Running head: FMRI: IT'S NOT ALL BAD

Common ground for behavioural and neuroimaging research.

Scott D. Brown

School of Psychology, University of Newcastle

Abstract

The division between cognitive research that employs neuroimaging data and cognitive research that excludes such data has led to extreme views on both sides. As in most conflicts, these extreme views reinforce each other. They are further deepened by some poor science and a cavalier attitude toward the interpretation of neuroimaging data by some practitioners. I discuss two of the most problematic aspects of the interpretation of neuroimaging data: the logical fallacy of confirming the consequent; and the difficulties of linking assumptions which link the data with theoretical constructs. Both of these problems are not unique to neuroimaging research, and errors with them certainly occur in standard behavioural psychology. Nevertheless, when executed carefully, and when interpreted cautiously, neuroimaging has the potential to contribute to the understanding of cognition.

Keywords: linking hypotheses; MRI; neuroimaging; mathematical psychology.

Common ground for behavioural and neuroimaging research.

A modern division within cognitive psychology is based on the kind of measurements one considers. Some cognitive psychologists focus on behavioural measurements, while others focus on neural measurements (most commonly from function magnetic resonance imaging – fMRI – or electroencephalogram – EEG). As with many divisions, this one can be quite blurry. Almost all of those in the "cognitive neuroscience" camp also examine behavioural data, and some of those in the "behavioural" camp, including myself, also examine neural data. The differences between the camps often come down to a greater emphasis on one or the other kind of data, rather than exclusion. For example, almost every article reporting a cognitive fMRI experiment also reports behavioural data from the experiment, at least at a cursory level such as mean reaction time and accuracy. Similarly, even the staunchest critics of fMRI research might sometimes find their theoretical developments guided by consideration of data from patients with brain lesions.

Division along this continuum has been deepened by the promotion of extreme views on both sides. For example, a recent discussion by Insel (2010) –published in the very prominent outlet *Scientific American* – goes so far as to suggest that behavioural measurements will become obsolete as non-invasive physiological measurements become more prevalent. The extreme version of this view characterises most behavioural data as "subjective, mental" measurements, which are inferior to the more objective data from neuroimaging. In a prominent example of the counterargument, Coltheart (2006), claimed that more than a decade of functional neuroimaging had not contributed at all to the understanding of cognition. Even more extreme views have been promulgated by some, for example Miller (2010) claims not just that neuroimaging has not contributed so far, but that it will never contribute, at least in its current form. As with most deeply partisan debates, the truth probably lies somewhere in between these extremes; that both

behavioural and physiological data can be useful and informative, and that when well-executed, functional neuroimaging data can provide another useful data stream to help us understand cognition. Detailed accounts of precisely this middle ground have been set out before by many others (e.g., for recent versions, see Gonsalves & Cohen, 2010 or Forstmann, Wagenmakers, Eichele, Brown, & Serences, 2011).

It is easy to list examples of very bad research using functional neuroimaging (by which I mainly mean fMRI and EEG). Such research typically makes errors of inference which lead to grand claims that are both unsupported by the data and which fly in the face of well-established results from behavioural experiments. However, lists of bad research are not helpful. These lists merely illustrate that - as with all measurement tools - there is both good and bad research work carried out using neuroimages.

The most common inferential error in neuroimaging research is sometimes known as the "fallacy of confirming the consequent", which I discuss in more detail below. This fallacy leads some researchers to mistakenly construe weak evidence in favour of a theory as very strong evidence that logically compels the theory – this deepens the divide by really annoying the better scientists from the behavioural camp, who are always careful to justify all of their own claims. However, this particular kind of bad research is also not confined to cognitive neuroimaging, because one can commit the fallacy of confirming the consequent whenever one considers whether the data are *consistent* with the theory, without also considering what other theories might also have given rise to the data. For instance, Dunn and Kirsner (1988) provide a detailed account of the way that cognitive psychological theories, when specified at a verbal level, can be linked to empirical data, and how failure to consider the nuances of these links has resulted in unsupported claims about many different theories. I am guilty of making such claims myself, and in at lest one case proponents of the alternative theory did not think my claims were as strongly supported as I thought they were (Brown, Marley, Dodds, & Heathcote, 2009; Stewart &

Matthews, 2009). My point here is that it is a mistake to assume that poor logic and unsupported conclusions are new problems, or are unique to neuroimaging research.

The most vehement elements of this argument are restricted to research involving just one sort of physiological data – blood-oxygenation level dependent (BOLD) responses measured during fMRI scans. Other physiological measurements, such as spike trains measured from the neurons of behaving monkeys or cats are not usually subject to the same criticisms, although I'm not sure why this should be. Nor are other techniques closely related to BOLD analyses, such as electro-encephalograms (EEG), or the use of purely structural data measured using MRI machines. The particular kind of research that raises criticism most often is when these BOLD measurements are used to investigate cognitive hypotheses. These studies often examine differences in BOLD activity in certain brain regions between conditions, and try to relate those differences, and those regions, to the cognitive theory. It seems to be this link, between neural data and cognitive theory, that contributes most to the division (for examples of the criticisms, see, e.g., Coltheart, 2006, and for examples of the criticised type of research, see the ACT-R work described below, such as that by Qin et al., 2003).

The difficulties of using neuroimaging data to test cognitive theories are well-known, and have been discussed at length and in great detail, by people much better informed than myself. An easy place to begin reading about this topic is recent special in *Perspectives on Psychological Science*, where thoughtful arguments and counter-arguments were presented by both sides. Some of the most nuanced perspectives on the flaws of fMRI research are the ones discussed by those active in the field, such as Miller's (2010) in that special section. In the same volume, Beck (2010) makes an interesting case that fMRI research should be subject to *more* stringent quality control than other research, because its appeal to the general public can sometimes lead to decreased skepticism. Although it is not topic of this article, there is much to be said for Beck's argument – it seems possible

that the cause of at least some of the angst about neuroimaging research is the apparent ease with which credulous journal editors and grant agencies support it.

Given what has already been said, I do not hope to contribute much new to this argument. Instead, I hope to make some links between behavioural and neuroimaing research, by discussing problems that are often associated with one type of research and showing how those problems are tractable, and are also common to both kinds of research. The two problems I focus on are:

- Many cognitive theories have little or no neural content. How should physiological data be related to such theories? Such difficulties often encourage the confirmation of the consequent, and therefore "blobology".
- The proper place and consideration of "linking functions" between physiological measurements, cognitive theories, and behavioural data.

Blobology and Confirmation of the Consequent

From the point of view of behavioural scientist, some of the most infuriating research using fMRI data is the kind of investigation that might be dismissed as "blobology". At its most basic, this is apparently empty research that attempts to connect cognitive processes with brain regions. Such research is open to criticism on several fronts, and these criticisms have been made forcefully and clearly by others (e.g., by Coltheart, 2006 and Miller, 2010). Blobological research is empty if it goes no further than identifying a putative mapping between regions and processes - it is not clear whether anything new could ever be learnt about the cognitive processes from this effort. Blobology also begins with the powerful assumption that finding such a mapping is even possible. For example, to link the cognitive processing of speech production to Broca's area suggests a one-to-one link: that Broca's area is differentially activated whenever speech production processes are engaged, and that whenever Broca's area is differentially activated, then speech

production processes must be engaged. Such a link is highly questionable, because even the simplest brain structures have been hypothesised to be involved in more than one cognitive process, and even the simplest cognitive process presumably recruits more than one brain region. For example, searching the automated meta-analysis website for fMRI data, *NeuroSynth* (`http://neurosynth.org/`), shows that, from 48 studies, four different brain regions have been associated with the term "speech production". There are good reasons why one might expect multiple regions to be engaged by even the simplest cognitive processes (e.g., because of all the concomitant processes that are engaged, such as attention, perception, and movement). For these reasons, perhaps the goal of finding links at the level of regions is wrong. Instead, perhaps a better level of analysis might be of *networks* of regions; for a detailed discussion of these matters, see McIntosh (2000)

Given the many-to-many mapping between brain regions and cognitive processes, it is simplistic in the extreme to assume that a simple mapping can be identified by blobology. In fact, to assert such a mapping from most data sets is to commit the fallacy of confirming the consequent. In terms of formal logic, this is just faulty reasoning of the form: *If A, then B; I observe B; Therefore A.* Translated into an example from neuroscience, this might go: *The anterior cingulate is differentially activated by conflict monitoring; I observe differential activation in the anterior cingulate; Therefore, conflict monitoring must be occurring.* The logic is wrong because it overlooks the many other possible causes of differential activity in the anterior cingulate.

Because of this logical problem, observing differential activity in some region does not logically compel the conclusion that some associated cognitive process is engaged. However, this does not mean that the data are worthless. The same kind of conclusion really can be drawn from such data, but just at a weaker and more carefully qualified level. As with almost all data analysis in psychology, it is a matter of probability and the balance of evidence. To continue the above example, observing differential activity in the

anterior cingulate does not necessarily compel the conclusion that conflict monitoring is occurring, but it does make that hypothesis more likely, simply because the vast majority of cognitive processes have *not* been associated with the anterior cingulate.

In a study highly critical of much fMRI-based research, Poldrack (2006) took this line of argument out of the realm of armchair philosophy, and into the realm of empiricism. By searching another large database of fMRI studies (*BrainMap*, `http://www.brainmap.org/`), and applying a relatively simple statistical analysis, he was able to quantify the amount of additional evidence in favour of the engagement of some specific cognitive processes provided by the observation of differential activation in certain brain regions. The amount of additional evidence depended on the size of the brain region (more specific regions of interest provide stronger evidence) and the diversity of findings associated with it (more selective influence leads to stronger evidence). In one example, Poldrack examined the link between differential activity in a region centred on Brodmann's area 44 – part of the left-side inferior frontal gyrus – and the hypothesis that language processing is occurring. In the best case that he examined for this area (a 4x4x4 millimetre region of interest), the Bayes Factor in favour of the cognitive process occuring was 4.5. This means, for example, that if you were previously completely unsure about whether language processing was occurring (i.e., you gave the hypothesis a 50% chance of being true) observation of differential activity in the region would increase your belief to 82%. Viewed in these more careful terms, the fMRI data certainly do not support the simple link between activation and cognition that some people propose, but they do go someway towards that goal.

Can blobology be useful?

In its most basic form, there is little doubt that blobology is close to vacuous. The mere (and uncertain) identification of brain regions with cognitive processes does not

advance knowledge about cognition at all, and does not advance neurological knowledge very much. However, there might be certain cases in which this exercise is not only useful, but is the only reasonable way to proceed. The most obvious such case is when dealing with purely cognitive theories – those that do not purport to describe neurophysiological processes or implementation in any detail. These theories are very common and they cannot be tested using physiological data from imaging experiments (at least, not without making some very tenuous linking hypotheses between the data and the theory; see next section). Nevertheless, linking the various components of these theories with brain regions, through blobology, can be the first step in extending the cognitive theory to include a more detailed account of physiological implementation.

Consider the example of the ACT-R production theory, which began as a purely cognitive account (see, e.g., Anderson, 1992, among hundreds of other papers). Despite its many shortcomings, one strength of the ACT-R framework has been its continuing extension to cover new paradigms and explain data from different kinds of experiments and different measurements, from education to memory, decision-making to car driving, and language to human-computer interaction. One way in which ACT-R has been extended is to develop a neurophysiological interpretation of its proposed cognitive processes; this is an extension because it makes predictions for yet another kind of data, and it provides an account at another level of explanation (Marr, 1982).

It is difficult to imagine how the extension of a purely cognitive theory to include a neurophysiological account could begin without some kind of exploratory mapping between the basic elements of cognition and neurophysiology: blobology, in the case of fMRI data. In its initial stages, this investigation is vacuous in just the way described earlier. The results are also tenuous just because they are exploratory: the theoretical conclusions are direct re-descriptions of the data, and if the data had turned out differently the theory would simply have been re-made in their image. For example, Qin et al. (2003)

used a symbolic manipulation task and found that the time series of the BOLD response in left pre-frontal cortex could be predicted accurately if it was linked to a memory retrieval component of ACT-R. This exploratory finding is open to the criticism that the theory would also have been consistent with almost any other data pattern from the fMRI measurements – presumably, the ACT-R components would simply have been linked with whichever brain regions displayed the appropriate differential activations. However, this criticism is only valid the first time the link is made. In subsequent studies, if the fMRI data turn out differently then the hypothesised link is not supported. Conversely, if the fMRI data turn out similarly, the new neurophysiological component of the ACT-R theory is supported – for examples of this kind of support, see Sohn, Goode, Stenger, Carter, and Anderson (2003) and Danker, Gunn, and Anderson (2008). Once a link between some element of a cognitive theory and some brain region (or network) has been proposed and supported, it becomes part of the larger cognitive-neurophysiological theory. Just like any other theory, it is then open to empirical evaluation. The new neurophysiological theoretical elements can even be used to set up discriminatory tests to evaluate competing accounts, if those accounts make different predictions for neurophysiological data. In the simplest case, competing theories might predict the involvment of different brain regions in some particular cognitive task. Empirical evaluation of such predictions is possible with carefully designed experiments, in just the same way that competing predictions of purely cognitive theories can be evaluated with behavioural experiments. In the next section, I describe an example of this process regarding the "cortico-striatal" account decision caution (e.g., Lo & Wang, 2006), which makes different predictions about neural activity from competing accounts based on response urgency (Cisek, Crammond, & Kalaska, 2003; Cisek & Kalaska, 2005). Another example of this process maturing relates to the diffusion model for decision making (Ratcliff, 1978; Ratcliff & Rouder, 1998). That model described decision making in purely cognitive terms, and was successfully extended to many aspects

of behaviour. In more recent work, Ratcliff, Segraves, and Cherian (2003) used single-cell neural recordings to support a link between a core cognitive component of the model ("evidence accumulation") and the firing rates of particular cells in the superior colliculi of monkeys. Once made, this link can be contrasted with other theories that make the link to different brain regions (as with Cisek et al.'s account). Even though the links between the diffusion model and neurophysiology were made using single cell recordings rather than neuroimaging, similar arguments apply – had the recordings from different structures than the superior colliculus matched the data, the theory would have been revised. Importantly, once the initial link is made it is open to test and even extension to neuroimaging data (for a related evaluation using ERP data, see: Ratcliff, Philiastides, & Sajda, 2009).

Linking Hypotheses

Drawing inferences from neural data about a cognitive theory requires a set of "linking hypotheses". These hypotheses form an adjunct to the cognitive theory by specifying how elements of the cognitive theory should be mapped onto neurophysiological data. Linking hypotheses have received decades of study both at the philosophical level and the empirical level (e.g., Teller, 1984, provides an early, but comprehensive review of the philosophical side). A canonical example concerns colour vision in humans. One might want to make inferences about psychological states (such as "the perception of redness") from neurophysiological data (such as the firing rates of red-green cone cells in the retina). This inference requires a fairly obvious linking hypothesis: the perception of red-green is caused by the activities of the relevant retinal cone cells. Linking hypotheses can also be made in the opposite direction. Indeed it originally occurred that way when Young and Helmholtz drew inferences about the neural structure of colour vision from behavioural data.

The need for linking hypotheses is not particular to cognitive neuroscience. Even

when only behavioural data are involved, linking hypotheses are required to relate observed data to the latent constructs of cognitive theories. This problem is one of the first things we teach undergraduate students in research methods courses: the difference between a latent variable and its operationalised counterpart. These links are sometimes well understood and carefully tested (e.g., the relationship between behavioural data from the complex digit span task and the latent construct of "working memory") but other times they are not (e.g., the relationship between confidence ratings from recognition memory experiments and hypothesised states of "familiarity" or "recollection"). Dunn and Kirsner (1988) provide an excellent review of many ways in which untested linking hypotheses have led to unsupported inferences in behavioural studies of cognition. Once again, my point is that the problems that arise in neuroimaging research are closely related to those that arise in almost all other psychological research. For this reason, we should be careful not to single out neuroimaging research for special censure.

An Example

As an example of a study using fMRI data to test a cognitive (and neural) theory, I will use the study reported by Forstmann et al. (2008). We investigated the influential group of theories that assume simple decisions are made by "sequential sampling". These theories have been specified at both the neural and cognitive levels (see, e.g.: Bogacz & Gurney, 2007; Mink, 1996; Usher & McClelland, 2001). Cognitively, the theories propose that simple decisions are made by sequential sampling, which is the dominant theoretical paradigm for models of choice response time (e.g.: Ratcliff, 1978; Smith & Ratcliff, 2004; Van Zandt, Colonius, & Proctor, 2000). Sequential sampling models assume that decisions are made by accumulating evidence in favour of the different choices, and then triggering a decision as soon as the amount of evidence for any one choice exceeds some "response threshold". These models have been studied by psychologists for more than half a century

(at least since Stone, 1960). More recently, the neural instantiations of sequential sampling models have been investigated, initially using data from awake, behaving monkeys (Churchland, Kiani, & Shadlen, 2008; Gold & Shadlen, 2007; Smith & Ratcliff, 2004; Shadlen & Newsome, 2001).

Forstmann et al. (2008) examined one particular aspect of sequential sampling models: the speed-accuracy tradeoff. In almost any task, observers can be encouraged to make very careful decisions, at the expense of also making very slow decisions; or they can be encouraged to make very fast, but also less accurate decisions. This ability is captured at the cognitive level by the concept of a "response threshold". If an observer sets a low response threshold, not much evidence is accumulated in favour of a choice before a response is triggered, so the predicted responses are fast but inaccurate (and vice versa for a high response threshold). Neural accounts of the mechanism that implements this threshold adjustment are still a matter of debate, but an influential group of theories is the cortico-striatal account (e.g. Lo & Wang, 2006). These theories were developed in view of the anatomical connections between brain structures, including the various excitatory and inhibitory connections between efferent nerves and the central nervous system via lower brain structures. The cortico-striatal accounts focus on the role of the basal ganglia, which are a collection of small nuclei. In their default state, the output nuclei of the basal ganglia send tonic inhibition to lower brain regions, preventing the premature execution of any action (Deniau & Chevalier, 1985). When cortical processes start to favour one response (an action), this leads to selective activation of the input nuclei of the basal ganglia (mainly the striatum), which, in turn, leads to selective suppression of the output nuclei, releasing the lower brain from inhibition and allowing the appropriate action to be executed (Mink, 1996).

Forstmann et al. (2008) tested the cortico-striatal account of the speed-accuracy tradeoff in humans by having observers make simple perceptual decisions while an MRI

scanner recorded changes in the BOLD signal from their brains. The observers were sometimes encouraged to make fast (but possibly inaccurate) decisions, and other times to make slow but careful decisions: we called these the "speed-emphasis" and "accuracy-emphasis" conditions, respectively. A prediction of the cortico-striatal account for this experiment is that BOLD signals recorded from the basal ganglia and related cortical areas should be different between the speed-emphasis and accuracy-emphasis conditions.

This hypothesis is easy to test using imaging equipment and a carefully designed experiment. However, drawing conclusions about the underlying theories requires making several assumptions about the relationships between the observations, the experimental conditions, and the theories; these are the linking hypotheses. For example, a linking hypothesis between the empirical imaging data and the neural theory is that activation in the striatum is monotonically associated with BOLD measurements from that brain region. Another linking hypothesis, between the experimental design and the cognitive theory, is that cues encouraging fast (or careful) responses change the response thresholds. A related linking hypothesis connects changes in a cognitive response threshold with changes in the activity of neural structures which implement this threshold. Some of these hypotheses are so tightly integrated with the neural theory itself that it is unclear whether they should be considered "linking hypotheses", or rather elements of the theory itself. For example, the height of the response threshold described in the cognitive theories of decision making is assumed to be - at least partially - instantiated in the neural theories by the level of tonic excitation arising from cortical areas (mostly the pre-supplementary motor area - preSMA) and corresponding tonic inhibition from the striatum. This might be considered a linking hypothesis between the neural and cognitive theories, or it might be considered an integral element of a larger, unified theory; either way, it is an interesting research question in its own right.

Model-Based Cognitive Neuroscience

Linking hypotheses are pervasive in cognitive research, because a primary goal is to make inferences about unobservable, latent quantities (such as memory strength) on the basis of empirical data (such as decision accuracy or response time). In general, the role of these linking hypotheses is well accepted, and their existence does not halt research progress. However, when cognitive research uses fMRI data, the entailed linking hypotheses cause greater consternation than in many other instances. A possible reason for this is that the linking hypotheses are sometimes very poorly explicated in cognitive neuroimaging studies, or perhaps the linking hypotheses are not examined closely enough. These problems are ones of practice, not of principle: the linking hypotheses *can* be well specified, and carefully evaluated. For this reason, it seems sensible to treat cognitive neuroimaging research in the same manner as other cognitive research, and the divide described earlier seems artificial.

Many of the linking hypotheses assumed in cognitive neuroimaging can be tested through careful research, and especially through the application of quantitative models. A common linking hypothesis in almost all neuroimaging research is that the average firing rate from any group of spatially-contiguous neurons is monotonically related to the BOLD response measured from that location. This hypothesis has received a great deal of investigation for many years. As one might expect from hundreds of studies on the same topic, there are findings both for and against this crucial linking hypothesis. The balance of evidence seems to support the link, at least when well-defined brain structures with mostly homogenous neuronal populations are of interest. For example, an influential study by Rees, Friston, and Koch (2000) found a simple relationship between the signal recorded by fMRI and the neural firing rates recorded directly from the visual cortex of monkeys. Indeed, the relationship was so direct that it was possible to quantify the linking hypothesis mathematically: a 1% change in the BOLD response corresponded to a change

in neural firing rate of about 9 spikes per second (see Heeger, Huk, Geisler, & Albrecht, 2000 for a corresponding result from a different region of visual cortex). Other research has identified conflicting evidence. For example, Sirotin and Das (2009) demonstrated at least one instance in which BOLD responses were dissociated from spiking activity by showing that spatially global changes in the BOLD response preceded the onset of a periodic stimulus were not correlated with changes in spiking activity. Fortunately, most fMRI studies do not use regularly paced trials, and in most analysis methods the subtraction of spatially global modulations has a negligible impact on the results (see, e.g.: Serences & Boynton, 2007).

Conclusions

My claim is not that cognitive neuroimaging studies are uniformly good, or that they should be preferred to cognitive studies that do not use neural data; neither of these statements are even approximately true. It is also an unhelpful generalisation to tar all neuroimaging studies with one brush. To speak against such generalisation, I have tried to highlight some common ground between the two kinds of studies. I have attempted this by re-describing two of the most common and well-understood pitfalls in the analysis and interpretation of neuroimaging experiments. I have shown – as have others before me – that these problems are not insurmountable, and that related problems can arise in purely behavioural research. Indeed, cognitive psychologists are guilty of making just these sorts of errors. As in behavioural research, when the analysis assumptions are taken seriously, and the logical arguments are made cautiously, neuroimaging experiments have the potential to make scientific contributions. Because of the similarities, it does not seem useful to draw a hard distinction between cognitive psychology based on behavioural data and cognitive psychology that also considers neuroimaging data.

References

Anderson, J. R. (1992). Automaticity and the ACT* theory. *American Journal of Psychology*, *105*.

Beck, D. (2010). The appeal of the brain in the popular press. *Perspectives on Psychological Science*, *5*, 762–766.

Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Computation*, *19*, 442–477.

Brown, S. D., Marley, A., Dodds, P., & Heathcote, A. (2009). Purely relative models cannot provide a general account of absolute identification. *Psychonomic Bulletin & Review*, *16*, 583–593.

Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision–making with multiple alternatives. *Nature Neuroscience*, *11*, 693–702.

Cisek, P., Crammond, D. J., & Kalaska, J. F. (2003). Neural activity in primary motor and dorsal premotor cortex in reaching tasks with the contralateral versus ipsilateral arm. *Journal of Neurophysiology*, *89*, 922–942.

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

Coltheart, M. (2006). What has functional neuroimaging told us about the mind (so far)? *Cortex*, *42*(3), 323–331.

Danker, J., Gunn, P., & Anderson, J. R. (2008). A rational account of memory predicts left prefrontal activation during controlled retrieval. *Cerebral Cortex*, *18*, 2674–2685.

Deniau, J., & Chevalier, G. (1985). Disinhibition as a basic process in the expression of striatal functions. ii. the striato-nigral influence on thalamocortical cells of the ventromedial thalamic nucleus. *Brain Research*, *334*.

Dunn, J. C., & Kirsner, K. (1988). Discovering functionally independent mental processes: The principle of reversed association. *Psychological Review*, *95*, 91–101.

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., et al. (2008). Striatum and pre–SMA facilitate decision–making under time pressure. *Proceedings of the National Academy of Sciences*, *105*, 17538–17542.

Forstmann, B. U., Wagenmakers, E.-J., Eichele, T., Brown, S., & Serences, J. T. (2011). Reciprocal relations between cognitive neuroscience and cognitive models: Opposites attract? *TRENDS in Cognitive Sciences*, *15*(6), 272–279.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574.

Gonsalves, B., & Cohen, N. J. (2010). Brain imaging, cognitive processes, and brain networks. *Perspectives on Psychological Science*, *5*, 744–752.

Heeger, D., Huk, A., Geisler, W., & Albrecht, D. (2000). Spikes versus bold: what does neuroimaging tell us about spikes versus bold: What does neuroimaing tell us about neuronal activity? *Nature Neuroscience*, *3*(7), 631–633.

Insel, T. (2010). Faulty circuits. *Scientific American*, *302*, 44–51.

Lo, C.-C., & Wang, X.-J. (2006). Cortico-basal ganglia circuit mechanism for a decision thresh- old in reaction time tasks. *Nature Neuroscience*, *9*, 956–963.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information.* New York: Freeman.

McIntosh, A. (2000). Towards a network theory of cognition. *Neural Networks*, *13*, 861–870.

Miller, G. (2010). Mistreating psychology in the decades of the brain. *Perspectives on Psychological Science*, *5*, 716–743.

Mink, J. W. (1996). The basal ganglia: Focused selection and inhibition of competing motor programs. *Progress in Neurobiology*, *50*.

Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data? *TRENDS in Cognitive Sciences*, *10*(2), 59–63.

Qin, Y., Sohn, M.-H., Anderson, J. R., Stenger, V., Fissell, K., Goode, A., et al. (2003). Predicting the practice effects on the blood oxygenation level-dependent (bold) function of fmri in a symbolic manipulation task. *Proceedings of the National Academy of Sciences*, *100*(8), 4951–4956.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.

Ratcliff, R., Philiastides, M. G., & Sajda, P. (2009). Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the eeg. *Proceedings of the National Academy of Sciences*, *106*, 6539–6544.

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two–choice decisions. *Psychological Science*, *9*, 347–356.

Ratcliff, R., Segraves, M., & Cherian, A. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of simple two–choice decisions. *Journal of Neurophysiology*, *90*, 1392–1407.

Rees, G., Friston, K., & Koch, C. (2000). A direct quantitative relationship between the functional properties of human and macaque v5. *Nature Neuroscience*, *3*, 716–723.

Serences, J. T., & Boynton, G. (2007). Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron*, *55*(2), 301–312.

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, *86*, 1916–1936.

Sirotin, Y., & Das, A. (2009). Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity. *Nature*, *457*(475–479).

Smith, P. L., & Ratcliff, R. (2004). The psychology and neurobiology of simple decisions. *Trends in Neurosciences*, *27*, 161–168.

Sohn, M.-H., Goode, A., Stenger, V., Carter, C. S., & Anderson, J. R. (2003). Competition and representation during memory retrieval: Roles of the prefrontal cortex and the posterior parietal cortex. *Proceedings of the National Academy of Sciences*, *100*, 7412–7417.

Stewart, N., & Matthews, W. (2009). Relative judgment and knowledge of the category structure. *Psychonomic Bulletin & Review*, *16*, 594–599.

Stone, M. (1960). Models for choice–reaction time. *Psychometrika*, *25*, 251–260.

Teller, D. (1984). Linking propositions. *Vision Research*, *24*(10), 1233–1246.

Usher, M., & McClelland, J. L. (2001). On the time course of perceptual choice: The leaky competing accumulator model. *Psychological Review*, *108*, 550–592.

Van Zandt, T., Colonius, H., & Proctor, R. W. (2000). A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin & Review*, *7*.

Author Note

This research was supported by ARC Discovery Project DP0878858.

Correspondence concerning this article may be addressed to: Scott Brown, School of Psychology, University of Newcastle, Callaghan NSW 2308, Australia; Email: scott.brown@newcastle.edu.au.