



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

## 4. Catecholamines Modulate Intrinsic Long-Range Correlations in the Human Brain

### Abstract

Brain activity fluctuates intrinsically, even in the absence of changes in sensory input and motor output. These fluctuations are correlated across large-scale networks of brain regions, and their strength and topography changes dynamically. Such dynamic changes in functional connectivity may be induced by brainstem neuromodulatory systems: in particular the locus coeruleus, which projects widely to the forebrain where it co-releases the catecholamines norepinephrine and dopamine. In the current study we examined whether catecholamines change the strength or the spatial structure (topology) of intrinsic long-range correlations, or both. Using a double-blind placebo-controlled crossover design, we pharmacologically increased central catecholamine levels in healthy human participants by administering atomoxetine. We used two complementary analysis approaches to examine the effect of catecholamines on fine-grained strength and topology of intrinsic functional connectivity patterns: ‘dual regression’ and ‘spatial mode decomposition’. Both approaches provided converging evidence for an atomoxetine-related reduction in correlation strength between distributed brain regions. Importantly, the pre-dominant effects of the drug were quantitative changes of correlations within existing functional networks that left the spatial structure of these networks intact, rather than reconfigurations of the topology of these networks. We conclude that catecholamines modulate dynamic changes in the strength of intrinsic inter-regional correlations.

*This chapter is based on:*

van den Brink RL, Rombouts SARB, Donner TH, and Nieuwenhuis S (*under review*). Catecholamines Modulate Intrinsic Long-Range Correlations in the Human Brain

## 4.1 Introduction

In the absence of changes in sensory input and motor output, brain activity fluctuates in intrinsically organized correlated networks (Biswal et al., 1995; Fox and Raichle, 2007). The strength and spatial structure of these intrinsic correlations predict task-based brain activation (Cole et al., 2016; Tavor et al., 2016), behavior (De Luca et al., 2005; Seeley et al., 2007), and are useful to study neural dysfunction in clinical populations (Greicius et al., 2004; De Luca et al., 2005; Rombouts et al., 2005; Di Perri et al., 2016). The topology of this ‘functional connectivity’ is constrained by the (largely) fixed structural connectivity between brain regions (Deco et al., 2011; Deco et al., 2013), which determines the anatomical backbone along which functional connectivity patterns can change dynamically (Allen et al., 2014; Zalesky et al., 2014; Barttfeld et al., 2015). It has been proposed that such changes in functional connectivity patterns may be induced by brainstem neuromodulatory systems (Leopold et al., 2003; Drew et al., 2008; Schölvinck et al., 2010). An important example is the locus coeruleus, which sends diffuse, ascending projections to the forebrain, where noradrenergic terminals release the catecholamines norepinephrine (NE) and dopamine (DA) (Devoto and Flore, 2006).

Several lines of evidence suggest that catecholamines might shape intrinsic correlations in activity between brain regions, possibly in diverse ways, changing either the strength or the topology of these correlation patterns, or both. First, at the single-cell level, catecholamines enhance neuronal responses to excitatory synaptic input (Rogawski and Aghajanian, 1980; Seamans et al., 2001b; Wang and O'Donnell, 2001) and can amplify GABAergic inhibition (Moises et al., 1979; Seamans et al., 2001a). Such enhanced synaptic efficacy results in system-level signal amplification (Berridge and Waterhouse, 2003). Second, direct pharmacological manipulations of synaptic catecholamine levels have been shown to alter the global strength of inter-regional co-fluctuations (Guedj et al., 2016; van den Brink et al., 2016; Warren et al., 2016). Third, evidence from small-scale circuits in crustaceans suggests that (catecholaminergic) neuromodulation can dynamically reconfigure functional networks, despite a constant structural connectome (Marder, 2012; Bargmann and Marder, 2013; Marder et al., 2014). Analogous effects have been suggested to underlie fast “resets” of brain network dynamics in the mammalian brain (Bouret and Sara, 2005), but direct evidence for catecholaminergic reconfiguration of cortical networks in humans has been lacking so far.

Here, we investigated catecholaminergic modulations of large-scale patterns of intrinsic fMRI signal correlations in the human brain at ‘rest’. To this end, we re-analyzed data from a double-blind placebo-controlled crossover study (van den Brink et al., 2016) of central catecholamine effects using atomoxetine, a selective NE transporter blocker. Our previous study quantified atomoxetine-induced modulations of the global strength of intrinsic correlations (van den Brink et al., 2016). Here, by contrast, we examined finer-grained patterns of intrinsic correlations, in order to test

for possible atomoxetine-induced quantitative changes in existing correlation patterns versus reconfiguration of correlation patterns.

We used two complementary analysis approaches. The first approach, known as ‘dual regression’, has been widely used to study the effects of pharmacological manipulations on fMRI signal correlations during the resting state (Beckmann, 2009; Filippini et al., 2009). The dual regression approach first delineates patterns of intrinsically correlated brain regions, so-called functional networks, and then compares voxel-level cofluctuation strength with those networks between conditions or groups. This method has proven useful for elucidating pharmacological effects on fMRI functional connectivity (Chamberlain et al., 2007; Klumpers et al., 2012; Cole et al., 2013; Klaassens et al., 2015; Guedj et al., 2016; Schranter et al., 2016; Klaassens et al., 2017). The second approach was linear decomposition of intrinsic signal correlation matrices into so-called ‘spatial modes’, again constituting patterns (or ‘networks’) of cofluctuations in brain activity (Mitra and Pesaran, 1999; Friston and Büchel, 2004; Donner et al., 2013). We used a generalization of the spatial mode decomposition that, different from dual regression, directly delineated networks showing the strongest drug-related changes in correlations, without prior selection of certain candidate networks (Friston & Büchel, 2004; Donner et al., 2013).

Both approaches provided converging evidence for the notion that catecholamines reduce the strength of fine-grained cofluctuation between diverse brain regions (including sensory- and motor-related networks). In addition, spatial mode decomposition revealed an atomoxetine-related shift from left to right-lateralized frontoparietal dominance in cofluctuation strength. Importantly, the predominant changes of correlation patterns we detected, all reflected quantitative changes in existing correlations, rather than a qualitative reconfiguration of network topology.

## 4.2 Materials and Methods

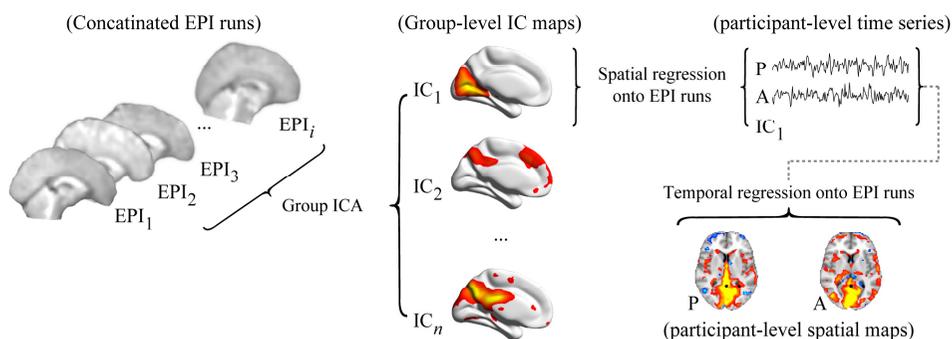
*Design and MRI preprocessing.* We reanalyzed data from van den Brink et al. (2016). This dataset comprised eyes open ‘resting-state’ (blank fixation) fMRI scans of 28 participants who received either placebo or atomoxetine (40 mg) on two separate sessions, scheduled one week apart. Atomoxetine is a selective NE transporter blocker that increases synaptic catecholamine levels (Bymaster et al., 2002; Devoto et al., 2004; Swanson et al., 2006; Koda et al., 2010). The study had a double-blind placebo-controlled crossover design, and was approved by the Leiden University Medical Ethics Committee. All participants gave written informed consent before the experiment, in accordance with the declaration of Helsinki.

Salivary markers of central catecholamine levels confirmed drug uptake (Warren et al., 2017). A full description of scan parameters and preprocessing details can be found in van den Brink et al. (2016). In brief, we applied the following preprocessing steps to the fMRI data (TR = 2.2 s; voxel size = 2.75 mm isotropic): realignment and

motion correction; B0 unwarping; high-pass filtering at 100 s; prewhitening; smoothing at 5 mm FWHM; coregistration of the functional scans with an anatomical T1 scan to 2 mm isotropic MNI space; artifact removal using FMRIB's ICA-based X-noiseifier (Griffanti et al., 2014; Salimi-Khorshidi et al., 2014); and retrospective image correction to account for differences in heart and breath rate between the atomoxetine and placebo conditions (Glover et al., 2000). In the current article, we focus on the runs following atomoxetine / placebo ingestion.

**Dual regression analysis.** We estimated a set of independent components (ICs) that were representative of the combined set of resting-state runs (i.e., runs from all participants and both the atomoxetine and placebo conditions) by applying a spatial independent component analysis (ICA) to all temporally concatenated data using FSL's MELODIC. The number of ICs to be detected (51) was automatically estimated from the data. Each IC represented a statistical parametric map and corresponding time series of consistent spatio-temporal dynamics. Next, we spatially correlated each IC spatial map with the 10 intrinsic connectivity networks reported by Smith et al. (2009) and selected the ICs that showed the highest correlation coefficient. The selected components showed an average correlation coefficient of 0.48 (range: 0.28 - 0.70), which indicated that the ICs as expressed in our data corresponded relatively well to previously reported intrinsic connectivity networks (Smith et al., 2009).

The 10 selected ICs represented spatial maps of ICs that were reliably expressed across the combined set of resting-state runs. They were thus representative of group-level spatiotemporal dynamics, but did not necessarily represent spatiotemporal dynamics within individual runs. To produce a time series and a spatial map for the individual resting-state runs and for each IC, we used dual regression (Beckmann, 2009; Filippini et al., 2009). Figure 1 shows a schematic overview of this analysis



**Figure 1.** Schematic overview of the dual regression method. First, a group-level independent component analysis is run to produce spatial maps. A selection of these maps is subsequently regressed onto the individual runs to produce participant- and component-specific time series. Finally, these time series are used in temporal regression to produce participant- and component-specific spatial maps, which can then be compared between conditions. P: placebo; A: atomoxetine; EPI: echo planar imaging; IC: independent component

approach. In a first step, we used the group-level IC spatial maps in multiple spatial regression onto the individual runs. This produced a time series for each IC as expressed within the individual runs. Then, in a second step, we used the participant-level IC time series as temporal regressors to produce spatial maps of regression coefficients for each IC and each run. Thus, this two-stage regression approach resulted in a spatial map for each participant, condition, and IC, that indicated the degree of covariation between individual voxels and the IC time series.

Finally, we collected the IC spatial maps of the individual runs into single 4D files (one per condition, per IC). This allowed us to compare these spatial maps to zero across participants to examine which brain regions cofluctuated with the IC time series, and compare them between conditions to assess which voxels displayed changes in cofluctuation strength with the IC, using non-parametric permutation testing (10,000 iterations) as implemented in FSL's Randomise. The  $\alpha$  level was set at 0.05, family-wise error (FWE) corrected for multiple comparisons using threshold-free cluster enhancement.

*Brain parcellation and inter-regional covariance analysis.* We extracted the fMRI time series of individual brain regions using the Automated Anatomical Labeling (AAL; Tzourio-Mazoyer et al., 2002) atlas, which contained 90 regions (cf. van den Brink et al., 2016). In addition, we used an alternate atlas that was based on a functional parcellation (Craddock et al., 2012). This atlas contained 140 individual brain regions. After averaging across voxels within each brain region (for each atlas separately), we Z-scored the multivariate time series ( $M$ , with dimensionality imaging volumes by brain regions) for each run  $i$  and then computed the group-averaged covariance matrices ( $C$ ) for the placebo and atomoxetine conditions (subscript P and A, respectively) via the following:

$$(1) C_P = N^{-1} \sum_{iP=1}^N \frac{M_{iP}^T \cdot M_{iP}}{nTR - 1}, C_A = N^{-1} \sum_{iA=1}^N \frac{M_{iA}^T \cdot M_{iA}}{nTR - 1}$$

where  $nTR$  is the number of volumes (211),  $N$  is the number of participants (24), and  $T$  denotes a matrix transposition. The matrices  $C_P$  and  $C_A$  represented the covariance between the BOLD time series of all brain regions, averaged across participants. Note that by variance normalizing (Z-scoring) the time series, the units of  $C$  (covariance) are equivalent to the Pearson correlation coefficient.

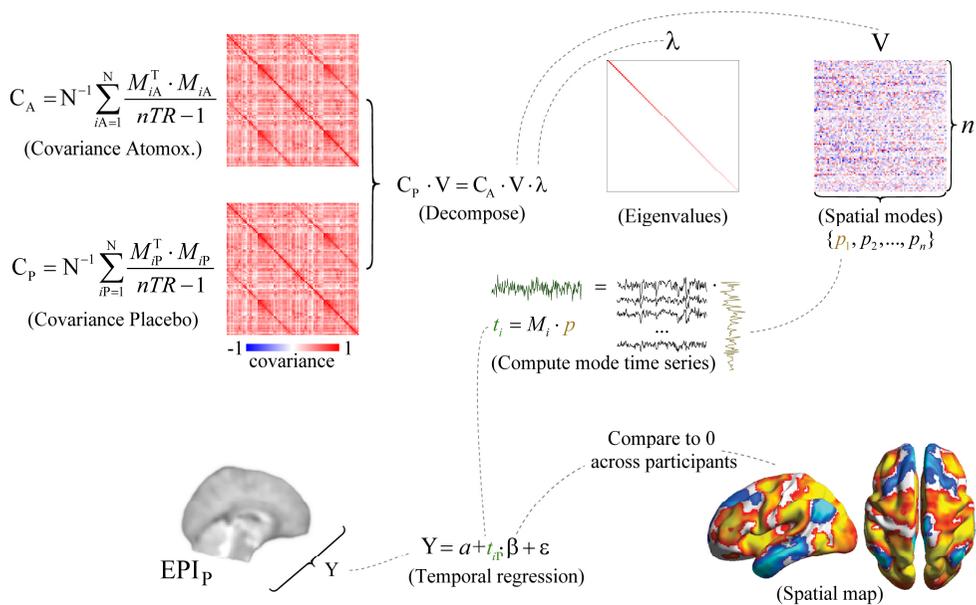
*Singular value decomposition of covariance matrices.* The dual regression analysis described at the beginning of this section relies on a linear decomposition of the data (ICA). An alternative linear decomposition, eigenvalue decomposition (SVD) as used in principal component analysis, can be extended to comparisons of correlation patterns between two conditions. Singular value decomposition (SVD) is a multivariate linear decomposition that identifies spatial modes of signal cofluctuations; each of these spatial modes can be conceptualized as a 'network' of correlated (or anti-correlated) brain regions (Mittra and Pesaran, 1999; Friston and Büchel, 2004; Donner et al., 2013). The decomposition can be generalized to extract spatial modes

that are more strongly expressed in one experimental condition than in the other, in other words: maximize the ratios of explained variance between conditions (Friston and Büchel, 2004; Donner and Nieuwenhuis, 2013).

Our first objective was to determine if SVD (specifically: principal component analysis) identified similar correlation patterns as ICA. We thus submitted the covariance matrices ( $C$ ) of the placebo and atomoxetine conditions to SVD:

$$(2) C = V \cdot \lambda \cdot V^T$$

where  $T$  denotes transposition,  $\lambda$  is an  $n$ -by- $n$  matrix with eigenvalues on its diagonal, and  $V$  is an  $n$ -by- $n$  matrix of corresponding eigenvectors in which rows are brain regions ( $n = 90$ ) and columns define individual modes ( $p$ ).



**Figure 2.** Schematic overview of the spatial mode decomposition method. The covariance matrices  $C_A$  and  $C_P$  are submitted to generalized eigenvalue decomposition to produce a matrix of eigenvalues ( $\lambda$ ) and eigenvectors ( $V$ ). The decomposition equation as given here delineated modes that were more strongly expressed in the placebo condition than in the atomoxetine condition. To identify modes that were more strongly expressed in the atomoxetine condition, the covariance matrices  $C_A$  and  $C_P$  were swapped. After decomposition, the participant-level time series ( $t$ ) corresponding to each individual spatial mode ( $p$ ) can be computed for each run  $i$  by projecting the mode onto the data ( $M$ ). The number of brain regions in the parcellation scheme is denoted by  $n$ . A spatial map of brain regions that consistently covaried with the mode time series is computed by regressing the spatial mode time series for the atomoxetine (A) and placebo (P) conditions onto the voxel-level fMRI time series, and comparing the regression coefficients to zero across participants.

The overall sign of the elements in  $p$  is arbitrary but the sign of one element with respect to another indicates their relative co-variation, with equal signs indicating positive correlation and unequal signs indicating negative correlation.

For each run  $i$ , separately for the atomoxetine and placebo condition, we calculated participant-level time series  $t$  corresponding to each mode by projecting the mode onto the participant-level multivariate time series  $M$  via:

$$(3) \quad t_i = M_i \cdot p$$

The so-computed  $t$  described the time-varying strength of the expression of the spatial mode (functional network) in each individual participant's data. Next, we obtained voxel-level spatial maps for each mode and each run by regressing the vectors  $t_i$  onto the corresponding voxel-level BOLD data using multiple linear regression. We then selected modes based on maximal spatial correlation with the 10 intrinsic connectivity networks reported by Smith et al. (2009), similar to the selection of ICA components described above. For the placebo condition, the average correlation coefficient was 0.41 (SD 0.12, min 0.16, max 0.56), and for the atomoxetine condition the average correlation coefficient was 0.40 (SD 0.12, min 0.15, max 0.53), indicating that SVD was able to identify networks of intrinsically co-fluctuating activity reasonably well (Figure S1 and S2). Similar results were obtained with the Craddock atlas. Next, we describe the generalization of SVD to extract modes that are more strongly expressed in one condition relative to the other.

*Generalized eigenvalue decomposition of covariance matrices.* We used generalized eigenvalue decomposition to decompose the covariance matrices into spatial modes that maximized the ratio of explained variance in the placebo condition relative to the atomoxetine condition (Mitra and Pesaran, 1999; Friston and Büchel, 2004; Donner et al., 2013). Figure 2 shows a schematic overview of this analysis approach. We refer to previous work for experimental validation of generalized eigenvalue decomposition for use on fMRI data (Donner et al, 2013). For simplicity, we here refer to this method as 'spatial mode decomposition'. Using the 'eig' function in MATLAB 2012a, we decomposed the participant-averaged atomoxetine covariance matrix  $C_A$  and placebo covariance matrix  $C_P$  by solving the equation:

$$(4) \quad C_P \cdot V = C_A \cdot V \cdot \lambda$$

where  $\lambda$  is an  $n$ -by- $n$  matrix with generalized eigenvalues on its diagonal, and  $V$  is an  $n$ -by- $n$  matrix of corresponding eigenvectors in which rows are brain regions ( $n = 90$  for the AAL atlas, and  $n = 140$  for the Craddock atlas) and columns define individual modes ( $p$ ). Here,  $p$  were spatial patterns that maximized the variance accounted for in one condition relative to the other (as measured by the corresponding  $\lambda_p$ ). The above equation identified spatial modes that were more strongly expressed in the placebo condition than in the atomoxetine condition. To identify spatial modes that were more strongly expressed in the atomoxetine condition, the covariance matrices  $C_A$  and  $C_P$  were swapped. We arranged  $V$  and  $\lambda$  such that their first entries

corresponded to the modes that explained most variance. In other words, we sorted  $\lambda$  in descending order and then sorted  $V$  by  $\lambda$ .

We next calculated participant-level time series  $t$  corresponding to  $p$  for each individual run  $i$  via:

$$(5) \quad t_i = M_i \cdot p$$

Here,  $t_i$  was a vector with length 211 (the number of volumes), and  $M_i$  was a matrix of Z-scored fMRI time series from the run, with size 211 by  $n$  (volumes by brain regions).

To examine the spatial distribution of each mode, we used the corresponding time series in multiple temporal regression. Specifically, for each participant and condition separately, we regressed the mode time series onto the single-voxel time series from the corresponding run. This yielded one spatial map of regression coefficients per participant, condition, and mode. For each mode and for each condition, we could then compare the regression coefficients to zero using non-parametric permutation testing (10,000 iterations). The  $\alpha$  level was set at 0.05, FWE-corrected for multiple comparisons using threshold-free cluster enhancement. The resulting statistical parametric maps indicated which voxels (if any) covaried with the mode time series consistently across participants, and were thus indicative of the spatial distribution of the modes.

*Quantifying the across-subject consistency and reliability of spatial modes.* The spatial modes were computed such they explained more variance in the group-average data, in the atomoxetine condition than in the placebo condition (or the converse). We aimed to quantify, in a cross-validated fashion, how consistently the fluctuation strength of these group-average spatial modes distinguished between conditions within individual subjects. The fluctuation amplitude  $s_i$  corresponding to each mode's time series in each individual run from each participant quantified the amount of variance that the mode explained in the data, and was calculated via:

$$(6) \quad s_i = t_i^T \cdot t_i$$

Note that this is equivalent to:

$$(7) \quad p^T \cdot M_i^T \cdot M_i \cdot p = p^T \cdot C_i \cdot p = s_i$$

We then divided  $s_i$  by the sum of eigenvalues ( $\lambda$ ) to convert it to units of percentage variance explained. In contrast to the eigenvalues, which capture the group-level mode's ratio of explained variance between conditions,  $s_i$  captured the amount of variance that the mode captured in the condition-specific runs at the individual participant-level. For cross-validation, we defined modes (using eq. 4) based on the group-average covariance matrices  $C_A$  and  $C_P$  that were generated from the first half of volumes in  $M_i$  (using eq. 1). Then, each mode was projected onto the (independent) remaining half of volumes in  $M_i$  as described above (eq. 5) and their corresponding fluctuation amplitudes were calculated (via eq. 6). We then used the second half of volumes to define the modes and projected them onto the first half, and averaged the two values of  $s_i$ . The percentage variance explained by each mode could then be

compared between conditions with non-parametric permutation testing (10,000 iterations).

We used receiver operating characteristic (ROC) analysis to quantify the reliability of the spatial modes in discriminating between experimental conditions, at the level of short segments (25% of volumes, ~114 s) of individual fMRI runs. ROC analysis performs more accurately with densely populated distributions of measurements. Thus, we defined spatial modes based on the group average covariance matrices calculated from a smaller subset of volumes (25%), as described above (using eq. 1 and eq. 4). We subdivided the remainder of volumes into 20 equal-sized bins, and computed  $s_i$  for each of them. We cross-validated the fluctuation amplitude calculation by computing modes and projecting them onto the remaining data four times such that eventually all data were used to define the modes. This yielded four distributions of  $s_i$  per condition and participant that were submitted to ROC analysis, resulting in four ROC-curves per participant. We calculated the area under the ROC-curve (AUC) and averaged the resulting AUC values across the four ROC-curves of each participant. This AUC value could then be interpreted as the probability with which we could predict the condition from the mode's fluctuation strength in a given data segment. The AUC values were tested for significance by comparing them to chance level (0.5) using non-parametric permutation testing (10,000 iterations). In order to exclude the possibility that the significance of the ROC results depended on the number (25%) of volumes on which the mode was defined, we repeated the ROC analyses for modes defined on ~14%, 20%, and ~33% of the data, and found identical results in terms of direction and significance.

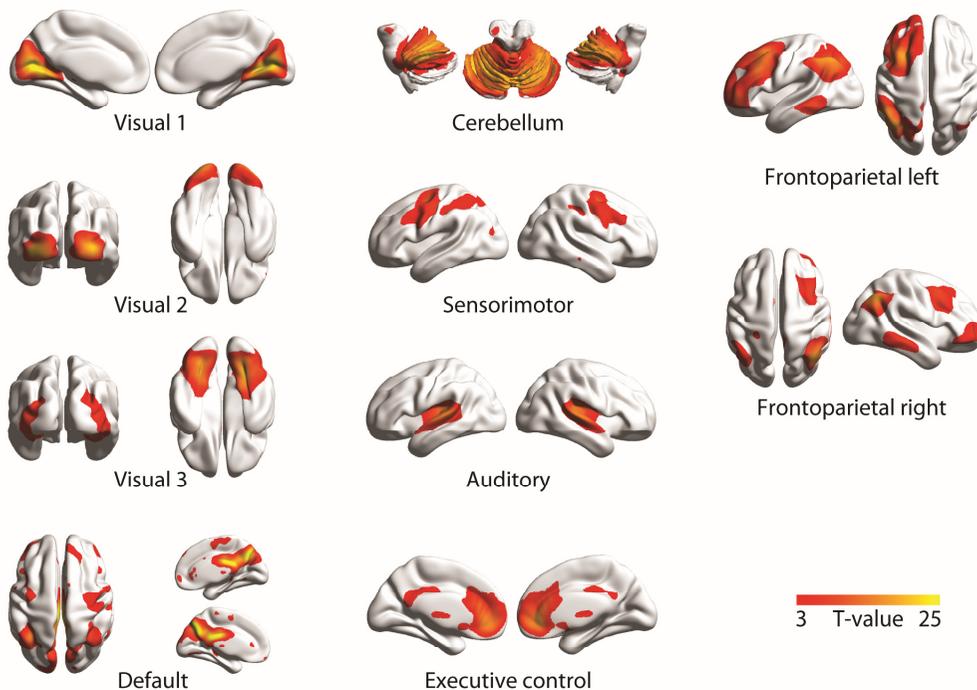
*Correlation between mode spatial maps and independent components.* To determine if the mode spatial maps depended on the parcellation scheme, we used spatial correlation. Specifically, for each individual participant and condition, we correlated the (unthresholded) spatial maps of regression coefficients of the modes that were generated with the AAL atlas, and those that were generated with the Craddock atlas. We then compared the distribution of Fisher-transformed correlation coefficients to zero using a two-tailed  $t$ -test. Similarly, we characterized the correspondence in mode spatial maps between the individual conditions by correlating the unthresholded spatial maps at the individual participant level, and comparing the resulting distribution of Fisher-transformed correlation coefficients to zero using a two-tailed  $t$ -test.

To characterize correspondence between the mode spatial maps and well-characterized intrinsic connectivity networks, we first created a mode spatial map by temporally concatenating the mode time-series of the atomoxetine and placebo conditions, and regressing this concatenated time series onto the temporally concatenated BOLD time-series data for each participant. The purpose of this concatenation procedure was to create spatial maps that were independent of drug condition, similar to ICA components that were identified in temporally concatenated EPI data. We could then correlate these condition-invariant unthresholded participant-

level mode spatial maps with the IC spatial maps that were selected for dual regression analysis, and compare the distribution of Fisher-transformed regression coefficients to zero using a two-tailed  $t$ -test. In all cases where we report average correlation coefficients, we applied Fisher's  $r$ -to- $Z$  transform prior to averaging, and subsequently applied the  $Z$ -to- $r$  transform.

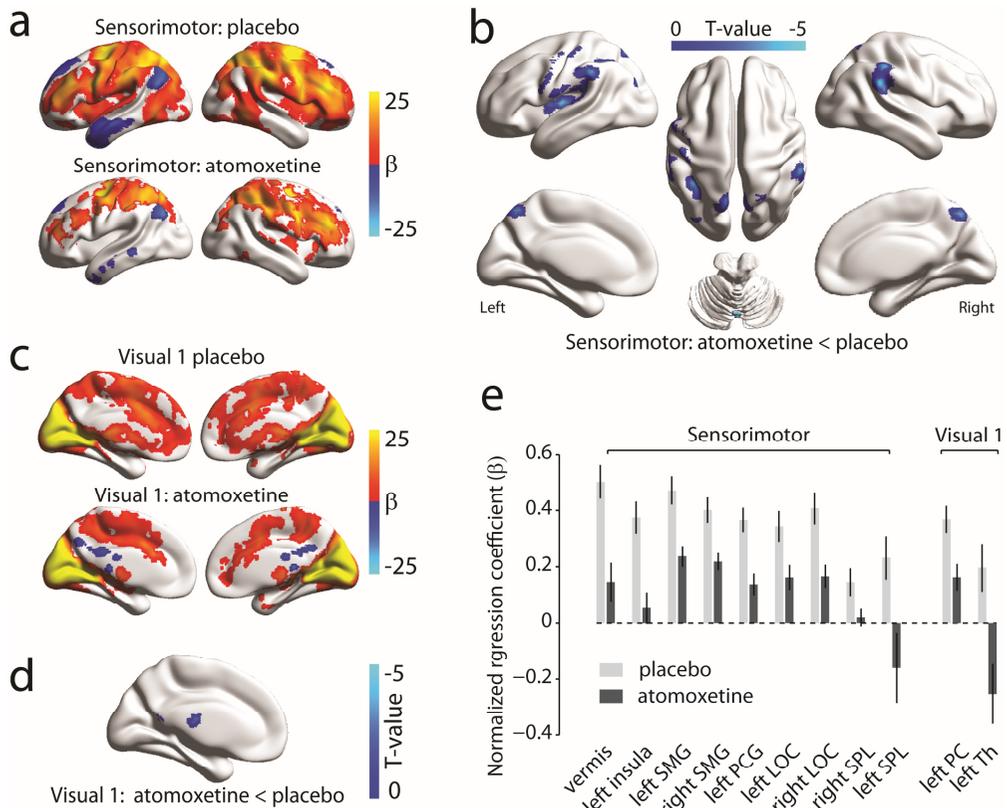
### 4.3 Results

Our first aim was to characterize atomoxetine-induced changes in fine-grained (voxel-level) co-fluctuation strength with a set of 10 well-characterized intrinsic connectivity networks (Smith et al., 2009) using conventional methods for the analysis of pharmacological resting-state fMRI: dual regression (Figure 1) (Beckmann, 2009; Filippini et al., 2009). Second, we report the results of an alternative analysis approach that is targeted at finding spatial patterns (“modes”) of correlated activity that maximize the ratio of explained variance between conditions in a fully data-driven manner (Figure 2). Thus, instead of testing if any of an *a priori* selection of networks showed



**Figure 3.** Spatial maps of the independent components that were selected for dual regression analysis. Components were selected based on spatial correlation with the 10 canonical resting-state networks presented by Smith et al. (2009).

drug-induced changes in the cofluctuation strength (as with dual regression), spatial mode decomposition directly yielded the networks that exhibited drug-induced changes in cofluctuations, in terms of cofluctuation strength or spatial pattern, or both. In contrast to singular value decomposition of condition-level data (see Materials and Methods) or linear decomposition of the data using ICA, spatial mode decomposition directly reveals those patterns that cofluctuate more/less in one condition than in the other. Moreover, this analysis allows us to characterize to what extent atomoxetine-



**Figure 4.** Results of the dual regression analysis. **a**) Brain areas that significantly ( $p < 0.05$ , FWE-corrected) cofluctuated with the sensorimotor network. **b**) Atomoxetine-induced changes in cofluctuation strength with the sensorimotor network. **c**) Brain areas that significantly cofluctuated with the visual 1 network. **d**) Atomoxetine-induced changes in cofluctuation strength with the visual 1 network. Blue colors indicate reduced cofluctuation strength following atomoxetine compared to placebo. **e**) Cofluctuation strength (range-normalized across participants to between -1 and 1 for illustrative purposes only) for each condition and each significant ( $p < 0.05$ , FWE-corrected) cluster of atomoxetine-induced changes in cofluctuation strength. Error bars show the SEM. Abbreviations: PC: precuneous cortex; Th: thalamus; SMG: supramarginal gyrus; PCG: precentral gyrus; LOC: lateral occipital cortex; SPL: superior parietal lobule.

related modulations of covariance reflect reconfigurations in the topological structure of cofluctuations by directly comparing the mode's spatial structure between the atomoxetine and placebo conditions.

*Results dual regression analysis.* We first computed group-level ICA spatial maps and then submitted a selection to dual regression analysis (Figure 1). The spatial maps of the included components are shown in Figure 3.

The sensorimotor network and first visual network both showed significant ( $p < 0.05$ , FWE-corrected) atomoxetine-induced changes in cofluctuation strength (Figure 4a-d). In all significant clusters, atomoxetine reduced the strength of cofluctuation, consistent with our earlier findings obtained at coarser levels of spatial granularity (van den Brink et al., 2016). All clusters that showed a significant atomoxetine-related reduction in cofluctuation strength also covaried positively with the component time series in the placebo condition (Figure S3), and were thus functionally linked to the networks. In most cases, atomoxetine moved cofluctuations from positive towards zero, and in some cases cofluctuations reversed polarity, at least numerically (Figure 4e; Figure S3). Cofluctuations with the component time series in the clusters that numerically reversed polarity, however, were not significantly negative in the atomoxetine condition (Figure S4). Thus, the primary effect of atomoxetine on cofluctuation strength was to move positive cofluctuations towards zero. The MNI coordinates and peak T-statistics of all significant clusters are summarized in Table 1. Together, these results suggest that atomoxetine attenuated voxel-level cofluctuation strength between brain regions that cofluctuated positively with sensory- and motor-related networks, and the fluctuations of those networks.

**Table 1. Clusters that showed an atomoxetine-induced change in cofluctuation strength with resting-state networks.**

Component	Location	MNI Coordinates (x y z)	Peak T-statistic
Visual 1	Left PC	-4 -56 12	-4.91
	Left Th	0 -22 12	-4.21
Sensorimotor	Vermis	0 -72 -14	-6.04
	Left insula	-40 -12 8	-4.65
	Left SMG	-66 -34 34	-5.71
	Right SMG	58 -44 28	-4.57
	Left PCG	-50 -10 56	-4.85
	Left LOC	-34 -82 26	-4.36
	Right LOC	18 -66 72	-4.73
	Right SPL	38 -54 68	-3.30
	Left SPL	-34 -42 70	-3.43

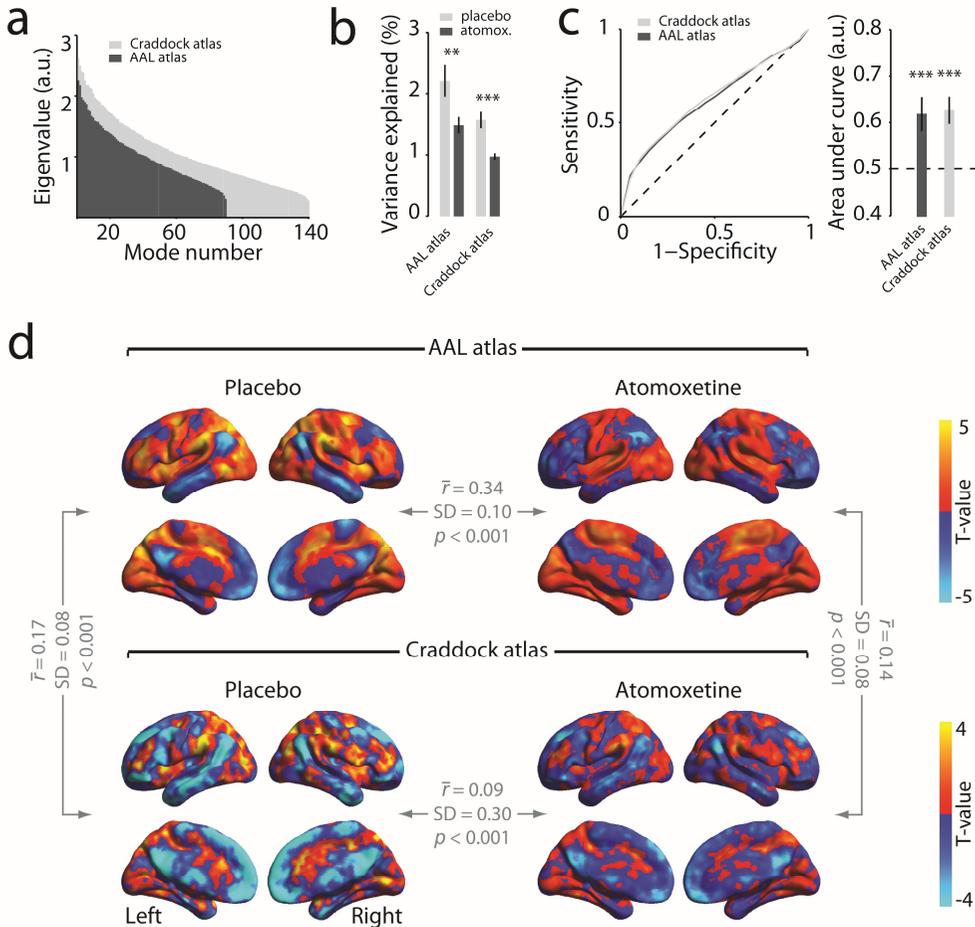
Cluster locations were assessed using the Harvard-Oxford structural atlas. Peak MNI coordinates are indicated in mm. Abbreviations: PC: precuneous cortex; Th: thalamus; SMG: supramarginal gyrus; PCG: precentral gyrus; LOC: lateral occipital cortex; SPL: superior parietal lobule.

*Spatial modes that are less strongly expressed in the atomoxetine condition relative to placebo.* The above reported atomoxetine-related reductions in cofluctuation strength with sensory- and motor-related networks resulted from dual regression analysis. In the following, we report the results of an alternative analysis approach that directly reveals those patterns that cofluctuate more/less in one condition than in the other, and that allows us to characterize to what extent atomoxetine-related modulations of covariance reflect reconfigurations in the topological structure of cofluctuations. Given that dual regression analysis identified only atomoxetine-reductions in the strength of cofluctuations, we first focus on spatial modes that reflected an atomoxetine-related reduction in cofluctuations.

The eigenvalues of the modes that were less strongly expressed in the atomoxetine condition are shown in Figure 5a. We focused on mode number 1 because it had the largest eigenvalue and thus accounted for most variance in the data, and because mode orthogonality can obscure the interpretation of modes with higher ranks (c.f. Donner et al., 2013).

This first spatial mode robustly differed in its fluctuation strength (i.e., variance explained) between the atomoxetine and placebo conditions. We first tested, using cross-validation, if the first spatial mode consistently explained less variance in the atomoxetine condition than in the placebo condition across subjects: we computed the mode based on covariance in each half of the volumes in each participant's runs, projected the mode onto the remaining half of the volumes, and calculated the mode's proportion of explained variance in each condition (see Materials and Methods). Indeed, the first spatial mode accounted for significantly less variance in the atomoxetine condition than in the placebo condition, for both parcellation schemes (AAL:  $p = 0.003$ ; Craddock:  $p < 0.001$ ; Figure 5b). Further, ROC analysis showed that even at the level of short individual data segments (~114 s), the first spatial mode's fluctuation strength reliably discriminated between drug and placebo conditions, with AUC-values larger than 0.6 for both parcellation schemes (Figure 5c). This indicates that the spatial mode identified by our analysis reflected a robust effect of the pharmacological intervention on brain-wide intrinsic correlations.

We next compared the spatial distributions of the expressions of the first spatial mode, between different parcellation schemes and with the spatial maps obtained from the dual regression analysis. The unthresholded spatial map of mode 1 (reflecting voxel-level covariation with the mode's time series) is shown in Figure 5d, separately per condition and for modes that were generated using the AAL atlas, and for modes that were generated using the Craddock atlas. Figure S5 shows thresholded ( $p < 0.05$ , FWE-corrected) mode spatial maps. Despite using parcellation schemes that differed both in the number of brain regions and in the way the brain regions were defined (anatomical parcellation and functional clustering, respectively), the mode spatial maps generated with the two atlases corresponded robustly across



**Figure 5.** Spatial modes that are less strongly expressed in the atomoxetine condition. **a**) Eigenvalues of all modes. **b**) A comparison between conditions of the percentage of variance explained by the first mode. **c**) ROC curves to distinguish conditions based on the fluctuation amplitude of the first mode. **d**) Spatial map of the first mode. Colored regions show covariation with the mode time series. Error bars show the SEM. \*\*:  $p < 0.01$ ; \*\*\*:  $p < 0.001$ . The  $r$  values indicate the average correlation coefficient across participants.

participants (placebo:  $t(23) = 10.43$ ,  $p < 0.001$ ; atomoxetine:  $t(23) = 9.54$ ,  $p < 0.001$ ; Figure 5d).

To determine if the mode corresponded to any of the intrinsic connectivity networks that were used for dual regression analysis, we correlated the mode spatial map with the ICA component spatial maps at the individual participant level. The spatial map of mode 1 that was generated with the AAL atlas correlated most strongly with the left-lateralized frontoparietal ICA component (mean  $r = -0.15$ , SD 0.05;  $t(23) = -16.33$ ,  $p < 0.001$ ). The spatial map of mode 1 that was generated with the Craddock atlas also correlated significantly across participants with the left-lateralized

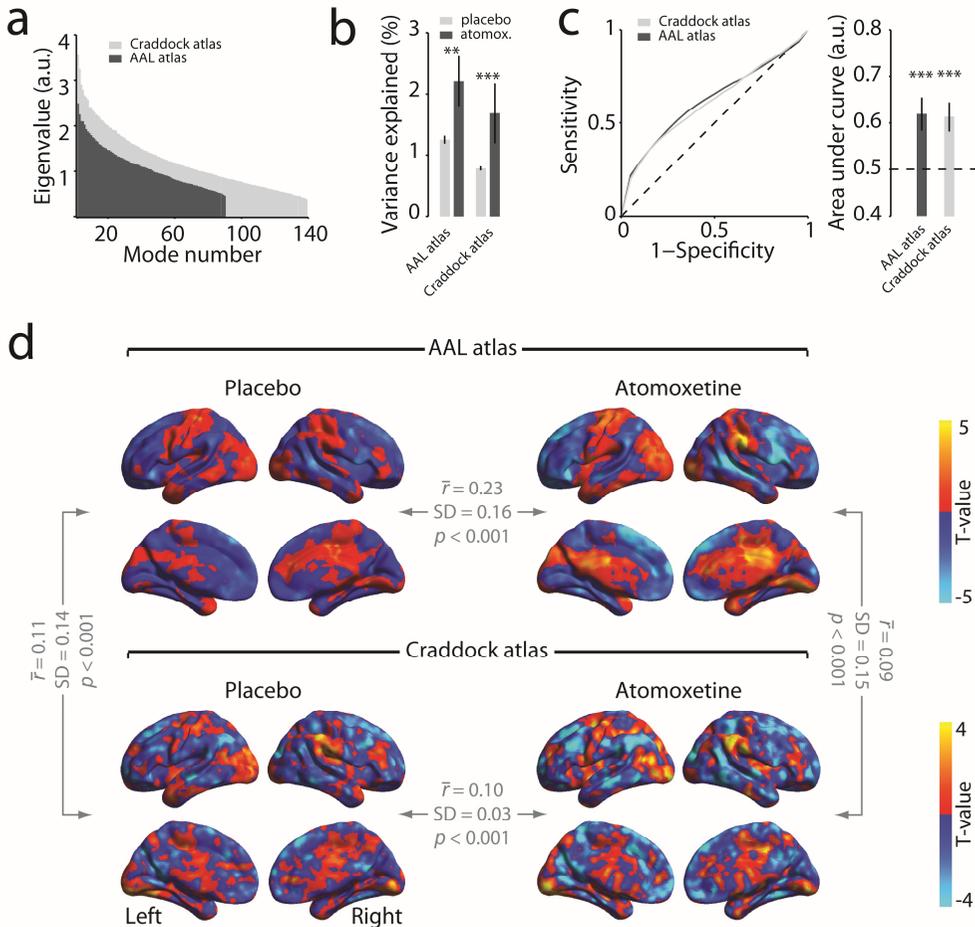
frontoparietal component (mean  $r = -0.07$ , SD 0.04;  $t(23) = -8.68$ ,  $p < 0.001$ ). Moreover, for both atlases the spatial map of mode 1 correlated significantly with the sensorimotor component (AAL: mean  $r = 0.13$ , SD 0.04;  $t(23) = 17.41$ ,  $p < 0.001$ ; Craddock: mean  $r = 0.07$ , SD 0.03;  $t(23) = 14.36$ ,  $p < 0.001$ ), suggesting that atomoxetine reduced the strength of cofluctuations in a network that resembled the ICA-identified sensorimotor network (Figure S6). If this is indeed the case, then the regions that showed an atomoxetine-related reduction in cofluctuations in the dual regression analysis (Figure 4) should show a similar sign in the mode spatial map (i.e. be part of the same cofluctuating network). We therefore masked the thresholded spatial map of mode 1 in the placebo condition (Figure S5) with the significant clusters in Figure 4. All clusters showed the same sign (Figure S7), indicating that the spatial mode reflected a reduction in cofluctuation strength across brain regions that showed similar reductions in the dual regression analyses.

Spatial correlation also enabled us to examine if the mode reflected a reconfiguration of the spatial structure of cofluctuations or if it reflected a quantitative change in strength that left the structure of cofluctuations intact. To this end, we operationally defined reconfiguration as a change in mode topology, implying a spatial mode that was only expressed in one condition, but not in the other. By contrast, quantitative changes would entail the spatial mode to be expressed in both conditions, only to a different degree. Note that both scenarios might lead to a robust spatial mode maximizing the ratio between variance accounted for in both conditions.

We correlated the mode spatial map of the placebo condition with that in the atomoxetine condition (Figure 5d). Inconsistent with the notion of an atomoxetine-related reconfiguration of cofluctuation structure, the mode spatial map correlated between conditions for both atlases (AAL:  $t(23) = 15.57$ ,  $p < 0.001$ ; Craddock:  $t(23) = 14.89$ ,  $p < 0.001$ ). In other words, the spatial distribution of the expression of the first spatial mode that most discriminated between conditions, was in fact similar between conditions. This indicates that the predominant effect of atomoxetine was a quantitative reduction of the strength of cofluctuations rather than a topological reconfiguration of intrinsic functional connectivity networks (see Discussion).

*Spatial modes that are more strongly expressed in the atomoxetine condition relative to placebo.* Thus far we have shown, using dual regression analysis, that atomoxetine reduced cofluctuation strength between distributed brain regions (including sensory- and motor-related networks). Spatial mode decomposition revealed similar reductions in cofluctuation strength, and further indicated that these changes in cofluctuation strength left the spatial structure of cofluctuations intact. We now turn to spatial modes that reflected an atomoxetine-related increase in cofluctuations (see Materials and Methods), which may have not been identified by dual regression analysis.

The eigenvalues of the modes that were more strongly expressed in the atomoxetine condition are shown in Figure 6a. Again, we selected mode 1 because it accounted for most variance in the data. Similar to the above reported analysis of



**Figure 6.** Spatial modes that are more strongly expressed in the atomoxetine condition. **a**) Eigenvalues of all modes. **b**) A comparison between conditions of the percentage of variance explained by the first mode. **c**) ROC curves to distinguish conditions based on the fluctuation amplitude of the first mode. **d**) Spatial map of the first mode. Colored regions show covariation with the mode time series. Error bars show the SEM. \*\*:  $p < 0.01$ ; \*\*\*:  $p < 0.001$ . The  $r$  values indicate the average correlation coefficient across participants.

mode variance, we computed the mode based on covariance in each half of the volumes, and projected it onto the remaining half. For both atlases, the mode explained significantly more variance in the atomoxetine condition than in the placebo condition (AAL:  $p = 0.002$ ; Craddock:  $p < 0.001$ ; Figure 6b), and consistently throughout the resting-state runs (ROC values  $> 0.6$ ;  $p < 0.001$ ; Figure 6c). Thus, the mode reflected a pattern of brain regions in which activity cofluctuated more strongly following atomoxetine than following placebo.

Figure 6d shows the (unthresholded) spatial map of mode 1, separately per condition, for modes that were generated using the AAL atlas, and for modes that

were generated using the Craddock atlas. Figure S8 shows thresholded ( $p < 0.05$ , FWE-corrected) mode spatial maps. Again, the spatial map of the modes generated with the two atlases corresponded robustly across participants (placebo:  $t(23) = 3.96$ ,  $p < 0.001$ ; atomoxetine:  $t(23) = 3.98$ ,  $p < 0.001$ ).

The spatial map of mode 1 correlated most strongly with the right-lateralized frontoparietal component (AAL atlas: mean  $r = -0.05$ , SD 0.03;  $t(23) = -7.98$ ,  $p < 0.001$ ; Craddock atlas: mean  $r = -0.09$ , SD 0.03;  $t(23) = -14.44$ ,  $p < 0.001$ ). Together, these results suggest that atomoxetine increased the strength of cofluctuations in a distributed network that resembled the right-lateralized frontoparietal network.

Next, we again examined if the mode reflected a change in cofluctuation network structure, or if it reflected a modulation of cofluctuation strength alone. We thus correlated the mode spatial map in the atomoxetine condition and in the placebo condition, and found robust correlations (AAL atlas:  $t(23) = 6.93$ ,  $p < 0.001$ ; Craddock atlas:  $t(23) = 14.89$ ,  $p < 0.001$ ). Thus, the spatial structure of the first mode was similar between conditions, and therefore the most prominent atomoxetine-related increases in cofluctuations can be interpreted as a quantitative increase in the strength of those cofluctuations rather than a more profound reconfiguration of network topology.

In sum, dual regression analysis and spatial mode decomposition converge on the conclusion that atomoxetine weakens cofluctuation strength between distributed brain regions, including sensory- and motor-related networks. In addition, spatial mode decomposition revealed a shift from left- to right-lateralized frontoparietal network dominance, as confirmed by a significant interaction in the strength of correlation between mode polarity (atomoxetine-induced increase versus decrease) and component (frontoparietal left versus right) (repeated-measures ANOVA; AAL:  $F(1,23) = 163.14$ ,  $p < 0.001$ ; Craddock:  $F(1,23) = 56.15$ ,  $p < 0.001$ ). Lastly, spatial mode analysis revealed that the predominant effect of atomoxetine was the quantitative scaling of preexisting cofluctuation patterns whereby the overall spatial structure of these cofluctuating networks was left intact.

#### 4.4 Discussion

In the present study, we examined the effect of the selective NE transporter blocker atomoxetine on the fine-grained spatial structure of resting-state fMRI cofluctuations using dual regression and spatial mode decomposition. First, dual regression analysis revealed that atomoxetine reduced cofluctuation strength between a distributed set of brain regions that included sensory- and motor-related networks. Second, spatial mode decomposition provided converging evidence for such a reduction in sensory- and motor related coupling. Third, spatial mode decomposition revealed an atomoxetine-related shift in the dominance from left-lateralized to right-lateralized frontoparietal network activity. Importantly, spatial mode decomposition indicated that the most prominent atomoxetine-related changes in

cofluctuations did not alter the topology of the networks in which these changes occurred, but instead reflected quantitative modulations within these networks that left the overall cofluctuation structure intact.

The study of small neural circuits has revealed dynamical reconfigurations of functional networks through neuromodulators, including catecholamines (Marder, 2012; Bargmann and Marder, 2013; Marder et al., 2014). Yet, our results show that the total landscape of cofluctuation changes is dominated to a greater extent by quantitative catecholamine-related changes (i.e. mode 1 accounted for most variance relative to other modes, and it reflected quantitative changes). It is worth noting that more subtle catecholamine-related reconfigurations may have occurred, but were not detected by our current analyses (e.g. may have been captured by modes that accounted for less variance, which we did not examine). Additionally, our findings leave open the possibility that the rapid and transient (i.e. phasic) release of catecholamines has a more profound influence on the topological organization of intrinsic cofluctuations, given that such phasic catecholamine release can have qualitatively different effects on neural conductance properties (Rodgers et al., 2011b; Rodgers et al., 2011a) and behavioral performance (de Gee et al., 2017) than changes in tonic levels. Moreover, the dynamical structure of time-varying changes in network topology may be more susceptible to influence by neuromodulatory tone than the stationary network topology, given that such time-varying topological changes have been reported to covary with behavioral performance and pupillary indices of neuromodulation (Shine et al., 2016).

Our findings are broadly consistent with an earlier study (Guedj et al., 2016) that examined the effect of atomoxetine on resting-state cofluctuations in rhesus macaques using dual regression, and found widespread atomoxetine-related cofluctuation reductions within and between networks, including the somatomotor, somatosensory, (peripheral) visual, and a bilateral frontoparietal network. Interestingly, the authors report an atomoxetine-related reduction in cofluctuations between the somatomotor network and bilateral clusters that strongly resemble the bilateral clusters in the supramarginal gyrus reported here (Figure 4, Table 1), suggesting that the supramarginal gyrus is a particularly prominent target region of catecholaminergic neuromodulation. Moreover, the general correspondence between the findings reported by Guedj et al. (2016) and those reported here suggests that the modulation of cofluctuating networks, mediated by catecholamines, is a mechanism that occurs consistently across species. However, in contrast to the interpretation by Guedj et al. (2016) our findings suggest that these altered cofluctuation dynamics may reflect quantitative changes rather than broad network reconfigurations.

The finding that atomoxetine reduced cofluctuation strength with the sensorimotor network is noteworthy in light of the therapeutic profile of atomoxetine. Atomoxetine has been shown to improve inhibitory motor control in rats (Robinson et al., 2008), patients with ADHD (Chamberlain et al., 2007), and healthy humans (Chamberlain et al., 2006b). Our findings suggest that these response inhibition-enhancing effects of

atomoxetine may be the result of catecholaminergic action within the sensorimotor system. Nevertheless, the potential link between the here reported reduction in cofluctuation strength with the sensorimotor network and the response inhibition-enhancing effects of atomoxetine awaits further investigation.

Furthermore, we found that atomoxetine caused a shift in the dominance from left- to right-lateralized frontoparietal network dominance. Frontoparietal regions in the right hemisphere have been implicated in attentional reorientation and the regulation of goal-directed stimulus selection (Corbetta and Shulman, 2002; Corbetta et al., 2008; Thiebaut de Schotten et al., 2011). Interestingly, right-lateralized frontoparietal regions have also been suggested to be particularly susceptible to noradrenergic influences (Corbetta and Shulman, 2002; Corbetta et al., 2008), and atomoxetine has been reported to improve the precision of neural representations of stimuli (Warren et al., 2016). It is tempting to speculate that the here observed atomoxetine-related shift from left- to right-lateralized frontoparietal network dominance may indicate a shift towards goal-oriented stimulus processing. While our participants were not engaged in a task (other than active fixation), this speculation provides an interesting avenue for future research.

Our findings show the utility of spatial mode decomposition for the analysis of pharmacological resting-state fMRI data. One of its primary advantages over dual regression analysis is that it does not require an *a priori* selection of functional networks, but instead automatically yields the networks (spatial modes) that show the strongest drug-related effects. Thus, it reduces the chances of overlooking prominent drug-related changes in inter-regional cofluctuations, as evidenced by the atomoxetine-related increases in covariance that were identified by spatial mode decomposition, but not by dual regression. Moreover, spatial mode decomposition is computationally inexpensive when used in combination with an anatomical atlas, as we have done here. We should note, however, that even though our results demonstrate some robustness of the method to the particular parcellation scheme, it is not a certainty that the resulting networks will generalize to other parcellation schemes, in particular those of radically different densities. In addition, we only examined the first modes because they explained the largest amount of variance in the data, but modes with higher rank numbers may contain information regarding relevant changes in connectivity as well. Examining these, however, would require additional statistical corrections that could increase the false negative rate. Moreover, the interpretability of modes with higher ranks may be hindered by mode orthogonality. Lastly, the decomposition can only be used to compare two separate conditions (or groups), which limits its applicability in complex (e.g., longitudinal) study designs. Nevertheless, spatial mode decomposition offers a thorough characterization of drug-related changes in the structure of cofluctuating activity.

In sum, we have shown that dual regression and spatial mode decomposition converge on the conclusion that catecholamines reduce cofluctuation strength within and between distributed systems, including sensory- and motor-related networks. In

addition, spatial mode decomposition revealed an atomoxetine-related shift from left to right-lateralized frontoparietal network dominance. Importantly, however, these quantitative changes left the overall spatial structure of cofluctuations intact, suggesting that the predominant effect of increased synaptic catecholamine levels was to quantitatively scale cofluctuations in preexisting networks. Lastly, our findings lend support to the notion that catecholamines modulate dynamic changes in the strength of intrinsic inter-regional coupling, possibly to coordinate flexible modulations of network interactions to facilitate goal-directed behavior.

#### 4.5 Supplementary Materials

*Ruling out confounding artifacts in the global signal.* Recent findings have suggested that the global MRI signal may contain artifacts that are related to various non-neural sources, and these artifacts are not effectively removed by standard preprocessing techniques (Power et al., 2017). While the independent components that were used for dual regression analysis by definition do not contain such global artifacts (due to the spatial independence of components), these artifacts may have caused spurious differences between conditions in the structure of inter-regional covariance. We therefore applied global signal (the mean of all regional time series) regression to the regional BOLD time series prior to computing covariance matrices, and repeated our key spatial mode decomposition analyses.

For the decomposition placebo < atomoxetine, the percentage variance explained of mode 1 differed between conditions and in the expected direction (AAL:  $t(23) = 4.45$ ,  $p < 0.001$ , area under ROC curve = 0.64,  $t(23) = 6.88$ ,  $p < 0.001$ ; Craddock:  $t(23) = 4.55$ ,  $p < 0.001$ , area under ROC curve = 0.69,  $t(23) = 7.54$ ,  $p < 0.001$ ). Similarly, for the decomposition placebo > atomoxetine the percentage variance explained of mode 1 also differed between conditions and in the expected direction (AAL:  $t(23) = -5.15$ ,  $p < 0.001$ , area under ROC curve = 0.63,  $t(23) = 8.97$ ,  $p < 0.001$ ; Craddock:  $t(23) = -6.23$ ,  $p < 0.001$ , area under ROC curve = 0.63,  $t(23) = 7.06$ ,  $p < 0.001$ ). Moreover, the spatial distribution of mode 1 computed without global signal regression correlated significantly with the spatial distribution of mode 1 computed on global signal-removed time series, for both atlases and for the decomposition in both directions (all  $p$  values < 0.001). Thus, the findings as presented in the main text were unlikely to be driven by spurious differences between conditions relating to artifacts in the global signal.

## Placebo

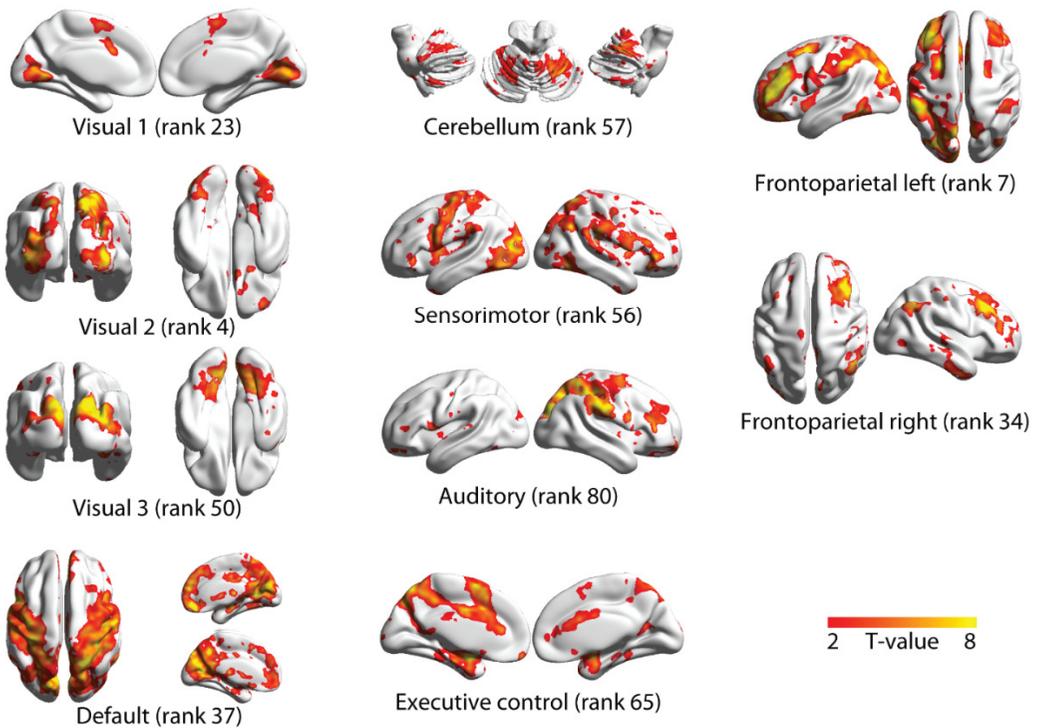


Figure S1. Modes resulting from singular value decomposition of AAL atlas covariance in the placebo condition. The modes were selected based on maximal spatial correlation with the independent component topographies presented by Smith et al. (2009). The rank number indicates the relative proportion of explained variance of each mode, where lower rank numbers account for relatively more variance in the data than high rank numbers.

## Atomoxetine

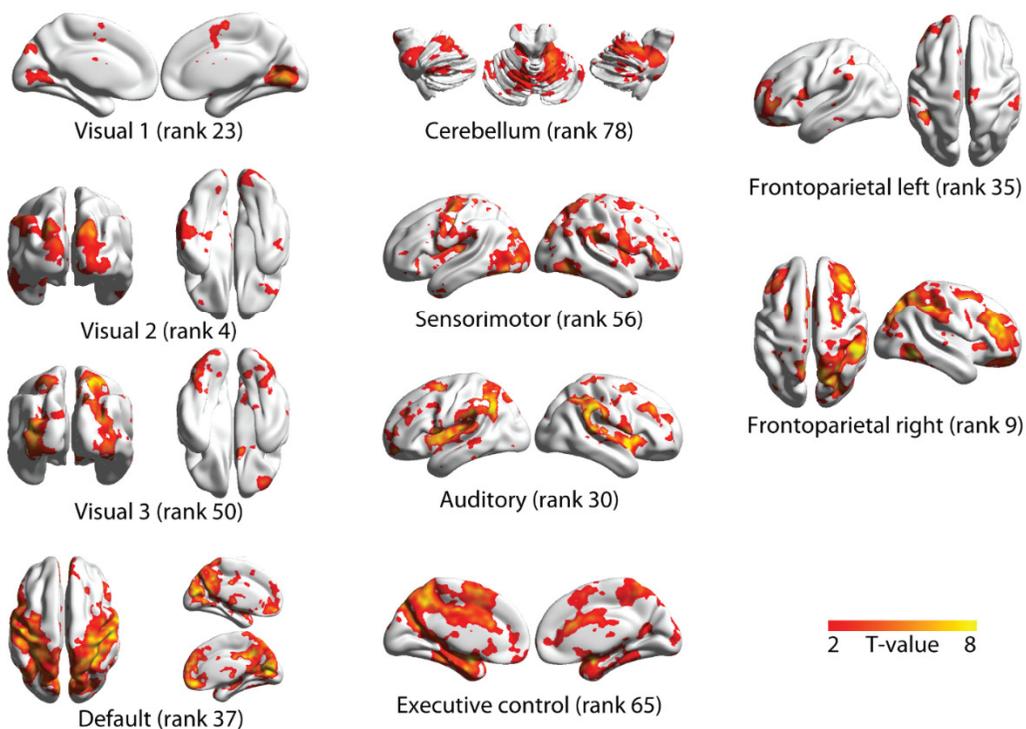


Figure S2. Modes resulting from singular value decomposition of AAL atlas covariance in the atomoxetine condition. The modes were selected based on maximal spatial correlation with the independent component topographies presented by Smith et al. (2009). The rank number indicates the relative proportion of explained variance of each mode, where lower rank numbers account for relatively more variance in the data than high rank numbers.

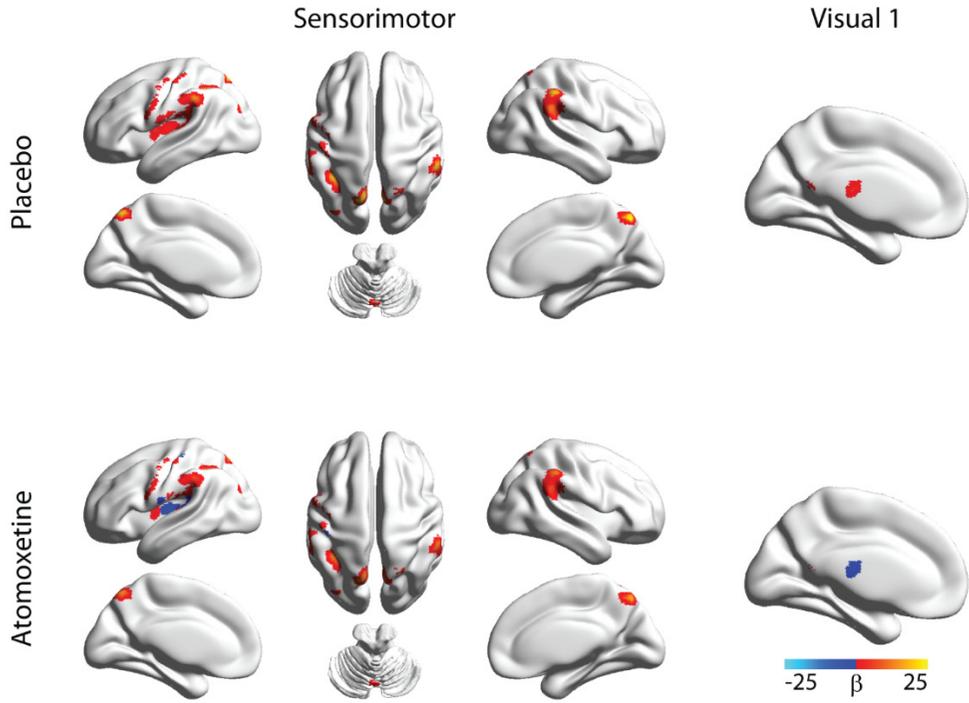


Figure S3. Covariation with component time-series in the individual conditions, only for clusters that showed a significant ( $p < 0.05$ , FWE-corrected) atomoxetine-induced reduction in coupling.

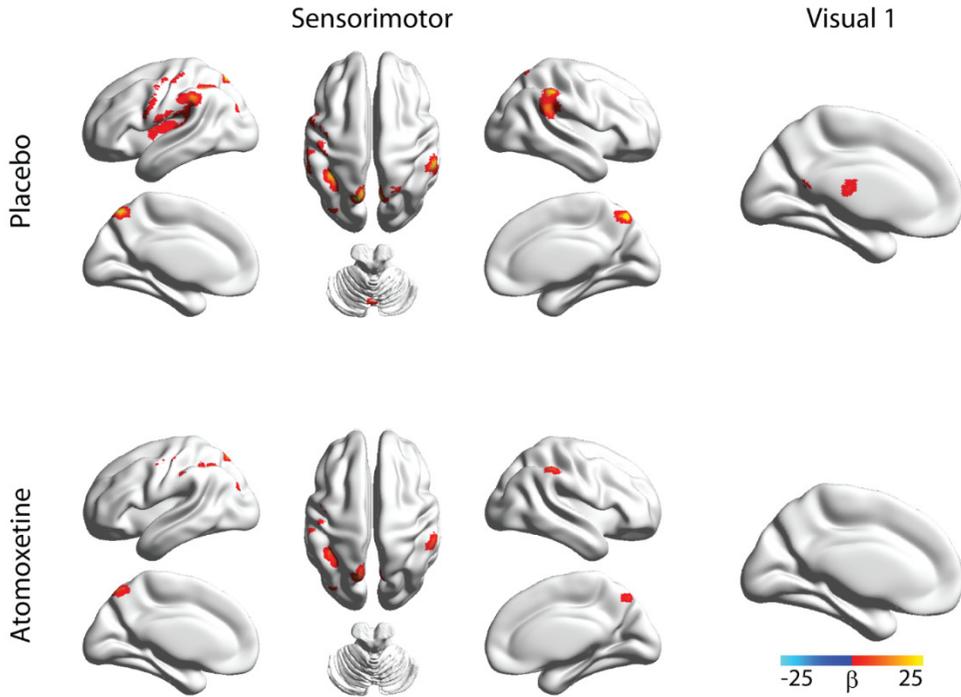


Figure S4. Covariation with component time-series in the individual conditions, only for clusters that both showed a significant ( $p < 0.05$ , FWE-corrected) atomoxetine-induced reduction in coupling, and significant ( $p < 0.05$ , FWE-corrected) covariation with the component time series. Note that brain regions that showed (numerically) an atomoxetine-induced polarity reversal are not significant.

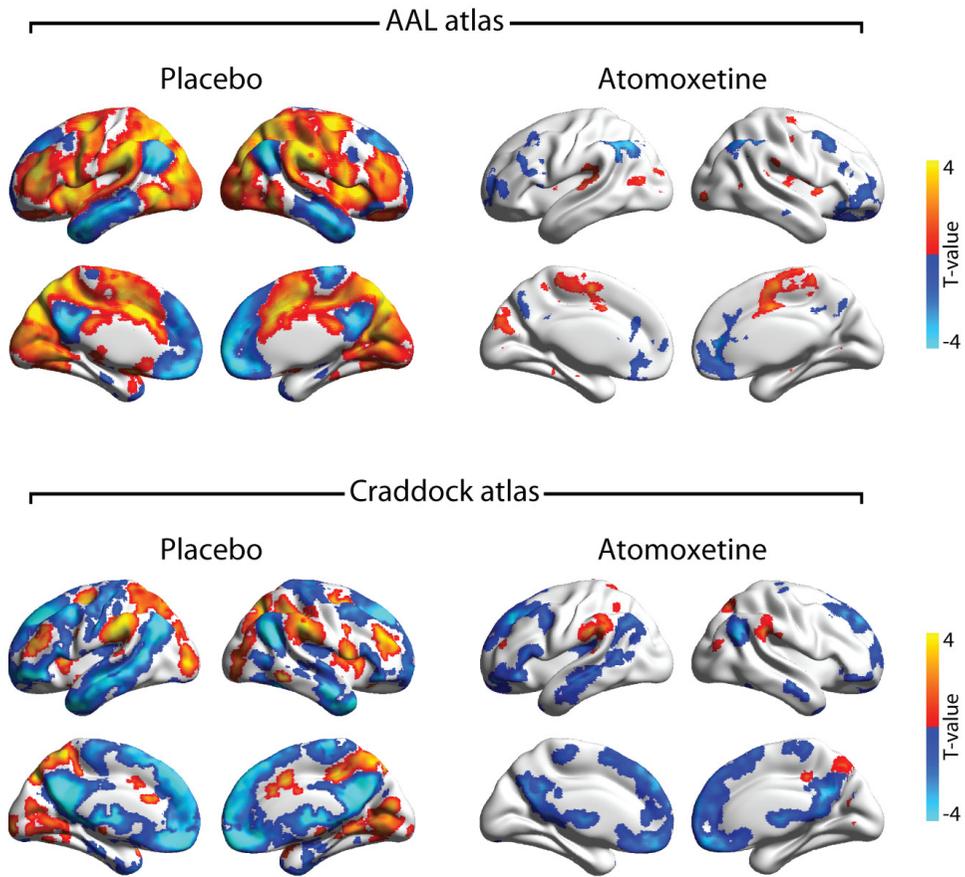


Figure S5. Spatial mode 1 for the decomposition placebo > atomoxetine, with FWE-corrected threshold of  $p < 0.05$ .

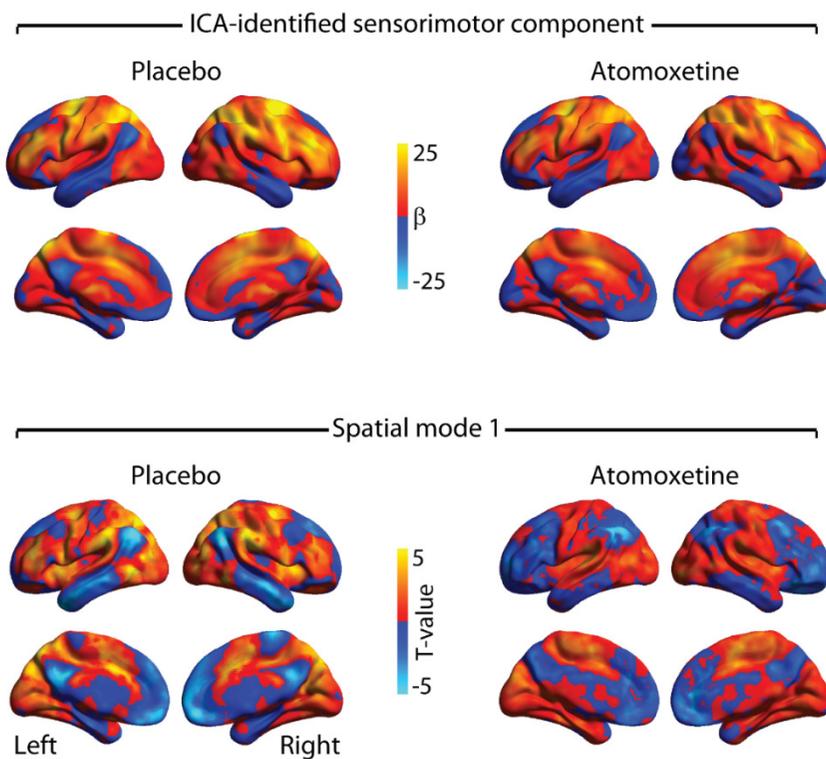


Figure S6. Unthresholded spatial maps of average regression coefficients of the ICA-identified sensorimotor component, and spatial maps of mode 1 generated using the AAL atlas, and for the decomposition placebo > atomoxetine.

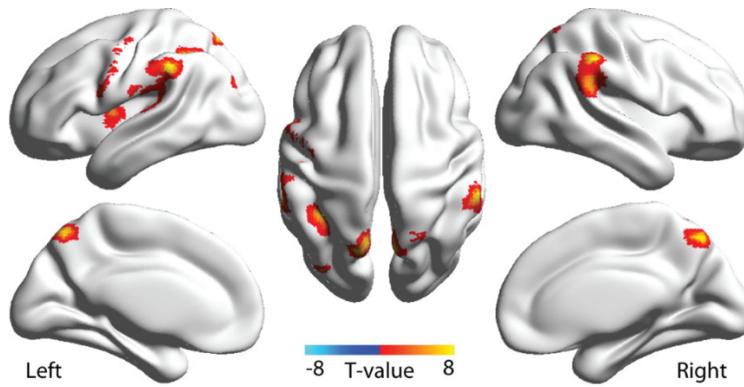


Figure S7. Spatial mode 1 in the placebo condition generated with the AAL atlas (top left panel of Figure S5), masked with the regions that showed a significant reduction in connectivity with the sensorimotor network in the dual regression analysis.

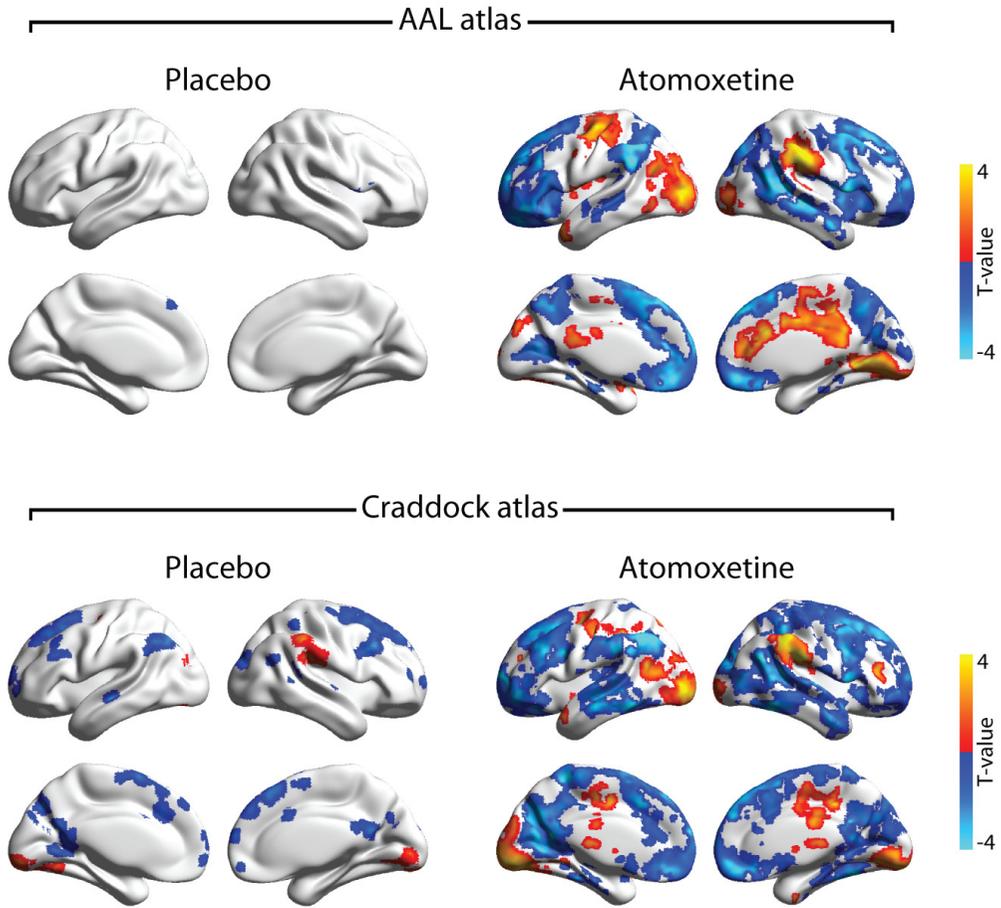


Figure S8. Spatial mode 1 for the decomposition placebo > atomoxetine, with FWE-corrected threshold of  $p < 0.05$ .