



Using Dual Regression to Investigate Network Shape and Amplitude in Functional Connectivity Analyses

Lisa D. Nickerson^{1,2*}, Stephen M. Smith³, Döst Öngür^{2,4} and Christian F. Beckmann^{3,5,6}

¹ Applied Neuroimaging Statistics Lab, McLean Hospital, Belmont, MA, USA, ² Department of Psychiatry, Harvard Medical School, Harvard University, Boston, MA, USA, ³ Nuffield Department of Clinical Neurosciences, Oxford University Centre for Functional MRI of the Brain, John Radcliffe Hospital, University of Oxford, Oxford, UK, ⁴ Schizophrenia and Bipolar Disorder Research Program, McLean Hospital, Belmont, MA, USA, ⁵ Department of Cognitive Neuroscience, Radboud University Medical Centre Nijmegen, Netherlands, ⁶ Center for Cognitive Neuroimaging, Donders Institute for Brain, Cognition and Behavior, Radboud University, Nijmegen, Netherlands

OPEN ACCESS

Edited by:

Bertrand Thirion,
Institut National de Recherche en
Informatique et Automatique, France

Reviewed by:

Xin Di,
New Jersey Institute of Technology,
USA
Chao-Gan Yan,
Chinese Academy of Sciences, China

*Correspondence:

Lisa D. Nickerson
lisa_nickerson@hms.harvard.edu

Specialty section:

This article was submitted to
Brain Imaging Methods,
a section of the journal
Frontiers in Neuroscience

Received: 19 September 2016

Accepted: 23 February 2017

Published: 13 March 2017

Citation:

Nickerson LD, Smith SM, Öngür D
and Beckmann CF (2017) Using Dual
Regression to Investigate Network
Shape and Amplitude in Functional
Connectivity Analyses.
Front. Neurosci. 11:115.
doi: 10.3389/fnins.2017.00115

Independent Component Analysis (ICA) is one of the most popular techniques for the analysis of resting state fMRI data because it has several advantageous properties when compared with other techniques. Most notably, in contrast to a conventional seed-based correlation analysis, it is model-free and multivariate, thus switching the focus from evaluating the functional connectivity of single brain regions identified a priori to evaluating brain connectivity in terms of all brain resting state networks (RSNs) that simultaneously engage in oscillatory activity. Furthermore, typical seed-based analysis characterizes RSNs in terms of spatially distributed patterns of correlation (typically by means of simple Pearson's coefficients) and thereby confounds together amplitude information of oscillatory activity and noise. ICA and other regression techniques, on the other hand, retain magnitude information and therefore can be sensitive to both changes in the spatially distributed nature of correlations (differences in the spatial pattern or "shape") as well as the amplitude of the network activity. Furthermore, motion can mimic amplitude effects so it is crucial to use a technique that retains such information to ensure that connectivity differences are accurately localized. In this work, we investigate the dual regression approach that is frequently applied with group ICA to assess group differences in resting state functional connectivity of brain networks. We show how ignoring amplitude effects and how excessive motion corrupts connectivity maps and results in spurious connectivity differences. We also show how to implement the dual regression to retain amplitude information and how to use dual regression outputs to identify potential motion effects. Two key findings are that using a technique that retains magnitude information, e.g., dual regression, and using strict motion criteria are crucial for controlling both network amplitude and motion-related amplitude effects, respectively, in resting state connectivity analyses. We illustrate these concepts using realistic simulated resting state fMRI data and *in vivo* data acquired in healthy subjects and patients with bipolar disorder and schizophrenia.

Keywords: ICA, dual regression, resting state networks, functional connectivity, amplitude

INTRODUCTION

Spatial independent component analysis (ICA) has emerged as a popular tool for investigating brain network functional connectivity measured using functional magnetic resonance imaging (fMRI) (McKeown et al., 1998; Cordes et al., 1999; Calhoun et al., 2001; Kiviniemi et al., 2003; Greicius et al., 2004; van de Ven et al., 2004; Beckmann et al., 2005). This technique aims to identify underlying hidden spatio-temporal processes that co-occur to give rise to the measured blood oxygenation level dependent (BOLD) fMRI signals from the brain. These spatio-temporal processes can correspond to spontaneous fluctuations in resting state networks (RSNs) during wakeful rest, to task-locked alterations in brain networks during task performance, to scanner artifacts that occur during the imaging period, or to physiological noise. The appeal of the spatial ICA is that no time-series model of brain activity is required for the analysis—in that sense, it is model-free—and as a consequence, the spatial ICA has the potential to identify spatio-temporal signal sources that are not well-characterized or well-understood. These properties are well-suited to the analysis of fMRI measurements using a resting state paradigm in which there is no task *per se* (for example, resting with eyes open or closed, during sleep, or during pharmacologic manipulations) to investigate the resting state functional connectivity (RSFC) of brain networks.

The standard approach for investigating RSFC of brain networks implements a multi-subject analysis in which a group-average spatial ICA, or group ICA (GICA), is done on the entire set of fMRI data concatenated across all subjects (Calhoun et al., 2001; Beckmann and Smith, 2004). This technique will identify the collection of networks that are common to the entire set of subjects. Because the output spatial maps are common to all subjects, further processing is necessary to compare functional connectivity of any given network between groups of subjects or conditions. For this purpose, the subject-specific network maps corresponding to each group ICA map must be identified in each subject, similar to a contrast map from a task fMRI analysis. These subject-specific network maps capture between-subject variability in the “shape,” or spatial pattern, of the network and can be used in a higher-level general linear model analysis to investigate group differences in functional connectivity.

In addition to the shape of an RSN, the amplitude of an RSN (e.g., the magnitude of the BOLD activity in the RSN) has been shown to convey important information regarding resting state activity. For example, positron emission tomography (PET) task activation studies have long reported decreases in cerebral blood flow (CBF) in a certain collection of brain regions when comparing task to passive rest conditions (Raichle et al., 2001). This collection of brain regions also has high cerebral metabolic rate of oxygen consumption (CMRO₂) and CBF at rest and, even early on, the magnitude of the decreases in CBF during brain activation were noted to be likely related to task difficulty, suggesting that magnitude of deactivation carries important information (Shulman et al., 1997). This collection of brain regions is now well-known as the default mode network (DMN), the physiology of which has been studied extensively using PET (Gusnard and Raichle, 2001; Raichle and Snyder,

2007). The function of the DMN has also been studied extensively with BOLD fMRI during wakeful rest (Greicius et al., 2003) using the seed-based connectivity analysis approach first presented by Biswal et al. (1995), and using ICA-based approaches (Beckmann et al., 2005; Smith et al., 2009) and other methods (Andrews-Hanna et al., 2010). BOLD fMRI depends on CBF, CMRO₂, and cerebral blood volume, thus results from PET studies would also predict effects on BOLD signal amplitudes in the DMN during task performance.

Importantly, now several studies have shown that the amplitude of BOLD activity within the DMN, measured as the standard deviation of the BOLD signal timecourse, is related to task-load during brain activation (McKiernan et al., 2003; Singh and Fawcett, 2008) and is sensitive to different resting state conditions (eyes open with and without fixation vs. eyes closed; Yan et al., 2009 using the power spectrum). Other brain regions also show similar amplitude-related effects. For example, Bianciardi et al. (2009) and Jao et al. (2013) have shown that the amplitude of resting state BOLD signal oscillations in the visual cortex is smaller with an eyes-open fixation resting condition relative to resting with eyes closed. Jao et al. (2013) also reported differences in amplitude of BOLD signals from nodes of the DMN during these two states. Resting fMRI amplitudes can be affected by pharmacologic stimulation (Kiviniemi et al., 2005), by disease (Zang et al., 2007; Vargas et al., 2013) and can also reliably predict task fMRI responses in activated brain regions, for example, during motor and breath hold tasks (Kannurpatti et al., 2012). In addition, Kannurpatti and Biswal (2008) proposed the use of resting state fluctuation amplitudes to scale task-induced BOLD responses for calibrated BOLD, as a proxy for breath holding or 5% CO₂ inhalation. Thus, the amplitude of the resting fMRI signal likely also reflects an important aspect of RSFC and preserving such amplitude information within any inferential procedure is important to more fully characterize functional connectivity within and between subjects.

The seed-based correlation approach (SBCA) for investigating RSFC originally relied on simple time series correlation (to detect similarities, e.g., synchronous fluctuations, between the timecourse from a seed-region and the timecourses from different parts of the brain). However, it is now common practice to implement SBCA as a voxel-wise multiple regression analysis that is capable of handling confounders such as physiological and subject motion. SBCA will reveal the full connectivity profile of a given seed region, showing all of the brain regions that are functionally connected to the seed in the connectivity map. As a result, it is challenging to differentiate the distinct networks with which a seed region may be functionally connected. One would have to visually “parse” in a purely qualitative way those blobs out into separate networks for interpretation in the context of which networks were connected with the seed. In addition, network amplitude effects will not be accurately reflected using such an approach, with (voxel-wise estimated) correlation coefficients not reflecting amplitude at all (all other things being equal), and regression coefficients in this framework being estimated relative to the seed-timecourse amplitude itself. To disentangle functional connectivity and amplitude of multiple networks, in e.g., the repertoire of brain networks known to be “active” at rest in

the human brain (Beckmann et al., 2005; Smith et al., 2009), multivariate regression approaches that analyse the behavior of all networks simultaneously, for example, dual regression [DR; Beckmann et al., 2009; Filippini et al., 2009; implemented in FMRIB Software Library (FSL)], are most suitable. In this case, a full set of brain network maps derived from a GICA are used with DR to derive the subject-specific spatial maps and timecourses corresponding to each GICA spatial map, which maintains sensible control over the contribution of amplitude effects while also providing information regarding network shape and disentangling the effects of multiple networks. Crucially, motion effects can also be a source of amplitude effects by causing abrupt and potentially large spikes or discontinuities in FMRI timeseries, introducing variance into the timecourse that can mimic an amplitude effect. Many studies have shown that motion is a major source of variability in functional connectivity studies that can lead to erroneous results when comparing groups of subjects (Power et al., 2012; Satterthwaite et al., 2012; Van Dijk et al., 2012). Thus, it is crucial to investigate functional connectivity within a framework that maintains control over any amplitude effects that may be present in the data, whether real network effects or motion-related.

In this report, we investigate the use of FSL's group ICA with dual regression (GICA-DR) approach for assessing group/condition differences in functional connectivity. In contrast to group SBCA, there is no need for specifying either subject-specific seed timecourses from regions of interest or confound regressors to capture subject-specific noise effects in the FMRI data. Instead, the GICA identifies the spatiotemporal signals associated with RSNs (e.g., networks of interest), artifacts, and noise that are inherent in the FMRI data. It does so by decomposing the FMRI data in a purely data-driven way into a set of independent components (ICs), with each IC consisting of a spatial map (that is statistically independent from the other IC maps), and a corresponding timecourse. For resting state data, the IC timecourses from a GICA are generally not interpretable because it is common to do data reduction prior to the GICA. However, the spatial maps identified from GICA characterize RSNs that are common to all the subjects, providing a more comprehensive view of the functional hierarchy of the brain (Beckmann et al., 2005), while also identifying artifacts that are, in principle, being identified as components separate from RSNs. Although there are some methodological challenges to ICA (Cole et al., 2010) there is now a consensus that spatial ICA is a powerful analytical tool for investigating functional connectivity across multiple functional domains, such as sensory, motor, cognitive, and limbic systems, simultaneously. Once the group IC maps have been identified by GICA, DR is used to estimate the subject-specific network timecourses and the subject-specific spatial maps for each corresponding GICA map. The resulting subject-specific timecourses and spatial maps can be used to compare functional connectivity between groups/conditions. Some studies have discussed or compared SBCA and ICA-based approaches for assessing group differences in brain network functional connectivity (Cole et al., 2010; Van Dijk et al., 2010; Smith et al., 2014) and have investigated the test-retest reliability and other aspects of GICA-DR (Zuo et al., 2010; Allen et al., 2012). In this

study, we focus on how to use GICA-DR to investigate RSN shape while maintaining control over amplitude effects.

Group ICA with Dual Regression (GICA-DR)

The GICA identifies a set of IC maps that are common to the entire population (e.g., FMRI data from every subject is included in the analysis) and the dual regression is a mathematical tool that utilizes these IC maps as network templates to identify the corresponding functional connectivity maps in each subject. Although it is generally recommended to use the output spatial maps of a GICA of the dataset under investigation in the DR, dual regression is a generic method in the sense that it can be used to assess the functional connectivity of networks (or parcellations) that are identified using any strategy—not just a GICA of the FMRI data. For example, a set of template networks could be derived from a GICA of an independent set of subjects, from an atlas, or using functional localizers, although these templates may not contain information related to specific artifacts or sources of noise in the data under study, as would a GICA of the original data.

Once the template maps (on the basis of GICA or else) have been identified, the dual regression analysis proceeds in two stages (**Figure 1**). In the first stage of the dual regression (**Figure 1A**), each subjects' 4D FMRI dataset, Y , is reorganized into a 2D (N voxels \times T timepoints) data matrix, and the unthresholded full set of template maps are regressed into this data matrix. e.g., let the template maps be the estimated IC spatial maps from a group ICA, \hat{S} , reorganized into a 2D (N voxels \times M components) matrix, then these \hat{S} are the independent variables or predictors in a multivariate multiple linear regression:

$$Y = \hat{S}B_{TC} + E_1 \quad (1)$$

with:

$$\hat{B}_{TC} = \text{pinv}(\hat{S})Y \quad (2)$$

pinv denotes the matrix pseudoinverse, $Y \in \mathcal{R}^{N \times T}$ is the subject's FMRI data, $E_1 \in \mathcal{R}^{N \times T}$ is the matrix of errors, and $\hat{B}_{TC} \in \mathcal{R}^{M \times T}$ is the matrix of stage 1 timecourses, one for each IC map.¹ Each timecourse in \hat{B}_{TC} reflects the average timecourse computed over voxels in the corresponding IC spatial map after taking into account the contributions of the other components to each voxel's timecourse. The regression parameters, \hat{B}_{TC} , are sensitive to the scale of the maps in \hat{S} . However, for dual regression, \hat{S} are the same across all subjects.

In the second stage of the dual regression (**Figure 1B**), the network-specific timecourses from the stage 1 regression are used as predictors in a second multivariate multiple linear regression into the individual subject's FMRI data². This second regression identifies the subject-specific spatial maps, B_{SM} :

$$Y' = \hat{B}'_{TC}B_{SM} + E_2 \quad (3)$$

¹These output timecourses are *not* the same as the timecourses that are output from a GICA of resting state data using FSL MELODIC. The GICA timecourses from MELODIC are derived by inverting the initial PCA reduction done prior to the group ICA.

²All predictors are mean centered such that the mean of each predictor is zero, but the units are still on the original scale unless otherwise transformed.

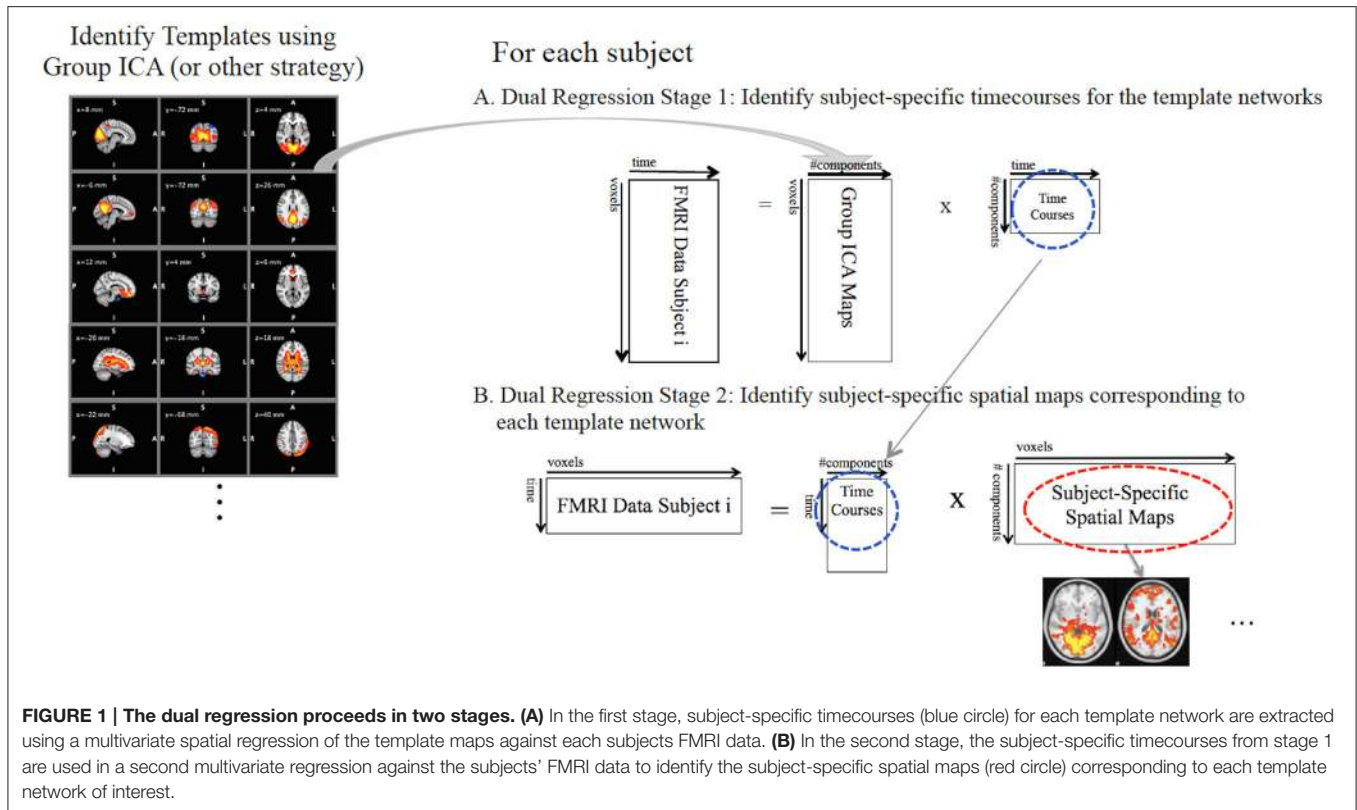


FIGURE 1 | The dual regression proceeds in two stages. (A) In the first stage, subject-specific timecourses (blue circle) for each template network are extracted using a multivariate spatial regression of the template maps against each subjects fMRI data. **(B)** In the second stage, the subject-specific timecourses from stage 1 are used in a second multivariate regression against the subjects' fMRI data to identify the subject-specific spatial maps (red circle) corresponding to each template network of interest.

Which gives:

$$\hat{B}_{SM} = Y * pinv(\hat{B}_{TC}) \tag{4}$$

The estimated $\hat{B}_{SM} \in \mathcal{R}^{N \times M}$ have a spatial map for each corresponding IC template map. These maps are the subject's functional connectivity maps for the networks in \hat{S} .

Generally, it is the case that RSN timecourses output from stage 1 of the dual regression (Equation 2) will have different amplitudes, measured as the standard deviation of the timecourse, by virtue of these timecourses being average timecourses (weighted sums) computed over the voxels in the RSN (after partialling out other network temporal effects). When these timecourses are regressed into the fMRI data in the second regression (Equation 3), the regression coefficients at each voxel (Equation 4) are estimated with respect to that subject's RSN amplitude or "scale". Comparing the voxel values in \hat{B}_{SM} (e.g., the regression coefficients) across subjects is confounded by this apparent difference in network amplitude across subjects. Note that this is not the case for the scale of the \hat{S} since these maps are used for dual regression in all subjects and thus any scaling in the spatial maps will be constant across subjects.

Standardizing the stage 1 timecourses by dividing by the standard deviation of the timecourse prior to estimating Equation 3 obviates the network amplitude (or scale) problem, with regression coefficients that are interpreted as a change in Y per unit standard deviation in X. Subsequently, the voxel-wise regression coefficient for the group analysis will be evaluated

in the same units of standard deviation in X across subjects. This is applicable for a network-wide amplitude effect, in which the amplitude of the whole network is affected (e.g., different across subjects). However, another potential problem arises when a region, or node, within a network has an amplitude that is different from the other nodes in the network. In the case of such a "within-network" amplitude effect, a different scale problem arises for the voxels within the network itself—with their regression coefficients being estimated on different scales with respect to each other. Standardizing to change in Y per unit standard deviation in X places all voxels on an equivalent scale, while also placing all subjects on an equivalent scale. As an aside, stage 1 timecourse normalization differs from normalizing the individual voxel's fMRI timecourses prior to the dual regression. In the latter case, even after normalization of the individual voxel-wise timecourses, the subject-specific timecourses from the first stage of the dual regression are still a weighted sum of the data time courses—even if all original time series are zero mean and unit standard deviation, their weighted sum does not remain normalized in either magnitude or variance upon dual regression.

As a consequence, to investigate amplitude effects—or to ensure validity in the case of unknown amplitude effects—the timecourses *must* be normalized by the amplitude prior to estimation of the \hat{B}_{SM} using Equation 3, which will give "semi-standardized" regression coefficients, \hat{B}_{SM}^* :

$$\hat{B}_{SM}^* = Y * pinv(A^{-1} * \hat{B}_{TC}) \tag{5}$$

With $A = \text{diag}\{a_{\hat{B}_{TC}}\}$, e.g., a diagonal matrix with the standard deviations, or amplitudes a_i , of each stage 1 timecourse on the diagonal. In FSL dual regression parlance, normalization of these timecourses is referred to as design normalization.

Importantly, \hat{B}_{SM}^* are sensitive to both network amplitude and shape effects and moreover, are the *only* accurate measure for assessing group differences in functional connectivity in the presence of within-network amplitude effects. See Appendix 1 (Supplementary Material) for a more detailed explanation of this using a simple two-voxel network example.

There are two issues to keep in mind when using \hat{B}_{SM}^* . The first is that there is a price for including the extra amplitude information. The spatial maps fed forward to the group analysis will reflect variations in shape, *and* at the voxel level, the between-subject variations in amplitude. The extra random effects variance from the amplitude will make it more difficult to detect an effect since both effects contribute to the random-effects between-subject variation. The second is that the \hat{B}_{SM}^* are sensitive to subject motion that manifests as large-amplitude spikes in the timecourses and, as with any RSFC analysis, careful attention must be given to any motion effects in the data. That being said, at least this type of amplitude effect will be accurately localized (which cannot be said for the impact of motion on \hat{B}_{SM} or when using SBICA). The fact that motion can mimic amplitude effects likely underlies the overall greater sensitivity of functional connectivity analyses to motion and may be (at least partly) the cause of very unpredictable effects that motion has been shown to have on functional connectivity outcomes. (Satterthwaite et al., 2012; Van Dijk et al., 2012; Power et al., 2014, 2015).

A summary of key points is that:

- 1) \hat{B}_{SM} , or the raw regression coefficients of Equation (4), do not contain information related to the network-wide amplitude and thus can't be used to investigate network amplitude.
- 2) \hat{B}_{SM} will be wrong if there is a within-network amplitude difference (e.g., when a node in a network has higher or lower amplitude than other nodes in the network).
- 3) \hat{B}_{SM} also will be wrong in the presence of motion effects that mimic amplitude effects by increasing variance of network timecourses (e.g., motion effects will not be localized only to truly affected voxels).
- 4) Amplitude information is a source of extra random effects variance.
- 5) With design normalization, *any* amplitude effect, whether a real network-related effect or one that arises due to motion effects, will be accurately localized in the dual regression stage 2 maps, \hat{B}_{SM}^* .

In this report, we illustrate these concepts using simulated and real data. In the simulations, we manipulate ground truth timecourses and spatial maps derived from real FMRI data to create known group amplitude and shape differences in simulated data. The resulting simulated data are then analyzed using DR to derive \hat{B}_{SM} and \hat{B}_{SM}^* to show how amplitude and shape differences are localized in each. Amplitude effects also arise through impulsive subject movements, thus an analysis

of *in vivo* data is used to illustrate the consequences and best practices in using dual regression in the presence of motion.

MATERIALS AND METHODS

Simulations: Overview

Eight spatial maps corresponding to well-established RSNs (previously reported by Beckmann et al. (2005) and available at <http://www.fmrib.ox.ac.uk/analysis/roysoc8/>) were used as ground truth spatial patterns, and timecourses derived from real BOLD FMRI resting state data acquired in 36 healthy subjects were used as ground truth timecourses to construct ground truth resting state FMRI data for 36 simulated subjects that had 8 “active” RSNs (plus noise). Manipulation of derived timecourses and/or ground truth spatial maps in half of the simulated subjects was done to create group differences in functional connectivity.

Ground Truth Spatiotemporal Processes

Resting state BOLD FMRI data collected in 36 healthy participants at the University of Oxford Centre for Clinical Magnetic Resonance Research using a 3T Siemens Trio scanner (Filippini et al., 2009) were used for the simulations. The University of Oxford Centre for Clinical Magnetic Resonance Research Ethics Committee approved the original study and all participants gave informed written consent to participate in the original study. Scan parameters were: Field of view = 224 mm, $3 \times 3 \times 3.5$ mm resolution, TE/TR/FA = 28/2000 ms/89°, total scan time = 6:04 min. The first four frames were discarded for magnetic field equilibration. FMRI data were pre-processed using FSL (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl, Smith et al., 2004) to perform head motion correction using FLIRT (Jenkinson et al., 2002); non-brain removal using BET (Smith, 2002); spatial smoothing by a Gaussian kernel (FWHM 6 mm); grand-mean intensity normalization of the entire 4D dataset by a single multiplicative factor; and highpass temporal filtering by subtraction of a Gaussian-weighted least-squares straight line fitting with $\sigma = 150$ s (Niazy et al., 2011). Finally, registration of each FMRI dataset to the corresponding high-resolution structural scan was carried out using FLIRT. Registration of the high-resolution structural images to MNI152 standard space was achieved using FLIRT, with further refinement using FNIRT non-linear registration (Smith et al., 2004). All subjects' four-dimensional timeseries data were transformed into standard space at $2 \times 2 \times 2$ mm³ resolution using the registration transformation matrices.

Ground truth spatial maps: The spatial network maps from Beckmann et al. (2005) were used as the ground truth spatial maps, hereafter referred to as the roysoc8 maps. These eight maps correspond to (A) medial visual network (MVN), (B) lateral occipital, (C) auditory, (D) sensorimotor, (E) default mode network (DMN), (F) executive control network (ECN), and (G–H) right- and left-lateralized fronto-parietal networks (RFPN and LFPN, respectively).

Ground truth network timecourses: The roysoc8 maps were regressed against one of the real pre-processed FMRI datasets to derive 8 corresponding “ground truth” RSN timecourses per subject. The voxel-wise standard deviations estimated from the residuals of the regression (std_{res}) for each subject were also calculated.

Simulated Resting State FMRI Data

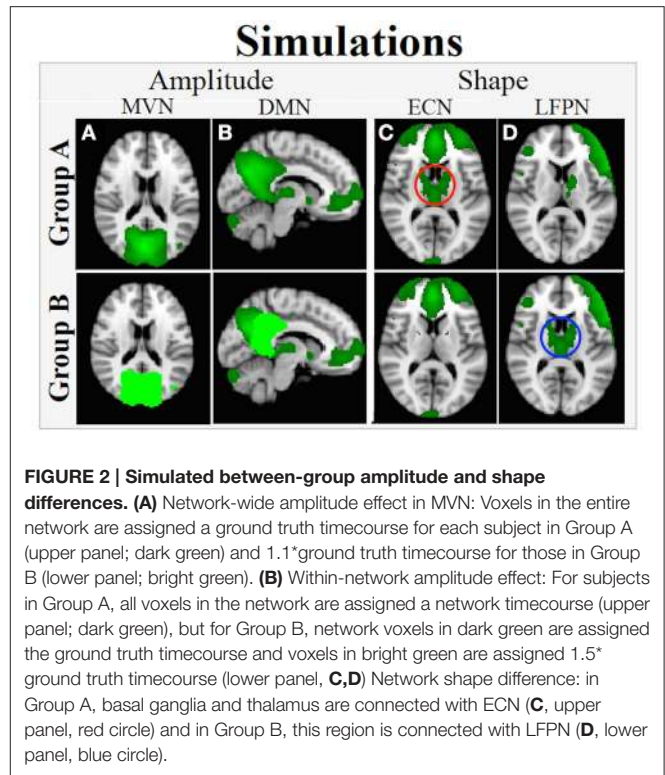
Simulated single subject resting state FMRI data were created by computing the product of the roysoc8 maps and one set of ground truth RSN timecourses, and adding (at each time point) voxel-wise random Gaussian noise distributed as $\mathcal{N}(0, std_{res}^2 I)$ plus the voxel-wise means of the real FMRI timeseries. Simulated data created from half of the real FMRI data (18 subjects) were designated as Group A data, the other half (18 subjects) as Group B.

For Group B data, manipulation of the timecourses and/or associated spatial maps prior to multiplication of timecourses with the spatial maps was done to simulate between-subject variability in RSN amplitude and shape as follows³:

- Network-Wide Amplitude Difference (**Figure 2A**): For each subject, the ground truth timecourse for the MVN was multiplied by 1.1 to simulate a 10% greater network-wide amplitude in Group B relative to Group A.
- Within-Network Amplitude Difference (**Figure 2B**): For each subject, the ground truth timecourses from voxels in posterior cingulate cortex (PCC) region of the DMN were multiplied by 1.5 to simulate a 50% greater amplitude of the PCC relative to other regions in the DMN. This enforces a within-network group amplitude difference, e.g., a within-network difference around the mean network amplitude.
- Shape (**Figures 2C,D**): Voxels in selected regions of the ECN (basal ganglia and thalamus) were “connected” to the LFPN in Group B subjects by giving that subset of voxels the timecourse for LFPN instead of the timecourse for the ECN. The net effect is that the basal ganglia and thalamus are connected to ECN in Group A (**Figure 2C**) and to LFPN in Group B (**Figure 2D**), resulting in a group difference in shape in both ECN and LFPN. Note that the timecourses used here were not normalized to unit standard deviation and thus there will also be a small amplitude difference in the overall network timecourses.

Note that each amplitude and shape effect impacts a single map when assessing differences in functional connectivity between Group A and Group B, with a network-wide amplitude difference in MVN, a within-network difference in DMN, and shape differences in ECN and LFPN. The other maps did not have any simulated group differences. All effects were included in a single simulated Group B dataset to also show that dual regression can separate out multiple different network amplitude and shape effects that may be present in any given FMRI dataset.

³Our specific choices of RSNs, amplitude, and shape differences were based simply on demonstrating the ability of different statistics described in Section Introduction to accurately reflect amplitude and shape differences.



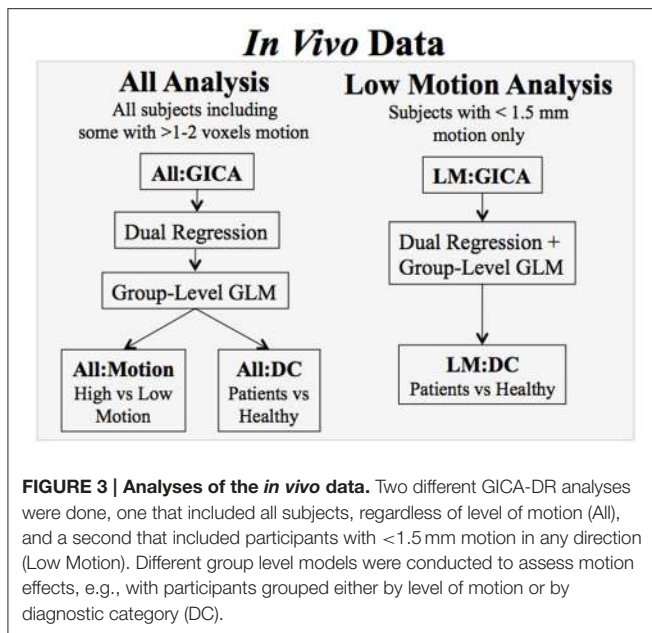
Group Comparisons of RSN Functional Connectivity using Dual Regression

Dual regression (<http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/DualRegression>) was done as follows. At stage 1, the roysoc8 spatial maps were regressed against the simulated FMRI datasets to obtain the subject-specific timecourses (SSTC). The SSTC outputs are text files, one per subject, containing a matrix of values that have one column per spatial map and one row per time point.

At stage 2 of the dual regression, the SSTC (for all networks) were regressed against each of the simulated datasets to obtain the subject-specific spatial maps (SSSM). Each stage 2 output is a 4D spatial maps data file per subject ($N = 8$ components in this case). Dual regression was run without and with normalization of the SSTC (each timecourse normalized to unit standard deviation), resulting in SSSM that were \hat{B}_{SM} or \hat{B}_{SM}^* , respectively.

Once the \hat{B}_{SM} and \hat{B}_{SM}^* were computed, estimation and inference on the two different sets of statistic maps was done using FSL Randomise with a two-group unpaired t -test to assess differences between Group A and B in each network and for each statistic (\hat{B}_{SM} and \hat{B}_{SM}^*), with 5,000 permutations and cluster-mass based thresholding (with cluster forming threshold $Z = 2.3$), to achieve a corrected family-wise error rate, $p < 0.05$.

SSTCs are evaluated for accuracy by comparing the SSTC with the ground truth timecourses. The correlation between each ground truth timecourse and the corresponding dual regression timecourse for every network should be close to 1. We also compared the amplitudes (standard deviations) of the ground truth timecourses with the amplitudes of the SSTC for the MVN (the component with the 10% amplitude effect) and the other



networks. For MVN, the amplitudes of SSTC and ground truth timecourses should be equal in Group A, and for Group B, the amplitudes of the SSTC should be equal to 1.1 times the amplitudes of the ground truth timecourse for each subject. There will also be an amplitude effect in the DMN because the SSTC is a weighted sum of all voxel's timecourses in the DMN and some voxels have a larger amplitude than others, although this is not easily estimated as is the network-wide effect. For the other five networks, SSTCs should be the same as the ground truth timecourses.

The results of the group-level analysis of the \hat{B}_{SM} and \hat{B}_{SM}^* spatial maps should show group differences in MVN, DMN, ECN, and LFPN that reflect the voxel-wise differences in amplitude/synchrony as shown in **Figure 2**, with no differences in functional connectivity in the other four networks.

In Vivo Resting State fMRI: Overview

An *in vivo* resting state fMRI dataset from 13 healthy controls (HC), 14 individuals with bipolar disorder (BD), and 12 individuals with schizophrenia (SZ), for a total of 39 subjects, is used to illustrate how motion can distort functional connectivity results. Analyses to assess differences in functional connectivity between the three diagnostic categories was done two ways: (1) using all of the data ($N = 39$), which included several subjects in each group with motion in excess of 1–2 voxels, and (2) using a subset of the data that included only subjects with less than 1.5 mm motion ($N = 25$). These two analyses demonstrate how excess motion can impact GICA maps and group differences in functional connectivity and how using strict motion criteria can improve findings even at the expense of sample size. In addition, we ran a group analysis that compared “high” vs. “low” motion subjects (with these two groups matched for the same number of participants from each diagnostic category) to show how motion can cause spurious group differences in functional connectivity.

As an aside, we are not interested in the neurobiological differences between individuals with different diagnoses *per se* and will not be discussing our findings in this context; details of the participants and neurobiological research findings have been reported previously (Öngür et al., 2010; Chai et al., 2011).

Data Acquisition

MRI data were collected at the McLean Imaging Center at McLean Hospital using a Siemens Trio 3T scanner. The local Institutional Review Board of McLean Hospital approved the original study and all participants in the original study gave written informed consent. High-resolution anatomical images were acquired for registration purposes using an MPRAGE sequence with $256 \times 256, 1 \times 1.3$ mm in-plane resolution, and 1.3 mm slice thickness. Resting state BOLD fMRI data were acquired while subjects rested quietly with eyes open. Scan parameters were: 42 slices, $3.5 \times 3.5 \times 3.5$ mm resolution, TE/TR/FA = 24/2,500 ms/82°, interleaved slice acquisition, 64×64 matrix, total scan time = 10 min. The first four frames were discarded for magnetic field equilibration.

Pre-processing

Data pre-processing was done similarly to the simulated data and included: Head motion correction using FLIRT; non-brain removal using BET; spatial smoothing by a Gaussian kernel (FWHM 5 mm); grand-mean intensity normalization of the entire 4D dataset by a single multiplicative factor; and highpass temporal filtering by subtraction of a Gaussian-weighted least-squares straight line fitting with $\sigma = 150$ s. Finally, registration of each subject's fMRI data to their high-resolution structural scan was carried out using FLIRT. Registration of the high-resolution structural images to MNI152 standard space was achieved using FLIRT, with further refinement using FNIRT non-linear registration. All subjects' four-dimensional timeseries data were transformed into standard space at $2 \times 2 \times 2$ mm³ resolution using the registration transformation matrices.

GICA-DR

Figure 3 shows an illustration of the analyses.

All Analysis: For this analysis, all participants' data ($N = 39$) were included in the All:GICA, dual regression, and group-level models, regardless of level of motion. In the **All:Motion** group-level analysis, a two group GLM was set up to compare functional connectivity between two artificially constructed groups based on participant's absolute mean displacement being either smaller than 1 mm (“low motion”; $N = 20$) or larger than 1 mm (“high motion”; $N = 19$), irrespective of diagnostic category. With 1 mm as the cutoff, nearly equal numbers of participants from each diagnostic category are included in each group ($N = 7$ HC in both, $N = 6$ SZ in both, $N = 7$ BPD in “low” and $N = 6$ BPD in “high”), thus diagnostic category was matched across the two groups. In the **All:DC** group-level analysis, a diagnostic category three-group GLM was done to compare functional connectivity between the patient groups and the healthy controls. This analysis included some participants with high levels of motion in each group.

Correlation Between Ground Truth and Dual Regression Timecourses

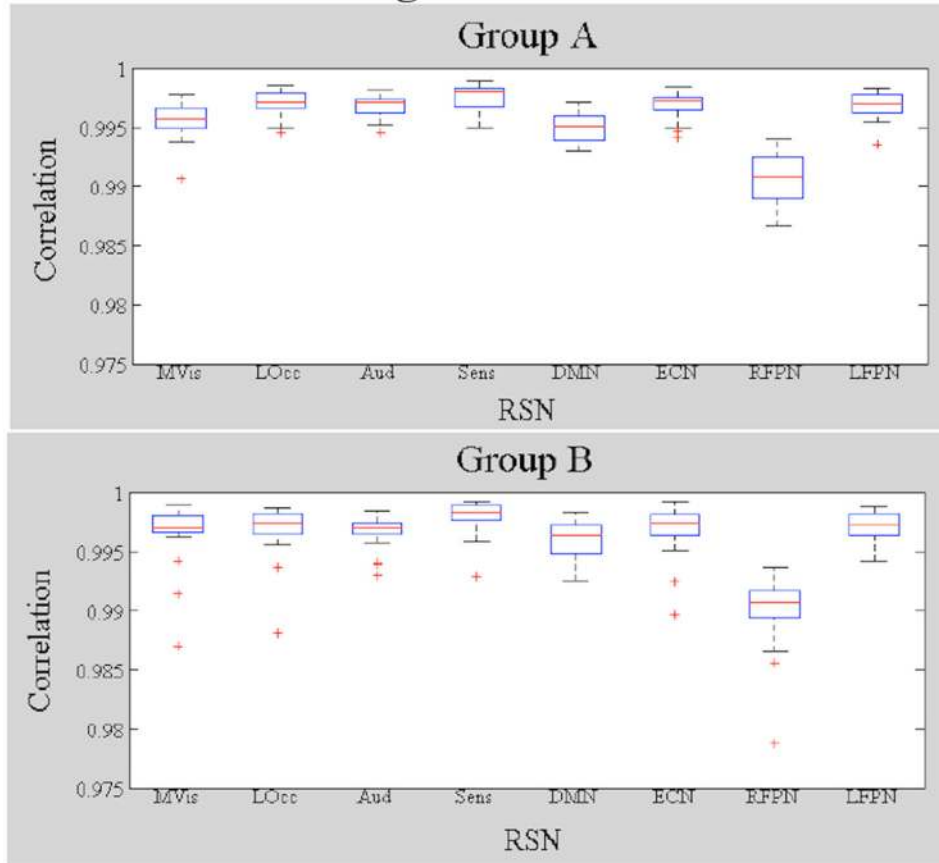


FIGURE 4 | Simulation Results. Boxplots of the correlations between the stage 1 and ground truth timecourses for each network (19 subjects in each group). Data in Group B had amplitude and shape differences in some networks.

Low Motion (LM) Analysis: This analysis used a stricter motion criteria, removing all participants with >1.5 mm motion prior to the LM:GICA and subsequent dual regression and group-level GLM (leaving 6 BD, 10 HC, and 9 SZ). An additional HC participant was also excluded from this analysis even though the max absolute mean displacement was <1 mm because a rotational movement of 0.006 radians around the z axis by the subject led to a large signal change in the dual regression timecourses (this is further discussed in the results section). A single group GLM was done, **LM:DC**, which was a diagnostic category-based three-group GLM to assess group differences between patient groups and the healthy controls. This analysis did not include any participants with high levels of motion and can thus be compared with All:DC to assess how motion can corrupt functional connectivity results. For All:DC and LM:DC, we use results for HC vs. SZ which have approximately the same number of subjects in each group for both analyses.

All:GICA and LM:GICA: The corresponding fMRI data for the subjects included in each GICA ($N = 39$ for All:GICA and $N = 25$ for LM:GICA) were concatenated together in

the temporal dimension and the group ICAs were done using FSL MELODIC (Multivariate Exploratory Linear Decomposition into Independent Components) Version 3.09 (Beckmann and Smith, 2004; Beckmann et al., 2005). The number of components was fixed to twenty for both analyses to give maps similar to those reported in Beckmann et al. (2005). To obtain a stable decomposition of the data, each GICA was run eight times followed by a meta-level GICA fed by all of the spatial maps from the 8 decompositions (Smith et al., 2009; Wisner et al., 2013). The meta-level analysis for the All:GICA estimated 17 components and for the LM:GICA estimated 20 components.

Dual Regression and Group-Level Statistical Analysis: The dual regressions for the All and Low Motion analyses were implemented as described in the simulations, to obtain the statistic maps (\hat{B}_{SM} , \hat{B}_{SM}^*) for each RSN for each subject for each analysis. Estimation and inference on group differences in functional connectivity using \hat{B}_{SM} and \hat{B}_{SM}^* was done using FSL Randomise (5,000 permutations, cluster-based thresholding, $Z = 2.3$ cluster forming threshold) for each network at $p < 0.05$, corrected.

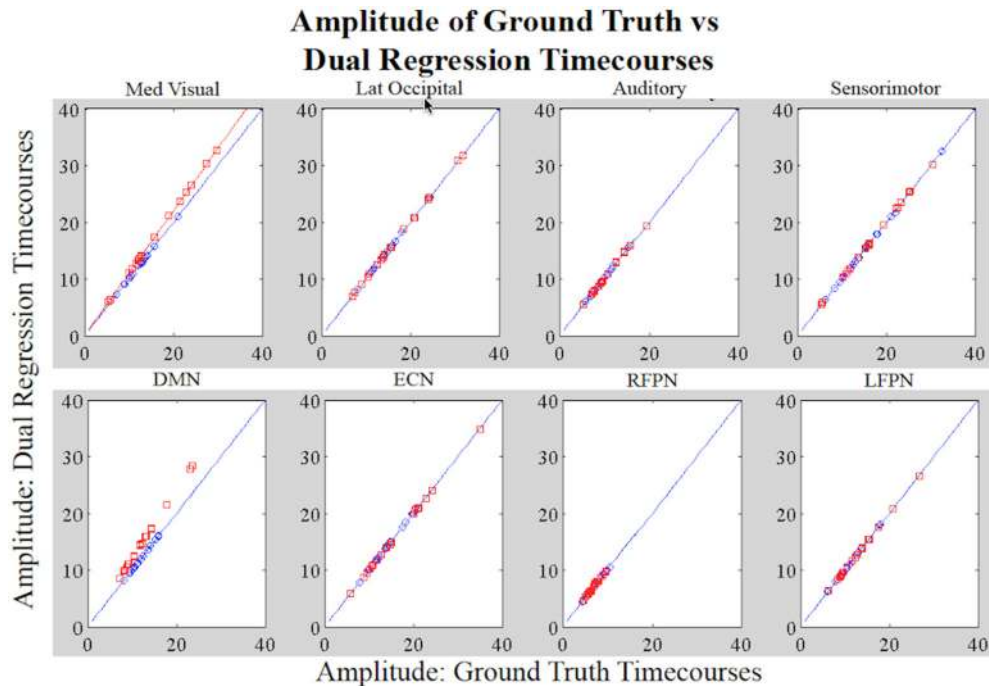


FIGURE 5 | Simulation Results. Each datapoint shown in the plots is the amplitude of the original RSN timecourse vs. the amplitude of the stage 1 timecourse for the subject assigned that timecourse (Group A: blue; Group B: red). For MVN, the network-wide amplitude difference between Groups A and B is reflected by the different slopes (Group A: $y = x$ and Group B: $y = 1.1 \cdot x$; $x =$ original timecourses). For the DMN, the within-network amplitude differences between groups are also reflected by datapoints in Group B not laying on the same line as those in A. All other networks show no group differences in amplitude.

RESULTS OF SIMULATIONS: GROUP DIFFERENCES IN RSN SHAPE AND AMPLITUDE

Accuracy of Stage 1 Timecourses

Figure 4 shows a boxplot of the correlation coefficients between each stage 1 timecourse and each ground truth timecourse for each network for subjects in Groups A and B. We did not alter the shape of the temporal trace for any of the RSNs in the simulated data and the dual regression is able to faithfully reproduce the temporal response of the ground truth timecourses, with the correlations between stage 1 timecourses and ground truth timecourses all near 1. Figure 5 shows the amplitudes of the stage 1 timecourses are as expected, with: A 10% difference in network amplitude in the MVN, reflected in the slopes of the lines of best fit for Group A and Group B, a difference in amplitude (e.g., slope) between the two groups in the DMN, and no difference in amplitudes for any of the other networks.

Group Differences in RSN Functional Connectivity: Amplitude Effects

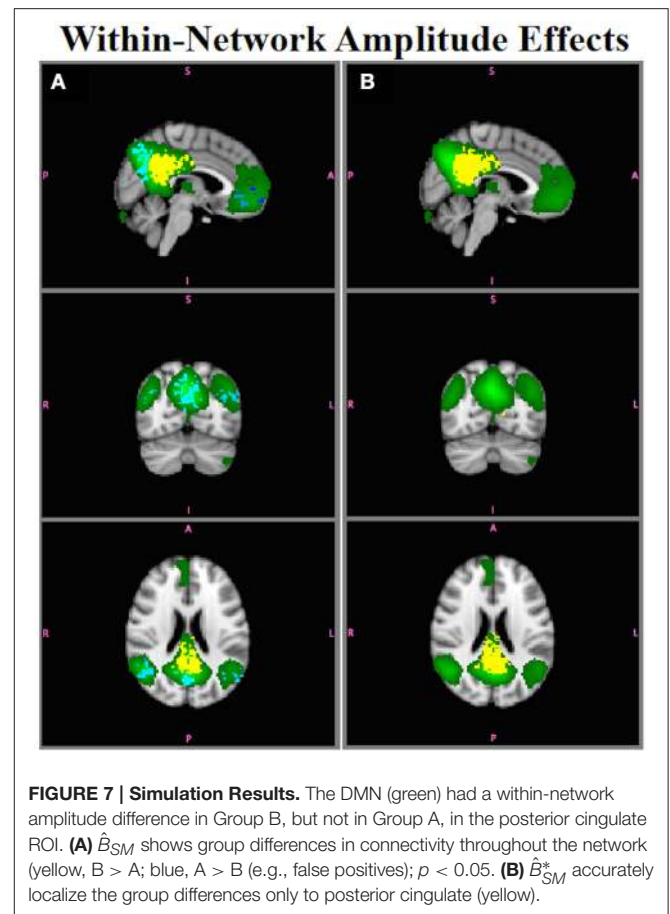
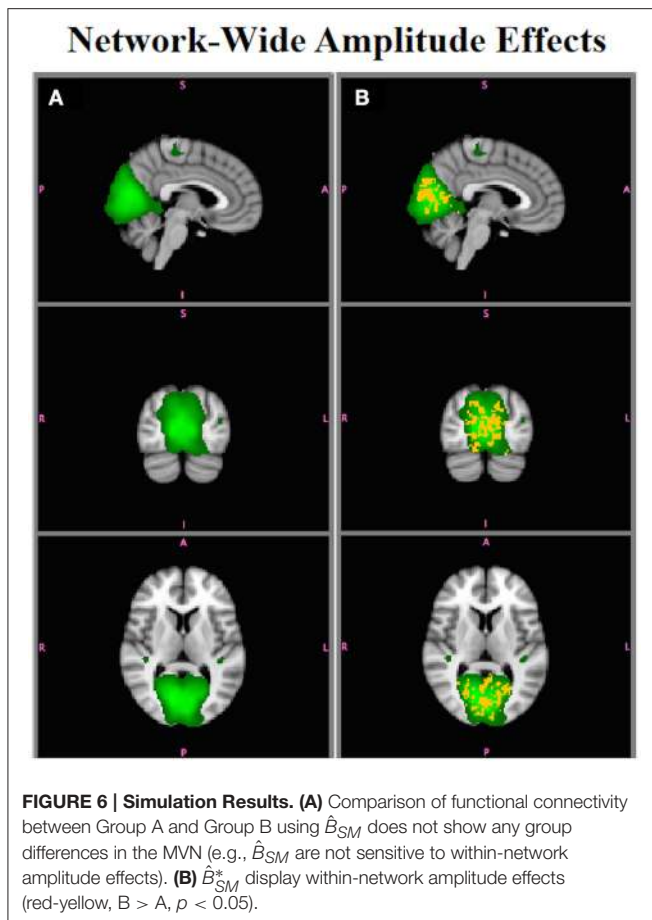
Figure 6 shows differences in functional connectivity between Group A and Group B for the MVN, which had simulated network-wide amplitude differences. In this and all subsequent brain image figures, three orthogonal slices are displayed in each vertical panel, with the royso8 RSN spatial map shown in green overlaid onto the MNI 2 mm standard brain image.

Statistically significant positive differences are shown in red-yellow and negative differences are shown in blue-lightblue. All results are significant at $p < 0.05$ corrected. Figure 6A shows differences assessed using \hat{B}_{SM} as inputs to Randomize, and Figure 6B shows those using \hat{B}_{SM}^* . \hat{B}_{SM} are not sensitive to network-wide amplitude effects, whereas \hat{B}_{SM}^* are sensitive to network-wide amplitude differences, with the voxels showing significant differences lying only within the network that were simulated to have an amplitude effect (Figure 2A).

Figure 7 shows results for comparison of Group A and Group B for the DMN. Figure 7A shows group differences assessed using \hat{B}_{SM} as inputs to Randomize, Figure 7B shows those using \hat{B}_{SM}^* . Although the PCC region was the only region with an amplitude difference between Groups A and B in the DMN (Figure 2B), using \hat{B}_{SM} would lead one to conclude that there were group differences throughout the entire network, with some regions showing greater connectivity in Group B relative to Group A (red-yellow) and others showing decreased connectivity in Group B relative to Group A (blue-light blue). \hat{B}_{SM}^* accurately localizes within-network amplitude differences to only the PCC region with the correct sign (e.g., $B > A$).

Group Differences in RSN Shape

Figures 8 show the results for RSNs with simulated shape differences (shown in Figures 2C,D). Both statistics are capable of accurately reflecting the differences in functional connectivity when those differences are due only to shape differences (or



synchrony between the timecourses from different areas in the network) and not amplitude. In **Figures 8A,B**, the ECN is more strongly connected to caudate/thalamus regions in Group A, using either \hat{B}_{SM} or \hat{B}_{SM}^* , respectively. **Figures 8C,D** show that the LFPN is more strongly connected with caudate/thalamus in Group B, with \hat{B}_{SM} and \hat{B}_{SM}^* , respectively. In both cases, \hat{B}_{SM}^* have a slightly lower p -values than \hat{B}_{SM} (e.g., brighter yellow colors indicate more significant differences than orange colors) because of the extra random effects variation due to including amplitude information in \hat{B}_{SM}^* , but not in \hat{B}_{SM} . No other networks showed group differences in shape.

Results of the *in Vivo* Analysis: Motion and Functional Connectivity

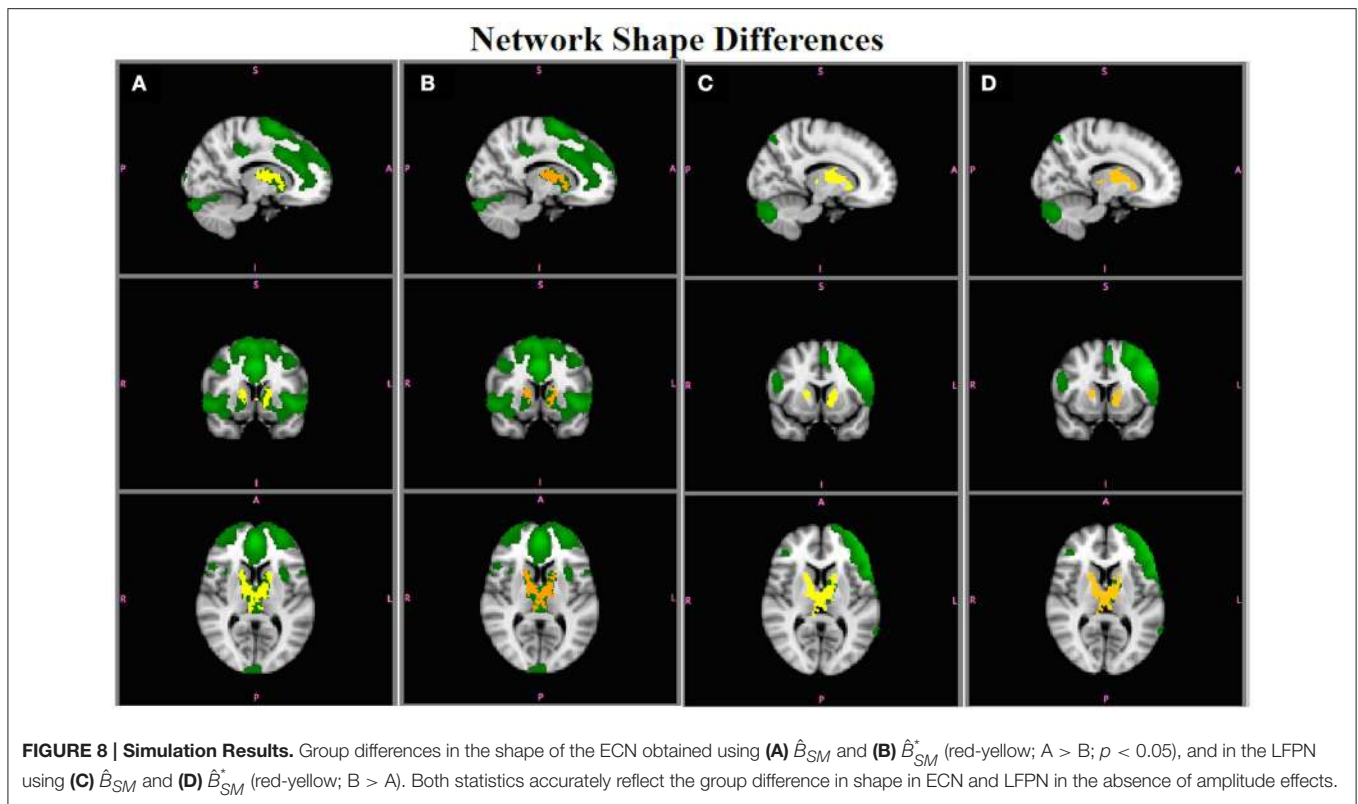
All:GICA vs. LM:GICA

Visual inspection of the independent component spatial maps shows that removing participants with greater motion generally resulted in improvements in the spatial maps, although the GICA maps from both analyses highly resembled (via spatial cross correlation) the roysoc8 maps for 7/8 major networks (all but LFPN, which showed some variability). This is not surprising given that ICA is effective for denoising fMRI data by identifying and separating out effects due to both physiological and subject motion (McKeown et al., 1998; Kiviniemi et al., 2003;

Beckmann et al., 2005; Salimi-Khorshidi et al., 2014; Du et al., 2016). However, there were some subtle differences in the two sets of maps that demonstrated that the LM:GICA generally produced the more reasonable maps. For example, **Figure 9** shows the MVN from the All:GICA (**Figure 9A**) and from the LM:GICA (**Figure 9B**). While the MVNs identified from both analyses do largely resemble each other, in the LM:GICA map, smaller bilateral regions in the lateral geniculate nucleus and the thalamus are identified in the MVN (**Figure 9B**) that are not seen in the All:GICA map (**Figure 9A**). The presence of these regions in MVN is consistent with known connections of LGN and thalamus with visual cortex (Castelo-Branco et al., 1998; Kwon and Jang, 2014), indicating the Low Motion MVN is more reasonable. The LGN were observed in the All:GICA map at a lower threshold of $Z = 1.63$, but this threshold also showed many other widespread regions (although we did not directly compare these two maps for statistically significant differences).

Spurious Group Differences in RSN Functional Connectivity Due to Motion

The All:Motion analysis, which compared RSN functional connectivity in individuals with <1 mm motion vs. those with >1 mm motion, showed several networks with spurious group differences in functional connectivity, for inference on both \hat{B}_{SM}^*




and \hat{B}_{SM} . For example, **Figure 10A** shows that the ECN (green) has widespread group differences in functional connectivity (red-yellow; using \hat{B}_{SM}^*). Since the two groups being compared are matched for diagnostic category, these spurious group differences are likely due to the effects of motion. Inspection of the results for All:DC (**Figure 10B**; SZ > HC) also shows differences in functional connectivity in this network in regions that overlap the map in **Figure 10A**, suggesting that motion is affecting the clinical group comparison as well. Indeed, if a strict motion criteria is applied as in the LM:DC analysis (**Figure 10C**), the group differences in functional connectivity between SZ and HC seen in **Figure 10B** are no longer observed, even at very low thresholds. Although we did attempt to keep the number of subjects in each analysis comparable, there were fewer subjects in the LM:DC analysis (9HC/9SZ) than in the All:DC analysis (14HC/12SZ), which could also reduce the ability to detect group differences, although lowering the statistical threshold still did not show any group differences in LM:DC.

Further evidence that motion-related amplitude effects underlie the observed differences in network functional connectivity observed in the All:Motion and All:DC analyses (**Figures 10A,B**, respectively) comes from inspection of the dual regression network timecourses for each subject in the high (>1.5 mm) vs. low (<1.5 mm) motion groups. The stage 1 timecourses (**Figure 11A** vs. **Figure 11B**) clearly show that many subjects in the high motion group have greater amplitudes (or standard deviations) in the network timecourses. Inspection of the stage 1 timecourses for SZ and HC included in the All:DC vs. LM:DC analyses (**Figures 11C,D**) shows that there

were several HC and SZ with spikes and other motion-related temporal features (identified by comparison with the motion timecourses) included in All:DC that were removed for LM:DC. Another point of note is that the timecourse with the large wide deviation over several TRs in panels **Figure 11A** (red line) and **Figure 11C** (blue line) closely resembles the timecourse of the McFLIRT estimated z-rotation for that subject ($r = -0.76$), however the estimated rotation was less than 0.006 radians, which did not result in any appreciable effect in the mean displacement (not shown). Inspection of the raw fMRI data for this subject clearly showed the rotation and the motion correction was not able to appreciably fix it. This suggests that the dual regression procedure can be sensitive to small levels of motion and that inspection of the stage 1 timecourses can provide key information as to whether or not motion may be corrupting functional connectivity results.

Loss of Sensitivity to Group Differences in RSN Functional Connectivity Due to Motion

Including participants with higher levels of motion (>1.5 mm) in All:DC obscured functional connectivity differences between SZ and HC in MVN that were revealed in the LM:DC with these subjects removed. **Figure 12** shows the differences in functional connectivity for SZ > HC ($p < 0.05$ corrected) using \hat{B}_{SM}^* for LM:DC. No statistically significant group differences in this network were observed in the All:DC. For LM:DC, group differences using \hat{B}_{SM} (not shown) were observed in some regions that overlapped with \hat{B}_{SM}^* . However, there were also differences in white matter and brainstem that were not observed with \hat{B}_{SM}^* ,



and some of the bilateral differences, for example, in the nucleus accumbens, that were observed using \hat{B}_{SM}^* were not observed in the \hat{B}_{SM} results, underscoring the fact that inference on \hat{B}_{SM}^* and \hat{B}_{SM} will generally not give the same results. Since \hat{B}_{SM}^* are able to accurately localize any amplitude effects, whereas \hat{B}_{SM} will mislocalize any within-network amplitude effects, it is recommended to only infer on \hat{B}_{SM}^* maps even though amplitude and synchronicity effects cannot be disambiguated without conducting additional analyses. Notably, the group differences in bilateral LGN and nucleus accumbens, that were observed using \hat{B}_{SM}^* are in agreement with what is known about visual disturbances in schizophrenia (Silverstein and Keane, 2011; Yoon et al., 2013), providing additional support that \hat{B}_{SM}^* are most valid.

Inspection of the stage 1 timecourses for the MVN (Figure 13) shows that, in All:DC, the amplitudes of MVN timecourses are generally higher in HC vs. SZ (Figures 13A,B). After removing high motion subjects from the analysis, the within-group variability in the amplitudes is reduced (Figures 13C,D).

DISCUSSION

The GICA-DR approach is a powerful tool for identifying group differences in brain network functional connectivity that reflects

differences in both network shape and network amplitude. Findings from our simulations based on creating ground truth group differences in network shape and amplitude from real resting state fMRI data show that:

1. Dual regression identifies subject-specific network timecourses with good temporal accuracy in the presence of amplitude and shape effects, assuming no subject variability in location of networks or network nodes.
2. \hat{B}_{SM} SSSMs (e.g., raw regression coefficient maps from stage 2 of the dual regression without normalization of stage 1 timecourses) *should not* be used for inference at the group level because they are insensitive to network-wide amplitude effects, yet will mis-localize any within-network amplitude effects.
3. Using \hat{B}_{SM}^* subject-specific spatial maps (SSSMs) from stage 2 of the dual regression (e.g., with normalization of stage 1 timecourses) for inference at the group level accurately localizes network-wide and within-network amplitude effects, as well as network shape differences.

While it may be tempting to use both \hat{B}_{SM} and \hat{B}_{SM}^* together to disambiguate amplitude effects (e.g., with the first being sensitive only to shape and the second sensitive to shape and network-wide amplitude effects), it is not recommended to do so because of the possibility of within-network amplitude effects. In this case, the \hat{B}_{SM} maps will simply be incorrect. Since one cannot rule out the possibility of within-network amplitude effects without applying additional analyses, the implication is that within the GICA-DR framework, only *combined* amplitude and shape can be assessed. \hat{B}_{SM} cannot be used to disambiguate the two effects. Thus, the important contribution of this paper is to show that design normalization of stage 1 timecourses is not a convenient method for teasing apart shape or network amplitude differences, it is a *necessary* step in the dual regression procedure to deal with any existing amplitude effects in a disciplined way to give group differences that are valid and localized correctly. That is, \hat{B}_{SM} can't be used to infer shape differences since there is the potential to obtain mis-localized amplitude effects masquerading as connectivity differences in these maps. Another consequence of timecourse normalization is that amplitude effects cannot be studied separately from shape effects and vice-versa within the dual regression framework. On the other hand, resting state network amplitude has been shown to be an important network property that may be more interpretable relative to task fMRI, for which amplitude conveys meaningful information. Thus, including amplitude information together with shape can be a strength of the dual regression with timecourse normalization.

Our findings from the *in vivo* analyses illustrate that motion can mimic amplitude effects and that such effects can dramatically impact the connectivity maps, which is well-documented for seed-based connectivity analyses but less well-characterized for GICA-DR. Fortunately, the stage 1 timecourses are useful for identifying subjects with motion effects that have propagated into the connectivity analysis that could drive group differences, even motion effects that are not necessarily large and remain undetected during standard quality assurance. Using \hat{B}_{SM}^* for inference on functional connectivity will not only accurately

All:Motion and All:DC vs LM:DC

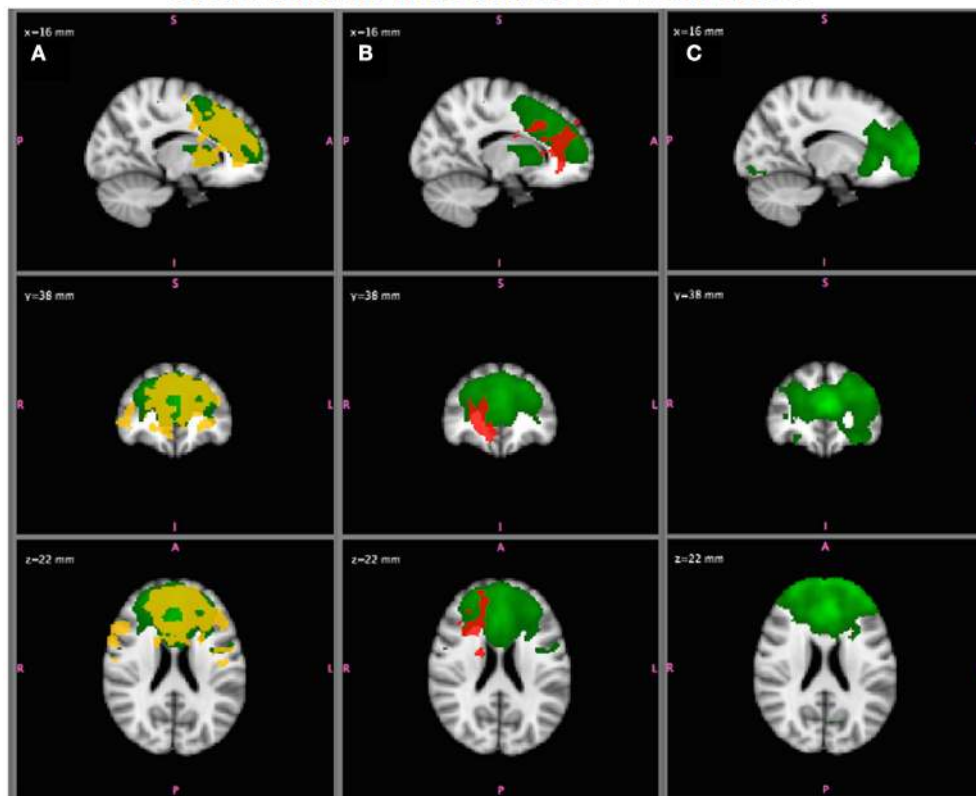


FIGURE 10 | In Vivo Results. (A) All:Motion analysis: the ECN (green) shows statistically significant differences in functional connectivity for high motion > low motion (red-yellow; $p < 0.05$). **(B)** All:DC analysis: the ECN shows group differences for SZ > HC ($p < 0.05$) in regions overlapping with the motion effects in **(A)**. **(C)** Low Motion analysis: group differences in ECN between SZ and HC are no longer observed, even at very low statistical thresholds. Note that the IC map shown in A and B is from the All:GICA. The IC maps shown in C is from the LM:GICA.

localize *real* network amplitude effects, because \hat{B}_{SM}^* are also sensitive to motion-related amplitude effects, any motion effects will at least be accurately localized and thus not lead to spurious mis-localized group differences in other regions or networks that may be attributed to group differences in connectivity. That is, it is easier to detect motion-related amplitude effects in the group difference maps when these effects are accurately localized (with \hat{B}_{SM}^*), e.g., as rings around the brain edges or large differences at tissue interfaces, vs. being mis-localized throughout the image (with \hat{B}_{SM}), not giving the same visual appearance that is suggestive of motion artifacts.

In the context of other published findings, there have been a few studies that have assessed dual regression for back-reconstructing SSTC and SSSM for test-retest reliability and to assess performance in the presence of shape and amplitude differences (Zuo et al., 2010; Allen et al., 2012; Wang et al., 2013; Wisner et al., 2013). In the study by Zuo et al. (2010), they found that GICA-DR produced subject-specific brain network spatial patterns that had moderate-to-high reliability and reproducibility. They also reported that GICA-DR showed superior performance over applying single-subject ICA combined with template-matching procedures to identify the

subject-specific brain networks. Wang et al. (2013) investigated the temporal signatures of RSNs that were generated from stage 1 of the dual regression and found that some univariate metrics such as entropy and the amplitude of low frequencies fluctuations (ALFF) of the stage 1 timecourses were reliable and reproducible. The neurometric profiles of RSNs identified using GICA-DR have also been shown to be highly reproducible and replicable across samples, and highly reliable within individuals for some RSNs (Wisner et al., 2013). For example, the right-lateralized fronto-parietal, auditory, and visuospatial networks had such a profile, suggesting that these networks are consistent with the profile of a trait. Other networks had a neurometric profile of high reproducibility and replicability, but low reliability in individual subjects. Left-lateralized fronto-parietal, emotion, and reward-related networks show this profile and, consequently, were interpreted as networks whose connectivity could change depending on the specific state of the individual.

Allen et al. (2012) investigated the ability of ICA combined with back-projection to investigate resting state networks in the presence of inter-subject variability in simulated FMRI data with spatial, amplitude, and temporal variability. They report that temporal concatenate GICA implemented with

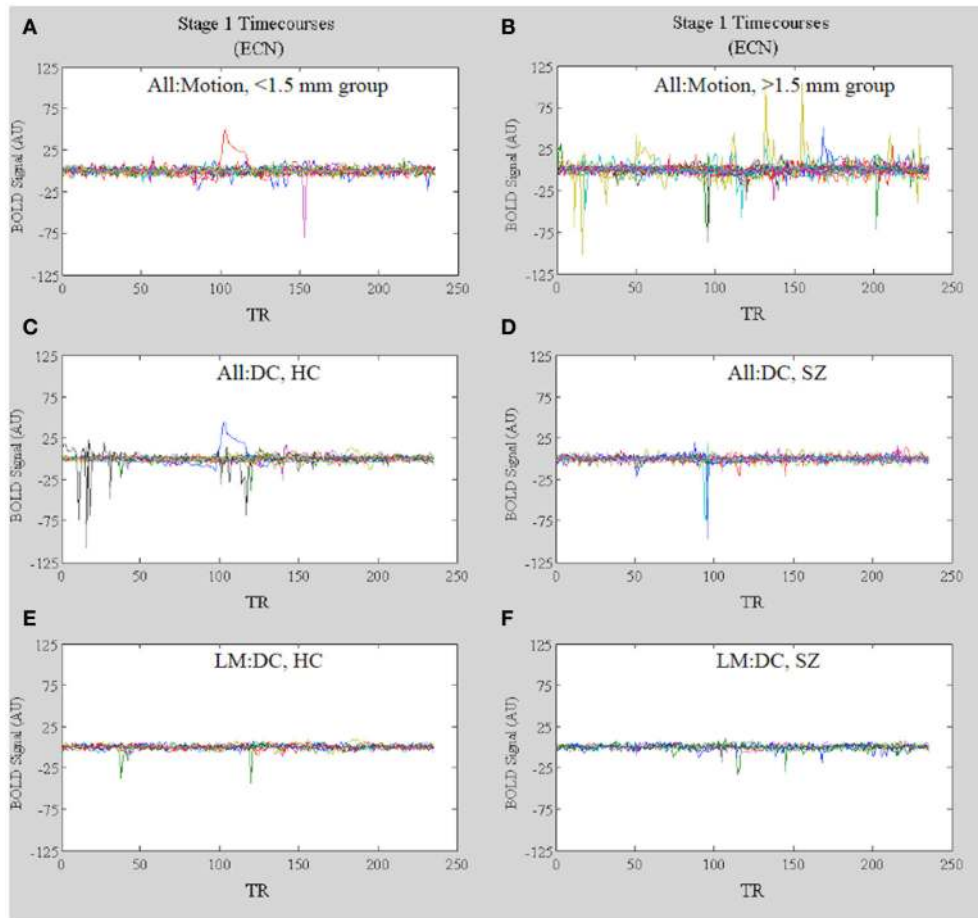


FIGURE 11 | *In Vivo* Results. Stage 1 timecourses for the ECN for: **(A)** All:Motion, lowest motion subjects (<1.5 mm) and **(B)** All:Motion, highest motion subjects (>1.5 mm); **(C)** HC included in the All:DC analysis and **(D)** SZ included in the All:DC analysis; **(E)** HC included in the LM:DC analysis and **(F)** SZ included in the LM:DC analysis. Note that the timecourses for the bipolar patients are included in B but not in **(D)** or **(F)**.

the GIFT toolbox combined with back-projection (Calhoun et al., 2001; <http://icatb.sourceforge.net>) and GICA-DR in FSL should generally perform extremely well for studying large-scale networks, but will be less effective when activations are very small in spatial extent in the presence of spatial variability. However, it is important to note that the general equivalence between the two approaches holds only for the case in which subject-level PCAs implemented in GIFT are not used for data reduction, only for whitening. If this is not the case, there is a potential for PCA bias (Beckmann, 2011).

While the GICA with back-projection in GIFT or dual regression in FSL may perform similarly under certain conditions in empirical studies, there is a very important distinction between back-projection techniques and dual regression. Back-projection techniques are heuristic approaches, [whose outcomes can depend strongly on the within-subject dimensionality reduction (Smith et al., 2014)] while the dual regression is a generative model that provides a statistical framework for assessing group differences in functional connectivity (Beckmann, 2011). As such, while back-projection approaches may be tuneable to give good

results in empirical studies using specific conditions in simulated data (Erhardt et al., 2011; Allen et al., 2012), we advocate the dual regression approach as being theoretically well-principled and practically “safer.” The issue is not one of empirical validation, but one of assessing how the choice of parameters for the dual regression impacts the interpretability of the results with respect to group differences in functional connectivity.

One limitation of the current study is that, in the simulations, we investigated only one level of difference between groups in amplitude effects and shape; namely, we selected a network-wide amplitude difference of 10%, a within-network difference of 50% on one region of the network, and a difference in shape that reflected an absent connection of a region in one network in one group, with the region being connected to a different network in the other group. As such, our results do not speak to the sensitivity and specificity of the statistics we tested over a range of effects and signal levels, nor do our simulations incorporate spatial variability at the subject level. Our primary purpose in this work was to show that design normalization is necessary for dual regression to retain and reflect amplitude (either real

LM:DC vs All:DC




FIGURE 12 | In Vivo Results. Results. LM:DC showed differences in the connectivity of MVN (green), with SZ > HC, in lateral and medial geniculate nuclei, bilateral thalamus, nucleus accumbens, insula, regions in PFC, and ACC, and precuneus (red-yellow, $p < 0.05$). There were no group differences for HC vs. SZ in this network for the ALL:DC analysis.

Stage 1 Timecourses for All:DC vs LM:DC

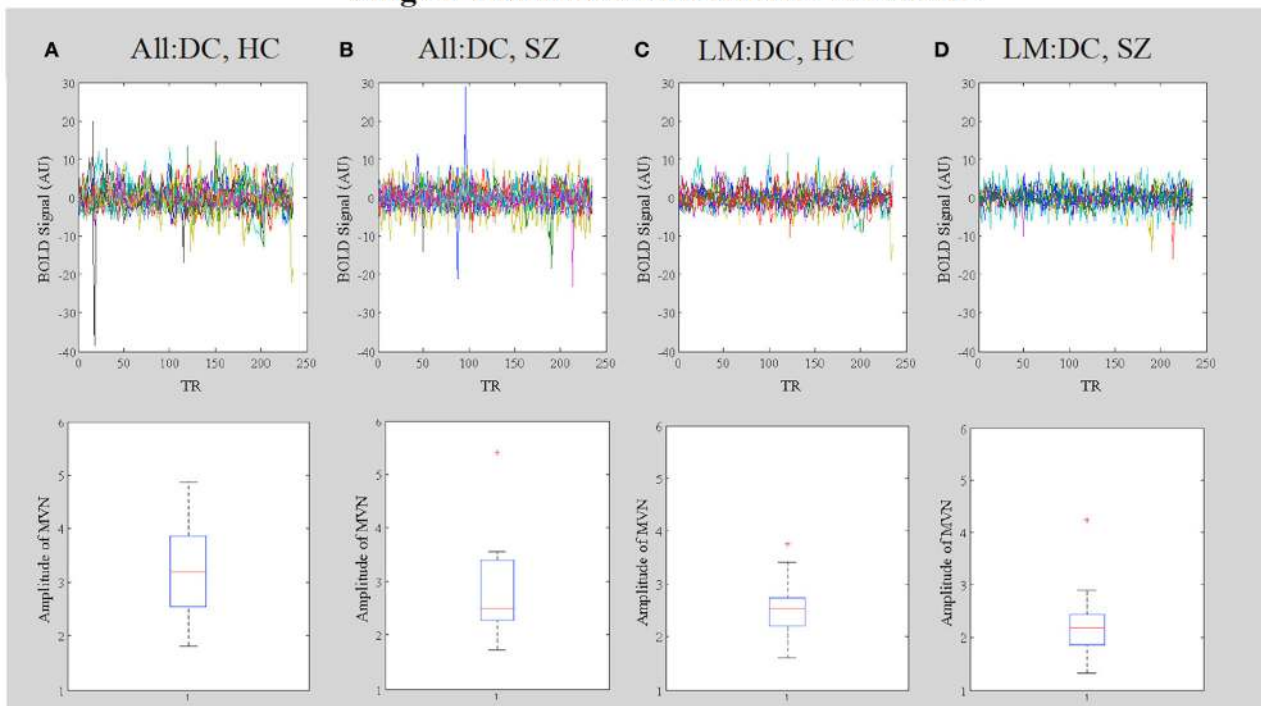


FIGURE 13 | In Vivo Results. Results. Stage 1 timecourses from dual regression can be used to identify corruption by motion. The upper four panels show the stage 1 timecourses for (A) All:DC HC, and (B) All:DC SZ subjects (which included subjects with larger motion effects in both groups). Timecourses for (C) LM:DC HC and (D) LM:DC SZ subjects, in which subjects with >1.5 mm motion were removed. The lower four panels show corresponding boxplots of the amplitudes of the timecourses for each group.

network or motion-related) information accurately. There is evidence that spatial variability across subjects degrades the performance of DR for identifying subject-specific timecourses

and spatial maps (Allen et al., 2012; Kim et al., 2012). While techniques such as iterated DR (Kim et al., 2012) may be necessary to derive accurate subject-specific timecourses and

network maps, especially in the case of high-dimensional GICA for finer functional parcellations, the issue of amplitude effects at the DR stage must still be addressed. Simulating realistic resting state fMRI data with various ground truth sources of inter-subject variability can be used to further investigate these issues.

In sum, \hat{B}_{SM}^* (from dual regression with normalization of stage 1 timecourses) accurately reflect both network amplitude and motion-related amplitude effects in group functional connectivity analyses. However, \hat{B}_{SM}^* may require slightly larger sample sizes given that when these spatial maps are forward to the group analysis, they will have additional random effects variance from the amplitude that will make it more difficult to detect group differences. Inference using raw multiple regression coefficients (without normalization of stage 1 timecourses) can lead to mislocalized group differences, spurious group differences, or loss of sensitivity to real group differences. A stricter motion criterion is recommended for RSFC analyses, and inspection of the subject-specific timecourses from stage 1 of the dual regression is also recommended as part of the quality assurance to assess the effects of small movements (less than a single voxel) that may have large effects on the group difference maps. While we simply removed “high” motion subjects from the analysis, denoising the single-subject fMRI data using ICA (Salimi-Khorshidi et al., 2014) or using other approaches such as data scrubbing and regression of nuisance timecourses out of the data (Power et al., 2012; Siegel et al., 2014) prior to the dual regression or during dual regression (e.g., by including a subject’s motion timecourses with stage 1 timecourses) may be desirable to remove motion effects and thus retain subjects. While these approaches have been shown to effectively remove motion, ICA-based denoising profoundly

improves the reproducibility of group-level RSN spatial maps relative to either scrubbing or nuisance regression (Pruim et al., 2015). This is because the latter two approaches may also affect signals of interest, resulting in no improvement in reproducibility of group-level RSN spatial maps over doing no additional motion denoising, even though both techniques do effectively remove motion effects. Moreover, the fact that motion effects can mimic amplitude effects may be the fundamental reason that functional connectivity analyses are so sensitive to motion effects.

AUTHOR CONTRIBUTIONS

LN, SS, and CB conceived and designed the study. LN conducted the simulations and *in vivo* data analysis, worked with co-authors to execute the study and interpret the findings, and wrote the paper. SS and DÖ contributed fMRI data for the simulations and the *in vivo* analysis, respectively. All authors contributed extensively to the interpretation of the findings and revising the manuscript.

FUNDING

LN received support for this work from NIDA K25 DA017712 and NIDA R01 DA037265. DO received support from NIHM K24 MH104449.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <http://journal.frontiersin.org/article/10.3389/fnins.2017.00115/full#supplementary-material>

REFERENCES

- Allen, E. A., Erhardt, E. B., Wei, Y., Eichele, T., and Calhoun, V. D. (2012). Capturing inter-subject variability with group independent component analysis of fMRI data: a simulation study. *Neuroimage* 59, 4141–4159. doi: 10.1016/j.neuroimage.2011.10.010
- Andrews-Hanna, J. R., Reidler, J. S., Sepulcre, J., Poulin, R., and Buckner, R. L. (2010). Functional-anatomic fractionation of the brain’s default network. *Neuron* 65, 550–562. doi: 10.1016/j.neuron.2010.02.005.
- Beckmann, C. F. (2011). “Robust and unbiased multisubject RSN analysis using ICA and Dual Regression,” in *Workshop on Group Inference For On-Going Activity: How to Compare Intrinsic Functional Connectivity Between Subjects, Annual Meeting Organization for Human Brain Mapping* (Québec).
- Beckmann, C. F., DeLuca, M., Devlin, J. T., and Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 360, 1001–1013. doi: 10.1098/rstb.2005.1634
- Beckmann, C. F., Mackay, C. E., Filippini, N., and Smith, S. M. (2009). Group comparison of resting-state fMRI data using multi-subject ICA and dual regression. *Neuroimage* 47(Suppl. 1), S39–S41. doi: 10.1016/S1053-8119(09)71511-3
- Beckmann, C. F., and Smith, S. M. (2004). Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE Trans. Med. Imaging* 23, 137–152. doi: 10.1109/TMI.2003.822821
- Bianciardi, M., Fukunaga, M., van Gelderen, P., Horowitz, S. G., de Zwart, J. A., and Duyn, J. H. (2009). Modulation of spontaneous fMRI activity in human visual cortex by behavioural state. *Neuroimage* 45, 160–168. doi: 10.1016/j.neuroimage.2008.10.034
- Biswal, 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. *Magn. Reson. Med.* 34, 537–541. doi: 10.1002/mrm.1910340409
- Calhoun, V. D., Adali, T., Pearlson, G. D., and Pekar, J. J. (2001). A Method for making group inferences from functional mri data using independent component analysis. *Hum. Brain Mapp.* 14, 140–151. doi: 10.1002/hbm.1048
- Castelo-Branco, M., Neuenschwander, S., and Singer, W. (1998). Synchronization of visual responses between the cortex, lateral geniculate nucleus, and retina in the anesthetized cat. *J. Neurosci.* 18, 6395–6410.
- Chai, X. J., Whitfield-Gabrieli, S., Shinn, A. K., Gabrieli, J. D., Nieto Castanon, A., McCarthy, J. M., et al. (2011). Abnormal medial prefrontal cortex resting-state connectivity in bipolar disorder and schizophrenia. *Neuropsychopharmacology* 36, 2009–2017. doi: 10.1038/npp.2011.88
- Cole, D. M., Smith, S. M., and Beckmann, C. F. (2010). Advances and pitfalls in the analysis and interpretation of resting-state fMRI data. *Front. Syst. Neurosci.* 4:8. doi: 10.3389/fnsys.2010.00008
- Cordes, D., Carew, J. D., Eghbalnia, H., Meyerand, M. E., Quigley, M. A., Arfanakis, K., et al. (1999). “Resting-state functional connectivity study using Independent Component Analysis,” in *Proceedings of the International Society for Magnetic Resonance in Medicine* (Philadelphia, PA).
- Du, Y., Allen, E. A., He, H., Sui, J., Wu, L., and Calhoun, V. D. (2016). Artifact removal in the context of group ICA: a comparison of single-subject and group approaches. *Hum. Brain Mapp.* 37, 1005–1025. doi: 10.1002/hbm.23086
- Erhardt, E. B., Rachakonda, S., Bedrick, E. J., Allen, E. A., Adali, T., and Calhoun, V. D. (2011). Comparison of multi-subject ICA methods for analysis of fMRI data. *Hum. Brain Mapp.* 32, 2075–2095. doi: 10.1002/hbm.21170

- Filippini, N., Macintosh, B., Hough, M., Goodwin, G., Frisoni, G. B., Smith, S. M., et al. (2009). Distinct patterns of brain activity in young carriers of the APOE-ε4 allele. *Proc. Natl. Acad. Sci. U.S.A.* 106, 7209–7214. doi: 10.1073/pnas.0811879106
- 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. *Proc. Natl. Acad. Sci. U.S.A.* 100, 253–258. doi: 10.1073/pnas.0135058100
- Greicius, M. D., Srivastava, G., Reiss, A. L., Menon, V. (2004). Default-mode network activity distinguishes Alzheimer's disease from healthy aging: evidence from functional MRI. *Proc. Natl. Acad. Sci. U.S.A.* 101, 4637–4642. doi: 10.1073/pnas.0308627101
- Gusnard, D. A., and Raichle, M. E. (2001). Searching for a baseline: functional imaging and the resting human brain. *Nat. Rev. Neurosci.* 2, 685–694. doi: 10.1038/35094500
- Jao, T., Vertes, P. E., Alexander-Bloch, A. F., Tang, I. N., Yu, Y. C., Chen, J. H., et al. (2013). Volitional eyes opening perturbs brain dynamics and functional connectivity regardless of light input. *Neuroimage* 69, 21–34. doi: 10.1016/j.neuroimage.2012.12.007
- Jenkinson, M., Bannister, P. R., Brady, J. M., and Smith, S. M. (2002). Improved optimisation for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* 17, 825–841. doi: 10.1006/nimg.2002.1132
- Kannurpatti, S. S., and Biswal, B. B. (2008). Detection and scaling of task-induced fMRI-BOLD response using resting state fluctuations. *Neuroimage* 40, 1567–1574. doi: 10.1016/j.neuroimage.2007.09.040
- Kannurpatti, S. S., and Rypma, B., Biswal, B. B. (2012). Prediction of task-related BOLD fMRI with amplitude signatures of resting-state fMRI. *Front. Syst. Neurosci.* 6:7. doi: 10.3389/fnsys.2012.00007
- Kim, Y. H., Kim, J., and Lee, J. H. (2012). Iterative approach of dual regression with a sparse prior enhances the performance of independent component analysis for group functional magnetic resonance imaging (fMRI) data. *Neuroimage* 63, 1864–1889. doi: 10.1016/j.neuroimage.2012.08.055
- Kiviniemi, V. J., Kantola, J.-H., Jauhainen, J., Hyvärinen, A., Tervonen, O. (2003). Independent component analysis of nondeterministic fMRI signal sources. *Neuroimage* 19, 253–260. doi: 10.1016/S1053-8119(03)00097-1
- Kiviniemi, V. J., Haanpää, H., Kantola, J. H., Jauhainen, J., Vainionpää, V., Alahuhta, S., et al. (2005). Midazolam sedation increases fluctuation and synchrony of the resting brain BOLD signal. *Magn. Reson. Imaging* 23, 531–537. doi: 10.1016/j.mri.2005.02.009
- Kwon, H. G., and Jang, S. H. (2014). Neural connectivity of the lateral geniculate body in the human brain: diffusion tensor imaging study. *Neurosci. Lett.* 578, 66–70. doi: 10.1016/j.neulet.2014.06.033
- McKeown, M. J., Makeig, S., Brown, G. G., Jung, T. P., Kindermann, S. S., Bell, A. J., et al. (1998). Analysis of fMRI data by blind separation into independent spatial components. *Hum. Brain Mapp.* 60, 160–188. doi: 10.1002/(SICI)1097-0193(1998)6:3<160::AID-HBM5>3.0.CO;2-1
- McKiernan, K. A., Kaufman, J. N., Kucera-Thompson, J., and Binder, J. R. (2003). A parametric manipulation of factors affecting task-induced deactivation in functional neuroimaging. *J. Cogn. Neurosci.* 15, 394–408. doi: 10.1162/089892903321593117
- Niazy, R. K., Xie, J., Miller, K., Beckmann, C. F., and Smith, S. M. (2011). Spectral characteristics of resting state networks. *Prog. Brain Res.* 193, 259–276. doi: 10.1016/B978-0-444-53839-0.00017-X
- Öngür, D., Lundy, M., Greenhouse, I., Shinn, A. K., Menon, V., Cohen, M. B., and Renshaw, P. F. (2010). Default mode network abnormalities in bipolar disorder and schizophrenia. *Psychiatry Res.* 183, 59–68. doi: 10.1016/j.psychres.2010.04.008
- Power, J. D., Barnes, K. A., Snyder, A. Z., Schlagger, B. L., and Petersen, S. E. (2012). Spurious but systematic correlations in functional connectivity MRI networks arise from subject motion. *Neuroimage* 59, 2142–2154. doi: 10.1016/j.neuroimage.2011.10.018
- Power, J. D., Mitra, A., Laumann, T. O., Snyder, A. Z., Schlaggar, B. L., and Petersen, S. E. (2014). Methods to detect, characterize, and remove motion artefact in resting state fMRI. *Neuroimage* 84, 320–341. doi: 10.1016/j.neuroimage.2013.08.048
- Power, J. D., Schlaggar, B. L., and Petersen, S. E. (2015). Recent progress and outstanding issues in motion correction in resting state fMRI. *Neuroimage* 105, 536–551. doi: 10.1016/j.neuroimage.2014.10.044
- Pruim, R. H. R., Mennes, M., Buitelaar, J. K., and Beckmann, C. F. (2015). Evaluation of ICA-AROMA and alternative strategies for motion artefact removal in resting state fMRI. *Neuroimage* 112, 278–287. doi: 10.1016/j.neuroimage.2015.02.063
- 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. *Proc. Natl. Acad. Sci. U.S.A.* 98, 676–682. doi: 10.1073/pnas.98.2.676
- Raichle, M. E., and Snyder, A. Z. (2007). A default mode of brain function: a brief history of an evolving idea. *NeuroImage* 37, 1083–1090; discussion 1097–1089. doi: 10.1016/j.neuroimage.2007.02.041
- Salimi-Khorshidi, G., Douaud, G., Beckmann, C. F., Glasser, M. F., Griffanti, L., and Smith, S. M. (2014). Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage* 90, 449–468. doi: 10.1016/j.neuroimage.2013.11.046
- Satterthwaite, T. D., Wolf, D. H., Loughhead, J., Ruparel, K., Elliott, M. A., Hakonarson, H., et al. (2012). Impact of in-scanner head motion on multiple measures of functional connectivity: relevance for studies of neurodevelopment in youth. *Neuroimage* 60, 623–632. doi: 10.1016/j.neuroimage.2011.12.063
- Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., et al. (1997). Common blood flow changes across visual tasks: II. decreases in cerebral cortex. *J. Cogn. Neurosci.* 9, 648–663. doi: 10.1162/jocn.1997.9.5.648
- Siegel, J. S., Power, J. D., Dubis, J. W., Vogel, A. C., Church, J. A., Schlaggar, B. L., et al. (2014). Statistical improvements in functional magnetic resonance imaging analyses produced by censoring high-motion data points. *Hum. Brain Mapp.* 35, 1981–1996. doi: 10.1002/hbm.22307
- Silverstein, S. M., and Keane, B. P. (2011). Vision science and schizophrenia research: toward a re-view of the disorder, Editors' Introduction to Special Section. *Schiz. Bull.* 37, 681–689. doi: 10.1093/schbul/sbr053
- Singh, K. D., and Fawcett, I. P. (2008). Transient and linearly graded deactivation of the human default-mode network by a visual detection task. *Neuroimage* 41, 100–112. doi: 10.1016/j.neuroimage.2008.01.051
- Smith, S. M. (2002). Fast robust automated brain extraction. *Hum. Brain Mapp.* 17, 143–155. doi: 10.1002/hbm.10062
- Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., et al. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proc. Natl. Acad. Sci. U.S.A.* 106, 13040–13045. doi: 10.1073/pnas.0905267106
- Smith, S. M., Hyvarinen, A., Varoquaux, G., Miller, K. L., and Beckmann, C. F. (2014). Group-PCA for very large fMRI datasets. *Neuroimage* 101, 738–749. doi: 10.1016/j.neuroimage.2014.07.051
- Smith, S. M., Jenkinson, M., Woolrich, M. W., Beckmann, C. F., Behrens, T. E. J., Johansen-Berg, H., et al. (2004). Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage* 23, 208–219. doi: 10.1016/j.neuroimage.2004.07.051
- van de Ven, V. G., Formisano, E., Prvulovic, D., Roeder, C. H., and Linden, D. E. (2004). Functional connectivity as revealed by spatial independent component analysis of fMRI measurements during rest. *Hum. Brain Mapp.* 22, 165–178. doi: 10.1002/hbm.20022
- Van Dijk, K. R. A., 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. *J. Neurophysiol.* 103, 297–321. doi: 10.1152/jn.00783.2009
- Van Dijk, K. R., Sabuncu, M. R., and Buckner, R. L. (2012). The influence of head motion on intrinsic functional connectivity MRI. *Neuroimage* 59, 431–438. doi: 10.1016/j.neuroimage.2011.07.044
- Vargas, C., Lopez-Jaramillo, C., and Vieta, E. (2013). A systematic literature review of resting state network—functional MRI in bipolar disorder. *J. Affect. Disord.* 150, 727–735. doi: 10.1016/j.jad.2013.05.083
- Wang, X., Jiao, Y., Tang, T., and Wang, H., Lu, Z. (2013). Investigating univariate temporal patterns for intrinsic connectivity networks based on complexity and low-frequency oscillation: a test-retest reliability study. *Neuroscience* 254, 404–426. doi: 10.1016/j.neuroscience.2013.09.009
- Wisner, K. M., Atluri, G., and Lim, K. O., MacDonald, I. I. I., A.W. (2013). Neurometrics of intrinsic connectivity networks at rest using fMRI: retest reliability and cross-validation using a meta-level method. *Neuroimage* 76, 236–251. doi: 10.1016/j.neuroimage.2013.02.066

- Yan, C., Liu, D., He, Y., Zou, Q., Zhu, C., Zuo, X., et al. (2009). Spontaneous brain activity in the default mode network is sensitive to different resting-state conditions with limited cognitive load. *PLoS ONE* 4:e5743. doi: 10.1371/journal.pone.0005743
- Yoon, J. H., Sheremata, S. L., Rokem, A., and Silver, M. A. (2013). Windows to the soul: vision science as a tool for studying biological mechanisms of information processing deficits in schizophrenia. *Front. Psychol.* 4:681. doi: 10.3389/fpsyg.2013.00681
- Zang, Y. F., He, Y., Zhu, C. Z., Cao, Q. J., Sui, M. Q., Liang, M., et al. (2007). Altered baseline brain activity in children with ADHD revealed by resting-state functional MRI. *Brain Dev.* 29, 83–91. doi: 10.1016/j.braindev.2006.07.002
- Zuo, X. N., Kelly, C., Adelman, J. S., Klein, D. F., Castellanos, F. X., and Milham, M. P. (2010). Reliable intrinsic connectivity networks: test-retest evaluation using ICA and dual regression approach. *Neuroimage* 49, 2163–2177. doi: 10.1016/j.neuroimage.2009.10.080

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2017 Nickerson, Smith, Öngür and Beckmann. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.