

The Brain at Rest: What it's Doing and Why That Matters

Colin Klein*

University of Illinois at Chicago
Australian National University

Abstract

Neuroimaging studies of the resting state continue to gather philosophical and scientific attention. Most discussions assume an identification between resting-state activity and activity in the so-called ‘default mode network.’ I argue we should resist this identification, structuring my discussion around a dilemma first posed by Morcom and Fletcher. I offer an alternative view of rest as a state dominated by long-term processes, and show how interaction effects might thereby let rest shed light on short-term changes in activation.

1 Why care about rest?

I’m interested in using fMRI to constrain theories of neural function—that is, in moving from fMRI data down to computational or functional theories of the brain. This is known as *forward inference*. It is usually contrasted with *reverse inference*, which goes from

*cvklein@gmail.com

brain data back up to cognition. Reverse inference is notoriously problematic (Poldrack 2006). Forward inference, on the other hand, is often treated as relatively unproblematic.

I think that forward inference is also trickier than it seems. fMRI data very often shows regions where there was a *change* in activation between two tasks. So if, say, you care about reading, you might have subjects read words and then visually similar non-words, and then look at regions where there's a significant difference in response between the two. If you've matched your tasks correctly, you might think that this should highlight reading-selective areas.

The fact that you're working with differences in activation, however, raises an interesting methodological problem. Here's an analogy. Suppose we were trying to figure out the function of the heart. Suppose further that all we had were a list of circumstances in which one could expect changes in heart rate (or, even better, a list of pairs of circumstances in which heart rate differed). We might find out that the heart generally beats faster when you run, slower when you sleep, and so on. We could then try to build up a theory of hearts based on just this. Perhaps we would decide that the heart has something to do with arousal, and then we could argue whether this link with arousal is cause or consequence. But that, I take it, would miss the really important part—which is that the heart is *always* pumping blood, and pumping blood at any rate has a crucial functional role. The fluctuations in how the heart plays this crucial functional role are an important, but relatively minor, part of the whole story.

As with the heart, so with the brain. You might think that fMRI would be more informative if we had additional background information that we could use to interpret

changes in activation. A natural place to look for such a baseline is rest. I mean ‘rest’ in the neuroimaging sense: not being asleep, but rather sitting quietly in the scanner between performances of cognitive tasks. ‘Rest’ in this sense is a neurally active state. Your brain uses approximately 20% of your body’s energy even though it’s only about 2% of its weight. Most of that activity seems to be devoted to transmitter reuptake and other neurally interesting processes. The fluctuations that fMRI tracks usually aren’t more than about 5% of that overall baseline (Raichle and Mintun 2006). So the brain seems to be doing a lot, even when it’s not engaged in a task. That’s why it might be useful to know just what it’s doing.

So here’s a suggestion: why not figure out what’s going on at rest, and use that to interpret changes in activation? There’s been a real explosion of interest in doing precisely this in the last decade, both among philosophers and scientists. I’d like to focus on a methodological problem with this suggestion, first raised by Morcom and Fletcher in their (2007). They pose an interesting dilemma for those who would study rest. Suppose we care about the resting state. Either the resting state is the sort of thing that we can study using task-based fMRI, or it’s not. If it’s not, then we don’t really have any way to study it. The problem isn’t just that sticking electrodes in people is hard and complicated. It’s rather that most of psychology is task-based rather than observational: like other experimental sciences, it makes specific interventions on the system it cares about and uses the results of those interventions to build theories. Tasks just are those interventions. If we can’t get at rest using tasks, then, we can’t get at it at all. On the other hand, if rest *is* the sort of thing that we can study using differences in tasks, then it seems that facts about tasks that are driving our theories, not the study of rest per se. So

either it's impossible to interpret the rest state, or else the interpretation comes via data about plain old increases in activation. Either way, rest ends up evidentially irrelevant.

I think this is a really interesting dilemma. I want to talk about one currently popular solution: the notion of a default mode network and the identification of activity in such a network with the resting state. I'm going to argue against that solution, and introduce an alternative notion of what's going on at rest. I think this alternative notion solves the evidential dilemma that Morcom and Fletcher raise. I conclude by pointing out an odd, but I think acceptable, feature of my solution.

2 The Default Mode Network

2.1 The Standard Thesis

Early in the history of brain imaging, researchers discovered that some brain areas that show *decreases* relative to rest when subjects are engaged in cognitive tasks. If you tell people to read words, most of the time you're looking for areas of increase because you're interested in reading. But if you analyze that task relative to just a resting state, you find many brain regions that also show *less* activation during the task relative to rest. A wide variety of tasks give rise to these task-related deactivations.

In an influential set of papers, Deborah Gusnard and Marcus Raichle observed that the regions that show these decreases are fairly consistent across both PET and fMRI

studies (Gusnard and Raichle 2001; Raichle et al. 2001). The core network composes the precuneus, posterior cingulate cortex, ventromedial cortex, more laterally some temporoparietal structures, and less consistently some midline or inferior temporal regions. Most task-related decreases are seen in this set of areas. So the original suggestion was that these regions jointly compose a “default mode network” (DMN)—“an organized mode of brain function that is present as a baseline or default state and is suspended during specific goal-directed behaviors.” (Raichle et al. 2001 676) This in turn gives a positive thesis: the DMN is the interpretive baseline we’re looking for. So if we knew what the DMN was doing, we’d have a better handle on how to interpret activations in other situations. Now, DMN activation doesn’t have to be *specific* to rest—there might be other task-driven ways to get increases in this network. Figuring out what it’s doing at rest, though, will give us a baseline from which we can interpret activations.

Further, there is often a specific cognitive gloss put on the activity in this network. Activation in the DMN shows up in other studies. In particular, the DMN shows up as strongly activated in studies of self-reflection, episodic and autobiographical memory, imagination, theory of mind, moral reasoning, and what are generally called ‘self-projection’ tasks (Buckner and Carroll 2007). Raichle and Gusnard — and many more since — thus made the following proposal. Rest is not actually a blank slate. What people do at rest is some variety of self-reflection: crudely, you spend your downtime in the scanner thinking about your date last night, or whether to cheat on your taxes, or the thing her father said. This activity drops off when you get to work, and that shows up as a decrease in activity in the DMN. That sort of self-reflection is thus a cognitive

default mode that you drop into without more pressing tasks to take up your time

An important qualification is necessary here. There is now an enormous amount of work on the DMN. Much of it is not concerned with the thesis that I just sketched—that is, with the thesis that the DMN has some specific cognitive gloss or other, and that this can be used as an interpretive baseline. While the original motivation for looking at the DMN was in part due to interest in finding a putative functional baseline (hence the ‘Default Mode’ part of the name), interest in the network itself now far outstrips that motivation. Further, many studies of the DMN aren’t explicitly concerned with rest at all: they study the network for its own sake. What follows, then, should be understood as a more limited attack on the utility of DMN studies for understanding the resting state.

2.2 Four Arguments Against the Standard Thesis

I have four arguments against the thesis that the default mode network should be taken as an interpretive baseline.

First, there’s an argument from oxygen extraction fraction (OEF). The OEF is a physiological measure. The OEF of a region of brain at a particular time tracks the proportion of oxygen that’s extracted from a volume of blood passing through that region. It’s an indirect marker of how hard neurons are working. More activity increases regional blood flow, which causes the OEF to drop (though the amount of oxygen delivered increases overall, fueling neural work (Nair 2005)). Changes in OEF drive changes in the BOLD signal, which is turn what fMRI tracks. OEF is thus a marker in

whether things are working *more or less hard than average*. It's around its baseline when a region is around its modal level of activation, and then changes as areas become more or less active.

So we might ask whether the default mode network is more active at rest. We could answer that by looking at the OEF. We can look at the OEF directly using PET, which is good for measuring physiological variables. What do you find? You find that OEF is more or less uniform at rest. That is, across the brain, in both grey matter and white matter, it doesn't look like there are *any* areas that are actually working harder than normal during the resting state (see (Raichle and Mintun 2006 462ff]) for a review). That suggests that whatever's going on at rest, there's nothing special about the DMN itself

It also suggests—and this is crucial—that the decreases in activation one sees relative to rest are actually the DMN being *deactivated*, perhaps positively inhibited, rather than it dropping from a high level of activity to a lower baseline. The two scenarios would both produce a decrease in activation, but only the former is compatible with a uniform resting OEF.

It is true, and worth mentioning, that DMN areas seem to have a higher resting metabolic rate. This is particularly true for the precuneus and posterior cingulate cortex (Cavanna and Trimble 2006). This is interesting, but not germane. Return to the heart. Suppose we found out that, when people lay quietly, the heart used up more ATP than the kidneys did. Does that mean that the heart is *working harder* than the kidneys? Not really. In fact, the question seems like a bit of a mistake. What the heart does at rest takes a certain amount of energy, and same with the kidneys. We can compare different

states of the heart, or different states of the kidney—but comparing across organs isn't sensible, because what they're doing is so fundamentally different. The best we can say is that the heart and kidney at rest are both using exactly as much energy as they need to be using to do what they do. So too with brains: the OEF tells us whether an area is working hard, as it were, *by its own lights*. The metabolic rate for an area, by contrast, tells you how hard, overall, that area has to work to do what it does. It's the former that's relevant to interpreting the DMN. And again, when we look at OEF, it suggests that the DMN is not working especially hard at rest.

Second, there's an argument from resting state correlations. Some background first: a lot of the interesting work on the resting state has involved correlations between fluctuations in activity in DMN areas over relatively long (10sec) timespans. The initial work on this showed that there are nice correlations between DMN areas over long timespans, and that the DMN also appears to be anticorrelated with another so-called 'task-positive' network that becomes active during tasks.¹ However, claims about what's *not present* at rest are tricky: there might be networks that are just easy to overlook.

More recent research suggests precisely that. If you look at resting state data and task data and do an independent component analysis, you find a few things. Even at rest you can discern multiple correlated networks, not just the DMN. Further, you can also discern the DMN during task data. Finally, the set of networks discernible in rest data and task data look like they are basically identical. So it doesn't seem like the DMN (or any other network) is especially associated with rest (Smith et al. 2009). I don't think

¹See (Fox et al. 2005). Note also, however, that the claims of strict anticorrelation might be an artifact of the statistical methods used; see (Cole, Smith, and Beckmann 2010) for a good review of the controversy.

this is terribly surprising, especially if you think that functional connectivity supervenes in part on relatively stable structural connections (Bullmore and Sporns 2012).

It's worth re-emphasizing that the whole brain is always active, even at rest. Indeed, Raichle and Snyder think this now too:

“While initially attributed specifically to a specific system, now often called the *default network*, we now appreciate that all parts of the brain exhibit a *default mode* of functioning that largely reflects their ongoing intrinsic activity.” (2007 1088)

Note, though, that we've jumped off the second horn of the dilemma back onto the first. If there really is this ongoing intrinsic activity, it becomes completely opaque how to study it using neuroimaging.

Third: a larger methodological problem. Return to Morcom and Fletcher's dilemma. In elaborating it, they claim that:

In fact, we can find nothing wrong with the simple hypothesis that the high level of neural activity at rest and in other 'low-level' tasks in some brain regions represents a greater level of processing than that engaged in the same regions during what experimenters might suppose to be more 'demanding' tasks. (2007 1080)

Here's a way to take this: All of this stuff about the DMN is great, but what do people

do when they want to tell you what the DMN does? They don't look at rest itself. They go back and look at task-related activations, and read that back into rest.

So one of the things that the DMN is supposed to do is self-reflection. That's plausible—people are narcissistic, so it wouldn't be surprising that when they're laying in the scanner they think about themselves a lot. But note the logic of this. Why do we think that rest involves self-reflection? Because of studies that look for a difference in activation between equally demanding self- and other-reflective activity (Gusnard et al. 2001). Those studies drive the interpretation of rest. So it looks like the interpretive arrow goes the wrong way. We wanted rest to help us understand activation, but when it comes down to the resting state networks, we go back to activation to interpret rest.

Fourth and finally, I want to suggest a big picture problem. I motivated all of this by the idea that rest might be an interpretive baseline. This itself might be a mistake. If you care about tasks, it might seem natural to look to a non-task state to understand them. This makes sense for many human artifacts. That's because artifacts are often *inert* at rest. Factory robots, say, don't perform their function when they're sitting idle. By studying the idle state we thus learn a lot about why *changes* from the idle state are important. As I've been emphasizing, however, the brain isn't really idle at rest. Neurally speaking, we shouldn't really expect a deep distinction between rest and tasks. The idea of finding this rest baseline is kind of a red herring.

To return one last time to the heart rate analogy: it's not as if when we go to find out what the heart is doing, we pick a state like sitting quietly and say that that's the special baseline back to which everything is referenced. Instead, what happens is that the heart

is doing something, what it's doing is appropriate to the task you're doing, and if you change your task it shifts to a different, still appropriate, level of activation. So maybe what's going on in the brain is like that. We want to find out why changes happen. That requires looking at what was going on *before* a task was performed and figuring out why there was a change, not looking for a privileged baseline. But if that's the case, then looking to the DMN to interpret activation is, unsurprisingly, a nonstarter.

3 The Positive View

3.1 Cognition at Multiple Temporal Bands

So much for the attack. I'd like to offer a positive suggestion.

First, I want to make a claim that I hope is uncontroversial. Mental processes can be both short- and long-term. Since mental processes are instantiated by brain processes, that means that brain processes can be both short- and long-term as well. Part of the job of psychology, in the broadest sense, is to disentangle processes that occur at what Newell called different temporal bands (1990). Of course, long-term processes are composed of short-term ones, and much of what task-based psychology does is study those short-term components. But it's also true that long-term processes provide the context for the short-term ones. This is a familiar point from perception. The sensations from being in a loud dark tube is going to strike you very differently if you know you're in an fMRI experiment versus if you're at a water park versus if you just woke up from a

bender.) So understanding long-term processes may also provide clues to understanding the short-term ones.²

The distinction between long- and short-term, as I'm using it here, is not a deep metaphysical distinction. I think it's best drawn relative to an instrument. In neuroimaging, short-term processes are on the timescale that you can investigate in an fMRI experiment. (They correspond to the timescale of the regressors used in the model). Long-term ones are things that are modeled, during experiments, as things like traits—or, in general, as stable over the experimental situation.

Given that, here's a proposal about how to understand rest. The resting state is just the state where activation is dominated by long-term processes. 'Dominated' here means that the long-term processes are explanatorily most relevant to behavior. As there are no particular short-term processes that are especially urgent, most of the short-term processes that occur are in the service of a long-term process. When you perform a task, the shift is to a state where the short-term processes that occur are all in the service of relatively short-term tasks.

Thinking of rest this way brings to the fore a number of usually overlooked features of rest in fMRI. Most obviously, note that even when subjects are laying quietly in the scanner, they're still doing something: they're *participating in an experiment*. In other words, 'rest' itself is a bit of a mischaracterization: there's not really a point where subjects are doing nothing. That depends on decisions made, and instructions given, in the span of minutes to hours before the rest block itself. Those instructions must be

²Here I have been influenced, albeit very loosely, by (Dewey 1896).

maintained. They also affect how subsequent perceptions are filtered, interpreted, and acted upon. Because they're in an experiment, they also have to suppress other desires they might have: to move their head, to get up and walk around, and so on.

While many long-term states can be understood in personal-level terms, 'long-term' shouldn't be restricted to personal-level processes. The process of keeping vigilant for task-related cues, for example, should be expected to recruit numerous subpersonal processes. Further, response contingencies can change due to unexpected events: the utility of staying in the scanner, say, drops radically if the fire alarm sounds. There are also numerous cognitive processes that are obviously long-term—skill acquisition or memory consolidation, for example, depend on subpersonal changes that occur on the span from hours to weeks.

Note that the distinction between short-term and long-term maps poorly onto the other distinctions that have been made. So it's not the same as intrinsic versus extrinsic—no interesting important sense of intrinsic here, because there's always a task that modulates. Similarly not the distinction between exogenous and endogenous motivation, or between stimulus-driven vs non-stimulus-bound attention, or to purposeful vs non-purposeful behavior. These distinctions make sense for short term mental processes, but get hopelessly muddled for long-term ones. When someone is laying quietly in a scanner, they are motivated *both* by the initial instructions *and* by their internal motivations. Given your long-term plans, you're *always* engaged in some sort of purposeful behavior (Even if my desire is to go for a relaxing-mind clearing walk on which I don't think about anything in particular, I'm still carrying out my intention to

go for a walk, and that intention shapes my subsequent actions). Similarly, there's no sense in which the activity at rest represents 'intrinsic' activity, at least if 'intrinsic' is taken to mean 'activity in the absence of *any* plans or stimuli.' Again, you're always doing *something*.

3.2 The Coordinative Role of the DMN

This also explains the intuition that motivated the standard DMN story. It is, I suspect, quite true that people think about their upcoming court dates and haircut appointments and whatnot while they lie bored in the scanner. That's not equivalent to the meaningless chugging noise that an idling engine throws off, however. People often think about things like that precisely to advance their long-term plans. Again, even on this conception rest is just the state where longer-term processes are *explanatorily* most relevant to a subject's behavior.

Suppose that the story I've told is right. I think this might also shed some light on why we see activity decreases in the DMN when subjects move from rest to focused tasks. The picture I've sketched is one on which we move back and forth between short-term activities that promote specific goals and longer-term processes that may or may not require immediate action. The key suggestion is that the DMN may play a tonically active coordination role that facilitates these switches.

The Posterior Cingulate Cortex in particular detects changes in the environment that signal changes in the rewards associated with certain actions (Pearson et al. 2011). The

posterior cingulate/precuneus also appears to play an important role in routing information around in the cortex in response to differing task demands (Fransson and Marrelec 2008). Similarly, medial prefrontal cortex plays a crucial role in suppressing many of the possible goals that we might pursue at one moment in order to promote one or a handful of other goals (Miller and Cohen 2001). The medial frontal cortex plays a well-known role in suppressing many of the possible goals that we might pursue at one moment in order to promote one or a handful of other goals (Miller and Cohen 2001). Evidence from so-called “warned reaction time” tasks—where subjects have to wait for a tone that signals an upcoming task—suggests the involvement of the medial PFC (Jaffard et al. 2008). Linking the medial PFC directly to the demands of neuroimaging experiments, Jaffard et al. note that

... the fact that, in passive viewing experiments, subjects are explicitly instructed to “refrain from moving and reacting” ... strongly supports our suggestion that the activity at rest may be partly due to an active and sustained process consisting of locked movement initiation mechanisms. The default mode of this “gating” mechanism would be a tonic inhibition state. ((Jaffard et al. 2008) 1204)

The posterior cingulate/precuneus similarly appears to play an important role in routing information around in the cortex in response to differing task demands (Fransson and Marrelec 2008). Decreased activity in the DMN, therefore, might represent the deactivation of a tonically active system which is necessary for coordinating the variety of different goals, plans, and projects that we might be engaged in at any particular

moment. Further, if this coordination were to break down, one would expect a variety of characteristic disorders—and indeed, various sorts of DMN dysfunction have been linked to conditions like autism, schizophrenia, and dementia (Broyd et al. 2009). A defining feature of many of these disorders is a failure of coordination: between fantasy and reality, or between immediate sensory input and longer-term desires and goals.

A full defense of this story is beyond the scope of this paper, and would require a detailed story about the expected relationship between excitatory/inhibitory balance in DMN regions and expected BOLD responses.³ DMN activity here need not be *distinctive* of rest, but it will play a crucial role in doing what needs to be done during rest. So to recap: rest is important precisely because it’s a state explanatorily dominated by long-term processes. Puzzles arise about rest only when we lost that perspective and think that psychological tasks are the most important unit for investigating mind.

3.3 Solving the Dilemma

If the above is sensible, here’s a nice answer to Morcom and Fletcher. They asked why we care about rest, and how we could possibly study it. This provides a bit of an answer. We care about rest because we care about long-term processes too, and we ought to care about them because they provide context for the short-term ones. This doesn’t require studying *either* the resting state or changes in activation over the short-term. Instead, to

³For an intriguing possibility along these lines, see (Carhart-Harris and Friston 2010). I’m neutral on the Freudian cast of the article, but it seems that the linkage to predictive coding frameworks might be a useful integrative principle framework for thinking about the interaction between shorter and longer temporal bands.

study long-term processes, the most telling data will be those which show *interaction effects* between short and long term processes. That is, we want to find evidence that long-term processes have a differential effect on the short-term ones in some domain.

Here's an example of that methodology in action. Price and Devlin investigated word reading (2011) . They were especially concerned with activation in the posterior lateral fusiform gyrus (PLF).⁴ This is a region of the brain where, if you read writing in a familiar script versus visually matched non-script characters, you find an increase in activation. Just focusing on reading, however, Price and Devlin looked at the interaction between reading and fluency.

If we just look at the short-term, PLF tends to show more activity during reading compared to matched non-reading tasks. Price and Devlin found an interaction effect between short- and long-term processes. In particular they found an inverted 'U' shape plotting relative activation as a function of fluency. What happens is that non-fluent readers show relatively low activation when they see words. During early learning, PLF shows the most activation, and then it drops down again as people become more fluent. Note that if you just divided readers versus non-readers (lumping the final two stages), you get an increase in reading over non-reading, supporting the idea that PLF just has reading-related function. But this short-term effect masks a more interesting interaction effect: as people become more fluent, the total amount of activation rises and then drops.

⁴Note that in Price and Devlin it is referred to as ventral occipitotemporal cortex; I refer to it as PLF for consistency with earlier work. PLF activation isn't specific to reading; it's seen in an enormous number of other non-reading contrasts (Price and Friston 2005). I assume that within-domain contrasts are reasonable bases, however, for theory-building.

The interaction effect allows for a more interesting functional ascription. Price and Devlin propose a so-called *interactive account* of PLF. The function of PLF is to integrate both bottom-up information and top-down error signals. The BOLD signal is a function of both sorts of signal.⁵ So activation in PLF is driven both by reading per se as well as feedback about reading from higher-level predictive areas in cortex. In early learning, then, you see the highest activation because error feedback is at its peak; fluency reduces the feedback signal.

Note that this allows an ascription of function to PLF that is more stable over time. The long-term function of PLF is something like tuning its own response properties to better predict the visual world and feed that to higher-order predictive centers. (Note too that this will apply not just to reading, but other processes in which the PLF is involved) This is not a process that is done over the short-term per se, though it does involve some relatively short-term changes. The full functional picture only emerges when you consider these short-term processes as they interact with longer-term context.

So that's the pitch: if you care about rest, for the reasons I've outlined, you should look for interaction effects. That may sound a bit puzzling. It has the odd consequence that if you care about rest, studying what's going on at rest isn't terribly important by itself (though 'rest' may be a conceptually foundational notion, as the state dominated by long-term effects.) But unlike on Morcom and Fletcher's dilemma, we have a story now about why rest matters. Further, it's not as if forward inference is just dominated by discussion of short-term effects. Rather, it is the interaction between short-term and

⁵See (Bartels, Logothetis, and Moutoussis 2008). For a similar model of binocular rivalry, as well as a discussion about the general applicability of such feedback models, see (Hohwy, Roepstorff, and Friston 2008).

long-term processes that provide a more satisfying basis for inference about function.

References

- Bartels, Andreas, Nikos K. Logothetis, and Konstantinos Moutoussis. 2008. "fMRI and its interpretations: An illustration on directional selectivity in area V5/MT." *Trends in Neurosciences* 31 (9): 444–453.
- Broyd, Samantha J, Charmaine Demanuele, Stefan Debener, Suzannah K Helps, Christopher J James, and Edmund JS Sonuga-Barke. 2009. "Default-mode brain dysfunction in mental disorders: A systematic review." *Neuroscience & biobehavioral reviews* 33 (3): 279–296.
- Buckner, Randy L., and Daniel C. Carroll. 2007. "Self-projection and the brain." *Trends in Cognitive Sciences* 11 (2): 49–57.
- Bullmore, Ed, and Olaf Sporns. 2012. "The economy of brain network organization." *Nature Reviews Neuroscience* 13 (5): 336–349.
- Carhart-Harris, Robin L, and KJ Friston. 2010. "The default-mode, ego-functions and free-energy: a neurobiological account of Freudian ideas." *Brain* 133 (4): 1265–1283.
- Cavanna, Andrea E., and Michael R. Trimble. 2006. "The precuneus: A review of its functional anatomy and behavioural correlates." *Brain* 129 (3): 564–583.
- Cole, David M, Stephen M Smith, and Christian F Beckmann. 2010. "Advances and pitfalls in the analysis and interpretation of resting-state FMRI data." *Frontiers in systems neuroscience* 4:1–15.
- Dewey, John. 1896. "The reflex arc concept in psychology." *Psychological review* 3 (4): 357–370.
- Fox, Michael D, Abraham Z Snyder, Justin L Vincent, Maurizio Corbetta, David C Van Essen, and Marcus E Raichle. 2005. "The human brain is intrinsically organized into dynamic, anticorrelated functional networks." *Proceedings of the National Academy of Sciences* 102 (27): 9673–9678.
- Fransson, Peter, and Guillaume Marrelec. 2008. "The precuneus/posterior cingulate cortex plays a pivotal role in the default mode network: Evidence from a partial correlation network analysis." *Neuroimage* 42 (3): 1178–1184.
- Gusnard, Debra A., Erbil Akbudak, Gordon L. Shulman, and Marcus E. Raichle. 2001. "Medial prefrontal cortex and self-referential mental activity: relation to a default mode of brain function." *Proceedings of the National Academy of Sciences* 98 (7): 4259–4264.
- Gusnard, Debra A., and Marcus E. Raichle. 2001. "Searching for a baseline: Functional imaging and the resting human brain." *Nature Reviews Neuroscience* 2 (10): 685–694.

- Hohwy, Jakob, Andreas Roepstorff, and Karl J Friston. 2008. "Predictive coding explains binocular rivalry: An epistemological review." *Cognition* 108 (3): 687–701.
- Jaffard, Magali, Marieke Longcamp, Jean-Luc Velay, Jean-Luc Anton, Muriel Roth, Bruno Nazarian, and Philippe Boulinguez. 2008. "Proactive inhibitory control of movement assessed by event-related fMRI." *Neuroimage* 42 (3): 1196–1206.
- Miller, Earl K., and Jonathan D. Cohen. 2001. "An Integrative Theory Of Prefrontal Cortex Function." *Annual review of Neuroscience* 24 (1): 167–202.
- Morcom, Alexa M, and Paul C Fletcher. 2007. "Does the brain have a baseline? Why we should be resisting a rest." *Neuroimage* 37 (4): 1073–1082.
- Nair, Dinesh G. 2005. "About Being BOLD." *Brain Research Reviews* 50:229–243.
- Newell, Allen. 1990. *Unified Theories of Cognition*. Cambridge: Harvard University Press.
- Pearson, John M, Sarah R Heilbronner, David L Barack, Benjamin Y Hayden, and Michael L Platt. 2011. "Posterior cingulate cortex: Adapting behavior to a changing world." *Trends in cognitive sciences* 15 (4): 143–151.
- Poldrack, Russell A. 2006. "Can cognitive processes be inferred from neuroimaging data?" *Trends in Cognitive Sciences* 10 (2): 59–63 (Feb).
- Price, Cathy J., and Joseph T. Devlin. 2011. "The Interactive Account of ventral occipitotemporal contributions to reading." *Trends in cognitive sciences* 15 (6): 246–253.
- Price, Cathy J., and Karl J. Friston. 2005. "Functional ontologies for cognition: The systematic definition of structure and function." *Cognitive Neuropsychology* 22 (3): 262–275.
- Raichle, Marcus E., Ann Mary MacLeod, Abraham Z. Snyder, William J. Powers, Debra A. Gusnard, and Gordon L. Shulman. 2001. "A default mode of brain function." *Proceedings of the National Academy of Sciences* 98 (2): 676–682.
- Raichle, Marcus E, and Mark A Mintun. 2006. "Brain work and brain imaging." *Annual Review of Neuroscience* 29:449–476.
- Raichle, Marcus E, and Abraham Z Snyder. 2007. "A default mode of brain function: A brief history of an evolving idea." *Neuroimage* 37 (4): 1083–1090.
- Smith, Stephen M, Peter T Fox, Karla L Miller, David C Glahn, P Mickle Fox, Clare E Mackay, Nicola Filippini, Kate E Watkins, Roberto Toro, and Angela R Laird. 2009. "Correspondence of the brain's functional architecture during activation and rest." *Proceedings of the National Academy of Sciences* 106 (31): 13040–13045.