

Mixtures and Psychological Inference with Resting State fMRI

Joseph McCaffrey and David Danks

Abstract

In this essay, we examine the use of resting state fMRI data for psychological inferences. We argue that resting state studies hold the paired promises of discovering novel functional brain networks, and of avoiding some of the limitations of task-based fMRI. However, we argue that the very features of experimental design that enable resting state fMRI to support exploratory science also generate a novel confound. We argue that seemingly key features of resting state functional connectivity networks may be artifacts resulting from sampling a ‘mixture distribution’ of diverse brain networks active at different times during the scan. We explore the consequences of this ‘mixture view’ for attempts to theorize about the cognitive or psychological functions of resting state networks, as well as the value of exploratory experiments.

1. *Introduction*
2. *Functional Network Discovery in Resting State fMRI*
3. *Psychological Inference and Mixture Artifacts*
4. *Arguments For and Against the Mixture View*
 - 4.1 *Objections to the mixture view*
5. *Finding Value for the Resting State*
 - 5.1 *New tools and techniques*
 - 5.2 *Exploration via the resting state*
6. *Conclusion*

1 Introduction

Functional magnetic resonance imaging (fMRI) is an immensely popular tool for studying human cognition. In standard task-based fMRI studies, neuroscientists manipulate tasks and measure resulting changes in the blood-oxygen-level-dependent (BOLD) signal within voxels, or small volumes of brain analogous to pixels in a digital image). In a standard subtraction design (Posner *et al.* [1988]; Friston *et al.* [1996]), neuroscientists compare BOLD measurements from a psychological task of interest (reading a meaningful word like ‘DOG’) to those of a matched control task that ideally differs in only one psychological component or process (reading a non-word like ‘BLORT’). Significant differences in BOLD signals are used to localize a given psychological process (semantic processing) to particular voxels or brain regions.¹ This kind of function-to-structure inference is known as ‘forward inference’ (Henson [2006]; Machery [2012]). Given mappings established by these experiments, neuroscientists who observe similar activation in a subsequent experiment often infer that the new experiment involves that psychological process. This form of structure-to-function inference is known as ‘reverse inference’ (Poldrack [2006]; Nathan and Del Pinal [2016]).

¹ Cognitive subtraction in neuroimaging was originally developed for early positron emission tomography (PET) studies of language (as in Posner *et al.* [1988]) before being applied to fMRI (as in Kanwisher *et al.* [1997]).

Philosophers and neuroscientists have extensively debated the experimental and inferential logics employed in subtractive, task-based neuroimaging (Friston *et al.* [1996]; Van Orden and Paap [1997]; Coltheart [2006]; Poldrack [2006]; Roskies [2009]; Klein [2010]; Machery [2012], [2013]; Glymour and Hanson [2015]; Nathan and Del Pinal [2016]). One debate is whether the assumption of ‘pure insertion’—namely, that task pairs differ solely in the recruitment of one identifiable, independently modifiable psychological component—is valid (Friston *et al.* [1996]; Van Orden and Paap [1997]; Roskies [2010]). Another is whether region-based reverse inference is legitimate given that brain regions appear to be highly multi-functional (Poldrack [2006]; Anderson [2010]; Klein [2012]; Machery [2013]; McCaffrey [2015]; Nathan and Del Pinal [2016]). In addition, new experimental designs and analytic tools (such multi-voxel pattern analysis, MVPA) are emerging that seemingly have different strengths and limitations than traditional designs, and so avoid some philosophical questions while prompting others (Glymour [2015]; Del Pinal and Nathan [2017]; Wright [forthcoming]). More generally, as novel neuroimaging methods arise, we need to ask whether they present novel challenges or issues, and also whether they escape or inherit the challenges of older methods.

Over the past twenty years, a new type of experimental design—resting state fMRI—has emerged from the puzzling observation that many brain regions exhibit correlated, low frequency (~0.01-0.10 Hz) BOLD fluctuations while participants are ‘resting’ between task blocks (Biswal *et al.* [1995]; Binder *et al.* [1999]; Gusnard and Raichle [2001]; Raichle *et al.* [2001]). This surprising, accidental finding in the 1990’s led to the subsequent characterization of the so-called ‘default mode network’ (—a set of regions whose activity is correlated at rest, but uncorrelated during specific tasks (Raichle *et al.* [2001]; Greicius *et al.* [2003])—and explorations of the brain’s ‘intrinsic’ dynamics in the absence of explicit experimental tasks (Snyder and Raichle [2012]). In general, resting state fMRI designs (i) measure BOLD correlations between brain areas in participants given no instructions other than to stay awake; and (ii) use those correlations, perhaps after some processing, to infer functional connectivity networks’—that is, networks of highly correlated brain regions (Power *et al.* [2014]).

Resting state fMRI is a rapidly burgeoning area of clinical and basic research. Power, Schlagger, and Petersen ([2014], p. 692) observe that ‘resting state fMRI has grown from an unexpected observation in fMRI ‘noise’ to a major area of human neuroimaging.’ At the same time, it is not obvious what resting state studies reveal about the mind or brain. Modestly, they tell us which voxels or regions show correlated BOLD activity during this type of rest. More ambitiously, these correlations could be a way to discover possible anatomical connections (as suggested by Biswal *et al.* [1995]; Bray *et al.* [2015]). And many theorists (Snyder and Raichle [2012]; Andrews-Hanna *et al.* [2014]; Klein [2014]) want to go further and use resting state fMRI to make inferences about the psychological functions of particular resting state networks, despite the lack of experimental control (of psychological states or processes) in these studies.

In this paper, we examine psychological inferences drawn from resting state fMRI research, such as ‘network *X* is for autobiographical memory’; ‘changes in network *Y* reflect sensorimotor learning’; or ‘differences in network *Z* underlie the cognitive deficits observed in schizophrenia.’ In short, our targets are the various forward or reverse inferences made using resting state networks. We make two arguments. First, resting state fMRI has the potential to reveal new aspects of the brain’s functional architecture in a (somewhat) bottom-up fashion. Resting state designs are potentially a kind of exploratory experiment (Steinle [1997]; Franklin [2005]) that enables discovery of both cognitive functions and relevant functional brain units without strict task designs (Biswal *et al.* [2010]; Snyder and Raichle [2012]). This approach

contrasts with cognitive subtraction designs, which rely on either strict control of functions/tasks to discover relevant brain units (forward inference), or assume knowledge of the unit-to-function mappings to discover what functions are part of a designated task (reverse inference) (Van Orden and Paap [1997]; Roskies [2010]). Resting state experiments are not, we will suggest, entirely bottom-up or theory free, as they aim to contribute to cognitive theorizing. Nonetheless, they require many fewer top-down commitments or constraints than most fMRI studies.

Second, and more importantly, we develop a novel challenge for interpretations of resting state fMRI data. Resting state research often aims to discover the function(s) of large-scale (correlation) networks that are consistently identified across individuals. We argue, however, that these networks plausibly contain sampling artifacts that do not correspond to underlying brain connections. More precisely: (1) resting state data plausibly result from sampling (over time) a mixture distribution composed by multiple smaller, truly functional networks engaged at different times; and so (2) some (though surely not all) observed correlations are spurious and do not correspond to an underlying causal or functional relation; but (3) the data alone do not tell us which correlations are spurious and which are not. Thus, the very experimental features that reveal new networks in resting state fMRI—namely, participants engaging in many different cognitive processes over a period of time—create a new problem for psychological inferences about resting state data. There are other challenges for cognitive inferences from resting state data that are independent of our central concern, but we focus here on the novel problem of mixtures.

We begin in Section 2 with an introduction to resting state fMRI, including its methodology and possible uses. Section 3 presents the mixture view that resting state networks may involve sampling artifacts rather than only genuine features of the brain's functional anatomy. This view thus presents novel complications that have not been discussed in previous philosophical work on the possibility of artifacts in neuroimaging research (Roskies [2007]). Section 4 examines the mixture view in more depth, including evidence in its favor (though novel experimental validation falls outside of the scope of the present paper) and prominent objections. Section 5 turns to the more positive project of examining two ways that neuroscientists might respond to the mixture challenge, both of which reopen the possibility of using resting state fMRI as a 'discovery science' for functional brain anatomy (Biswal *et al.* [2010]). Resting state fMRI is a novel methodology with intriguing results, but with significant challenges to its use for psychological inferences.

2 Functional Network Discovery in Resting State fMRI

Task-based fMRI studies employ a clear strategy: conduct controlled psychological experiments, measure BOLD signal changes, and then draw functional inferences about regions with changed activation. This strategy could enable one to discover what regions are involved in recognizing a face (Kanwisher *et al.* [1997]) or whether two different memory tasks recruit the same brain regions (Henson [2006]). The success of these projects may be debatable (Coltheart [2006]; Roskies [2009]), but the goals are clear. In contrast, the aims of resting state fMRI (rsfMRI) are less immediately clear, as those studies are 'uncontrolled according to the usual conventions that apply to cognitive neuroimaging' (Snyder and Raichle [2012], p. 902). We propose that they can advance a 'discovery science' of functional brain anatomy (Buckner *et al.* [2008]; Biswal *et al.* [2010]), but that claim requires more detail about the methodology.

In a typical rsfMRI study, participants lie passively in the scanner, and researchers examine how the BOLD time series for different seed regions (predefined regions of interest) correlate with one another over timescales of several minutes to an hour, perhaps using various

thresholds or network detection algorithms to extract robust network structure (Beckman *et al.* [2005]; Fox *et al.* [2005]; Power *et al.* [2014]). The resulting functional connectivity networks are typically reported as pairwise correlation coefficients between seed regions, and are often depicted as graphical networks superimposed on an anatomical image of the brain (see Bullmore and Sporns [2009]). We use the term ‘resting state network’ (RSN) to refer to these functional connectivity patterns observed in rsfMRI experiments. Importantly, the term ‘functional connectivity’ is potentially misleading here, as these networks only capture correlational (as opposed to causal or anatomical) structure. Causal relationships are encoded in so-called effective connectivity networks (Friston [2011]), which require different inference techniques.

There is ample evidence that RSNs reflect real neural activity, and are not merely the product of head motion, cardiac cycle, or other obvious confound (Glover and Lee [1995]; Lowe *et al.* [1998]; He *et al.* [2008]; Johnston *et al.* [2008]). RSNs appear to be altered in clinical and psychiatric populations (Zhou *et al.* [2007]; Wu *et al.* [2009]), and can be used to predict the severity of stroke-related cognitive deficits (Warren *et al.* [2014]). They may thus have prognostic or diagnostic value, independently of any other value. But we might hope for more, as different studies have consistently identified RSNs corresponding to known anatomical or functional brain networks. For example, there are RSNs corresponding roughly to the visual system, motor system, executive system, and so forth (Damoiseaux *et al.* [2006]; Fox and Raichle [2007]; Smith *et al.* [2009]; Cole *et al.* [2014]; Bray *et al.* [2015]). More generally, areas one would expect to be coordinated during tasks, such as different hemispheres of the motor cortex, tend to correlate during rest (Biswal *et al.* [1999]; Fox *et al.* [2005]). This suggests some as-yet-unknown relationship between RSNs and the causal/functional brain networks that perform sensory, motor, and cognitive functions.

Many researchers are thus interested in the potential cognitive functions of RSNs (Andrews-Hanna *et al.* [2008]; Vahdat *et al.* [2011]; Uddin [2015]), such as imagination (Mason *et al.* [2007]), mindwandering (Christoff *et al.* [2016]), top-down attention (Markett *et al.* [2014]), and so on. But how can we establish these mappings when we have no control over the participants’ psychological processes? Philosophically, it is useful to think of rsfMRI studies as exploratory experiments (Steinle [1997]; Franklin [2005]; Biswal *et al.* [2010]) that search for meaningful patterns in data without definite prior hypotheses about what might cause those patterns. In genetics, for example, researchers sometimes aim to learn about gene regulatory networks by measuring large numbers of mRNA transcripts in target cells rather than using background knowledge to design targeted experiments (Basso *et al.* [2005]). Similarly, in rsfMRI, researchers seek to learn about brain activation patterns without a strong background cognitive theory, and have thereby seemingly revealed novel networks not reliably found in task-based studies, such as the original default mode network (DMN) (Biswal *et al.* [1995]; Shulman *et al.* [1997]; Gusnard and Raichle [2001]; Greicius *et al.* [2003]; Power *et al.* [2013]). rsfMRI thus might be useful as a discovery tool (Damoiseaux *et al.* [2006]; Biswal *et al.* [2010]) that provides a new way to identify functional brain networks.

Meaningful functional divisions typically cannot be identified solely based on anatomical divisions in the brain (sulci, gyri, cytoarchitectural divisions, and so on). Instead, we have historically used lesion studies and low-level neurophysiology studies, and more recently task-based fMRI studies, to find evidence of functional localization. For example, various fMRI studies have suggested (though not without debate) plausible psychological functions for the fusiform face area (FFA) (Kanwisher *et al.* [1997]), the right temporoparietal junction (Saxe and Kanwisher [2003]), the visual word form area (Cohen *et al.* [2000]), and many other cortical

regions. Discovery of such mappings in task-based fMRI is parasitic on our ability to design carefully controlled task pairs that target a single psychological process (Van Orden and Paap [1997]; Roskies [2009]). Thus, the search for new brain mappings in task-based imaging is highly constrained by our existing cognitive models and tasks (Van Orden and Paap [1997]; Price and Friston [2005]). In contrast, rsfMRI seems to have the potential to reveal functional brain networks without carefully controlled task pairs inspired by existing cognitive models. rsfMRI seems to provide a ‘bottom-up’ (unconstrained by cognitive theorizing) way of doing human brain mapping.

As an example, the DMN was initially observed during rsfMRI experiments, but with little idea about its potential functions (Greicius *et al.* [2003]). Subsequent task-based studies now implicate the DMN in numerous cognitive functions related to internally-guided thoughts such as imagination or mindwandering (Mason *et al.* [2007]; Spreng and Grady [2010]; Andrews-Hanna *et al.* [2014]; Christoff *et al.* [2016]). We thus have a reversal of the usual approach to brain mapping: researchers first identified a potential ‘functional’ brain network, and only then asked what function(s) it performs. Of course, this hypothesis-suggesting role for rsfMRI is quite limited, as seen in a dilemma posed by Morcom and Fletcher ([2007], see also Klein [2014]): if rsfMRI reveals RSNs that cannot be studied with tasks, then we seemingly cannot determine their function; but if these networks can be studied with tasks, then rsfMRI studies are ultimately superfluous. We contend that adjustments to our view of exploratory experiments can enable us to escape their dilemma.

In general, we argue that a key dimension characterizing exploratory experiments is (lack of) experimental control. Experiments can be more or less controlled, and thus there is not a strict distinction between ‘hypothesis-driven’ and ‘exploratory’ experiments. Moreover, loosening experimental control can sometimes permit the discovery of novel patterns that would be time-consuming, difficult, or perhaps impossible to observe in more controlled settings. For example, O’Keefe and Dostrovsky ([1971]) observed correlations between neural firing rates and freely-performed naturalistic behaviours such as walking, and thereby found initial evidence for rat hippocampal place cells. Hypotheses played a role in this exploratory work, as previous lesion work suggested the hippocampus was a reasonable place to find encodings of environmental features. However, their experiments did not force the rats to perform a particular task, and so could reveal more neural-behaviour correlations than a more strictly controlled experiment.

For reasons of space, we do not provide a full philosophical analysis of ‘control’ in this paper, but we need only a high-level characterization for our present purposes. Experimental control is a multidimensional notion, including variation in terms of (at least): magnitude of control; precision of control; specificity of the target of control; scope or breadth of control; and context sensitivity of control abilities. For our present purposes, it suffices to note that rsfMRI involves reduced control of the participant’s cognitive processes along every one of these dimensions; for example, the lack of any specific task means that rsfMRI involves small magnitude or extent of control, as the experimenter does not have any substantial influence on the participant’s cognitive processing.

Prior philosophical work has often characterized exploratory experiments as those that: (1) are not hypothesis-driven, and (2) involve wide instrumentation or simultaneous measurements of numerous variables (Steinle [1997]; Franklin [2005]). However, rsfMRI studies are not wholly ‘theory free’, and are often constrained by prior cognitive theory. For instance, experimental explorations of the DMN are guided by theoretical considerations about what participants are likely doing during their time in the scanner (Christoff *et al.* [2016]). Indeed, linking RSNs to

cognitive functions necessarily requires some degree of background theory. Nonetheless, rsfMRI experiments should still be understood as exploratory studies, precisely because they are far less controlled than subtraction designs. This lack of control can be used to identify novel networks (as with the DMN), or to broaden the search for correlations between behavioural measures and brain network topology. Some concrete examples can help to demonstrate that reduced experimental control, not elimination of theory, is the key to exploratory experiments (at least, for rsfMRI).

Vahdat *et al.* ([2011]) conducted rsfMRI scans of arm-related somatosensory and motor regions before and after participants learned (outside of the scanner) a novel reaching task with sensory and motor components. They determined correlations between neuroplasticity in RSNs and changes in either motor or perceptual performance, and thereby found a novel network corresponding to perceptual changes accompanying motor learning. In a standard task-based design, participants would perform the same reaching task before and after learning, and so BOLD activation levels would not distinguish changes due to motor execution versus perceptual learning. By contrast, comparison of RSNs to performance enabled the researchers to decouple these effects. Similarly, Markett *et al.* ([2014]) found correlations between rsfMRI-derived topological features of the frontoparietal attention network (FPAN) and performance on various attentional tasks (such as low centrality in two regions predicted better performance on alerting attention tasks). The focus on FPAN was driven by prior theory, but the novel potential connection between brain network topology and task performance depended precisely on not controlling participants in a task-based design way.

Neither of these rsfMRI studies was wholly bottom-up or theory free; they both involved many theory-laden decisions, particularly about focal brain regions and behaviours.² However, in each case, neuroscientists identified novel patterns of correlations by relaxing the degree of experimental control. That is, rsfMRI studies functioned as exploratory experiments that uncovered correlation patterns that would have been difficult or perhaps even impossible to discover on a task-by-task basis. However, this conclusion raises a crucial question: do the functional connectivity (correlation) patterns observed at rest really signify the kinds of brain structures onto which psychological functions map? Unfortunately, there are significant grounds for skepticism, despite the tantalizing correlations uncovered so far.

3 Psychological Inference and Mixture Artifacts.

rsfMRI studies reliably reveal a set of large-scale (correlational) RSNs that often mirror known functional brain systems (Damoiseaux *et al.* [2006]; Fox and Raichle [2007]), and sometimes suggest the existence of novel ones (Power *et al.* [2011]; Uddin [2015]). These findings suggest that RSNs are somehow (cognitively) functionally relevant, and thus there is increasing interest in characterizing the cognitive functions of RSNs (Mason *et al.* [2007]; Spreng and Grady [2010]; Vahdat *et al.* [2011]; Markett *et al.* [2014]). In this section, though, we argue that RSNs obtained in rsfMRI plausibly include some connections (though we do not know which) that are sampling artifacts produced by sampling from a mixture distribution of smaller networks active at different times during the scan. That is, seemingly key features of RSN *X* may reflect sampling artifacts rather than real features of the brain's underlying causal or cognitive organization, and this discrimination requires knowledge about which brain networks have coherent psychological functions. That is, reliable inferences about the psychological functions of RSNs require prior

² Thanks to an anonymous reviewer for emphasizing this point.

knowledge of the very psychological functions we aim to infer. We will use the term ‘mixture view’ to refer to this collection of concerns.

Typical psychological inferences in neuroimaging—in particular, that network X performs function Y —require a clear target region (Poldrack [2006]), network (Klein [2012]; Glymour and Hanson [2015]), or activation pattern (Del Pinal and Nathan [2017]). We need to know the ‘structure’ whose cognitive function is being inferred to (in the case of forward inference) or from (in the case of reverse inference). RSNs seem to provide a new set of targets for functional inference, such that researchers could propose ‘network X is for mindwandering’ (a forward inference) or ‘network Z is currently engaged, so the participant is likely engaged in mental imagery’ (a reverse inference). Non-mixture concerns have previously been raised about whether RSNs provide such targets. First, RSNs are sometimes thought to reflect baseline activity that is unrelated to conscious cognition (Raichle *et al.* [2001]), and so are not appropriate targets for psychological inferences.³ This possibility might be surprising, but consider an analogy: when a truck is idling, its engine and drivetrain components are not doing exactly what they do when the truck is driving, but are still informative about how a truck operates. Of course, this possibility depends on whether networks observed during rest are active in other contexts (Klein [2012]). Recent studies report a great deal of similarity in the networks observed at rest versus during tasks (Cole *et al.* [2014]; Yeo *et al.* [2015]), but the precise relationship between occurrent cognitive processes and RSNs remains unclear (Van Calster *et al.* [2016]).

A second concern about RSNs as targets for psychological inference is that they are often fairly large networks associated with many functions. The DMN, for instance, spans much of the cortex and is linked to mental imagery, autobiographical memory, mind wandering, and more (Mason *et al.* [2007]; Andrews-Hanna *et al.* [2014]; Christoff *et al.* [2016]). The DMN might therefore be a ‘brain system’ (a collection of regions or networks that serve distinct, but related functions, such as the visual system) rather than a network performing a unique single function (Poldrack [2006]; Klein [2012]; McCaffrey [2015]). For example, Andrews-Hanna *et al.* ([2014]) recently identified three sub-networks within the DMN, and found that each sub-network was associated (in other experiments) with different cognitive domains (such as semantic or episodic memory). Sub-components of RSNs may be better targets for psychological inference than whole RSNs.

These are legitimate concerns, but we have a different focus here. The debate to date (and most of the neuroscientists we have mentioned) has assumed that RSNs are coherent structures, or can be fully decomposed into coherent structures, that plausibly have distinct psychological functions. The mixture view challenges exactly this assumption. RSNs are statistical (correlational) structures that, we contend, can contain ‘spurious’ or ‘artifactual’ connections in the sense that: (a) the connection does not correspond to any underlying causal relation; and (b) it would not obtain under different sampling or measurement processes. We emphasize that these spurious correlations are truly present in the data; they are ‘spurious’ because they do not reflect any underlying causal or neural connection. The issue is thus not just whether and how we can pin cognitive functions to RSNs, but also whether RSNs faithfully capture the structure of the networks underlying cognition in the first place. Full understanding of this concern requires a bit of a detour into some experimental and statistical details.

³ Of course, the ‘resting state’ may consist of both stable baseline metabolic activity and more transient, heterogeneous activity related to cognition.

Resting state fMRI involves measuring BOLD correlations⁴ over long periods of time, typically more than ten minutes. Furthermore, scanning protocols for rsfMRI sometimes involve longer acquisition times (TRs) than task-based fMRI, since the BOLD fluctuations at rest are slower than most task-related changes (see Power *et al.* [2014]). The sum total of these methodological choices is that rsfMRI captures the brain’s activity over a relatively long time scale. At the same time, participants are plausibly engaged in a number of different cognitive processes during rest. As Morcom and Fletcher ([2007]) note, the ‘resting’ brain is a complex, dynamic state in which various brain regions and networks interact with one another as participants imagine what they will have for dinner, remember going to the dentist last week, attend to the noises of the scanner, and so forth. The resting state signal may thus be significantly driven by dynamic network changes accompanying spontaneous thought or ‘mindwandering’ (Mason [2007]; Christoff *et al.* [2016]; Van Calster *et al.* [2016]). Despite this likely cognitive diversity, most rsfMRI analyses infer a single functional connectivity network from the whole scan/dataset: the RSN reflects correlations that obtain in the complete time series, even though data at different points were plausibly generated by different brain/causal networks. In statistical terms, the full time series is plausibly drawn from a mixture distribution.⁵

Suppose that we have two different univariate probability distributions, $P_1(X)$ and $P_2(X)$. A mixture distribution $P_M(X)$ results from mixing together samples from $P_1(X)$ and samples from $P_2(X)$. Mathematically, we have $P_M(X) = \alpha P_1(X) + (1-\alpha)P_2(X)$, where α is a mixture parameter that encodes the relevant proportion of samples from the component distributions P_1 and P_2 . More generally, a mixture distribution is a probability distribution that results from combining, in a single dataset, samples drawn from distinct probability distributions over the same variables (without encoding provenance of samples). In the case of rsfMRI, the whole time series plausibly includes samples generated by different cognitive processes or networks, each of which corresponds to (or produces) a distinct probability distribution. That is, the rsfMRI time series is generated by sampling across time from an unknown, heterogeneous mixture of region and causal network activation at different times.

It is important to be clear about the nature of these ‘mixtures’. At any given moment, the brain participates in many different psychological and physiological processes. For example, reading a word involves word form recognition, attention, eye movement control, and more, whether simultaneously or in short succession. Thus, it is extremely likely, perhaps necessary, that any functional connectivity network (not just resting state) will include sub-networks corresponding to multiple psychological functions. In fact, cognitive subtraction is intended partly to solve this very issue (Posner *et al.* [1988]), and provides a strategy to find the relevant sub-networks. In rsfMRI, however, there are no tasks that we can use to isolate psychologically functional coherent sub-networks. We might nonetheless hope that we could (somehow) use correlations or topological features to decompose the full RSN into sub-networks, and then determine the psychological processes for those sub-networks (see Andrews-Hanna *et al.* [2014]). Unfortunately, that strategy will not necessarily work due to the possibility of spurious or artifactual edges. These additional edges, we contend, are the product of computing one

⁴ Some recent RSN inference methods use associations more generally (rather than the more specific measure of correlations), particularly to capture non-linear relationships. For convenience, we refer here only to correlations, but everything that we say applies to associations in general.

⁵ Throughout this section, we talk only about distributions, but everything we discuss generalizes straightforwardly to probability densities.

correlation value for each pair of seed regions over a time period in which the underlying causal connections are likely shifting.

It is well-known that sampling a mixture distribution can result in spurious correlations (Hammell and O'Connell [1975]; Redner and Walker [1984]; Zheng and Fray [2004]). Consider the following non-neuroscientific example. Assume there is no correlation between height and beard growth in men; a short man is just as likely to sport a beard as a tall one. Under these conditions, adequate sampling from the population of men will likely yield a low height-beardedness correlation in the data. Women will similarly have a low height-beard correlation, even though they are, on average, both shorter than men and also less likely to have a beard. If we now take measurements of men and women (but do not measure sex), then we will have data from a mixture distribution. Moreover, we will likely find a correlation between height and beard growth, even though there is no correlation within either group: a short (tall) individual is more likely to be a woman (man) and so less (more) likely to have a beard. That is, knowledge of an individual's height will be informative about his or her beardedness, even though such knowledge is uninformative within each group. In this toy example, the correlation is spurious because (by assumption) these traits are uncorrelated in each population. The correlation is a sampling artifact that results from drawing individuals from a mixture of two heterogeneous populations.

This toy example involves a spurious correlation resulting from the mixture of populations in which the variables are uncorrelated. The opposite phenomenon can also occur: two variables can be independent in the mixture distribution, even though they are correlated within each subpopulation. More generally, even if the qualitative existence and direction of correlations are the same in every subpopulation, the mixture distribution can have quite different correlations, as seen in classic statistical puzzles such as Simpson's paradox and the Berkeley graduate admissions case of the 1970s (Bickel *et al.* [1975]). The standard methodological advice is to analyse only statistically or causally homogeneous subpopulations, but this advice is of limited use when our investigation is intended exactly to discover the causal relations that define the subpopulations (Cartwright [1979], [1989]).

We can be more precise about the potential impact of mixtures on functional connectivity networks, as some changes in correlations (and so changes in inferred edges) are more plausible than others (Ramsey *et al.* [2011]).⁶ In particular, if two variables are correlated in many or all of the component distributions, then typically only a few mixture parameter settings will yield independence (or zero correlation). In contrast, two variables that are independent in every component distribution will typically be correlated in the mixture distribution for a wide range of mixture parameters. Simply put, mixture distributions typically create correlations rather than eliminate them.⁷ The exception is when a variable has roughly the same probability distribution in each of the component distributions; in that case, mixing typically does not create new correlations. Moreover, mixture distributions can be quite stable when measured over long periods of time, and so the RSNs that we infer from rsfMRI data can also be highly stable. Thus, stability and consistency of inferred resting state networks across multiple subjects (Power *et al.* [2014]; Hurlburt *et al.* [2015]; Laumann *et al.* [2016]) is no counter-example to the mixture view (see also Objection (1) in the next section).

⁶ This paragraph can be made precise by reading 'implausible' as Lebesgue measure zero. We do not make stronger claims about probability or likelihood, as the standard Lebesgue measure might be inappropriate. We suspect that probability distributions over the (mixture and component) parameters will exhibit significant context-dependence.

⁷ At a high level, independence results only when the mixture parameters and mixture components are 'balanced', but there are many more ways to have an unbalanced equation than to have a balanced one.

Our mixture view leads to qualitative expectations about the RSNs for resting state data, even without knowledge of the particular mixture in a particular experimental participant or group of participants. First, we should not expect the whole-scan resting state network to be complete (that is, an edge between every pair of variables), since some brain areas presumably have similar activation probabilities across different cognitive processes, whether because the brain regions are engaged in functions that are not tied to particular cognitive processes, or because of shared functions across the participant's diverse cognitive activities. Data uniformity across time is even more likely to arise, given the enormous amount of smoothing and processing that is done to try to remove other types of artifacts (Laumann *et al.* [2016]). Second, since sampling a mixture tends to generate correlations rather than remove them, we should expect to find that inferred RSNs are 'supergraphs'—that is, have strictly more edges—of the superposition of the component graphs for the (relevant) underlying neural processes. That is, resting state analyses will tend to yield statistical structures with more edges than the union of those constituent networks. For example, sampling from a mixture of (data from) the two causal connectivity graphs in Figure 1(a) could yield data with a correlation structure represented by something like the functional connectivity network in Figure 1(b), though this is just one possibility of many, depending on the network parameters and the mixture parameter. Notice that every connection in one of the component graphs appears in the inferred graph (solid black for those in just the first graph, dotted black for only the second, and dashed black for both), and there are two spurious connections indicated with gray dashed lines (A — E and B — C) reflecting data correlations induced by the mixture. These kinds of changes can happen for resting state functional connectivity networks, except that the edges obviously do not come marked in this way. Instead, we learn Figure 1(b) or something similar from data, and then must (somehow) determine that it represents a mixture of the two causal connectivity networks in Figure 1(a).

[Insert Figure 1 about here]

According to the mixture view, any given RSN likely consists of both 'real' edges—correlations that would obtain in at least one unmixed component—and 'spurious' edges—correlations that obtain only in the mixed sample. However, we do not know which edges are real or spurious, so we cannot reliably identify sub-networks as targets for (reverse) psychological inference. For example, some studies associate the DMN with autobiographical memory (Spreng and Grady [2010]; Andrews-Hanna *et al.* [2014]), so one might hope that the DMN, or some portion of it, would support a network-based reverse inference (Poldrack [2006]; Klein [2012]) to autobiographical memory. But since aspects of the RSN may result from sampling a mixture distribution, critical connections might be merely artifactual; more generally, this large-scale RSN cannot be reliably decomposed into relevant sub-networks (Andrews-Hanna *et al.* [2014]). Similarly, even if a topological feature of an RSN correlates with a behavioural measure—for example, the centrality measure of nodes in the FPAN correlates with alerting attention measures (Markett *et al.* [2014])—we cannot reliably determine whether the edges supporting that correlation are merely spurious, and so cannot draw interesting psychological inferences. The possibility of mixtures thoroughly complicates attempts to draw psychological inferences from resting state studies.

4 Arguments For and Against the Mixture View

The previous section gave theoretical reasons to think that the mixture view is possible, and so resting state studies may not provide the right targets for psychological inferences such as ‘network X performs function Y ’. In this section, we provide empirical evidence that the mixture view is truly plausible, not merely possible. A full defense and confirmation of the mixture view would obviously require significant novel experiments and other empirical tests that we cannot conduct in the scope of this paper. Nonetheless, we think that it is important to establish that this is not a ‘merely theoretical’ worry for resting state research.

First, the mixture view predicts that the inferred RSN should dynamically shift over time if we look at only short windows of data. The analyses of Allen *et al.* ([2014]) provide direct support for this prediction, as they inferred functional connectivity networks for relatively short segments of resting state fMRI data from each of a large number of participants (see also Eavani *et al.* [2015]).⁸ They then used these inferred networks to compute the ‘network timecourse’ for each participant in their database, and showed that most participants had multiple networks over the course of data collection (typically changing every five or ten seconds during the five-minute scan). Moreover, their ‘changing networks’ model fit the data significantly better than a ‘static network’ model with a single unchanging RSN, even accounting for the additional degrees of freedom in the former model. Allen *et al.* ([2014]) thus provide direct evidence for one of the fundamental presuppositions of the mixture view—namely, rsfMRI time series data are generated by different networks over time. This study is part of a larger research trend focusing on dynamic shifts in resting state functional connectivity patterns (Hutchinson *et al.* [2013]).⁹ Relatedly, there is evidence that inferred RSNs can change depending on the instructions given to participants before the scan (Benjamin *et al.* [2010]; Van Dijk *et al.* [2010]), which is expected if the RSN is (partly) a mixture of networks for different cognitive processes, but not if the RSN is a single stable, constant background network.

A second major prediction of the mixture view is that RSNs will resemble the union of various task-based networks, plus additional inferred edges (though we cannot a priori predict which extra or spurious edges will result). Cole *et al.* ([2014]) directly compared a whole-scan RSN and various task-based networks, and generally found that the RSN was quite similar to the ‘multitask’ network inferred from multiple task-based (or not resting state) studies. That is, the RSN was similar to the superposition or union of networks produced by either seven or sixty-four different tasks (including working memory, motor, and decision making tasks). They additionally compared the RSN against the network for each specific task, and found that the majority of differences between the resting dataset and various task-based ones were edges present in the RSN but absent in the task-evoked network. In other words, they found that the whole-scan RSN is (almost) the union of task-evoked networks for multiple tasks, plus some additional edges. The mixture view predicts exactly this pattern of results since sampling from mixture distributions tends to generate correlations rather than eliminate them. Interestingly, Cole *et al.* do not endorse, or even seriously consider, the mixture view, but instead argue that we should interpret the resting state network as describing the ‘fundamental architecture’ of possible connections within the brain.

⁸ Specifically, Allen *et al.* estimated covariance matrices for moving windows of approximately forty-four seconds; clustered those matrices; and then used the cluster centroid covariance matrices to infer functional connectivity networks.

⁹ There is debate over whether the fluctuations observed in sliding window analyses reflect measurement noise or genuine dynamic shifts in functional connectivity patterns. However, recent studies support the idea that the resting state has dynamics that are not explained by noise (Hindriks *et al.* [2016]).

Lastly, the mixture view provides an explanation of why RSNs are often associated with many different cognitive functions, even after attempts at statistical decomposition. We previously noted that Andrews-Hanna *et al.* ([2014]) performed a meta-analysis of task-based studies of the DMN and found consistent differences in the cognitive functions attributed to three subsystems identified by Yeo and colleagues ([2011]): a ‘core’ network, a dorsal medial subsystem, and a medial temporal subsystem. For example, the dorsal medial subsystem was strongly associated with semantics and theory of mind, while the medial temporal subsystem was associated with episodic memory. These findings suggest that perhaps the DMN is not a single functional network, but a collection of somewhat related networks. Perhaps the dorsal medial system appears to have two quite different functions (semantics and mindreading) because different causal networks are being connected only through spurious edges. To the extent that distinct functional networks share some nodes and frequently occur closely in time, this situation is likely to arise. Of course, it is also possible that the DMN actually is a single network with multiple, context-sensitive cognitive functions. Or perhaps it performs some single function that is poorly captured by our current ‘cognitive ontology’ (Klein [2012]). The mixture view challenge is that we cannot distinguish between these options on statistical grounds alone, but rather require additional information or guidance. As this conclusion suggests, the mixture view does not imply that rsfMRI data are completely useless. Rather, it points towards the need to refine or improve our methodologies, or to bring additional knowledge to bear on the psychological inferences. In the next section, we consider just what advances might be required.

4.1 Objections to the mixture view

We conclude our defense of the plausibility of the mixture view by considering three objections, grounded in (1) consistency of resting state functional connectivity networks; (2) timing and magnitude of resting state activity; and (3) persistence of resting state activity in the absence of consciousness. We address these in turn, but emphasize that, even if more research is needed, the plausibility of spurious edges suffices to undercut our epistemic warrant that a particular RSN is a good candidate for psychological inference.

Objection (1): The first objection centers on the widespread reports that RSNs are highly consistent both between- and within-subjects (Damoiseaux *et al.* [2006]; Laumann *et al.* [2015]), sometimes to quite stunning degrees (Laumann *et al.* [2016]). In fact, Laumann *et al.* ([2016]) explicitly argue that RSNs are highly stable in a way that precludes the interpretation that ‘moment-to-moment changes in cognitive content’ (p. 1) are driving the signal. Moreover, roughly similar functional connectivity patterns emerge between scanning sessions for particular humans (Damoiseaux *et al.* [2006]; Laumann *et al.* [2015]); between humans (Damoiseaux *et al.* [2006]); and between individual non-human mammals such as mice (White *et al.* [2011]) or monkeys (Vincent *et al.* [2007]). These reports appear to be inconsistent with the mixture view, as the ratios of cognitive processes (that are supposedly being mixed) need not be consistent. For example, a participant might spend one resting scan thinking about an upcoming trip to Disney World, vividly imagining the food, riding the roller coasters, and so forth. On the next scanning session, the same participant might spend the entire time replaying songs in her head from a concert she attended. These very different cognitive sequences should produce very different mixtures of conscious cognitive processes, and so (on the mixture view) very different patterns of brain activation, seemingly contrary to the actual empirical findings.

However, this objection depends on the claim that there is little-to-no consistency in the psychological (and so brain) networks activated during the ‘free range’ cognition that occurs in a

rsfMRI experiment, and no evidence has been provided (beyond anecdotal introspection) in defense of this claim. Different people, and possibly other mammals, are plausibly engaged in a common set of background functions during the scanning session: the proprioceptive system is monitoring changes in limb position, the oculomotor system is controlling eye movements, the auditory system is responding to scanner noises, and so forth. And plausibly, a number of ‘typical’ cognitive processes, perhaps quite basic ones, occur during resting scans in a relatively predictable fashion or distribution. This possibility is even more likely for long scans, which are exactly the ones that show the greatest within- and between-individual consistency (Laumann *et al.* [2016]).

Moreover, the mixture view predicts that the mixtures are likely of related cognitive processes, as those will be temporally mixed more often. For example, the different cognitive functions associated with the default mode network (including social cognition, introspection, and autobiographical memory) seem to be connected in just this way (Andrews-Hanna *et al.* [2014]). More generally, free range cognition over the course of a long scan session can plausibly generate between- or within-subject consistency in functional connectivity patterns, as cognitive processes plausibly occur, co-occur, and follow one another, in a regular fashion. Moreover, recent studies suggest that participants exhibit similar cognitive patterns whether in the scanner or in the world (Hurlburt *et al.* [2015]), and so we should expect there to be significant within-subject RSN consistency that is grounded in these cognitive patterns.

Objection (2): The second objection arises from reports that the timing and magnitude of resting state BOLD fluctuations differ from those of task-evoked BOLD fluctuations (Raichle [2009]; Snyder and Raichle [2012]), which suggests that the former are not just mixtures of the latter. For example, Snyder and Raichle ([2012]) argue that resting state BOLD fluctuations are both slower and larger in magnitude than task-evoked BOLD changes, and so ‘unconstrained cognition alone does not account for the greatest part of intrinsic activity’ (p. 903). Their inference is too quick, however, as the mixture view can easily account for these qualitative differences. In terms of timing, if functional connectivity patterns arise from mixtures of cognitive processes operating at different time scales, then the overall mixture will involve correlations over longer time scales than the component processes that neuroscientists typically measure using cognitive tasks. And in terms of magnitude, the resting state data reveal conflicting results (for example, Snyder and Raichle [2012] vs. Damoiseaux *et al.* [2006]), and it is not known whether this difference is due to features of the experimental protocol or differences in the mixture. In fact, Section 5 will suggest a way to use this disagreement to possibly test some implications of the mixture view.

Objection (3): Perhaps the most serious objection to the mixture view comes from the finding that resting state activity persists in heavily sedated monkeys (Vincent *et al.* [2007]) and humans (Greicius *et al.* [2008]; Hutchinson *et al.* [2013]). This result appears inconsistent with a mixture composed of conscious cognitive processes. Since functional connectivity patterns measured during rest in alert participants seem to persist even in unconscious individuals, rsfMRI is plausibly thought to measure intrinsic activity.

There are two points to make here. First, as we noted earlier, the mixture view is not committed to the implausible claim that all RSNs arising solely from mixtures of networks for conscious cognition. The brain undoubtedly exhibits some degree of baseline metabolic activity, and essentially all cognitive theories leave room for ongoing background activity. For example, predictive coding theories (Hohwy [2014]) hold that the brain is constantly generating predictive models of the environment, and testing these models against incoming information, whether from

external sensory input or other sources. The mixture view is entirely consistent with RSNs being a mix of baseline metabolic processes, background cognitive processes, and free-ranging conscious cognition. That is, we view the stability of certain resting state networks in sedated individuals (Rosazza and Minati [2011]) as providing insight into one of these process-types, rather than implying that the mixture view is false.

Second, the problem of mixtures is a sampling problem that could also arise for purely unconscious cognitive or metabolic processes. Even if resting state protocols measure intrinsic neural activity, this activity could result from many distinct processes that are mixed together in the sampling process. Insofar as there are dynamic changes in the causal networks engaged during the scan, mixtures become a real possibility (whether or not these changes result from occurrent cognitive processes or merely changing baseline regimes). Some proponents of resting state fMRI argue that resting state analyses measure intrinsic activity—activity that would persist in the absence of conscious cognition—that reflects preparatory or anticipatory activity in known functional systems, such as motor or visual areas (Damoiseaux *et al.* [2006]). But inferences to these psychological functional subsystems similarly require a solution to the possibility of mixtures.

5 Finding Value for the Resting State

The mixture view challenges whether RSNs as such are appropriate targets of psychological inferences from rsfMRI, and so raises questions about the value for cognitive neuroscience of this (relatively new) methodology, particularly given advances in task-based methods, such as time-varying tasks and dynamic analyses. In this section, we aim to provide positive answers to those questions: we suggest that rsfMRI studies nonetheless have value in their possibility as a ‘discovery science’ that discovers new correlational brain networks (Biswal *et al.* [2010]), whether because participants do things during the scanning session that neuroscientists have not designed tasks to explicitly target (Mason *et al.* [2007]; Christoff *et al.* [2016]), or because there is background activity in previously uncharacterized functional brain networks during rest (Greicius *et al.* [2003]). Unfortunately, the very features of resting state studies that provide this benefit—participants engage in uncontrolled cognition over long periods of time—are the ones that lead to the problem of mixtures, so we need to advance resting state methodologies to avoid without falling victim to the problem of mixtures.

5.1 New tools and techniques

The most obvious potential methodological change would be to shift from reverse inference (taking an RSN inferred from a whole scan as given and searching for its function) to forward inference (attaching occurrent functions to RSNs). If we determine the cognitive process(es) occurring at relevant points in time, then we can use that information to attribute cognitive function to the active connectivity network at that moment. The mixture view holds that at least some RSNs result from occurrent cognition; they are the networks responsible for remembering the grocery list, planning one’s route from the lab to the store, deciding whether to stop on the way home, and so on. Resting state experiments do not normally collect reports about participants’ introspective mental activities, which is precisely why reverse inference (or inferences from learning paradigms) is normally required. In principle, information about participants’ internal cognition could be used for forward inference as in task-based studies, albeit without control tasks and granting that conscious, introspectively accessible cognition is

informative about only some of the brain's activity. If an individual is visualizing at particular times, for example, then we can (try to) learn the connectivity network for those times.

There are long-standing worries about the reliability of introspection and self-reporting (Jones and Harris [1967]; Nisbett and Ross [1991]; Schwitzgebel [2008]; Engelbert and Carruthers [2010]). If participants are instead asked to report their conscious cognition in real-time, then the study will involve a constant task (namely, to remember one's conscious states) and so will no longer capture unconstrained cognition. Even if participants give reports only afterwards, knowledge that they will be asked to do so can plausibly induce constant memory and metacognition tasks throughout the experiment. Thus, it seems that the request for (retrospective) self-reports must come as a surprise at the end of the study, which raises worries about whether participants will report whatever information was ad hoc encoded, perhaps with significant error, in memory. Instead, we can consider more sophisticated methods to learn about cognitive activities (see also Section 5.2 below). Introspective sampling techniques like Descriptive Experience Sampling (DES) ask participants to report their introspective experiences when prompted by a random beeper, and so can provide a principled basis for dividing the rsfMRI time series, thereby potentially revealing interesting correspondences between patterns of thought and resting state signals (Hurlburt *et al.* [2015]). Some neuroscientists have recently adopted this methodology; for instance, Van Calster *et al.* ([2016]) use various experience sampling techniques to link occurrent episodes of top-down and bottom-up attention with activity in the dorsal and ventral attention networks. Alternately, one could use experimental instructions that are more open-ended than a task-based study without being quite as open-ended as in traditional rsfMRI studies, such as 'think about your trip to the lab' or 'every 30 seconds, imagine a different object'.

A different strategy is to employ statistical methods that can (defeasibly) identify candidate networks. The results of Allen *et al.* ([2014]) and Andrews-Hanna *et al.* ([2014]) demonstrate that sophisticated analyses or time-windowing can reveal better targets for psychological inference. Of course, as we argued earlier, simple decomposition strategies are not sufficient since temporal mixtures can produce spurious edges and the component networks are not necessarily known in advance. But we can potentially use more sophisticated methods, coupled with network discovery algorithms, to simultaneously infer both the networks and their mixing parameters. For example, standard techniques such as independent components analysis (ICA) can extract sources or signals that are mixed in a data stream, as the multiple networks are mixed. However, successful use of ICA requires that the mixture be relatively stable, and it is unclear whether mindwandering will lead to a stable medium-term mixture. Alternately, there are statistical methods that can segment a time series based on (likely) changes in the generative distribution, such as changepoint detection algorithms (Desobry *et al.* [2005]; Adams and MacKay [2007]) or other methods (Gregory *et al.* [1996]; Scargle [1998]). Those methods require strong assumptions, however, so may not be suitable for use with rsfMRI data. The reliability of these advanced statistical methods depends on open mathematical and data questions that have not yet been answered. And even if they do suffice to address the mixture problem, those methods will not solve other analysis challenges for fMRI data.¹⁰

Finally, we can explore technological approaches to the mixture challenge. Many resting state experimental protocols do not collect sufficiently fine-grained data to adequately separate

¹⁰ For example, the challenge of inferring networks from correlations given complications such as time-averaging of neural activity in the BOLD signal, undersampling of fMRI measurements (Seth *et al.* [2013]), and many other issues.

the mixture components or changing mixture parameters.¹¹ Most researchers have used longer repetition times (TRs) in their resting state studies (though there has been a recent shift towards shorter 2000 – 800 ms TRs), as these scanner settings yield improved signal-to-noise ratios without impairing inference about the slower changes that have been the principal focus of resting state research. The mixture view argues, however, that significant components of the ‘resting’ state time series are due to cognitive functional networks that presumably change more rapidly. Thus, we should arguably use faster scan times so we have the temporal resolution to adequately separate the mixture components. Using shorter TRs may help with these challenges, and such studies are part of ongoing research on dynamics in the resting state.

The main overall message is that the mixture challenge is a serious one, but matters are far from hopeless. There are methodological, statistical, and technological routes that could reduce the plausibility of spurious connections in RSNs, and thereby increase the chances of successful psychological inference from resting state fMRI data. Key issues for further theorizing and experimentation are the degree to which RSNs correlate with occurrent cognition (Van Calster *et al.* [2016]) and the extent to which RSNs change dynamically during the scanning session (Hutchinson *et al.* [2013]).

5.2 Exploration via the resting state

Given the inferential challenges that we have identified, one might suggest that we simply stop attaching cognitive functions to RSNs, perhaps reserving rsfMRI data for diagnostic or prognostic uses. We contend that this would be a mistake, as rsfMRI studies do have significant potential for cognitive brain mapping. However, making good on this potential (if the mixture view is on the right track) will require reassessing the role of rsfMRI studies in brain mapping. The core difference between task-based and resting state studies lies not in targeting different kinds of networks (intrinsic versus task-evoked networks), but rather in different degrees of experimental control over brain and cognitive activity. Standard task-based studies involve significant control, as the participants’ cognition is presumably driven principally by experimental demands. In contrast, resting state studies are relatively uncontrolled, often quite explicitly so. Precision about the exact differences in control would require a full philosophical analysis of ‘control’; as we noted earlier, we do not provide such an account here. However, even our earlier brief comments are sufficient to say more about the potential value of resting state studies.

A common view is that greater experimental control is always better, but this claim is not correct in general; rather, it depends on what one is trying to learn. Significant control (including randomization) can be the best way to discover whether some particular factor *A* causes some particular *B*, as control can enable one to exclude possible confounding factors. But if one is instead trying to characterize or understand the typical behaviour of a system in its natural, unmanipulated environment, then significant experimental control can actually impair learning, precisely because that control can change the system from its usual states.

Scientists often conduct exploratory experiments in order to map how a system operates in a ‘natural’ setting, which may be different than the operation in any particular, controlled experimental context (Franklin [2005]). The mixture view implies that rsfMRI studies could similarly reveal ‘natural’ brain behaviour.¹² Task-based studies are, by design, artificial in certain respects, as they push the individual to engage in certain cognitive activities, typically to the

¹¹ Thanks to David Plaut for emphasizing this worry.

¹² One might wonder whether laying in a scanner is a particularly ‘natural’ environment. At the very least, though, resting states involve more natural cognition than task-based studies.

exclusion of others. Potentially, there are significant brain networks that have not been observed, simply because those networks underlie some task that has not been isolated in any particular experiment. Precisely because resting state studies are relatively uncontrolled, they hold forth the promise of revealing previously unknown or understudied brain networks. People's free-ranging cognition plausibly traverses a wider space than experimenters have previously thought to isolate. By studying more naturalistic cognition, we potentially will find new networks to be studied and understood. Of course, we emphasize that significant methodological advances, including those discussed in Section 5.1, are required before we can fully realize these possibilities.

Moreover, if we understand resting state versus task-based fMRI studies as simply endpoints in a continuum of experimental control, then many other possibilities arise. This perspective shift suggests a number of intermediate experimental designs in which the experimenter only partially controls the participants' cognitive activities. For example, one could conduct an fMRI study that asks people to 'think about the last three meals that you ate'. This instruction is not as completely open-ended as in a traditional resting state study, but also does not aim for the very tight control of traditional task-based studies. More generally, resting state studies can be incredibly useful and powerful, but they must be interpreted with appropriate care. On the mixture view, these studies do not necessarily reveal some intrinsic, task-free, omnipresent network. They can, however, potentially reveal new brain networks and help us better understand naturalistic brain activity, including the ways in which different networks co-occur and interact when external conditions do not dictate a particular goal, task, or cognition. That is, resting state fMRI studies can be (we suggest) a full-blooded realization of the possibilities of exploratory science. In order to do so, however, we must ensure that we account for mixtures by adopting the methodological and analytic techniques outlined earlier.

6 Conclusion

Recent work in the philosophy of experiment suggests that exploratory experiments can reveal interesting patterns about a system in the absence of substantial experimental control (Steinle [1997]; Franklin [2005]). We propose there is sometimes a tradeoff between the degree of control one exerts over a target system and the possibility of discovering novel patterns. For a science such as human brain mapping, in which our knowledge of what structures are cognitively interesting and what functions they perform is somewhat limited, it may be particularly useful to explore both ends of this experimental spectrum. Resting state fMRI fits this view of exploratory experimentation: it putatively identifies functional brain networks (via functional connectivity patterns) in a more naturalistic setting. As we have argued in this paper, though, this very freedom creates a novel inferential challenge for any attempt to connect the observed functional connectivity networks back to psychological function(s).

Thus, rsfMRI reflects a shift away from the tightly controlled paradigms that characterized earlier cognitive neuroimaging experiments (Biswal *et al.* [2010]). This shift has many potential benefits, such as generating the prospect of discovering new functional networks, but also raises new problems, such as how to attach functions to the networks so discovered. Resting state fMRI is typically interpreted as revealing features of persistent, underlying functional networks (and not 'mere' correlation structure) that can be interpreted in a number of different ways. A common theme is that resting state analyses can be used to identify large-scale, cognitive and metabolic, functional networks in advance of knowing what exactly those networks are doing. There is much more agreement on this point than on what networks there are, or what precise functions these networks may be performing.

Our account challenges this widespread way of interpreting resting state data. The mixture view proposes that the appearance of a single or few large-scale networks in resting state experiments may involve various sampling artifacts rather than genuine features of the brain's functional organization. According to this view, the so-called 'resting' brain is a mixture of different kinds of processes—background metabolic, background cognitive, and free-ranging conscious cognition—that occur at different times and timescales. Mixtures of these diverse processes yield functional connectivity networks in which there is no way to know which inferred edges correspond to connections in cognitive or metabolic networks, and which to spurious connections. We agree that resting state analyses hold the promise of identifying novel functional networks, since free-ranging cognition can traverse a broader range of brain processes than researchers typically target in task-based studies. To find these networks, though, we must account for the possibility of mixtures, thereby improving the ability of resting state analyses to disentangle these diverse functional signals. With these methodological and analytic adjustments, we can potentially make good on the promise of identifying candidate cognitive networks in an exploratory fashion.

Acknowledgements

Thanks to two anonymous referees for their valuable comments and feedback. Thanks also to Mazviita Chirimuuta, Clark Glymour, Steve Hanson, Edouard Machery, Steve Petersen, Carl Craver, David Plaut, and Russ Poldrack for invaluable feedback on earlier versions of this paper.

Joseph McCaffrey
Department of Philosophy and Philosophy-Neuroscience-Psychology Program
Washington University in St. Louis
St. Louis, MO
jbmccaffrey@wustl.edu

David Danks
Departments of Philosophy and Psychology
Carnegie Mellon University
Pittsburgh, USA
ddanks@cmu.edu

References

- Anderson, M. L. [2010]: 'Neural Reuse: A Fundamental Organizational Principle of the Brain', *Behavioral and Brain Sciences*, **33**, pp. 245-66.
- Adams, R. P. and MacKay, D. J. C. [2007]: 'Bayesian Online Change-point Detection', Technical report, University of Cambridge, Cambridge, UK. arXiv:0710.3742v1.
- Allen, E. A., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T. and Calhoun, V. D. [2014]: 'Tracking Whole-brain Connectivity Dynamics in the Resting State', *Cerebral Cortex*, **24**, pp. 663-76.
- Andrews-Hanna, J. R., Smallwood, J. and Spreng, R. N. [2014]: 'The Default Network and Self-generated Thought: Component Processes, Dynamic Control, and Clinical Relevance', *Annals of the New York Academy of Sciences*, **1316**, pp. 29-52.

- Basso, K., Margolin, A. A., Stolovitzky, G., Klein, U., Dalla-Favera, R. and Califano, A. [2005]: 'Reverse Engineering of Regulatory Networks in Human B Cells', *Nature Genetics*, **37**, pp. 382-90.
- Benjamin, C., Lieberman, D. A., Chang, M., Ofen, N., Whitfield-Gabrieli, S., Gabrieli, J. D. and Gaab, N. [2010]: 'The Influence of Rest Period Instructions on the Default Mode Network', *Frontiers in Human Neuroscience*, **4**, p. 218.
- Bickel, P. J., Hammel, E. A. and O'Connell, J. W. [1975]: 'Sex Bias in Graduate Admissions: Data from Berkeley', *Science*, **187**, pp. 398-404.
- Binder, J. R., Frost, J. A., Hammeke, T. A., Bellgowan, P. S., Rao, S. M. and Cox, R. W. [1999]: 'Conceptual Processing during the Conscious Resting State: A Functional MRI Study', *Journal of Cognitive Neuroscience*, **11**, pp. 80-95.
- Biswal, B. B., Yetkin, F. Z., Haughton, V. M. and Hyde, J. S. [1995]: 'Functional Connectivity in the Motor Cortex of Resting Human Brain using Echo-planar MRI', *Magnetic Resonance Imaging*, **34**, pp. 537-41.
- Biswal, B. B., Mennes, M., Zuo, X. N., Gohel, S. Kelly, C., Smith, S. M., Beckmann, C. F., Adelstein, J. S., Buckner, R. L., Colcombe, S., Dogonowski, A. M., Ernst, M., Fair, D., Hampson, M., Hoptman, M. J., Hyde, J. S., Kiviniemi, V. J., Kotter, R., Li, S. J., Lin, C. P., Lowe, M. J., Mackay, C., Madden, D. J., Madsen, K. H., Margulies, D. S., Mayberg, H. S., McMahon, K., Monk, C. S., Mostofsky, S. H., Nagel, B. J., Pekar, J. J., Peltier, S. J., Petersen, S. E., Riedl, V., Rombouts, S. A., Rympa, B. Schlagger, B. L., Schmidt, S., Seidler, R. D., Siegle, G. J., Sorg, C., Teng, G. J., Veijola, J., Villringer, A., Walter, M., Wang, L., Weng, X. C., Whitfield-Gabrieli, S., Williamson, P., Windischberger, C., Zang, Y. F., Zhang, H. Y., Castellanos, F. X. and Milham, M. P. [2010]: 'Toward Discovery Science of Human Brain Function', *Proceedings of the National Academy of Sciences*, **107**, pp. 4734-9.
- Bray, S., Arnold, A. E., Levy, R. M. and Iaria, G. [2015] 'Spatial and Temporal Functional Connectivity Changes between Resting and Attentive States', *Human Brain Mapping*, **36**, pp. 549-65.
- Buckner, R. L., Andrews-Hanna, J. R. and Schacter, D. L. [2008]: 'The Brain's Default Network: Anatomy, Function, and Relevance to Disease', *Annals of the New York Academy of Sciences*, **1124**, pp. 1-38.
- Bullmore, E. and Sporns, O. [2009]: 'Complex Brain Networks: Graph Theoretical Analysis of Structural and Functional Systems', *Nature Reviews Neuroscience*, **10**, pp. 186-98.
- Cartwright, N. [1979]: 'Causal Laws and Effective Strategies', *Noûs*, **13**, pp. 419-37.
- Cartwright, N. [1989]: *Nature's Capacities and Their Measurement*. Oxford: Oxford University Press.
- Christoff, K., Irving, Z. C., Fox, K. C., Spreng, R. N. and Andrews-Hanna, J. R. [2016]: 'Mind-wandering as Spontaneous Thought: A Dynamic Framework', *Nature Reviews Neuroscience*, **17**, pp. 718-31.
- Cohen, L., Dehaene, S., Naccache, L., Lehéricy, S., Dehaene-Lambertz, G., Hénaff, M. A., and Michel, F. [2000]: 'The Visual Word Form Area', *Brain*, **123**, pp. 291-307.
- Cole, M. W., Bassett, D. S., Power, J. D., Braver, T. S. and Petersen, S. E. [2014]: 'Intrinsic and Task-evoked Network Architectures of the Human Brain', *Neuron*, **83**, pp. 238-51.
- Coltheart, M. [2006]: 'Perhaps Functional Neuroimaging Has Not Told Us Anything about the Mind (So Far)', *Cortex*, **42**, pp. 422-7.

- Damoiseaux, J. S., Rombouts, S. A. R. B., Barkhof, F., Scheltens, P., Stam, C. J., Smith, S. M. and Beckmann, C. F. [2006]: ‘Consistent Resting-state Networks across Healthy Subjects’, *Proceedings of the National Academy of Sciences*, **103**, pp. 13848-53.
- Del Pinal, G. and Nathan, M. J. [2017]: ‘Two Kinds of Reverse Inference in Cognitive Neuroscience’, in J. Leefman and E. Hildt (eds), *The Human Sciences after the Decade of the Brain*, Elsevier, pp. 121-139.
- Desobry, F., Davy, M. and Doncarli, C. [2005]: ‘An Online Kernel Change Detection Algorithm’, *IEEE Transactions on Signal Processing*, **8**, pp. 2961–74.
- Eavani, H., Satterthwaite, T. D., Filipovych, R., Gur, R. E., Gur, R. C. and Davatzikos, C. [2015]: ‘Identifying Sparse Connectivity Patterns in the Brain Using Resting-state fMRI’, *Neuroimage*, **105**, pp. 286-99.
- Engelbert, M. and Carruthers, P. [2010]: ‘Introspection’, *Wiley Interdisciplinary Reviews: Cognitive Science*, **1**, pp. 245-53.
- Fox, M. D. and Raichle, M. E. [2007]: ‘Spontaneous Fluctuations in Brain Activity Observed with Functional Magnetic Resonance Imaging’, *Nature Reviews Neuroscience*, **8**, pp. 700-11.
- Fox, M. D., Snyder, A.Z., Vincent, J. L., Corbetta, M., Van Essen, D. C. and Raichle, M. E. [2005]: ‘The Human Brain is Intrinsically Organized into Dynamic, Anticorrelated Functional Networks’, *Proceedings of the National Academy of Sciences*, **102**, pp. 9673–8.
- Franklin, L. R. [2005]: ‘Exploratory Experiments’, *Philosophy of Science*, **72**, pp. 888-99.
- Friston, K. J., Price, C. J., Fletcher, P., Moore, C., Frackowiak, R. S. J., and Dolan, R. J. [1996]: ‘The Trouble with Cognitive Subtraction’, *Neuroimage*, **4**, pp. 97-104.
- Friston, K. J. [2011]: ‘Functional and Effective Connectivity: A Review’, *Brain Connectivity*, **1**, pp. 13-36.
- Glover, G. H. and Lee, A. T. [1995]: ‘Motion Artifacts in fMRI: Comparison of 2DFT with PR and Spiral Scan Methods’, *Magnetic Resonance in Medicine*, **33**, pp. 624-35.
- Glymour, C. and Hanson, C. [2015]: ‘Reverse Inference in Neuropsychology’, *The British Journal for the Philosophy of Science*, **67**, pp. 1139-53.
- Gregory, A. W., Nason, J. M. and Watt, D. G. [1996]: ‘Testing for Structural Breaks in Cointegrated Relationships’, *Journal of Econometrics*, **71**, pp. 321-41.
- Greicius, M. D., Kiviniemi, V., Tervonen, O., Vainionpää, V., Alahuhta, S., Reiss, A. L. and Menon, V. [2008]: ‘Persistent Default-mode Network Connectivity during Light Sedation’, *Human Brain Mapping*, **29**, pp. 839-47.
- Greicius, M. D., Krasnow, B., Reiss, A. L. and Menon, V. [2003]: Functional Connectivity in the Resting Brain: A Network Analysis of the Default Mode Hypothesis’, *Proceedings of the National Academy of Sciences*, **100**, pp. 253-8.
- Gusnard, D. A., Akbudak, E., Shulman, G. L. and Raichle, M. E. [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**, pp. 4259–64.
- Gusnard, D. A. and Raichle, M. E. [2001]: ‘Searching for a Baseline: Functional Imaging and the Resting Human Brain’, *Nature Reviews Neuroscience*, **2**, pp. 685–94.
- He, B. J., Snyder, A. Z., Zempel, J. M., Smyth, M. D. and Raichle, M. E. [2008]: ‘Electrophysiological Correlates of the Brain's Intrinsic Large-scale Functional Architecture’, *Proceedings of the National Academy of Sciences*, **105**, pp. 16039-44.
- Henson, R. [2006]: ‘Forward Inference using Functional Neuroimaging: Dissociations versus Associations’, *Trends in Cognitive Sciences*, **10**, pp. 64-9.

- Hindriks, R., Adhikari, M. H., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N. K. and Deco, G. [2016]: ‘Can Sliding-window Correlations Reveal Dynamic Functional Connectivity in Resting-state fMRI?’, *Neuroimage*, **127**, pp. 242-56.
- Hohwy, J. [2013]: *The Predictive Mind*. Oxford: Oxford University Press.
- Hurlburt, R. T., Alderson-Day, B., Fernyhough, C. and Kühn, S. [2015]: ‘What Goes on in the Resting-state? A Qualitative Glimpse into Resting-state Experience in the Scanner’, *Frontiers in Psychology*, **6**, p. 1535.
- Johnston, J. M., Vaishnavi, S. N., Smyth, M. D., Zhang, D., He, B. J., Zempel, J. M., Shimony, J. S., Snyder, A. Z. and Raichle, M. E. [2008]: ‘Loss of Resting Interhemispheric Functional Connectivity after Complete Section of the Corpus Callosum’, *The Journal of Neuroscience*, **28**, pp. 6453-8.
- Jones, E. E. and Harris, V. A. [1967]: ‘The Attribution of Attitudes’, *Journal of Experimental Social Psychology*, **3**, pp. 1–24.
- Kanwisher, N., McDermott, J. and Chun, M. M. [1997]: ‘The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception’, *The Journal of Neuroscience*, **17**, pp. 4302-11.
- Klein, C. [2010]: ‘Philosophical Issues in Neuroimaging’, *Philosophy Compass*, **5**, pp. 186-98.
- Klein, C. [2012]: ‘Cognitive Ontology and Region-versus Network-oriented Analyses’, *Philosophy of Science*, **79**, pp. 952-60.
- Klein, C. [2014]: ‘The brain at Rest: What It Is Doing and Why that Matters’, *Philosophy of Science*, **81**, pp. 974-85.
- Laumann, T. O., Gordon, E. M., Adeyemo, B., Snyder, A. Z., Joo, S. J., Chen, M. Y., Gilmore, A. W., McDermott, K. B., Nelson, S. M., Dosenbach, N. U. F., Schlaggar, B. L., Mumford, J. A., Poldrack, R. A. and Petersen, S. E. [2015]: ‘Functional System and Areal Organization of a Highly Sampled Individual Human Brain’, *Neuron*, **87**, pp. 657-70.
- Laumann, T. O., Snyder, A. Z., Mitra, A., Gordon, E. M., Gratton, C., Adeyemo, B., Gilmore, A. W., Nelson, S. M., Berg, J. J., Greene, D. J., McCarthy, J. E., Tagliazucchi, E., Laufs, H., Schlaggar, B. L., Dosenbach, N. U. F. and Petersen, S. E. [2016]: ‘On the Stability of BOLD fMRI Correlations’, *Cerebral Cortex*, pp. 1-14.
- Lowe, M. J., Mock, B. J. and Sorenson, J. A. [1998]: ‘Functional Connectivity in Single and Multislice Echoplanar Imaging Using Resting-state Fluctuations’, *Neuroimage*, **7**, pp. 119-32.
- Machery, E. [2012]: ‘Dissociations in Neuropsychology and Cognitive Neuroscience’, *Philosophy of Science*, **79**, pp. 490-518.
- Machery, E. [2013]: ‘In Defense of Reverse Inference’, *The British Journal for the Philosophy of Science*, **65**, pp. 251-67.
- Markett, S., Reuter, M., Montag, C., Voigt, G., Lachmann, B., Rudolf, S., Elger, C. E. and Weber, B. [2014]: ‘Assessing the Function of the Fronto-parietal Attention Network: Insights from Resting-state fMRI and the Attentional Network Test’, *Human Brain Mapping*, **35**, pp. 1700-9.
- Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T. and Macrae, C. N. [2007]: ‘Wandering Minds: The Default Network and Stimulus-independent Thought’, *Science*, **315**, pp. 393-5.
- McCaffrey, J. B. [2015]: ‘The Brain’s Heterogeneous Functional Landscape’, *Philosophy of Science*, **82**, pp. 1010-22.
- Morcom, A. M. and Fletcher, P. C. [2007]: ‘Does the Brain Have a Baseline? Why We Should Be Resisting a Rest’, *Neuroimage*, **37**, pp. 1073–82.

- Nathan, M. J. and Del Pinal, G. [2016]: ‘Mapping the Mind: Bridge Laws and the Psycho-neural Interface’, *Synthese*, **193**, pp. 637-57.
- Nisbett, R. E. and Ross, L. [1991]: *The Person and the Situation: Perspectives of Social Psychology*. New York: McGraw-Hill Publishing.
- O’Keefe, J. and Dostrovsky, J. [1971]: ‘The Hippocampus as a Spatial Map. Preliminary Evidence from Unit Activity in the Freely Moving Rat’, *Brain Research*, **34**, pp. 171-5.
- Poldrack, R. A. [2006]: ‘Can Cognitive Processes Be Inferred from Neuroimaging Data?’, *Trends in Cognitive Sciences*, **10**, pp. 59-63.
- Posner, M. I., Petersen, S., Fox, P. and Raichle, M. [1988]: ‘Localization of Cognitive Operations in the Human Brain’, *Science*, **240**, pp. 1627-31.
- Power, J. D., Cohen, A. L., Nelson, S. M., Wig, G. S., Barnes, K. A., Church, J. A., Vogel, A. C., Laumann, T. O., Miezin, F. M., Schlaggar, B. L. and Petersen, S. E. [2011]: ‘Functional Network Organization of the Human Brain’, *Neuron*, **72**, pp. 665-78.
- Power, J. D., Schlaggar, B. L., Lessov-Schlaggar, C. N. and Petersen, S. E. [2013]: ‘Evidence for Hubs in Human Functional Brain Networks’, *Neuron*, **79**, pp. 798-813.
- Power, J. D., Schlaggar, B. L. and Petersen, S. E. [2014]: ‘Studying Brain Organization via Spontaneous fMRI Signal’, *Neuron*, **84**, pp. 681-96.
- Purdon, P. L. and Weisskoff, R. M. [1998]: ‘Effect of Temporal Autocorrelation Due to Physiological Noise and Stimulus Paradigm on Voxel-level False-positive Rates in fMRI’, *Human Brain Mapping*, **6**, pp. 239-49.
- Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A. and Shulman, G. L. [2001]: ‘A Default Mode of Brain Function’, *Proceedings of the National Academy of Sciences*, **98**, pp. 676–82.
- Raichle, M. E. [2009]: ‘A Paradigm Shift in Functional Brain Imaging’, *The Journal of Neuroscience*, **29**, pp. 12729-34.
- Ramsey, J. D., Spirtes, P. and Glymour, C. [2011]: ‘On Meta-analyses of Imaging Data and the Mixture of Records’, *NeuroImage*, **57**, pp. 323-30.
- Redner, R. A. and Walker, H. F. [1984]: ‘Mixture Densities, Maximum Likelihood, and the EM algorithm’, *SIAM Review*, **26**, pp. 195-239.
- Rosazza, C. and Minati, L. [2011]: ‘Resting-state Brain Networks: Literature Review and Clinical Applications’, *Neurological Sciences*, **32**, pp. 773-85.
- Roskies, A. L. [2007]: ‘Are Neuroimages like Photographs of the Brain?’, *Philosophy of Science*, **74**, pp. 860-72.
- Roskies, A. L. [2009]: ‘Brain-Mind and Structure-Function Relationships: A Methodological Response to Coltheart’, *Philosophy of Science*, **76**, pp. 927-939.
- Roskies, A. L. [2010]: ‘Saving Subtraction: A Reply to Van Orden and Paap’, *The British Journal for the Philosophy of Science*, **61**, pp. 635-65.
- Saxe, R. and Kanwisher, N. [2003]: ‘People Thinking about Thinking People: The Role of the Temporo-parietal Junction in “Theory of Mind”’, *Neuroimage*, **19**, pp. 1835-42.
- Scargle, J. D. [1998]: ‘Studies in Astronomical Time Series Analysis. V. Bayesian Blocks, a New Method to Analyze Structure in Photon Counting Data’, *The Astrophysical Journal*, **504**, pp. 405-18.
- Schwitzgebel, E. [2008]: ‘The Unreliability of Naive Introspection’, *Philosophical Review*, **117**, pp. 245-73.

- Seth, A. K., Chorley, P. and Barnett, L. C. [2013]: 'Granger Causality Analysis of fMRI BOLD Signals is Invariant to Hemodynamic Convolution but not Downsampling', *Neuroimage*, **65**, pp. 540-55.
- Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R. and Beckmann, C. F. [2009]: 'Correspondence of the Brain's Functional Architecture During Activation and Rest', *Proceedings of the National Academy of Sciences*, **106**, pp. 13040-5.
- Snyder, A. Z. and Raichle, M. E. [2012]: 'A Brief History of the Resting State: The Washington University Perspective', *Neuroimage*, **62**, pp. 902-10.
- Sprengh, R. N. and Grady, C. L. [2010]: 'Patterns of Brain Activity Supporting Autobiographical Memory, Propection, and Theory of Mind, and Their Relationship to the Default Mode Network', *Journal of Cognitive Neuroscience*, **22**, pp. 1112-23.
- Steinle, F. [1997]: 'Entering New Fields: Exploratory Uses of Experimentation', *Philosophy of Science*, **64**, pp. S65-S74.
- Uddin, L. Q. [2015]: 'Salience Processing and Insular Cortical Function and Dysfunction', *Nature Reviews Neuroscience*, **16**, pp. 55-61.
- Vahdat, S., Darainy, M., Milner, T. E. and Ostry, D. J. [2011]: 'Functionally Specific Changes in Resting-state Sensorimotor Networks After Motor Learning', *The Journal of Neuroscience*, **31**, pp. 16907-15.
- Van Den Heuvel, M. P. and Pol, H. E. H. [2010]: 'Exploring the Brain Network: A Review on Resting-state fMRI Functional Connectivity', *European Neuropsychopharmacology*, **20**, pp. 519-534.
- Van Dijk, K. R., Hedden, T., Venkataraman, A., Evans, K. C., Lazar, S. W. and Buckner, R. L. [2010]: 'Intrinsic Functional Connectivity as a Tool for Human Connectomics: Theory, Properties, and Optimization', *Journal of Neurophysiology*, **103**, pp. 297-321.
- Van Calster, L., D'Argembeau, A., Salmon, E., Peters, F. and Majerus, S. [2016]: 'Fluctuations of Attentional Networks and Default Mode Network During the Resting State Reflect Variations in Cognitive States: Evidence from a Novel Resting-state Experience Sampling Method', *Journal of Cognitive Neuroscience*, **29**, pp. 95-113.
- Van Orden, G. C. and Paap, K. R. [1997]: 'Functional Neuroimages Fail to Discover Pieces of Mind in the Parts of the Brain', *Philosophy of Science*, **64**, S85-S94.
- Vincent, J. L., Patel, G. H., Fox, M. D., Snyder, A. Z., Baker, J. T., Van Essen, D. C., Zempel, J. M., Snyder, L. H., Corbetta, M. and Raichle, M. E. [2007]: 'Intrinsic Functional Architecture in the Anaesthetized Monkey Brain', *Nature*, **447**, pp. 83-6.
- White, B. R., Bauer, A. Q., Snyder, A. Z., Schlaggar, B. L., Lee, J. M. and Culver, J. P. [2011]: 'Imaging of Functional Connectivity in the Mouse Brain', *PloS One*, **6**, p. e16322.
- Wright, J. [forthcoming]: 'The Analysis of Data and the Evidential Scope of Neuroimaging Results', *British Journal for the Philosophy of Science*.
- Wu, T., Wang, L., Chen, Y., Zhao, C., Li, K. and Chan, P. [2009]: 'Changes of Functional Connectivity of the Motor Network in the Resting State in Parkinson's Disease', *Neuroscience Letters*, **460**, pp. 6-10.
- Yeo, B. T., Krienen, F. M., Sepulcre, J., Sabuncu, M. R., Lashkari, D., Hollinshead, M., Roffman, J. L., Zöllei, L., Polimeni, J. R., Fischl, B., Liu, H. and Buckner, R. L. [2011]: 'The Organization of the Human Cerebral Cortex Estimated by Intrinsic Functional Connectivity', *Journal of Neurophysiology*, **106**, pp. 1125-65.

- Zheng, J. and Frey, H. C. [2004]: 'Quantification of Variability and Uncertainty Using Mixture Distributions: Evaluation of Sample Size, Mixing Weights, and Separation Between Components', *Risk Analysis*, **24**, pp. 553-71.
- Zhou, Y., Liang, M., Jiang, T., Tian, L., Liu, Y., Liu, Z., Liu, H. and Kuang, F. [2007]: 'Functional Dysconnectivity of the Dorsolateral Prefrontal Cortex in First-episode Schizophrenia Using Resting-state fMRI', *Neuroscience Letters*, **417**, pp. 297-302.

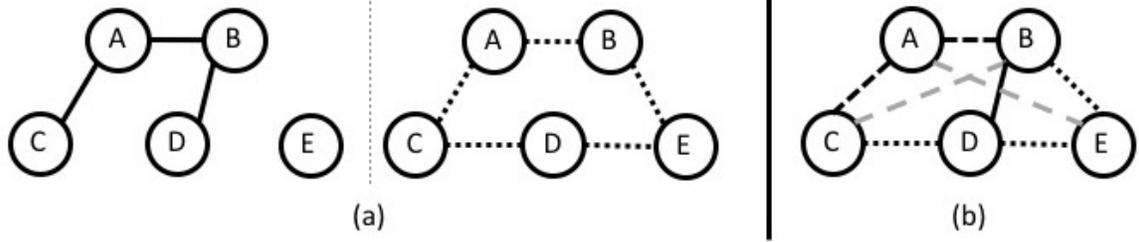


Figure 1: (a) Two possible 'cognitive function' graphs; (b) One possible inferred functional connectivity graph