

# Supporting Information

Vidaurre et al. 10.1073/pnas.1705120114

## SI Methods

**Behavioral Traits.** The considered behavioral traits, enumerated in Table S1 and further described on the HCP website (<https://wiki.humanconnectome.org/display/PublicData/HCP+Data+Dictionary+Public+500+Subject+Release>), are divided into three groups: well-being, intelligence, and personality. Our motivation for choosing these, and excluding, for example, health-related, sociological, and demographical variables is to focus on the pure cognitive aspects of the individuals, which can be reasonably well described using these three categories. There is abundant literature on the neuroscience of each of these three cognitive domains. In particular, the field of affective neuroscience investigates which neurophysiological features underpin pleasure and happiness (or the absence of) (43). A large body of research has been devoted to identify the brain circuits that sustain the various forms of intelligence, and the extent to which different forms of intelligence are segregated in different brain mechanisms (44). Finally, some advances have been carried through in analyzing how different aspects of personality modulate the response to emotional experience (45). There is, however, scarce research about how each of these relates to dynamic functional connectivity.

**Hidden Markov Modeling.** The HMM assumes that the time series data can be described using a hidden sequence of a finite number of states. More explicitly, if vector  $x_t$  represents the data and  $s_t$  represents the hidden state at time point  $t$ , we assume that

$$x_t | s_t = k \sim \text{Multivariate Gaussian}(\mu_k, \Sigma_k),$$

where  $\mu_k$  is a vector with (number of channels) elements containing the mean blood oxygen level-dependent (BOLD) activation and  $\Sigma_k$  is the (number of channels by number of channels) covariance matrix codifying the variances and covariances between channels when state  $k$  is active. This is referred to as the observational model, which, in this case, characterizes the distribution of each state  $k$  by parameters  $(\mu_k, \Sigma_k)$ . Although we use a multivariate Gaussian here, other choices for the observation model are possible; for example, one could use a multivariate autoregressive (MAR) observation model, where the states are defined and identified by their distinct spectral signature (20). Furthermore, the state sequence is regularized by modeling the probability of transition between all pairs of brain states; that is, before observing the data, the probability,  $\text{Pr}$ , of a given state at time point  $t$  depends on which state was active at time point  $t - 1$ :

$$\text{Pr}(s_t = k) = \sum_l \Theta_{l,k} \text{Pr}(s_{t-1} = l),$$

where  $\Theta_{l,k}$  refers to the transition probabilities. Within matrix  $\Theta$ , we can further distinguish between the on-diagonal elements,  $\Theta_{kk}$ , which control the persistence of each state, and the off-diagonal elements,  $\Theta_{kl}$  (with  $k \neq l$ ), which refer to the actual transitions. Finally, a parameter  $\eta$  encodes the initial state probabilities for each scanning session. Therefore, according to this formulation, the observed data at each time point are effectively modeled as a mixture of Gaussian distributions, with weights given by  $w_{tk} = \text{Pr}(s_t = k)$ .

Although it is possible to apply the HMM on each subject independently, we applied the HMM on all concatenated subjects here, such that we obtained a group estimation of the states. Therefore, whereas the states were at the group level, the information of when a state becomes active (the state time course) is specific to each subject. On these premises, an inference algorithm is used to estimate the parameters defining the posterior

distribution of each state (what are the states  $\mu_k$  and  $\Sigma_k$ ), the probability of each state being active at each time point (when the states occur,  $s_t$ ), and the transition probabilities between each pair of states from the available data (probabilities,  $\Theta_{l,k}$ ). In particular, based on the principles of variational Bayes (VB), we use an approach providing an analytical approximation of the model posterior distribution at a reasonable cost by assuming certain additional factorizations in such posterior distribution (42). In practice, VB inference is typically based on the idea of updating groups of parameters in turn, iterating through the different groups of parameters until convergence is attained. Because of the large amount of data, however, VB is still computationally costly, such that we use a strategy based on stochastic learning, which greatly alleviates both time and memory requirements (46). In a nutshell, the basic idea is to update the parameters on a subset of subjects at a time only, such that the update is noisy but inexpensive. Comprehensive details about VB and HMM modeling can be found elsewhere (19, 20, 42, 46). One important byproduct of VB is the free energy, an approximation to the Bayesian model evidence that provides an estimation of how well the model fits the data and can be used for model selection purposes (42). Before the main inference procedure, we estimated five shorter (independent) inference runs, each starting from a random initialization; the best in terms of the free energy was then used as an initial point for the main inference run.

**Estimation and Statistical Testing on the Metastates.** The metastates were apparent in two different ways. The first method was using the Louvain community detection algorithm (22) on the transition probability matrix (TPM). This method aims to find communities or nodes (here, metastates) in a graph (here, a directed, weighted graph representing the TPM), such that the connectivity between the nodes (here, states) within a community is strong with respect to the connectivity across communities. We embedded the estimation into a bootstrapped subject resampling procedure from which we can obtain intervals of confidence for the metastate separation and, then, a  $P$  value. Second, if we look at the structure of the FO correlation matrix, which contains the correlation for each pair of states' FO across subjects, the metastates naturally emerge from the fact that this matrix has a strong mode of covariance (or eigenvector) that contains most of the information and reflects the metastate distribution. This eigenvector mainly captures that the correlation between states that belong to the same metastate is very high and positive and that the correlation between states that belong to different metastates is also very high and negative. Hierarchical clustering using the Ward's algorithm confirms this result.

**Assessing the Relation Between the States and the Behavioral Traits.** To assess how strongly the state distribution relates to the considered behavioral traits, we use Bayesian partial least squares (BPLS), which aims to predict the state distribution from the behavioral traits using an intermediate low-dimensional space (23). To rule out any possible sex or age influence, we regressed them out, along with motion, from both the FO and the traits. To test the significance of the prediction, BPLS was embedded into permutation testing (1,000 permutations), with the permutations respecting the family structure of the subjects (47). To evaluate the relationship of the behavioral traits with each state and metastate separately, we use regularized regression (42) in a 10-fold, cross-validation setting, again regressing out sex, age, and

motion. To assess how much the metastates' FO is better predicted than the individual states' FO, we repeated this cross-validation analysis for 1,000 bootstrap examples. Finally, we used canonical correlation analysis (CCA) (24) to quantify the individual significance of each variable (either FO or trait). CCA estimates the linear relation between two sets of variables by finding an optimal linear combination for each of the sets in such a way that these linear combinations (referred to as a canonical covariates) are maximally correlated. The difference with the CCA model used by Smith et al. (24) is, essentially, that we use the nonstationary dynamics encoded in the FO instead of stationary functional connectivity. Because we have 12 states (12 measures of FO) only, as opposed to the 19,900 functional connectivity edges in the study by Smith et al. 24, we do not need to perform principal component analysis or any other regularization mechanism to control the dimensionality of the data; hence, our use of CCA here does not require any tuning parameter. Once the CCA model is trained, we can assess each variable (FO or behavioral trait) independently by looking at the correlation between the variable and the opposite canonical covariate. To assess the statistical significance of these correlations, this procedure was embedded into permutation testing (1,000 permutations).

## SI Results

**Relation Between the HMM Dynamics and FO Correlations.** What is the relation between the TPM (Fig. 2A) and the FO correlation matrix (Fig. 2B)? Does the latter follow from the former? We next show through simulations that the strong FO correlation (metastate) structure can be caused by corresponding structure in the state transition probabilities (probability of switching from one state to another) and persistencies (probability of remaining in the current state). While this effect is stronger for the state transition probabilities, we find that either factor on its own is sufficient to bring about state time courses with a strong FO correlation (metastate) structure. Also, to a large extent, this effect is caused by between-subject differences in these statistics.

More specifically, we calculated the subject-specific TPMs from the estimated HMM state time courses, and simulated new state time courses of the same size as the original (estimated) state time courses. (Note that in the HMM results from the real data, there is no subject-specific modeling of state transition probabilities, which are modeled only at the group level). Then, we recomputed the TPM and the FO correlation. In a second set of experiments, we simulated state time courses using the group-level TPM (as provided by the HMM), thus ignoring between-subject differences. With this experiment, we aim to investigate which aspects of the TPM (on- or off-diagonal elements; *Methods*) are linked to the FO correlations, and to what extent this effect is due to between-subject differences. For each of these two cases, we carry out six variations of this scheme. In the first three, we suppress the influence of one of the three elements that define the HMM dynamics (initial state probabilities, on-diagonal elements of the TPM, and off-diagonal elements of the TPM; *Methods*) in turn. More specifically, we do so by setting the corresponding parameters to be equal to the prior distributions of the model, which are the same for all states and subjects and do not contain any information about metastates or hierarchies. In the other three, we keep just one of the three elements and remove the information from the other two.

First, the fact that we can faithfully reproduce the behavior seen in real data when we account for subject-specific differences in the TPM (Fig. S3A, first column) means that the strong structure in the FO correlation matrix can be explained, to a large extent, by Markovian statistics (i.e., transition probabilities and state persistence). Further, as observed in Fig. S3A, any two of these three elements of the HMM dynamics are sufficient to reproduce the block-wise FO correlation matrix (Fig. S3A; second,

third, and fourth columns). On the contrary, the metastate structure vanishes completely from the FO matrix when we suppress both the on- and the off-diagonal elements of the TPM (Fig. S3A, fifth column), indicating that the initial probabilities are not sufficient to produce this effect. Importantly, both the state persistencies (on-diagonal elements of the TPM) and the transition probabilities (off-diagonal elements of the TPM) can, in isolation, produce the block-wise FO correlation matrix (Fig. S3A, sixth and seventh columns), although more strongly in the case of the transitions (as reflected when comparing the third and fourth or sixth and seventh columns in Fig. S3A). When ignoring the between-subject differences in the state dynamics by using a single TPM (Fig. S3B), we observe that the correlations in the FO correlation matrix are massively decreased, although some hints of the metastates are still present when the transition probabilities are preserved (Fig. S3B; first, second, third, and sixth columns). Therefore, these simulations show that (i) the block-wise structure of the FO correlation matrix can be caused to a large extent by the between-subject differences in the TPM and (ii) both the on-diagonal (state persistencies) and off-diagonal (state transition probabilities) elements of the TPM are, on their own, sufficient to produce such structure. Hence, both the state persistencies and transition probabilities reflect the hierarchical nature of the state time courses.

**Reproducibility.** We assessed the reproducibility of the metastates and the sensitivity of the result to the choice of the model parameters, particularly to the number of states. To answer the first question, we divided the data into four blocks, where, for each of them, we take one-fourth of each subject's sessions (i.e., the first block will contain 1–400 s, the second 401–800 s, and so on) and then repeat the analysis. The two metastates clearly emerge in all four blocks, as illustrated in Fig. S6A. We then estimated new models from scratch using eight and 16 states. The correlation of the FO across subjects reflects a consistent two-metastate pattern (Fig. S6B). Altogether, these results demonstrate that the metastates are a reproducible and robust pattern of brain activity.

**The Autoregressive Model and the Metastates.** We now analyze the relationship between the autoregressive (AR) model and the HMM in the context of the metastates. To get some insight about whether the metastate information is contained in an AR model representation of the data, we designed a new set of simulations, where we generated four different simulated datasets. In each case, we fitted a specific model to the real data and then simulated data from the fitted model to generate the same amount of data as in the real dataset (i.e., 1 h of data in four sessions, with repetition time = 0.73 s). In the first scenario, we fitted a single MAR model of order 1 for all subjects. In the second simulation, we fitted a MAR model of order 1 for each subject separately. In the third simulation, for each subject, we fitted a univariate AR model of order 7 to each channel. In the fourth simulation, we fitted a multivariate Gaussian distribution for each subject. In summary, the first, second, and fourth simulations account for cross-region interactions; the second, third, and fourth simulations account for subject-specific differences; and the fourth simulation matches the observation model used by the HMM. We then fitted an HMM on the simulated data for each of the four types of simulations, with exactly the same configuration of parameters as in the real experiments. Note that whereas three of the simulations use 820 different models (one per subject), the HMM is itself a group-level model containing only 12 (Gaussian) models.

Fig. S7 summarizes the results. Importantly, the first simulation reveals that a group-level MAR is unable to generate data with metastates, providing confidence that the metastate result is not a trivial consequence of the average spectral properties of the BOLD signal captured by a global MAR. Further, neither the subject-specific univariate AR model (third simulation) nor

the subject-specific multivariate Gaussian distribution (fourth simulation) can generate data from which the HMM can find metastates, suggesting that (i) the two metastates can only be identified when we allow for differences in functional connectivity (which the univariate AR disregards) and (ii) the subject-specific multivariate Gaussian distributions are unable to capture the dynamics that also characterize each metastate. Interestingly, the HMM metastates do emerge to some extent when we use subject-specific MAR models to generate the simulated data (second simulation). This is because the metastate profile is very subject-specific (Fig. 4C), and the subject-specific features that (at least partly) drive the metastates are reasonably captured by the subject-level MAR models. This is not surprising, as the MAR model (containing 2,500 parameters per subject) has sufficient explanatory power to capture similar dynamics (35) to those captured by the HMM in a more parsimonious manner (using a few simpler observation models). Importantly, although a subject-specific multivariate AR model contains the information that is relevant to the metastates, these are not straightforwardly revealed in the parameters of this model.

**Using the HMM to Assess Nonstationarity.** We can assess the nonstationarity in the data using the simulations described in the preceding section. In particular, in the fourth simulation, we used a single (subject-specific) multivariate Gaussian distribution (inferred from the real data) to generate the simulated data and applied the HMM using the same configuration as in the real data. We then compared real and simulated data using the maximum FO per session, which reflects the amount of switching between states: If this value is close to 1.0 for a given session, then a single state is taking up most of the time; if lower, then different states are needed to explain the session. The distribution of this statistic is shown in Fig. S7C. For the real data (Fig. S7C, *Top*), the maximum FO distribution indicates that most sessions need more than one state to be described; in contrast, in the simulated data (Fig. S7C, *Bottom*), a large proportion of sessions are explained well with just one HMM state. Altogether, this shows that, when we account for estimation uncertainty, there is still evidence of nonstationarity in the sense of the data being generated from a multivariate Gaussian distribution.

**Relation to Activity Propagation.** Another interesting perspective on the data is the analysis of activity propagation (36, 37). This approach estimates, for each pair of regions, the lag at which the cross-covariance is maximal. The matrix decomposition of the resulting lag matrix yields the main “propagation threads” in the data. This approach has the benefit of offering an interpretable and robust characterization of this type of dynamics. However, could the metastates be captured by this type of analysis? The metastates have dwell times that range from 5 s to more than 1 min (Fig. S5C), and network dynamics at slow time scales can be accessed through post hoc analysis of the HMM state time course dynamics. It is not clear that such slow time scales would be accessible to activity propagation analysis. On the other hand, there may well be other characteristics of brain dynamics that can be better captured using activity propagation analysis compared with the HMM.

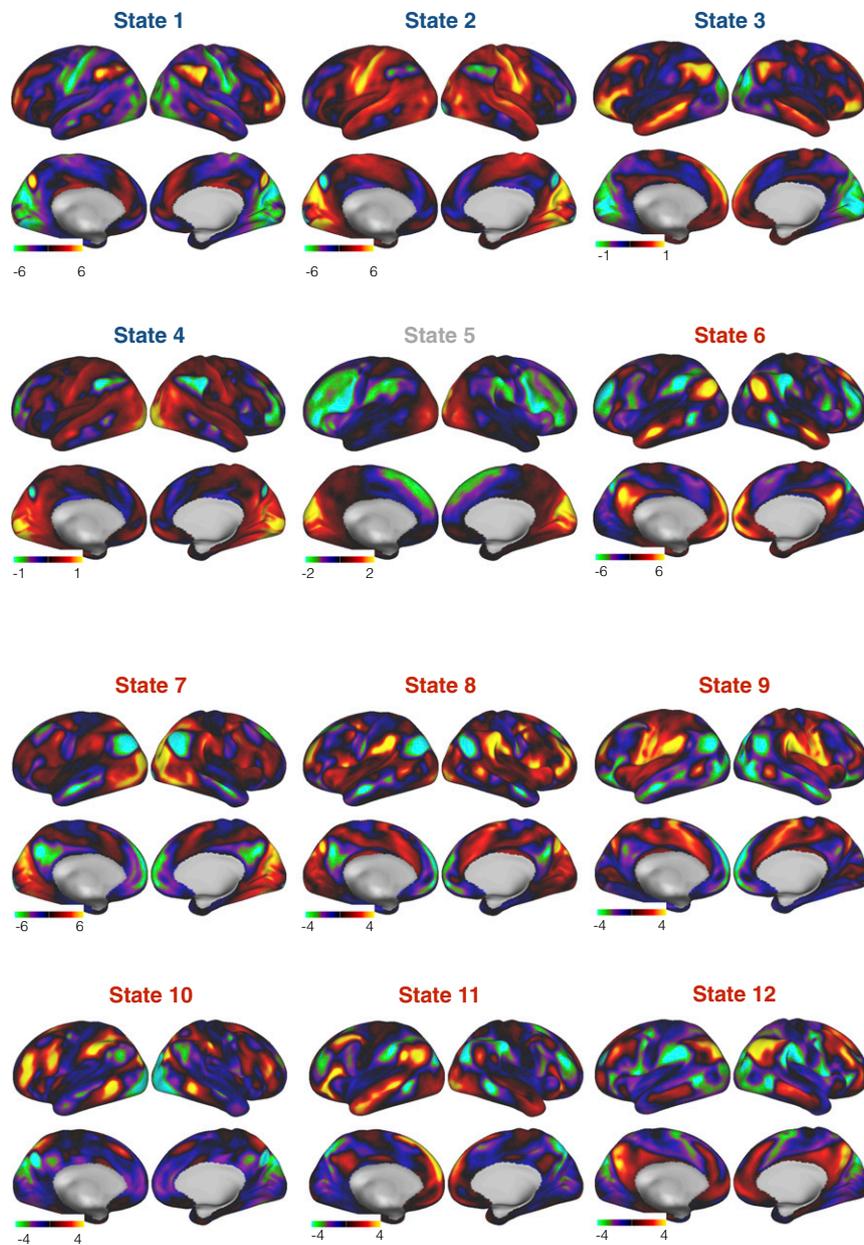
**Methodological Considerations.** The HMM makes two assumptions that require some consideration when applied to brain activity data. First, while it is the case that the inferred state assignments are probabilistic, such that the states can share the responsibility of explaining every time point (Fig. 1), the HMM does assume that the states are mutually exclusive (i.e., at each given time

point, only one state is assumed to be active). In theory, this might limit the number of possible expressible patterns to the number of HMM states. For example, how well could the HMM (using a discrete and not very high number of states) capture the brain activity in a cognitive task with varying difficulty, such that the magnitude of the activity varied continuously? One way in which graded activity could be represented in the HMM is through different state FOs (i.e., the proportion of time that the brain visits a particular state) at different time scales or, perhaps, through higher order state relationships such as the prevalence of certain sequences of states. Clearly, an alternative would be to dispense with the state exclusivity assumption. However, it is worth noting that this would require other assumptions (with their own limitations) to be made in order for the decomposition to remain identifiable.

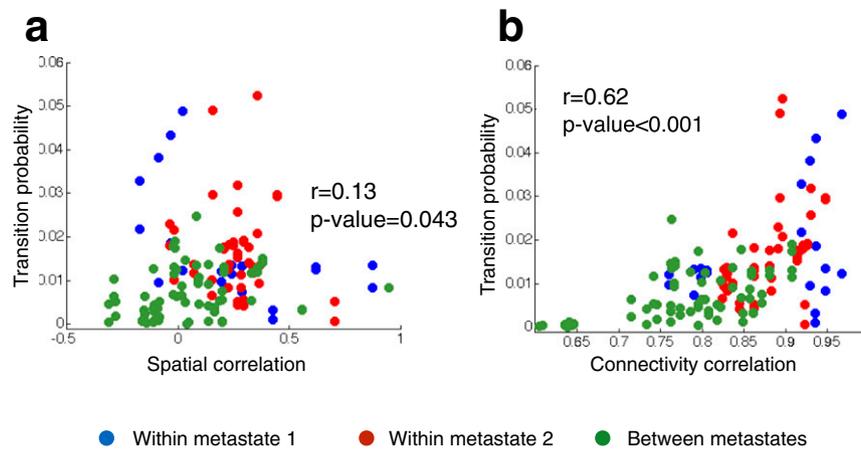
A second important assumption in the HMM is the Markovian assumption. This assumes that provided that we know the state at time  $t - 1$ , then we can predict what the state will be in at time  $t$  without needing to know the state time courses before time  $t - 1$ . This results in what is known as short-range dependency, where the dependency between state occurrences decays exponentially with time and does not reach very temporally distant data. This is potentially limiting, since there is good evidence that brain activity exhibits long-range dependency (8, 48). However, even though the HMM does not model long-range dependency, it does not prohibit it; in other words, the inferred HMM state sequence is free to discover long-range dependencies inherent to the data. Indeed, in this study, we have shown that the HMM is able to infer a certain type of long-range time dependencies in the form of metastates, even when these are not explicitly parameterized in the model.

An important question to bear in mind is the potential variability of the hemodynamic function across subjects and regions, or even within sessions. While the within-state, zero-lag, multivariate Gaussian distribution assumption is robust to this confound, the state time course dynamics may not be. Hence, despite the intrinsic regularization mechanisms imposed by the HMM on the dynamics, the hemodynamic remains a possible limitation of the current method. It is, however, unlikely that the hemodynamic variability explains the metastate structure reported in this study, given that the metastate dwell times (duration of the metastates' visits) are around 30 s, on average, with some visits being longer than 1 min (Fig. S5C), while the time scales of possible hemodynamic variation are much shorter [time to peak on the order of 2–6 s (49)].

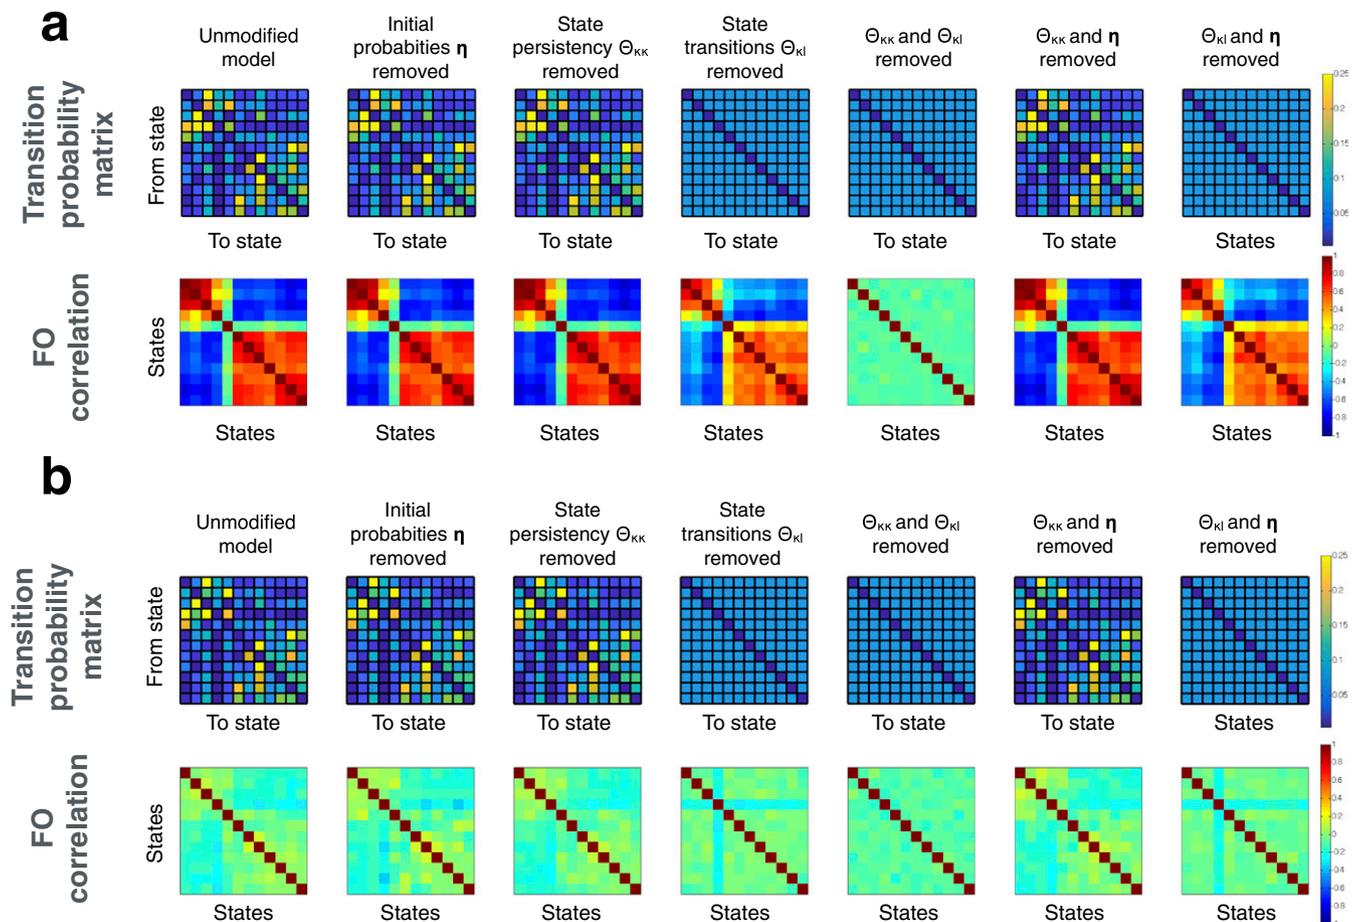
Finally, we can observe there are pairs of states that are significantly anticorrelated in their mean activity (Fig. S1). The most prominent example is state 1 and state 2, which also have relatively similar connectivity profiles and have larger amplitude in comparison to the other states within metastate 1. Insofar as the signal has a dominant pink noise (1/f) spectral component and the data have been demeaned, this is to be expected, given that the mean of the states (one of the parameters that characterize the states, together with the covariance matrix) is signed. Further, there is evidence that the brain exhibits alternating periods of high and low general synchronization and activity (10), which would explain why some of the states have a larger amplitude than others within each metastate. An alternative configuration of the HMM (not explored in this study) is to let the states be characterized only by functional connectivity (covariance matrix) by fixing the mean activity to zero. By focusing only on functional connectivity, the HMM results would be more directly comparable with standard sliding window analyses (21).



**Fig. S1.** Each metastate is composed of different brain states, with each defined by mean activation and functional connectivity. Each panel shows the mean activation for one state (functional connectivity is not shown). The first four states correspond to the sensorimotor/perceptual metastate. State 5 is considered to be independent of the metastates. States 6–12 correspond to the cognitive metastate.



**Fig. S2.** Relationship between the states' spatial characteristics and the HMM dynamics. Each dot in the panels corresponds to a pair of states. Whereas similarity in the mean activation (measured in terms of the correlation of their mean parameter) is mostly unrelated to the transition probability (A), similarity in functional connectivity between states (measured as the Pearson correlation of the off-diagonal elements of their covariance matrices) highly correlates with the probability of transitioning between them (B). This indicates that changes at the level of connectivity are smoother than changes at the level of absolute activation.



**Fig. S3.** Simulation experiments showing the relation between the TPM and the FO correlations. (A) In the first column, we estimated subject-specific HMM parameters from the empirical state time courses, and we generate synthetic state courses using these parameters for each subject; then, using such synthetic state courses, we recompute the TPM (Top) and the FO correlation matrix (Bottom) in the same manner as on the real data. Considering that the HMM dynamics are ruled by three different elements (initial state probabilities, on-diagonal elements of the TPM or state persistencies, and off-diagonal elements of the TPM or between-state transition probabilities), we repeated this procedure by removing the influence of one or two of these elements at a time; the resulting recomputed TPM and the FO correlation matrices are shown in the second to seventh columns. (B) We repeat this procedure generating data from the group HMM instead of from subject-specific parameters, such that all between-subject differences in the HMM dynamics are ignored; the metastates are still present in the FO correlation matrix when preserving the transition probabilities, but the effect is greatly reduced.









**Table S1. Behavioral traits**

Well-being	Intelligence	Personality
Anger-affect	Episodic memory	Agreeableness
Anger-hostility	Episodic memory (age-adjusted)	Openness to experience
Anger-aggression	Cognitive flexibility	Conscientiousness
Fear-affect	Cognitive flexibility (age-adjusted)	Neuroticism
Fear-somatic	Inhibitory control	Extraversion
Sadness	Inhibitory control (age-adjusted)	
Life satisfaction	Fluid intelligence accuracy	
Mean purpose	Fluid intelligence speed	
Positive affect	Reading	
Friendship	Reading (age-adjusted)	
Loneliness	Vocabulary	
Perceived hostility	Vocabulary (age-adjusted)	
Perceived rejection	Processing speed	
Emotional support	Processing speed (age-adjusted)	
Instrumental support	Spatial orientation	
Perceived stress	Attention TP	
Self-efficacy	Attention TN	
	Verbal episodic memory	
	Working memory	
	Working memory (age-adjusted)	
	Emotion recognition	

TN, true negatives; TP, true positives.