



Universiteit  
Leiden  
The Netherlands

## **Catecholamine function, brain state dynamics, and human cognition**

Brink, R.L. van den

### **Citation**

Brink, R. L. van den. (2017, November 7). *Catecholamine function, brain state dynamics, and human cognition*. Retrieved from <https://hdl.handle.net/1887/54947>

Version: Not Applicable (or Unknown)

License: [Licence agreement concerning inclusion of doctoral thesis in the Institutional Repository of the University of Leiden](#)

Downloaded from: <https://hdl.handle.net/1887/54947>

**Note:** To cite this publication please use the final published version (if applicable).

Cover Page



Universiteit Leiden



The handle <http://hdl.handle.net/1887/54947> holds various files of this Leiden University dissertation.

**Author:** Van den Brink R.L.

**Title:** Catecholamine function, brain state dynamics, and human cognition

**Issue Date:** 2017-11-07

### 3. Catecholaminergic Neuromodulation Shapes Intrinsic MRI Functional Connectivity in the Human Brain

#### Abstract

The brain commonly exhibits spontaneous (i.e., in the absence of a task) fluctuations in neural activity that are correlated across brain regions. It has been established that the spatial structure, or topography, of these intrinsic correlations is in part determined by the fixed anatomical connectivity between regions. However, it remains unclear which factors dynamically sculpt this topography as a function of brain state. Potential candidate factors are subcortical catecholaminergic neuromodulatory systems, such as the locus coeruleus-norepinephrine (LC-NE) system, which send diffuse projections to most parts of the forebrain. Here, we systematically characterized the effects of endogenous central neuromodulation on correlated fluctuations during rest in the human brain. Using a double-blind placebo-controlled cross-over design, we pharmacologically increased synaptic catecholamine levels by administering atomoxetine, an NE transporter blocker, and examined the effects on the strength and spatial structure of resting-state MRI functional connectivity. First, atomoxetine reduced the strength of inter-regional correlations across three levels of spatial organization, indicating that catecholamines reduce the strength of functional interactions during rest. Second, this modulatory effect on intrinsic correlations exhibited a substantial degree of spatial specificity: the decrease in functional connectivity showed an anterior-posterior gradient in the cortex, depended on the strength of baseline functional connectivity, and was strongest for connections between regions belonging to distinct resting-state networks. Thus, catecholamines reduce intrinsic correlations in a spatially heterogeneous fashion. We conclude that neuromodulation is an important factor shaping the topography of intrinsic functional connectivity.

*This chapter is based on:*

van den Brink RL, Pfeffer T, Warren CM, Murphy PR, Tona KD, van der Wee NJ, Giltay E, van Noorden MS, Rombouts SA, Donner TH, and Nieuwenhuis, S (2016). Catecholaminergic Neuromodulation Shapes Intrinsic MRI Functional Connectivity in the Human Brain. *The Journal of Neuroscience*, 36(30): 7865-7876.

### 3.1 Introduction

The resting-state – here defined as periods during which a participant is not engaged in a complex explicit task – is characterized by fluctuations in neural activity that are correlated across brain regions (Biswal et al., 1995; Leopold et al., 2003; Fox and Raichle, 2007; Hiltunen et al., 2014). Such spontaneous, correlated fluctuations exhibit a rich spatial (Yeo et al., 2011) and temporal (Allen et al., 2014; Zalesky et al., 2014) structure that is reflective of the brain's functional organization (Tavor et al., 2016). The strength and spatial distribution of these correlated fluctuations are predictive of behavior and pathological conditions (Greicius et al., 2004; De Luca et al., 2005). Moreover, the global structure, or topography, of correlated activity changes dynamically with alterations in conscious state (Barttfeld et al., 2015) and task conditions (Nir et al., 2006; Sepulcre et al., 2010). While the existence and overall spatio-temporal structure of the spontaneous inter-regional correlations are well-established (Fox and Raichle, 2007), uncertainty remains regarding the underlying physiological mechanisms. It has been proposed that correlations across distant brain regions could be induced by brainstem neuromodulatory systems – and in particular the locus coeruleus-norepinephrine (LC-NE) system, which sends diffuse, ascending projections to the forebrain (Leopold et al., 2003; Drew et al., 2008; Schölvinck et al., 2010), where noradrenergic terminals co-release dopamine (DA; Devoto and Flore, 2006). Here, we examined if and how the catecholaminergic neuromodulators NE and DA shape correlated fluctuations during rest in the human brain.

A number of observations suggest that catecholamines should generally increase the strength of functional connectivity. Both iontophoretic NE application and DA agonism enhance neuronal responses to excitatory synaptic input (Rogawski and Aghajanian, 1980; Seamans et al., 2001b; Wang and O'Donnell, 2001). Furthermore, NE and DA can amplify synaptic GABAergic inhibition (Moises et al., 1979; Seamans et al., 2001a). These and other findings have led to the view that catecholamines boost the efficacy of synaptic interactions between neurons (Berridge and Waterhouse, 2003; Winterer and Weinberger, 2004), resulting in an increased difference in firing rates between strongly and weakly active neurons. Such signal amplification yields a system-wide facilitation of signal transmission (Waterhouse et al., 1998). Recent computational work suggests that this effect of catecholamines should boost both positive and negative temporal correlations between the activities of local groups of neurons, resulting in stronger and increasingly clustered network connectivity (Donner and Nieuwenhuis, 2013; Eldar et al., 2013). Putative behavioral and pupillary indices of heightened NE activity have accordingly been shown to co-occur with stronger functional coupling throughout the brain (Eldar et al., 2013). A first consideration of the anatomy of the LC-NE system suggests that these changes in functional connectivity might show little spatial specificity. LC neurons exhibit tightly synchronous firing and collateralize broadly, resulting in largely homogeneous catecholaminergic

innervation throughout the brain (Swanson and Hartman, 1975; Aston-Jones et al., 1984; Ishimatsu and Williams, 1996; Berridge and Waterhouse, 2003).

In the present study we systematically characterized catecholamine effects on the strength and spatial structure of resting-state inter-regional correlations, measured with functional magnetic resonance imaging (fMRI). Using a double-blind placebo-controlled cross-over design, we manipulated catecholamine activity by administering a single dose of atomoxetine, a selective NE transporter (NET) blocker. Within the cortex NET is also responsible for DA reuptake, due to the cortical paucity of DA transporters (Devoto and Flore, 2006). Thus, NET blockers increase both central NE and cortical DA availability (Bymaster et al., 2002; Devoto et al., 2004; Swanson et al., 2006; Koda et al., 2010). We systematically quantified catecholamine effects on functional connectivity—globally, between brain networks, and at the level of individual connections between brain regions. In contrast to the notion of a catecholamine-induced homogeneous increase in functional connectivity, we found that atomoxetine *reduced* correlations across most pairs of brain regions. Most remarkably, atomoxetine altered the strength of inter-regional correlations in a highly spatially specific manner. These results have important ramifications for our understanding of resting-state activity and central catecholaminergic function.

### 3.2 Materials and Methods

*Participants.* Neurologically healthy right-handed individuals (N=24, age 19-26, 5 male) were recruited and medically screened by a physician for physical health and drug contraindications. Exclusion criteria included: standard contraindications for MRI; current use of psychoactive or cardiovascular medication; a history of psychiatric illness or head trauma; cardiovascular disease; renal failure; hepatic insufficiency; glaucoma; hypertension; drug or alcohol abuse; learning disabilities; poor eyesight (myopia  $\leq -6$  diopters); smoking more than 5 cigarettes a day; and pregnancy. All participants gave written informed consent prior to the experiment and screening, and were compensated with €135 or course credit.

*Design and functional MRI data.* We used a double-blind placebo-controlled cross-over design. In each of two sessions, scheduled one week apart at the same time of day, participants received either a single oral dose of atomoxetine (40 mg) or placebo (125 mg of lactose monohydrate with 1% magnesium stearate, visually identical to the drug). Elsewhere we report data showing that the atomoxetine treatment significantly increased salivary levels of cortisol and alpha amylase, reliable markers of sympathetic nervous system and hypothalamus-pituitary-adrenal axis activation, respectively (Warren et al., in preparation), thus confirming drug uptake. In both sessions participants were scanned once before pill ingestion ( $t = -20$  min) and once at  $t = 90$  min, when approximate peak-plasma levels are reached. The interaction contrast (post atomoxetine – pre atomoxetine) minus (post placebo – pre placebo)

allowed us to examine the effects of atomoxetine while controlling for other session-related differences. Each scan comprised 8 minutes of eyes-open resting-state fMRI. During scanning the room was dark, and participants fixated on a black fixation cross presented on a gray background.

*MRI data collection and preprocessing.* All MRI data were collected with a Philips 3T MRI scanner. In each of the scanning sessions we collected a T2\*-weighted echo planar imaging (EPI) resting-state image (echo time 30 ms, repetition time 2.2 s, flip angle 80°, field of view 80 x 80 x 38 voxels of size 2.75 mm isotropic, and 216 volumes). To allow magnetic equilibrium to be reached, the first 5 volumes were automatically discarded.

In addition, each time the participant entered the scanner we collected a B<sub>0</sub> field inhomogeneity scan (echo time 3.2 ms, repetition time 200 ms, flip angle 30°, and field of view 256 x 256 x 80 voxels with a reconstructed size of 0.86 x 0.86 mm with 3 mm thick slices). Finally, at the start of the first session we collected a high-resolution anatomical T1 image (echo time 4.6 ms, repetition time 9.77 ms, flip angle 8°, and field of view 256 x 256 x 140 voxels with size 0.88 x 0.88 mm with 1.2 mm thick slices).

We used tools from the FMRIB Software Library (FSL) for preprocessing of the MRI data (Smith et al., 2004; Jenkinson et al., 2012). EPI scans were first realigned using MCFLIRT motion correction and skull-stripped using BET brain extraction. We used B<sub>0</sub> unwarping to control for potential differences in head position each time the participant entered the scanner and resulting differences in geometric distortions in the magnetic field. The B<sub>0</sub> scans were first reconstructed into an unwrapped phase angle and magnitude image. The phase image was then converted to units radians per second and median-filtered, and the magnitude image was skull-stripped. We then used FEAT to unwarped the EPI images in the y-direction with a 10% signal loss threshold and an effective echo spacing of 0.332656505.

The unwarped EPI images were then high-pass filtered at 100 s, prewhitened, smoothed at 5 mm FWHM, and co-registered with the anatomical T1 to 2 mm isotropic MNI space (degrees of freedom: EPI to T1, 3; T1/EPI to MNI, 12). Any remaining artifacts (e.g., motion residual, susceptibility-motion interaction, cardiac and sinus artifacts) were removed using FMRIB's ICA-based X-noiseifier (Griffanti et al., 2014; Salimi-Khorshidi et al., 2014) with pre-trained weights (Standard.RData). Noise classification performance was checked afterwards, by manually classifying components as either 'signal', 'noise', or 'unknown'. Then, the accuracy of the automated artifact detection algorithm was quantified as the percentage of components that had the label 'noise' in both classifications. The accuracy was found to be 96.4% correct. All subsequent analyses were conducted in MATLAB 2012a.

*Physiological recordings and correction.* We recorded heart rate using a pulse oximeter and breath rate using a pneumatic belt at 500 Hz during acquisition of each EPI scan. We used these time series for retrospective image correction [RETROICOR; (Glover et al., 2000)]. This method assigns cardiac and respiratory phases to each volume in each individual EPI time series which can then be removed

from the data. The physiological time series were first down-sampled to 100 Hz. Next, the pulse oximetry data were band-pass filtered between 0.6 and 2 Hz, and the respiration data were low-pass filtered at 1 Hz, using a two-way FIR filter. We then extracted peaks in each time series corresponding to maximum blood oxygenation and maximum diaphragm expansion. The inter-peak intervals were then converted to phase-time by linearly interpolating across the intervals to between 0 and  $2\pi$ . Next, we used these phase-time series to extract the sine- and co-sine components of the dominant and first harmonic Fourier series of each signal. After down-sampling to the EPI sample rate this yielded 8 regressors (4 cardiac and 4 respiratory) that could then be used to remove cardiac and respiratory effects from the blood-oxygen-dependent (BOLD) time series using multiple linear regression. The findings reported here were based on non-corrected data, but we replicated all of our results using the RETROICOR-corrected data (see Results).

*Pupillometry.* Pupil size was measured from the right eye at 500 Hz with an MRI-compatible Eyelink 1000 eye tracker. Blinks and other artifacts were interpolated offline using shape-preserving piecewise cubic interpolation. Pupil data were low-pass filtered at 5 Hz to remove high-frequency noise and Z-scored across conditions. Five participants were excluded from pupil-related analyses due to poor signal quality (>50% of continuous time series interpolated) or missing data. Of the remaining participants, on average 20% (SD 9%) of the data were interpolated.

*Brain parcellation and connectivity.* Time series of brain regions were extracted for the 90 regions of the Automated Anatomical Labeling [AAL, (Tzourio-Mazoyer et al., 2002)] atlas (Fig. 1a). We did not include the cerebellum because it was not fully inside the field of view for all participants. Following averaging across voxels within each brain region, time series ( $M$ ) for each run  $i$  were Z-scored and correlation matrices ( $R$ ) were computed between them via:

$$R_i = \frac{M_i' \cdot M_i}{nTR - 1}$$

where  $'$  denotes transposition and  $nTR$  is the number of volumes (211). Because positive and negative correlations jointly determine a network's functional organization (Fox et al., 2005), many prior studies have used the absolute value of the correlation coefficient to describe functional interactions (Achard and Bullmore, 2007; Eldar et al., 2013; Li et al., 2013). Moreover, computational work suggests that catecholamines should boost temporal correlations regardless of their sign (Donner and Nieuwenhuis, 2013; Eldar et al., 2013). We therefore used the absolute correlation coefficient as our measure of connectivity strength. The signed and absolute matrices were very similar, because anti-correlations were rare (mean 3.4% of all connections, SD 3.5%), as is common when no global signal regression has been performed. In the group- and condition-averaged correlation matrix, 0.28% were anti-correlations (11 out of 4005 unique connections, Fig. 1b). To facilitate comparisons of values across participants, we range-normalized each participant's absolute correlation matrices between 0 and

1 across the 4 conditions. This procedure discarded the between-participant variance while leaving the spatial structure and between-condition variance intact.

In addition, for the post atomoxetine condition time-resolved connectivity (Allen et al., 2014) was computed for 189 tapered windows  $w$  of length  $nw$  (22 volumes) via:

$$R_{wi} = \frac{M_{wi}^T \cdot M_{wi}}{nw - 1}$$

The taper was created by convolving a Gaussian (SD 3 TRs) with a rectangle.  $R_{wi}$  was Fisher-transformed to stabilize variance across windows. We then again used the absolute value as our measure of connectivity strength. An identical sliding window was applied to the pupil diameter data in the post atomoxetine condition such that for each window in  $R_{wi}$  there was a corresponding value of pupil size during that window. Then, we divided up pupil size into 3 equal-sized bins, and averaged the corresponding values in  $R_{wi}$  for each pupil bin separately. To rule out the possibility that the results depended on the choice of bin size, we also tried alternate bin sizes (2, 5, and 7 bins) and found similar effects.

*Graph-theoretical analysis of global correlation structure.* For each condition, we constructed a binary undirected (adjacency) matrix  $A$ . We did this by first concatenating the correlation matrices across participants such that for each condition we had a brain region by brain region by  $N$  (90 by 90 by 24) matrix of connectivity. We then assessed with a  $t$ -test across the participant dimension for each element  $y, x$  in the connectivity matrix whether its value differed significantly from the average of its row  $y$  or column  $x$  (Hipp et al., 2012). In other words, for each connection we obtained a distribution across participants of weighted values, and two distributions corresponding to the mean weighted values of each brain region that was linked by that particular connection. The connection distribution was then compared to each of the brain region distributions with a  $t$ -test. If either of the two comparisons was significant, the connection was scored as 1, and otherwise it was scored as 0. The alpha level was set to 0.01, Bonferroni-corrected for two comparisons to 0.005 (Hipp et al., 2012).

Note that this procedure – as opposed to simply applying a fixed-percentage threshold – results in adjacency matrices that can differ in the number of connections between conditions, and therefore allows the assessment of correlation structure, or degree. We thus quantified the global degree  $k$  in each condition as the average across the adjacency matrix (Hipp et al., 2012) via:

$$k = n^{-1} \sum_{x=1}^n n^{-1} \sum_{y=1}^n A(x, y)$$

where  $n$  is the number of brain regions in the AAL atlas.

To test the prediction that increased catecholamine levels should result in stronger functional connectivity, we used  $k$  as our measure of connectivity strength rather than relying on the mean weighted values (i.e. the average of  $R_i$ ). The binarization of weighted graphs is common in functional network analysis (Achard and Bullmore,



2007; Rubinov and Sporns, 2010; Hipp et al., 2012; Li et al., 2013) and is intended to preserve only the strongest (most probable) connections. This ensures that weak edges, which are more likely to be spurious (Rubinov and Sporns, 2010), do not convolute the global mean. Given that these edges are less likely to reflect true neurophysiological interactions, they are less likely to be sensitive to any experimental manipulation that is specifically intended to alter neurophysiology (in our case, drug intake). Thus, excluding these connections decreases the likelihood of false negatives in between-condition comparisons of the global mean. In addition, by treating each connection equally (either present or absent) the global mean is not disproportionately influenced by extremely strong connections that are more likely to decrease in strength after an experimental manipulation by virtue of regression towards the mean.

Furthermore, by defining adjacency matrices using a statistical test across participants, each connection that is present in the adjacency matrix is ensured to be reliably expressed across the group of participants for a given condition. Thus, the adjacency matrices are representative of the group-level topography of connectivity. We used two measures of clustering, defined using these group-level adjacency matrices, to test the prediction that an increase in central catecholamine levels should be accompanied by more strongly clustered network connectivity. The clustering coefficient  $C$  was quantified as the average fraction of triangles  $\tau$  around a node, the latter given by:

$$\tau_x = 2^{-1} \sum A(x, y)A(x, z)A(y, z) \text{ , where } y, z \in N$$

and  $N$  represents the total set of nodes.  $C$  was then given by:

$$C = n^{-1} \sum_{y=1}^n n^{-1} \sum_{x=1}^n \frac{2\tau_x}{k_x(k_x - 1)}$$

The clustering coefficient here is equivalent to the average proportion of the node's neighbors that are in turn neighbors to each other (Watts and Strogatz, 1998; Rubinov and Sporns, 2010). Thus, the clustering coefficient represents the mean fraction of clustering around each node.

Because  $C$  is normalized by degree ( $k$ ) individually per node, it may be biased by nodes with a relatively low  $k$ . We therefore also included a measure of clustering that is normalized by  $k$  collectively and hence does not suffer from the same potential bias. This measure is known as transitivity ( $T$ ), and is given by:

$$T = \frac{\sum_{x=1}^n 2\tau_x}{\sum_{x=1}^n k_x(k_x - 1)}$$

Note that this is equivalent to the ratio of triangles to triplets in the network. Both clustering coefficient and transitivity capture the extent to which the network is segregated in terms of processing, because a large number of triangles implies

functional clustering. These two measures were computed using the Brain Connectivity Toolbox (Rubinov and Sporns, 2010). Note that both clustering and transitivity are (partially) dependent on global degree (van Wijk et al., 2010).

To test statistically whether degree, clustering coefficient and transitivity differed between conditions, we used non-parametric permutation testing. We shuffled the condition labels for each participant prior to computing the adjacency matrices and then computed the graph-theoretical measures. This was done for 10,000 iterations to produce a null distribution. We then derived a  $p$  value for each contrast by dividing the number of null observations more extreme than the observed contrast by the total number of null observations, and subtracting this value from 1.

*Network identification via community detection.* We used the Louvain method for community detection optimized for stability (Blondel et al., 2008; Le Martelot and Hankin, 2013) to classify each brain region as belonging to a particular network, or module. This method works by maximizing the number of within-group connections (edges) while minimizing the number of between-group connections via greedy optimization. We first defined an adjacency matrix  $A_s$  by concatenating the condition-averaged correlation matrices across participants, and then statistically comparing each element  $y,x$  to the average of its row  $y$  or column  $x$ , similar as described above. However, to accurately classify networks we needed to retain only those connections that were most informative about community structure. We therefore promoted sparsity in the condition-averaged adjacency matrix by defining it using a one-tailed  $t$ -test with a conventional alpha level (0.05) and a correction for multiple comparisons using the false discovery rate (FDR). This preserved only those connections that were consistently the strongest across participants (16.9% of all possible connections). We then submitted this sparse condition-averaged adjacency matrix to the Louvain community detection algorithm. The optimization procedure (Le Martelot and Hankin, 2013) ensured a stable solution across multiple runs of the algorithm. In the optimization procedure, the Markov time acts as a resolution parameter that determines the community scale, and thus the number of modules that the algorithm will return. This parameter was set to 0.9, resulting in 6 separate modules. We set the number of modules to be detected to 6 because, given the relatively coarse anatomical layout of the AAL atlas, this number yielded a relatively reliable modular organization. The community detection and optimization resulted in a 'module number' for each AAL brain region indicating to which module it belonged, and a single Q-value indicating the strength of modularity.

We first verified whether the Q-value was significantly higher than chance. To do so, we generated 10,000 randomized null networks with an identical size, density and degree distribution as  $A_s$  (Maslov and Sneppen, 2002), and submitted them to Louvain community detection and optimization to produce a null-distribution of Q-values. We then derived a  $p$ -value for the observed modularity by dividing the number of null Q-values more extreme than the empirical Q-value by the total number of null Q-values, and subtracting it from 1.

The observed Q-value of 0.46 was significantly higher than chance ( $p < 0.001$ ), showing that group-average connectivity was strongly modular. We then visualized the modular structure by re-arranging the condition-averaged correlation matrix by module. The assignment of brain regions to modules corresponded closely to a number of well-characterized intrinsic connectivity networks, indicating that the modular structure reflected a functionally meaningful grouping of brain regions.

*Graph-theoretical analysis of network structure.* The procedure described above allowed us to group brain regions into modules of intrinsically coupled AAL brain regions. We could then use these modules to assess changes in the structure of intrinsic correlations at the within- and between-network level, rather than as a function of the system in its entirety. To do this, we first re-arranged the condition-specific adjacency matrices by their module number, and computed average degree of elements within and between modules via:

$$k_m = n_a^{-1} \sum_{x_a=1}^{n_a} n_b^{-1} \sum_{y_b=1}^{n_b} A_s(x_a, y_b)$$

where  $n_a$  is the number of brain regions belonging to module  $a$  and  $n_b$  is the number of brain regions in module  $b$ . This yielded, for each condition, a symmetric and module-by-module matrix of continuous average degree values, in which values on the diagonal indicated the average number of connections within each module, and each value around the diagonal indicated the average number of connections between a combination of modules.

We could then use these ‘module matrices’ to test for atomoxetine-related changes in degree of the connections within modules, and the connections linking different modules. This allowed us to characterize changes in connectivity in a spatially more specific way than for global degree. We again used non-parametric permutation testing, similar as described for global degree, except that it was done for individual elements within the module matrices.

*Control analyses using an alternate atlas and multiple thresholds.* To rule out the possibility that our results were specific to the use of the AAL atlas, we repeated all of our key analyses using the atlas made available by Craddock et al. (2012), which comprised 87 distinct regions after excluding the cerebellum, and found similar effects in terms of both direction and significance. Moreover, to verify that our results were independent of the statistical threshold used to define the adjacency matrices, we conducted a control analysis in which a range of adjacency matrices was created per condition with varying condition-averaged connection densities (40-75%). This was done by progressively raising / lowering the alpha level of the  $t$ -test that was used to determine whether a connection is present or absent (see above). Then, for each threshold we computed the graph-theoretical measures, and for each condition and measure separately calculated the area under the curve (AUC) across thresholds. This allowed us to compare the AUC between conditions with permutation testing

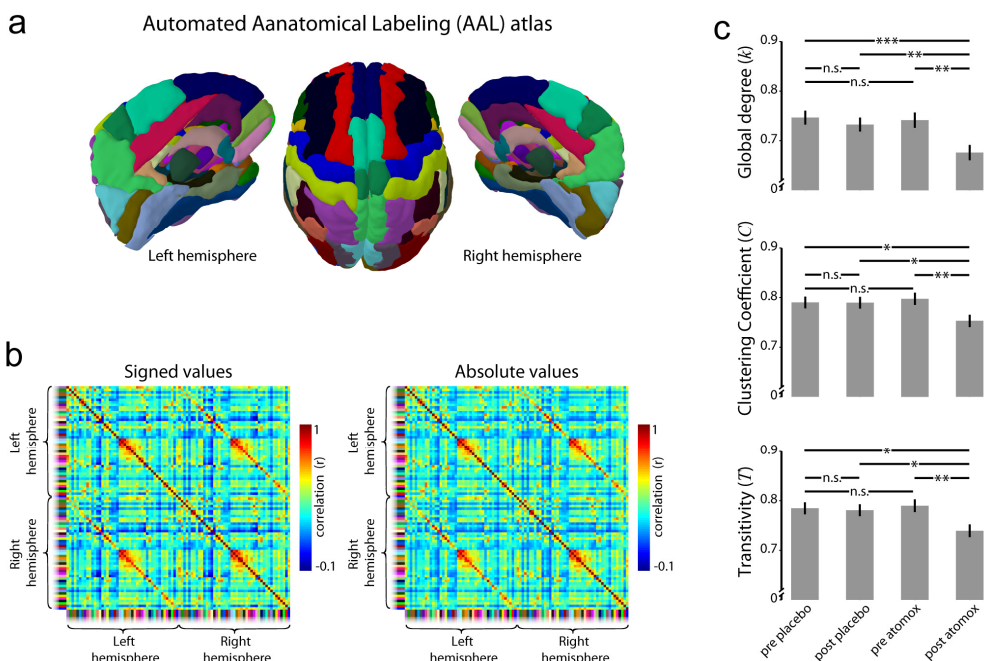
(10,000 iterations). For all measures the critical interaction contrast was significant and in the same direction as our original findings (see Results).

*Controlling for regression towards the mean.* The correlation between baseline coupling strength and the atomoxetine-related change in coupling strength (Fig. 3e, see Results) is confounded by regression towards the mean. That is, if two particular brain regions show strong baseline coupling, then simply by chance they are more likely to show a reduction under atomoxetine, and so a negative correlation is likely to occur. We therefore controlled for regression towards the mean using permutation testing. For 10,000 permutations we shuffled the condition labels across participants prior to computing the atomoxetine-related change in coupling strength. We then computed the correlation between baseline coupling and atomoxetine-related change in coupling to produce a distribution of correlation coefficients under the null hypothesis of regression towards the mean. Finally, we derived a  $p$  value for the empirical correlation coefficient by dividing the number of null observations more extreme than the correlation coefficient by the total number of null observations, and subtracting this value from 1. This  $p$  value indicated the significance of the observed correlation coefficient beyond regression towards the mean.

*Analysis of BOLD signal variance.* We calculated for each participant and each AAL brain region the fractional amplitude (i.e., variance) of low-frequency fluctuations in the non-Z-scored BOLD time series [fALFF (Zou et al., 2008)]. This measure indexes the relative contribution of low-frequency (0.01-0.08 Hz) fluctuations to the total amplitude spectrum. We compared fALFF between conditions using repeated measures analysis of variance. Additionally, for each participant we correlated the atomoxetine-related change in fALFF with the atomoxetine-related change in inter-regional correlation strength across AAL brain regions. We then compared the distribution of Fisher-transformed correlation coefficients to zero using a two-tailed  $t$ -test. Very similar results were obtained using alternative measures of variance (e.g., average 0.01-0.08 Hz amplitude or the signal standard deviation rather than fractional amplitude).

### 3.3 Results

*Atomoxetine reduces global degree and clustering.* In a first set of analyses, we examined the effect of atomoxetine on graph-theoretical summary measures of functional connectivity strength. We parcellated each participant's brain into 90 separate regions according to the Automated Anatomical Labeling (AAL) atlas [(Tzourio-Mazoyer et al., 2002), Fig. 1a] and computed the correlation between the Z-scored time series of all pairs of regions (Fig. 1b). We then took the absolute correlation coefficient as our measure of functional connectivity strength (see Materials and Methods). In general, functional connectivity was strongest between



**Figure 1.** Inter-regional correlation and global graph-theoretical results. **a)** Topography of the AAL atlas. Each brain region within hemispheres has a unique color. **b)** Condition-averaged inter-regional correlation. Both the signed and absolute values are shown. Color labels on the left and bottom axes correspond to brain regions in a). **c)** Atomoxetine effects on global graph-theoretical measures. Error bars represent the SD of the bootstrapped null-distribution. n.s.: non-significant; \*:  $p < 0.05$ ; \*\*:  $p < 0.01$ , \*\*\*:  $p < 0.001$ .

visual cortical areas and between homologue areas in both hemispheres (Fig. 1b), consistent with a host of previous work (Fox and Raichle, 2007).

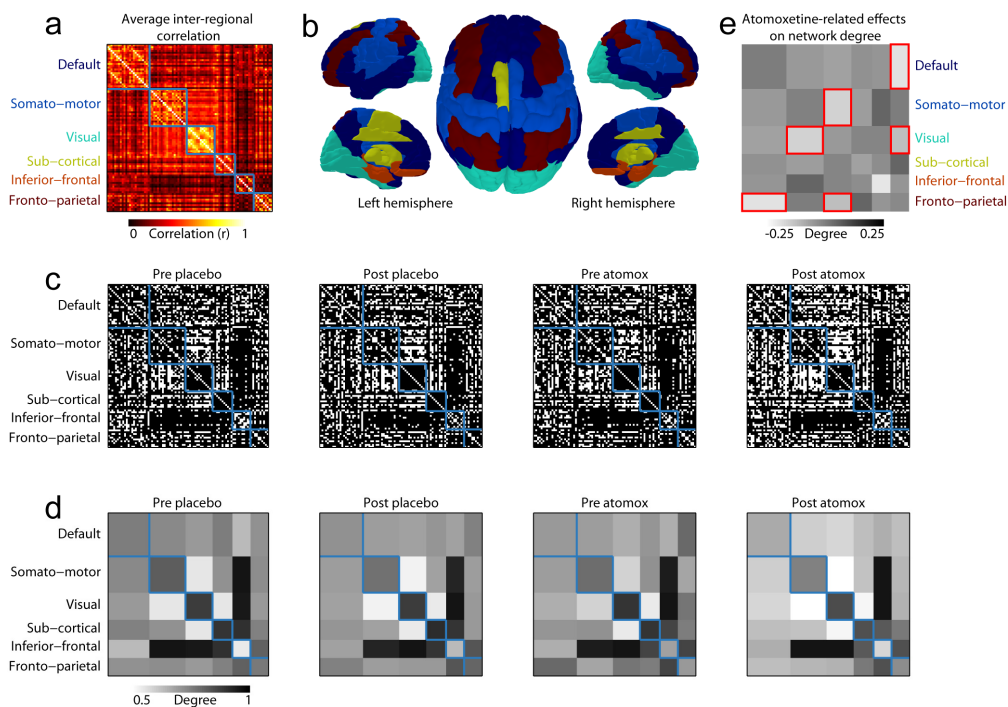
For each condition (pre placebo, post placebo, pre atomoxetine, post atomoxetine), we constructed a binary matrix of connections (*edges*) between pairs of brain regions that consistently differed in strength across participants from the average of other connections involving either of the two brain regions [c.f. (Hipp et al., 2012)]. Graph theory allowed us to capture different properties of these matrices of intrinsic correlations in a small number of diagnostic scalar quantities (Bullmore and Sporns, 2009; Rubinov and Sporns, 2010). Specifically, we assessed three such measures: the *global degree*, which indexes the number of strongly correlated regions (above a certain threshold; see Materials and Methods) in the network, and two descriptors of the extent to which network connectivity is clustered in segregated local groups of brain regions: *clustering coefficient* and *transitivity*, both of which are (partially) dependent on the strength of connectivity (van Wijk et al., 2010; Eldar et al., 2013). If catecholamines increase global functional connectivity, then atomoxetine should increase all three measures.

Figure 1c shows that atomoxetine significantly reduced the number of strong correlations present in the network, as indicated by lower global degree. This was reflected in a significant interaction between treatment and time ( $p = 0.039$ ). A similar pattern of results was found for the two measures of clustering, both of which decreased in magnitude (Fig. 1c): clustering coefficient ( $p = 0.043$ ) and transitivity ( $p = 0.048$ ). Thus, atomoxetine reduced the number of strongly correlated brain regions, as well as the extent to which correlated brain regions formed local functional ensembles. Together, these results show that atomoxetine decreases, rather than increases, overall inter-regional correlations in the brain at rest.

*Atomoxetine reduces inter-network degree.* Many studies of resting-state activity in humans have revealed a consistent set of groups or *modules* of brain regions that are characterized by strong coupling between brain regions belonging to the same module, and weaker coupling between brain regions belonging to different modules (Bullmore and Sporns, 2009). These modules are often referred to as “intrinsic functional connectivity networks” (Fox and Raichle, 2007). In a next set of analyses, we investigated atomoxetine-related changes in the strength of functional connectivity within and between these networks.

To do this we arranged the connectivity matrix by network [Fig. 2a; see (Blondel et al., 2008)]. This resulted in 6 functional networks that correspond closely to previously reported resting-state networks (Yeo et al., 2011; Zalesky et al., 2014). Based on their topography (Fig. 2b) we termed them: default; somato-motor; visual; sub-cortical; inferior-frontal; and fronto-parietal networks. We then calculated the average number of connections within and between these networks, resulting in a 6 by 6 *network degree* matrix for each condition (Fig. 2c,d). Finally, we examined atomoxetine-related changes in within- and between-network degree using permutation testing. This allowed us to explore if changes in functional connectivity occurred in intra- or inter-network connections. Note that the atomoxetine-related reduction in global degree (Fig. 1c) is visible in the network degree matrices as an overall increase in ‘brightness’ in the right panel of Fig. 2d. Consistent with the decrease in global degree reported above, we observed only atomoxetine-related *reductions* in network degree (Fig. 2e). The interaction between treatment and time was significant for the connections between the visual and somato-motor networks ( $p < 0.001$ ); between the visual and fronto-parietal networks ( $p = 0.044$ ); and between the fronto-parietal and default networks ( $p < 0.001$ ). After using the FDR ( $q = 0.05$ ) to correct for multiple comparisons, all connections except the connection between the visual and fronto-parietal networks remained significant. However, when comparing the area under the curve across a range of thresholds (see Materials and Methods), all connections remained significant after FDR correction.


Interestingly, all significant reductions in network degree were in connections between (as opposed to within) functional networks. Thus, the most robust decreases in functional coupling occurred for connections linking functionally dissociable groups of brain regions. These results corroborate the conclusion drawn above, that



**Figure 2.** Intrinsic connectivity networks and changes in graph-theoretical measures of network structure. **a)** Condition-averaged inter-regional correlation arranged by network. The networks are outlined in blue. **b)** Topography of functional networks. Colors correspond to the labels in a). **c)** Condition-specific adjacency matrices arranged by network. Black elements indicate that a connection is present. **d)** Average degree for within- and between-network connections. To facilitate visual comparison, the size of each network is the same as in c). Note, however, that all statistical comparisons were conducted on symmetrically sized matrices in which each network contributed equally to the global mean. **e)** Atomoxetine-induced changes in degree for connections within and between networks. Shades of grey reflect the value of the interaction contrast (post atomoxetine - pre atomoxetine) minus (post placebo - pre placebo). Significant ( $p < 0.05$ ) changes in degree are outlined in red.

atomoxetine decreased inter-regional correlations in the brain at rest. More importantly, these results provide a first indication that this reduction in inter-regional correlations is not spatially homogeneous across the brain. In the following, we further characterize the spatial heterogeneity of the atomoxetine-induced reductions in inter-regional correlations.

*Regionally-specific reductions and baseline-dependent changes in connectivity with atomoxetine.* Having assessed the topographical changes induced by atomoxetine at the global level and the level of functional networks, we next assessed differences in the strength of inter-regional correlations at the level of individual connections between brain regions, using the absolute correlation coefficient. We



**Figure 3.** Atomoxetine-related effects on inter-regional correlation. **a)** Region-by-region matrix of atomoxetine-related changes in inter-regional correlation strength. Colors represent the value of the interaction contrast (post atomoxetine - pre atomoxetine) minus (post placebo - pre placebo). Blue colors indicate reduced correlation following atomoxetine. The matrices are organized following Figure 1b. **b)** Atomoxetine-related effect on the absolute inter-regional correlation coefficient, rendered in 3D with an arbitrary threshold applied. The threshold is indicated by the white dashes in the color bar. Spheres are placed in the center of mass of their respective AAL atlas regions. Both the size and color indicate the average atomoxetine-related effect on coupling [i.e the average across rows or columns in a)]. **c)** Transverse (top is anterior) and sagittal (right is anterior) view on 3D rendering of significant correlation changes, resulting from the whole-brain two-step analysis. Individual connections that changed significantly with atomoxetine are plotted as cylinders between the corresponding regions. **d)** Inter-regional correlation in each condition, averaged across the significant connections shown in c). Error bars represent the SEM. **e)** Correlation between baseline inter-regional correlation strength (collapsed across pre placebo and pre atomoxetine) and change with atomoxetine. Each dot represents a unique region-by-region connection. Self-connections were excluded.

found that atomoxetine altered correlation strength in a strikingly structured fashion (Fig. 3a): In general, connectivity was reduced by atomoxetine, especially in posterior brain regions (Fig. 3b). These observations align with our findings of reductions in inter-network degree involving the visual system. To quantify these effects, we used a two-step procedure. Specifically, we first derived a set of data-driven hypotheses by identifying, in the first half of the fMRI volumes, the limited number of individual connections that exhibited an atomoxetine-related change in connectivity that was reliable across participants ( $p < 0.05$ , using a two-tailed  $t$ -test), thereby reducing the number of comparisons for the subsequent step. We then re-tested those connections




using the (independent) second half of the volumes, and selected those that again showed a systematic atomoxetine-related change in correlation strength ( $p < 0.005$ , two-tailed). Atomoxetine significantly lowered correlation strength in a cluster of occipital brain regions (Fig. 3c,d), specifically correlations between left calcarine cortex and right calcarine cortex / bilateral lingual gyrus; between left cuneus and right calcarine cortex / lingual gyrus; between left lingual gyrus and right calcarine cortex / lingual gyrus; and between right lingual gyrus and right calcarine cortex / right fusiform gyrus. These contiguous connections remained significant after applying a highly conservative cluster size threshold ( $p < 0.0001$ ), obtained by generating a distribution of maximum cluster sizes under the null hypothesis with permutation testing (Nichols and Holmes, 2001). Thus, the cluster involved significantly more connections than would be expected by chance.

Note that we did not find significant changes in connectivity between structures of the basal ganglia, which have been widely studied in relation to catecholaminergic drug effects (Sulzer et al., 2016). This lack of an atomoxetine-related effect in the human basal ganglia is consistent with the observation that the basal ganglia receive relatively sparse noradrenergic innervation (Aston-Jones et al., 1984), and with findings that atomoxetine has little effect on DA levels within the basal ganglia of rodents (Bymaster et al., 2002). Indeed, unlike in the cortex, in the basal ganglia there is an abundance of DA transporter (Sulzer et al., 2016), so DA reuptake is not dependent on the NE transporter. Thus, our finding that atomoxetine reduced the strength of inter-regional correlations in (predominantly visual) cortical areas is consistent with the specific effect of atomoxetine on synaptic catecholamine levels within the cortex.

In sum, atomoxetine lowered the strength of correlations between visual cortical brain regions (Fig. 3c), regions that on average showed strong connectivity (Fig. 1b, 2a). This raises the question whether the atomoxetine-induced change in connectivity was dependent on the baseline level of connectivity between any pair of brain regions. To address this question we correlated weighted coupling strength collapsed across the pre placebo and pre atomoxetine conditions (i.e., baseline connectivity) with the coupling change following atomoxetine (controlling for regression towards the mean with permutation testing). The observed correlation between baseline connectivity strength and the change with atomoxetine was significant ( $r = -0.22$ ,  $p = 0.029$ , Fig. 3e). This indicates that the strongest functional connections tended to show the largest connectivity reductions after atomoxetine, and vice versa.

*Atomoxetine induces de-coupling of early visual cortex from the rest of the brain.*

As noted above, atomoxetine reduced correlations between occipital brain regions. To establish whether these occipital regions also showed reduced coupling to the rest of the brain, as suggested by visual inspection (Fig. 3b), we computed a summary statistic (median) of correlation strength between early visual (pericalcarine) cortex and all other AAL atlas regions. There were no differences between left and right early visual cortex, so we collapsed the data across hemispheres.




**Figure 4.** Atomoxetine reduces correlation strength between early visual (pericalcarine) cortex and the rest of the brain. **a)** Topography of condition-averaged correlation seeded from left and right early visual cortex. **b)** Topography of atomoxetine-related effects on correlation seeded from left and right early visual cortex. Colors represent the value of the interaction (post atomoxetine - pre atomoxetine) minus (post placebo - pre placebo). **c)** Median correlation values across the brain seeded from left and right early visual cortex. Error bars represent the SEM.

Average connectivity seeded from early visual cortex is shown in Figure 4a. Atomoxetine reduced connectivity between early visual cortex and the rest of the brain as reflected by a significant interaction between treatment and time ( $F(1,23) = 5.31$ ,  $p = 0.031$ , Fig. 4b,c). The only significant pairwise comparison was post atomoxetine versus pre atomoxetine ( $t(23) = 2.34$ ,  $p = 0.028$ ). Together these results suggest that the early visual cortical areas not only de-coupled from each other following atomoxetine (Fig. 3c) but also from the rest of the brain.

The results of our analyses at the level of individual connections between brain regions converge with those at the global level and at the level of networks of brain regions, showing that atomoxetine decreased functional connectivity. In addition, the results show that atomoxetine modulated functional connectivity in a highly regionally specific fashion, with more robust changes in visual cortex than in other brain areas. *Excluding alternative explanations.* In five sets of control analyses, we ruled out the possibility that the atomoxetine-related changes in inter-regional correlations were driven either by local changes in BOLD variance, by retinal effects due to pupil dilation associated with atomoxetine, by head motion, by saccade-related retinal transients, or by atomoxetine-induced changes in physiology (heart rate and breath rate). First, the correlation coefficient between two signals is their covariance normalized by the signals' variances. Thus, it is possible that the observed changes in inter-regional correlations are caused by local changes in variance alone (Haynes et al., 2005; Freeman et al., 2011), rather than by changes in covariance (i.e., the degree to which the BOLD signals in two regions fluctuated together). If this is the case, then the atomoxetine-related change in average inter-regional correlation and the atomoxetine-related change in BOLD signal variance should be negatively correlated across brain regions.

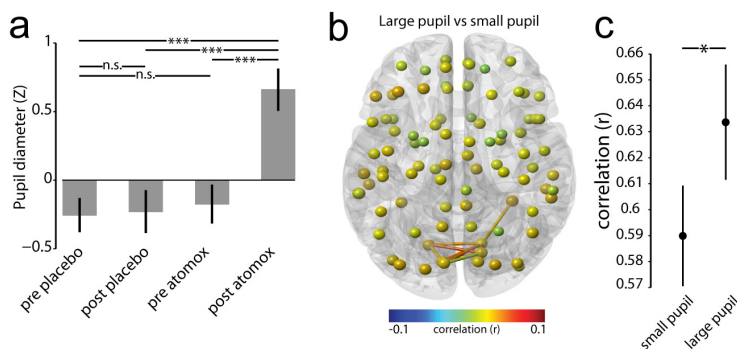
**Figure 5.** Spectral BOLD characteristics and the relation with inter-regional correlations. (Left) atomoxetine-induced changes in spectral amplitude for AAL brain regions that showed an atomoxetine-induced increase (red) and decrease (blue) in inter-regional correlation strength. Brain-region and condition-averaged amplitude is shown in black. (Right) Mean correlation between the region-averaged atomoxetine-induced change in coupling strength and fractional amplitude of low-frequency BOLD fluctuations. Error bars show the SEM. \*\*:  $p < 0.01$ .



Instead, we found a positive relationship between changes in inter-regional correlation and changes in BOLD variance, which was consistent across participants ( $t(23) = 3.36$ ,  $p = 0.003$ , Fig. 5), ruling out variance as a confound. Moreover, there was no interaction between treatment and time in overall BOLD variance ( $F(1,23) = 0.71$ ,  $p = 0.40$ ), or in variance for only the occipital brain regions that showed reduced atomoxetine-related inter-regional correlation ( $F(1,23) = 0.41$ ,  $p = 0.53$ ).

Second, because atomoxetine increased the size of the pupil (Fig. 6a), it is conceivable that this peripheral effect, rather than the effect of atomoxetine on central catecholamine levels, was driving the changes in inter-regional correlation in visual cortex (Haynes et al., 2004). To examine this potential confound, we binned inter-regional correlation by pupil size in the post atomoxetine condition, focusing on those correlations that showed a significant reduction under atomoxetine. If larger pupil size is responsible for the reduction in correlations, then time

**Figure 6. a)** Atomoxetine effect on pupil diameter. **b,c)** Correlation strength in the post atomoxetine condition binned by pupil size, only for connections that showed an atomoxetine-related reduction in inter-regional correlation. Error bars represent the SEM. n.s.: non-significant; \*:  $p < 0.05$ .



periods during which the pupil is large should be associated with weaker correlations than time periods during which the pupil is small. Interestingly, we found the opposite pattern (Fig. 6b, c): stronger correlations for large pupil ( $t(18) = 2.84$ ,  $p = 0.010$ ), ruling out an interpretation in terms of pupil size.

Third, head motion can have a strong influence on the strength of inter-regional correlations (Van Dijk et al., 2012). To rule out the possibility that our key finding of atomoxetine-related changes in inter-regional correlation was driven by head motion, we first compared head motion between conditions. Neither mean head motion nor mean absolute head motion differed between conditions (all  $ps > 0.05$ ). No participant's head motion exceeded 2 mm, indicating that overall there was little head motion. However, general mild head motion tends to increase correlations between proximate areas and decrease connectivity between distant areas (Van Dijk et al., 2012). Thus head motion can potentially lead to spatially heterogeneous effects on connectivity in a manner that is related to the distance between brain areas. To rule out the possibility that the spatial structure of atomoxetine-related changes in connectivity was driven by subtle (non-significant) differences in head motion between conditions, we correlated Euclidean distance between the center of mass of each pair of AAL brain areas and the strength of functional connectivity between those areas, for each participant and each condition. We then compared the distribution of Fisher-transformed correlation coefficients between conditions. If head motion is responsible for the observed change in connectivity between conditions, then the correlation between Euclidean distance and strength of connectivity should also differ between conditions. However, we did not find any differences between conditions (all  $ps > 0.05$ ), ruling out head motion as a confound.

Fourth, it is possible that the atomoxetine-related reduction in the strength of correlation between visual cortical areas occurred due to differences between conditions in saccade-related retinal transients. In order to rule out this possibility, we extracted several eye movement metrics from the eye tracker gaze position data using the EYE-EEG toolbox (Dimigen et al., 2011). There was no interaction between treatment and time for any of the metrics: the number of saccades ( $F(1,18) = 0.47$ ,  $p = 0.50$ ), median saccade amplitude ( $F(1,18) = 0.45$ ,  $p = 0.51$ ), median saccade duration ( $F(1,18) = 0.11$ ,  $p = 0.74$ ), or median saccade peak velocity ( $F(1,18) = 3.32$ ,  $p = 0.085$ ). This latter trend was driven by a numeric difference between the pre placebo and post placebo conditions. Pre and post atomoxetine did not differ significantly in saccade peak velocity ( $t(18) = -0.43$ ,  $p = 0.67$ ). Together, these results show that our key result of an atomoxetine-related reduction in the strength of correlation between visual cortical regions was unlikely to be driven by saccade-related retinal transients.

Finally, atomoxetine significantly increased breath rate ( $F(1,23) = 8.96$ ,  $p = 0.007$ ) and heart rate ( $F(1,23) = 4.66$ ,  $p = 0.041$ ), as reflected by a significant interactions between treatment and time. We therefore corrected the BOLD time series using the RETROICOR method [see Materials and Methods; (Glover et al., 2000)]. The average

$R^2$  of the physiology regressors was relatively low (0.034), indicating that physiology accounted for a small proportion of the total BOLD variance [which was likely the result of artifact removal by FMRIB's ICA-based X-noiseifier (Griffanti et al., 2014; Salimi-Khorshidi et al., 2014)]. Nevertheless, to conclusively rule out atomoxetine-related changes in physiology as confounds, we repeated the key analyses on the physiology-corrected data. All three global graph-theoretical measures remained significant and in the same direction as reported above (all  $p$ 's < 0.05). We also found significant reductions in network degree in the same inter-network connections (all  $p$ 's < 0.05). Lastly, we observed a similar contiguous cluster of significantly reduced inter-regional correlations within visual cortex (all  $p$ 's < 0.005, and cluster-corrected at  $p$  < 0.0001). Thus, our key results were unlikely to be driven by atomoxetine-related changes in physiology.

### 3.4 Discussion

Using a pharmacological manipulation, we examined the effects of increased extracellular levels of the catecholamines NE and DA on resting-state fMRI connectivity in the human brain. First, we found that our manipulation reduced the strength of inter-regional correlations across three levels of spatial organization, indicating that catecholamines reduce the strength of functional interactions during rest. Second, this modulatory effect on the structure of resting-state correlations exhibited a substantial degree of spatial specificity, indicating that catecholamines differentially reduce spontaneous correlations between select brain regions. These two key findings are surprising in light of the common understanding of the neurophysiology and computational function of catecholaminergic systems. They also identify catecholaminergic neuromodulation as an important factor shaping the spatial structure and strength of intrinsic functional connectivity in the human brain.

Our first key finding is that atomoxetine, a selective NET blocker that increases synaptic NE and DA levels (Bymaster et al., 2002; Devoto et al., 2004; Invernizzi and Garattini, 2004; Swanson et al., 2006; Koda et al., 2010), reduced the strength of inter-regional correlations. Specifically, atomoxetine reduced the strength of connectivity globally (Fig. 1c), between nodes belonging to distinct intrinsic connectivity networks (Fig. 2e), and between individual brain regions within the visual system (Fig. 3c). This consistent pattern of results seems to be at odds with the notion of a facilitative effect of catecholamines on brain-wide signal transmission (Aston-Jones and Cohen, 2005; Eldar et al., 2013). One possible explanation for this discrepancy lies in the fact that in our experiment participants did not actively respond to incoming sensory information. According to a recent theory, the effects of NE on neural activity strongly depend on interactions with local glutamate release (Mather et al., 2015). Accordingly, enhanced NE may have qualitatively different effects during task processing, associated with relatively high glutamate activity, than during states of relative cortical

quintessence (i.e., at rest), associated with relatively low glutamate activity. In line with this possibility, Coull *et al* showed dissociable effects of the  $\alpha_2$  adrenergic agonist clonidine on positron emission tomography effective connectivity obtained during task performance and during rest (Coull *et al.*, 1999). Whereas during task performance clonidine increased connectivity between frontal and parietal cortical regions, during rest clonidine reduced connectivity from frontal cortex to thalamus, and in connections to and from visual cortex. Two other studies that employed NE drugs also provided evidence for regional reductions in connectivity strength during rest (McCabe and Mishor, 2011; Metzger *et al.*, 2015). These studies, however, only used a small number of seed regions to assess connectivity, and hence did not examine large-scale topographical changes.

Our second key finding is that atomoxetine resulted in spatially heterogeneous changes in inter-regional correlations. For example, atomoxetine caused a reduction in the number of strongly correlated brain regions between (but not within) distinct resting-state networks (Fig. 2e). Furthermore, the effect of atomoxetine on inter-regional correlations was dependent on the baseline level of coupling: the strongest functional connections tended to show the largest connectivity reductions after atomoxetine (Fig. 3d). How can such spatially structured effects of catecholamines come about? First, recent anatomical tracing work has suggested that the projection profile of the LC is more heterogeneous than once thought (Schwarz and Luo, 2015). For example, even though on the whole there is broad collateralization within the LC-NE system, sub-populations of LC neurons selectively innervate distinct brain regions (Chandler *et al.*, 2014; Schwarz *et al.*, 2015). Moreover, sub-populations of LC neurons that differ in their afferent projection profile also show marked differences in their firing characteristics (Chandler *et al.*, 2014). The firing modes of LC neurons in turn have differentiable effects on neuronal synchronization within the cortex (Safaai *et al.*, 2015). Importantly, LC neurons have been reported to co-release DA (Devoto and Flore, 2006). Thus, spatially selective effects of catecholamines on correlated fluctuations in the brain can be achieved via a heterogeneous cortical innervation by the LC.

Second, heterogeneity in the effect of catecholamines on inter-regional correlations could be achieved by regional differences in the expression of different receptor types. For example, expression of the  $\alpha_2$  receptor roughly follows an anterior to posterior gradient (Nahimi *et al.*, 2015), with particularly strong expression in primary visual cortex (Zilles and Amunts, 2009). Interestingly, we observed an anterior to posterior gradient in the effect of atomoxetine on the strength of correlations (Fig. 3b). Moreover, we found a pronounced reduction in the strength of correlations between regions within visual cortex, and between early visual cortex and the rest of the brain (Fig. 4). The similarity between the spatial distributions of  $\alpha_2$  receptors and the effects of atomoxetine thus warrants further investigation into the relationship between specific NE receptor types and their influence on correlated activity across the brain.

A number of limitations of the present study should be acknowledged. First, we examined the effects of only one dose (40 mg) of atomoxetine. Dose-dependent pharmacological effects of catecholaminergic drugs on neural function are not uncommon (Berridge and Waterhouse, 2003). Future work on the neurochemical basis of functional connectivity will need to examine dose-dependent effects of atomoxetine, and other catecholaminergic drugs, with different pharmacokinetic profiles. Second, we do not know whether atomoxetine would have similar effects on functional connectivity in clinical populations characterized by disturbed catecholaminergic function (e.g., attention deficit hyperactivity disorder and depression). Third, although we used BOLD activity as a proxy for neural activity, the link between neuronal interactions and BOLD activity is not entirely clear (Logothetis, 2008). Models of catecholamine function make predictions about how NE and DA should affect neural communication. However, the translation of these predictions to BOLD correlations is not straightforward. Lastly, we used an atlas-based brain parcellation to investigate inter-regional correlations. Thus, the spatial resolution of our analyses was restricted by the resolution of the atlas. Future work, using voxel-level approaches, is needed to investigate more fine-grained spatial effects of catecholamine levels on functional connectivity.

The synaptic effects of catecholamines have been relatively well charted (Berridge and Waterhouse, 2003; Winterer and Weinberger, 2004). However, there is considerable uncertainty about how these low-level effects translate to system-wide functional interactions. Recently, a study by Safaai *et al* provided an important first glimpse into how the LC-NE system modulates spontaneous cortical activity and how this modulation in turn affects sensory processing in anesthetized rats (Safaai et al., 2015). Specifically, they showed that LC bursts can both attenuate and enhance processing of sensory stimuli depending on their timing relative to the stimulus and the cortical activity state. However, the effects of catecholamines on the large-scale communication *between* distant brain areas and their neurophysiological underpinnings remain exceedingly unexplored. Our finding that atomoxetine reduced inter-regional correlations in a spatially structured manner thus calls for novel work on the neural mechanisms that produce such effects.

Theory and evidence indicate that the topography of intrinsic fMRI correlations is dictated to an important extent by the fixed anatomical connectivity of each brain region (Deco et al., 2011; Deco et al., 2013). That is, brain regions that are anatomically strongly connected are more likely to show strong functional coupling than those that are connected weakly or only indirectly. However, within the constraints of physical connectivity there is substantial room for state-dependent movement in functional topological space (Allen et al., 2014; Zalesky et al., 2014; Barttfeld et al., 2015). Our results identify NE and DA as important factors driving these movements, and thus suggest that spontaneous fluctuations of catecholamine levels can serve to flexibly alter the structure of spontaneous correlations both globally and in specific brain regions, around the anatomical backbone.

