

Free Choice and Voluntary Action

Abstract

A preliminary to any valid theory of voluntary action is a conceptual framework that permits it to be tested empirically. Where deficits in the conceptual framework make this impossible, the empirical data become uninterpretable. Here I show that “free choice” and “conflict” tasks exhibit such deficits, casting doubt on the testability of any theory that depends on them. I suggest that a reevaluation of the tasks used to study voluntary action is necessary.

Parashkev Nachev¹

*Institute of Neurology
University College London*

Few problems in neuroscience are as challenging as the biological nature of free will. For to say that an action is freely chosen is to say that no-one but oneself is its author: To give an explanation of the former would therefore seem to require an explanation of the latter. And if the nature of indexicals, the domain to which questions of self-ascription belong, is hard enough for a philosopher to make sense of, it is doubly so for the scientist, who not only has to get the conceptualization right but also has to find a felicitous way of building an empirical edifice on it.

Here I do not wish to examine the deeper conceptual difficulties, which is not to suggest that a scientist can afford to ignore such matters. The distinction between sense and nonsense—the business of philosophy—naturally precedes the distinction between truth and falsehood, the business of science. Something that makes sense can be true or false, something that makes no sense can be neither. But the conceptual constraints on empirical discovery are matters best addressed by philosophers, not scientists.

A constraint that is within the competence of a scientist, however, is what can and cannot be tested by experiment. Indeed, it would seem the duty of every scientist is to establish the boundaries of the empirical space he needs to explore, and to determine his capacity to reach them. If large parts of the space are inaccessible for one reason or another, the picture he will obtain will inevitably be incomplete. How incomplete he may never know, for he may also have no knowledge of how much is obscured. It may be that what is missing is precisely the critical part without which a synoptic view is impossible.

¹ Queen Square, London, WC1N 3BG, UK
p.nachev@ion.ucl.ac.uk

Let us therefore consider how much of the boundary of the empirical space in relation to the neuroscience of voluntary action is visible to us, given the experimental tools we have at our disposal. If we can see all of it, then we know that we cannot be misled in the long run, if we cannot, then we have to reconsider our approach to the problem. Specifically, we must see what we can do to constrain the space so as to bring its boundaries into view, and, therefore, to make it surveyable.

Empirical Boundaries

To be clear, by empirical space I mean the totality of empirically determinable possibilities of the physical states-of-affairs under investigation. To use a simple biological example, the link between a hormone and a target organ has to be mediated by a receptor molecule that is present in the organ and has the right binding properties: These are some obvious boundaries on the empirical space a researcher interested in finding the receptor has to explore. Hormone/receptor pair possibilities are generally confined within a relatively narrow space since biological receptors are highly stereotyped: They are usually proteins; they usually have membrane-binding elements, and so forth. The problem, from the point of view of the size of the space of possibilities, is thus relatively straightforward.

Now consider the form of empirical enquiry here (Figure 1). We are trying to differentiate between a limited series of models (e.g. the receptor may be a membrane-bound protein, a nuclear-binding protein, etc.) We construct an experiment whose outcome allows us to make a differentiation: to strengthen one model and/or weaken the others. This evidence is then used to update our belief about which model is closest to the truth. Since the space of empirically determinable models is relatively small, constructing an experiment that will usefully discriminate between them is easy, or at least feasible.

When we attempt to apply the same process to the question of the neural basis of voluntary action (or indeed most other serious questions in cognitive neuroscience), we are immediately in difficulty, for the space of possible models is absolutely enormous here (Figure 2). Even a conservative estimate of the number of neurons in the brain (e.g. Pakkenberg & Gundersen, 1997) and their interconnections produces an infeasibly large number of possible physiological models. Clearly, the empirical space here will not be surveyable unless we can find a way of constraining it. But how can we legitimately do this?

One way is to look at what data we already have. Known facts about the brain may allow us to eliminate a subsection of possibilities. For example, although the brain exhibits a high level of interconnectivity, we know connections are both spatially inhomogeneous and similarly so across individuals: Clearly, the anatomical structure will exclude models that are, say, uniformly distributed.

Unfortunately, our knowledge of the underlying physiology is extremely sparse. The two major tools in the cognitive neuroscientists' laboratory, functional neuroimaging and single cell recordings, are only readily applicable at two extremes of the anatomical scale: the macroscopic and the microscopic. What might be called the mesoscopic level of architecture in the motor brain, the level corresponding to the columnar organization in primary visual cortex, is rarely explored. It is perhaps

unsurprising that neurophysiological models of action are usually either macroscopic or microscopic at their level of explanation; clearly, the reason for this may not be the data but the bias in our sampling of it. Indeed, given what we know about the visual cortex, it would seem likely that by ignoring the mesoscopic scale our models can only be impoverished.

Moreover, most neurophysiological data is correlational: It, therefore, does not and cannot establish the necessity of a neuronal substrate for a given function. This serious objection is often brushed aside by an appeal to parsimony: Why would a neuronal substrate be active if it were not required in a given set of circumstances? There are at least two reasons why such objections are weak. First, we do not know what aspect of a task correlated brain activity is necessary for: It may have nothing to do with our conceptualization of what the task involves. For example, enhanced activity in the cingulate when performing a difficult task compared with an easy one is routinely taken to imply a cognitive function (e.g. Botvinick, 2007) even though the same contrast can be obtained when the difficulty is non-cognitive, and the correlation is actually better explained by the enhanced autonomic activation that inevitably accompanies difficult tasks (Critchley et al., 2003). Second, neurons in the brain do not “know” what the subject is doing because knowledge of the task is a property of the subject, not of any isolated part of his brain. A great deal of the activation we observe in any particular circumstances may be related to what the subject **could** be doing (e.g. Cisek & Kalaska, 2002) but is not, perhaps because of suppressive activity from elsewhere. In sum, although experimental data are helpful to some extent, they do not constrain the space of possible physiological models to a sufficient degree.

In fact, such constraints as empirical data provide grounds for themselves dependent on a unique feature of our experimental domain: the requirement for a behavioural task. Indeed, it is generally thought that the task provides most of the constraints that model building in behavioral neuroscience requires. We, therefore, need to examine the extent to which the tasks used to probe the neural basis of free will do indeed constrain the empirical space sufficiently for the models arising from them to have any value. It is to this question that we now turn.

The Nature of Task Constraints

Experiments in behavioral neuroscience differ from experiments in other branches of science in that they involve the performance of a behavioral task of some kind. The task generally involves contrasting behavior under two or more conditions assumed to be dissociable in neurophysiological terms. I say assumed because if an experiment demonstrates a neurophysiological dissociation, it is taken to be a confirmation of the dissociation in reality. If it does not, it is generally taken to imply an imperfection in the empirical design (e.g. signal-to-noise) rather than in the conceptualization of the task on which the design is based. Moreover, the task construction has to come from somewhere other than the experiment itself, for it is what the experiment is predicated on.

Now a great deal can be hidden in this seemingly innocuous step. For to say that someone performs a movement (say) in the absence of external stimuli is not like saying that a chemical reaction (say) occurs in the absence of oxygen. The task manipulation may radically alter what is plausibly argued to be going on in ways

remote from the aims of the manipulation itself. Yet, remarkably little space is devoted to the study of tasks and what their features imply, certainly not beyond the obvious points such as the artificiality and over-rehearsal of classical behavioral paradigms. Critically, often very little is said about what effect, other than that due to the intended factor, a task manipulation might involve. If such confounds exist, and their presence extinguishes the constraining power of the task, little if anything can be made of the empirical data produced from it.

We should, therefore, consider whether the tasks commonly employed to study voluntary action do indeed offer the constraints on which any empirical inferences from them depend. We start by considering the ways in which voluntary action is conventionally parameterized.

Voluntary Action Tasks

We should first consider what aspects of an action we can manipulate, and second, the kinds of manipulation thought to isolate its distinctively voluntary forms.

Condition-action Associations

Voluntary action has two aspects: first, the action itself, and second, the circumstances that attend it. Clearly pressing a button and pulling a lever in response to a visual cue are two different actions, even if their prompt, the visual cue, is the same. Similarly, pressing a button in response to a visual cue and pressing a button in response to an auditory cue are two different actions even if the mechanics of the movement are the same. A minimal parameterization of voluntary action, therefore, has to specify both the conditions for an action and the action itself, in short, the condition-action association.

Note that by “condition” here, we do not need to confine ourselves to an external stimulus. Anything that can materially affect the neural activity underlying an action is plausibly a condition of it, irrespective of whether it is external, such as a visual cue, or internal, such as a desired goal or the passage of time.

Isolating Voluntary Action

So characterized, let us see how we can manipulate a set of condition-action associations in order to extract what is most distinctive about voluntary action. Here, three theoretical assumptions are commonly made, two of them uncontroversial (Hacker, 2007; Passingham, 1993).

First, as we have said already, for an action to be voluntary it must be possible to say that the agent is its author. The movements of an alien limb are not voluntary because their agent, by definition, disowns them; the same is true of partial seizures. An identical movement, for example, in choreiform movement disorders, may be designated as voluntary or involuntary depending on whether the subject reports believing himself the owner. This assumption is conceptually unobjectionable.

Second, an action may be voluntary only if it may be said to have been chosen. Where no alternative exists, for example, in the movement features of a deep tendon

reflex, the only voluntariness we can speak of is that implicit in being able to withhold it. This assumption is also unobjectionable.

Third, it is argued that an action is voluntary primarily, or even only, if it may be said to have been intended. Intention is conceived as an event that precedes and determines the nature of the action: either directly or by virtue of its temporal co-localization. This assumption is conceptually objectionable; here we shall not deal with it other than to point to the paradigms that draw upon it.

These assumptions allow us to compare circumstances of supposedly differing voluntariness: what we may call volitional contrasts. The first two of these form the basis of free choice and conflict tasks, the last of Libet-style dissociations of intention and action (Figure 3). We shall discuss only the first two here.

Free Choice

The Paradigm

If a distinctive feature of a voluntary act is that no-one but the agent is its owner, voluntary action should be most “volitional” when neither the conditions nor the action itself is specified by anything but the agent executing it: in other words, when purely “free choice” obtains. Let us see how such “free-choice” paradigms are implemented.

With regard to the action itself, free choice is difficult to parameterize. If the distinctive feature of a movement is that it is maximally free, one would not expect it to be a repetition of any movement in the past, still less a copy of any movement the experimenter has shown to the subject. Unfortunately, neither success nor failure in “freeing” such an action would be very easy to gauge in any particular circumstances. Free choice paradigms, therefore, do not manipulate the action itself, but keep that aspect of the condition-action association highly stereotyped: a choice of button presses, or joystick twists, and so forth.

The conditions, however, are readily amenable to manipulation. We can impoverish the environment of sensory cues such that nothing immediately biases responding towards one action or another (Deiber et al., 1991; Frith, Friston, Liddle, & Frackowiak, 1991; Jahanshahi et al., 1995; Jenkins, Jahanshahi, Jueptner, Passingham, & Brooks, 2000). Since something has to determine the choice between two equally weighted movements, if nothing in the external environment does, it has to come from the agent himself. Free choice tasks therefore typically ask the subject to choose “freely” from two action alternatives in circumstances where neither is externally prompted in any shape or form.

The contrast of such “free choice” with situations where the choice is externally specified is then supposed to be reflected in differential engagement of the mechanism on which voluntary action depends. Crudely speaking, the substrate engaged by “free choice” is thought as being closer to the substrate of voluntary action than that engaged by “directed choice”. In short, the free choice paradigm is argued to expose a demarcation line in the empirical model space: between, on the

one hand, neural substrates that deal with voluntary actions and, on the other, those that deal with actions less so.²

The Confound

The critical move here is the assumption that the task manipulation, between externally-guided and freely-chosen actions, is the only material thing that changes. If this assumption is incorrect, and the task boundary may not be reliably translated into a physiological boundary, our model space remains wide open. Let us, therefore, see if it is so.

If someone attempted to contrast a single finger movement with a performance of Ravel's Piano Concerto for the Left Hand in D major, we would naturally judge the contrast imbalanced, and, therefore, of limited value: One simple movement cannot be compared with a multitude, delivered in a complex sequence. Now if this is true of the action itself it is also true of the conditions that prompt it. Thus the contrast between "free choice" conditions and "directed choice" conditions works only if nothing but the freedom to choose is contrasted. But this is clearly not so, as we shall see.

First, when one makes the choice "free" one does not make it unconditioned. Something has to determine the subject's decision, by asking the subject to choose freely we are merely asking him not to give us an explicit criterion: in essence, to conceal (perhaps even from himself) the reason for his choice. We know that subjects cannot respond randomly (Brugger, 1997; Tune, 1964; Wagenaar, 1972); indeed, pseudorandom responding has a behavioral cost. "Free" responding here, therefore, means following a complex pseudorandomization schedule. Now some experimenters attempt to circumvent this problem by asking the subject not to take his response history into account. But this is rather like asking someone actively not to think about something, not to think about not thinking about it, and so on, recurrently. Others do not instruct subjects not to pseudorandomize, and instead simply drop those who show evidence of it in their behavior (e.g. Soon, Brass, Heinze, & Haynes, 2008): This clearly amounts to much the same thing because the subjects who are included are simply selected for naturally interpreting the instructions so and pseudorandomizing well.

In any event, whether or not the response is consciously pseudorandomized, we still have the condition part of the condition-action association here that is in need of explanation. We know the choice is never stochastic, in the sense in which radioactive decay is stochastic, so something must bias it in one direction or another at each trial. We must therefore determine the plausible sources of this bias and how they might differ in comparison with the directed condition.

In any circumstance, the behavior exhibited at any one time must be the outcome of a vast range of biases, some of which are more heavily weighted than others. If we conceive this as competition between neuronal ensembles encoding various condition-action associations (though it does not matter much to our argument whether we do or not) (e.g. Cisek, 2007), we might imagine that some ensembles are stimulated and others are suppressed, with the behavioral outcome reflecting the most stimulated, winning ensemble. In the directed condition, two ensembles may be

² It is worth noting that some researchers draw the distinction between externally and internally triggered action without taking this contrast to imply a differences in the voluntariness of the action concerned. Clearly, this does not vitiate the criticisms of the contrast itself.

thought to be dominant: that encoding making (say) a left movement in response to a left arrow and another with the converse association. Clearly, other ensembles will also be in competition (e.g. to get out of the experimental chair and do something more rewarding), but their weighting will be lower (until some other event happens, such as a signal for the end of the experiment).

In the “free” condition, however, the range of ensembles likely to be co-activated is undefined: deliberately so. Moreover, we have every reason to believe that it is much larger than in the directed condition, for no readily identifiable set of stimuli or conditions are dominant, and it is also more complex because a simple rule is precisely what the subject is encouraged (or selected) to avoid. In any case, we have no grounds whatever for believing that the characteristics of the conditions are matched in the “free” and “directed” settings (Nachev, Kennard, & Husain, 2008; Nachev, Wydell, O'Neill, Husain, & Kennard, 2006).

The notion that the free/directed conceptual plane must have a neurophysiologic reflection is, therefore, false. Data derived from experiments that employ free choice paradigms to this end are, therefore, uninterpretable because the constraint on the empirical space imposed by the task is illusory. The neurophysiological differences could just as well be explained by differences in the range of co-activated ensembles or their complexity: We have no way of knowing which, since neither of these things is parameterized in the free choice task, indeed, its very essence is to leave them open.

This criticism is not novel. It has been suggested that because responding in free choice is “under-determined,” there is enhanced conflict neuronal populations encoding response alternatives, and, therefore, enhanced recruitment of a putative system for detecting or resolving response conflict (Botvinick, Braver, Barch, Carter, & Cohen, 2001). Since the areas of the brain activated by conflict and by free choice are closely overlapping, this seems a plausible suggestion. But the alternative it offers, a partial explanation of the nature of free choice in terms of conflict, also lacks sufficient constraint. We turn to this question next.

Conflict

The Paradigm

We have seen that choice is central to the notion of voluntary action: An action is voluntary only if it may be said to have been chosen. But such a choice need not be balanced: One of the action alternatives may be considered less open to choice than another. For example, making a reflexive saccade to a sudden peripheral onset may be thought to involve less choice than making a saccade to a static target when prompted by a central cue. This is because in the former case the parameters of the action, the gain, and direction of the saccade are, arguably, not chosen but determined by the parameters of the peripheral stimulus: This is what a reflexive saccade is. The only thing that appears to be under the control of the subject is whether or not the saccade is executed at all. Other actions, such as sneezing or yawning, exhibit this feature even more strikingly: One cannot make oneself sneeze or yawn, but one can suppress it. A key component of voluntary action is, therefore, widely thought to be the capacity to suppress automatic, or at least less sophisticated, response tendencies;

a capacity often conceptualized as (a part of) “executive control” (Logan, 1985; Norman & Shallice, 2000; Schall, Stuphorn, & Brown, 2002). Clearly, such suppression need only be invoked where the “primitive” response gets in the way of the more voluntary one: in other words, where there is *conflict* between responses. This notion has given rise to a class of task said to generate conflict.

Although conflict tasks vary widely in their implementation, they have two features in common. First, they contrast a more natural or habitual action with one less so. For example, the classic Stroop task (Stroop, 1935) contrasts reading a color word (which is natural) with reporting the color of the ink it is written in (which is not). Second, in order to generate conflict it is necessary for the conditions for each of the actions being conflicted to overlap to some degree: that is, for their features to be shared at least to some degree. Thus, in the Stroop task, a color is always named on viewing a word; the difference is the particular aspect of the word attended to: either the word itself or the color of the typeface. It is obviously necessary for the rival condition-action association to be stimulated if any conflict is to occur. Without conditional overlap (even if only by virtue of performing the trials within the same block), there cannot be any conflict to speak of.

The contrast such tasks purportedly extract, then, is between situations of high response conflict and situations where no or low response conflict is present: The former is said to engage executive control to a greater extent than the latter. Once again, we have to consider whether or not the contrast is a valid one: that is, whether or not the conceptual division has an inevitable neurophysiological correlate. And once again, all depends on what else the contrast extracts other than any supposed necessity for executive control.

The Confound

We have already seen that in order to generate a conflict contrast at least two things are incidentally manipulated: the degree to which one response is more natural or habitual, and the degree of conditional overlap. Neither of these things is easily corrected for, even if one recognizes their significance, as we shall see.

A difference in naturalness or habit is difficult to parameterize because it is defined by experience outside the experimental setting that is hard to quantify. Moreover, since the degree of conflict is directly related to this difference, it cannot be dissociated from it by any behavioral manipulation. This is true even for supposedly more balanced conflict tasks such as the Eriksen flanker task (Eriksen & Eriksen, 1974) or the change task (Husain, Parton, Hodgson, Mort, & Rees, 2003; Logan, 1994), where subjects are asked to follow one of a set of congruent or incongruent arrows. Arrows are ecologically presented in congruent fashion; to point in one direction is what an arrow is generally used for. Indeed, if sets of arrows in the real world were as commonly incongruent as congruent, no conflict would be produced at all, or the conflict would not be between outcomes that differ with respect to the supposed aims of the “executive”: to suppress more automatic responses in favor of less automatic ones.

Differences in conditional overlap can be parameterized, but not independently of the conflict itself. Any task with conditional overlap will always

generate conflict: To the extent to which the conditions are overlapping, their associated condition-action associations will be co-activated, and if they are incompatible, then conflict will inevitably arise. If they are compatible, then the conditional overlap will count for nothing: a non-conflicting task with conditional overlap is, in every sense, redundant.

Far from isolating conflict, such contrasts are therefore inevitably contaminated by ineliminable confounds. The data derived from conflict experiments cannot be framed within a conflict/no conflict dichotomy: the constraint on which any model building here depends.

Counter-arguments

Let us pause to summarize the logic of what we have said. To be in a position to argue in favor of a physiological model of voluntary action, we need to have some constraint on the space of possible models. Since the extant data is insufficient to provide an adequate constraint, we need to rely on the task to provide this. We have shown that two key classes of behavioral task, free choice and conflict tasks, cannot offer such constraint because they are fatally confounded. The data derived from such tasks, therefore, cannot be used to support one model, for a myriad of other models could fit it just as well.

The form of this criticism is really quite old-fashioned: We are simply bringing attention to a confound that shows an experimental manipulation does not do what it purports to do. Nonetheless, it has been suggested that because the criticisms cannot be proved or disproved empirically with the tools currently at our disposal, they can be safely ignored.

This unusual argument has been used before in a similar context (see Nachev & Husain, 2007; Owen et al., 2007). Here, as there, it ought to be easy to see that it is misguided. What we have shown is that the conceptual framework of an experiment admits a much larger set of possibilities than has been realized, and, therefore, that the data cannot be used to favor one physiological model over another. That this conclusion cannot be countered by more of the same data is precisely the point: The problem is not the data but the framework on which making sense of it depends. If the framework is shown to be insecure, the data count for nothing.

As an analogy, imagine that someone produced a cosmological theory about some unobserved part of the universe, extrapolating from observational evidence of the seen part through some mathematical theory. If we showed that an error in the mathematics rendered his predictions meaningless, could he reasonably dismiss our criticisms just because we have no empirical evidence to support them? Could he really say: "I know my equation involves a dividing by zero but I am not worried"?

The analogy is helpful because it illustrates a critical point about conceptual criticisms generally: They matter precisely because they are relevant to what we do not observe, but we could theoretically observe, given better experimental tools. We may be able to show empirically that the physiological organization of the brain does not have as a component a macroscopic division between externally and internally

triggered action: It is just that our tools currently make it very difficult. But rather than showing empirically what is conceptually obvious, our time is better employed constructing experiments where the data does tell us something about the underlying physiology. Failing to appreciate the significance of conceptual criticisms can only impede us in that task.

It is also tempting to counter such criticism by demanding a better explanation for the data. But where the criticism is that the framework does not allow us to use the data such a counter-attack would seem silly: We would be guilty of self-contradiction if we insisted on any one model in such circumstances. So in suggesting one possible model for the role of medial frontal cortex in voluntary action, one that offers a better fit to the data—an organization in terms of a rostrocaudal gradient of conditional complexity (Nachev *et al.*, 2008)—we are simply providing a hypothesis to be tested in future experiments. No stronger position can be taken here.

Finally, we should remind ourselves that by rejecting the specific task constraints we deal with here, we are not saying it is impossible to constrain the empirical space with any task: There are circumstances where this may be possible. Furthermore, there are neglected theoretical constraints other than the task, for example the argument against functional pleomorphism (Nachev *et al.*, 2008) that can be very useful. To make criticisms of this kind is therefore not to be nihilistic about behavioral neuroscience.

Conclusion

I have suggested that two classes of behavioral task, “free choice” and “conflict,” are of little use in exploring the neural basis of voluntary action. The reasons for this are to be found in careful examination of the nature of the tasks themselves. Indeed, a subsidiary purpose of this article is to illustrate the importance of examining tasks much more closely than is usual in neuroscientific studies, where the transition from concepts to data often skims over this critical terrain. If the task is wrong, then whether or not the data are right has no bearing on the question they are supposed to answer.

Where does this leave the empirical study of voluntary action? I have said nothing about alternative empirical strategies precisely because our task now, it seems to me, is to pause and to consider carefully how we can constrain the empirical space we need to explore here. Attempting to build the roof before the foundations are in place can only have one outcome.

Figure captions

Figure 1. A simple sketch of scientific inference. A range of hypothetical models is narrowed down on the basis of data generated from experiments designed to discriminate between the models. Note that confidence in the superiority of any one model depends on tacit assumptions about the space of possible models.

Figure 2. Scientific inference in the cognitive neuroscientific domain. The pattern here is the same as in Figure 1, except that the combination, of our ignorance and the complexity of the brain, is such that the model space is vastly greater than is usual in science. Experimentation here is, therefore, dependent on narrowing the space of possible models with the aid of constraints derived from the behavioral task. Whether or not such constraints are successful, therefore, depends on the true nature of the task. [Note to Editor: this figure is meant to be reproduced at a scale such that the text in the boxes is invisible: it is meant to convey the disparity with the size of the model space in Figure 1; the identity of the boxes does not matter because they are the same as in Figure 1].

Figure 3. The structure of volitional contrasts. In order to isolate what is distinctively voluntary about an action we need to compare behavioral circumstances that differ in voluntariness. This may be accomplished by supposedly varying the degree of voluntariness—manipulating the extent of ownership of an action (externally-guided vs. internally-guided) and manipulating the nature of the choice it results from (automatic vs. controlled)—or by comparing times when volition is believed to be exercised versus times when it is not (intention vs. action.) This article only discusses the first two.

Acknowledgements

This article is an abridged version of a talk given at the ICN “Workshop on Volition” in London, on 28 March 2009, organized by Professors Patrick Haggard and Masud Husain.

The author is funded by the National Institute for Health Research. Thanks are due to Catherine Siviter for comments on earlier versions of the text.

References

- Botvinick, M. M. (2007). Conflict monitoring and decision making: reconciling two perspectives on anterior cingulate function. *Cogn Affect Behav Neurosci*, 7(4), 356-366.
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychol Rev*, 108(3), 624-652.
- Brugger, P. (1997). Variables that influence the generation of random sequences: an update. *Percept Mot Skills*, 84(2), 627-661.
- Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philos Trans R Soc Lond B Biol Sci*, 362(1485), 1585-1599.
- Cisek, P., & Kalaska, J. F. (2002). Simultaneous encoding of multiple potential reach directions in dorsal premotor cortex. *J Neurophysiol*, 87(2), 1149-1154.
- Critchley, H. D., Mathias, C. J., Josephs, O., O'Doherty, J., Zanini, S., Dewar, B. K., et al. (2003). Human cingulate cortex and autonomic control: converging neuroimaging and clinical evidence. *Brain*, 126(Pt 10), 2139-2152.

- Deiber, M. P., Passingham, R. E., Colebatch, J. G., Friston, K. J., Nixon, P. D., & Frackowiak, R. S. (1991). Cortical areas and the selection of movement: a study with positron emission tomography. *Exp Brain Res*, *84*(2), 393-402.
- Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, *16*(1), 143-149.
- Frith, C. D., Friston, K., Liddle, P. F., & Frackowiak, R. S. (1991). Willed action and the prefrontal cortex in man: a study with PET. *Proc Biol Sci*, *244*(1311), 241-246.
- Hacker, P. M. S. (2007). *Human Nature: The Categorial Framework*: Blackwell Publishing.
- Husain, M., Parton, A., Hodgson, T. L., Mort, D., & Rees, G. (2003). Self-control during response conflict by human supplementary eye field. *Nature neuroscience*, *6*, 117-118.
- Jahanshahi, M., Jenkins, I. H., Brown, R. G., Marsden, C. D., Passingham, R. E., & Brooks, D. J. (1995). Self-initiated versus externally triggered movements. I. An investigation using measurement of regional cerebral blood flow with PET and movement-related potentials in normal and Parkinson's disease subjects. *Brain*, *118* (Pt 4), 913-933.
- Jenkins, I. H., Jahanshahi, M., Jueptner, M., Passingham, R. E., & Brooks, D. J. (2000). Self-initiated versus externally triggered movements. II. The effect of movement predictability on regional cerebral blood flow. *Brain*, *123* (Pt 6), 1216-1228.
- Logan, G. D. (1985). Executive control of thought and action. *Acta Psychologica*, *60*(2-3), 193-210.
- Logan, G. D. (1994). On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. *Inhibitory processes in attention, memory, and language*, 189-239. London: Academic Press.
- Nachev, P., & Husain, M. (2007). Comment on "Detecting awareness in the vegetative state". *Science*, *315*(5816), 1221.
- Nachev, P., Kennard, C., & Husain, M. (2008). Functional role of the supplementary and pre-supplementary motor areas. *Nature Reviews Neuroscience*, *9*(11), 856-869.
- Nachev, P., Wydell, H., O'Neill, K., Husain, M., & Kennard, C. (2007). The role of the pre-supplementary motor area in the control of action. *NeuroImage*, *36*(Supplement 2), T155-T163.
- Norman, D. A., & Shallice, T. (2000). Attention to action: Willed and automatic control of behavior. In M. Gazzaniga (Ed), *Cognitive neuroscience: a reader*, (pp. 376-390). Oxford: Blackwell Publishers, Ltd.
- Owen, A. M., Coleman, M. R., Boly, M., Davis, M. H., Laureys, S., Jolles, D., et al. (2007). Response to Comments on " Detecting Awareness in the Vegetative State". *Science*, *315*(5816), 1221c.
- Pakkenberg, B., & Gundersen, H. J. (1997). Neocortical neuron number in humans: effect of sex and age. *J Comp Neurol*, *384*(2), 312-320.
- Passingham, R. E. (1993). *The frontal lobes and voluntary action*: Oxford University Press, USA.
- Schall, J. D., Stuphorn, V., & Brown, J. W. (2002). Monitoring and control of action by the frontal lobes. *Neuron*, *36*(2), 309-322.
- Soon, C. S., Brass, M., Heinze, H. J., & Haynes, J. D. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, *11*(5), 543.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology: General*, *18*, 643-662.
- Tune, G. S. (1964). A brief survey of variables that influence random-generation. *Percept Mot Skills*, *18*, 705-710.
- Wagenaar, W. A. (1972). Generation of random sequences by human subjects - critical survey of literature. *Psychological Bulletin*, *77*(1), 65-72.