

THROUGH A GLASS DARKLY: SCHIZOPHRENIA AND FUNCTIONAL BRAIN IMAGING

DAN LLOYD



ABSTRACT: Functional neuroimaging frequently assumes anatomically and functionally specific “modules” in the brain. Schizophrenia, which is challenging in many ways, also challenges modularity. A prominent example of a modular hypothesis links working memory (WM) deficits in schizophrenia with decreased function in the prefrontal cortex, a result confirmed in several studies. Nonetheless, these studies are systematically blind to the alternative possibility that schizophrenia reflects dysfunction in broader, more distributed systems of the brain. However, to establish hypotheses about distributed cognition in health and illness requires new interpretations of functional neuroimaging. Modularity can be assessed by large-scale meta-analyses of neuroimaging studies, illustrated here through the on-line database, Brainmap (www.brainmap.org). The apparent association of WM and prefrontal cortex fades when one considers all areas engaged in WM, and all tasks associated with prefrontal cortex. In addition, several global measures of connectivity based in functional neuroimaging may provide measures of distributed dysfunction. For example, Independent Component Analysis reveals unusual oscillations in distributed networks in patients with schizophrenia.

KEYWORDS: functional neuroimaging, fMRI, schizophrenia, Independent Component Analysis, WM, prefrontal cortex.

TO WILLIAM JAMES, conscious life was a stream; to Edmund Husserl, a flow. These metaphors point to the marvelous continuity of experience as it weaves through the world of thought and things. We might similarly talk about the flow of the body, as I reach for my cup of coffee. A physiologist could decompose the action, isolating the contribution of each muscle and joint to the whole. This functional analysis would constitute one form of explanation of the movement. As we replace “I grab the cup” with the physiologist’s account, there is a shift in level of description and a turn toward underlying processes, but the physiologist has added nothing to nature. Explanation of this sort, functional reduction, is simply a form of redescription that isolates and clarifies the details. After this sort of translation, there remains no dedicated cup-grabbing bodily component, but instead an integrated flow of overlapping micro-actions that jointly orchestrate and constitute the grasp of the mug.

Functional reduction is also a common form of explanation in cognitive science. Like the physiologist’s reduction of bodily motion, cognitive re-

duction tries to identify the neuronal muscles and joints that collaborate in the flux of experience. Here, explanation begins with functional decomposition, refracting experience into a simultaneous counterpoint of cognitive faculties, which are then reduced to implementation, ultimately zooming in on the interplay of neural circuits. We may not recognize our moment-to-moment experience as a flow of neurotransmitters, but that is what it is. In recent years, cognitive neuroscience has shown that the two reductions, of bodily motion and of cognition, are surprisingly similar in more than logical form, insofar as blood flow in specific areas of the brain varies with types of task. Neuroimaging exploits this regularity, resulting in an explosion of research in functional brain imaging. The collective data of these experiments seem to validate the overall project of cognitive reduction.

Schizophrenia, however, is a challenge. Successful reductive explanations require explicit “translations” from entities at one level of explanation (e.g., behavior) to those at another (e.g., functions of brain regions; Wimsatt 1976). Some psychological and neurological disorders, especially those originating in specific brain lesions, point to brain–behavior correlations—these are the classic success stories of neurology and neuropsychology. An explicitly cognitive theory of a mental disorder pursues functional decomposition and reduction within a broadly computational framework, conceiving of the disorder as a specific deviation from information processing in the healthy brain. Accordingly, a cognitive theory of schizophrenia follows the pathophysiology of the illness, but at a higher level, describing how physiological alterations affect perceptual or cognitive representations and processes that manipulate them. Confirmation of such a theory would come from two directions. Like any theory of any illness, a computational theory needs to predict some or all of the symptoms of the disorder. At the same time, observations of the neural mechanisms could confirm that the computational theory is in fact implemented in the brain.

Schizophrenia undercuts this agenda at every turn. The causes of the disorder are uncertain and probably disjunctive, including multiple genetic factors along with unknown environmental

influences. The alterations in the brains of schizophrenic individuals are multiple and various, and may reflect both primary effects of the illness mixed with chronic secondary effects and effects of treatment, along with differences that may have no functional significance (Niznikiewicz et al. 2003). Most important, the symptomatology of schizophrenia and related illnesses is complex, conceptually challenging, and highly variable over time and among individuals (Crow 1980; Liddle 1987). Thus, any attempted reductive explanation of schizophrenia is almost guaranteed to omit some variant of the disorder. Indeed, many authors have proposed that “schizophrenia” is a crude umbrella term for a set of disorders that have not been well distinguished (see Lloyd 2007, for discussion).

Thus, schizophrenia presents a test case to the strategy of cognitive reduction. In the present paper, schizophrenia is the mirror held up to a particular method for cognitive reduction, namely, functional neuroimaging. I argue that the experimental strategy shared by many neuroimaging studies is partially blind to the mechanisms that derail in this disorder. Thus, schizophrenia is not just a puzzle, but a deeper enigma that reveals and undermines some of the prevalent assumptions at work in neuroimaging and cognitive neuroscience in general. Methods based on different assumptions may lead to an enlarged understanding, not only of schizophrenia, but of the flow of healthy cognition as well.

The empirical challenge of schizophrenia (along with its urgency as a matter of public health) has made it a main research target in several disciplines. A recent search of the PubMed database identified more than 68,000 papers on schizophrenia, of which close to 2,000 involved functional neuroimaging (positron emission tomography [PET] or functional magnetic resonance imaging [fMRI]). A full review of this literature is well beyond the scope of this paper. Instead, I consider a methodological case study, extracting just one strand of research in the cognitive neuroscience of the illness. This example, however, can demonstrate both the standard logic of neuroimaging experiments and illustrate its stresses and breakdowns in the face of schizophrenia.

WORKING MEMORY IN SCHIZOPHRENIA

WM is the ability to maintain and manipulate items in memory over short time periods. Individuals with schizophrenia are generally slower and less accurate on WM tasks than healthy controls, suggesting that WM deficits may be related to negative symptoms like distractibility or perhaps to disordered thought (Kuperberg and Heckers 2000). WM has been probed in thousands of studies and in many experimental paradigms, of course. Among them, one experimental WM task is the “*n*-back” task. In the *n*-back task, subjects listen to or watch a sequence of letters, numbers, or other stimuli, and respond if the current item matches the *n*th preceding item. The immediately preceding stimulus is 1-back from the current target; the stimulus before that is 2-back, and so forth. Stimuli can be of any type, but the majority of *n*-back experiments use letters or numbers. Certainly the *n*-back task requires a cognitive “scratch pad” capable of repeated and rapid updating (most people find 3-back very difficult).

n-Back matching has been scanned, and often activates the dorsolateral prefrontal cortex (DLPFC), encompassing Brodmann areas 9 and 46 (Figure 1). This activity increases with the difficulty of the task. Other areas respond too, but these vary with the kind of stimuli (see the discussion in Barch [2005, 325].) Accordingly, DLPFC is often associated with aspects of WM distinct from the initial perception and encoding of stimuli, namely, the “executive function” of manipulating representations, including the rapid update of the scratch pad. In short, WM seems to fit well into the pattern of functional decomposition and reduction.

Several studies relate impaired DLPFC activity to deficits in WM, so the nexus of the region and the cognitive capacity (CC) serves as a methodological case in point for functional neuroimaging. I begin with a single study, and slowly enlarge the circle of research around the issues it raises. Menon et al. (2001) studied subjects with schizophrenia in comparison with healthy controls using a 2-back task consisting of auditory numbers. From a study of eleven men with schizophrenia and thirteen

control subjects, the authors observed that several brain areas were less active in the schizophrenia group: DLPFC, the frontal operculum, and the inferior and superior parietal lobes. In further discussion, the authors refine their conclusions, pointing more explicitly to the DLPFC:

The largest brain activation differences between the groups occurred in the DLPFC, a subregion that has been implicated in executive functions involved in verbal WM. Our results suggest that DLPFC deficits occur in both hemispheres. . . . This is consistent with fMRI studies indicating that both hemispheres, rather than just the left hemisphere, are involved in verbal WM processing. . . . Further, several PET and fMRI studies have shown a correlation between left and right DLPFC activation and increased memory load. (Menon et al. 2001, 438–9)

The passage is typical of neuroimaging research in that it moves one of the four brain areas detected in the experiment onto center stage. It does this in light of a body of other research referenced in this passage and throughout the article. This background sustains and supports the specific claims in the Menon et al. (2001) paper, and vice versa. Some of the supporting background is found in other schizophrenia studies, but a great deal is found in a broader literature of neuroimaging linking DLPFC and WM. I consider these contexts in order.

Several schizophrenia studies have examined WM, and have reported a decrease in activity in DLPFC in those with the illness (Barch et al. 2002; Callicott et al. 1998; Callicott et al. 1999; Callicott et al. 2000; Callicott et al., 2003; Carter et al. 1998; Jansma et al. 2004; Mendrek et al. 2004; Menon et al. 2001; Meyer-Lindenberg et al. 2002; Perlstein et al. 2001; Wykes et al. 2002). These so-called hypofrontality results are among the most replicated observations in schizophrenia research (Callicott 2003). This is good news with respect to the specific claims of the Menon et al. paper. However, not every study conforms, and it is among the variations of the *n*-back theme that troubles for hypofrontality may be found.¹

In a very useful article, Deanna Barch (2005) reviewed fourteen neuroimaging studies of schizophrenia using the *n*-back task. Of these, nine showed decreased activity in DLPFC in at




FIGURE 1. THE DORSOLATERAL PREFRONTAL CORTEX (BRODMANN AREAS 9 AND 46).

least one hemisphere. As Barch states, “the modal finding has been of decreased activation in DLPFC in schizophrenia” (p. 333). However, this result was not confirmed in five other studies, which showed either no change in DLPFC or (in three papers) increased activation—hyperfrontality rather than hypofrontality (Callicott et al. 2000; Honey et al. 2002; Honey, Fletcher et al. 2002; Jansma et al. 2004; Sabri et al. 2003; Walter et al. 2003). Moreover, seven papers from the set reported significant changes in other areas, often the posterior parietal cortex. Not surprisingly, the hyperfrontality results have provoked an ongoing discussion in the literature.

Rather than enter this debate, the goal here is to probe its assumptions. Foremost among these is the supposition that DLPFC has a fixed role in WM in healthy cognition and that it is the exclusive implementer of that role. In this case, the issue of hypo- versus hyperfrontality is understood as a reflection either of group or experiment differ-

ences, rather than as challenges to the presumed function of this region of the brain. The papers reviewed in Barch (2005) pass the assumption back and forth among them, and together form a web of assurance that WM is correctly decomposed into certain functions, and that one of these is the dedicated job of one part of the brain. This assurance in turn intersects with a large body of research in healthy subjects that also mutually affirms this instance of modularity. The DLPFC certainly seems to be involved in WM. But in nearly all the papers in this research literature, this “involvement” understates or omits qualified claims like the following: the DLPFC, *along with other areas*, is most definitely involved in WM, *as well as other functions*. The implications of these clauses turn out to be significant: DLPFC is not wholly responsible for WM, nor is that its only function. Indeed, in this case the association between region and function turns out to be surprisingly weak.

SCHIZOPHRENIA CONFOUNDED

Our understanding of schizophrenia, such as it is, is framed by a more general understanding of normal brain function, informed by an expanding background of neuroimaging research. With respect to this methodological study, a background constant is the normative function of the DLPFC as a site for the implementation of some aspect of WM. The strong form of the “DLPFC–WM” hypothesis holds that the function of executive control in WM is localized in DLPFC, and conversely, that DLPFC is dedicated to the function of executive control. (A specific WM deficit in schizophrenia is accordingly attributed to a dysfunction of DLPFC.) Together, localization and dedicated function comprise a hypothesis of *modularity* for this function and, ipso facto, this region of the brain.

In this context, then, modularity is a hypothesis about implementation of cognitive functions. In general, if a cognitive function is modular, then two conditions are met: first, the function is localizable in distinct regions of the brain (which need not be contiguous); second, the regions are dedicated to the function they implement, and do not contribute to other functions. Modules in this sense are quite unlike the modules proposed by Fodor (1983), which are domain-specific input processors defined via their computational properties. Fodorian modularity is noncommittal about implementation, whereas modularity in neuroimaging arises from observations of functional localization. This sort of modularity attempts to determine what specific areas of the brain are doing, expressed in cognitive terms.

The alternative to modularity is some form of distributed processing: cognitive functions are served by collective activity of multiple brain regions, or conversely, particular regions of the brain each contribute to several distinct cognitive functions. Distributed processing in this context is adapted from the corresponding concept in connectionism. Connectionism, also known as Parallel Distributed Processing or Artificial Neural Network Modeling, has emphasized the computational powers of networks of neuron-like processors, where each contributes in small ways to the

function of the whole (McClelland and Rumelhart 1986; Smolensky 1988). Functional neuroimaging methods cannot discern individual neurons, but they can see larger regions of the brain; thus, distributed processing here refers to processing across and between regions. Between strict modularity and maximal distribution (equipotentiality), there is likely a continuum of partial divisions and overlaps. Somewhere on this continuum lies the human brain.

Most functional neuroimaging studies leave modularity to the imagination. Explicit conclusions are more cautious, linking brain regions with cognitive functions loosely. The qualified verbs used above by Menon et al.—“implicated,” “involved,” and “correlated”—are typical of this research, and reflect a reasoned awareness of the limits of logic and method in neuroimaging. Nonetheless, I argue that even these careful conclusions underestimate a methodological blindness in the standard interpretations of many neuroimaging experiments. By their design, some of the most popular methods cannot discriminate modular from non-modular hypotheses, and consequently leave behind a powerful impression of modularity, despite their cautious language.

Herein I outline one experimental strategy often deployed in PET and fMRI. fMRI, the technology involved in 53,000 papers about the brain indexed in the PubMed database, is barely into its second decade. Within the cornucopia are hundreds of methodological discussions and data processing innovations. Yet there is a strong central tendency (Jessard et al. 2001). To explore some CC, here is the recipe for many studies.

1. Decompose: Analyze the CC in terms of component information processing functions.
2. Operationalize: Adapt or develop a task that exemplifies a subset of the functions composing the CC.
3. Image the task and more: Scan subjects performing the task repeatedly and performing different tasks, or “resting.”
4. Average the images: For each of the experimental conditions (tasks, non-tasks), create a single mean image.
5. Contrast task versus non-task images: Compute a “difference image,” for example, the task image minus the non-task image (often called “cognitive subtraction”).

6. Interpret the blobs: Find and group significant regions of the difference image, “areas of activation,” and identify their anatomical locations.
7. Conclude: Activity in the identified brain regions implements the task, and therefore partly or wholly implements the CC.

In several respects, the protocol reflects limits of the technology. Scanners constrain movement and immobilize a subject’s head for the duration of the scan, so tasks must be adapted accordingly. More important, the metabolic signal detected by the scanner is embedded in noise. Averaging groups of images allows the random variations in individual images to cancel out. Even after averaging, the signal difference from engaging in a task is small, around one percent or less. These are visible (and significant) only after the subtractive contrast.

The last point is important for all consumers of images of glowing blobs in the brain. Most of these pictures in every media—journals included—are contrast images. All show just the differences between two or more states of global metabolic activity. Every contrast image, then, should provoke the question, “Compared with what?” Not surprisingly, the baseline for the comparison makes a difference in the conclusion for any imaging experiment. In fact, imaging experiments in this research paradigm encounter a dilemma in the preliminary stages of decomposition and operationalization. The validity of an experimental contrast depends on an assumption that the capacity under study can be decomposed into *separate and exclusive* components, or in other words that the component functions engaged in each experimental condition (tasks and non-tasks) are implemented in nonoverlapping regions of the brain. As a consequence, by this assumption the addition of an experimental task and its correlated brain activation has no effect on the activation already provoked by the baseline—an assumption of “pure insertion” of the additional task (Friston et al. 1996). Modularity is therefore assumed in the interpretive strategy.²

A schematic example can illustrate the problem. Suppose that our analysis of some CC suggests that it comprises component functions *A*, *B*, and *C*, and these are implemented by some unknown combination of just three regions in a (small) brain.

Further suppose that the exclusivity assumption is violated, and that the functions are implemented in overlapping brain regions: *A* depends on areas 1 and 2; *B* depends on areas 1 and 3; *C* depends on areas 2 and 3. Figure 2 schematizes this functional architecture. These are the stipulated facts of the scenario. Will subtractive “brain imaging” be able to discover them? Following the subtractive recipe, we collect and average “raw” images during each of the three tasks, and derive the contrasts. Here is the full list of possible outcomes:




FIGURE 2. A SMALL BRAIN: THREE COGNITIVE FUNCTIONS—*A*, *B*, AND *C*—DISTRIBUTED OVER THREE NUMBERED “BRAIN” REGIONS. (HYPOTHETICAL FUNCTION *A* IS IMPLEMENTED IN REGIONS 1 AND 2, ETC.) ANY SUBTRACTION OF ONE FUNCTIONAL STATE FROM ANOTHER WILL EXCLUDE REGIONS THAT SUPPORT BOTH FUNCTIONS. SHARED OR DISTRIBUTED RESOURCES ARE INVISIBLE IN CONTRASTIVE IMAGES.

$$\begin{aligned}
 A(1,2) - B(1,3) &= \text{region 2} \\
 A(1,2) - C(2,3) &= \text{region 1} \\
 B(1,3) - A(1,2) &= \text{region 3} \\
 B(1,3) - C(2,3) &= \text{region 1} \\
 C(2,3) - A(1,2) &= \text{region 3} \\
 C(2,3) - B(1,3) &= \text{region 2}
 \end{aligned}$$

The math is strange, in part because only positive values are kept in typical image subtractions. (*A* – *B*, for example, leaves behind not only region 2, but a “deactivation” of 3. More on this below.) More important, the results are empirically wrong in the scenario: *none* of the possible subtractions correctly identify the regions involved. Two experiments about function *A* could yield contradictory results, yet both would be correct. Nor does it help to contrast *A* with the conjunction of *B* and *C*, despite the appearance that this would sharpen the result (the function versus everything else). In

this case, $A - (B \text{ and } C)$ excludes all of the areas supporting A . In short, when functions overlap in their implementations in the brain, the strong conclusion that a contrast image shows the areas dedicated to the function is invalid. What can be concluded is that the revealed areas are part of the implementation, but that none of the other areas can be excluded. Research reports based on the subtractive strategy always hedge their strongest conclusions, but they do not always acknowledge the potential for omission of areas no less involved.

In practice, a dilemma arises for neuroimaging. On the one hand, to understand the computational operation of specific brain regions requires a fine-grained characterization of the functions the regions compute, and equally fine anatomical distinctions. On the other, to make these distinctions the contrastive functions are quite similar in type, making it harder to suppose that they are exclusively implemented in nonoverlapping regions of the brain. This is illustrated, once again, in the 2-back experiment. The Menon and colleagues' conclusions rest on a double contrast. The first contrast differentiated an auditory numerical 2-back task from this baseline task: press a button every time you hear the number "3." Subjects practiced both the 2-back task and the simpler target detection control task several times, and then in the scanner the two tasks alternated in blocks, with eleven trials in each block, for a total of twelve blocks. So each block comprised a stream of digits; all that differed was what subjects were supposed to do. Then the contrasts derived for each group were further contrasted to reveal differences between the schizophrenia contrast and the healthy control contrast.

Should we expect these two tasks to draw on exclusive brain functions? Clearly, some areas of overlap are irrelevant. For example, both tasks involve hearing numbers and pushing buttons, but those functions are presumably distinct from the WM functions of interest. A trickier question arises around WM itself. The 2-back and the control task both involve WM in that subjects must maintain a scratch pad representation of which task is required over each block. The 2-back task

adds a significant WM load, and in healthy subjects it seems reasonable to suppose that the load pushes brain areas specific to WM into high gear. Is this contrast going to be the same in subjects with schizophrenia? Individuals with schizophrenia have trouble with a variety of memory tasks. It seems plausible that both the control task and the job of switching tasks from block to block (i.e., staying "on task") will be more difficult for them. In response to this background load, subjects with memory difficulties already taxing their WM capacities already. The "reduced" activation in the 2-back images might in fact be an artifact of increased activation in the control conditions.

In general, then, the contrast between control subjects and those with schizophrenia can be affected by differential activity in either or both of the baseline and the target task conditions. Part of the challenge of schizophrenia lies in its multiple and pervasive effects on cognition, and this makes every contrast less certain. Moreover, the experimental design also makes other assumptions of common ground between both subject groups. For example, it is assumed that subjects are all equally motivated to perform both tasks, that stimuli are equally salient and meaningful, that subjects from different groups do not employ different strategies, and that extraneous responses—for example, emotional responses—are similar in both groups. Given the breadth of symptoms possible in schizophrenia, one or more of these conditions may confound the contrast between the two groups. In short, the methodological issues that undermine the traditional approach magnify in any comparison between schizophrenic and normal cognition. In effect, symptoms will mingle with any task, and this uncertain mix will be reflected by an equally uncertain mix of regional activity in the brain, some of it reflecting differences in processing underlying the task, some reflecting the symptoms, and some unrelated to either (but still detectable in image contrasts). These confounds may be overcome in the group average, but on the other hand this is a science of small differences. The devil may be in the details.

WHERE DOES WORKING MEMORY WORK?

In the previous section, I argued that the subtractive method in neuroimaging favors apparent modularity, and that in one experiment in schizophrenia research the assumption of exclusivity may not be correct. However, distributed processing remains a mere possibility, and an unsubstantiated hazard of the traditional methods. A genuine empirical study of distributed processing in the brain will need new methods (e.g., Price and Friston 1997; Friston et al. 1996); in the next section, I consider some opportunities and preliminary results. Meanwhile, the majority of imaging studies are framed by the subtractive method. Can these studies be reconsidered in some way that enables us to assess the alternatives?

The small brain example (Figure 2) suggests two strategies for overcoming the limitations of individual contrasts. Both depend on increasing the variety of contrasts considered: first, we could begin with a brain region and consider all the contrasts that activate it. In the schematic example, region 1 was the residue of two contrasts, two notional “tasks.” Separately, neither contrast correctly identified the disjunctive function of the region, but taken together they correctly identified hypothetical functions *A* and *B*. In general, then, a “brain first” strategy works region by region through the brain, and identifies all the contrasts that “light up” that region. Collectively, the disjunction of all the functions targeted by these region-specific contrasts characterizes the overall function of the region more completely than any single contrast.

Second, we could begin with a function and collect all the contrasts that use tasks that implicate the function. In effect, a “function first” strategy diversifies the baseline conditions by which the contrast is measured. In the example, two contrasts probe function *A*, extracting two different contrastive functions. Neither of those contrasts completely identified the involved brain regions in the function, but conjointly they do correctly indicate regions 1 and 2.³

The small brain example affords a complete reduction of function to region(s) by expanding

and diversifying contrasts, and grouping them by region (the brain-first strategy) or by function (the function-first strategy). They illuminate a conceptual shift that occurs when distributed processing is a possibility. If modularity is assumed, then the various contrasts possible in the small brain data are inconsistent, and as a result the inconsistency must be explained away by further elaborations in the interpretation of the diverging results. Many discussion sections in this literature do exactly that. However, in the nonmodular framework, the apparent inconsistency is the predicted effect of diverging target or baseline tasks in various experiments. Differing contrast results triangulate the underlying distributed function. In this way of thinking, the discrepancy is itself the positive evidence for the underlying distributed processing. It does not need to be explained away. (However, if two experiments use the same target and baseline tasks, then their results should agree, just as individual subject results within one experiment should agree.)

The smallness of the example makes the two strategies look easy, but the net must be cast widely to overcome the illusion of exclusivity. In principle, every neuroimaging experiment yields results that constrain every cognitive function, as well as the function of every area of the brain. With thousands of studies in print, this is a tall order. Fortunately, however, several visionary researchers have seen the need for synoptic databases that compile the results of many imaging studies. One of these databases is BrainMap, maintained with a web interface at the University of Texas at San Antonio (available: www.brainmap.org). BrainMap records the loci of peaks of activation revealed by experiments using the traditional paradigm, along with brief descriptions of the experiment, citation details, the subject groups, the imaging modality, and the type of cognitive function probed. That is, it condenses many published results in functional brain imaging. WM is among the 37 behavioral domains distinguished, and listed in Figure 3, discussed below. (Another ambitious archive, the fMRI Data Center, offers summary data along with the raw image data itself, pixel-by-pixel. See www.frmidc.org.) As of January 2007, BrainMap archived 959 papers reporting 4,221 different

experiments. Collectively, this body of research reports 34,223 significant peaks of activation. (That is, even in the published reports guided by the expectations and the methods of modularity, on average each experiment revealed around eight regions of activation.)

Nine hundred papers make up a small and self-selected subset of imaging research, and it would be difficult to determine whether the sample is representative of neuroimaging as a whole. Nonetheless, these 4,000 experiments afford a far broader picture of the field than any review article I have encountered on any topic. As a first approximation, one can sift these records with the questions of modularity in mind. Continuing along the narrow path of this paper, we turn again to DLPFC and WM.

One tacit support of the hypo/hyperfrontality hypotheses in schizophrenia was a general view about the normal function of DLPFC, as a main component in the neural implementation of WM. BrainMap enables us to test this assumption with two questions:

1. What does DLPFC do in general? That is, what tasks (when contrasted with various baselines) reveal differential activity in this region of the brain? The modular view suggests that these will be cognitive tasks, heavily tilted toward WM or perhaps other functions of executive control.
2. What parts of the brain serve WM? Here a modular hypothesis leads to a strong expectation that DLPFC will be a constant.

A closer look at the studies in the database can cast some light on these questions. To this end, I conducted a meta-analysis of 478 papers based on fMRI, comprising 1,249 experiments reporting 8,179 peaks of activity, drawn from the Brainmap database in late 2006. This analysis included all studies that met the following criteria: fMRI was the imaging method, studying normal subjects (i.e., excluding patient groups, children, seniors, special subpopulations like musicians, for example; experiments separating male and female results were included if both were available). Finally, the analysis was limited to “main effects” (i.e., positive activations, the experimental result of the contrast between the task and a control condition. Accordingly, the review excludes published meta-analyses,

probes of modulatory effects, parametric studies, and any experiments where it was not clear that loci were main effects.) For present purposes (i.e., understanding the function of prefrontal cortex), the meta-analysis assigns each peak of fMRI signal activity to either the Brodmann area (and hemisphere) in which it occurs, or to right or left “sublobar” regions, a “none of the above” with respect to cortex. The analysis counts each significant reported locus as a single hit (for a full bibliography and more detailed discussion, see <http://www.trincoll.edu/~dlloyd/brodmann.html>). Henceforth, “database” refers to this subset of information from Brainmap.

We begin, then, with the question of DLPFC function. (As shown in Figure 1, the DLPFC comprises Brodmann areas 9 and 46 in both hemispheres.) In this database, there are 137 instances of DLPFC activity driven by WM tasks. *Prima facie*, this suggests a strong connection between region and function. However, this is immediately qualified by several considerations. The first qualification is not too disquieting: it turns out that attention tasks drive DLPFC more frequently than WM tasks: 169 hits (versus WM, 137). Nonetheless, this seems consistent with DLPFC as the seat of “executive function.” A second issue, however, is much more serious. We must guard against a base-rate fallacy and a confirmation bias that would follow. For example, suppose the authors contributing to the database were very interested in WM, and as a result 90% of the contributed experiments probed that function. Further, suppose that DLPFC activates completely at random, with no regular correlation with any particular type of cognition. In that case, DLPFC experiments are likely to implicate WM nine times out of ten, just because nine of ten experiments feature WM tasks—vivid, but completely illusory, confirmation of a region–function link. As it happens, this is exactly the situation here. The top two favorite experiments of this research community are probes of attention (173 experiments) and WM (143 experiments).

This base-rate bias is easily corrected by normalizing (dividing) the census of activation hits by the number of experiments in each domain. Figure 3, accordingly, shows the percentage of




FIGURE 3. Comparative DLPFC activity in thirty-seven behavioral domains distinguished in the BrainMap database. Bars indicate total dorsolateral prefrontal activity in the 1,249 experiments analyzed. Each bar is divided to show relative contributions of the specific DLPFC Brodmann areas. (The contribution of each area is measured as the fraction of experiments in each domain that reported activity. Thus the sum of contributions of multiple areas can exceed 1.) From left to right, each bar sums left area 9, right area 9, left area 46, right area 46. Working memory is marked with an asterisk (*). The figure suggests that working memory is not the only job for DLPFC, which is strongly activated by most cognitive domains and other tasks as well.

experiments in each domain that activate the four areas of DLPFC. After this adjustment, WM still trails attention, but not only that: six other broad behavioral domains generate more DLPFC action than WM, including action inhibition and motor learning, reasoning, phonology, temporal cognition, and even tasks evoking anger. Tasks in fourteen domains activate all four areas of DLPFC. Thirty-three domains activate at least one DLPFC region. That is, DLPFC is involved in just about every behavioral domain—only anxiety, hunger, sexuality, and speech are exempt. It is still true that DLPFC is activated by WM . . . along with other behaviors. Now that the extent of those other tasks is apparent, the limits to the claim of functional specificity are conspicuous.

The converse question can also be assessed: what regions of the brain support WM? Figure 4 condenses the 143 WM experiments in the database, totaling the number of reported activation loci in each cortical area. The four DLPFC areas are certainly in the running, but once again, the DLPFC areas are not the most prominent. The graphic suggests the widely distributed areas that seem to be involved in WM tasks. Indeed, 86% of all reported hits are not in DLPFC. They fall broadly in many (although not all) cortical regions, 71 of 84 Brodmann areas.⁴

These statistics are surprising in part because a different approach to the database could yield an apparently much different result. Armed with BrainMap alone, a fan of the marriage of prefrontal cortex and WM could correctly reference thirty-six papers reporting on eighty-four experiments in which the connection is confirmed via at least one activation peak in DLPFC. But this is to overlook the sea of activity and library of experiments from which these fortunate few were drawn. In particular, it overlooks fifty-eight WM experiments in which *no* DLPFC activity was detected. The results of these experiments are all valid, considered one by one. DLPFC is involved in WM. But this is neither the sole role of the region, nor is WM exclusively implemented through circuits including DLPFC. Even more dismaying, all of these results are contrasts, in which a baseline task was chosen in such a way as to isolate some aspect of WM function, often with the hypothesis that

prefrontal activity would be provoked thereby. As discussed, all those contrasts rest on an assumption of exclusivity, which the meta-analysis suggests is unlikely to be true.

What then is the functional job of the DLPFC? It seems to have no single function, but to contribute to many aspects of cognition.

What parts of the brain serve WM? As noted, forty percent of WM experiments in the database do not involve DLPFC. The areas that do serve WM are distributed, about seven regions on average. And they are disjunctive. As the discussion of this well-known region suggested, the network serving memory depends greatly on the task, and even more (and in unknown ways) on the baseline used in the contrast. The search for WM in 478 papers has not only failed to identify a single key component, but also has not identified a consistent network of activity serving this cognitive function.

INDICATORS OF DISTRIBUTED DYSFUNCTION

The traditional framework for functional neuroimaging is attuned to specific regions of the brain that reliably correlate with specific cognitive tasks. This dictates a parallel approach to psychological disorders, which will be ascribed to specific dysfunctions of specific brain regions. But schizophrenia did not fit. The attempt to localize a deficit in WM led to conflicting results. The inconsistencies may be due to several unknown variables. This paper considered the underlying assumption of modularity, in this case in the form of the hypothesis that one neural component, DLPFC, is specifically associated with WM function. Meta-analysis of experiments involving DLPFC, WM, or both challenged this assumption. Schizophrenia fails to fit in a modular scheme, but neither does healthy cognition. Indeed, the variability found in reviews of schizophrenia papers seems to approximate the variability hidden in a much larger set of functional neuroimaging results about DLPFC or WM.

It should be acknowledged that these inferences rest on a single, constrained meta-analysis. We considered everything in the realm of one brain region, and one cognitive function. But the link




FIGURE 4. BRAIN AREAS ACTIVATED BY WORKING MEMORY TASKS. BARS REPRESENT THE NUMBER OF EXPERIMENTS IN THE ANALYSIS THAT REPORT ACTIVITY IN EACH BRODMANN AREA, ARRANGED SYMMETRICALLY BY HEMISPHERE. DORSOLATERAL PREFRONTAL CORTEX AREAS ARE CIRCLED. THE FIGURE SUGGESTS THAT WORKING MEMORY IS SUPPORTED BY SEVERAL REGIONS IN ADDITION TO DLPFC.

considered, between DLPFC and WM, is exemplary, and multiply confirmed in both healthy subjects and as a specific locus of difference in individuals with schizophrenia. The literature of cognitive neuroscience often reasserts and rediscovers the assumption. That it should be in part an artifact of the method is surprising, and spreads a *prima facie* doubt across modular tendencies everywhere in neuroimaging. This more general survey is a subject of ongoing research, but is left open here.

An alternative approach to schizophrenia research would take distributed processing more seriously, and attempt to conceptualize the illness as the result of distributed dysfunction. Just as strong modularity was the stalking horse above, we can consider a strong form of distributed dysfunction. Where, then, is the dysfunctionality located? It is not to be found in one or even several specific regions of the brain, but rather in some more general characterization of computational function that applies to many areas of the brain. Twenty-two experiments in the database involve schizophrenia patients performing WM tasks (all of them *n*-back tasks). I compared the normalized activity in all Brodmann areas in patients with that of healthy controls. Three DLPFC areas (Brodmann area 9 left and Brodmann area 46 bilaterally) are depressed in individuals with schizophrenia. Area R 9, however, is elevated, and this is the largest difference between the two groups. But several areas also vary: seven others are reduced at least as much as the DLPFC areas, and four are elevated nearly as much as R 9. Altogether, at least ten areas are conspicuously altered (more than a 10% difference). With the net cast just slightly more broadly, there is *prima facie* reason to doubt both versions of frontality hypotheses with respect to the more general function of WM. At a minimum, it seems that even within the domain of WM, the task used in the experiment makes a large difference in the resulting contrast. Considering the intersection of schizophrenia, WM, and DLPFC, this survey leads to a conclusion similar to that of Barch and others: neuroimaging leans toward a role for DLPFC in schizophrenia, but with serious unresolved doubts.

Overall, schizophrenia presents many indicators of a distributed dysfunction. These include its diversity of symptoms and cognitive deficits,

its multiple sites of anatomical and functional difference, its multiple genetic and environmental determinants, and its relationship to abnormalities in one or more broadly projecting neurotransmitter systems (with an enduring focus on dopamine. Braver et al. 1999). Accordingly, one strand of schizophrenia research considers the illness not as the effect of one or several regional dysfunctions, but rather as a defect in the circuits connecting several brain areas (e.g., Niznikiewicz et al. 2003). Schizophrenia, on this view, is a “disconnection syndrome.”

Several versions of disconnection hypotheses have been proposed, with no consensus as yet (e.g., Andreasen et al. 1999; Friston 1999; Frith 1992). How can neuroimaging contribute to this research? Connectivity can be indicated by the co-activation of brain regions, on the assumption that if two regions of the brain are strongly correlated in their activity over time, their correlation must reflect either an effective connection between the areas, or some third source of effective input to both areas. Functional neuroimaging data can inform this research, in conjunction with techniques for measuring correlations across time and space in the brain. Some of these techniques are rather technical, and will not be surveyed here. One method in particular, however, has shown initial promise. That method is independent component analysis (ICA; Calhoun et al. 2002; McKeown et al. 1998). ICA, a form of “blind source separation,” is a statistical method for identifying hidden patterns of sources as they blend together in mixed signals. Applied to functional neuroimaging, ICA extracts ensembles of brain regions that co-activate in a series of images over time. ICA is blind in the sense that it requires no specification of the timing of the task(s) during an experiment, and no prior isolation of regions of interest. Instead, it discovers both the anatomy and time course of a set of components that jointly account for a very large part of the variation in the raw data.

Application of ICA to schizophrenia has only begun. Conceptually, it could intersect with a disconnection hypothesis as follows. First, ICA could identify components—networks—in common between schizophrenia subjects and healthy controls. Then those networks could be compared

to identify differences in function and functional anatomy that could predict the manifestations of the illness. For example, Garrity and colleagues (2007) used ICA to compare twenty-one schizophrenia patients with twenty-two healthy subjects during an “auditory oddball” task. The task is simple: identify (with a button press) beeps at a certain pitch, scattered at random among many beeps of a second pitch (which should be simply ignored). Infrequently, the beep stream is interrupted with an auditory oddball, another noisy sound effect that is distinct from both beeps. This too is to be ignored. In this case, the researchers were not interested in components that activated at the target tones, but rather just the opposite—that is, those components that deactivated during the task. In recent years, a steady network has been identified that tends to be active off task (i.e., it is deactivated during the experimental task). This network, consisting of the posterior cingulate cortex, medial prefrontal lobe, and parahippocampal and inferior parietal cortices (among others), is similar regardless of the task that flanks it. Raichle and associates named it the “default mode” (DM) network, because the brain seemed to activate it by default as attention to demanding tasks subsided (Raichle et al. 2001; McKiernan et al. 2003). (Traditional methods paid little heed to the network because it was deactivated by most tasks.)

ICA easily detects DM networks in both healthy individuals and those with schizophrenia. (In this case, DM components were selected by their anatomical match with a standard DM network. Subsequent analysis revealed that they were equally picked out by the time course of the task, which deactivated just this set of regions.) So, both anatomically and functionally, the DM components are roughly similar in the two groups. They differed, however, in one functional respect: in patients with schizophrenia, this network tended to oscillate more at a higher frequency. The component activity in those with schizophrenia traced a more erratic time course than in those without the illness. The interpretation of this observation is speculative at this point, depending as it does on our very incomplete understanding of the functions of the DM (Gusnard and Raichle 2001; Fransson 2005; Greicius 2003; Greicius

and Menon 2004). But, even with the bare bones characterization of DM as the state one settles toward when attention is not being absorbed by some task, the schizophrenia result is suggestive. Individuals with schizophrenia seem to engage and disengage DM more often than healthy controls. Compared with the controls, the patients’ brains were either slipping more often from attention to their task, or finding their attention more frequently engaged by some non-task stimuli. Some of the cognitive deficits observed in schizophrenia seem to fit the pattern suggested by this analysis of one independent component. It invites further exploration using this nontraditional method.

The observations from Garrity and colleagues and the disconnection hypotheses in general invite attention to two avenues of further research into schizophrenia. It is striking, first, that Garrity and associates found a difference in schizophrenic cognition in a component not directly implicated in the experimental task. This is consistent with the overall regard for distributed processing. Second, the observations identified a difference in the temporal dynamics of the brain. Like every aspect of schizophrenia that has been surveyed herein, the data of schizophrenia refract through the lens of healthy cognition. Cognition requires getting the timing right. More broadly, cognition requires the coordination of information and behavior over time as well as space. WM is just one manifestation of this temporal juggling act. Other forms of memory along with the routine processing of information in time all require a highly organized, temporally inflected cognitive facility. A minimally necessary condition for temporal coordination, then, is the ability to retain information from one moment to the next, and the complementary ability to anticipate what is about to happen. (In Lloyd [2002, 2003], I argue that this capacity is also fundamental to an understanding of human consciousness.)

Temporality seems to be the infrastructure of cognition, implicated in tasks of every sort, and schizophrenia seems to assault the brain in several ways at once. This suggests that our distributed perspective on schizophrenia should seek global signs of health and illness, and seek interpretations of those signs within the framework of temporal

cognition. Global measures of cognition would consider dynamic properties of all the components together, namely, the conjoint function of all regions of the brain. Functional neuroimaging supports this sort of analysis as well. For a final illustration of the nontraditional application of neuroimaging, I briefly discuss one measure of temporal cognition in schizophrenia.

Twenty subjects from each of the two groups in the Garrity and co-workers' (2007) study were reanalyzed, thanks to the generosity of the authors in sharing these data. Specifically, I used the independent components that this team had already extracted from the image series. The goal was to identify a marker of the capacity to retain information about the immediate past and anticipate the immediate future. This capacity must be translated into some measure detectable in an image series. In an image series, then, each image should contain information about more than just the present environment. In principle, in the image there should be detectable overlays of both the past and the future. At the neural level, information about the past is represented in information about prior images in the series, and likewise about the future. That is, from each image it should be possible to extract information to partly specify the prior and next images in the series. To test these hypotheses, I used simulated neural networks to measure the information available in each image about the images flanking it in the time series. In normal subjects engaged in a variety of tasks, this hypothesis is confirmed. (See Lloyd [2002] for a full description of methods and results.) Interestingly, brain images represent their immediate past in more detail than their anticipated future. Turning to the Garrity and associates' data, in both subject groups the same preservation and prediction could be detected in the image series. Both groups displayed a continuous capacity for short-term information storage and short-term anticipation. However, in the patient group, this capacity was diminished. Specifically, subjects with schizophrenia anticipated less information to accurately specify the immediate future (+1.5 seconds).

In sum, individuals with schizophrenia differ from healthy subjects in one global measure of

change and continuity over time. Consistent with a general hypothesis of distributed dysfunction, the observation suggests that a generalized temporal discontinuity characterizes schizophrenic cognition. This discontinuity in turn implies an impaired ability to maintain the manifold constructs that jointly constitute the perception of a stable world and stable agency within that world.

Regardless of interpretation, there is a methodological point here that invites further exploration. The measure described is truly global, in that it is a general characterization of the total ensemble of brain components. All the activity detected in functional neuroimaging (fMRI, in this case), contributes to these measures. This is one way to take distributed processing seriously. Moreover, this measure is more data driven and, by the same token, less reliant on *a priori* hypotheses about neural components and their functions. Following the critique of the traditional paradigm, this also seems like a good direction to explore. At this time, secondary analysis is underway of Alzheimer's patients, healthy older subjects, and children. As more comparisons emerge, it is hoped that the understanding of global function of the brain will increase.

CONCLUSION

Brain injuries and disorders are inevitably tragic for those who suffer them. But for more than a century, these arbitrary experiments of nature have been a principal source of understanding of the function of the brain in both health and injury. The view of brain function that seemed to come into focus first was modular. The specific deficits that attended specific lesions pointed strongly to anatomically distinct components with characteristic functions. The modular view also fit well with classical cognitive science and the various computational analogies that informed research for a generation. Subsequently, functional neuroimaging developed interpretive methods that assumed modularity. But, being an assumption, modularity was not itself tested in any of these experiments. The contrastive methods of neuroimaging remain systematically blind to nonmodular dynamics in the brain.

Schizophrenia fails to fit readily into the modular scheme. Contradictory observations and a shifting set of symptoms puts pressure on modular explanations. There are proposals to save the modular appearances, but in this paper I followed a deeper strand of doubt, leading to the possibility that schizophrenia is a global dysfunction of a highly distributed neural system. Schizophrenia overflows modularity mainly because healthy cognition is itself highly distributed. In a companion paper (Lloyd 2007), I discuss the consequences of distributed dysfunction further, and with reference to a specific dynamic system architecture. Here, considering the contributions of neuroimaging, some possibilities for new empirical interpretations have been proposed. The best interpretive methods are doubtless yet to come. It will take both imaging and imagination to unravel the mysteries of schizophrenia.

ACKNOWLEDGMENTS

The author thanks Vince Calhoun of the Olin Neuropsychiatric Research Center for sharing experimental fMRI data discussed in section IV; Abigail Garrity of the Trinity Neuroscience Program for sharing the results of her ICA analyses; Kara Carvalho of the Trinity Psychology Department for research assistance; and Matthew Broome and anonymous reviewers for helpful comments on earlier drafts. Research on this paper was supported by a Howard Hughes Medical Institute Faculty Development Grant to Trinity College.

NOTES

1. One question concerns the connection between *n*-back performance, WM, and schizophrenia. Although WM seems impaired in schizophrenia, performance measures like *n*-back are not considered to be diagnostic markers of the disease, because subjects without schizophrenia may also have trouble with the *n*-back test (and also show *n*-back hypofrontality). First-degree relatives of schizophrenia patients, in particular, score more like patients on various WM tasks, despite not having the disease. Accordingly, an important strand of schizophrenia research seeks correlations between brain observations like hypofrontality and diagnostic symptoms of the disorder. One example, once again, is the Menon et al. paper, which found that decreased DLPFC activation correlated with unusual thought

content. Other examples include Green (1996), Green and Nuechterlein (1999), Norman et al. (1999), Addington and Addington (1999), Dickerson et al. (1999), and Beck and Rector (2005).

2. Neuroimaging researchers are aware of these assumptions, and have proposed several alternative methods; I consider a few of these herein. The assumptions of “pure insertion” discussed herein is not belief but rather a methodological exclusion implicit in the logic of the standard method

3. Comparing contrasts is exactly what an informed neuroimaging research paper does, of course. But there are different strategies for combining the results of multiple subtractions. For example, Price and Friston’s “cognitive conjunction” proposes that functions be probed in contrasts using different baseline control conditions (Price and Friston 1997). What is constant in the multiple contrasts, then, is essential to the function under study. This method excludes regions that may support a function if either baseline shares resources with the task of interest. In effect, cognitive conjunction identifies the intersection of multiple contrasts. However, if brain resources are distributed, then different contrasts reveal different components of a network of multitasking regions, the union of the contrastive results. In short, “cognitive disjunction” is a method for discovering distributed neural processing.

4. Just as Brainmap sorts experiments by behavioral domain, it also sorts by experimental paradigm, of which *n*-back experiments are one example. This analysis distinguished 51 experimental paradigms, and compared them with respect to cortical areas of activation. Not surprisingly, the outcomes for *n*-back resemble those for WM overall: the database contains 88 *n*-back experiments with 652 identified loci of activation. Only visual pursuit experiments are more numerous. Accordingly, *n*-back dominates the scoreboard for all the DLPFC areas. But when the playing field is leveled, and all paradigms assigned an equal weight, *n*-back falls to tenth place or worse, considering DLPFC overall or considering any of the four DLPFC areas individually. Its near neighbors are all attention demanding tasks (like Stroop or the Wisconsin Card Sorting task), but no special task emerges as the overwhelming driver of DLPFC activity. In fact, twenty of the fifty-one paradigms involve activity in all four of the DLPFC regions. And forty-four paradigms (86%) involve at least one DLPFC component. That is, the database warrants the conclusion that DLPFC is involved in forty-four behavioral paradigms. Likewise, *n*-back cognition is supported by most cortical areas, not just DLPFC. Fully 79% of all Brodmann areas are reported to be active in one or more *n*-back experiments.

REFERENCES

- Addington, J., and D. Addington. 1999. Neurocognitive and social functioning in schizophrenia. *Schizophrenia Bulletin* 25:173–82.
- Andreasen, N. C., P. Nopoulos, D. S. O'Leary, D. D. Miller, T. Wassink, and M. Flaum. 1999. Defining the phenotype of schizophrenia: Cognitive dysmetria and its neural mechanisms. *Biology Psychiatry* 46:908–20.
- Barch, D. 2005. Cognitive neuroscience of schizophrenia. *Annual Review of Clinical Psychology* 1:321–53.
- Barch, D. M., J. G. Csernansky, T. Conturo, and A. Z. Snyder. 2002. Working and long-term memory deficits in schizophrenia: Is there a common prefrontal mechanism? *Journal of Abnormal Psychology* 111:478–94.
- Beck, A., and N. Rector. 2005. Cognitive approaches to schizophrenia: Theory and therapy. *Annual Review of Clinical Psychology* 1:577–606.
- Braver, T., D. Barch, and J. Cohen. 1999. Cognition and control in schizophrenia: A computational model of dopamine and prefrontal function. *Biological Psychiatry* 46:312–28.
- Calhoun, V. D., T. Adali, G. D. Pearlson, P. C. van Zijl, and J. J. Pekar. 2002. Independent component analysis of fMRI data in the complex domain. *Magnetic Resonance in Medicine*, 48:180–92.
- Callicott, J. H. 2003. An expanded role for functional neuroimaging in schizophrenia. *Current Opinions in Neurobiology*, 13:256–60.
- Callicott, J. H., A. Bertolino, V. S. Mattay, F. J. Langheim, J. Duyn, R. Coppola, et al. 2000. Physiological dysfunction of the dorsolateral prefrontal cortex in schizophrenia revisited. *Cerebral Cortex* 10:1078–192.
- Callicott, J. H., V. S. Mattay, A. Bertolino, K. Finn, R. Coppola, J. A. Frank, et al. 1999. Physiological characteristics of capacity constraints in working memory as revealed by functional MRI. *Cerebral Cortex* 9:20–6.
- Callicott, J. H., V. S. Mattay, B. A. Verchinski, S. Marenco, M. F. Egan, and D. R. Weinberger. 2003. Complexity of prefrontal cortical dysfunction in schizophrenia: More than up or down. *American Journal of Psychiatry* 160:2209–15.
- Callicott, J. H., N. F. Ramsey, K. Tallent, A. Bertolino, M. B. Knable, R. Coppola, T. Goldbert, P. van Gelderen, V. S. Mattay, J. A. Frank, C. T. Moonen, and D. R. Weinberger. 1998. Functional magnetic resonance imaging brain mapping in psychiatry: Methodological issues illustrated in a study of working memory in schizophrenia. *Neuropsychopharmacology* 18:186–96.
- Carter, C. S., W. Perlstein, R. Ganguli, J. Brar, M. Mintun, and J. D. Cohen. 1998. Functional hypofrontality and working memory dysfunction in schizophrenia. *American Journal of Psychiatry* 155:1285–7.
- Crow, T. J. 1980. Molecular pathology of schizophrenia: more than one disease process? *British Medical Journal* 280:6207:66–8.
- Dickerson, F., J. J. Boronow, N. Ringel, and F. Parente. 1999. Social functioning and neurocognitive deficits in outpatients with schizophrenia: A 2-year follow-up. *Schizophrenia Research* 37:13–20.
- Fodor, J. 1983. *Modularity of mind*. Cambridge: MIT Press.
- Fransson, P. 2005. Spontaneous low-frequency BOLD signal fluctuations: An fMRI investigation of the resting-state default mode of brain function hypothesis. *Human Brain Mapping* 26:15–29.
- Friston, K., C. Price, P. Fletcher, R. Moore, R. Frackowiak, and R. Dolan. 1996. The trouble with cognitive subtraction. *Neuroimage* 4:97–104.
- Friston, K. J. 1999. Schizophrenia and the disconnection hypothesis. *Acta Psychiatrica Scandinavica Supplement* 395:68–79.
- Frith, C. 1992. *Cognitive neuropsychology of schizophrenia*. Hove, UK: Lawrence Erlbaum.
- Garrity, A., G. D. Pearlson, K. McKiernan, D. Lloyd, K. A. Kiehl, and V. D. Calhoun. 2007. Aberrant 'default mode' functional connectivity in schizophrenia. *American Journal of Psychiatry* 164:450–7.
- Green, M. F. 1996. What are the functional consequences of neurocognitive deficits in schizophrenia? *American Journal of Psychiatry* 153:321–30.
- Green, M. F., and K. H. Nuechterlein. 1999. Should schizophrenia be treated as a neurocognitive disorder? *Schizophrenia Bulletin* 25:309–19.
- Greicius, M. D., B. Krasnow, A. L. Reiss, and V. Menon. 2003. Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings of the National Academy of Sciences of the United States of America* 100:253–8.
- Greicius, M. D., and V. Menon. 2004. Default-mode activity during a passive sensory task: Uncoupled from deactivation but impacting activation. *Journal of Cognitive Neuroscience* 16:1484–92.
- Gusnard, D. A., and M. E. Raichle. 2001. Searching for a baseline: Functional imaging and the resting human brain. *National Review of Neuroscience* 2:685–94.
- Honey, G. D., T. E. Bullmore, and T. Sharma. 2002. De-coupling of cognitive performance and cerebral functional response during working memory in schizophrenia. *Schizophrenia Research* 53:45–56.
- Honey, G. D., P. C. Fletcher, and T. E. Bullmore. 2002. Functional brain mapping of psychopathology. *Journal of Neurology, Neurosurgery & Psychiatry* 72:432–9.

- Jansma, J. M., N. F. Ramsey, N. J. van der Wee, and R. S. Kahn. 2004. Working memory capacity in schizophrenia: A parametric fMRI study. *Schizophrenia Research* 68:159–71.
- Jessard, P., P. Matthews, and S. Smith. 2001. *Functional MRI: An introduction to the methods*. Oxford: Oxford University Press.
- Kraepelin, E. 1896. Dementia Praecox. In *The clinical roots of the schizophrenic concept*, ed. J. Cutting, and M. Shepherd. Cambridge: Cambridge University Press.
- Kuperberg, G., and S. Heckers. 2000. Schizophrenia and cognitive function. *Current Opinion in Neurobiology* 10:205–10.
- Liddle, P. F. 1987. The symptoms of chronic schizophrenia. A re-examination of the positive-negative dichotomy. *British Journal of Psychiatry* 151:145–51.
- Lloyd, D. 2002. Functional MRI and the study of human consciousness. *Journal of Cognitive Neuroscience* 14:818–31.
- . 2003. *Radiant cool: A novel theory of consciousness*. Cambridge, MA: MIT Press.
- . (2007). Civil schizophrenia. In *Distributed cognition and the will: Individual Volition and Social Context*, ed. D. Ross, D. Spurrett, H. Kincaid, & G. Lynn Stephens. Cambridge, MA: MIT Press.
- McClelland, J., and D. Rumelhart. 1986. *Parallel distributed processing: Explorations in parallel distributed processing*. Cambridge, MA: MIT Press.
- McKeown, M. J., S. Makeig, G. G. Brown, T. P. Jung, S. S. Kindermann, A. J. Bell, and T. J. Sejnowski. 1998. Analysis of fMRI data by blind separation into independent spatial components. *Human Brain Mapping* 6:160–88.
- McKiernan, K. A., J. N. Kaufman, J. Kucera-Thompson, and J. R. Binder. 2003. A parametric manipulation of factors affecting task-induced deactivation in functional neuroimaging. *Journal of Cognitive Neuroscience* 15:394–408.
- Mendrek, A., K. R. Laurens, K. A. Kiehl, E. T. Ngan, E. Stip, and P. F. Liddle. 2004. Changes in distributed neural circuitry function in patients with first-episode schizophrenia. *British Journal of Psychiatry* 185:205–14.
- Menon, V., R. Anagnoson, D. Mathalon, G. Glover, and A. Pfefferbaum. 2001. Functional neuroanatomy of auditory working memory in schizophrenia: Relation to positive and negative symptoms. *Neuroimage* 13:433–46.
- Meyer-Lindenberg, A., R. S. Miletich, P. D. Kohn, G. Esposito, R. E. Carson, M. Quarantelli, D. R. Weinberger, and K. F. Berman. 2002. Reduced prefrontal activity predicts exaggerated striatal dopaminergic function in schizophrenia. *Nature Neuroscience* 5:267–71.
- Niznikiewicz, M., M. Kubicki, and M. Shenton. 2003. Recent structural and functional imaging findings in schizophrenia. *Current Opinion in Psychiatry* 16:123–147.
- Norman, R. M., A. K. Malla, L. Cortese, S. Cheng, K. Diaz, E. McIntosh, T. S. McLean, A. Rickwood, and L. P. Voruganti. 1999. Symptoms and cognition as predictors of community functioning: A prospective analysis. *American Journal of Psychiatry* 156:400–5.
- Perlstein, W. M., C. S. Carter, D. C. Noll, and J. D. Cohen. 2001. Relation of prefrontal cortex dysfunction to working memory and symptoms in schizophrenia. *American Journal of Psychiatry* 158:1105–13.
- Price, C. J., and K. J. Friston. 1997. Cognitive conjunction: A new approach to brain activation experiments. *Neuroimage* 5:261–70.
- Raichle, M. E., A. M. MacLeod, A. Z. Snyder, W. J. Powers, D. A. Gusnard, and G. L. Shulman. 2001. A default mode of brain function. *Proceedings of the National Academy of Sciences of the United States of America* 98:676–82.
- Sabri, O., A. Owega, M. Schreckenberger, L. Sturz, B. Fimm, P. Kunert, P. T. Meyer, d. Sander, and J. Klingerhöfer. 2003. A truly simultaneous combination of functional transcranial Doppler sonography and H(2)(15)O PET adds fundamental new information on differences in cognitive activation between schizophrenics and healthy control subjects. *Journal of Nuclear Medicine* 44:671–81.
- Smolensky, P. 1988. On the proper treatment of connectionism. *Behavioral and Brain Sciences* 11: 1–23.
- Walter, H., A. P. Wunderlich, M. Blankenhorn, S. Schafer, R. Tomczak, M. Spitzer, and G. Grön. 2003. No hypofrontality, but absence of prefrontal lateralization comparing verbal and spatial working memory in schizophrenia. *Schizophrenia Research* 61:175–84.
- Wimsatt, W. C. 1976. Reductive explanation: A functional account, in *PSA-1974, Boston studies in the philosophy of science*, vol. 32, ed. R. S. Cohen, C. A. Hooker, A. C. Michalos, and J. van Evra, 671–710. Dordrecht: D. Reidel.
- Wykes, T., M. Brammer, J. Mellers, P. Bray, C. Reeder, C. Williams, and J. Corner. 2002. Effects on the brain of a psychological treatment: Cognitive remediation therapy: Functional magnetic resonance imaging in schizophrenia. *British Journal of Psychiatry* 181:144–52.