



Contents lists available at ScienceDirect

Neuroscience and Biobehavioral Reviews

journal homepage: www.elsevier.com/locate/neubiorev



Review

Bayesian modeling of flexible cognitive control

Jiefeng Jiang^{a,b}, Katherine Heller^{a,c}, Tobias Egner^{a,b,*}

^a Center for Cognitive Neuroscience, Duke University, United States

^b Department of Psychology & Neuroscience, Duke University, United States

^c Department of Statistical Science, Duke University, United States

ARTICLE INFO

Article history:

Received 28 June 2013

Received in revised form 31 March 2014

Accepted 3 June 2014

Keywords:

Cognitive control

Bayesian models

Conflict

Congruency sequence effect

Conflict adaptation

Proportion congruency effect

ABSTRACT

“Cognitive control” describes endogenous guidance of behavior in situations where routine stimulus-response associations are suboptimal for achieving a desired goal. The computational and neural mechanisms underlying this capacity remain poorly understood. We examine recent advances stemming from the application of a Bayesian learner perspective that provides optimal prediction for control processes. In reviewing the application of Bayesian models to cognitive control, we note that an important limitation in current models is a lack of a plausible mechanism for the flexible adjustment of control over conflict levels changing at varying temporal scales. We then show that flexible cognitive control can be achieved by a Bayesian model with a volatility-driven learning mechanism that modulates dynamically the relative dependence on recent and remote experiences in its prediction of future control demand. We conclude that the emergent Bayesian perspective on computational mechanisms of cognitive control holds considerable promise, especially if future studies can identify neural substrates of the variables encoded by these models, and determine the nature (Bayesian or otherwise) of their neural implementation.

© 2014 Elsevier Ltd. All rights reserved.

Contents

1. Cognitive control as statistical prediction.....	00
1.1. Cognitive control as ‘guided’ information processing.....	00
1.2. The conflict monitoring model.....	00
1.3. The dual mechanisms of control model.....	00
1.4. Cognitive control as statistical inference.....	00
2. Modeling cognitive control using Bayesian models.....	00
2.1. Overview of Bayesian methods.....	00
2.2. Bayesian modeling of speed-accuracy trade-off.....	00
2.3. Bayesian modeling of the Eriksen flanker effect.....	00
2.4. Bayesian modeling of response inhibition.....	00
2.5. Toward modeling the flexibility of cognitive control.....	00
3. Modeling the flexibility of cognitive control using Bayesian methods.....	00
3.1. A Bayesian model of flexible conflict-control.....	00
3.2. Simulating the short-term and long-term trial history effects of conflict-control.....	00
3.2.1. Subjects.....	00
3.2.2. Stimuli and procedure.....	00
3.2.3. Experimental design.....	00
3.2.4. Data analysis.....	00
3.2.5. Results and discussion.....	00

* Corresponding author at: Center for Cognitive Neuroscience, Duke University, LSRC Box 90999, Durham, NC 27708, United States. Tel.: +1 919 684 1049.
E-mail address: tobias.egner@duke.edu (T. Egner).

3.3.	Simulating the flexibility of conflict-control	00
3.3.1.	Subjects	00
3.3.2.	Stimuli and procedure	00
3.3.3.	Experimental design	00
3.3.4.	Data analysis	00
3.3.5.	Results and discussion	00
4.	Future directions and concluding remarks	00
4.1.	How “Bayesian” is cognitive control?	00
5.	Conclusion	00
	Acknowledgments	00
	Appendix A	00
	References	00

1. Cognitive control as statistical prediction

“Cognitive control” describes the ability to guide one’s behavior in line with internal goals. A key characteristic of cognitive control is thought to be flexibility: control processes must be capable of dynamically adapting (both qualitatively and quantitatively) to ongoing changes in the environment. How this type of contextual regulation of control occurs (in the absence of an all-knowing *homunculus*) is a key question in current cognitive psychology and neuroscience research. In the present paper, we review recent attempts of modeling the “control of control”, with a particular focus on the increasingly popular idea that prediction using Bayesian algorithms, which behave similar to reinforcement learning algorithms with varying learning rates (see below), may furnish a potent means for flexibly adapting control settings to contextual changes in the environment. In Section 1, we review two influential (non-Bayesian) models of cognitive control and highlight some limitations in their ability to adapt control to changing circumstances, specifically with respect to integrating contextual information across different time scales. We then suggest that Bayesian methods can achieve time-varying, self-adaptive integration of control-relevant contextual information. In Section 2, we review recent efforts to use Bayesian models to simulate various aspects of cognitive control. In Section 3, we outline a novel Bayesian model of conflict-control and demonstrate how it can account for various key behavioral phenomena. In Section 4, possible directions for future research regarding the application of Bayesian models to cognitive control are discussed.

1.1. Cognitive control as ‘guided’ information processing

In interacting with our environment, we transform sensory input into internal representations and select cognitive or motor actions based on these representations and our current goals. Given the fact that there is an enormous amount of sensory information and many possible actions available in contrast to only a few desired responses, appropriate action selection is a difficult task. To simplify this task, stimuli and actions that are frequently paired become mnemonically associated (e.g., via Hebbian learning) into stimulus–response (S-R) ensembles (or pathways) or more complex and extended action schemas (Norman and Shallice, 1986) that facilitate prompt reaction. Because much sensory information is processed in different pathways in parallel but only few actions can (or should) be taken simultaneously, stimulus representations and S-R pathways are believed to compete for being selected to drive behavior (Desimone and Duncan, 1995; Miller and Cohen, 2001; Norman and Shallice, 1986). The results of this competition are largely driven by the strength of associative pathways: stronger (i.e., more frequently activated) pathways are more likely to win the competition than weaker or novel ones. Once selected (and executed), the strength of a particular pathway may be reinforced or reduced depending on the assessment of how well the selected

actions have fulfilled the organism’s intended goals (Balleine and Dickinson, 1998).

This competition mechanism (or “contention scheduling”, see Norman and Shallice, 1986) can generate appropriate behavior in many situations, but strong, stereotyped pathways can also result in suboptimal and even hazardous actions in some situations. For example, a US citizen’s habitual driving on the right side of the road may have serious consequences when performed in the UK. In this case, a set of weaker or even novel associations (e.g., driving on the left side of the road) must be biased to win the competition in order to achieve the organism’s goals. This “top-down” biasing of information processing to favor goal-directed stimuli and actions is the essence of cognitive control (e.g., Norman and Shallice, 1986; Botvinick et al., 2001; Miller and Cohen, 2001). In present-day neuroanatomical models, cognitive control is closely tied to the prefrontal cortex (PFC), which is proposed to harbor temporary representations of current goals, goal-relevant stimuli and strategies (Badre, 2008; Botvinick et al., 2001; Braver and Barch, 2002; Duncan, 2001; Fuster, 2008; Koechlin et al., 2003; Miller and Cohen, 2001; Norman and Shallice, 1986). To implement control, representations of goals, context and related methods (like rules) are thought to be actively maintained in the PFC, which sends biasing signals to posterior brain regions to guide the information flowing through the desired pathways and reach the selection of appropriate actions (e.g., Miller and Cohen, 2001).

In the laboratory, cognitive control is traditionally tested in interference (or “conflict”) tasks such as the Stroop task (MacLeod, 1991), which entail conditions that require subjects to overcome a stronger habitual response in favor of a weaker (but correct) response. Consider, for instance, a variant of the Stroop task we employ in the empirical section of this paper (Section 3). This task requires a subject to respond to the gender of a face image, while ignoring a word label (either “male” or “female”) that is overlaid on the image and which can be either congruent (e.g., “male” overlaid on a male face) or incongruent (e.g., “female” overlaid on a male face) with the face image (Egner et al., 2008). In order to arrive at the correct response during an incongruent trial, the subject has to overcome the highly automatic processing of the word-meaning in favor of categorizing the face’s gender. Correct response selection on incongruent trials therefore requires the application of cognitive control in the PFC, strengthening the information flowing through the task-relevant processing pathway to win out over the task-irrelevant (though more habitual) one (Botvinick et al., 2001; Braver and Barch, 2002; Cohen et al., 1990). Accordingly, many neuroimaging studies of these types of tasks have documented higher activation in the PFC associated with higher conflict and control levels (Barch et al., 2001; Botvinick et al., 2004; Ridderinkhof et al., 2004), and modulated activity in brain regions related to processing task relevant- and irrelevant stimuli (Egner and Hirsch, 2005; King et al., 2010; Liu et al., 2004; Wittfoth et al., 2006).

One crucial question regarding this account, however, is how cognitive control itself is controlled. For example, when does

cognitive control engage to bias competition of pathways? How does it change strength when more or less control is needed? And how is control withdrawn? In this review, we argue (as have others before us, see Botvinick et al., 2001) that the regulation of cognitive control relies on the prediction of processing demands (e.g., anticipated conflict or congruency levels), which is derived from previous experience. In the following, we first review two influential models: the conflict monitoring model (Botvinick et al., 2001) and the dual mechanisms of control model (Braver, 2012). Both models adjust the level of cognitive control based on previous experience. Yet, as we describe below in detail, neither model can explain how the brain flexibly incorporates and combines information across different time scales (short-term and long-term) to predict conflict. We argue that this flexibility can be modeled using a Bayesian approach. In Section 2, we review previous work using Bayesian models to account for various aspects of cognitive control. In Section 3, we outline a new Bayesian model of conflict-control and demonstrate how it can account for various key behavioral phenomena of cognitive control. In Section 4, directions for future research regarding the application of Bayesian models to cognitive control are discussed.

1.2. The conflict monitoring model

The conflict monitoring model (Botvinick et al., 2001) treats the intervention of cognitive control as a reactive processing adjustment following the detection of conflict. This adjustment is achieved by incorporation of two systems: a conflict monitoring system that estimates the levels of conflict and sends signals to a control system, which in turn delivers biasing signals to information processing pathways. It is not entirely clear in the model whether control is originally recruited for dealing with conflict in the ongoing trial or for subsequent trials only (for discussion, see Egner et al., 2010), but the effects of conflict-driven control that are seen to support the model are typically measured by observing performance on the subsequent trial(s).

The specific mechanisms of the conflict monitoring system are made explicit in a neural network implementation (Botvinick et al., 2001), in which RT was simulated as the time-point when the Hopfield energy (Hopfield, 1982) of one output node (out of two or more) reached a pre-defined threshold. This neural network implementation successfully simulated various landmark behavioral effects found in interference tasks. For example, the congruency sequence (or *conflict adaptation*) effect – a smaller interference effect (measured by subtracting mean RT of congruent trials from mean RT of incongruent or neutral trials) following an incongruent trial than after a congruent trial (Gratton et al., 1992), and the *proportion congruency* effect, which describes the pattern that the larger the proportion of congruent trials is in a block, the higher the average interference effect is in that block (Logan and Zbrodoff, 1979; Tzelgov et al., 1992), have both been simulated successfully by the conflict-monitoring model using a reinforcement learning algorithm that updates the prediction of congruency by incorporating (in)congruency at the current trial via a fixed learning rate α . Specifically, the prediction for the forthcoming trial is a linear combination of the (in)congruency at the current trial and the prediction concerning the current trial, with the rates of α and $(1 - \alpha)$, respectively. The model further proposes that the conflict monitoring system is housed in the anterior cingulate cortex (ACC) and the control system in the lateral PFC. These propositions have been supported by neuroimaging findings showing elevated activation in the ACC under conditions where conflict is high and control is assumed to be low (Barch et al., 2001; Botvinick et al., 1999; Carter et al., 1998; Kerns et al., 2004; MacDonald et al., 2000; MacLeod and MacDonald, 2000) and enhanced activation in lateral PFC under conditions where conflict is low and control is assumed to be high

(Egner and Hirsch, 2005; Kerns et al., 2004; MacDonald et al., 2000), as well as increased functional connectivity between the lateral PFC and regions supporting task-relevant stimulus information in the posterior brain (Egner and Hirsch, 2005).

Although the conflict monitoring model is able to simulate the phenomena of conflict adaptation and proportion congruency effects (Botvinick et al., 2001, simulation 2A and 2B) separately, a closer look at the simulation results suggests the model is not able to replicate these two effects using the same set of parameters. Specifically, in the simulation of conflict adaptation (simulation 2A), the best model has a learning rate of 0.5; while the learning rate is dramatically reduced to 0.05 when simulating the decreasing interference effect as the proportion of incongruent trials increases in the simulation of proportion congruency effects (simulation 2B). The 0.5 learning rate in simulating conflict adaptation effects essentially represents a *phasic* or *transient* mechanism relying more on recent experience, while the 0.05 learning rate reflects a more *tonic* or *sustained* mechanism incorporating temporally more remote or extended information that allows for the proportion of incongruent trials to be learnt. The fact that the conflict-monitoring model cannot simulate both of these effects simultaneously is problematic, given that they are supposed to reflect the same basic phenomenon (conflict-driven control) and that conflict adaptation and proportion congruency effects do in fact co-occur in a single task-setting (e.g., Torres-Quesada et al., 2013), a finding which the conflict-monitoring model is clearly unable to capture.

1.3. The dual mechanisms of control model

The more recent dual-mechanisms of control model may have the potential to overcome this problem, as it specifically accommodates control effects that operate over different times scales, by incorporating both a “reactive” and a “proactive” control mechanism (Braver, 2012; Braver et al., 2007; De Pisapia and Braver, 2006). The key difference between these two mechanisms lies in their time scales and their relation to stimulus onsets. Specifically, the reactive mechanism accounts for transient changes of cognitive control after a stimulus has been encountered (e.g., following conflict), whereas the proactive mechanism monitors long-term changes of conflict density and applies changes to cognitive control before the onsets of incoming stimuli. Although operating on different time scales, these two mechanisms cooperate to modulate cognitive control. To test the feasibility of this model (De Pisapia and Braver, 2006) conducted a color-naming Stroop fMRI study, which included three types of blocks with varying proportions of incongruent trials. The authors found that in ACC and left DLPFC the conflict-related (i.e., incongruent–congruent, at trial level) activity was highest when most trials were congruent (and proactive control presumably low), suggesting a reactive, short-term/phasic type of control being applied; whereas in the right DLPFC, the sustained, block-wise activation was the highest when most trials were incongruent, suggesting the wielding of a proactive, long-term/tonic type of control. The authors furthermore found that a model in which both ACC and the DLPFC units had a reactive and a proactive component could simulate both the phasic and tonic activation patterns found in the fMRI data. This dual mechanisms of control model represents a novel approach to understanding cognitive control, but there is presently little empirical evidence to support the idea of two conflict monitoring units working on different time scales in the ACC. It is also unclear whether this model can simulate both long-term (e.g., proportion congruency) and short-term (e.g., conflict adaptation) regulation of control simultaneously, and it would be more parsimonious if both types of control were integrated into a single mechanism. In the following, we aim to sketch out how such integration can be achieved.

1.4. Cognitive control as statistical inference

In the computer simulations of both conflict-monitoring and dual-mechanism models, short-term information (e.g., congruency at the current trial) and long-term information (e.g., congruency at earlier trials) were integrated using a *fixed* weight. Other computational cognitive control models using reinforcement learning (Blais et al., 2007) and Hebbian learning (Verguts and Notebaert, 2008, 2009) have also used fixed parameters in their simulations of various behavioral phenomena of conflict-control. Although these simulations matched empirical data well, the use of a fixed weight for information integration elicits two important, yet unanswered questions: (1) how is the weight determined? And (2) (How) does the weight change when the (experimental) environment changes? To answer these questions, we argue that the weight should be self-adapting based on how reliable the short-term and long-term information is. The idea of a self-adapting learning rate is not a new concept: in classical conditioning, there have been models that use the novelty of stimuli to affect learning rate (Pearce and Hall, 1980; Rescorla and Wagner, 1972; Schmajuk et al., 1996). In these models, novelty guides changes in learning rate, which in turn updates the association between conditioned and unconditioned stimuli. Since there are no conditioned and unconditioned stimuli in typical interference tasks, these models cannot be directly applied to simulating cognitive control processes. However, Bayesian models provide a natural solution of dynamically updating predictions based on integrating prior, temporally remote (long-term information) with recent observations (short-term information). Accordingly, several recent studies have employed Bayesian methods to model aspects of cognitive control (Ide et al., 2013; Mozer et al., 2002; Shenoy et al., 2010; Shenoy and Yu, 2011; White et al., 2012; Yu et al., 2009), which are reviewed in the next section.

2. Modeling cognitive control using Bayesian models

2.1. Overview of Bayesian methods

Bayes' theorem can be written as follows:

$$P(Y|X) = \frac{P(X|Y)P(Y)}{P(X)} \tag{1}$$

where X, Y are random variables (e.g., sensory input, internal states, motor output, etc.). Unlike conventional variables, the value of a random variable can vary due to randomness. Thus, a random variable is often represented in the probability distribution of its possible values. This equation means that the conditional probability of Y given X could be calculated using the probabilities of X, Y , and the conditional probability of X given Y . This equation is especially useful when $P(Y|X)$ (posterior probability) is difficult to estimate but $P(X|Y)$ is relatively easy to obtain. For example, when X is an observation and Y is an internal state which cannot be observed directly, one can infer the state of Y based on X and $P(X|Y)$ using the Bayes' theorem. Thus, Bayes' theorem can be used to infer, for instance, the distributions of conflict-control, based on the congruency observed. The estimated internal states can then be used to predict congruency in forthcoming trials. Bayesian methods have been widely applied in cognitive neuroscience studies (e.g., Bach and Dolan, 2012; Vilares and Kording, 2011), and a comprehensive review of studies using Bayesian methods is beyond the scope of this article. Instead, we focus on Bayesian models that employed a graphical representation, because it provides a natural representation of dependence on previous information.

A graphical representation of a Bayesian model (Koller and Friedman, 2009; Pearl, 1988) consists of a set of nodes and a set of edges connecting pairs of nodes. A node represents a variable in a Bayesian model, such as conflict, or observed congruency. In




Fig. 1. A basic Bayesian model of conflict-control. The model entails 2 variables, conflict (f) and observation (o), shown in gray indicating this variable is observable for each trial. The directed edges indicate the information flow.

addition, a node is associated with the probability distributions of the variable represented. An edge represents a relation (reflecting conditional independencies) between two nodes and can be either directed or undirected. For example, a directed edge from node A to node B means that the distribution of B depends on A. The edge is also associated with a distribution on which the estimation of parameters and Bayesian inference is based. This distribution encodes interactions between the two variables connected. Specifically, a directed edge from node A to node B is associated with a conditional probability distribution $P(B|A)$, which encodes how variable A influences variable B.

For example, the temporal dependency of conflict between trials can be formally represented using a Bayesian model (Fig. 1, see Yu and Cohen (2009) for a similar model). In this model f_i , and o_i denote predicted conflict level and observed congruency at trial i , respectively. f_i is quantified as the probability that the forthcoming trial is incongruent, ranging from 0 to 1. o_i is a binary variable in which 0 and 1 encode a congruent and an incongruent trial, respectively. The temporal dependency is represented by edges from the states at the current trial to the states at the next trial. An edge from predicted conflict level to observation is added to estimate f_i using Bayesian inference.

This Bayesian model does not only allow dependency between variables to be incorporated, but also significantly reduces the amount of computation needed to infer the states of these variables. Based on the structure of the graphical representation and the Markov property which states that each variable's future value is conditionally independent of the past, given its present value, the joint distribution in this model of cognitive control can be factorized as:

$$p(f_{i+1}, o_{i+1} | o_1, \dots, o_i) = \int p(f_i) p(f_{i+1} | f_i) p(o_{i+1} | f_{i+1}) df_i \tag{2}$$

The conditional probability distribution on the left side contains many variables. Without any prior knowledge of the model structure, it is intractable to calculate, or even store this distribution. However, the Markov property decomposes this distribution into a product of three much simpler distributions, each of which is easy to store and compute. Specifically, the Markov property shows that the prediction (states at trial $i + 1$) based on all previous information is equivalent to the prediction based only on the previous trial. In other words, relevant historical information is integrated into the states of the most recent trial. Thus, storage and computation involving older trials are not necessary.

In sum, the graphical representation incorporates dependency between model variables; and its structure greatly reduces the computational and storage burden of estimating posterior probabilities. Recently, Bayesian models with graphical representation have demonstrated great potential in modeling cognitive functions using behavioral data (Mozer et al., 2002; Reynolds and Mozer, 2009; Shenoy et al., 2010; Shenoy and Yu, 2011; Tenenbaum et al., 2006; Tenenbaum and Xu, 2000; Vossel et al., 2013; Yu et al., 2009) and brain imaging data (Behrens et al., 2007; den Ouden et al., 2010; Ide et al., 2013). In the following, we review recent studies using Bayesian models with graphical representation to model various aspects of cognitive control, including speed-accuracy trade-off

(Section 2.2), conflict effects in the Eriksen flanker task (Section 2.3), and response inhibition (Section 2.4). The Bayesian models reviewed below all attempted to account for decisions/behavior using within-trial and/or inter-trial simulations. In the within-trial simulations, those models accumulated evidence from independent sources via Bayesian integration. The decision/behavior best supported by the evidence was then selected by the models. In the between-trial simulations, predictions of stimuli were based on trial-history information integrated via Bayes' rule. The predictions were then used as the initial evidence for within-trial simulations.

2.2. Bayesian modeling of speed-accuracy trade-off

Some recent studies demonstrate the feasibility of using generative Bayesian models to simulate both within-trial dynamics and across-trial sequential effects in cognitive control. One such study applied a Bayesian model to explaining the dynamics of the speed-accuracy tradeoff and its dependency on trial-history from two speeded discrimination tasks (Mozer et al., 2002). Subjects responded to a target letter by pressing one button and responded to other letters by either pressing a second button (discrimination task) or doing nothing (go/no-go task). The Bayesian model used in this study operates at both within- and across-trial levels. At the within-trial level, the posterior distribution at a given time point encoded the probability of making one response vs. the other. This posterior distribution of responses depended on the prior distribution of responses and the sensory input. At the beginning of each trial, the posterior distribution is the same as the prior distribution. This posterior distribution of responses is then subject to (Bayesian) updating after each time point to favor the response suggested by visual input. Because this updating was performed at every time point, the influence of visual input accumulated with time, guiding the posterior distribution of response to gradually shift from the prior distribution to a distribution which is biased toward the correct response to the visual stimulus. Therefore, the effect of cognitive control is reflected by the change of the posterior distribution based on accumulation of sensory input. The simulated RT was the time that optimized the cost of sensory input accumulation against the probability of an incorrect response. The simulated accuracy was estimated from the posterior distribution at the simulated RT. For both the discrimination task and the go/no-go task, within-trial simulation successfully replicated accuracy and RT patterns from empirical data under different target to non-target ratios. Based on these simulations, the authors argued that speed-accuracy trade-off is optimal in these tasks in that it minimizes a cost that combines time pressure and the certainty of perception. At the across-trial level, the initial prior was also updated after each trial using a similar rule as used in within-trial simulation to account for sequential effects of response priming. The across-trial simulation successfully captured complex RT and accuracy patterns when trials were grouped based on the trial history, up to 4 trials preceding the current trial. In this study, then, a Bayesian model was used to simulate perceptual decision-making as an integration of prior information and visual input, which could naturally be modeled using the prior distribution and the likelihood distribution, respectively.

2.3. Bayesian modeling of the Eriksen flanker effect

To investigate different mechanisms that may account for generating conflict in the Eriksen flanker task (Eriksen and Eriksen, 1974), a study by Yu et al. (2009) applied two rival Bayesian models to behavioral data from a "deadline version" of the flanker task (Gratton et al., 1988; Servan-Schreiber et al., 1998). Here, subjects were pressured to make fast responses (in order to beat an experimenter-imposed deadline) to a target (center) letter (either

"S" or "H") that is flanked by distractors (either "S" or "H"). Thus a trial could be either congruent (e.g., "SSSSS" or "HHHHH") or incongruent (e.g., "HSHH" or "SSHSS"). Bayesian models were used to simulate two different potential sources of conflict, namely "compatibility bias" (a prior assuming that more than half of the trials were congruent) and "spatial uncertainty" (where perception of one letter was interfered with by nearby letters). Both models adopted the same hierarchical design with 3 levels: The highest level encoded the (in)congruency of a trial; below the congruency level was the stimulus level, in which there were 3 nodes, each representing a letter; at the bottom was a level of 3 nodes, each encoding activity of a group of neurons whose receptive fields were centered on a particular letter. The difference between the two models lay in how prior knowledge was applied: in the compatibility bias model, the prior assumed there were more congruent trials than incongruent trials, and each node at the level of neuronal activity was only influenced by the letter it represented and random noise. This model simulated a situation in which the conflict monitoring process was biased to an expectation of low conflict and receptive fields of neuron groups were narrow. By contrast, the prior of the spatial uncertainty model assumed a 50/50 distribution of congruent and incongruent trials, but here a neuron's activity was influenced by not only the letter its receptive field centered on, but also the neighboring letter(s). This model simulated a situation in which neurons have large receptive fields and the interference is caused by the ambiguous neuronal signals containing information of different letters.

Within-trial and between-trial simulations were then conducted using both models. To simulate within-trial dynamics, the simulation was divided into multiple time steps. The two models operated as neural decoders: they estimated visual input and congruency based on simulated neuronal activity. Thus, the key part of the simulation is the joint posterior distribution of letter and congruency conditioned on neuronal activity, which started with the prior distribution of congruency and visual input, and was then updated at every time step based on Bayes' rule. A simulated response was made when the marginal posterior distribution of one letter exceeded a pre-defined threshold. Both models successfully replicated RT distributions acquired from empirical studies using deadline Eriksen flanker tasks. The two models were also extended to allow for across-trial updates of the prior distribution of congruency. The extended models were able to simulate conflict adaptation and proportion congruency effects in Eriksen flanker tasks (Yu et al., 2009).

In another study, a Bayesian spotlight diffusion model was proposed to account for various aspects of the Eriksen flanker task (White et al., 2012). Specifically, a spotlight diffusion model (White et al., 2011) was used to simulate attentional mechanisms, and Bayesian belief-updating was employed to account for the information processing mechanisms involved in task performance (i.e., how evidence for response selection was accumulated within a trial). A spotlight was used to simulate the locus of attention, within which all information was selected as evidence for the decision-making process. At the beginning of a trial, the spotlight covered both the center target and the flankers, such that the overall evidence was driven predominantly by the flanker stimuli. The spotlight then gradually narrowed to only cover the central letter, resulting in the evidence being biased toward the target information. At each time point, beliefs were updated to incorporate new evidence about what the correct response was, using Bayesian evidence integration. Accordingly, the model predicts that responses are biased to the flankers at the beginning of a trial and then gradually shift to be dominated by target information. By fitting the model to empirical data, it was shown that this Bayesian spotlight diffusion approach could successfully account for the relationship between RT and accuracy in the Eriksen flanker task (White et al., 2012).

2.4. Bayesian modeling of response inhibition

Bayesian models have also been employed in investigating inhibitory control (Ide et al., 2013; Shenoy et al., 2010; Shenoy and Yu, 2011). In these studies, subjects performed a stop-signal task, in which a habitual response based on a (frequent) go signal needed to be suppressed when a (rare) stop signal was presented at varying intervals after the go signal. Bayesian models were used to provide a rational account (i.e., behavior guided by optimizing a cost function) for various behavioral patterns observed in the stop-signal task. Here, one Bayesian model was used to simulate beliefs about the appropriate action to take, and a second Bayesian model was used to simulate beliefs about perceiving a stop signal. Within each trial, both beliefs started with the (true) prior probability of go/stop trials in the task, and were then updated based on visual input using Bayes' rule at every time step. For each time step, a cost function was calculated based on the beliefs and possible actions available. An action (i.e., button-press vs. withholding response) was selected by minimizing the cost function. This within-trial Bayesian model successfully simulated the pattern of increased error rates as the interval between the onset of the go signal and the onset of the stop-signal increased. It could also simulate the commonly observed faster responses in error stop trials compared to successful go trials (Shenoy et al., 2010; Shenoy and Yu, 2011). A third Bayesian model was employed for simulating between-trial effects, predicting the likelihood of encountering a stop-signal in the forthcoming trial. This prediction was a linear combination of the prior probability of encountering a stop-trial and the posterior probability of encountering a stop-trial based on trial history. These two probabilities were integrated via fixed weights. The prediction significantly correlated with RTs, where high probability of encountering a stop signal predicts slower responses, suggesting more inhibitory control being exerted (Ide et al., 2013). Additionally, this model successfully simulated post-stop-trial response slowing, and increased RTs and error rates when the proportion of stop-trials increased (Shenoy et al., 2010; Shenoy and Yu, 2011). Based on these Bayesian predictions, fMRI data recorded during the stop-signal task (Ide et al., 2013) further revealed that the dorsal ACC encodes prediction error (presence/absence of stop signal – prediction).

2.5. Toward modeling the flexibility of cognitive control

Although the models reviewed above were successful in simulating various effects of cognitive control, the *fixed* parameters controlling learning used in these simulations raise several concerns. First, it is unclear whether the way in which these parameters were determined in the models reflects the mechanisms of parameter-selection in the brain. This is especially unlikely (or impossible) in cases where optimal parameters were fit from data: here, a model determines the parameters *after* acquiring all data, in contrast to the brain having to determine the parameters on the fly. Second, even with a model that could potentially provide on-the-fly simulation (e.g., the model in Fig. 1), the lack of a mechanism for online adjustment of cognitive control based on adaptively integrated information from various time scales (e.g., long-term and short-term history) would result in sub-optimal performance in a non-stationary environment, because without such a mechanism, the parameters that determine the level of cognitive control (e.g., the learning rate) have to stay constant throughout the experiment. Given the high flexibility required of cognitive control, it is unclear that those fixed parameters are globally optimal across various experimental configurations. In fact, it has been shown that fixed learning rates are suboptimal in non-stationary environments (Behrens et al., 2007; den Ouden et al., 2010; Vossel et al., 2013). In the next section, we therefore propose a Bayesian model




Fig. 2. The graphical representation of the Bayesian model of flexible conflict-control. The model uses 3 variables, volatility (v), conflict (f), and observation (o), shown in gray indicating this variable is observable for each trial. The directed edges indicate the information flow.

that resolves these two concerns. This model simulates the on-the-fly selection of parameters in the brain by making estimates only based on previous trial information. It also models the flexibility of cognitive control by incorporating a component that accounts for changes in the experimental environment. By applying this model to empirical data, we demonstrate that it can simulate various key phenomena of cognitive control in a Stroop task.

3. Modeling the flexibility of cognitive control using Bayesian methods

3.1. A Bayesian model of flexible conflict-control

Here, we propose a Bayesian model that can account for the flexibility of cognitive control over conflict in a non-stationary environment. The modeling done relies on the ability to perform statistical inference, taking the perspective that the regulation of cognitive control should be considered as a process of *predicting* the optimal amount of cognitive control required in a given context. To achieve this contextual flexibility, the model estimates future conflict from previous experience and, importantly, it does so via a weighed integration of longer-term and short-term estimations of conflict distributions, with the integration weights being adjusted on the basis of the (belief about) volatility of the environment. For instance, in a stable environment (e.g., when most trials are congruent/incongruent), the weights are biased to historically remote/long-term information because an occasional oddball trial (e.g., an incongruent stimulus in a largely congruent trial history context) is unlikely to reflect a true change in the environmental statistics. When the environment is fast-changing (e.g., when the proportion of congruent trials varies frequently over time), however, the weights are biased to more recent information. This is because older information is likely to be outdated, and an unexpected trial type may indicate a true change of conflict likelihood in the environment. In order to assess the stability of the environment, we extend the model in Fig. 1 by adding a volatility variable (denoted by v), the belief of which in turn determines the weights of integration (Fig. 2). The structure of this model is identical to the model of Behrens et al. (2007). The mathematical formulation of the model is described in details in the appendix. This model is also an example of a *hierarchical model*, in which the information flows in one direction, that is, there are no reciprocal edges from nodes at lower levels to nodes at higher levels. Hierarchical Bayesian models have been widely used in modeling cognitive functions such as categorization (Tenenbaum and Xu, 2000; Xu and Tenenbaum, 2007) and visual cognition (Lee and Mumford, 2003; Summerfield et al., 2011). This model yields a probability distribution over the predicted conflict level variable. In the implementation of this model, we approximate predicted conflict level using the probability of encountering an incongruent trial. In other words, the predicted conflict level is higher if the next trial is deemed more likely to be incongruent. Both variables are then used to determine the

amount of control needed, which is reflected in sequential effects such as conflict adaptation and longer-term effects like proportion congruency. Another way to understand the relation between volatility and predicted conflict level can be drawn from Yu and Dayan (2005): predicted conflict level encodes the probability (distribution) that the forthcoming trial is incongruent. The variance of this distribution (i.e., change of probability) is determined by the volatility. Compared to previous studies that model flexible control (e.g., Reynolds and Mozer, 2009; Yu and Cohen, 2009), the proposed Bayesian model integrates information sampled from multiple temporal scales, and does not assume any specific structure of the environment (e.g., eight types of blocks in Reynolds and Mozer (2009), or a fixed rate of change in the environment (Yu and Cohen, 2009)). The present model also differs from another Bayesian model that integrates information from sources of different temporal scales (Kording et al., 2007) in that our model infers the weights of integrating those sources, whereas the model of Kording et al. (2007) assumes a constant integration across those sources and infers the state of each source. By applying this model to empirical datasets, we demonstrate that it can account for classic short-term and long-term effects of conflict-control (Section 3.2), and more importantly, it successfully simulates the flexibility of conflict-control (Section 3.3).

3.2. Simulating the short-term and long-term trial history effects of conflict-control

In this section, we first confirm that short-term and long-term trial history effects occur simultaneously in a single empirical dataset (e.g., Torres-Quesada et al., 2013). Then we demonstrate that both effects can be simulated simultaneously using our Bayesian model. Importantly, no post hoc optimization of parameters was necessary in the simulation.

3.2.1. Subjects

Fifty-six healthy volunteers (mean age = 26.1, 30 females) gave informed consent in accordance with institutional guidelines. All subjects were native or highly proficient English speakers and had normal or corrected-to-normal vision.

3.2.2. Stimuli and procedure

Stimulus delivery and behavioral data collection were carried out using Presentation software (<http://www.neurobs.com/>). Stimuli were presented on a 19 in. LCD screen with a refresh rate of 60 Hz. Stimuli consisted of a collection of 24 black and white photographs of male and female faces (12 each) of neutral expression that were overlaid with red gender word labels (“male” and “female”), which could be printed in lower or upper case lettering. On each trial, one face-word compound stimulus (subtending approximately 3° of horizontal and 4° of vertical visual angle) was presented against a gray background in the center of the screen. Stimuli were presented for 500 ms, followed by a jittered inter-stimulus interval ranging from 2 to 3 s in uniformly distributed steps of 500 ms, during which a fixation cross remained on screen. Subjects performed a speeded button response that categorized the gender of the face stimulus with either index finger (for example, left-hand response to male faces, right-hand response to female faces, counterbalanced across subjects), while trying to ignore the task-irrelevant gender labels and stimulus locations. Face stimuli never repeated across adjacent trials, and the lettering alternated between lower- and upper-case across trial. A practice run was conducted before the main task to ensure subjects comprehended the task requirements.

3.2.3. Experimental design

This task consisted of 7 runs of 4 blocks each. Each block contained 41 trials with pseudo-randomized congruency. Across all blocks, the proportion of congruent trials followed the order of approximately (deviated by 1 trial, or ~2.4%) 15%, 35%, 65%, 85%, 75%, 50% and 25%, repeated for 4 times over the 7 runs. Within each block, the proportion congruency remained constant. The starting proportion congruency and the order of the sequence were counter-balanced across subjects. To model both the conflict adaptation and proportion congruency effects, we analyzed response time (RT) data using a 7 (proportion congruency) × 2 (previous congruency) × 2 (current congruency) factorial design.

3.2.4. Data analysis

For the behavioral data, the mean RT was computed in each subject for each of the experimental cells, excluding incorrect and post-error trials, as well as RT values that deviated >2 standard deviations from an individual subject's grand mean. The trimmed RT values were then averaged across subjects and entered into repeated measures 3-way analyses of variance (ANOVAs) with the factors described above. For the simulation data, the trial sequences observed by the subjects were fed to the model to produce trial-by-trial estimates of volatility and predicted conflict level. Then, for each experimental cell, the mean volatility and predicted conflict level were computed, excluding trials that were excluded in empirical data analysis. Finally, to link model predictions to the empirical data, a general linear model (GLM) was constructed using group means of the parameter estimates (28 conditions), which were then fit to the group mean of RTs. Specifically, this GLM contained 6 regressors (i.e., free parameters), namely the volatility, the predicted conflict level, and the grand mean, separately for congruent and incongruent trials. Note that this fitting procedure is not geared at finding the optimal parameters for our model. Rather, the purpose of this fitting was similar to within-trial simulation, or in other words, to quantify how predictions made prior to a trial influence the information processing during that trial, as reflected in the RTs.

3.2.5. Results and discussion

The subjects performed the task with high accuracy (mean = 92.9%). The 3-way ANOVA on empirical RTs showed a significant effect of current congruency ($F_{1, 55} = 57.07, P < 0.001$), due to longer RTs in incongruent trials (592 ± 11 ms) than in congruent trials (569 ± 10 ms). An interaction between proportion congruency and current congruency was also found (the proportion congruent effect, $F_{1, 55} = 2.73, P < 0.03$), driven by a decrease in interference effects as the proportion of incongruent trials increased (Fig. 3A). There was also an interaction between previous and current trial congruency (the conflict adaptation effect, $F_{1, 55} = 5.03, P < 0.03$), driven by a larger interference effect (26 ± 3 ms) in post-congruent trials than in post-incongruent trials (16 ± 3 ms; Fig. 3B). Thus, the behavioral results replicated a large literature on the proportion congruency (for review, see Bugg and Crump, 2012) and the conflict adaptation effect (for review, see Egner, 2007), as well as previous findings of these two effects occurring simultaneously in the same data set (Torres-Quesada et al., 2013). As can be seen in Fig. 3C, both effects were successfully simulated using our Bayesian model. Specifically, the model predicted an interference effect (congruent trials: 569 ms; incongruent trials: 592 ms), a decreased interference effect as the proportion of incongruent trials increased (Fig. 3C), and a higher interference effect in post-congruent trials (27 ms) than in post-incongruent trials (18 ms). These simulation results suggest that our Bayesian model is able to simultaneously account for both long-term and short-term effects of cognitive control. These fits were achieved using a single control mechanism with a flexible learning rate

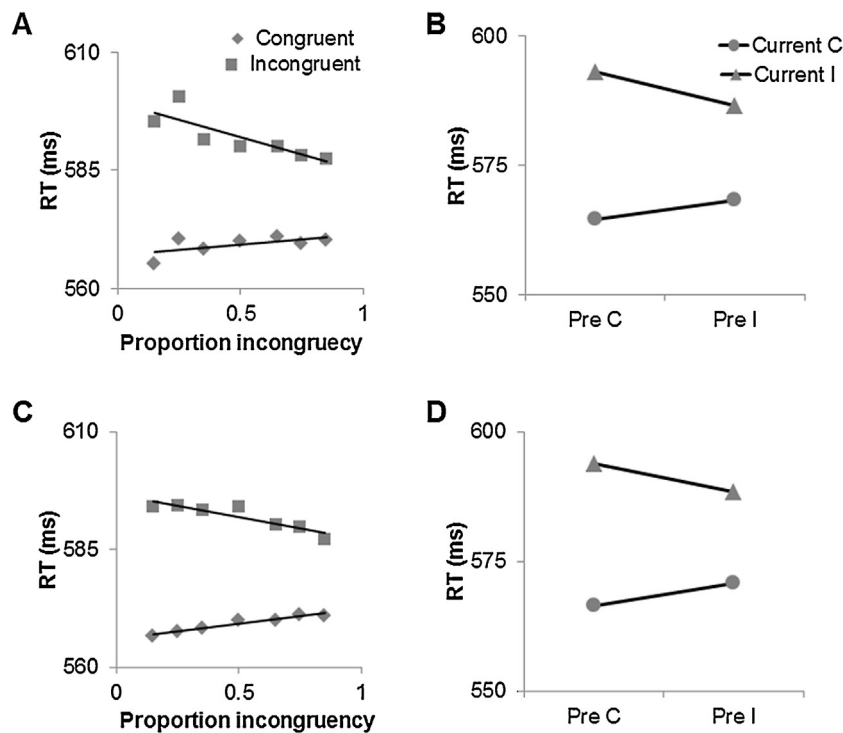


Fig. 3. Empirical and simulated effects of congruency, proportion congruency, and conflict adaptation. (A) Empirical proportion congruency effect, with RT plotted as a function current trial congruency and the block-wise proportion of incongruent trials. (B) Empirical conflict adaptation effect, with RT plotted as a function of current and previous trial congruency. (C) Simulated proportion congruency effect, plotted in the same way as in (A). (D) Simulated conflict adaptation effect, plotted in the same way as in (B). Pre C/Pre I = preceded by a congruent/incongruent trial; Current C/Current I = current trial is congruent/incongruent.

rather than the dual mechanism structure of De Pisapia and Braver (2006) or the separate fits with different learning rates as applied by Botvinick et al. (2001). Moreover, our data fits were derived from on-the-fly simulations and not based post hoc setting of learning rate parameters.

3.3. Simulating the flexibility of conflict-control

To further demonstrate the model's ability to simulate the flexibility of cognitive control, we conducted a second experiment, in which we created two environments with different dependence on short-term and long-term information. We show that our model can successfully simulate the behavioral patterns observed in the empirical data.

3.3.1. Subjects

Forty-six healthy volunteers (mean age = 19.9, 33 females) gave informed consent in accordance with institutional guidelines. All subjects were native or highly proficient English speakers and had normal or corrected-to-normal vision.

3.3.2. Stimuli and procedure

The same stimuli and basic task procedure was used as the one described above in Section 3.2.

3.3.3. Experimental design

This task consisted of 4 runs of 9 blocks each. Each block contained 20 trials with pseudo-randomized congruency. The first block had 50% congruent trials and served as a "reset" (or "burn-in") block to bring the predictions to the same baseline at the beginning of each run. To create experimental environments that differ in their dependence on long-term and short-term trial history, a run could be either volatile (the proportion congruency alternated

between 20% and 80% every block) or stable (the proportion congruency remained either 20% or 80% for all 8 post-reset blocks). Subjects were given no indication that a new block was beginning. The order of volatile and stable runs was counter-balanced across subjects. This manipulation resulted in a 2 (volatile/stable) \times 2 (proportion congruency) \times 2 (current trial congruency) factorial design.

3.3.4. Data analysis

We applied the same analyses as described in Section 3.2. Note that trials in reset blocks were also given to the model, so as to also generate a reset of trial-by-trial estimates of volatility and predicted conflict level in the model at the beginning of each run. However, reset block trials were excluded from further analyses.

3.3.5. Results and discussion

Participants performed the task with high accuracy (mean = 92.8%). The 3-way ANOVA on empirical RTs again revealed a significant effect of current trial congruency ($F_1, 45 = 49.3, P < 0.001$), due to longer RTs in incongruent trials (549 ± 13 ms) than in congruent trials (522 ± 11 ms). The proportion congruency effect was also found, reflected in a significant interaction between proportion congruency and current trial congruency ($F_1, 45 = 10.5, P = 0.002$). This effect was driven by a larger interference effect in 80% congruency blocks (33 ± 5 ms) than in 20% congruency blocks (22 ± 4 ms). Importantly, we also observed a significant main effect of volatility ($F_1, 45 = 4.5, P = 0.04$), due to longer RTs in volatile runs (540 ± 12 ms) than in stable runs (531 ± 12 ms). Note that this main effect was not driven by "outliers" in experimental cells, because no interactions involving volatility and any of the other factors were found. Furthermore, a trend for longer RTs in volatile compared to stable environments can be observed in all 4 current trial congruency \times proportion congruency conditions (Fig. 4A). This main effect of volatility may reflect a cost of frequently adjusting

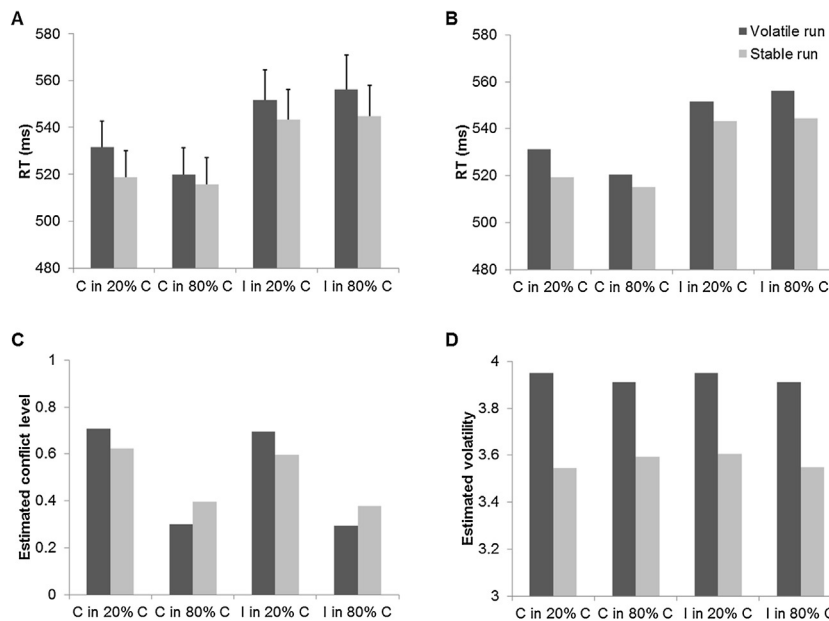


Fig. 4. Empirical and simulated effects of congruency, proportion congruency in an environment of changing volatility. All quantities are plotted as a function of volatility, proportion congruency and congruency at the current trial. (A) Empirical reaction times and their standard errors. (B) Simulated reaction time. (C) Estimated predicted conflict level from the Bayesian model. (D). Estimated volatility from the Bayesian model in arbitrary units. C/I in 20%/80% C = congruent/incongruent trials in a block of 20%/80% congruent trials.

the strategy of cognitive control (e.g., adjusting the dependence on long-term vs. short-term information) in volatile runs.

This pattern of RTs was again successfully simulated using our model (Fig. 4B): firstly, the model recapitulated the main effect of current congruency (congruent trials: 522 ms; incongruent trials: 549 ms); secondly, it simulated the proportion congruency effect (inference effect in high proportion congruency blocks: 33 ms; inference effect in low proportion congruency blocks: 22 ms); and lastly (and most importantly), it predicted longer RTs in volatile runs (540 ms) than in stable runs (531 ms). We contend that this “volatility cost” performance pattern can in fact only be accounted for by using a model that estimates the volatility of the environment. Consider a model with no information about volatility (e.g., having a fixed learning rate): in such a model, each trial’s contribution to the prediction of predicted conflict level is a constant. Furthermore, the trial’s contribution to the prediction of all forthcoming trials is also a nearly constant, because its influence decays continuously with a constant discounting rate after each trial, and approaches zero in a relatively short period of time (except for extremely low learning rates, which are unrealistic given the commonly observed short-term effects of cognitive control). As a consequence, using a fixed learning rate, two sequences of trials with the same proportion congruency and number of trials will produce the same sum of conflict-level estimates across all trials, regardless of the volatility of those sequences. Indeed, even as shown in our model, estimates of predicted conflict level displayed an interaction between volatility and proportion congruency (Fig. 4C). Thus, it is impossible to account for the pattern of slower RTs in volatile compared to stable runs in the empirical dataset when using only predicted conflict level estimates; they would have to be combined with estimates of volatility (Fig. 4D) to simulate the empirical data pattern. To further illustrate this point, we performed an additional analysis to simulate the results in this experiment using a reinforcement learning algorithm. Specifically, we created 3 learners with different learning rates: 0.05 (high dependence on long-term information), 0.5 (balanced dependence on long-term and short-term information) and 0.95 (high

dependence on short-term information). The learning was conducted using the following equation:

$$f_{i+1} = f_i \times (1 - \alpha) + o_i \times \alpha \quad (3)$$

where f_{i+1} is the predicted conflict level at trial $i + 1$, α is the learning rate, and o_i is the observed congruency. If any of these learners were able to account for the behavioral pattern observed in experiment 2, there should be a significant difference in f_{i+1} between the 2 volatility levels, because f_{i+1} is the only output of these learners. However, in none of the 3 learners did we observe such a significant difference (all P s > 0.2). Thus, in the absence of a (volatility-modulated) flexible learning rate, the model is unable to account for the behavioral patterns obtained across different task settings.

An alternative explanation to the main effect of volatility is that, similar to the proportion congruency effect where trials with better predicted (in)congruency were associated with faster RTs, less frequent alternation of the underlying proportion congruency in stable runs was associated with better predicted (in)congruency, and thus faster RTs, compared to volatile runs. Hence, a reinforcement learning model with fixed individual learning rate (and without an explicit variable encoding volatility), paired with a linear mapping from prediction error of the reinforcement learning model to RT (Jones et al., 2013) could in principle also account for the main effect of volatility. In this type of model, both the main effect of volatility and the proportion congruency effect only depend on a single underlying variable, the learning rate. It follows that there should be a correlation between these two effects across subjects with varying learning rates. In fact, using a reinforcement learning algorithm and 100 randomly generated trial sequences, we found an almost perfectly linear negative correlation ($r = -0.99$) between these two effects across a wide range of learning rates (from 0.05 to 0.95, step size = 0.01). However, our empirical data produced no support for this model, as we observed no correlation between effects of volatility and proportion congruency effects (correlation $r = -0.07$, $P = 0.62$), thus suggesting that an account that ascribes the effects of volatility and proportion congruency to a single underlying variable cannot be accurate. By

contrast, the Bayesian model, which incorporates volatility as an explicit additional variable, can account for both the group level RT pattern as well and the uncorrelated main effect of volatility and proportion congruency effect across subjects, because the two effects are here attributed primarily to two distinct factors (the proportion congruency effect is attributable primarily to prediction error, and the main effect of volatility is attributable primarily to volatility).

Note that the predicted conflict levels display the inverse pattern of empirical RTs, e.g., for incongruent trials higher estimated (or predicted) conflict levels are predictive of faster RTs. This pattern essentially corresponds to the classic, intuitive explanation of the empirical proportion congruency effect, namely, that control is higher in conditions where conflict is frequently encountered (e.g., Carter et al., 1998).

A few additional points should be noted in the interpretation of these data. First, both conflict adaptation and proportion congruency effects have at times been argued to exclusively reflect associative processes related to the particular stimulus and response features of the task (e.g., Mayr et al., 2003; Hommel et al., 2004; Schmidt and Besner, 2008). By contrast, the present results show that both of these effects can be faithfully captured by a model that does not consider specific stimulus or response features at all—it only learns about the incidence of congruent and incongruent stimuli. This documents that, at least in principle, learning of specific physical stimulus and response properties is not a necessary precondition for producing these effects. Second, while the main effect of volatility can be quantitatively accounted for by our Bayesian model (the manipulation of volatility was captured by the volatility variable, as can be seen in Fig. 4C), the model architecture itself does not necessitate such an effect. In other words, the model could equally well fit behavioral data in the absence of a main effect of volatility (unpublished observations).

Nevertheless, this does of course not mean that the empirical data themselves were not potentially subject to such lower-level learning effects. We consider it unlikely that such processes contributed in a substantial manner to the present results, however, for the following reasons. First, in order to prevent trial-by-trial priming effects at the level of physical stimulus features (Hommel et al., 2004; Mayr et al., 2003), face stimuli in the present experiments never repeated across successive trials, and the lettering of the distracter labels alternated between lower- and upper-case across trials. Second, in order to minimize the possibility that proportion congruency effects in our protocol would be mediated by subjects associating specific face stimuli with a particular response (e.g., the gender-congruent response in high proportion congruency blocks), we included a large number (24) of unique facial identities (cf. Bugg and Hutchison, 2013). Nonetheless, this leaves the possibility that subjects may use the contingency between the distracter (in our case, the gender word) and the response to guide their action selection (Bugg, 2012; Bugg and Chanani, 2011; Schmidt and Besner, 2008). For example, in an environment of high proportion congruency, the gender word is highly predictive of the correct response. However, note that in volatile runs, this contingency changes every 10 occurrences for each word (on average), leading to a less predictive word-response association than in stable runs, where the contingency remains unchanged. Therefore, if distracter-response contingency were a major contributing factor to the proportion congruency effect in our data, this effect should be modulated by the contingency's predictive power (i.e., volatility), resulting in a 3-way interaction between volatility, proportion congruency and current-trial congruency. However, we did not observe such an interaction ($F_1, 45 = 2.9$, n.s.). More specifically, the contingency account predicts that contingency with higher predictive power (i.e., the stable runs) should evoke larger proportion congruency effects than low-contingency conditions (i.e., the

volatile runs). However, numerically, the opposite is true for our data (volatile runs: proportion congruency effect = 16 ± 5 ms; stable runs: proportion congruency effect = 4 ± 4 ms; $t_{45} = 1.7$, n.s.). Thus, contingency learning seems highly unlikely to have contributed to the empirical proportion congruency effects that our model simulated.

In sum, we showed that a Bayesian model that learns to predict control demand using a flexible, volatility-driven learning rate, can account for simultaneously occurring conflict adaptation, proportion congruency, and volatility effects, without the need for multiple controllers or post hoc fit-derived learning rate parameters. We conclude that this model represents a promising new application of a Bayesian approach to exploring computational mechanisms of cognitive control, in particular with respect to simulating the flexibility that is required of control processes wielded in a changing environment.

4. Future directions and concluding remarks

4.1. How “Bayesian” is cognitive control?

One important avenue for future work in this context is to evaluate to what extent (or in what sense) cognitive control might actually be Bayesian in nature. According to Bowers and Davis (2012), there are three levels at which one can use Bayesian methods in modeling cognitive processes: as computational tools, for generating “optimal” benchmarks for cognitive processes, and for modeling the actual neural computations carried out by the brain. The Bayesian models reviewed in this paper, along with our own model, all operate at the second level: these Bayesian models were treated as “optimal observers”, and produced optimal predictions, which in turn were used to account for behavior (and sometimes, neuroimaging data). A crucial next step then is to examine more widely whether it is possible to find neural substrates for the variables expressed in these Bayesian models, such as the brain regions/systems that store information concerning the prediction of conflict level and volatility. Although to the best of our knowledge no prior study has investigated this topic in the context of cognitive control, related work may be able to nevertheless shed some light on this question by indicating whether the kinds of variables involved in the work reviewed above are likely to be encoded in the brain, and what particular brain regions might be candidates for mediating these kinds of computations.

Statistically speaking, the prediction of conflict level supplies a probability (also, ‘expected value’, if a fixed maximum amount of control is assumed). In the literature on risky decision-making, evidence of neural representations of reward probability and expected reward (for reviews, see Platt and Huettel, 2008; Rushworth and Behrens, 2008) has been found in the midbrain (Fiorillo et al., 2003, 2005), the ventral striatum (Knutson et al., 2005; Preusschoff et al., 2006), the ACC (Amiez et al., 2006), and the medial (Knutson et al., 2005; Matsumoto et al., 2003; McClure et al., 2004; Volz et al., 2003, 2004) and lateral PFC (Matsumoto et al., 2003).

Compared to probability and expected value, fewer studies have examined the neural substrates of volatility. Using a gambling task that manipulated the frequency of altering the underlying probability distribution of winning options, Behrens et al. (2007) found that volatility predicted in a Bayesian model correlated with activity in the dorsal ACC when subjects observed the outcome of the gamble. The ACC activation observed in Behrens et al. (2007) was also accounted for using another model that separately encoded positive and negative prediction error (the difference between expected outcome and actual outcome) to model ACC activity (Silvetti et al., 2011). This model furthermore successfully simulated ACC activation patterns in another independent fMRI dataset

that manipulated volatility (Silvetti et al., 2013), suggesting that the ACC may indeed be a prime candidate for encoding this type of information on environmental variability also in the context of steering cognitive control. Another computational model of the medial PFC captured the signal variations evoked by changes in volatility through simulated negative surprise signals (the absence of a predicted outcome) (Alexander and Brown, 2011). However, these studies also suggest that the ACC's response to expected uncertainty could be higher than its response to unexpected uncertainty. Thus, future study designs should attempt to de-confound the manipulation of volatility from the expectancy of uncertainty.

Volatility can also be considered as second-order uncertainty (e.g., the deviation of probability, see Yu and Dayan, 2005), and thus it may also be computed in brain regions that have been found to encode second-order statistics, such as deviation. Deviation of reward from its mean has been found to be encoded in posterior cingulate cortex neurons (McCoy and Platt, 2005). Another study has shown that in the primate anterodorsal septum, neural firing rates display an “inverted-U” shape as a function of the probability of reward, peaking at 50% (Monosov and Hikosaka, 2013). Given that the standard deviation of reward increases starting from 0 reward probability, then peaks at 50%, and finally drops until the probability of reward reaches 100%, this inverted-U pattern of neural firing may also represent potential neural substrates of deviation and/or volatility. A similar neural firing pattern was also found in the mid-brain (Fiorillo et al., 2003) and the ventral striatum (Preusschoff et al., 2006) in monkeys. Volatility is also associated with confidence (or uncertainty). For example, in a stable environment, prediction of conflict is usually more confident because the expectation is less likely to be violated. Previous studies have shown that confidence of decision is encoded in the lateral intraparietal cortex (Kiani and Shadlen, 2009) as well as the ventral medial and rostralateral prefrontal cortices (De Martino et al., 2013). Finally, high volatility also arguably makes decision-making more ambiguous; hence neural substrates of ambiguity, as found in orbital frontal cortex (Hsu et al., 2005) and the lateral PFC (Huettel et al., 2006) may also contribute to encoding volatility.

To what degree the putative computation of volatility and the consequent titration of predicted conflict and recruitment of control processes outlined in our own model map onto these anatomical territories remains an interesting topic for a future neuroimaging study. Moreover, apart from testing the neural representations of the variables' distributions in a Bayesian model, another direction of future research is to examine whether the brain implements cognitive control in a Bayesian fashion as far as neural computations are concerned. To test this hypothesis, one would need to identify the neural processes of information integration, for example, how posterior distributions are calculated based on prior and likelihood.

5. Conclusion

The flexibility of cognitive control enables the brain to adaptively adjust the degree of top-down biasing in a dynamically changing environment. This adjustment can be cast as a prediction of control demand, which can be optimally achieved via Bayesian belief propagation. Using Bayesian models, previous studies have successfully modeled various phenomena of cognitive control. Yet, those models usually depend on specific (and post hoc) selection of parameters to achieve optimal performance. This approach is not only unlikely to mimic the true mechanism of parameter selection in the brain, but is also likely incapable of accounting for the flexible adjustments in control strategies when the environment changes. To solve this problem, we propose a Bayesian model that takes into account the environment's volatility, which in turn adjusts

the model's dependence on short-term and long-term information to produce a prediction of control demands (i.e., predicted conflict level of the forthcoming trial). Our model, along with other Bayesian models, exhibits great potential for improving our understanding of cognitive control processes in a principled, formal framework. A prime target for future work is to identify possible neural substrates of the key variables and processes implied by these models, and further to determine the specific neural implementation (whether Bayesian or not) of cognitive control.

Acknowledgments

We thank Tim Behrens for sharing code and Chris Summerfield for discussion of the model. This work was funded in part by National Institute of Mental Health (NIMH) grants R01MH087610 (T.E.) and R01MH097965 (T.E.).

Appendix A.

The generative model is shown as follows (see below for definition of distribution parameters):

$$\begin{aligned} v_{i+1} &\sim N(v_i, \sigma_v) \\ f_{i+1} &\sim \text{Beta}(\alpha, \beta) \\ o_{i+1} &\sim \text{Bernoulli}(E(f_{i+1})) \end{aligned} \quad (\text{A1})$$

This Bayesian model simulates a trial sequence of an experiment on a trial-by-trial basis. Within each trial, the simulation takes the form of a predict-update algorithm that is used in Kalman filters (Johan Masrleiez and Martin, 1977). In other words, the simulation of each trial $i + 1$ contains two steps. The first step makes prediction of the states of v_{i+1} , f_{i+1} and o_{i+1} before the stimulus is presented. These predictions were used to account for the behavioral patterns observed in the empirical studies reported above. The second step updates/filters the belief of the states of v_{i+1} and f_{i+1} given the observed congruency o_{i+1} . These two steps were repeated for each trial to generate trial-by-trial estimates.

In the first step, the model initially predicts a joint distribution based on the model's previous states and 2 transition distributions:

$$\begin{aligned} p(\sigma_v, v_{i+1}, f_{i+1} | o_1, \dots, o_i) &= \int \int p(v_{i+1} | v_i, \sigma_v) \\ &\times p(f_{i+1} | f_i, v_{i+1}) p(\sigma_v, v_i, f_i | o_1, \dots, o_i) dv_i df_i \end{aligned} \quad (\text{A2})$$

where σ_v constrains how v can change over time. Specifically, $p(v_{i+1} | v_i, \sigma_v) \sim N(v_i, \sigma_v)$. In other words, the transition distribution $p(v_{i+1} | v_i, \sigma_v)$ is a Gaussian distribution with the mean of v_i and the standard deviation of σ_v . This transition distribution assumes that v is most likely to remain in its previous states, although it can also possibly drift to another state. σ_v determines how likely it is for v_i to shift to a new state. Because in experiment 2 the volatility altered between conditions, we set σ_v as a variable and used the Bayesian model to infer its state. This transition distribution determines how the estimate of v_i is used to predict its future state, which is in turn employed to compute the predicted conflict level, as described below.

$p(f_{i+1} | f_i, v_{i+1})$ describes how f can change over time. $p(f_{i+1} | f_i, v_{i+1}) \sim \text{Beta}(\alpha, \beta)$. This distribution is a beta distribution, with parameters α and β taking the following form:

$$\begin{cases} f_i = \frac{\alpha - 1}{\alpha + \beta - 2} \\ v_{i+1} = \log_2(\alpha + \beta) \end{cases} \quad (\text{A3})$$

There are two main reasons for using a beta distribution: first, we defined the predicted conflict level as the probability of

encountering an incongruent stimulus in the upcoming trial, so the possible value of f_{i+1} should be limited to a range of 0–1, which is also the range of values that a beta distribution is defined on. Second, and more importantly, with the current set-up, the probability density function of $p(f_{i+1}|f_i, v_{i+1})$ takes the following form:

$$p(f_{i+1}|f_i, v_{i+1}) = \text{constant} \cdot f_{i+1}^{\alpha-1} \cdot (1 - f_{i+1})^{\beta-1} \quad (\text{A4})$$

which can be interpreted as the likelihood function of observing $\alpha - 1$ incongruent trials and $\beta - 1$ congruent trials with the underlying proportion incongruency f_i . Thus, a larger $p(f_{i+1}|f_i, v_{i+1})$ suggests the prediction of conflict is better supported by temporally more-extended trial-history information. Following this interpretation, according to the Eqs. (A3) and (A4), v_{i+1} controls the length of the trial sequences of this likelihood function. A larger v_{i+1} suggests a longer trial sequence to take into account, which in turn indicates more dependence on long-term information, or a more stable condition. In other words, a larger v_{i+1} leads to a narrower spread of $p(f_{i+1}|f_i, v_{i+1})$, which constrains f_{i+1} from drifting far from its previous state. Another interpretation of $p(f_{i+1}|f_i, v_{i+1})$ can be linked to the learning rate used in many reinforcement learning models: a narrower spread of $p(f_{i+1}|f_i, v_{i+1})$ results in smaller difference between f_i and f_{i+1} that resembles the effect of a smaller learning rate, compared to a wider spread of $p(f_{i+1}|f_i, v_{i+1})$ determined by a smaller v_{i+1} . However, it is counterintuitive to have a larger volatility value when the environment is more stable. Thus when reporting volatility, we use a linear transform to make more volatile task settings correspond to larger volatility estimates while preserving the quantitative patterns of our results (see below). Furthermore, f_i is the mode of $p(f_{i+1}|f_i, v_{i+1})$, indicating that the predicted conflict level is most likely to reflect its previous state. Thus, $\int \int p(v_{i+1}|v_i, \sigma_v) p(f_{i+1}|f_i, v_{i+1}) dv_i df_i$ can be viewed as defining how the prediction is made by incorporating various learning rates. The third term in the integral, $p(\sigma_v, v_i, f_i|o_1, \dots, o_{i-1})$ is the belief in the previous trial and also represents the weights $p(\sigma_v, v_{i+1}, f_{i+1}|o_1, \dots, o_i)$ in making the prediction (Eq. (A1)). After the joint distribution $p(\sigma_v, v_{i+1}, f_{i+1}|o_1, \dots, o_i)$ is calculated, the estimates of volatility and conflict are computed as the mean of their corresponding marginalized distributions. The observed congruency o_{i+1} was then predicted to have a Bernoulli distribution with a probability of $E(f_{i+1})$ of incongruency, where $E(f_{i+1})$ denotes the mathematical expectation of f_{i+1} .

In the first step, we adopted a numeric implementation for this Bayesian model to avoid the complexity of developing an analytical implementation: the range of each of the variables of σ_v, v and f was divided into multiple segments with equal length. For example, f_i was represented using an array ranging from 0 to 1 and with a step size of 0.02 (that is, 51 cells). The value of each cell represented the probabilistic density at that point. Similarly, the joint probability distribution was represented by a 3D array with 3 dimensions of σ_v, v_{i+1} and f_{i+1} . $p(f_{i+1}|f_i, v_{i+1})$ was represented using a 3D array with 3 dimensions of f_{i+1}, f_i and v_{i+1} . And $p(v_{i+1}|v_i, \sigma_v)$ was represented using a 3D array with 3 dimensions of v_{i+1}, v_i and σ_v . All the aforementioned calculations were performed on these arrays. Specifically, step 1 took the form of:

$$p(\sigma_v, v_{i+1}, f_{i+1}|o_1, \dots, o_i) = \sum_{v_i} \sum_{f_i} p(v_{i+1}|v_i, \sigma_v) \times p(f_{i+1}|f_i, v_{i+1}) p(\sigma_v, v_i, f_i|o_1, \dots, o_i) \quad (\text{A5})$$

After the first step, $p(\sigma_v, v_{i+1}, f_{i+1}|o_1, \dots, o_i)$ was divided by its sum across σ_v, v_{i+1} and f_{i+1} so that the cells in $p(\sigma_v, v_{i+1}, f_{i+1}|o_1, \dots, o_i)$ summed to 1. The marginalization was done by collapsing the other dimensions. The mean was approximated using a weighed sum, $\sum xp(x)$.

In the second step (i.e., after the congruency of trial $i+1$ is observed), the belief of variables is updated using the observed congruency in the following manner:

$$p(\sigma_v, v_{i+1}, f_{i+1}|o_1, \dots, o_{i+1}) = p(\sigma_v, v_{i+1}, f_{i+1}|o_1, \dots, o_i) p(f_{i+1}|o_{i+1}) \quad (\text{A6})$$

where

$$p(f_{i+1}|o_{i+1}) = \begin{cases} 1 - f_{i+1}, & \text{if } o_{i+1} \text{ is incongruent} \\ f_{i+1}, & \text{if } o_{i+1} \text{ is congruent} \end{cases} \quad (\text{A7})$$

is the prediction error between the predicted conflict level and the true congruency.

After feeding the trial sequences experienced by each participant to the Bayesian model, we obtained estimates of volatility and predicted conflict level for each trial (i.e., the estimates from step 1). Then the trial-wise estimates were grouped based on experimental conditions to compute condition-specific means, which were further averaged across participants to produce group-level means of volatility and predicted conflict level for each experimental condition. These group-level means of parameter estimates were then employed to simulate group-level mean RTs using a linear model. For condition j , the linear model states:

$$RT_j = c_j(\beta_1 v_j + \beta_2 f_j + \beta_3) + (1 - C_j)(\beta_4 v_j + \beta_5 f_j + \beta_6) \quad (\text{A8})$$

where v_j and f_j are the group mean of volatility and predicted conflict level for condition j , respectively. $C_j = 0$ if condition j is congruent; and $C_j = 1$ if condition j is incongruent. After obtaining the coefficients for the linear model (i.e., the betas) using linear regression across all conditions, we calculated the condition-specific simulated RTs using the coefficients:

$$\widetilde{RT}_j = C_j(\beta_1 v_j + \beta_2 f_j + \beta_3) + (1 - C_j)(\beta_4 v_j + \beta_5 f_j + \beta_6) \quad (\text{A9})$$

where \widetilde{RT}_j is the simulated RT for condition j .

As mentioned before, our definition of v would result in a counterintuitive representation of having a lower v in a more volatile setting. Thus, to avoid confusion, we transformed the condition-specific volatility estimates before presenting them:

$$\widetilde{v}_j = \max(v) + \min(v) - v_j \quad (\text{A10})$$

where $\max(v)$ and $\min(v)$ are the maximum and minimum group-level mean of volatility estimates across all conditions, respectively. After this transform, more volatile conditions have higher \widetilde{v}_j s. Note that because all v_j s were transformed using the same constants, this transform had no effect on \widetilde{RT}_j s and hence our simulation results.

References

Alexander, W.H., Brown, J.W., 2011. Medial prefrontal cortex as an action-outcome predictor. *Nat. Neurosci.* 14, 1338–1344.
 Amiez, C., Joseph, J.P., Procyk, E., 2006. Reward encoding in the monkey anterior cingulate cortex. *Cereb. Cortex* 16, 1040–1055.
 Bach, D.R., Dolan, R.J., 2012. Knowing how much you don't know: a neural organization of uncertainty estimates. *Nat. Rev. Neurosci.* 13, 572–586.
 Badre, D., 2008. Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends Cogn. Sci.* 12, 193–200.
 Balleine, B.W., Dickinson, A., 1998. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology* 37, 407–419.
 Barch, D.M., Braver, T.S., Akbudak, E., Conturo, T., Ollinger, J., Snyder, A., 2001. Anterior cingulate cortex and response conflict: effects of response modality and processing domain. *Cereb. Cortex* 11, 837–848.
 Behrens, T.E., Woolrich, M.W., Walton, M.E., Rushworth, M.F., 2007. Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221.
 Blais, C., Robidoux, S., Risko, E.F., Besner, D., 2007. Item-specific adaptation and the conflict-monitoring hypothesis: a computational model. *Psychol. Rev.* 114, 1076–1086.
 Botvinick, M., Nystrom, L.E., Fissell, K., Carter, C.S., Cohen, J.D., 1999. Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature* 402, 179–181.

- Botvinick, M.M., Braver, T.S., Barch, D.M., Carter, C.S., Cohen, J.D., 2001. Conflict monitoring and cognitive control. *Psychol. Rev.* 108, 624–652.
- Botvinick, M.M., Cohen, J.D., Carter, C.S., 2004. Conflict monitoring and anterior cingulate cortex: an update. *Trends Cogn. Sci.* 8, 539–546.
- Bowers, J.S., Davis, C.J., 2012. Bayesian just-so stories in psychology and neuroscience. *Psychol. Bull.* 138, 389–414.
- Braver, T.S., 2012. The variable nature of cognitive control: a dual mechanisms framework. *Trends Cogn. Sci.* 16, 106–113.
- Braver, T.S., Barch, D.M., 2002. A theory of cognitive control, aging cognition, and neuromodulation. *Neurosci. Biobehav. Rev.* 26, 809–817.
- Braver, T.S., Gray, J.R., Burgess, G.C., 2007. Explaining the many varieties of working memory variation: dual mechanisms of cognitive control. In: Conway, A., Jarrold, C., Kane, M., Miyake, A., Towse, J. (Eds.), *Variation in Working Memory*. Oxford University Press, Oxford, pp. 76–106.
- Bugg, J.M., 2012. Dissociating levels of cognitive control: the case of Stroop interference. *Psychol. Sci.* 21, 302–309.
- Bugg, J.M., Chanani, S., 2011. List-wide control is not entirely elusive: evidence from picture-word Stroop. *Psychon. Bull. Rev.* 18, 930–936.
- Bugg, J.M., Crump, M.J., 2012. In support of a distinction between voluntary and stimulus-driven control: a review of the literature on proportion congruent effects. *Front. Psychol.* 3, 367.
- Bugg, J.M., Hutchison, K.A., 2013. Converging evidence for control of color-word Stroop interference at the item level. *J. Exp. Psychol. Hum. Percept. Perform.* 39, 433–449.
- Carter, C.S., Braver, T.S., Barch, D.M., Botvinick, M.M., Noll, D., Cohen, J.D., 1998. Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science* 280, 747–749.
- Cohen, J.D., Dunbar, K., McClelland, J.L., 1990. On the control of automatic processes: a parallel distributed processing account of the Stroop effect. *Psychol. Rev.* 97, 332–361.
- De Martino, B., Fleming, S.M., Garrett, N., Dolan, R.J., 2013. Confidence in value-based choice. *Nat. Neurosci.* 16, 105–110.
- De Pisapia, N., Braver, T.S., 2006. A model of dual control mechanisms through anterior cingulate and prefrontal cortex interactions. *Neurocomputing* 69, 1322–1326.
- den Ouden, H.E., Daunizeau, J., Roiser, J., Friston, K.J., Stephan, K.E., 2010. Striatal prediction error modulates cortical coupling. *J. Neurosci.: Off. J. Soc. Neurosci.* 30, 3210–3219.
- Desimone, R., Duncan, J., 1995. Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222.
- Duncan, J., 2001. An adaptive coding model of neural function in prefrontal cortex. *Nat. Rev. Neurosci.* 2, 820–829.
- Egner, T., 2007. Congruency sequence effects and cognitive control. *Cogn. Affect. Behav. Neurosci.* 7, 380–390.
- Egner, T., Ely, S., Grinband, J., 2010. Going, going, gone: characterizing the time-course of congruency sequence effects. *Front. Psychol.* 1, 154.
- Egner, T., Etkin, A., Gale, S., Hirsch, J., 2008. Dissociable neural systems resolve conflict from emotional versus nonemotional distracters. *Cereb. Cortex* 18, 1475–1484.
- Egner, T., Hirsch, J., 2005. Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nat. Neurosci.* 8, 1784–1790.
- Eriksen, B.A., Eriksen, C.W., 1974. Effects of noise letters upon identification of a target letter in a non-search task. *Percept Psychophys* 16, 143–149.
- Fiorillo, C.D., Tobler, P.N., Schultz, W., 2003. Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299, 1898–1902.
- Fiorillo, C.D., Tobler, P.N., Schultz, W., 2005. Evidence that the delay-period activity of dopamine neurons corresponds to reward uncertainty rather than backpropagating TD errors. *Behav. Brain Funct.: BBF* 1, 7.
- Fuster, J.M., 2008. *The Prefrontal Cortex*. Academic Press, London.
- Gratton, G., Coles, M.G., Donchin, E., 1992. Optimizing the use of information: strategic control of activation of responses. *J. Exp. Psychol. Gen.* 121, 480–506.
- Gratton, G., Coles, M.G., Sirevaag, E.J., Eriksen, C.W., Donchin, E., 1988. Pre- and post-stimulus activation of response channels: a psychophysiological analysis. *J. Exp. Psychol. Hum. Percept. Perform.* 14, 331–344.
- Hommel, B., Proctor, R.W., Vu, K.P., 2004. A feature-integration account of sequential effects in the Simon task. *Psychol. Res.* 68, 1–17.
- Hopfield, J.J., 1982. Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. Biol.* 79, 2554–2558.
- Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., Camerer, C.F., 2005. Neural systems responding to degrees of uncertainty in human decision-making. *Science* 310, 1680–1683.
- Huettel, S.A., Stowe, C.J., Gordon, E.M., Warner, B.T., Platt, M.L., 2006. Neural signatures of economic preferences for risk and ambiguity. *Neuron* 49, 765–775.
- Ide, J.S., Shenoy, P., Yu, A.J., Li, C.S., 2013. Bayesian prediction and evaluation in the anterior cingulate cortex. *J. Neurosci.: Off. J. Soc. Neurosci.* 33, 2039–2047.
- Johan Masreliez, C., Martin, R.D., 1977. Robust Bayesian estimation for the linear model and robustifying the Kalman filter. *IEEE Trans. Automatic Control* 22, 361–371.
- Jones, M., Curran, T., Mozer, M.C., Wilder, M.H., 2013. Sequential effects in response time reveal learning mechanisms and event representations. *Psychol. Rev.* 120, 628–666.
- Kerns, J.G., Cohen, J.D., MacDonald 3rd, A.W., Cho, R.Y., Stenger, V.A., Carter, C.S., 2004. Anterior cingulate conflict monitoring and adjustments in control. *Science* 303, 1023–1026.
- Kiani, R., Shadlen, M.N., 2009. Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* 324, 759–764.
- 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. *J. Neurosci.: Off. J. Soc. Neurosci.* 30, 12759–12769.
- Knutson, B., Taylor, J., Kaufman, M., Peterson, R., Glover, G., 2005. Distributed neural representation of expected value. *J. Neurosci.: Off. J. Soc. Neurosci.* 25, 4806–4812.
- Koechlin, E., Ody, C., Kouneiher, F., 2003. The architecture of cognitive control in the human prefrontal cortex. *Science* 302, 1181–1185.
- Koller, D., Friedman, N., 2009. *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press.
- Kording, K.P., Tenenbaum, J.B., Shadmehr, R., 2007. The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat. Neurosci.* 10, 779–786.
- Lee, T.S., Mumford, D., 2003. Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A: Opt. Image Sci. Vision* 20, 1434–1448.
- Liu, X., Banich, M.T., Jacobson, B.L., Tanabe, J.L., 2004. Common and distinct neural substrates of attentional control in an integrated Simon and spatial Stroop task as assessed by event-related fMRI. *Neuroimage* 22, 1097–1106.
- Logan, G.D., Zbrodoff, N.J., 1979. When it helps to be misled: facilitative effects of increasing the frequency of conflicting stimuli in a Stroop-like task. *Mem. Cognit.* 7, 166–174.
- MacDonald 3rd, A.W., Cohen, J.D., Stenger, V.A., Carter, C.S., 2000. Dissociating the role of the dorsolateral prefrontal and anterior cingulate cortex in cognitive control. *Science* 288, 1835–1838.
- MacLeod, C.M., 1991. Half a century of research on the Stroop effect: an integrative review. *Psychol. Bull.* 109, 163–203.
- MacLeod, C.M., MacDonald, P.A., 2000. Interdimensional interference in the Stroop effect: uncovering the cognitive and neural anatomy of attention. *Trends Cogn. Sci.* 4, 383–391.
- Matsumoto, K., Suzuki, W., Tanaka, K., 2003. Neuronal correlates of goal-based motor selection in the prefrontal cortex. *Science* 301, 229–232.
- Mayr, U., Awh, E., Laury, P., 2003. Conflict adaptation effects in the absence of executive control. *Nat. Neurosci.* 6, 450–452.
- McClure, S.M., Laibson, D.I., Loewenstein, G., Cohen, J.D., 2004. Separate neural systems value immediate and delayed monetary rewards. *Science* 306, 503–507.
- McCoy, A.N., Platt, M.L., 2005. Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat. Neurosci.* 8, 1220–1227.
- Miller, E.K., Cohen, J.D., 2001. An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 24, 167–202.
- Monosov, I.E., Hikosaka, O., 2013. Selective and graded coding of reward uncertainty by neurons in the primate anterodorsal septal region. *Nat. Neurosci.* 16, 756–762.
- Mozer, M.C., Colagrosso, M.D., Huber, D.E., 2002. A rational analysis of cognitive control in a speeded discrimination task. *Advances in Neural Information Processing Systems 14*, vols. 1 and 2, pp. 51–57.
- Norman, D.A., Shallice, T., 1986. Attention to action: willed and automatic control of behavior. In: Schwarz, G.E., Shapiro, D. (Eds.), *Consciousness and Self-regulation*. Plenum Press, New York.
- Pearce, J.M., Hall, G., 1980. A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychol. Rev.* 87, 532–552.
- Pearl, J., 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, San Mateo, CA.
- Platt, M.L., Huettel, S.A., 2008. Risky business: the neuroeconomics of decision making under uncertainty. *Nat. Neurosci.* 11, 398–403.
- Preuschoff, K., Bossaerts, P., Quartz, S.R., 2006. Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* 51, 381–390.
- Rescorla, R.A., Wagner, A.R., 1972. A theory of Pavlovian conditioning: variations in the effectiveness of reinforcement and nonreinforcement. In: Black, A.H., Prokasy, W.F. (Eds.), *Classical Conditioning II*. Appleton-Century-Crofts, New York, pp. 64–99.
- Reynolds, J., Mozer, M., 2009. Temporal dynamics of cognitive control. In: Koller, D., Schuurmans, D., Bengio, Y., Bottou, L. (Eds.), *Advances in Neural Information Processing Systems*, pp. 1353–1360.
- Ridderinkhof, K.R., Ullsperger, M., Crone, E.A., Nieuwenhuis, S., 2004. The role of the medial frontal cortex in cognitive control. *Science* 306, 443–447.
- Rushworth, M.F., Behrens, T.E., 2008. Choice, uncertainty and value in prefrontal and cingulate cortex. *Nat. Neurosci.* 11, 389–397.
- Schmajuk, N.A., Lam, Y.W., Gray, J.A., 1996. Latent inhibition: a neural network approach. *J. Exp. Psychol.: Anim. Behav. Process.* 22, 321–349.
- Schmidt, J.R., Besner, D., 2008. The Stroop effect: why proportion congruent has nothing to do with congruency and everything to do with contingency. *J. Exp. Psychol.: Learn. Mem. Cognit.* 34, 514–523.
- Servan-Schreiber, D., Bruno, R.M., Carter, C.S., Cohen, J.D., 1998. Dopamine and the mechanisms of cognition: Part I. A neural network model predicting dopamine effects on selective attention. *Biol. Psychiatry* 43, 713–722.
- Shenoy, P., Rao, R., Yu, A.J., 2010. A rational decision making framework for inhibitory control. In: Lafferty, J., Williams, C.K.I., Shawe-Taylor, J., Zemel, R.S., Culotta, A. (Eds.), *Advances in Neural Information Processing Systems*. MIT, Boston, pp. 2146–2154.
- Shenoy, P., Yu, A.J., 2011. Rational decision-making in inhibitory control. *Front. Hum. Neurosci.* 5, 48.
- Silvetti, M., Seurinck, R., Verguts, T., 2011. Value and prediction error in medial frontal cortex: integrating the single-unit and systems levels of analysis. *Front. Hum. Neurosci.* 5, 75.

- Silvetti, M., Seurinck, R., Verguts, T., 2013. Value and prediction error estimation account for volatility effects in ACC: a model-based fMRI study. *Cortex: J. Devoted Study Nerv. Syst. Behav.* 49, 1627–1635.
- Summerfield, C., Behrens, T.E., Koechlin, E., 2011. Perceptual classification in a rapidly changing environment. *Neuron* 71, 725–736.
- Tenenbaum, J.B., Griffiths, T.L., Kemp, C., 2006. Theory-based Bayesian models of inductive learning and reasoning. *Trends Cogn. Sci.* 10, 309–318.
- Tenenbaum, J.B., Xu, F., 2000. Word learning as Bayesian inference. In: *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*, pp. 517–522.
- Torres-Quesada, M., Funes, M.J., Lupianez, J., 2013. Dissociating proportion congruent and conflict adaptation effects in a Simon–Stroop procedure. *Acta Psychol. (Amst.)* 142, 203–210.
- Tzelgov, J., Henik, A., Berger, J., 1992. Controlling Stroop effects by manipulating expectations for color words. *Mem. Cognit.* 20, 727–735.
- Verguts, T., Notebaert, W., 2008. Hebbian learning of cognitive control: dealing with specific and nonspecific adaptation. *Psychol. Rev.* 115, 518–525.
- Verguts, T., Notebaert, W., 2009. Adaptation by binding: a learning account of cognitive control. *Trends Cogn. Sci.* 13, 252–257.
- Vilares, I., Kording, K., 2011. Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Ann. N.Y. Acad. Sci.* 1224, 22–39.
- Volz, K.G., Schubotz, R.I., von Cramon, D.Y., 2003. Predicting events of varying probability: uncertainty investigated by fMRI. *Neuroimage* 19, 271–280.
- Volz, K.G., Schubotz, R.I., von Cramon, D.Y., 2004. Why am I unsure? Internal and external attributions of uncertainty dissociated by fMRI. *Neuroimage* 21, 848–857.
- Vossel, S., Mathys, C., Daunizeau, J., Bauer, M., Driver, J., Friston, K.J., Stephan, K.E., 2014. Spatial attention, precision, and bayesian inference: a study of saccadic response speed. *Cereb. Cortex*, 1436–1450.
- White, C.N., Brown, S., Ratcliff, R., 2012. A test of Bayesian observer models of processing in the Eriksen flanker task. *J. Exp. Psychol. Hum. Percept. Perform.* 38, 489–497.
- White, C.N., Ratcliff, R., Starns, J.J., 2011. Diffusion models of the flanker task: discrete versus gradual attentional selection. *Cogn. Psychol.* 63, 210–238.
- Wittfoth, M., Buck, D., Fahle, M., Herrmann, M., 2006. Comparison of two Simon tasks: neuronal correlates of conflict resolution based on coherent motion perception. *Neuroimage* 32, 921–929.
- Xu, F., Tenenbaum, J.B., 2007. Word learning as Bayesian inference. *Psychol. Rev.* 114, 245–272.
- Yu, A.J., Cohen, J.D., 2009. Sequential Effects: Superstition or Rational Behavior? *Advances in Neural Information Processing Systems*. MIT Press, Cambridge, MA, pp. 1873–1880.
- Yu, A.J., Dayan, P., 2005. Uncertainty, neuromodulation, and attention. *Neuron* 46, 681–692.
- Yu, A.J., Dayan, P., Cohen, J.D., 2009. Dynamics of attentional selection under conflict: toward a rational Bayesian account. *J. Exp. Psychol. Hum. Percept. Perform.* 35, 700–717.