necessary (i.e., during task performance). One way of avoiding this concern is to monitor for errors relative to richer sets of predictions, for instance, maintaining temporally specific predictions about the occurrence of certain conjunctions of responses and outcomes (Alexander & Brown, 2011). Another alternative to relying solely on outcome-related feedback (whether internal or external) is to take a more direct and online estimate of task difficulty. One such estimate is provided by the current level of processing conflict (Berlyne, 1957), for instance, how much competition exists between different potential responses (Botvinick, Braver, Barch, Carter, & Cohen, 2001). As a proxy for difficulty, levels of conflict on one trial can be used as an indication that greater control should be allocated on a subsequent trial (by, for instance, attending more to the task-relevant stimulus dimension), resulting in lower error rates and faster response times (conflict adaptation; Gratton, Coles, & Donchin, 1992; see Chapter 4 by Egner in this volume). Not only can these conflict signals act as indicators of difficulty, they can also in certain cases obviate error detection altogether by providing an earlier and more sensitive index of potential errors or confidence in one's response in advance of and/or in the absence of feedback (Yeung, Botvinick, & Cohen, 2004).

Cues to Control Incentives

Finally, the amount of overall reward being offered for successful performance of a controlled task can provide a straightforward signal indicating how beneficial it might be to increase control. Accordingly, a number of studies have found that individuals exert greater cognitive effort with higher incentives for performing a cognitive control task (review in Pessoa, 2015; Chapter 22 by Pessoa in this volume; Chapter 24 by Krebs & Woldorff in this volume). For instance, participants performing a Stroop-like task are faster, more accurate, and less influenced by both congruent and incongruent distractors when they are rewarded compared to when they are not rewarded (Etzel, Cole, Zacks, Kay, & Braver, in press; Padmala & Pessoa, 2011). Similarly, participants motivated to either obtain reward or avoid loss are better able to detect a target face appearing against a noisy background (Engelmann, Damaraju, Padmala, & Pessoa, 2009; Engelmann & Pessoa, 2007).

Why Not Exert Control?

There are clearly a number of sources of information that an individual can draw on to determine that there is a demand for greater control. And it seems intuitive that control recruitment *should* be based on such cues to task difficulty and reward. But this intuition hides a thorny problem: Why does control need to be 'recruited' at all? Why must the engagement of control be calibrated to task demands and rewards? Why is control not maximally engaged under all circumstances? Why, in other words, do people not exert maximal mental effort on everything, all the time?

Consider the study by Padmala & Pessoa (2011) cited above. Participants performed a selective attention task in which they had to classify an image while disregarding words that were either unrelated to the image category, or were related and either congruent or incongruent with the classification. Importantly, when a cue indicated that correct performance on the current trial would lead to rewards, performance improved (i.e., reaction times, error rates, and interference effects were reduced). An even more striking example of this effect was provided by Duckworth, Quinn, Lynam, Loeber, and Stouthamer-Loeber (2011), who showed that laboratory participants performing standardised intelligence tests scored more than half a standard deviation higher when performance was incentivised. These findings suggest that performance in the non-incentivised conditions in these studies was marked by less-than-maximal control exertion.

Why did people not exert the maximal amount of cognitive control to improve performance on these tasks? One intuitive answer focuses on the possibility that there is an intrinsic cost attached to the exertion of mental effort. That is, seemingly suboptimal exertion of cognitive control might be the normative result of a rational cost-benefit analysis, weighing the potential rewards of each option against its anticipated costs (Anderson, 1990; H. A. Simon, 1956; Stephens & Krebs, 1986).

Analogous Costs in the Domain of Physical Effort

A similar notion is widely accepted in research on physical-effort-based decision making, where it is assumed that, all else being equal, actions are selected to minimise demands for physical labour. This idea was famously formulated in Hull's (1943) Law of Less Work: all else being equal, agents will prefer lines of action with the smallest demands for physical effort. This basic notion has held currency since at least the 1920s (for a review, see Solomon, 1948) and remains widely influential in modern studies on the neural mechanisms underlying

cost-benefit analyses involving physical costs (Salamone, Correa, Farrar, & Mingote, 2007; Walton, Kennerley, Bannerman, Phillips, & Rushworth, 2006). The law of least work is supported by extensive empirical evidence. The most frequent and direct approach has been to place subjects (animals or humans) in situations where they must choose between two courses of action associated with different exertional demands. When rewards are equated, a bias is typically observed towards the less demanding course of action (Solomon, 1948; Thompson, 1944).

Evidence for Cognitive Effort Avoidance

Many psychologists have posited that an analogous law is at work in situations that involve demands for cognitive processing, that is, that mental effort inheres costs just as physical effort does. For example, Allport (1954) famously explained social prejudice as the product of avoidance of effortful controlled processes required for accurate social judgement. Similar assertions have been made in connection with the fields of behavioural economics (Bonner & Sprinkle, 2002; Camerer & Hogarth, 1999), judgement and decision making (Shah & Oppenheimer, 2008; Smith & Walker, 1993), and executive functioning (De Jong, 2000; Engelmann et al., 2009; Westbrook, Kester, & Braver, 2013). In all of these works, the idea that mental effort is costly has been treated as an explanatory principle. However, until very recently, the existence of control costs had not been subjected to a direct empirical test.

Initial progress towards validating the idea of effort costs was made by Kool, McGuire, Rosen, and Botvinick (2010), who introduced the demand-selection task (DST) to test the hypothesis that people avoid the exertion of mental effort. In the DST, participants face a recurring choice between two alternative lines of action, associated with different levels of demands for cognitive control. The most often-used version of the DST manipulates control demands by varying the frequency of shifts between parity and magnitude judgements of numerical digits. Across several versions of this task (using multiple cognitive control tasks), participants are found to consistently favour the option that commits them to fewer demands for cognitive control (Kool et al., 2010; Kool, McGuire, Wang, & Botvinick, 2013; McGuire & Botvinick, 2010; Schouppe, Ridderinkhof, Verguts, & Notebaert, 2014). Follow-up analyses confirmed that this behavioural tendency could not be fully explained by a motivation to avoid errors, or to minimise time on task. The results appear consistent with a law of least *mental* effort, the idea that, all else being equal, actions tend to be selected to minimise cognitive demand.

The idea that demand registers as a cost predicts, additionally, that people should be more willing to exert effort when appropriate incentives are provided. Evidence for this prediction has come from multiple sources. Kool et al. (2010) introduced a new demand-selection paradigm in which participants tended to avoid an effortful task switch, even if this would lead to increased task completion time. However, when monetary incentives were supplied for each completed trial, participants' propensity to avoid task switches disappeared. This finding suggests that the exertional cost of task options is weighed against associated rewards in a cost-benefit analysis. In a series of convergent studies, Westbrook et al. (2013) presented participants with a series of choices between performing a low-demand working memory task for a certain reward and a more demanding task for a higher wage. The latter payment was titrated until an indifference point was reached between the two tasks. Importantly, across participants, at this indifference point, the pay for the high-demand task was consistently higher than the pay for the low-demand task. As discussed before, other data consistent with this idea shows that performance on attention-demanding tasks is higher in the face of appropriate incentives (Aarts et al., 2010; Engelmann et al., 2009; Engelmann & Pessoa, 2007;

see also Chapter 24 by Krebs & Woldorff in this volume). Together, these studies suggest that people discount the value of reward by the amount of cognitive effort that needs to be exerted (see also Dixon & Christoff, 2012; Vassena et al., 2014).

Why Is Control Costly?

So far, we have explored the question of why people do not exert maximal control all of the time. The research that we reviewed provides a straightforward answer to this question, namely, that there is an intrinsic cost associated with the expenditure of mental effort. However, this observation in turn raises a new question: *Why* is the exertion of cognitive control costly? Over the last decades, several researchers have attempted to answer this question from at least two broad perspectives, one focussing on potential limits on control resources, and the other focussing on the opportunity costs associated with the exertion of cognitive control.

Resource Limits

One class of theories that helps address why cognitive effort appears to be underutilised proposes that the exertion of cognitive control depends on a limited resource, and that the amount of available resource determines to what degree control can be implemented (Baumeister, Bratslavsky, Muraven, & Tice, 1998; Bijleveld, Custers, & Aarts, 2009; Kahneman, 1973; Navon & Gopher, 1979). This theory has been supported by a collection of experiments showing that people are less inclined towards controlled behaviour after they have already engaged in a control-demanding activity, a phenomenon referred to as 'ego depletion' (see Hagger, Wood, Stiff, & Chatzisarantis, 2010; but see also Carter, Kofler, Forster, & McCullough, 2015; Blain, Hollard, & Pessiglione, 2016; Hagger & Chatzisarantis, 2016). According to this framework, control can be viewed as costly insofar as it depletes a resource that could otherwise be conserved or husbanded (Muraven, Shmueli, & Burkley, 2006).

However, recent research on self-control has initiated a trend away from this resource-based account of control, focussing instead on motivational explanations (Inzlicht & Schmeichel, 2012; Job, Dweck, & Walton, 2010). This trend has been supported by, among other things, findings suggesting that depletion effects disappear in the face of increased incentives (e.g., Muraven & Slessareva, 2003). These results place doubt on its potential as a mechanistic explanation and cast doubt on the conservation hypothesis more generally (but see also Holroyd, in press).

Opportunity Costs

In an attempt to offer an alternative to resource-based accounts, Kurzban, Duckworth, Kable, and Myers (2013) have recently proposed an alternative account according to which the cost of effort depends directly on the foregone reward from other options. In their opportunity cost framework, the cost of control is directly attached to the presence of task options that are more valuable than the current task. In other words, effort is the perceived output of ongoing cost/benefit computations that monitor for more profitable lines of actions and is not directly linked to the amount of top-down control that is demanded to successfully encode task representations. The opportunity cost framework proposes that the evolutionary purpose of this mechanism is to signal when to engage in adaptive alternative mental activities in order to gain maximal net benefit from the environment (see also Hockey, 1997).

Feng, Schwemmer, Gershman, and Cohen (2014) have used a computational modelling approach to formulate an alternate version of this idea. They argue that the cost of control may arise from cross-talk that occurs when multiple tasks are performed simultaneously. Assuming that these tasks rely on shared rather non-overlapping neural representations (an approach that leads to maximal flexibility and efficiency), they show that even relatively small levels of cross-talk in such a system can lead to large performance decrements. In this context, one can imagine that the cost of control serves as a penalty militating against this level of multitasking to avoid such decrements, while promoting efficient and flexible processing.

Putting It All Together: A Cost-Benefit Trade-Off in Cognitive Control

The work reviewed above suggests that control allocation is jointly determined by the payoffs and the intrinsic costs associated with cognitive control. However, it leaves open the question of how these factors are combined to produce effort-based decisions about control allocation. There clearly needs to be a principled way of integrating these two factors to determine when and how much to adjust control.

The Overall Expected Value of Control

We recently proposed a normative account of control allocation that adapts approaches from computational reinforcement learning for optimal action selection to the goal of determining the value of exerting control (Shenhav, Botvinick, & Cohen, 2013). Specifically, our theory suggests that we select the amount of cognitive control at a given moment that maximises a quantity we refer to as the overall expected value of control (EVC). The EVC represents a combination of the two factors we described earlier: the expected payoff and the cost of control (Figures 10.1 and 10.3). Payoffs consist of both positive and negative outcomes experienced after implementing a given setting of control signals. Positive outcomes that can result from control allocation include intrinsic rewards associated with task performance (Cacioppo & Petty, 1982; Eisenberger, 1992) and extrinsic rewards like money or social approval. Negative outcomes include monetary losses, physical pain, negative peer evaluations, and other states that we may experience as aversive (e.g., conditions of increased uncertainty or conflict; Fritz & Dreisbach, 2013; Inzlicht, Bartholow, & Hirsh, 2015). The salience of both types of outcome could additionally be modulated by the time required to obtain the outcome, as has been found in the sensitivity of humans and other animals to *reward rate* (i.e., rewards normalised by time to reward receipt) rather than simply rewards per se (Niv, Daw, & Dayan, 2006; Simen et al., 2009). In other words, individuals tend to maximise their net reward per unit time.

Deciding How Much Control to Exert

According to the EVC framework, an increase in the intensity of control has two simultaneous effects. The first is to change the probabilities of performance outcomes (success, failure, or more graded intermediate levels of outcome quality). This, in turn, can lead to greater expected payoffs without necessarily influencing the payoffs themselves. For instance, attending more or less to colour naming in the Stroop task influences the likelihood of giving a correct response, but not the rewards associated with giving such a response (versus an incorrect response). Control can also influence how long it takes you to arrive at

(a)

Control signal intensity ⟶
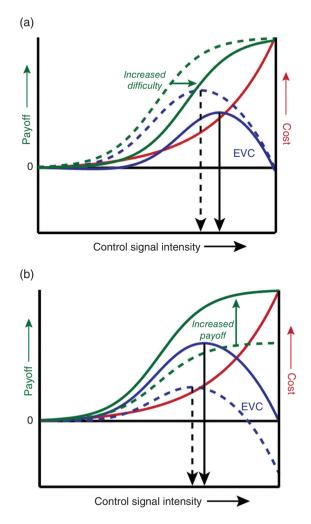
(b)

Control signal intensity ⟶

Figure 10.1 Weighing cognitive demands and payoff in the allocation of control. The expected value of control (EVC; blue) reflects the difference between expected payoffs (green) and intrinsic costs (red) associated with increasing control intensity, assuming a task in which the individual makes a binary choice. Control intensity can be selected on the basis of the maximum of the EVC curve. **Panel a** provides an example of increasing task difficulty (e.g., 3-back rather than 2-back), operationalised as a reduction in the efficacy of a given control intensity at achieving a correct response, resulting in a rightward shift of the payoff curve (dashed to solid lines). **Panel b** provides an analogous example of an increase in the payoff for a correct response. These examples show how both manipulations can, under certain conditions, lead to a decision to increase control. Source: Shenhav 2013. Reproduced with permission of Elsevier.

that state (i.e., the denominator in the reward rate calculation), for instance, by modulating response times. Accordingly, it is important to separate the outcomes described above (e.g., amount of reward for correct and loss for incorrect) from the factors that govern the relationship between control investment and the likelihood of obtaining a given outcome. These latter factors are typically what is referred to by the term *cognitive demand*. Signals of task

difficulty, like errors or amount of conflict, presage the likelihood of reaching a rewarding state given the current allocation of control, versus a situation in which one increases or decreases their control allocation (e.g., on the next trial).

The second consequence of intensifying control is to produce a subjective 'effort-like' cost for the individual. We assume that this cost increases monotonically with the amount of control being allocated. While the exact function relating these two is still a matter of speculation, there is reason to believe that cost may increase nonlinearly (e.g., exponentially) with control intensity (Fudenberg & Levine, 2006, 2011), as has been observed in research on physical effort (Klein-Flügge, Kennerley, Saraiva, Penny, & Bestmann, 2015).

Because both the payoffs and the costs for control are a function of the intensity of control, one can ascertain the EVC of a given intensity of control by taking the difference between these two quantities (see Figure 10.1 for illustrative examples of this cost-benefit analysis). The optimal set of control signals are those that maximise this difference, that is, that maximise the EVC. Because the EVC integrates both payoffs and costs, those factors that indicate increased demand for control (e.g., error feedback, conflict) can lead to greater control allocation but only up to a point (Figure 10.1a), and likewise for factors that indicate increasing reward for successful task performance (Figure 10.1b).

Deciding How Long to Exert Control

The cost-benefit analysis just described focuses primarily on how intensely to exert control. However, many effort allocation problems incorporate another crucial dimension, namely, how to allocate control over periods of time. Imagine, for example, a graduate student sitting in her office on a weekday. Across time, the student has to repeatedly make a choice between two categories of tasks. She can spend time on activities that will further her research, or she can take a break to daydream or peruse social media networks. The decision here is not just related to intensity of control, but can also be characterised as a graded and time-based choice between cognitive labour and cognitive leisure, or mental work and mental rest. In recent work, we tested a formal framework for understanding such temporal allocation problems. This work draws upon a model from labour economics that addresses the question of how workers choose to allocate time between labour and leisure.

The model, labour supply theory (LST), was originally formulated to account for how workers choose to allocate time between labour and leisure (for an introduction, see Nicholson & Snyder, 2008). Here, the benefits from work (i.e., wage) are weighed against its intrinsic costs in a nonlinear fashion, such that the more hours one is working, the more costly each additional unit of work becomes (Figure 10.2). Through this feature of the model, optimal-effort-based decision making results in time allocations that reflect a balance between the payoffs of work, and the intrinsic value of leisure (i.e., the intrinsic cost of work). LST has been validated in numerous experimental and field studies on labour markets (Charness & Kuhn, 2010; Dickinson, 1999; Fehr & Goette, 2007), and has been extended to decision making based on physical effort in animal learning research (Chen, Lakshminarayanan, & Santos, 2006; Conover & Shizgal, 2005; Kagel, Battalio, & Green, 1995).

We recently provided evidence that LST can be applied to decisions involving allocation of cognitive effort over extended periods (Kool & Botvinick, 2014). In these studies, participants divided their time between a mentally demanding task (labour) and a trivially easy task (leisure) within each of a series of experimental sessions. By varying the incentive structure across sessions, we tested a set of critical predictions from LST and found that these held in the case of cognitive effort. The results mirror findings in real-life labour markets and reveal that the cost of mental effort scales in a nonlinear fashion with time
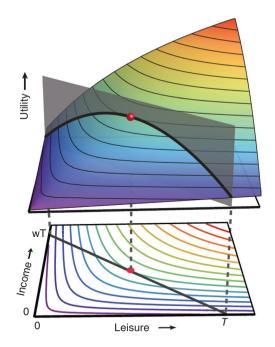
Figure 10.2 The utility function of labour supply theory (LST). As discussed by Kool and Botvinick (2014), LST is a theory of how workers should allocate time between work and leisure. Work is valuable because it yields income via a wage (*w*). Leisure, on the other hand, carries inherent value. Note that this value can also be viewed as the utility of avoiding the effort costs associated with work. Each combination of income and leisure corresponds to a point on the surface at the top, and is associated with one of the iso-utility contours in the bottom of the figure. Allocation decisions are constrained by a time budget, or the total interval of time (*T*) available for the two activities. This time budget is represented by the grey plane that intersects the utility surface in the top part of the figure top, and the diagonal line segment in the lower part. LST assumes that the iso-utility curves of the utility function are convex to the origin, resulting in the inverted-u shaped function on top. The optimal time allocation decision maximises utility over the time budget line. Because the utility is convex to the origin, LST predicts a preference for combinations of income and leisure that balance between the two extreme allocations. Source: Adapted from Kool 2014. Reproduced with permission of APA. Abbreviations: *T,* maximum time available for leisure; *wT,* maximum attainable income.

spent exerting control. This nonlinearity has the effect of motivating individuals to strike a balance between mental labour and leisure.

Deciding What Kinds of Control to Exert

Note that in this last example, the individual is faced with a decision not only over how much control to exert but also to which task or tasks to allocate this control. This highlights a crucial dimension of control that must be selected simultaneously with the amount of control: the type of control that needs adjusting (e.g., attention to different tasks, locations, or stimulus properties; threshold for sampling from episodic memory; affective responses to downregulate). We refer to these as the *identity* and *intensity* of potential control signals, which are the basic unit over which value is maximised within the EVC model described earlier (Figure 10.3). By operationalising cognitive control
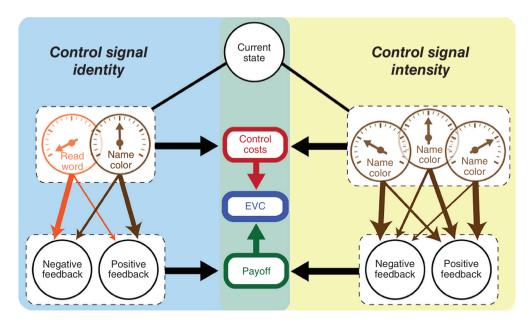
Figure 10.3 Specifying the optimal identity and intensity of control. EVC is simultaneously maximised across possible intensity settings (right) of available control identities (left), for instance, when deciding how much to attend to the colour versus word form when performing the Stroop task. Vertical arrows between gauges and feedback denote likelihoods of obtaining a given outcome under a given control setting. Source: Adapted from Shenhav 2013. Reproduced with permission of Elsevier.

adjustments in this way, the EVC framework is able to extend beyond simple models of reinforcement learning and draw explicit connections to the field of optimal control theory, an area of engineering focused on how to simultaneously optimise the settings of multiple control variables within a dynamical system (e.g., gears and steering-wheel positions on a car) on the basis of feedback from the environment (Kirk, 2012; Wolpert & Landy, 2012).

Because the EVC is simultaneously maximised over the intensities of multiple potential control signal identities, as EVC-relevant variables like difficulty and payoff change, optimal control adaptations can occur within as well as between control signals. For example, increasing difficulty within a certain range may encourage greater engagement of one control signal (e.g., greater attention to task-relevant stimuli), but further increasing task difficulty may lead to engaging a different control signal entirely (e.g., adjustment of decision threshold, or attending to an alternate task).

Under the Hood: Neural Circuitry for Control Evaluation and Specification

According to the EVC framework, describing the mechanisms for control requires delineating (at least) three distinct functions related to control: monitoring, specification, and regulation. *Monitoring* refers to the process of detecting and signalling control-relevant changes in the environment, including changes in the control demands described earlier (e.g., payoff, difficulty). *Specification* refers to the process of integrating these monitored-for

quantities into a determination of which control signals currently maximise the EVC. *Regulation* refers to the process of implementing these control signals. These three functions therefore roughly correspond to the inputs to control decision making (monitoring), the decision making itself (specification), and execution of control (regulation). A fourth relevant function can be described upstream of monitoring, related to *valuation* of individual events and objects or states in one's environment.

On the basis of this taxonomy, we have recently proposed that the dorsal ACC sits at the intersection of monitoring and specification functions of control, a proposal that resonates with a wide array of theories of the dorsal ACC (Cavanagh & Frank, 2014; Holroyd & Yeung, 2012; Kouneiher, Charron, & Koechlin, 2009; O'Reilly, 2010; Ullsperger, Danielmeier, & Jocham, 2014; see also Chapter 15 by Brown in this volume). Specifically, we suggest that the dACC is responsible for evaluating how best to allocate control (i.e., calculating the EVC) and that it signals the output of this decision (i.e., the set of control signals that maximises the EVC) to downstream regions (Shenhav et al., 2013). As we discuss later, this latter function can be seen as effectively licensing or 'motivating' the optimal allocation of control. This account draws on a number of lines of converging evidence surrounding the dACC's role in cognitive control and in the broader neuroscientific literature, the broad strokes of which will be summarised next.

dACC and Monitoring

A number of lines of evidence support our proposed role for the dACC. First, dACC activity has been shown to track each of the variables relevant to evaluating potential control signals within one's current state. This region differentiates between available states and control signals, such as rules (Durstewitz, Vittoz, Floresco, & Seamans, 2010; Johnston, Levin, Koval, & Everling, 2007; Matsuzaka, Akiyama, Tanji, & Mushiake, 2012; Womelsdorf, Johnston, Vinck, & Everling, 2010), actions (Cai & Padoa-Schioppa, 2012; Hayden & Platt, 2010; Isomura, Ito, Akazawa, Nambu, & Takada, 2003; Morecraft & Tanji, 2009), and task sets (Forstmann, Brass, Koch, & Von Cramon, 2006; Momennejad & Haynes, 2013; Wisniewski, Reverberi, Tusche, & Haynes, 2015). Moreover, the dACC tracks a variety of sources of cognitive demands (Cavanagh & Frank, 2014; Duncan, 2010; Paus, Koski, Caramanos, & Westbury, 1998; Venkatraman & Huettel, 2012), from errors (Emeric et al., 2008; Holroyd & Coles, 2002; Ito, Stuphorn, Brown, & Schall, 2003) to conflict (Botvinick et al., 2001; Ebitz & Platt, 2015; Sheth et al., 2012) to signals indicating increased surprise or unpredictability in one's environment (Behrens et al., 2007; Bryden, Johnson, Tobia, Kashtelyan, & Roesch, 2011; Cavanagh & Frank, 2014; Wessel, Danielmeier, Morton, & Ullsperger, 2012). dACC activity has also been found to anticipate and/or mark the occurrence of positive outcomes like juice or monetary gain (Bartra, McGuire, & Kable, 2013; Kouneiher et al., 2009; Wallis & Kennerley, 2011) as well as negative outcomes like electric shock or monetary loss (Botvinick, 2007; Shackman et al., 2011), though it should be noted that it is not always easy to distinguish these activations from signals related to salience or prediction error (Bartra et al., 2013; Cavanagh & Frank, 2014; Engelmann et al., 2009; Litt, Plassmann, Shiv, & Rangel, 2011) and therefore may sometimes reflect attentional demands rather than outcomes per se.

Importantly, feedback-related signals in the dACC (e.g., errors, reward) have properties that suggest that they are in some way tied to control requirements rather than simply being part of an all-purpose mechanism for learning about states in the environment. First, the presence and strength of feedback-related dACC signals is influenced by the amount of information that can be gained as well as the effort required of the task. When monkeys (Quilodran, Rothé, & Procyk, 2008) or humans (Amiez, Sallet, Procyk, & Petrides, 2012)

perform a task involving phases of exploration (attempting to discover the correct response) versus exploitation (repeating this response until it ceases to provide correct feedback), dACC activity is greatest when the feedback indicates discovery of the correct response or a need to explore once again, and is lowest in response to correct feedback during the exploit phase (Khamassi et al., 2010). These findings are consistent with the observation that reward-related activity in the rodent dACC is also diminished when no effort is required to obtain that reward (Hillman & Bilkey, 2010).

Moreover, although a number of findings tie the dACC to levels of surprise or unpredictability, it has also been shown that these surprise signals depend critically on the relevance of the surprising outcome to future control states. For instance, O'Reilly and colleagues (2013) asked participants to saccade to a target that would tend to occur in a constrained region of the screen, with that region changing over the course of the experiment. They found that the dACC was more active when the target appeared in an unexpected location and the target's colour indicated a general shift in possible future target locations, but not when the colour of that surprising target indicated that this trial had no predictive power. In other words, dACC responses to equally unexpected events are modulated by the degree to which that event signalled the need for an update of an internal model (see also Grundy & Shedden, 2014). In further support of our account of dACC monitoring for the purpose of adjusting relevant control signals, the monkey dACC has been found not only to contain separate and overlapping neuronal populations that track positive and negative outcomes (including errors and now also conflict; see below), but this region further distinguishes between different *sources* of errors. For instance, Shen and colleagues (2015) found that error signals in the dACC differentiated between errors that occurred during different stages of a task trial—related to putative failures in sustaining attention, in inhibiting distractors, or in execution of the incorrect response—consistent with a potential role in assigning credit to the appropriate control signal(s) that need to be adjusted (see also Quilodran et al., 2008).

Although some of these findings might simply argue for a role for the dACC in learning and/or selecting between the values of actions (e.g., Hare, Schultz, Camerer, O'Doherty, & Rangel, 2011), recent findings place limitations on such an account. Despite substantial evidence for action values in the dACC, recent studies suggest that the types and timing of value-related responses in the dACC may make this region better suited for post-choice evaluation and control than for direct comparison between action values. For instance, Cai and Padoa-Schioppa (2012) found that during economic choice, dACC action values tend to reflect post-decision variables (e.g., the value of a chosen option or action) rather than pre-decision variables (e.g., the value of the offered options, which are encoded earlier in the orbitofrontal cortex [OFC]). Similarly, Blanchard and Hayden (2014) found that dACC activity in a sequential foraging-like task signalled one post-decision variable (the value of the unchosen option) and one putatively pre-decision variable (the ratio of the offered reward to the delay to reward receipt), but its encoding of the latter interacted with the choice that was made (accept or reject). In other words, the findings support a role for the dACC in using the outcome of a previous decision to modify future strategies and/or control (see also Cowen, Davis, & Nitz, 2012).

Finally, in addition to studies linking the dACC to the encoding of physical effort demands (Hillman & Bilkey, 2010; Hosokawa, Kennerley, Sloan, & Wallis, 2013; Prévost, Pessiglione, Météreau, Cléry-Melin, & Dreher, 2010; Wallis & Kennerley, 2011) and correlates of cognitive demands discussed above, other findings suggest that the dACC may also track the subjective costs associated with exerting cognitive effort (Botvinick, 2007). For example, in one study, Botvinick, Huffstetler, and McGuire (2009) showed that ventral striatal (vStr) responses to both positive and negative feedback were diminished after completing a more rather than less cognitively demanding block of trials (i.e., one that required more rather than

less task switching), and that the vStr response to feedback on high-demand blocks was smaller the greater the dACC responded while performing the task for that block. In a follow-up study, McGuire and Botvinick (2010) found that activity in the dACC and the lateral PFC predicted preferences for avoiding a similarly demanding block of an experimental session (again, involving task switching), though in the dACC this association could be explained in terms of objective performance on the task (i.e., RT, error rate; see also Magno, Foxe, Molholm, Robertson, & Garavan, 2006).

dACC and Control Specification

The aforementioned evidence points to a role for the dACC in evaluating (or at least having access to) quantities relevant to control allocation, and suggests that this region is particularly responsive to those values when they are most relevant. A wide literature further implicates this region in exerting an influence on ongoing control on the basis of these monitored signals (see also Chapter 5 by Brown in this volume).

First, several studies have found overlapping representations in the dACC of specific tasks or rules to be attended (i.e., potential targets of control allocation) and their associated values. This includes the rewards associated with specific saccade directions (Cai & Padoa-Schioppa, 2012; Hayden & Platt, 2010) and with targets of covert attention allocation (Kaping, Vinck, Hutchison, Everling, & Womelsdorf, 2011). Cowen and colleagues (2012) found similar overlaps between the encoding of specific paths along a track and their associated effort demands in the rat dACC (though these responses appeared to be involved in regulating chosen actions rather than planning/choosing). Similarly, while participants performed a task that allowed them to choose which of three task sets to engage on a given trial, and the difficulty of those tasks varied over time (a chosen task grew more difficult over time while the unchosen tasks grew easier), Wisniewski and colleagues (2015) found overlapping regions of the dACC encoded the identities of potential tasks during the choice phase and encoded the difficulty of the chosen task during task performance.

Moreover, changes in dACC activity are associated with and often anticipate changes in task performance. Indirect evidence for this comes from a large body of findings showing that the dACC's responses to salient task events predict subsequent adjustments in control. For instance, the dACC has been implicated in error adaptation, predicting post-error adjustments on the following trial such as slowing (Cavanagh & Frank, 2014; Danielmeier, Eichele, Forstmann, Tittgemeyer, & Ullsperger, 2011; Narayanan, Cavanagh, Frank, & Laubach, 2013; Narayanan & Laubach, 2008) and reallocation of attention between low-level regions representing task-relevant versus task-irrelevant stimulus properties (Danielmeier et al., 2011; King, Korb, von Cramon, & Ullsperger, 2010). The dACC has also been implicated in conflict adaptation, for instance, predicting more efficient resolution of conflict within a high-conflict trial (Sohn, Albert, Jung, Carter, & Anderson, 2007) or on a trial that follows (Forster, Carter, Cohen, & Cho, 2011; Horga et al., 2011; Kerns et al., 2004; Oehrn et al., 2014), often through its connectivity with the lateral PFC. In the first demonstration of conflict-related signals in the monkey dACC, Ebitz and Platt (2015) found that dACC signals reflecting either errors or high-conflict stimuli (distractors that were either incongruent or congruent with a target stimulus direction, evoking what they refer to as *task conflict*) predicted pupillary adjustments on subsequent trials that were separately associated with reduced distractor interference on those trials (Shenhav & Botvinick, 2015; see also Amemori & Graybiel, 2012). Other studies have further linked dACC activity to subsequent choice reversals (Boorman, Rushworth, & Behrens, 2013; Shenhav & Buckner, 2014), changes in learning rate (Behrens et al., 2007; Bryden et al., 2011), and both within- and between-trial adjustments

in decision threshold in reaction not only to errors (as described above) but to increased choice conflict as well (Cavanagh & Frank, 2014; Cavanagh et al., 2011).

Tying the findings above to specific representations of control signals in the dACC, some studies have explicitly decoded representations of different task rules from neural activity in this region and examined the relationship between the strength of those representations and performance on a task. For instance, Johnston and colleagues (2007) showed that neuronal activity in the dACC could be classified on the basis of the current task rule (pro- or anti-saccade) and that the accuracy of these classifications predicted accuracy during the task (see also Rothé, Quilodran, Sallet, & Procyk, 2011; Womelsdorf et al., 2010): When dACC activity in preparation of engaging a given task rule was less selective for the correct rule, an error was more likely.

These findings collectively support the idea that the dACC has the representational capacity that would be expected of a region that maps states and potential control signals to their associated value. It also provides correlational evidence that such representations in the dACC are associated with downstream adjustments in control states. In addition to this, there is also mounting causal evidence that the dACC exerts an influence on control. For instance, patients with lesions in (anterior regions) of the dACC are impaired at tasks that require shifting between different task sets but not, for instance, at a value-based decision-making task (Gläscher et al., 2012). dACC lesions have also been found to impair sequential adaptation in response to control-relevant stimuli. Sheth and colleagues (2012) found that patients who had undergone cingulotomy showed diminished conflict adaptation following their surgery (relative to pre-lesion performance), while Newman, Creer, and McGaughy (2015) found analogous impairments in conflict adaptation in ACC-lesioned rats (but see also Mansouri, Tanaka, & Buckley, 2009). Narayanan and colleagues (2013) found that dACC-inactivated rats exhibited analogous impairments in adapting (i.e., slowing) to an error on the previous trial.

dACC lesions have more generally been found to impair behavioural adjustment to recent reversals in reinforcement (Kennerley, Walton, Behrens, Buckley, & Rushworth, 2006) and to produce motivational deficits such as global response slowing (Stuss & Alexander, 2007) and reduced willingness to overcome effortful obstacles (i.e., increased effort discounting; Walton et al., 2006). Far more dramatic consequences of lesioning dACC and surrounding regions of the medial prefrontal cortex (especially the supplementary motor area) include disorders like abulia, anergia, and akinetic mutism, with which patients are technically able to perform effortful behaviours but lack the motivation or drive/energy to engage in such activities (Stuss, 2011; Stuss & Alexander, 2007). Such findings are broadly consistent with our proposed role for this region in specifying the EVC-maximising control (akin to a motivational or 'willingness to pay' signal; see also Holroyd & Yeung, 2012; Kouneiher et al., 2009).

Our account of the dACC as overseeing the decision to engage control, and outputting the quantity of control deemed worthwhile, has also been borne out by a few studies that employ the opposite causal approach, stimulating rather than inhibiting the dACC. Parvizi, Rangarajan, Shirer, Desai, and Greicius (2013) used depth electrodes to electrically stimulate the dACC in two awake human patients as those patients verbalised their subjective experiences. A common feature of their reports was a feeling of an increased 'willingness to persevere' in the face of perceived environmental challenges. Using noninvasive methods (transcranial direct current stimulation [tDCS]) in healthy participants, Reinhart and Woodman (2014) were similarly able to provide evidence consistent with our account. They found that excitatory stimulation (anodal tDCS) of the dACC and surrounding regions on the medial surface (including SMA) increased monitoring-related ERPs (error- and feedback-related negativities) and led to improved overall accuracy, faster learning, and greater post-error slowing on a stop-signal target discrimination

task, relative to sham controls. The opposite pattern was observed when this region was instead inhibited (cathodal tDCS; see also Reinhart & Woodman, 2015; Rollnik et al., 2004).

By focussing so far on the dACC's role in evaluating and specifying control, we have failed to emphasise a key point, which is that the dACC is central but by no means singular in its involvement in control decision making. A number of structures are presumed to play critical roles in supporting the valuation of control-relevant processes (acting as inputs to the dACC) and in executing control itself (performing the output functions specified by the dACC). At the intersection of valuation and monitoring functions lie regions that include the insula, ventral striatum, orbitofrontal cortex, and ventromedial PFC (Craig, 2009; Haber & Knutson, 2010). On the output end, the regions that implement the necessary control signals vary according to the type and specificity of control required, for instance, whether it involves attention to particular rules or task sets (e.g., by LPFC; Miller & Cohen, 2001; Oehrn et al., 2014; Rothé et al., 2011; Shen et al., 2015); adjustments to decision thresholds to behave more cautiously (e.g., by the subthalamic nucleus; Aron & Poldrack, 2006; Cavanagh et al., 2011; Jahfari et al., 2011); or a change in exploratory versus exploitative modes (e.g., by the locus coeruleus; Aston-Jones & Cohen, 2005).

Conclusions

Throughout this chapter, we have reviewed a number of important domains that encompass the intersection between decision making and cognitive control. At the core of this overlap lies the observation that there are several motivational factors, such as performance-contingent rewards, that have been shown to modulate the degree to which mental effort is deployed (Botvinick & Braver, 2015; Padmala & Pessoa, 2011; Shenhav et al., 2013). This observation suggests that people are reluctant to exert cognitive control, a claim supported by data we reviewed above (Kool et al., 2010; Schouppe et al., 2014). We have also reviewed a decision-making framework that proposes a set of computational and neural mechanisms by which the costs and benefits of control are integrated in order to decide whether and how cognitive control should be deployed.

These findings have obvious ties with other areas of research in which reward-based decision making and executive functioning are implicated. Research on self-control, the important though fallible ability to resist immediate pleasures in favour of longer-term goals (Mischel, Shoda, & Rodriguez, 1989), seems to fit this description. In fact, recent findings suggest that the ability for self-control is related to the cost for cognitive activity (Kool et al., 2013; Westbrook et al., 2013), and recruits the same neural circuitry that is typically associated with cognitive control (Hare, Camerer, & Rangel, 2009; Heatherton, 2011; McClure, Laibson, Loewenstein, & Cohen, 2004). Other fields of research that could benefit from the current approach include the study of voluntary task selection (Arrington & Logan, 2004; Orr & Weissman, 2011) and perseverance towards goals over longer time scales (i.e., 'grit'; Duckworth, Peterson, Matthews, & Kelly, 2007; Duckworth & Seligman, 2005).

In order to gain traction in these related fields, and to facilitate communication across them, there needs to be a common language for describing the computational mechanisms that underlie the evaluation of tasks in terms of their demands for cognitive control. Throughout this chapter, we have provided a broad framework and particular implementations for such an analysis. Understanding cognitive control requires understanding the particular costs and benefits associated with this process. One way to do this is by employing the kinds of tasks reviewed in this chapter that can be used to measure and test the size of cost for cognitive control (DSTs; Kool & Botvinick, 2014; Kool et al., 2010; Westbrook et al., 2013).

Another way is to establish a comprehensive and empirical comparison of the different models that have been offered on how effort-based cost-benefit analyses are implemented (Holroyd & McClure, 2015; Kurzban et al., 2013; Shenhav et al., 2013; Silvetti, Alexander, Verguts, & Brown, 2014). With these and related opportunities at hand, we foresee a period of rapid progress towards understanding the connections between motivation and cognitive control.

References

Aarts, E., Roelofs, A., Franke, B., Rijpkema, M., Fernandez, G., Helmich, R. C., & Cools, R. (2010). Striatal dopamine mediates the interface between motivational and cognitive control in humans: Evidence from genetic imaging. *Neuropsychopharmacology*, *35*, 1943–1951.

Alexander, W. H., & Brown, J. W. (2011). Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience*, *14*, 1338–1344.

Allport, G. W. (1954). *The nature of prejudice*. New York, NY: Addison Wesley.

Amemori, K.-I., & Graybiel, A. M. (2012). Localized microstimulation of primate pregenual cingulate cortex induces negative decision-making. *Nature Neuroscience*, *15*, 776–785.

Amiez, C., Sallet, J., Procyk, E., & Petrides, M. (2012). Modulation of feedback related activity in the rostral anterior cingulate cortex during trial and error exploration. *NeuroImage*, *63*, 1078–1090.

Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.

Aron, A. R. (2011). From reactive to proactive and selective control: Developing a richer model for stopping inappropriate responses. *Biological Psychiatry*, *69*, e55–e68.

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to Stop signal response inhibition: Role of the subthalamic nucleus. *Journal of Neuroscience*, *26*, 2424–2433.

Arrington, C., & Logan, G. (2004). The cost of a voluntary task switch. *Psychological Science*, *15*, 610–615.

Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: Adaptive gain and optimal performance. *Annual Review of Neuroscience*, *28*, 403–450.

Bartra, O., McGuire, J. T., & Kable, J. W. (2013). The valuation system: A coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *NeuroImage*, *76*, 412–427.

Baumeister, R. F., Bratslavsky, E., Muraven, M., & Tice, D. M. (1998). Ego depletion: Is the active self a limited resource? *Journal of Personality and Social Psychology*, *74*, 1252–1265.

Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, *10*, 1214–1221.

Berlyne, D. E. (1957). Uncertainty and conflict: A point of contact between information-theory and behavior-theory concepts. *Psychological Review*, *64*, 329–339.

Bijleveld, E., Custers, R., & Aarts, H. (2009). The unconscious eye opener: Pupil dilation reveals strategic recruitment of resources upon presentation of subliminal reward cues. *Psychological Science*, *20*, 1313–1315.

Blain, B., Hollard, G., & Pessiglione, M. (2016). Neural mechanisms underlying the impact of daylong cognitive work on economic decisions. *Proceedings of the National Academy of Sciences of the United States of America*, *113*, 6967–6972.

Blanchard, T. C., & Hayden, B. Y. (2014). Neurons in dorsal anterior cingulate cortex signal postdecisional variables in a foraging task. *Journal of Neuroscience*, *34*, 646–655.

Bonner, S. E., & Sprinkle, G. B. (2002). The effects of monetary incentives on effort and task performance: Theories, evidence, and a framework for research. *Accounting, Organizations and Society*, *27*, 303–345.

Boorman, E. D., Rushworth, M. F., & Behrens, T. E. (2013). Ventromedial prefrontal and anterior cingulate cortex adopt choice and default reference frames during sequential multi-alternative choice. *Journal of Neuroscience*, *33*, 2242–2253.

Botvinick, M. M. (2007). Conflict monitoring and decision making: Reconciling two perspectives on anterior cingulate function. *Cognitive, Affective, & Behavioral Neuroscience*, *7*, 356–366.

Botvinick, M. M., & Braver, T. (2015). Motivation and cognitive control: From behavior to neural mechanism. *Annual Review of Psychology*, *66*, 83–113.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*, 624–652.

Botvinick, M. M., & Cohen, J. D. (2015). The computational and neural basis of cognitive control: Charted territory and new frontiers. *Cognitive Science*, *38*, 1249–1285.

Botvinick, M. M., Huffstetler, S., & McGuire, J. (2009). Effort discounting in human nucleus accumbens. *Cognitive, Affective, & Behavioral Neuroscience*, *9*, 16–27.

Bryden, D. W., Johnson, E. E., Tobia, S. C., Kashtelyan, V., & Roesch, M. R. (2011). Attention for learning signals in anterior cingulate cortex. *Journal of Neuroscience*, *31*, 18266–18274.

Cacioppo, J. T., & Petty, R. E. (1982). The need for cognition. *Journal of Personality and Social Psychology*, *42*, 116–131.

Cai, X., & Padoa-Schioppa, C. (2012). Neuronal encoding of subjective value in dorsal and ventral anterior cingulate cortex. *Journal of Neuroscience*, *32*, 3791–3808.

Camerer, C. F., & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *The Journal of Risk and Uncertainty*, *19*, 7–42.

Carter, E. C., Kofler, L. M., Forster, D. E., & McCullough, M. E. (2015). A series of meta-analytic tests of the depletion effect: Self-control does not seem to rely on a limited resource. *Journal of Experimental Psychology: General*, *144*(4), 796–815.

Cavanagh, J. F., & Frank, M. J. (2014). Frontal theta as a mechanism for cognitive control. *Trends in Cognitive Sciences*, *18*, 414–421.

Cavanagh, J. F., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., & Frank, M. J. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nature Neuroscience*, *14*, 1462–1467.

Charness, G., & Kuhn, P. (2010). What can labor economists learn from the lab? *IZA Discussion Paper no. 4941.*

Chen, M. K., Lakshminarayanan, V., & Santos, L. R. (2006). How basic are behavioral biases? Evidence from capuchin monkey trading behavior. *Journal of Political Economy*, *114*, 517–537.

Christoff, K., Prabhakaran, V., Dorfman, J., Zhao, Z., Kroger, J. K., Holyoak, K. J., & Gabrieli, J. D. E. (2001). Rostrolateral prefrontal cortex involvement in relational integration during reasoning. *NeuroImage*, *14*, 1136–1149.

Conover, K. L., & Shizgal, P. (2005). Employing labor-supply theory to measure the reward value of electrical brain stimulation. *Games and Economic Behavior*, *52*, 283–304.

Cowen, S. L., Davis, G. A., & Nitz, D. A. (2012). Anterior cingulate neurons in the rat map anticipated effort and reward to their associated action sequences. *Journal of Neurophysiology*, *107*, 2393–2407.

Craig, A. D. B. (2009). How do you feel—now? The anterior insula and human awareness. *Nature Reviews Neuroscience*, *10*, 59–70.

Danielmeier, C., Eichele, T., Forstmann, B. U., Tittgemeyer, M., & Ullsperger, M. (2011). Posterior medial frontal cortex activity predicts post-error adaptations in task-related visual and motor areas. *Journal of Neuroscience*, *31*, 1780–1789.

Danielmeier, C., & Ullsperger, M. (2011). Post-error adjustments. *Frontiers in Psychology*, *2*, 233.

Dayan, P. (2012). How to set the switches on this thing. *Current Opinion in Neurobiology*, *22*, 1068–1074.

De Jong, R. (2000). An intention-activation account of residual switch costs. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance* (pp. 357–376). Cambridge, MA: MIT Press.

Dickinson, D. L. (1999). An experimental examination of labor supply and work intensities. *Journal of Labor Economics*, *17*, 638–670.

Dixon, M. L., & Christoff, K. (2012). The decision to engage cognitive control is driven by expected reward-value: Neural and behavioral evidence. *PLoS ONE*, *7*, e51637.

Duckworth, A. L., Peterson, C., Matthews, M. D., & Kelly, D. R. (2007). Grit: Perseverance and passion for long-term goals. *Journal of Personality and Social Psychology*, *92*, 1087–1101.

Duckworth, A. L., Quinn, P. D., Lynam, D. R., Loeber, R., & Stouthamer-Loeber, M. (2011). Role of test motivation in intelligence testing. *Proceedings of the National Academy of Sciences of the United States of America*, *108*, 7716–7720.

Duckworth, A. L., & Seligman, M. E. P. (2005). Self-discipline outdoes IQ in predicting academic performance of adolescents. *Psychological Science*, *16*, 939–944.

Duncan, J. (2010). The multiple-demand (MD) system of the primate brain: Mental programs for intelligent behaviour. *Trends in Cognitive Sciences*, *14*, 172–179.

Durstewitz, D., Vittoz, N. M., Floresco, S. B., & Seamans, J. K. (2010). Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron*, *66*, 438–448.

Ebitz, R. B., & Platt, M. L. (2015). Neuronal activity in primate dorsal anterior cingulate cortex signals task conflict and predicts adjustments in pupil-linked arousal. *Neuron*, *85*, 628–640.

Eisenberger, R. (1992). Learned industriousness. *Psychological Review*, *99*, 248–267.

Emeric, E. E., Brown, J. W., Leslie, M., Pouget, P., Stuphorn, V., & Schall, J. D. (2008). Performance monitoring local field potentials in the medial frontal cortex of primates: Anterior cingulate cortex. *Journal of Neurophysiology*, *99*, 759–772.

Engelmann, J. B., Damaraju, E., Padmala, S., & Pessoa, L. (2009). Combined effects of attention and motivation on visual task performance: Transient and sustained motivational effects. *Frontiers in Human Neuroscience*, *3*, 4.

Engelmann, J. B., & Pessoa, L. (2007). Motivation sharpens exogenous spatial attention. *Emotion*, *7*, 668–674.

Etzel, J. A., Cole, M. W., Zacks, J. M., Kay, K. N., & Braver, T. S. (in press). Reward motivation enhances task coding in frontoparietal cortex. *Cerebral Cortex*.

Fehr, E., & Goette, L. (2007). Do workers work more if wages are high? Evidence from a randomized field experiment. *American Economic Review*, *97*, 298–317.

Feng, S. F., Schwemmer, M., Gershman, S. J., & Cohen, J. D. (2014). Multitasking versus multiplexing: Toward a normative account of limitations in the simultaneous execution of control-demanding behaviors. *Cognitive, Affective, & Behavioral Neuroscience*, *14*, 129–146.

Forster, S. E., Carter, C. S., Cohen, J. D., & Cho, R. Y. (2011). Parametric manipulation of the conflict signal and control-state adaptation. *Journal of Cognitive Neuroscience*, *23*, 923–935.

Forstmann, B. U., Brass, M., Koch, I., & Von Cramon, D. (2006). Voluntary selection of task sets revealed by functional magnetic resonance imaging. *Journal of Cognitive Neuroscience*, *18*, 388–398.

Fritz, J., & Dreisbach, G. (2013). Conflicts as aversive signals: Conflict priming increases negative judgments for neutral stimuli. *Cognitive Affective, & Behavioral Neuroscience*, *13*, 311–317.

Fudenberg, D., & Levine, D. K. (2006). A dual-self model of impulse control. *American Economic Review*, *96*, 1449–1476.

Fudenberg, D., & Levine, D. K. (2011). Risk, delay, and convex self-control costs. *American Economic Journal: Microeconomics*, *3*, 34–68.

Gläscher, J., Adolphs, R., Damasio, H., Bechara, A., Rudrauf, D., Calamia, M.,…Tranel, D. (2012). Lesion mapping of cognitive control and value-based decision making in the prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *109*, 14681–14686.

Goldman-Rakic, P. S. (1987). Circuitry of primate prefrontal cortex and regulation of behavior by representational memory. In F. Plum (Ed.), *Handbook of physiology* (pp. 373–417). Bethesda, MD: American Physiological Society.

Gratton, G., Coles, M. G., & Donchin, E. (1992). Optimizing the use of information: Strategic control of activation of responses. *Journal of Experimental Psychology: General*, *121*, 480–506.

Grundy, J. G., & Shedden, J. M. (2014). Support for a history-dependent predictive model of dACC activity in producing the bivalency effect: An event-related potential study. *Neuropsychologia*, *57*, 166–178.

Haber, S. N., & Knutson, B. (2010). The reward circuit: Linking primate anatomy and human imaging. *Neuropsychopharmacology*, *35*, 4–26.

Hagger, M. S., & Chatzisarantis, N. L. (2016). *Perspectives on Psychological Science*, *11*(4), 546–573.

Hagger, M. S., Wood, C., Stiff, C., & Chatzisarantis, N. L. D. (2010). Ego depletion and the strength model of self-control: A meta-analysis. *Psychological Bulletin*, *136*, 495–525.

Hare, T. A., Camerer, C. F., & Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. *Science*, *324*, 646–648.

Hare, T. A., Schultz, W., Camerer, C. F., O'Doherty, J. P., & Rangel, A. (2011). Transformation of stimulus value signals into motor commands during simple choice. *Proceedings of the National Academy of Sciences of the United States of America*, *108*, 18120–18125.

Hayden, B. Y., & Platt, M. L. (2010). Neurons in anterior cingulate cortex multiplex information about reward and action. *The Journal of Neuroscience*, *30*, 3339–3346.

Heatherton, T. F. (2011). Neuroscience of self and self-regulation. *Annual Review of Psychology*, *62*, 363–390.

Hillman, K. L., & Bilkey, D. K. (2010). Neurons in the rat anterior cingulate cortex dynamically encode cost-benefit in a spatial decision-making task. *Journal of Neuroscience*, *30*, 7705–7713.

Hockey, G. R. J. (1997). Compensatory control in the regulation of human performance under stress and high workload: A cognitive-energetical framework. *Biological Psychology*, *45*, 73–93.

Holroyd, C. B. (in press). The waste disposal problem of effortful control. In T. Braver (Ed.), *Motivation and cognitive control.* New York, NY: Psychology Press.

Holroyd, C. B., & Coles, M. G. H. (2002). The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity. *Psychological Review*, *109*, 679–709.

Holroyd, C. B., & McClure, S. M. (2015). Hierarchical control over effortful behavior by rodent medial frontal cortex: A computational model. *Psychological Review*, *122*, 54–83.

Holroyd, C. B., & Yeung, N. (2012). Motivation of extended behaviors by anterior cingulate cortex. *Trends in Cognitive Sciences*, *16*, 121–127.

Horga, G., Maia, T. V., Wang, P., Wang, Z., Marsh, R., & Peterson, B. S. (2011). Adaptation to conflict via context-driven anticipatory signals in the dorsomedial prefrontal cortex. *The Journal of Neuroscience*, *31*, 16208–16216.

Hosokawa, T., Kennerley, S. W., Sloan, J., & Wallis, J. D. (2013). Single-neuron mechanisms underlying cost-benefit analysis in frontal cortex. *Journal of Neuroscience*, *33*, 17385–17397.

Hull, C. L. (1943). *Principles of behavior.* New York: Appleton Century.

Inzlicht, M., Bartholow, B. D., & Hirsh, J. B. (2015). Emotional foundations of cognitive control. *Trends in Cognitive Sciences*, *19*, 126–132.

Inzlicht, M., & Schmeichel, B. (2012). What is ego depletion? Toward a mechanistic revision of the resource model of self-control. *Perspectives on Psychological Science*, *7*, 450–463.

Isomura, Y., Ito, Y., Akazawa, T., Nambu, A., & Takada, M. (2003). Neural coding of 'attention for action' and 'response selection' in primate anterior cingulate cortex. *Journal of Neuroscience*, *23*, 8002–8012.

Ito, S., Stuphorn, V., Brown, J. W., & Schall, J. D. (2003). Performance monitoring by the anterior cingulate cortex during saccade countermanding. *Science*, *302*, 120–122.

Jahfari, S., Waldorp, L., Van Den Wildenberg, W. P. M., Scholte, H. S., Ridderinkhof, K. R., & Forstmann, B. U. (2011). Effective connectivity reveals important roles for both the hyperdirect (fronto-subthalamic) and the indirect (fronto-striatal-pallidal) fronto-basal ganglia pathways during response inhibition. *Journal of Neuroscience*, *31*, 6891–6899.

Job, V., Dweck, C., & Walton, G. (2010). Ego depletion—Is it all in your head? *Psychological Science*, *21*, 1686–1693.

Johnston, K., Levin, H., Koval, M., & Everling, S. (2007). Top-down control-signal dynamics in anterior cingulate and prefrontal cortex neurons following task switching. *Neuron*, *53*, 453–462.

Kagel, J. H., Battalio, R. C., & Green, L. (1995). *Economic choice theory: An experimental analysis of animal behavior.* Cambridge: Cambridge University Press.

Kahneman, D. (1973). *Attention and effort.* Englewood Cliffs, NJ: Prentice-Hall.

Kaping, D., Vinck, M., Hutchison, R. M., Everling, S., & Womelsdorf, T. (2011). Specific contributions of ventromedial, anterior cingulate, and lateral prefrontal cortex for attentional selection and stimulus valuation. *PLoS Biology*, *9*, e1001224.

Kennerley, S. W., Walton, M. E., Behrens, T. E. J., Buckley, M. J., & Rushworth, M. F. S. (2006). Optimal decision making and the anterior cingulate cortex. *Nature Neuroscience*, *9*, 940–947.

Kerns, J. G., Cohen, J. D., MacDonald, A. W., Cho, R. Y., Stenger, V. A., & Carter, C. S. (2004). Anterior cingulate conflict monitoring and adjustments in control. *Science*, *303*, 1023–1026.

Khamassi, M., Wilson, C. R. E., Rothé, M., Quilodran, R., Dominey, P. F., & Procyk, E. (2010). Meta-learning, cognitive control, and physiological interactions between medial and lateral prefrontal cortex. In R. Mars, J. Sallet, M. F. S. Rushworth & N. Yeung (Eds.), *Neural bases of motivational and cognitive control* (pp. 351–370). Cambridge, MA: MIT Press.

King, J. A., Korb, F. M., von Cramon, D. Y., & Ullsperger, M. (2010). Post-error behavioral adjustments are facilitated by activation and suppression of task-relevant and task-irrelevant information processing. *Journal of Neuroscience*, *30*, 12759–12769.

Kirk, D. E. (2012). *Optimal control theory: An introduction*. Englewood Cliffs, NJ: Prentice-Hall.

Klein-Flügge, M. C., Kennerley, S. W., Saraiva, A. C., Penny, W. D., & Bestmann, S. (2015). Behavioral modeling of human choices reveals dissociable effects of physical effort and temporal delay on reward devaluation. *PLoS Computational Biology*, *11*, e1004116.

Kool, W., & Botvinick, M. M. (2014). A labor/leisure trade-off in cognitive control. *Journal of Experimental Psychology: General*, *143*, 131–141.

Kool, W., McGuire, J. T., Rosen, Z. B., & Botvinick, M. M. (2010). Decision making and the avoidance of cognitive demand. *Journal of Experimental Psychology: General*, *139*, 665–682.

Kool, W., McGuire, J. T., Wang, G. J., & Botvinick, M. M. (2013). Neural and behavioral evidence for an intrinsic cost of self-control. *PLoS ONE*, *8*, e72626.

Kouneiher, F., Charron, S., & Koechlin, E. (2009). Motivation and cognitive control in the human prefrontal cortex. *Nature Neuroscience*, *12*, 939–945.

Kurzban, R., Duckworth, A. L., Kable, J. W., & Myers, J. (2013). An opportunity cost model of subjective effort and task performance. *Behavioral and Brain Sciences*, *36*, 661–726.

Laming, D. (1968). *Information theory of choice-reaction times*. New York: Academic Press.

Litt, A., Plassmann, H., Shiv, B., & Rangel, A. (2011). Dissociating valuation and saliency signals during decision-making. *Cerebral Cortex*, *21*, 95–102.

Magno, E., Foxe, J. J., Molholm, S., Robertson, I. H., & Garavan, H. (2006). The anterior cingulate and error avoidance. *Journal of Neuroscience*, *26*, 4769–4773.

Mansouri, F. A., Tanaka, K., & Buckley, M. J. (2009). Conflict-induced behavioural adjustment: A clue to the executive functions of the prefrontal cortex. *Nature Reviews Neuroscience*, *10*, 141–152.

Matsuzaka, Y., Akiyama, T., Tanji, J., & Mushiake, H. (2012). Neuronal activity in the primate dorsomedial prefrontal cortex contributes to strategic selection of response tactics. *Proceedings of the National Academy of Sciences of the United States of America*, *109*, 4633–4638.

McClure, S. M., Laibson, D., Loewenstein, G., & Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science*, *306*, 503–507.

McGuire, J. T., & Botvinick, M. M. (2010). Prefrontal cortex, cognitive control, and the registration of decision costs. *Proceedings of the National Academy of Sciences of the United States of America*, *107*, 7922–7926.

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Reviews in Neuroscience*, *24*, 167–202.

Mischel, W., Shoda, Y., & Rodriguez, M. L. (1989). Delay of gratification in children. *Science*, *244*, 933–938.

Momennejad, I., & Haynes, J.-D. (2013). Encoding of prospective tasks in the human prefrontal cortex under varying task loads. *Journal of Neuroscience*, *33*, 17342–17349.

Morecraft, R. J., & Tanji, J. (2009). Cingulofrontal interactions and the cingulate motor areas. In B. A. Vogt (Ed.), *Cingulate neurobiology and disease* (pp. 113–144). Oxford: Oxford University Press.

Muraven, M., Shmueli, D., & Burkley, E. (2006). Conserving self-control strength. *Journal of Personality and Social Psychology*, *91*, 524–537.

Muraven, M., & Slessareva, E. (2003). Mechanisms of self-control failure: Motivation and limited resources. *Personality and Social Psychology Bulletin*, *29*, 894–906.

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

Narayanan, N. S., & Laubach, M. (2008). Neuronal correlates of post-error slowing in the rat dorsomedial prefrontal cortex. *Journal of Neurophysiology*, *100*, 520–525.

Navon, D., & Gopher, D. (1979). On the economy of the human-processing system. *Psychological Review*, *86*, 214–255.

Newman, L. A., Creer, D. J., & McGaughy, J. A. (2015). Cognitive control and the anterior cingulate cortex: How conflicting stimuli affect attentional control in the rat. *Journal of Physiology-Paris*, *109*, 95–103.

Nicholson, W., & Snyder, C. M. (2008). *Microeconomic theory: Basic principles and extensions*. Boston, MA: Cengage Learning.

Niv, Y., Daw, N., & Dayan, P. (2006). How fast to work: Response vigor, motivation and tonic dopamine. *Advances in Neural Information Processing Systems*, *18*, 1019.

O'Reilly, J. X., Schüffelgen, U., Cuell, S. F., Behrens, T. E. J., Mars, R. B., & Rushworth, M. F. S. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *110*, E3660–3669.

O'Reilly, R. C. (2010). The What and How of prefrontal cortical organization. *Trends In Neurosciences*, *33*, 355–361.

Oehrn, C. R., Hanslmayr, S., Fell, J., Deuker, L., Kremers, N. A., Do Lam, A. T.,…Axmacher, N. (2014). Neural communication patterns underlying conflict detection, resolution, and adaptation. *Journal of Neuroscience*, *34*, 10438–10452.

Orr, J. M., & Weissman, D. H. (2011). Succumbing to bottom-up biases on task choice predicts increased switch costs in the voluntary task switching paradigm. *Frontiers in Psychology*, *2*, 31.

Padmala, S., & Pessoa, L. (2011). Reward reduces conflict by enhancing attentional control and biasing visual cortical processing. *Journal of Cognitive Neuroscience*, *23*, 3419–3432.

Parvizi, J., Rangarajan, V., Shirer, W. R., Desai, N., & Greicius, M. D. (2013). The will to Persevere induced by electrical stimulation of the human cingulate gyrus. *Neuron*, *80*, 1359–1367.

Paus, T., Koski, L., Caramanos, Z., & Westbury, C. (1998). Regional differences in the effects of task difficulty and motor output on blood flow response in the human anterior cingulate cortex: A review of 107 PET activation studies. *Neuroreport*, *9*, R37–R47.

Pearce, J., & Hall, H. (1980). A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, *87*, 532–552.

Pessoa, L. (2015). Multiple influences of reward on perception and attention. *Visual Cognition*, *23*, 272–290.

Prévost, C., Pessiglione, M., Météreau, E., Cléry-Melin, M., & Dreher, J. (2010). Separate valuation subsystems for delay and effort decision costs. *Journal of Neuroscience*, *30*, 14080–14090.

Quilodran, R., Rothé, M., & Procyk, E. (2008). Behavioral shifts and action valuation in the anterior cingulate cortex. *Neuron*, *57*, 314–325.

Rabbitt, P. (1966). Errors and error correction in choice-response tasks. *Journal of Experimental Psychology*, *71*, 264–272.

Reinhart, R. M. G., & Woodman, G. F. (2014). Causal control of medial-frontal cortex governs electrophysiological and behavioral indices of performance monitoring and learning. *Journal of Neuroscience*, *34*, 4214–4227.

Reinhart, R. M. G., & Woodman, G. F. (2015). Enhancing long-term memory with stimulation tunes visual attention in one trial. *Proceedings of the National Academy of Sciences of the United States of America*, *112*, 625–630.

Rollnik, J. D., Schröder, C., Rodríguez-Fornells, A., Kurzbuch, A. R., Däuper, J., Möller, J., & Münte, T. F. (2004). Functional lesions and human action monitoring: Combining repetitive transcranial magnetic stimulation and event-related brain potentials. *Clinical Neurophysiology*, *115*, 145–153.

Rothé, M., Quilodran, R., Sallet, J., & Procyk, E. (2011). Coordination of high gamma activity in anterior cingulate and lateral prefrontal cortical areas during adaptation. *Journal of Neuroscience*, *31*, 11110–11117.

Salamone, J. D., Correa, M., Farrar, A. M., & Mingote, S. M. (2007). Effort-related functions of nucleus accumbens dopamine and associated forebrain circuits. *Psychopharmacology*, *191*, 461–482.

Schouppe, N., Ridderinkhof, K. R., Verguts, T., & Notebaert, W. (2014). Context-specific control and context selection in conflict tasks. *Acta Psychologica*, *146*, 63–66.

Shackman, A. J., Salomons, T. V., Slagter, H. A., Fox, A. S., Winter, J. J., & Davidson, R. J. (2011). The integration of negative affect, pain and cognitive control in the cingulate cortex. *Nature Reviews Neuroscience*, *12*, 154–167.

Shah, A., & Oppenheimer, D. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, *134*, 207.

Shen, C., Ardid, S., Kaping, D., Westendorff, S., Everling, S., & Womelsdorf, T. (2015). Anterior cingulate cortex cells identify process-specific errors of attentional control prior to transient prefrontal-cingulate inhibition. *Cerebral Cortex*, *25*(8), 2213–2228.

Shenhav, A., & Botvinick, M. (2015). Uncovering a missing link in anterior cingulate research. *Neuron*, *85*, 455–457.

Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, *79*, 217–240.

Shenhav, A., & Buckner, R. L. (2014). Neural correlates of dueling affective reactions to win-win choices. *Proceedings of the National Academy of Sciences of the United States of America*, *111*(30), 10978–10983.

Sheth, S. A., Mian, M. K., Patel, S. R., Asaad, W. F., Williams, Z. M., Dougherty, D. D.,…Eskandar, E. N. (2012). Human dorsal anterior cingulate cortex neurons mediate ongoing behavioural adaptation. *Nature*, *488*, 218–221.

Silvetti, M., Alexander, W., Verguts, T., & Brown, J. (2014). From conflict management to reward-based decision making: Actors and critics in primate medial frontal cortex. *Neuroscience and Biobehavioral Reviews*, *46*(Pt. 1), 44–57.

Simen, P., Contreras, D., Cara, B., Hu, P., Holmes, P., & Cohen, J. D. (2009). Reward rate optimization in two-alternative decision making: Empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance*, *35*, 1865–1897.

Simon, D. A., & Daw, N. D. (2011). Neural correlates of forward planning in a spatial decision task in humans. *Journal of Neuroscience*, *31*, 5526–5539.

Simon, H. A. (1956). A comparison of game theory and learning theory. *Psychometrika*, *21*, 267–272.

Smith, V. L., & Walker, J. M. (1993). Monetary rewards and decision cost in experimental economics. *Economic Inquiry*, *31*, 245–261.

Sohn, M.-H., Albert, M. V., Jung, K., Carter, C. S., & Anderson, J. R. (2007). Anticipation of conflict monitoring in the anterior cingulate cortex and the prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *104*, 10330–10334.

Solomon, R. L. (1948). The influence of work on behavior. *Psychological Bulletin*, *45*, 1–40.

Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.

Stuss, D. T. (2011). Functions of the frontal lobes: Relation to executive functions. *Journal of the International Neuropsychological Society*, *17*, 759–765.

Stuss, D. T., & Alexander, M. P. (2007). Is there a dysexecutive syndrome? *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *362*, 901–915.

Thompson, M. E. (1944). Learning as a function of the absolute and relative amounts of work. *Journal of Experimental Psychology*, *34*, 506–515.

Ullsperger, M., Danielmeier, C., & Jocham, G. (2014). Neurophysiology of performance monitoring and adaptive behavior. *Physiological Reviews*, *94*, 35–79.

Vassena, E., Silvetti, M., Boehler, C. N., Achten, E., Fias, W., & Verguts, T. (2014). Overlapping neural systems represent cognitive effort and reward anticipation. *PLoS ONE*, *9*, e91008.

Venkatraman, V., & Huettel, S. A. (2012). Strategic control in decision-making under uncertainty. *European Journal of Neuroscience*, *35*, 1075–1082.

Wallis, J. D., & Kennerley, S. W. (2011). Contrasting reward signals in the orbitofrontal cortex and anterior cingulate cortex. *Annals of the New York Academy of Sciences*, *1239*, 33–42.

Walton, M. E., Kennerley, S. W., Bannerman, D. M., Phillips, P. E. M., & Rushworth, M. F. S. (2006). Weighing up the benefits of work: Behavioral and neural analyses of effort-related decision making. *Neural Networks*, *19*, 1302–1314.

Wessel, J. R., Danielmeier, C., Morton, J. B., & Ullsperger, M. (2012). Surprise and error: Common neuronal architecture for the processing of errors and novelty. *Journal of Neuroscience*, *32*, 7528–7537.

Westbrook, A., Kester, D., & Braver, T. S. (2013). What is the subjective cost of cognitive effort? Load, trait, and aging effects revealed by economic preference. *PLoS ONE*, *22*, e68210.

Wisniewski, D., Reverberi, C., Tusche, A., & Haynes, J.-D. (2015). The neural representation of voluntary task-set selection in dynamic environments. *Cerebral Cortex*, *25*(12), 4715–4726.

Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision making. *Current Opinion in Neurobiology*, *22*, 996–1003.

Womelsdorf, T., Johnston, K., Vinck, M., & Everling, S. (2010). Theta-activity in anterior cingulate cortex predicts task rules and their adjustments following errors. *Proceedings of the National Academy of Sciences of the United States of America*, *107*, 5248–5253.

Yeung, N., Botvinick, M., & Cohen, J. (2004). The neural basis of error detection: Conflict monitoring and the error-related negativity. *Psychological Review*, *111*, 931–959.