

Angus Bror Andersen Stevner

Jesus College, Oxford

University Department of Psychiatry

DPhil Thesis

Department of Psychiatry

Supervisors

Professor Morten L. Kringelbach

Professor Mark W. Woolrich

Professor Eus Van Someren

Whole-brain spatiotemporal characteristics of functional connectivity in transitions between wakefulness and sleep

This thesis is submitted to the Department of Psychiatry, University of Oxford, in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Trinity Term 2017

Contents

Contents ..... **iii**

List of figures ..... **vi**

List of supplementary figures ..... **viii**

List of tables ..... **ix**

List of abbreviations ..... **x**

Acknowledgements ..... **xii**

Abstract ..... **xiii**

Prelude ..... **xiv**

Chapter 1 Introduction **1**

1.1 **Sleep and the brain – the research programme** **1**

1.2 **Brain activity during sleep** **3**

1.3 **Low-level brain mechanisms of sleep** **5**

1.4 **Why it is important to investigate whole-brain activity during sleep** **10**

1.5 **Directions and challenges in the study of whole-brain activity of sleep** **14**

1.6 **Summary of the argument** **41**

Chapter 2 fMRI functional connectivity on top of structural connectivity during sleep **44**

2.1 **Introduction** **44**

2.2 **Methods** **45**

2.3 **Results** **51**

2.4 **Discussion** **60**

Chapter 3 Frequency-resolved functional connectivity of sleep using MEG **67**

3.1 **Introduction** **67**

3.2 **Methods** **68**

3.3 **Results** **78**

3.4 **Discussion** **88**

Chapter 4 Large-scale cross-frequency coupling in sleep using MEG **96**

4.1 **Introduction** **96**

4.2 **Methods** **97**

4.3 **Results** **101**

4.4 **Discussion** **107**

Chapter 5 Transient events in sleep functional connectivity using MEG.....111
5.1 **Introduction 111**
5.2 **Methods 111**
5.3 **Results.....112**
5.4 **Discussion 118**
Chapter 6 Dynamic functional connectivity of sleep using a hidden Markov model of fMRI…..
.....121
6.1 **Introduction 121**
6.2 **Methods 122**
6.3 **Results.....131**
6.4 **Discussion 150**
Chapter 7 Conclusion 160
7.1 **New insights.....160**
7.2 **Perspectives 164**
Appendix.....168
Datasets included in thesis 168
Supplementary figures 171
References.....181

List of figures

Figure 1-1 *EEG and polysomnography*……..4
Figure 1-2 *Overview of the low-level sleep-wake regulation system*…….. ...7
Figure 1-3 *Flip-flop switching between wake and sleep stages* 8
Figure 1-4 *Two-process model of sleep regulation* 9
Figure 1-5 *Examples of key findings from the "traditional" framework of sleep neuroimaging*....17
Figure 1-6 *Overview of resting-state networks.*....19
Figure 1-7 *Methods for deriving resting-state networks in fMRI.* 20
Figure 1-8 *Preservation of resting-state networks across NREM sleep stages.*....22
Figure 1-9 *Functional connectivity and graph theory*..…… 25
Figure 1-10 *Whole-brain FC network changes from wakefulness through NREM sleep*....28
Figure 1-11 *Construction of whole-brain SC networks*....30
Figure 1-12 *Whole-brain computational modelling.* 32
Figure 1-13 *Signal origins of fMRI BOLD and MEG.* 33
Figure 1-14 *Electrophysiological underpinnings of whole-brain functional connectivity.* 36
Figure 1-15 *Dynamic (as opposed to static) functional connectivity.* 40
Figure 2-1 *Group-averaged vigilance-specific FC matrices*....52
Figure 2-2 *Comparative FC network analyses between vigilance conditions* 55
Figure 2-3 *Global SC-FC similarities for two independent fMRI datasets*....56
Figure 2-4 *Regional determinants of SC-FC similarity for wakefulness and SWS.* 58
Figure 2-5 *Dynamic tracking of SC-FC relationship and mean FC level*....59
Figure 3-1 *Examples of artefactual ICs removed by pre-processing* 69
Figure 3-2 *Spectrogram of sensor recordings included in functional connectivity analysis.* 72
Figure 3-3 *Overview of analysis workflow for computation of vigilance-specific MEG FC.*. 75
Figure 3-4 *Group-averaged functional connectomes from wakefulness, N2, and SWS* 79
Figure 3-5 *Similarity of FC patterns within versus across vigilance conditions.* 80
Figure 3-6 *Similarity of envelope power distributions and functional connectivity degree distributions within versus across vigilance conditions* 82
Figure 3-7 *Degree maps from each of the three vigilance conditions*....83
Figure 3-8 *Vigilance-dependent differences in mean FC* 84
Figure 3-9 *Difference matrices from paired t-tests between vigilance conditions.*....85
Figure 3-10 *Highlights of significantly different sub-network identified with NBS* 87
Figure 4-1 *Deriving cross-frequency FC and multi-layer interactions in the super-adjacency matrix*....100
Figure 4-2 *Within-condition cross-frequency super-adjacency matrices*....103
Figure 4-3 *Paired t-tests between vigilance-specific cross-frequency correlations*....105
Figure 4-4 *Interactions between intra- and inter-layer FC.*....106
Figure 5-1 *Overview of relationships between sensor spectral power and mean levels of FC in six frequency bands*....113

Figure 5-2 *Time-frequency representations of relationship between sensor power and mean level of FC in representative participant*.....116
Figure 5-3 *Sensor timecourses of N2-specific transients from representative participant*. 117
Figure 6-1 *Variance explained in PCA decomposition*. 124
Figure 6-2 *Overview of HMM analysis pipeline*.....127
Figure 6-3 *Summary measures of HMM solutions across a range of model orders* 133
Figure 6-5 *Summary statistics of HMM states for K = 19 grouped within vigilance conditions* 137
Figure 6-6 *Summary of mean life times, mean FC, and SC-FC similarity*. 139
Figure 6-7 *HMM state transitions* 142
Figure 6-8 *First and second division of wakefulness represented as largest mean activity changes and differential FCs* 144
Figure 6-9 *N1-, N2-, and SWS-related states represented as mean activity distributions and differential FCs* 145
Figure 6-10 *Large-scale dynamics of wakefulness and sleep stages in the HMM state space* 149

List of supplementary figures

Figure S 1 *Matrix organisation of AAL regions-of-interest*. 171

Figure S 2 *Time-frequency representations of relationship between sensor power and mean level of FC in wakefulness from participant 1 to 4 (*Supplement to Figure 5-2*)*. 172

Figure S 3 *Time-frequency representations of relationship between sensor power and mean level of FC in wakefulness from participant 5 to 8 (*Supplement to Figure 5-2*)* 173

Figure S 4 *Time-frequency representations of relationship between sensor power and mean level of FC in N2 sleep from participant 1 to 4 (*Supplement to Figure 5-2*)*. 174

Figure S 5 *Time-frequency representations of relationship between sensor power and mean level of FC in N2 sleep from participant 5 to 8 (*Supplement to Figure 5-2*)*. 175

Figure S 6 *Time-frequency representations of relationship between sensor power and mean level of FC in SWS from participant 1 to 4 (*Supplement to Figure 5-2*)*. 176

Figure S 7 *Time-frequency representations of relationship between sensor power and mean level of FC in SWS from participant 5 to 8 (*Supplement to Figure 5-2*)*. 177

Figure S 8 *Mean activity maps and differential FC networks of HMM states 1 to 9 (*Supplement to Figure 6-8 and Figure 6-9*)*. 178

Figure S 9 *Mean activity maps and differential FC networks of HMM states 10 to 19 (*Supplement to Figure 6-8 and Figure 6-9*)*. 180

List of tables

Table 2-1 *Overview of indexing of SC and FC matrices based on the AAL parcellation.* 53

List of abbreviations

AAL	Automated Anatomical Labeling
ACN	Anti-Correlation Network
BOLD	Blood-Oxygenation-Level-Dependent
DAN	Dorsal Attention Network
DMN	Default Mode Network
DTI	Diffusion Tensor Imaging
ECN	Executive Control Network
EEG	Electroencephalogram
FC	Functional Connectivity
fMRI	functional Magnetic Resonance Imaging
FSL	FMRIB Software Library
HCP	Human Connectome Project
HMM	Hidden Markov Model
ICA	Independent Component Analysis
MANOVA	Multivariate Analysis of Variance
MEG	Magnetoencephalography
MNI	Montreal Neurological Institute
N1	NREM sleep stage 1
N2	NREM sleep stage 2
N3	NREM sleep stage 3 or SWS
NBS	Network-Based Statistic
NREM	Non-Rapid-Eye Movement
PCA	Principal Component Analysis
PET	Positron Emission Tomography
REM	Rapid-Eye Movement
ROI	Region-Of-Interest
RSN	Resting State Network
SC	Structural Connectivity
SWS	Slow Wave Sleep
VAN	Ventral Attention Network

Acknowledgements

Many times over the course of my doctorate I have felt somewhat detached from the world around me, fully absorbed in things related to my own work, interests, ambitions, and degree. Yet, a special group of individuals have managed to break through my isolation to show their support and care. I want to take this opportunity to express my gratitude for all of these professional and personal encouragements.

First, I have to thank my supervisor, Morten, to whom I owe the opportunity to pursue my dream of studying the human brain. His vision and scientific versatility continues to spark ideas, and he has always been compassionate and supportive whenever I have struggled. These qualities are reflected in his research group, Hedonia, which is composed of extraordinary individuals of far-reaching interests, talents, and spirits. Christine, Katie and Maria at the beginning of my time and later Henrique, Patricia, Joana, Eloise, L.D., Alexander, Marina, and Søren have all contributed to an inspirational and welcoming atmosphere that I have always enjoyed as my scientific home. A special thanks should go to Tim, the 'duct-tape' of the Hedonia lab, whose help and guidance I will not even attempt to summarise here.

My co-supervisor, Mark, should be highlighted for providing me with clear and rigorous advice, and his ideas for approaching data analyses in original ways. I admire his team, the OHBA analysis group, for their insight and skills as well as for their abilities to communicate and teach these with people like myself of less technical background. Many of my findings and discoveries would not have been possible without the helpful and inspiring interactions with members like Diego Vidaurre, Giles Colclough, Hamid Mohseni, and Adam Baker.

Eus, my second co-supervisor in Amsterdam has, besides providing me with unique data, been highly influential for my ideas regarding the nature of sleep. I consider myself privileged to have had the opportunity to discuss such matters with a leading expert in the field, and people from his group, of which particular mention should be given to Gio Piantoni.

I have benefited from Morten's close collaboration with Gustavo Deco in Barcelona. The unparalleled advances made by him and his group within whole-brain computational modelling and theoretical concepts is a continuing inspiration to me and a driver for my future ambitions.

None of my work would have been possible if Frederikke Rosendal, then at Aarhus University, had not offered me to come and work in Oxford as part of her project, funded by the Lundbeck foundation, or if Leif Østergaard of CFIN/MindLAB at Aarhus University had not decided to generously fund my DPhil stipend.

While I have strong feelings for Oxford, it has not always been easy to be away from 'home' in Copenhagen. Almost four years has past since I left, and this time has taught me the value of family and friendships. Damien, Carolyn, and Dan have all done their part to fill the void, but I have always longed for the love and care of my mother, Lene, father, Michael, sister, Katinka, and step-dad, Erik. I would not have made it far without your support and belief in me.

To Tine – we have endured four years of living apart, not at least owing to you patience – it feels like nothing can stop us going forward, and I cannot wait!

Abstract

Whole-brain spatiotemporal characteristics of functional connectivity in transitions between wakefulness and sleep

Angus Bror Andersen Stevner

Jesus College

This thesis provides a novel dynamic large-scale network perspective on brain activity of human sleep based on the analysis of unique human neuroimaging data. Specifically, I provide new information based on integrating spatial and temporal aspects of brain activity both in the *transitions between* and *during* wakefulness and various stages of non-rapid-eye movement (NREM) sleep. This is achieved through investigations of inter-regional interactions, functional connectivity (FC), between activity timecourses throughout the brain. Overall, the presented findings provide new important whole-brain insights for our current understanding of sleep, and potentially also of sleep disorders and consciousness in general.

In Chapter 2 I present a robust global increase in similarity between the structural connectivity (SC) and the FC in slow-wave sleep (SWS) in almost all of the participants of two independent fMRI datasets. This could point to a decreased state repertoire and more rigid brain dynamics during SWS. Chapter 2 further identifies the changes in FC strengths between wakefulness and individual stages of NREM sleep across the whole-brain fMRI network. I report connectivity in posterior parts of the brain as particularly strong during wakefulness, while connections between temporal and frontal cortices are increased in strength during N1 and N2 sleep. SWS is characterised by a global drop in FC.

In Chapter 3 I take advantage of rare MEG recordings of NREM sleep to show, for the first time, the feasibility of constructing source-space FC networks of sleep using power envelope correlations. The increased temporal information of MEG signals allows me to identify the specific frequencies underlying the FC differences identified in Chapter 2 with fMRI. The beta band (16 – 30 Hz) thus stands out as important for the strong posterior connectivity of wakefulness, while a range of frequency bands from delta (0.25 – 4 Hz) to sigma (13 – 16 Hz) all appear to contribute to N2-specific FC increases. Consistent with the fMRI results, slow-wave sleep shows the lowest level of FC. Interestingly, however, the MEG signals suggest a fronto-temporal network of high connectivity in the alpha band, possibly reflecting memory processes.

In Chapter 4 I expand the within-frequency FC analysis of Chapter 3 to explore potential cross-frequency interactions in the MEG FC networks. It is shown that N2 sleep involves an abundance of frequency cross-talk, while SWS includes very little. A multi-layer network approach shows that the gamma band (30 – 48 Hz) is particularly integrated in wakefulness.

Chapter 5 addresses the identified MEG FC findings from the perspective of traditional spectral sleep staging. By correlating temporal changes in spectral power at the sensor level to fluctuations in average FC, a specific type of transient events is found to underlie the strong N2-specific coupling in static FC values.

Lastly, in Chapter 6 I make the leap out of the constraints of traditional low-resolution sleep staging, and extract dynamic states of FC from fMRI timecourses in a completely unsupervised fashion. This provides a novel representation of whole-brain states of sleep and the dynamics governing them. I argue that data-driven approaches like this are necessary to fully characterise the spatiotemporal principles underlying wakefulness and sleep in the human brain.

This thesis is submitted for the degree of Doctor of Philosophy at the University of Oxford.
Trinity term 2017

Prelude

The mystifying nature of sleep stems from the fact that all knowledge pertains to waking life. When I am awake I know things about the outside world and myself, but when asleep that sense of knowledge is gone, or at least impossible to recollect. The concurrent lack of mental access, or at least the absence of recollection of any such access, during sleep counteracts the sense of knowledge or understanding of what sleep actually is. The function of my hands, my mouth, or even more complex things like my memory, seem intuitive and familiar to me, because I have gained knowledge and learned about them through my repeating use of these during wakefulness. I may not understand how memory works, but I can feel it working away, assisting me in my daily activities. Sleep is not amenable to similar introspection. It can be used as 'an efficient cure against sleepiness', but in many ways that is where comprehension of sleep stops. This is truly remarkable, given that a third of an average human life is 'spent' in this state.

Chapter 1
Introduction

"Sleep is of the brain,
by the brain
and for the brain "
– J. Allan Hobson

" Those who speak
Do not know
Those who know
Do not speak "
– Lao Tse

1.1 ***Sleep and the brain – the research programme***

Let us perhaps begin by thinking about sleep from a classical behaviourism perspective. This can be over with fairly quickly, since sleep is essentially characterised by the absence of behaviour. There are subtleties to this behavioural absence, however, that make sleep different from mere resting. The responsiveness to the external environment is decreased, and characteristic body postures are assumed (Zepelin et al., 2005). Relative to states of similar disconnect from the external world, such as hibernation, coma, and general anaesthesia, sleep can be reversed relatively easily. Further, sleep is not an exceptional state, but occurs regularly during the 24-hour day, and if it for some reason, gets hindered it is compensated by more or rather intensified subsequent sleep – a phenomenon referred to as sleep homeostasis. These behavioural characteristics have been used to identify sleep in a very wide range of animal species, spanning most of the phylogenetic tree (Campbell and Tobler, 1984, Siegel, 2005, Allada and Siegel, 2008, Joiner, 2016). This suggests that sleep as a behavioural contrast to wakefulness has survived the long history of evolution. Meanwhile, sleep does not seem to represent any overt competitive advantage, compared to features of wakefulness like feeding and mating. These observations have led to scientific bewilderment. As Allan Rechtschaffen expressed it in 1978: "If sleep does not serve an absolutely vital function, then it is the biggest mistake the evolutionary process has ever made" (Walker, 2009b). The *vital function* that Rechtschaffen was after, can be referred to as the primordial function of sleep, and points to a fundamental purpose that is shared among species, from fruit flies to humans. While the search for this has occupied many and resulted in a range of proposals (Siegel, 2005, Cirelli and Tononi, 2008, Mignot, 2008), it is important to note that sleep may very well have evolved opportunistically to serve many different functions in different species (Krueger et al., 2016). It is therefore possible to speak about a function of sleep, without referring to *the* function.

While behaviour during sleep itself is sparse, it is possible to derive the importance of sleep by studying what happens when it is restricted. As such, evidence from humans and animals have established that sleep is important for a wide range of peripheral functions within immune, endocrine, autonomic, cardiovascular, and thermoregulatory systems (Stickgold and Walker, 2009). However, a majority of research has focussed on the specific effects that sleep exerts on the brain. Lack of sleep has thus been linked to a range of brain-related dysfunctions in cognition (Banks and Dinges, 2007), memory and learning (Stickgold, 2005), emotional regulation and mood (Walker, 2009a), as well as development (Hobson and Pace-Schott, 2002).

Psychiatric disorders are very commonly accompanied by disturbances in sleep. While these were previously thought to be side effects of the mental conditions, it is becoming increasingly clear that sleep may in fact be central to the underlying pathogeneses (Wulff et al., 2010, Jagannath et al., 2013).

It is thus natural to ask what the brain might be doing in the behaviourally quiescent state of sleep, and how the brain is able to support it. It first became possible to address such questions in the 1920's, following Berger's invention of the electroencephalogram (EEG) (Berger, 1929) (Figure 1-1A), which allowed investigators to non-invasively record the electrical activity of the brain – an event that is often referred to as the birth of modern sleep research (Stickgold and Walker, 2009). It was soon realised that wakefulness and sleep were associated with markedly different oscillatory patterns expressed in the EEG traces, and by the 1930's the group led by Loomis had made the first attempts to objectively describe sleep based on such differences in brain activity (Loomis et al., 1935a, b, Davis et al., 1937, Loomis et al., 1937) (Figure 1-1B).

Today our understanding of sleep remains intimately linked with brain recordings. This has allowed us to move beyond the behavioural characterisation, in humans and other animals with developed central nervous systems. This is exemplified by polysomnography, which combines EEG with recordings of other physiological parameters to classify sleep into a number of stages (see 1.2.1). The oscillatory signatures of the EEG that are used for such sleep staging have been studied extensively through intracranial recordings in non-human animals, aiming to identify their generative mechanisms (see 1.3.1). We also have a very good understanding of how areas located deep in the brain are able to control and regulate transitions between wakefulness and sleep, as well as their maintenance (see 1.3.2). In sum these advances have firmly established that sleep is a phenomenon *of* the brain, *by* the brain, while the effects of sleep deprivation suggests that sleep is also *for* the brain (Hobson, 1989, Hobson, 2005).

Sleep has thus become a central topic of modern neuroscience. Does this mean that we, as a scientific community, have evaded the difficult explanatory situation of our imagined classical behaviourist? Not quite. The lack of willed behaviour and reportability from a sleeping participant represents an inherent barrier to scientific enquiry. What we *have* gained, are some very good correlates of sleep. Nevertheless, translating these correlates into something meaningful for sleep's functions remains a difficult challenge.

While EEG recordings are reliable markers of sleep the patterns of stage-specific amplitude and spectral content, typically obtained from a limited number of EEG electrodes, represent a crude representation of the brain's immense amount of activities. Recently, new concepts, driven by the advances in modern neuroimaging techniques, have been emerging in which whole-brain activity is characterised and modelled in spatial detail as large-scale networks of both structure and function. This framework has proven promising for expanding our understanding and characterisation of the brain during wakefulness, and more recently begun to provide a richer picture of the brain's activities during sleep.

In this Introduction I will summarise the current knowledge about brain activity during sleep in terms of its classification in polysomnography, thalamocortical mechanisms of generation, and low-level regulations. I will then move on to a discussion of why a whole-brain view is necessary, and review recent progress towards this. Finally, I will point to outstanding issues and unresolved matters, which this thesis is trying to address.

1.2 ***Brain activity during sleep***

1.2.1 *Polysomnography*

Polysomnography denotes a consensus-based classification of sleep into stages (Rechtschaffen and Kales, 1968, Iber et al., 2007). It relies on measures of brain activity, eye movements, respiration, muscle tone, and heart rate, and on the observation that this set of physiological variables co-vary in characteristic patterns as a function of vigilance. The measurement tools include a few channels of EEG placed in strategic positions on the scalp, electrooculography (EOG), breathing counts, electromyography (EMG), and electrocardiography (ECG). The resulting characteristic patterns are commonly categorised into two main components, rapid eye movement (REM) sleep and non-REM (NREM) sleep. NREM sleep is further divided into either four (NREM sleep stage 1 to 4) (Rechtschaffen and Kales, 1968) or three (N1, N2, and slow wave sleep (SWS)) (Iber et al., 2007) stages. Staging is typically performed 'off-line' by a trained sleep scorer, who examines the multimodal recordings in consecutive windows of 30 seconds, labelling each of them as a particular stage. The EEG derivations are examined in terms of frequency content and phasic patterns, such as spindles, K-complexes, and delta waves. Overall, the EEG becomes more synchronised through NREM sleep, while it regains its desynchronised appearance, and is quite similar to wakefulness, during REM sleep. A more detailed description of the EEG characteristics and a summary of accompanying physiological changes are given in Figure 1-1.

Behaviourally, sleep is often placed on a scale from 'light' to 'deep', based on arousal thresholds (Carskadon and Dement, 2005). As such, it becomes gradually more difficult to wake a person from N1, over N2, to SWS. Arousal thresholds during REM sleep are comparable to those of N2 sleep. The amount of mental content that a person will report upon awakening also depends on the sleep stage. The border between wakefulness and sleep, experienced during sleep onset, can include a highly fluctuating amount of mentation and awareness of the external, which often leads to sleep misperception (Ogilvie, 2001, Goupil and Bekinschtein, 2012, Nir et al., 2013). The vivid imaginative, and sometimes hallucinating experiences, commonly referred to as dreams, are associated with REM sleep. If mental content is reported from NREM sleep, it is of a more musing kind, with a lack of clear visual-emotional quality (van Someren and Cluydts, 2013). The fact that the physiologically defined patterns are followed by behavioural changes, in terms of arousal, responsiveness, and reportability, is an important argument for the traditional staging of sleep.

When sleep is scored over a full night, the overview of the visited stages is referred to as a hypnogram, which describes the sleep architecture. The sleep architecture in the young and healthy human is characterised by a mechanistic cycling between "lighter" and "deeper" stages of NREM sleep (see Figure 1-1C). The entry into REM- from NREM sleep marks the end of an ultradian sleep cycle, which normally expresses a period of approximately 90 – 110 minutes. Left undisturbed sleep will typically cycle through NREM and REM about 5 times. The distribution of REM and NREM sleep changes over the course of sleep. As such, SWS is more prevailing early in the night, while REM sleep epochs become gradually extended towards the end.

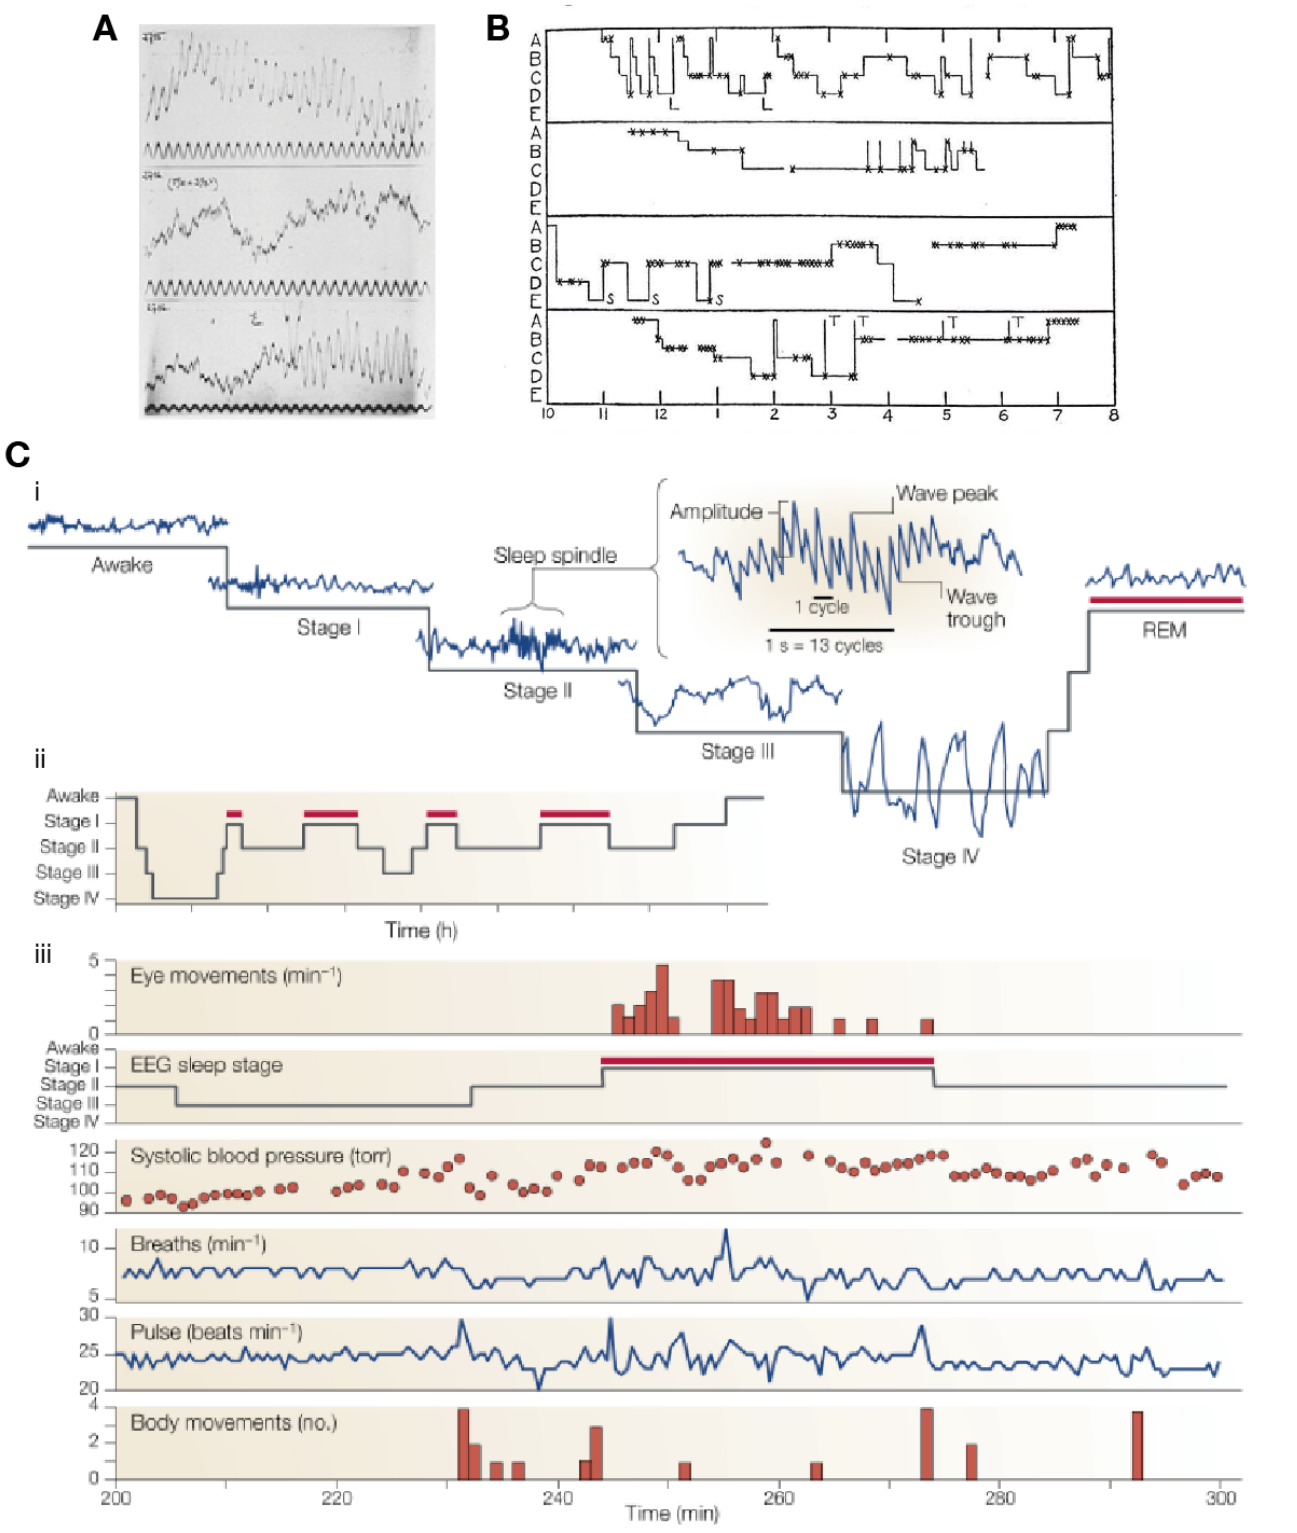

Figure 1-1 ***EEG and polysomnography*** **A)** An example of an early EEG trace obtained by Berger in the 1930's. **B)** Already in the 1930's Loomis and his colleagues suggested a division of human sleep into five stages, based on EEG traces. (from Loomis et al. (1937)) **C)** Overview of modern polysomnography. **i**: Windows are scored as *wakefulness* if they show a generally desynchronised EEG with predominance of power in the beta (13 – 25 Hz) and gamma (25 – 80 Hz) ranges, and when eyes closed a presence of a waxing and waning alpha (8 – 13 Hz) oscillation. *NREM stage 1* or *N1 sleep* is identified based on a disappearing alpha oscillation and a general slowing of the predominant spectral components. Theta activity (4 – 8 Hz) at relatively modest amplitudes becomes the principal pattern of the EEG. *NREM stage II* or *N2 sleep* is associated with a gradual increase in amplitude and lowered frequencies of the EEG waveforms compared to N1. The occurrence of K-complexes and sleep spindles is defining for N2. A K-complex is a large negative deflection in the EEG followed immediately by a smaller positive deflection, resembling the letter "K". A sleep spindle is a 1 – 2 seconds burst of activity in the sigma frequency range (12 – 16 Hz), with a waxing and waning pattern. *NREM stages III and IV* or *SWS* is characterised by slow waves in the EEG, fewer spindles and K-complexes, a

predominance of high amplitude oscillations with a main frequency component in the delta range (1 – 4 Hz). *REM sleep* (depicted as red bars in **ii** and **iii**) EEG becomes more similar to that of wakefulness. While NREM sleep can be recognised more or less solely on the basis of the EEG, REM-sleep scoring requires information from the supplementing physiological measures (see below). The most prominent frequency in REM is theta, with higher frequencies intermixed. Muscle tone, as measured with EMG, decreases from wakefulness through NREM sleep, while REM sleep is characterised by complete muscle atonia (apart from eye, ear, and respiratory). Eye movements decrease during NREM sleep, albeit present in the form of slow rolling during N1, but reappear in REM as saccadic rapid eye movements. Heart rate and respiration are also decreased during NREM sleep, and increased to wake-like levels during REM Figure adapted from: Pace-Schott and Hobson (2002).

1.3 ***Low-level brain mechanisms of sleep***

1.3.1 *Cortex, thalamus and sleep-specific oscillations*

A move closer to the neuronal sources through intra-cranial as well as intra-cellular recordings in non-human mammals, has contributed with a more detailed level of description. The so-called *slow oscillation* is a prime example of a fundamental principle of brain activity not immediately recognisable in extracranial EEG recordings (Achermann and Borbély, 1997). Representing a slow (< 1 Hz) alternation between neuronal states of depolarised activity and hyperpolarised quiescence, termed UP and DOWN states, respectively, the slow oscillation was originally discovered in thalamic and cortical neurons of cats during anaesthesia and SWS (Steriade et al., 1993a, Steriade et al., 1993b, c). Soon thereafter it was re-discovered in humans, and the slow oscillation is now known not only as a hallmark of SWS, but also as an orchestrator of most of the rhythmic EEG signatures of all NREM sleep. K-complexes have thus been shown to represent DOWN states of isolated slow oscillations (Cash et al., 2009), while delta oscillations, sleep spindles as well as faster oscillations occur in phase-locked relationships to the UP state (Steriade, 2006). The exact mechanisms underlying the generation of the slow as well as the other sleep-defining oscillations are still being discussed. A unifying view, however, is that they all, to some extent, derive from network interactions between the thalamus, with its specific and non-specific nuclei, and the various layers of the cortex (Datta, 2012, David and Schmiedt, 2014, Lemieux et al., 2014, Luthi, 2014, Crunelli et al., 2015, Piantoni et al., 2016a).

Moreover, human as well as animal intra-cranial recordings have established that sleep oscillations can occur very locally and independently throughout the cortex (Vyazovskiy et al., 2009, Andrillon et al., 2011, Nir et al., 2011), echoing the peculiar phenomenon of uni-hemispheric sleep seen in certain marine mammals, such as the dolphin (Lyamin et al., 2008). While local sleep cannot be resolved in a few EEG electrodes, evidence from high-density EEG have shown that low-frequency activity is not synchronously distributed throughout the cortex. “Slow waves”, a common term covering frequencies in the range of both slow and delta oscillations (Olbrich et al., 2011), have thus been found concentrated towards frontal topographies of the EEG (Finelli et al., 2001). Furthermore, slow oscillations have been modelled as so-called travelling waves across the brain (Massimini et al., 2004, Murphy et al., 2009b). Sleep spindles have been split into ‘fast’ (13 – 16 Hz), and ‘slow’ (11 – 13 Hz), based on their EEG topographical distributions (De Gennaro and Ferrara, 2003), and MEG source-reconstructions have demonstrated the asynchronous appearance of spindles throughout the cortical source-space (Dehghani et al., 2010a, Dehghani et al., 2010c).

1.3.2 *Basic mechanisms of sleep-wake regulation*

Since sleep is normally a highly regular phenomenon it may be deduced that its stages must be under some form of regulation. Throughout the past century lesion (Moruzzi and Magoun, 1949, Kerkhofs and Lavie, 2000) and pathological (Von Economo, 1930) studies have pointed to structures, located deep in the brain, facilitating this regulation through modulations of cortico-thalamic networks. These structures, primarily located in the brainstem, hypothalamus, and basal forebrain have been investigated extensively in rodents and cats. With modern techniques (see Weber and Dan (2016)) including regional lesions, electrical stimulation, pharmacological manipulations, immunohistochemistry, and more recently with cell-type specificity, optogenetics, pharmacogenetics, and calcium-imaging, this has led to a detailed mapping of low-level mechanisms of sleep regulation.

It is common to structure current knowledge on the topic into wake promotion, NREM sleep promotion, and REM sleep promotion (Saper et al., 2010, Brown et al., 2012, Weber and Dan, 2016). The identified sleep and wake-promoting neuronal populations are intricately coupled through mutual inhibitions and interactions, which serve to stabilise the cortical brain states associated with the different stages of sleep and wakefulness, as well as the switching between them. The term 'promotion' should here be understood as a term describing areas increasing their neural activity in association with the specific stage of vigilance.

Wake-promoting areas are located in the brainstem, hypothalamus, and basal forebrain. Neurons in these areas act through their neurotransmitters, which include acetylcholine, serotonin, noradrenalin, dopamine, histamine, and orexin/hypocretin. An overview is given in Figure 1-2. The modulating effects of these areas on the thalamus and the cortex directly, act to produce the desynchronised EEG of wakefulness (see Figure 1-2A).

NREM sleep is promoted by neuronal populations in the anterior hypothalamus, including the ventrolateral preoptic nucleus (VLPO) and the median preoptic nucleus (MnPO). The neurons involved express the inhibitory neurotransmitters, GABA and galanin. The sleep-inducing effects are mainly exerted through inhibition of the wake-promoting circuits (see Figure 1-2B).

Compared to that of NREM sleep, the EEG of REM sleep is more similar to that of wakefulness. One reason for this is that while all of the wake-promoting neurotransmitters are lowered in concentration during NREM sleep, acetylcholine is in fact high during REM sleep. REM sleep generators have been identified in the pons and medulla of the brainstem. These engage in inhibitory feed-back loops with the wake as well as the NREM-promoting areas (see Figure 1-2D-F).

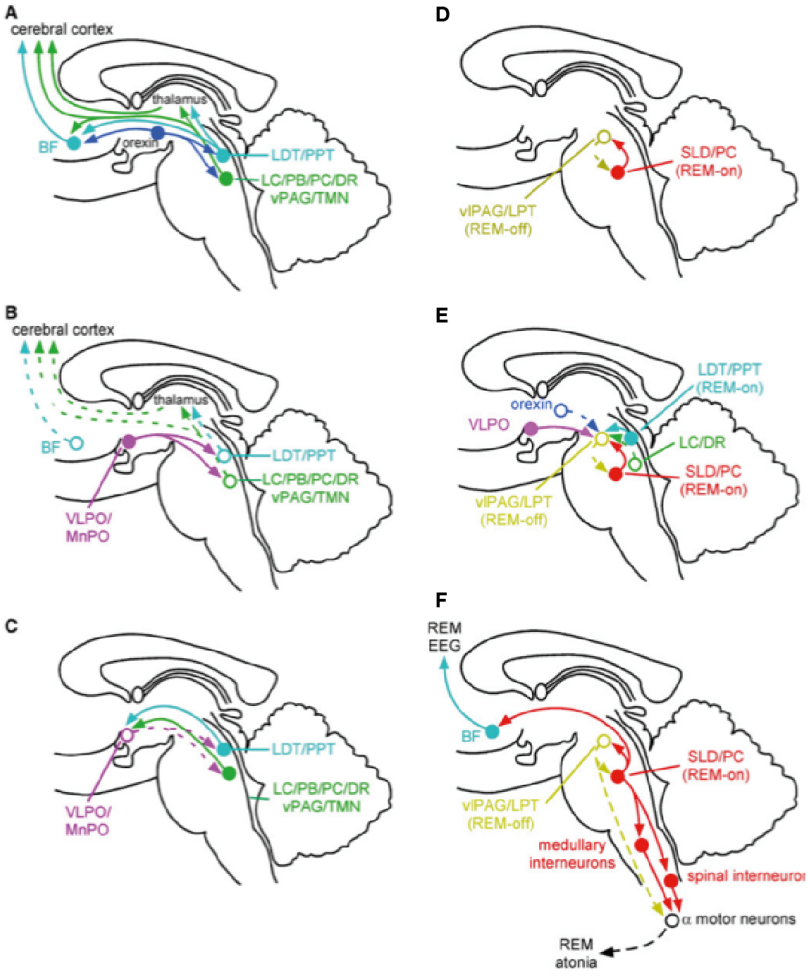

Figure 1-2 ***Overview of the low-level sleep-wake regulation system*** The *wake-promoting* areas (**A**) may be grouped according to their neurotransmitter, and include cholinergic neural populations in the basal forebrain, the pedunculopontine tegmental nucleus (PPT) and the laterodorsal tegmental nucleus (LDT), both in the pons; noradrenergic neurons in the locus coeruleus (LC) also at the pontine level; serotonergic neurons in the medial and dorsal raphe nucleus (DR) between the medulla and the pons; dopaminergic neurons around the ventral periaqueductal gray close to the third ventricle (vPAG); glutamatergic neurons in the parabrachial nuceus (PB) and the precoeruleus area (PC) at the junction between medulla and pons; histaminergic neurons in the tuberomamillary nucleus in the posterior hypothalamus; and finally orexin/hypocretin-producing neurons in the posterior part of the lateral hypothalamus, containing glutamate. The influence and control over the cortical state and therefore vigilance/arousal of the listed wake-promoting areas act through projections directly to the cortex and thalamus, or indeed to the other wake-promoting areas. The general effect is a desynchronised EEG associated with wakefulness. *NREM sleep promotion* (**B-C**) involves fewer areas than wake-promotion. Specifically, the ventrolateral preoptic nucleus (VLPO) and the median preoptic nucleus (MnPO) have been found to have strong sleep-inducing effects (see Saper et al. (2010)). These neural populations contain inhibitory neuromodulators, such as GABA and galanin, which act to supress the activity in the wake-promoting areas, in particular the monoaminergic neurons in the TMN, DR and LC. Sleep promotion is thus facilitated mainly through inhibition of wake-promoting areas. Conversely, wakefulness is dependent on inhibition of the sleep-promoting regions in the anterior hypothalamus. Hence, many of the wake-promoting areas have been shown to project to and inhibit the VLPO. *REM sleep promotion* (**D-F**) involves neurons located caudally, in the sublateral dorsal nucleus (SLD) and the precoeruleus expressing glutamate. While the latter acts on the basal forebrain to produce EEG desynchronisation, the former sends descending projections to the medulla and spinal chord to produce the muscle atonia that is a hallmark of the REM state. The SLD receives strong GABAergic inhibitory inputs from the vlPAG. Some of these medullary areas, in the ventral medulla (vM) send rostral GABAergic projections to the vlPAG and the LC, inhibiting these, and adding an important component to the loop of REM sleep promotion (not shown here, but see Weber and Dan (2016)). Finally, a group of neurons in the lateral and posterior hypothalamus producing the peptide melanin-concentrating hormone (MCH), which innervate many of the same areas as the wake-promoting orexin/hypocretin-producing neurons that are located in close proximity. Unlike the orexin/hypocretin neurons, which contain the excitatory transmitter glutamate, the MCH neurons contain the inhibitory GABA. MCH neurons have been found to increase their firing during REM, but are also active during NREM sleep. (from Saper et al. (2010))

The inhibitory relationships between the state-promoting groups of neurons also offer a possibility to formulate models explaining the transitions between the respective vigilance stages. A particularly well-established model for these transitions with its machinery centred in the hypothalamus and brain stem, uses of a flip-flop switch analogy to describe how the mutual inhibitions facilitate transitions (see Figure 1-3). The flip-flop switching may be thought of as a balancing scale with factors promoting different states weighing on each side of it (van Someren and Cluydts, 2013).

Importantly, models for REM-control have to account for switches to both NREM and wakefulness, while mainly from NREM, since these are the most common types of transitions in the hypnogram. Computational simulations have implemented mathematical formulations of the flip-flop mechanisms to validate their influence on the thalamocortical system and the production of the fundamental oscillations of wakefulness and NREM sleep (Robinson et al., 2011).

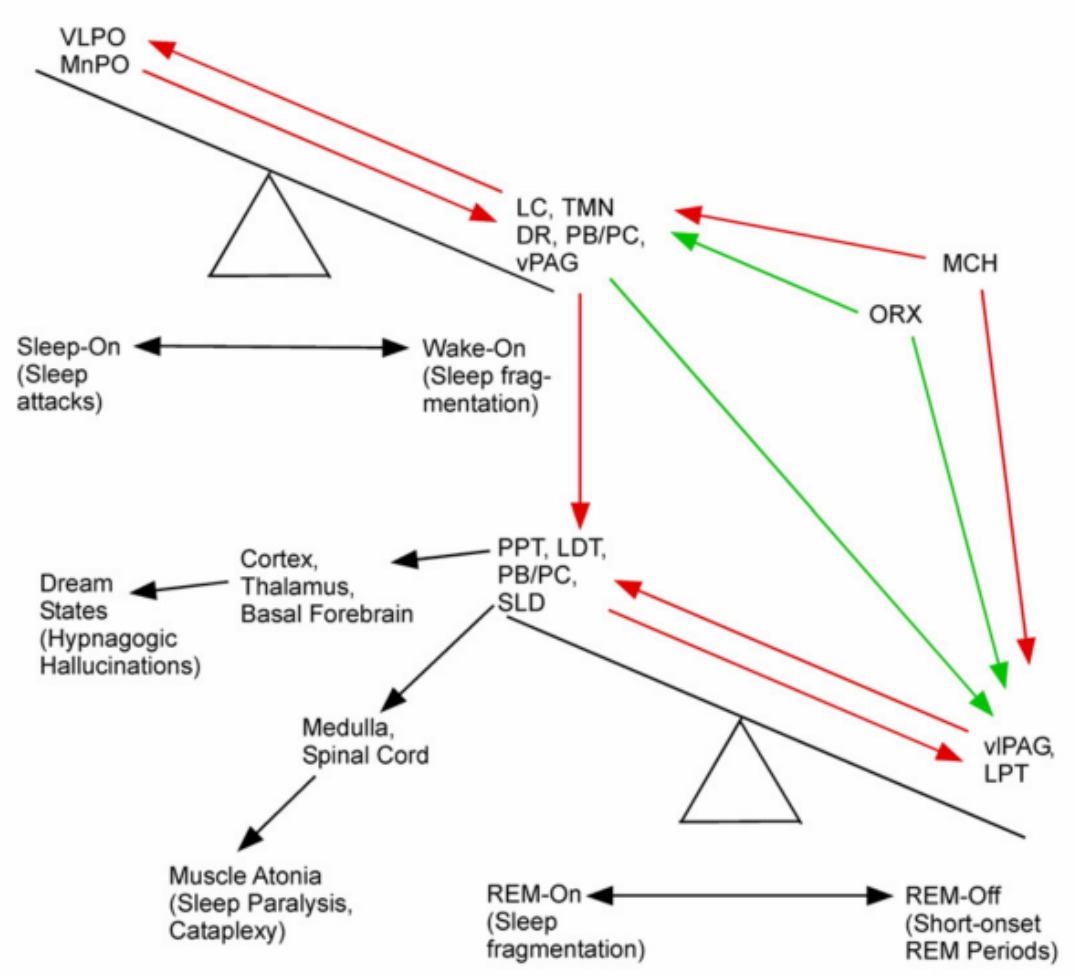

Figure 1-3 ***Flip-flop switching between wake and sleep stages*** The wake-to-NREM transition is especially facilitated by the inhibitory effect on VLPO of the neurotransmitters involved in wake-promotion and the presence of projections from these areas to the VLPO. VLPO in turn projects its GABA heavily to and inhibits the TMN neurons, thus counteracting their wake promotion and drive of other wake-active areas. In addition, the MnPO uses GABA to inhibit many of the wake-promoting areas, including DR, LC, and vPAG, Less is generally understood about the switching to REM sleep. As such, this phenomenon has both been modelled as a reciprocal interaction, or "predator-prey", model, involving the monoaminergic neurotransmitters as one component and acetylcholine as the other, and more recently, following the discovery of neural populations in SLD and their mutual inhibition with neurons in the vlPAG, as another flip-flop mechanism. Notice how the hypocretin/orexin (ORX) and the MCH neurons are situated on the "outside" of the switching mechanism, and thus ideally placed to act as state stabilisers (from Saper et al. (2010)).

1.3.3 *Homeostatic, circadian, and allostatic regulation*

A flip-flop switch can thus be used to model the rapid switching between wakefulness, NREM, and REM sleep. However, in order to explain the regular alternations between vigilance stages observed in the hypnogram at longer time scales, the model system has to be enriched with slower variables causing imbalances of the scales that can lead to these transitions.

Borbely proposed a two-process model of sleep regulation (Borbely, 1982), involving a circadian as well as a homeostatic factor. In synergy these two factors yield the observed patterns of wakefulness and sleep (Borbely et al., 2016) (see Figure 1-4). Briefly, the circadian "C" factor codes for time at which sleep is appropriate for the organism, while the homeostatic "S" factor is thought of as a sleep pressure that builds up during waking and dissipates during sleep. The differences between "S" and "C" defines the sleep propensity of the organism, and changes over time incorporating the relative timings of activities in the circadian and homeostatic systems. Ultimately the sleep propensity can be used to model wake-sleep behaviour (Achermann and Tarokh, 2014).

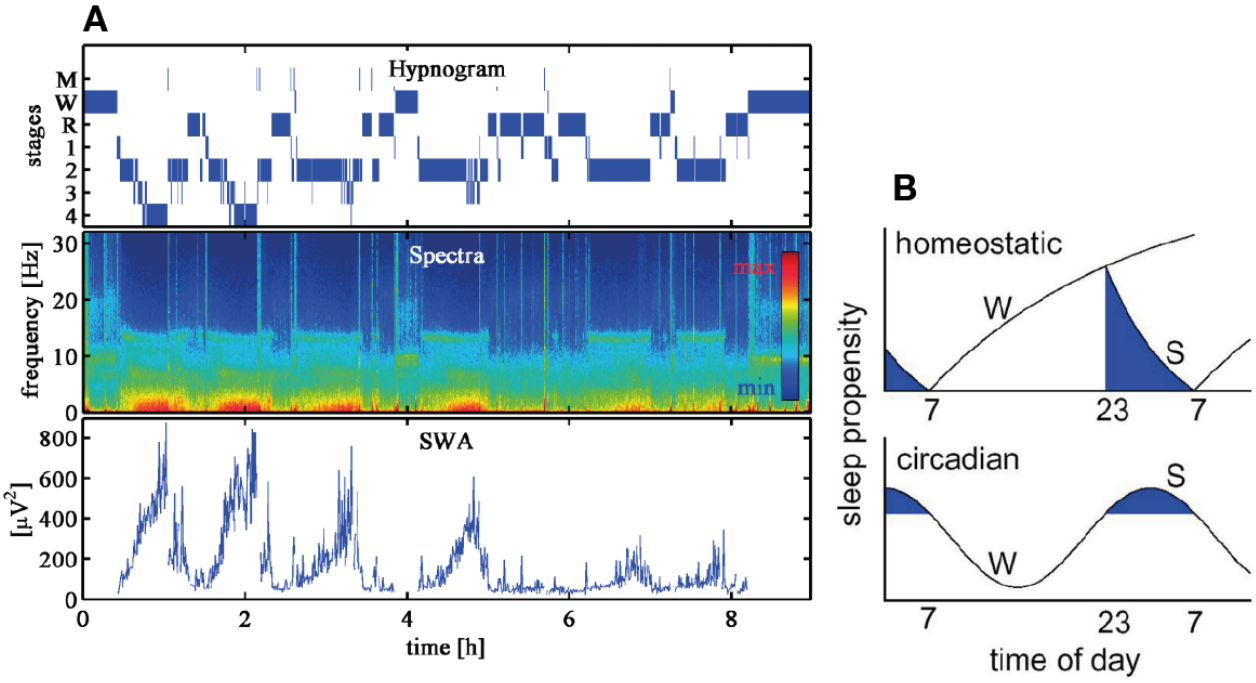

Figure 1-4 ***Two-process model of sleep regulation*** **A)** Illustration of the slow-wave activity (SWA) that captures the decrease in sleep pressure during the course of sleeping. Top row shows a traditional sleep staging or hypnogram. Middle row represents a spectral decomposition of the EEG measured for staging. Notice the bands of alpha activity during wakefulness, spindle activity during sleep stage 2, and slow-wave activity during slow-wave sleep. Bottom row shows how the power in the low-frequency band decreases with succeeding visits of slow-wave sleep. **B)** Schematic representation of the exponential increase of homeostatic sleep pressure during wakefulness (W) and exponential decrease during sleep (S). During sleep deprivation, sleep pressure will continue to rise (Top). The bottom row shows the oscillating circadian factor. The difference between the homeostatic and the circadian factors equals sleep propensity. (from: Achermann and Tarokh (2014))

The best neurobiological candidate known to represent the *circadian* "C" factor is the suprachiasmatic nucleus (SCN) located just above the optic chiasm in the anterior hypothalamus (Schulz and Steimer, 2009). The SCN express activity phase-locked to the environmental light-dark cycle, about which it receives information from photosensitive melanopsin-containing neurons embedded between the visual retinal cells in

the eye. The SCN communicates with the basic wake-sleep regulatory system through the subparaventricular zone (SPZ), dorsal and caudal to the SCN, and the dorsomedial nucleus of the hypothalamus (DMH). The latter projects to both the sleep-promoting VLPO and the wake-promoting orexin/hypocretin neurons in the lateral hypothalamus.

Intriguingly, the *homeostatic* "S" factor can be traced in the EEG as the power of slow wave activity (0.3 – 4 Hz) during NREM sleep. An increased level of slow wave activity can thus be observed during so-called recovery sleep, following sleep deprivation, compared to after normal durations of wakefulness. Adenosine has been proposed as a neurochemical encoder of sleep homeostasis (Brown et al., 2012). Recordings in rats have shown that the extracellular concentration of this metabolite of the energy-carrying ATP molecule increases during waking and decreases during sleep in the basal forebrain as well as in the cortex. How adenosine exerts its somnogenic effects is less well understood. A common hypothesis point to its excitatory A2a receptors that are abundant around the VLPO. Adenosine might also suppress wakefulness through its inhibitory A1 receptors in cholinergic neurons of the basal forebrain and in hypothalamic orexin/hypocretin neurons.

While most work on the homeostatic sleep factor has considered NREM sleep, it is no less true that REM sleep is governed by homeostatic regulation (Franken, 2002). As such, a rebound occurrence of REM sleep can be observed following selective REM sleep deprivation, and transitions into REM sleep become more prevalent.

It has become clear that other factors besides the homeostatic and circadian components are in play when it comes to sleep-wake regulation. Feeding patterns and stress have been described as efficient regulators of wake-sleep behaviour with the ability to override the circadian and homeostatic components (Saper et al., 2010).

1.4 ***Why it is important to investigate whole-brain activity during sleep***

The extensive advance in knowledge regarding the generation of EEG sleep rhythms and their regulation has successfully put to rest previous ideas about sleep as a passive process resulting from a lack of external stimulation. However, the resulting accounts have a tendency to focus on very localised areas of the brain, such as the oscillatory thalamus, and the modulatory brain stem, hypothalamus, and basal forebrain. The whole-brain perspectives of even modern studies within these fields are often limited to recordings from a few EEG electrodes. A long-standing view has developed, according to which sleep results from a collective "shutting-down" or "switching-off" of the cortex by these low-level mechanisms. In this sense, one could say that the "passive" role of the brain in relation to sleep has been moved "upwards" to the cortex.

1.4.1 *How to not sleep*

There are several reasons why a whole-brain perspective should not be neglected in the study of sleep. Beginning at the behavioural level, one thing that is common to human experience is the voluntary component involved in the initiation of sleep. Even when following strict circadian sleep schedules, and noticing the increasing sleep pressure during the course of wakefulness, we are still able to defy sleep at

quite some lengths (Gulevich et al., 1966). This ability is often useful when we are facing extraordinary tasks, or work a night shift. When sufficiently motivated we are thus able to override the circadian and homeostatic components of sleep regulation (Steenland, 2014). It seems unlikely that this kind of agency should be dependent solely on the low-level areas like the neuronal populations in the brain stem, hypothalamus, and basal forebrain. There is however evidence of substantial neural projections to some of these central sleep-wake systems from higher-order areas, such as the prefrontal cortex and limbic areas like the amygdala and the stria terminalis (Saper et al., 2010). A role for such higher-order areas have been hypothesised to serve a so-called "top-down" effect on sleep regulation, and their hyperactivity are thought to play a role in hyper-arousal and insomnia (Riemann et al., 2010). In line with this it is worth mentioning that no single brain area in the central sleep-wake system has led to complete absence of neither sleep nor wakefulness upon lesioning (Saper et al., 2010), suggesting that no single area is essential for sleep (Krueger et al., 2016). For some of these reasons, there is a general call for an incorporation of the (human) cortex in our understanding of sleep manifestation and related disorders (Larson-Prior et al., 2014).

1.4.2 *Local sleep*

A view gaining in popularity is that sleep is not a universal phenomenon in the brain, but rather acts in local and independent ways (Krueger and Tononi, 2011, Kurth and Huber, 2012, Vyazovskiy and Harris, 2013). In this way sleep may be serving its restorative and homeostatic functions locally, in line with the indications that sleep oscillations can occur in isolated patches of cortex and that slow wave activity is particularly intense over brain areas solicited during wake (Murphy et al., 2011). It has been claimed that slow waves represent the "price for plasticity" (Tononi and Cirelli, 2003, 2013), pointing to the role sleep is thought to serve for brain plasticity and maturation. It has even been argued that sleep is simply a fundamental property of neuronal assemblies, which following extensive use will inevitably require periods of quiescence (Krueger et al., 2008).

Such theories plead for a more heterogeneous view of whole-brain activity and a finer spatial detail, than what is offered by polysomnographic recordings.

1.4.3 *Memory and learning*

The enhancement of memory and learning is perhaps *the* phenomenon most consistently linked to sleep function (Stickgold, 2005, Diekelmann and Born, 2010). Several hypotheses exist for the link between sleep-dependent memory consolidation and sleep-related brain activity (Diekelmann and Born, 2010, Tononi and Cirelli, 2013). A curious observation in this respect is that sleep-dependent memory processing has been associated with virtually all aspects of brain activity during sleep, including individual stages, such as N2, SWS and REM, and specific oscillations, such as slow waves, slow oscillations, spindles and ripples (Stickgold, 2013). A more fine-grained view of whole-brain activity during sleep could potentially facilitate more specific insight.

1.4.4 *Features for sleep staging*

The final point above leads to a more general consideration about categorising brain activity during sleep. Polysomnographic criteria have suffered from extensive criticism since its introduction (Himanen and Hasan, 2000, Ogilvie, 2001). The traditional reliance on subjective expert scoring leads to low degrees of inter-scorer variability, while automated algorithms struggle to deal with inter-individual heterogeneity in sleep oscillatory EEG patterns. The traditional stages of sleep are typically based on a rather arbitrary epoch-splitting into 30 seconds windows (see 1.2.1), which imposes an artificial, non-physiological temporal constraint on the recorded brain activity. This kind of "eye-balling" has its obvious caveats, however. Applying different rules or emphasising on different aspects of the EEG might for instance produce a different view on staging. The so-called cyclic-alternating pattern (CAP) is one example of this (Terzano et al., 2001) as is the nine stages of N1 sleep introduced by Hori et al. (1994).

An important characteristic of sleep is that it evolves at many, very different, time-scales; from the slow circadian and homeostatic oscillations, to the short-lived phasic events such as spindles and K-complexes (Olbrich et al., 2011). While more data-driven approaches and computational models are being explored in important attempts to transcend the arbitrariness and subjectivity of polysomnography (Olbrich and Achermann, 2005, Abeysuriya et al., 2015, Abeysuriya and Robinson, 2016), they are commonly based on the same limited number of EEG electrodes, and as such fail to incorporate the multiple *spatial* scales of sleep. The question that arises is whether Loomis and his colleagues would have emphasised the same signatures of sleep had they had access to a modern-day view of brain activity.

1.4.5 *Information processing*

A more theoretically grounded approach directed towards information processing have hypothesised that sleep acts to reinstate a level of criticality in the brain's dynamics (Pearlmutter and Houghton, 2009, Meisel et al., 2013). Criticality should here be understood as a sign of a network system that is situated (in a given parameter space) at a transition point between two states. Importantly the critical state is characterised by optimal information processing in this given network (Beggs, 2008), and an often-highlighted candidate for the parameter space, along which criticality is controlled in brain networks, is the balance between excitation and inhibition (Beggs and Timme, 2012). This account of sleep especially considers the restorative effect that sleep appear to have on performance, such as cognition, relying on information processing in the brain. As such, extended wakefulness is thought to gradually shift the system of brain dynamics away from criticality towards a suboptimal state leading to the deleterious effects of sleep deprivation.

Criticality implies a network view of information processing and brain function, which again speaks in favour of a whole-brain perspective.

1.4.6 *Ideas of consciousness*

Apart from its many hypothesised functions, sleep is undeniably associated with marked changes in consciousness. Along with other states of altered consciousness, such as general anaesthesia hypnosis and coma, sleep has thus played an integral part in the quest to define the neural correlates of consciousness

(Koch et al., 2016). In this respect leading theories are collectively moving away from the concept of a central "switch" in low-level areas of the brain, as well as localised activities in general, and towards a distributed whole-brain perspective (Koch et al., 2016). Quite naturally, ideas of the neural basis of consciousness have depended on the neuroscientific evidence available at their time. As the philosopher John Searle once wrote "... because we do not understand the brain very well, we are constantly tempted to use the latest technology as a model for trying to understand it. In my childhood we were always assured that the brain was a telephone switchboard... Sherrington, the great British neuroscientist, thought the brain worked like a telegraph system. Freud often compared the brain to hydraulic and electromagnetic systems. Leibniz compared it to a mill... " (Searle, 1984).

The uniform EEG representation of the sleeping brain that is under a "bottom-up" control by the low-level brain areas and the thalamus represents a strong model with detailed evidence and has undoubtedly exerted its effects. Nevertheless, two of the currently most prevailing theories, such as The Global Neuronal Workspace Theory (Baars, 1993, Dehaene et al., 2003, Dehaene and Changeux, 2011) as well as the Integrated Information Theory (Tononi, 2008, Tononi et al., 2016) make predictions about large-scale distributed networks of information processing.

The Global Neuronal Workspace Theory postulates that conscious experiences require a "broadcasting" of neural information to an anatomically widespread network of cortical regions (Dehaene et al., 2003, Dehaene and Changeux, 2011). Consciousness then results from a selective incorporation of information from different modules of the brain into the workspace. In particular, the Global Neuronal Workspace Theory predicts that external stimuli of sensory organs will make it into consciousness, only if this stimuli is processed by higher-order and not only low-level sensory areas. Fronto-parietal communication has been suggested to play an important role for this global distribution of information and presentation to the conscious mind (Dehaene et al., 2003, Dehaene and Changeux, 2011).

The Integrated Information Theory is less specific in the anatomic sense. Instead it is based on information theory and often claimed to be born out of phenomenological considerations (Tononi, 2008, Tononi et al., 2016). The phenomenology of a conscious experience is described as being unique and differentiated while at the same time indivisible and integrated. The content of consciousness can hold a lot of information but will always be experienced as an integrated whole. According to the Integrated Information Theory consciousness arise in systems, which can hold high amounts of information, represented by many possible states, and which importantly is capable of integrating these specialised pieces of information through causal interactions. The brain is thought to encompass such a system with an ability to turn information into an integrated whole that equals more than the sum of its differentiated sub-parts.

Both of the above theories speak in favour of a whole-brain perspective of the brain when studying the consciousness-related characteristics of sleep (Tagliazucchi et al., 2013a).

1.5 ***Directions and challenges in the study of whole-brain activity of sleep***

1.5.1 *Intrinsic dynamics of whole-brain activity and networks*

Historically sleep was regarded a passive state, where the absence of external stimuli was thought of as the main sleep promoting mechanism, which would allow the brain to enter a state facilitating behavioural quiescence. As we have seen, sleep is now known to be an active process in the brain. Interestingly, a somewhat analogous development has occurred in our understanding of spontaneous brain activity, as measured with modern neuroimaging. The initial phase of neuroimaging focussed on task and stimulus related activity, largely regarding the brain as a passive processor of incoming information and its spontaneous activity as noise (Raichle, 2010). This view has been challenged, however, by the discovery that spontaneous brain activity is in fact non-random and non-noisy in its macro-scale organisation (Biswal et al., 1995), and that it influences behaviour (Valera et al., 1981, Arieli et al., 1996, Klimesch et al., 2007, Mazaheri and Jensen, 2010). This has led to a recent shift in focus towards understanding and characterising this intrinsic activity of the brain.

The development has been spearheaded by studies using functional magnetic resonance imaging (fMRI). It has been shown how the brain's spontaneous activity, as recorded in the absence of task or stimuli, during the so-called resting-state, is organised into large-scale patterns of co-variations, known as functional connectivity (FC) (Raichle and Mintun, 2006, Fox and Raichle, 2007). The concept of *resting-state networks* (RSNs) has emerged. Derived from the slowly fluctuating BOLD signal, these distributed networks have been shown to be consistent within and across individuals (Damoiseaux et al., 2006, Laumann et al., 2015), and to accurately overlap with task-related activation maps (Smith et al., 2009). The default mode network (DMN), originally identified with PET (Shulman et al., 1997, Raichle et al., 2001), has been awarded particular attention for its specific anti-correlation with externally oriented behaviour (task and stimuli), and been proposed as a correlate of a range of cognitive phenomena, such as self-directed awareness, autobiographical memory, and mind wandering (Buckner et al., 2008, Spreng et al., 2009, Andrews-Hanna et al., 2010). Cognitive scientists have welcomed RSNs, incorporating them into their models for explaining cognition from a network perspective of brain function (Bressler and Menon, 2010).

The network approach also resonated well with a general opposition to a longstanding tradition in brain sciences, referred to as *localizationism* (Friston, 2011, Rabinovich et al., 2012). Refuting a view according to which specific functions of the brain reside in specific and spatially confined areas or regions, the network view regards brain function as a result of the collective dynamics of whole-brain networks. Local functional specialisation is not disregarded, but rather embraced as segregated information, which has to be appropriately integrated into the continuously evolving dynamics for the brain to support its immense amount of functions (Sporns et al., 2000, Deco et al., 2011, Tognoli and Kelso, 2014). This level of description cannot be addressed by traditional parametric mapping of localised brain activities associated with specific tasks or stimuli.

The introduction of *whole-brain FC* networks may be seen as an expansion of the RSN framework, aiming to include all areas of the brain and not just the ones belonging to RSNs. This has allowed for the application and testing of concepts from network and dynamical systems theories. *Graph theory*, a

mathematically rooted description of networks, has provided tools to easily quantify levels of segregation and integration in whole-brain FC networks. Concepts like small-world topology, modular and hierarchical structures, as well as scale-free distributions have been attributed to the network organisation of spontaneous brain activity (Bullmore and Sporns, 2009).

The parallel development in other magnetic resonance imaging (MRI) techniques, such as diffusion weighted imaging, together with tractography analyses have allowed for the mapping of white-matter fibre tracts connecting grey-matter areas throughout the brain, giving rise to the concept of whole-brain *structural connectivity* (SC) and connectomics (Hagmann, 2005, Sporns et al., 2005). An intimate relationship between this structural connectome and FC patterns has been established, suggesting that the observed whole-brain FC networks are unfolding on a scaffold of underlying structural connections.

Whole-brain computational modelling has been used to explore this close SC-FC relationship further. By endowing the structural connectome with various dynamical regimes whole-brain FC networks can be simulated, and this has made it possible to test the dynamic principles underlying the generation of the spontaneous activity (Cabral et al., 2014a, Deco and Kringelbach, 2014).

The majority of discoveries within whole-brain FC have relied on the study of the low-frequency components of the BOLD signal as measured with fMRI. Meanwhile, the indirect hemodynamic relationship between BOLD and neural activity can make functional interpretations difficult. Electrophysiological studies have proven important for the link of the large-scale fMRI findings to direct measures of brain activity (Engel et al., 2013, Schölvinck et al., 2013). Through source-localisation, MEG, in particular, has been used to shed light on the *electrophysiological underpinnings* of the brain's spontaneous activity (Brookes et al., 2011c, Hipp et al., 2012, Siegel et al., 2012, Baker et al., 2014a).

Most recently, investigators have started to appreciate the non-stationary nature of the brain's spontaneous activity, and moved away from analysis strategies collapsing over long-interval recordings (Hutchison et al., 2013a, Calhoun et al., 2014). *Dynamical functional connectivity* explores temporal interactions between RSNs, the temporal evolution of whole-brain FC networks, and their relationships to SC, constantly providing new evidence for whole-brain computational models to explain (Breakspear, 2017, Cabral et al., 2017).

As it stands, the network approach to both structural and functional neuroimaging data, as well as their combination through whole-brain computational modelling, represents a promising and very broad framework for exploring brain function at the macroscale. As such, resting-state networks, functional and structural connectomes, as well as their dynamics have been implicated in a wide array of healthy and disordered behaviour (Buckner et al., 2008, Bassett and Bullmore, 2009, Zhang and Raichle, 2010, Deco and Kringelbach, 2014). The prosperity of this framework is reflected in numerous heavily-funded collaboratory initiatives, such as the human connectome project (HCP), and the UK Biobank (see e.g. Craddock et al. (2013) for an overview)

A majority of research within the whole-brain network perspective have focussed on spontaneous brain activity during wakefulness. This has resulted in tools to extract new features from brain recordings – a kind of new language to decode the brain without the use of sensory stimuli and tasks. Given that sleep represents

one of the most noticeable phenomena occurring spontaneously in the brain, it is sensible to explore these features not only during wake, but also during sleep.

In the following sections I will review the progress made so far within whole-brain studies of spontaneous activity during sleep that go beyond traditional EEG representations. While a few overview articles on the topic already exist and Spoormaker et al. (2011), see e.g. Picchioni et al. (2013), they are not quite up-to-date. While covering new knowledge published since then, I will organise the findings into the themes outlined in the beginning of this section and the methods employed by them. Given the severe underrepresentation of sleep studies in the field of whole-brain connectivity, I will occasionally draw on findings from sedation and general anaesthesia, while acknowledging that these behavioural states compare to sleep only in certain aspects (Brown et al., 2010).

1.5.2 *"Traditional" neuroimaging of sleep*

This section reviews advances made within traditional neuroimaging of sleep. The following findings are coined "traditional" mainly because they fall outside of my understanding of a network approach to brain activity. As such, most of these studies have used sleep stages, specific sleep oscillations, or external stimuli during sleep as event-related paradigms, regressing these onto the recorded activity timecourses. Studies falling within this category have been comprehensively reviewed elsewhere (Maquet, 2000, Dang-Vu et al., 2010b, Meyer et al., 2012, Nir et al., 2013). Here I will focus on the findings that I consider relevant to a network perspective of brain function. Some of these are presented in Figure 1-5.

Early studies, using PET, have shown that global levels of metabolism and blood flow are generally decreased during NREM sleep compared to wakefulness and REM sleep. Decreases are not homogenous across the brain, however. During NREM, blood flow is decreased in areas related to sleep regulation, such as the brainstem, basal forebrain, and thalamus, while a heterogeneous pattern of decreases can be found in the cortical areas. As such, blood flow was found preserved in primary sensory areas, while higher-order areas showed significant decreases. The heterogeneous picture was taken to be indicative of preserved information processing in the brain during sleep, as well as a sign of regional homeostasis in some of the most solicited areas during wakefulness (Dang-Vu, 2012, Meyer et al., 2012).

During REM sleep, blood flow increases compared to NREM, in the brainstem, thalamus, and extra-striate occipital cortex. This would appear in line with the observed desynchronised EEG, and the hypothesised mechanisms underlying ponto-geniculate-occipital (PGO) waves observed in animals (Datta, 2012). Increases have moreover been found in the hippocampus, amygdala, anterior cingulate, orbitofrontal cortex, and insula, and suggested to reflect emotional and memory processes and perhaps even dreams.

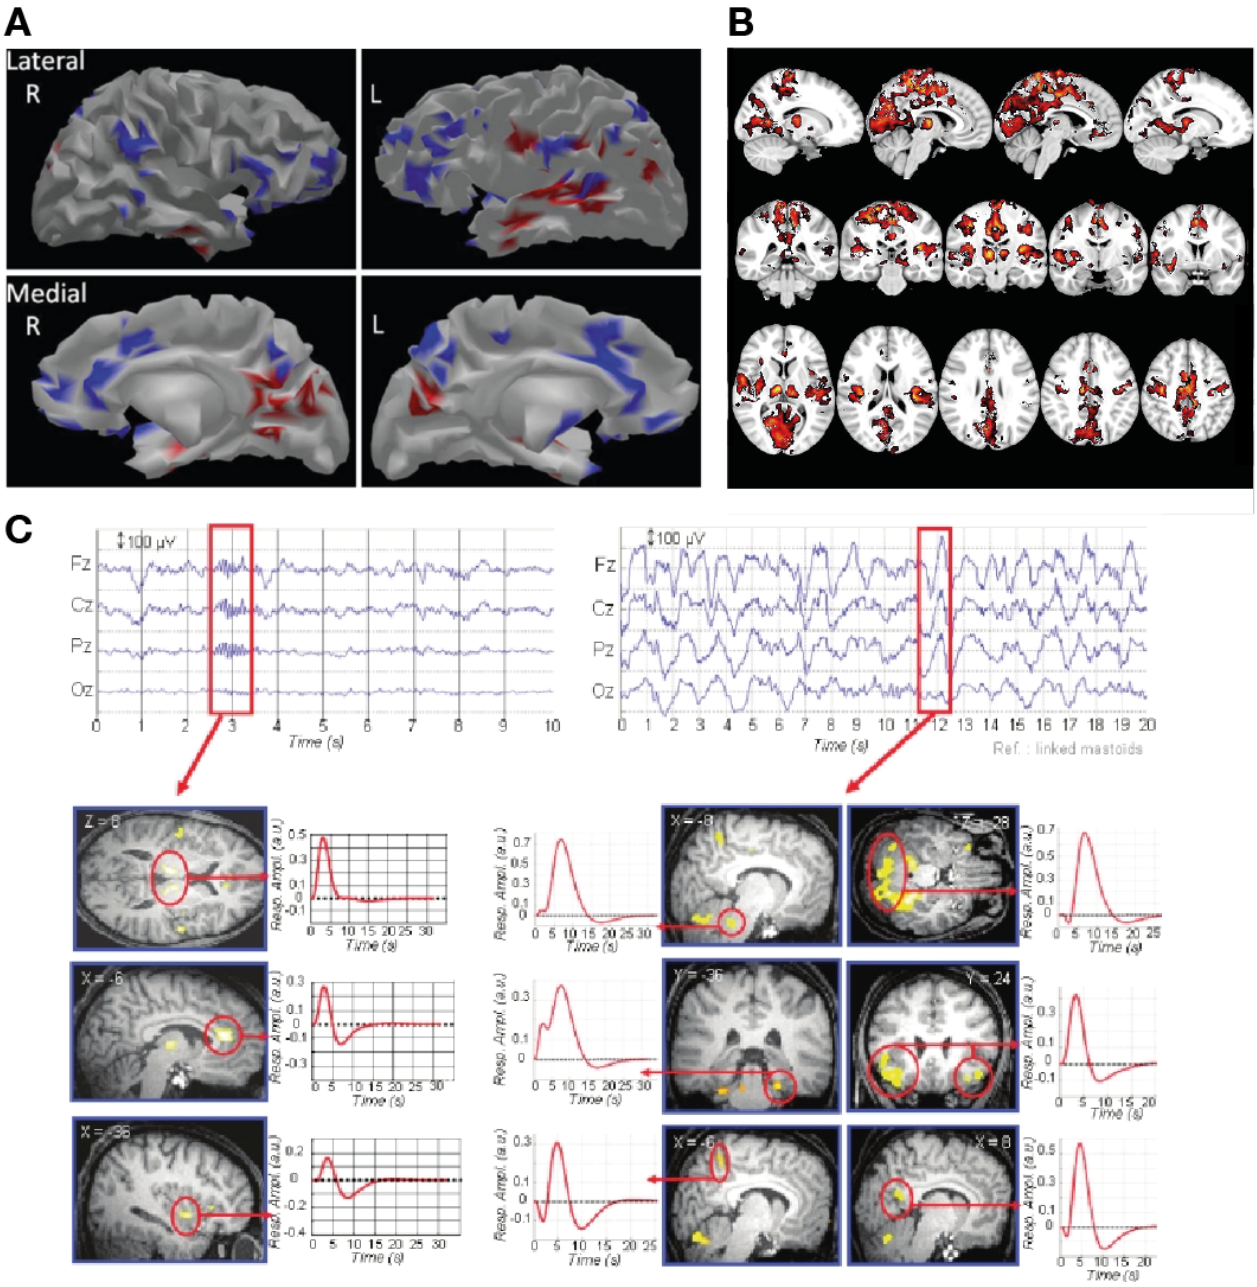

Figure 1-5 ***Examples of key findings from the "traditional" framework of sleep neuroimaging.*** **A)** A meta-analysis from Nir et al. (2013) of PET studies of regional increases (red) and decreases (blue) between wakefulness and NREMs sleep. **B)** Parametric mapping of fMRI activity time-locked to spontaneous K-complexes (adapted from Caporro et al. (2012)). **C)** Parametric mapping of spindles (left) and individual slow oscillations (right) with fMRI. (Adapted from Dang-Vu et al. (2010a); left originally from (Schabus et al., 2007); right originally from Dang-Vu et al. (2008b))

Taken together, these findings demonstrated that while brain activity is generally decreased during NREM sleep, it is not absent in either NREM nor REM. Sleep-resilient information processing has been examined through auditory and visual stimulation and the accompanying evoked responses measured with fMRI, taking advantage of its superior temporal resolution compared to PET. Initial studies showed that evoked responses were preserved in primary sensory cortices during NREM sleep (Portas et al., 2000), while contrary evidence pointed to decreased responses (Born et al., 2002, Czisch et al., 2002). This should be seen in light of findings from evoked response studies with scalp EEG, where certain aspects of the evoked

response have been shown to decrease with sleep, while others are in fact enlarged (Colrain and Campbell, 2007). K-complexes during N2 sleep can for instance be generated by sound stimulation (Colrain, 2005).

The non-evoked, spontaneous brain activity during sleep was also investigated using fMRI, through time-locking and averaging of BOLD responses to sleep oscillations. The hemodynamic activity changes related to specific EEG events were identified, such as slow oscillations (Dang-Vu et al., 2008a), sleep spindles (Laufs et al., 2007, Schabus et al., 2007, Tyvaert et al., 2008, Andrade et al., 2011, Caporro et al., 2012), K-complexes (Laufs et al., 2007, Caporro et al., 2012, Jahnke et al., 2012), and vertex sharp waves (Stern et al., 2011). In REM sleep, signals were averaged around the rapid eye movements (Wehrle et al., 2007).

Linking external stimuli and spontaneous EEG oscillations it was shown how specific oscillations can affect the regional hemodynamic patterns produced by sounds (Dang-Vu et al., 2011, Schabus et al., 2012). This led to the proposal that sleep spindles might represent a bottom-up "gate" mechanism for incoming stimuli exerted by the thalamus. This hypothesis was recently challenged by an intracranial study in rats (Sela et al., 2016). Evidence from evoked response studies in sedation and general anaesthesia generally showed preservation of activity in primary sensory regions and relative deactivations in higher-order frontal and parietal areas compared to wake. The absence of evoked activity in these higher-order areas has been hypothesised to underlie observed decreases in the late part of evoked potentials measured with EEG, and interpreted in light of the Global Neuronal Workspace Theory (Mashour, 2014, Uhrig et al., 2014).

Overall, findings within traditional paradigms of neuroimaging have taught us about global, and regional levels of activity in the contrast between wakefulness and traditional sleep stages and oscillations. Notably, they have increased the spatial resolution of the brain during sleep, and allowed for non-invasive localisation of sleep-specific oscillations in humans. The resulting non-uniform spatial patterns have confronted the view of whole-brain activity during sleep as reducible to a few EEG electrodes.

1.5.3 *Resting-state networks and sleep*

Compared to the sleep-related activations and deactivations, described above (1.5.2), which are mainly interpretable within a localizationist view of brain function, resting-state networks (RSNs) can be used to investigate changes in inter-areal communications from a network perspective. RSNs have become integral to the characterisation of whole-brain network activity during wakefulness (Biswal et al., 2010, Raichle, 2010, Snyder, 2016), and in the following I will review how they change during sleep.

RSNs are defined as highly structured patterns of co-variations in the slow component (< 0.1 Hz) of the raw BOLD signal. An RSN comprise areas distributed across the brain that are connected through functional connectivity (FC), where FC may be defined broadly as statistical interdependencies of activity timecourses (Friston, 1994, Friston, 2011). It is common to part the known RSNs into primary sensory networks (e.g. visual, auditory, and somatomotor), and higher-order networks (e.g. dorsal attention, fronto-parietal, central executive, language, and default mode). Examples of these are given in Figure 1-6.

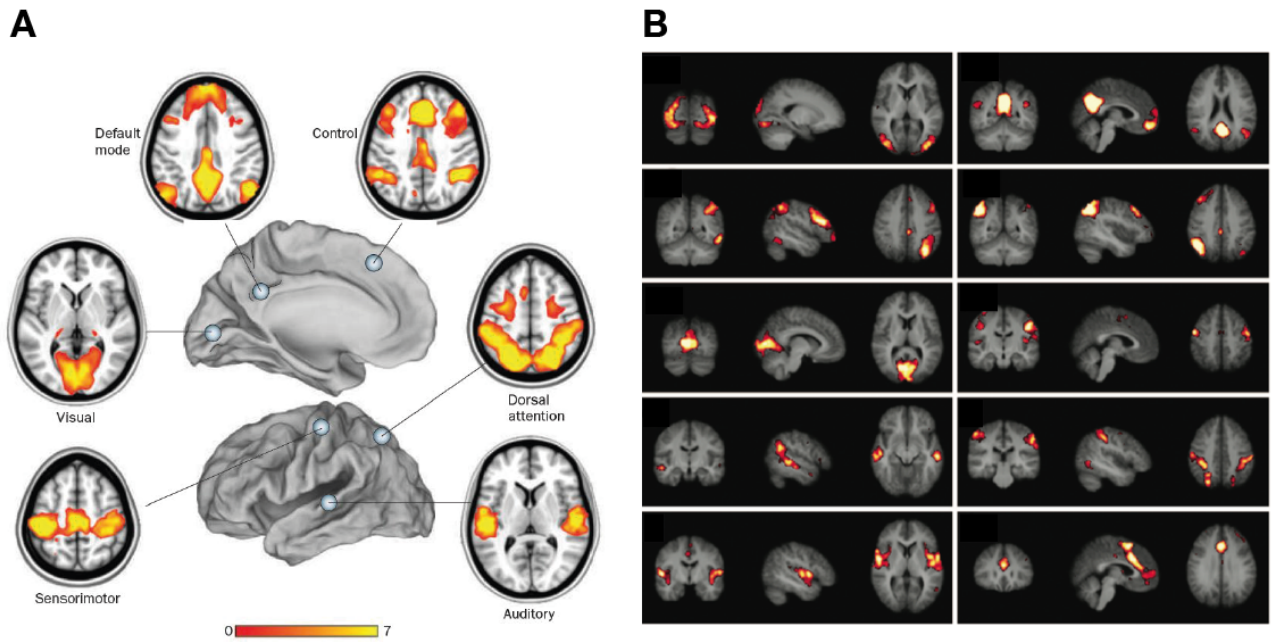

Figure 1-6 ***Overview of resting-state networks***. **A**) Anatomical representations of the main RSNs, encompassing sensory (visual, sensorimotor, and auditory) and higher-order (default mode, control, and dorsal attention) brain areas (from Zhang and Raichle (2010)) **B**) A number of these RSNs are found consistently across participants as shown in this study by Damoiseaux et al. (2006).

The two main methods for identifying RSNs in fMRI are seed-based regression analysis and spatial ICA. These are illustrated in Figure 1-7. In seed-based regression, covariance or linear correlation is typically used to calculate the dependency of timecourses from a given "seed" voxel in the brain to the remaining "target" voxels (Biswal et al., 1995, Fox et al., 2005, Fox and Raichle, 2007). If more than one voxel is used as seed it is commonly referred to as a region-of-interest (ROI), from which a representative timecourse can be derived. The resulting networks correspond to the target voxels that show particularly strong FC values with the seed, and are typically visualised as heat map on a brain volume. A commonly used pre-processing step in seed-based approaches is that of global signal removal, which entails regressing out the common, averaged signal across the entire brain volume from all individual voxel timecourses. This is necessary, for instance, for finding the strong negative correlation between the DMN and attention networks (Fox et al., 2005). The physiological soundness of this step is the topic of an on-going debate (Fox et al., 2009, Murphy et al., 2009a, Smith et al., 2011b, Murphy and Fox, 2016), which has particular relevance for the study of sleep, since, as we shall see, the global BOLD signal has been found sensitive to changes in vigilance (Fukunaga et al., 2006, Horovitz et al., 2008, Wong et al., 2013). This means that potentially important and non-artefactual information may be lost when the global signal is regressed out across all stages (Scholvinck et al., 2010).

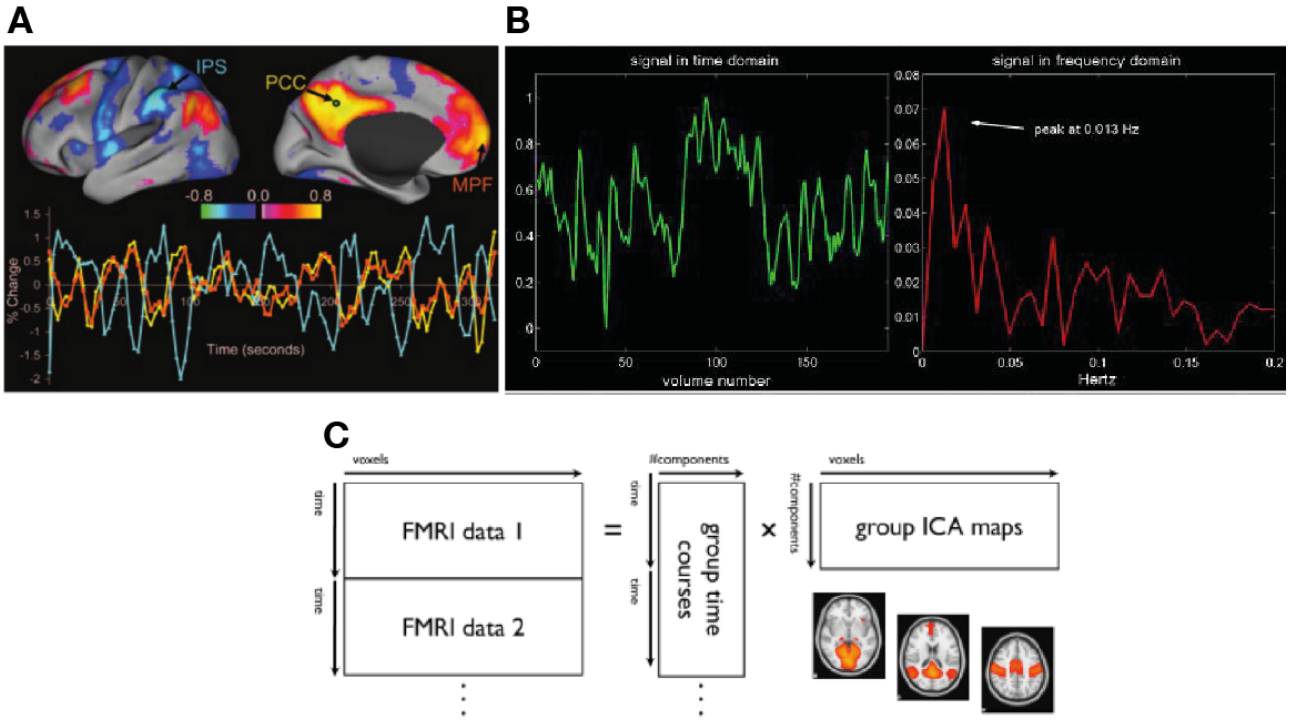

Figure 1-7 ***Methods for deriving resting-state networks in fMRI***. **A)** RSNs are connected through functional connectivity. Seed-based regression takes the timecourse of a given voxel in the brain (the seed), and calculates the covariance, or linear correlation, between this and the timecourses of the remaining voxels (the targets). Each voxel of the brain will thus be associated with a score corresponding to how similar its activity timecourse is to that of the seed. RSNs are then commonly presented as heat-maps on the brain volume, showing the voxels above a certain covariance or correlation threshold. By placing the seed in ROIs known to be involved in specific RSNs, it is possible to investigate the strength and integrity of the networks of interest. Note the anti-correlation (blue vs. red colors) between the DMN and the dorsal attention network (adapted from Fox et al. (2005)). **B)** Spectral composition of the raw BOLD signal, pointing to the slow component that is being investigated in RSNs and functional connectivity in general (adapted from (Mantini et al., 2007)). **C**) ICA is a Blind Source Separation algorithm, which decomposes a multivariate signal into a number of statistically independent components. In fMRI it is common to constrain the decomposition in the spatial domain. Compared to seed-based approaches, ICA is data-driven, although the number of components, the model order, has to be specified (adapted from Cole et al. (2010)).

An obvious limitation of the seed-based RSN analyses is the subjectivity associated with placement of the seed. Independent component analysis (ICA) overcomes this issue by deriving spatial maps from the full set of voxels in a data-driven manner (Kiviniemi et al., 2003, Beckmann et al., 2005, Damoiseaux et al., 2006, Cole et al., 2010). ICA is an algorithm that solves the Blind Source Separation problem, and produces an un-mixing of signals into a set of statistically independent components. Each of the independent components will thus be represented by a timecourse and a spatial map, where the latter shows the RSN pattern when displayed on the brain. A number of different implementations of ICA may be used. Of note is the choice of the dimension, in which the independence between components is maximised, i.e. space or time. In fMRI it is common to apply the spatial constraint, such that the resulting ICA maps become spatially independent. Compared to temporal independence, the spatial implementation does not enforce complete orthogonality between the timecourses, and arguably results in biologically more plausible components (Beckmann et al., 2005).

Importantly, the data-driven nature of the ICA approach comes at the cost of leaving the dimensionality, also known as the model order, as a free parameter. One has to define *a priori* how many components the

ICA algorithm should derive from the data. While certain methods exist, by which a 'suitable' number of components may be determined based on the data (Cole et al., 2010), the choice of model order remains an inherent challenge in the modelling of spontaneous brain activity. The long-range distributed organisation of 'canonical' RSNs are typically produced around a model order of 20. Increasing this number yields spatially more confined maps (Smith et al., 2011b).

The first studies to approach sleep in the framework of RSN activity established that the slow (< 0.1 Hz) fluctuations in the raw BOLD signal, from which RSNs are derived, could indeed be measured during sleep. Fukunaga and colleagues thus showed that the power of said fluctuations was in fact larger in participants who fell asleep during a resting-state recording compared to those who remained awake (Fukunaga et al., 2006), although it should be noted that sleep was defined rather loosely as participants who were non-responsive at the end of the scan. This was found at a whole-brain level as well as within the visual cortex (Fukunaga et al., 2006). The same group of researchers were later able to confirm this higher level of slow fluctuations in many cortical areas, during N1 and N2 sleep (Horovitz et al., 2008), and through seed-based regression demonstrated a preserved DMN during sleep (Horovitz et al., 2008), in line with ICA findings from a concurrent study by Boly et al. (2008).

Larson-Prior et al. used seed-based regression to investigate the development of most of the canonical RSNs (primary sensory as well as higher-order networks) from wakefulness to N2 sleep (Larson-Prior et al., 2009). FC values between the included ROIs were found largely unchanged in the descent to N2 sleep. These observations, together with their equivalents made in studies of alternative states of decreased vigilance, such as sedation and general anaesthesia in humans (Kiviniemi et al., 2005, Greicius et al., 2008, Martuzzi et al., 2010) as well as non-human primates (Vincent et al., 2007), caused confusion in the field. In particular the higher-order RSNs, which had been hypothesised to support complex behavioural and cognitive phenomena, appeared to be more or less insensitive to changes in vigilance and consciousness. Illustrations of the preservation of RSNs throughout sleep are given in Figure 1-8.

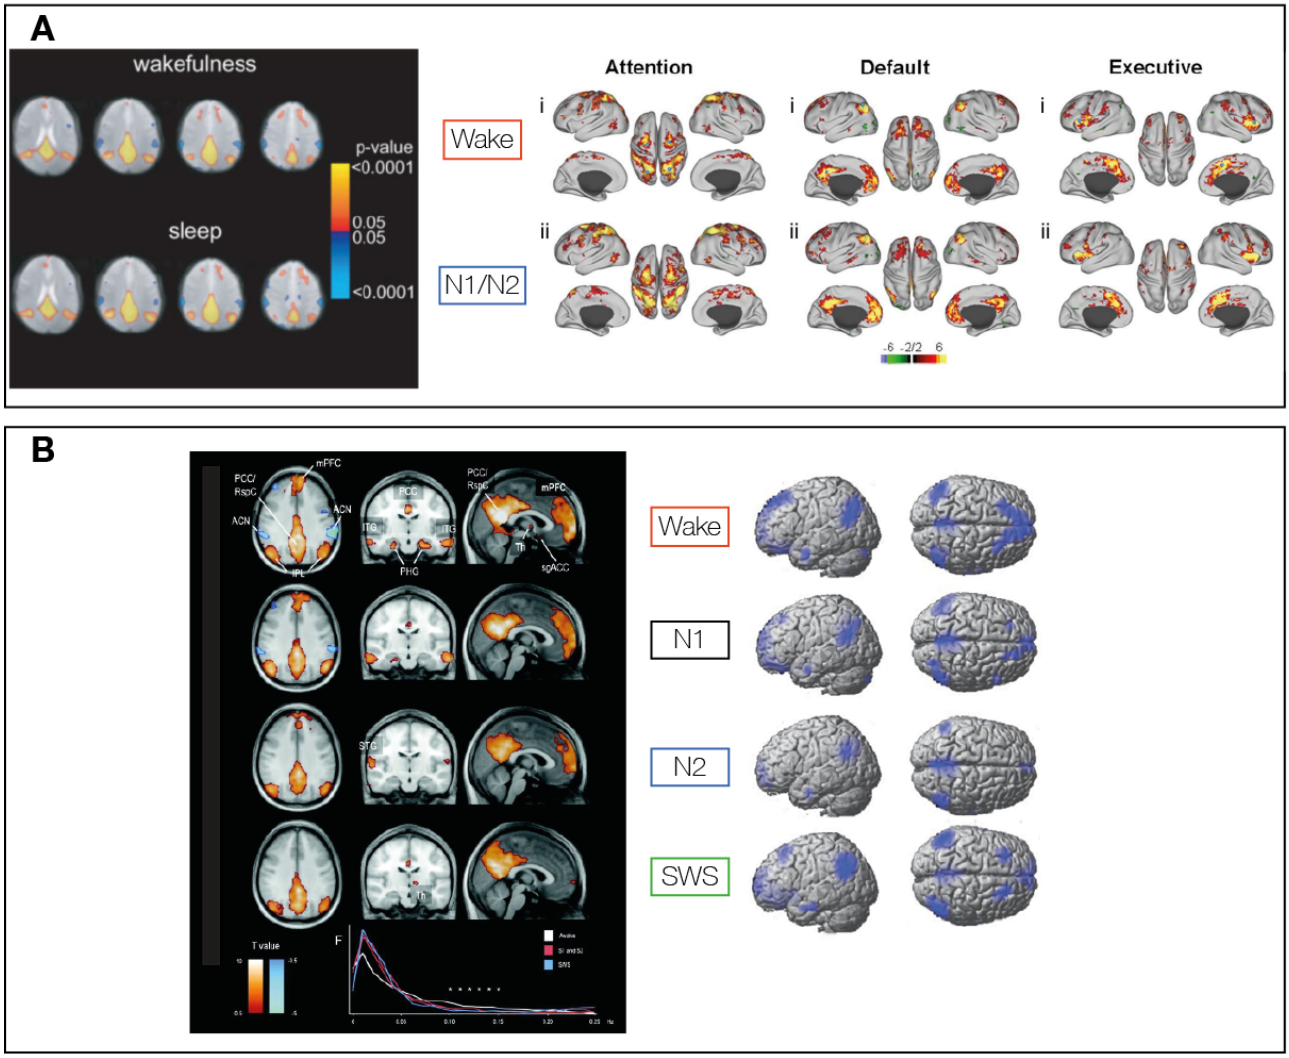

Figure 1-8 ***Preservation of resting-state networks across NREM sleep stages.*** **A**) The DMN as well as other higher-order networks remains intact in N1 and N2 sleep. This finding was in conflict with initial hypotheses assigning these networks important roles in wakeful cognition (left: from Horovitz et al. (2008); right: from Larson-Prior et al. (2009)). **B**) Equivalent observations were found for slow-wave sleep (SWS). To the left: Seed-based regression (together with global signal regression) suggested weakening connections within the DMN and a fading of anti-correlations (blue areas) as sleep deepens (from Sämann et al. (2011)). However, the DMN was intact nonetheless in all stages of NREM sleep, as seen from the ICA results on the right (from (Tagliazucchi et al., 2013c)).

Some reconciliation was offered as researchers began to consider not only early stages of NREM sleep, but also SWS. Horovitz and colleagues demonstrated a disintegration of the DMN when contrasting wakefulness and SWS (Horovitz et al., 2009). Specifically, they used seed-based regression to show how the frontal parts of the DMN decreased their FC values to the posterior parts. Sämann and colleagues used ICA to identify the DMN and a fronto-parietal attention network (Sämann et al., 2011). With seed-based regression they went on to track the strength of FC within, as well as the level of anti-correlation between, the ICA-derived networks, from wakefulness and through N1, N2, and SWS. While further establishing that both of these higher-order RSNs could be detected in all stages of sleep, they showed that all stages of NREM sleep were associated with a decoupling of the medial temporal parts of the DMN, thus confirming the finding of Horovitz and colleagues (2009).

In terms of the inter-network relationship between the DMN and the attention network, they demonstrated a disappearance of their anti-correlation in NREM sleep (Sämann et al., 2011). Chow and colleagues later confirmed this for SWS (Chow et al., 2013), as did Boveroux and colleagues for propofol-induced sedation (Boveroux et al., 2010).

Up to this point the majority of studies had focussed on NREM sleep and its contrast to wakefulness. The general lack of reports on REM sleep could be attributed to the fact that the employed sleep protocols rarely extended over more than a few hours. As we have seen, the two-process model of sleep regulation predicts that NREM sleep predominates during the early course of a sleep bout (see 1.3.3). The sleep restriction, commonly used as a facilitator of sleep in the mentioned studies, could also be thought to have increased the proportion of NREM sleep during the subsequent recordings. Koike and colleagues did not employ any sleep restriction, and were the first to report on RSNs in REM sleep (Koike et al., 2011). Once again the focus was directed towards the DMN, which was found preserved. While wakefulness was not directly included in the connectivity comparisons, the reported results did point to a vigilance effect on DMN integrity, with the highest level of intra-network connectivity found in N2 sleep, and the lowest in SWS. REM sleep was also included in the study by Chow and colleagues, who reported a highly intra-connected DMN during REM with FC levels similar to, or above, those observed for wakefulness (Chow et al., 2013).

To summarise, the study of RSNs during sleep has first and foremost been characterised by puzzlement over the surprising level of perseverance of RSNs despite the substantial loss of responsiveness, cognition, consciousness, and self-awareness commonly associated with sleep. The fact that e.g. the DMN can be found even during SWS still stands as an often-used argument against RSNs role in higher-order cognitive phenomena (Snyder, 2016). In this way, one might argue that RSN studies of sleep have had a higher impact on the understanding of brain activity during wakefulness, than it has taught us about sleep.

Some consistent, potentially sensitive markers of sleep were identified, however. FC *strengths* within as well as between RSNs appeared to change with sleep. This was supported by observations from RSN-studies of sedation and general anaesthesia (for reviews, see Bonhomme et al. (2012), Heine et al. (2012), and MacDonald et al. (2015)). While intriguing, the fading anti-correlations between the DMN and attention networks should be interpreted with caution, since all studies reporting this phenomenon relied on global signal removal. A third, often emphasised, observation was the decreased connectivity between frontal and parietal brain areas in sleep, particularly within the DMN, which has been interpreted as a sign of decreased feed-back or top-down communication in line with the Global Neuronal Workspace Theory of consciousness (Boly et al., 2012a, Tagliazucchi et al., 2013a, Mashour, 2014).

1.5.4 *Whole-brain functional connectivity networks, graph theory, and sleep*

Whole-brain FC networks are built by estimating the pairwise FC values between all areas of the brain, and not just the ones belonging to RSNs. To do this one has to first define these areas as a number of ROIs in what is commonly known as a parcellation.

Defining sub-parts of the brain in an informed way is no trivial task, and may be based on structural, anatomical, or indeed functional information (Guillery, 2000, Jbabdi and Johansen-Berg, 2011, Craddock et

al., 2013, Sporns, 2015, Glasser et al., 2016). In fact it could be argued that RSNs themselves, especially when estimated using the data-driven ICA, collectively represent a functional parcellation, with each network defining a parcel. It is more common, however to define parcels as spatially contiguous ROIs. Increasing the model order of an ICA decomposition splits up RSNs into smaller portions, and is one way of achieving a functional parcellation (Smith et al., 2011b). Another common approach is to parcellate the brain according to anatomically defined areas, such as those given by the Automated Anatomical Labeling (AAL) atlas, which consists of 116 cortical, sub-cortical, and cerebellar areas (Tzourio-Mazoyer et al., 2002).

Within a given parcellation, a whole-brain FC network can be represented as a square matrix including all pairwise FC values. This matrix can then be analysed in a variety of ways and used, for instance, to compare different experimental conditions or participant populations. Contrasting FC values, or strengths, over the full network, in such a comparative network analysis, comes with the challenge of dealing with multiple comparisons, as the number of connections increases quadratically with the number of ROIs. The inflated false positive rate can be controlled in a number of ways, including the conservative Bonferroni method, or non-parametric permutations within a general linear model (GLM) framework (Zalesky et al., 2010, Winkler et al., 2014).

Graph theory is a branch of mathematics used ubiquitously in the study of networks (Bullmore and Sporns, 2009), and can also be used to analyse whole-brain FC, as illustrated in Figure 1-9. Within its notation, a network is defined as a graph, consisting of nodes connected by edges. Graph-theoretical analyses operate with concepts central to information processing (Rubinov and Sporns, 2010, Fornito et al., 2013). Once a network has been properly defined (see Figure 1-9A–B), with nodes corresponding to ROIs and edges to functional connections, it is possible to derive topological characteristics reflective of integration and segregation properties (see Figure 1-9C–D).

At a global level, the overall topology of a graph may be characterised on a scale from order to randomness (see Figure 1-9C). In between these two extremes lies a range of topologies, balancing integration and segregation, resulting in optimal information processing. This is often referred to as complex topology or small-worldness (Watts and Strogatz, 1998, Sporns, 2011, Bullmore and Sporns, 2012).

Graph theory is a convenient and readily applicable framework for the study of whole-brain networks, as it quantifies concepts central to a network view of brain function (Tononi et al., 1998). It allows for network-related hypotheses about whole-brain neuroimaging data, and has been applied extensively to the study of spontaneous, resting-state activity in health and disease (Bassett and Bullmore, 2009, Fornito et al., 2015). As such, whole-brain networks have been found to express small-world topology, power-law-like degree distributions, and modular community structures (Bassett and Bullmore, 2006, Bullmore and Sporns, 2012). These topological characteristics have been interpreted as signs of criticality and metastability, which are thought of as fundamental properties of collective brain dynamics (Chialvo, 2010, Deco et al., 2011, Beggs and Timme, 2012, Tognoli and Kelso, 2014).

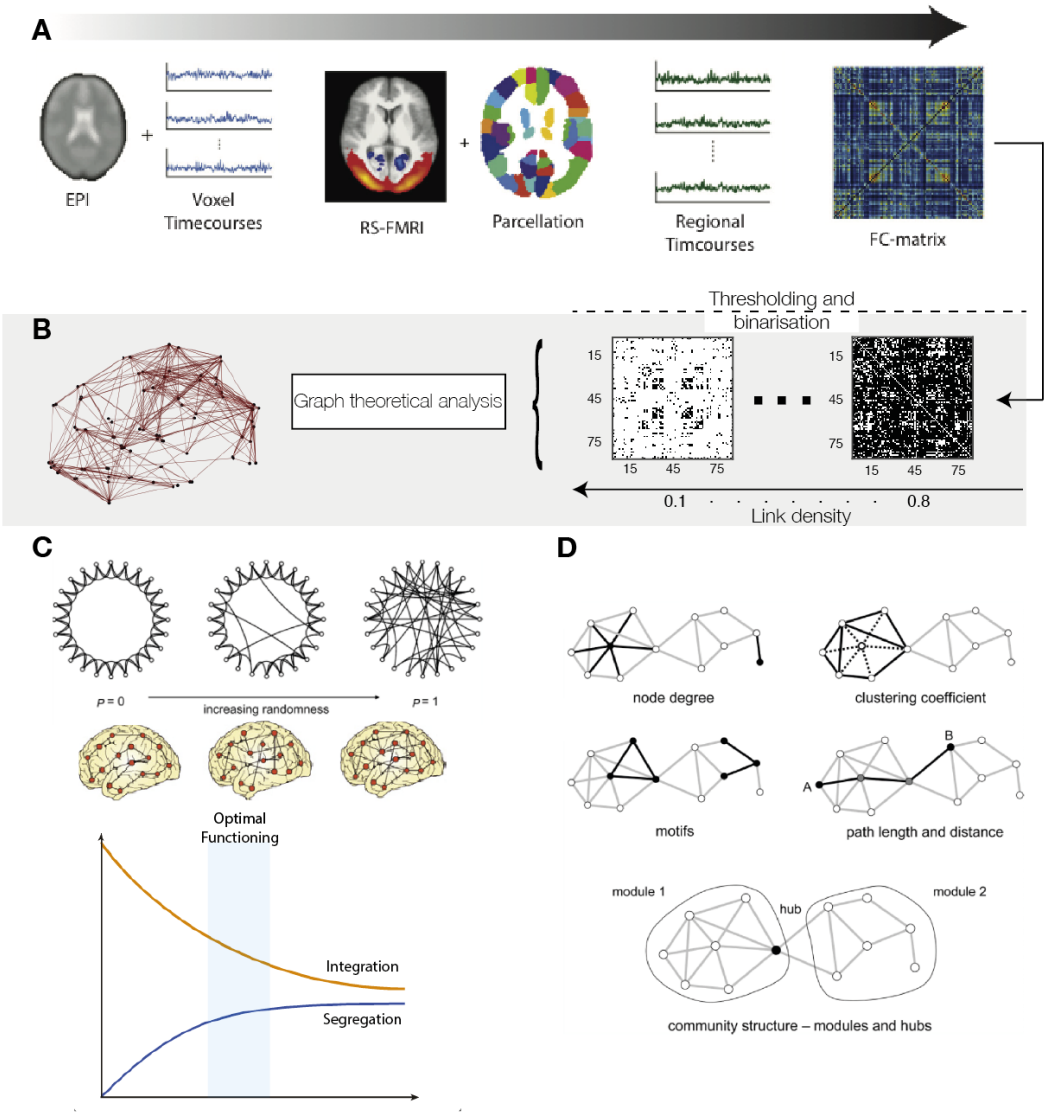

Figure 1-9 ***Functional connectivity and graph theory.*** **A**) *Construction of whole-brain FC networks.* Once a parcellation has been chosen FC values are commonly estimated in a pairwise manner between the chosen regions-of-interest (ROI) over a defined time window of the fMRI recording, using linear correlation or covariance. As each ROI will typically consist of a large number of voxels, a representative timecourse may be extracted instance by averaging or by taking the first principal component. The resulting whole-brain FC network can then be represented as a symmetric square matrix containing the total number of FC values, and offer a view of the brain's functional architecture not biased towards RSNs. RSN-specific information is not lost, but will rather be represented as sub-networks of the FC matrix with particularly strong intra-connectivity. Adapted from Deco and Kringelbach (2014). **B**) *Thresholding and Binarisation* Graph theoretical analyses work on binary networks, also referred to as adjacency matrices. Either an edge is present or absent. This does not immediately apply to functional connectivity, since statistical dependencies do not fall on a binary scale. A commonly used step to transform an FC matrix into an adjacency matrix is to choose a threshold, above which a functional connection is regarded a binary connection, i.e. a present edge, while the rest are taken as absent edges. Potential negative FC values by taking the absolute of the FC matrix or by considering only positive FC values (Fornito et al., 2013). **C**) *From order to randomness.* A completely ordered topology is characterised by clustered organisation of nodes, without long-range edges. Information is highly segregated, and the topological distances, or paths, are long. On the other end of the scale, randomness allows any type of edge, disregarding local structures of clusters or modules. Paths are therefore short, and integration high. Complex and small-world structures are found in between order and randomness and promote optimal balance of integration and segregation. Adapted from Bullmore and Sporns (2012) and from Watts and Strogatz (1998)). **D**) *Graph-theoretical concepts.* The topology of a graph is independent of geometrical space. Two nodes situated closely in physical space may thus be topologically distant if they are not connected through edges. Degree is equal to the total number of edges of a node, while clustering summarises a node's tendency to be part of local triangular groupings of nodes. From these basic concepts it is possible to calculate a range of local and global measures. Locally the graph may for example be parted into modules or communities supporting segregation, while certain nodes may act as global connecting hubs linking different modules together, and thus promoting integration in the graph. Small-worldness is typically estimated at the global level. Adapted from Sporns (2011).

Spoormaker and colleagues were the first to apply a whole-brain network perspective to spontaneous BOLD data during sleep. Parcellating the brain into the 90 cortical and subcortical regions of the AAL atlas, they computed whole-brain FC networks for wakefulness, N1, N2, and SWS, and investigated vigilance-specific changes in FC strengths. Compared to wakefulness they found widespread increases in N1 and N2, and decreases in SWS. The global increases during N1 and N2 were reserved mainly to primary sensory cortico-cortical connections, while FC values were particularly decreased for long-range connections during SWS. A specific "dropout" of functional connectivity between the cortex and the thalamus during N1 was highlighted. Notably, no correction was made for multiple comparisons in their comparative network analyses. However, the widespread increases in FC values during N1 and N2 were later confirmed by Tagliazucchi and Laufs (2014), and are also seen reflected in the RSN study by Larson-Prior and colleagues (Larson-Prior et al., 2009), where several networks increased their FC strength in N1.

The specific uncoupling of the thalamus in early NREM sleep has been interpreted in light of intracortical recordings showing thalamic changes in dynamics preceding those of the cortex at sleep onset (Magnin et al., 2010, Spoormaker et al., 2011). Later studies confirmed thalamocortical FC as sensitive to wake and NREM sleep stages (Picchioni et al., 2014, Tagliazucchi and Laufs, 2014, Hale et al., 2016), while emphasising the importance of considering the subdivisions of the thalamus, e.g. specific and non-specific nuclei. In general anaesthesia, where intracortical evidence point to an early change in cortical rather than thalamic dynamics during induction (Velly et al., 2007, Uhrig et al., 2014) (although see Baker et al. (2014b)), reports have been provided from whole-brain FC studies both for and against thalamocortical disconnections (Boveroux et al., 2010, Mhuircheartaigh et al., 2010, Liu et al., 2013).

Spoormaker and colleagues were also the first to apply graph-theoretical analyses to whole-brain FC networks during sleep (Spoormaker et al., 2010). Specifically they examined a global measure of small-worldness. The reported small-worldness indices followed a U-shaped curve across vigilance conditions, with high indices in wakefulness, low in N1 and N2, and high again in SWS. Together with the results from their FC comparisons this led the authors to propose that while information might be shared in the form of stronger connectivity during N1 and N2, it is done so in a network of lowered clustering and more random topology compared to wakefulness. Referring to the concomitant decrease in long-range (physical as opposed to topological) connections, the restored small-worldness, observed in SWS, was suggested as a sign of complexity in a spatially disconnected network. The random topology in N1 and N2 were interpreted as relevant to consciousness in the framework of the Integrated Information Theory (Spoormaker et al., 2010, Spoormaker et al., 2011), while the spatially constrained information processing in SWS was suggested as optimal for information re-processing and memory consolidation (Spoormaker et al., 2011).

Spoormaker and colleagues found decreased clustering, i.e. segregation, to be the driver of the decreased small-worldness in N1 and N2 (Spoormaker et al., 2010). While reporting a similar decrease in small-worldness in N1, Uehara and colleagues found an increased characteristic path length, i.e. integration, as the driver of this difference (Uehara et al., 2014). Uehara and colleagues used a much more fine-grained parcellation than the AAL used by Spoormaker and colleagues Spoormaker et al. (2010). Using a third parcellation, Larson-Prior and colleagues reported an increased level of small-worldness in N1 and N2

compared to wakefulness, thus contradicting the two other studies (Larson-Prior et al., 2011). These conflicting results illustrate the caution, with which one should interpret simple graph-theoretical measures (van Wijk et al., 2010, Zalesky et al., 2012, Fornito et al., 2013).

Modularity represents a more sophisticated way of estimating a network's segregation. Through iterative optimisation, a whole-brain network may be parted into modules of strong internal connectivity compared to external connectivity (Newman, 2006, Meunier et al., 2010). Global modularity reflects the degree to which a network can be split into such modules. Boly and colleagues tested the differences in global modularity between wakefulness and N2/SWS combined, and found a significant increase during sleep (Boly et al., 2012b). Tagliazucchi and colleagues examined all NREM stages, and in line with Boly and colleagues found increased values in N2 and SWS compared to wakefulness and N1 (Tagliazucchi et al., 2013b). Spoormaker and colleagues reported largely converging results, and at the regional level demonstrated that in particular frontoparietal areas became segregated from the network during NREM sleep (Spoormaker et al., 2012).

Taken together, the graph-theoretical investigations of modularity provided evidence for a more segregated whole-brain network during N2 and SWS. While these findings lent further support to the Integrated Information Theory of consciousness (Tononi, 2008), they also pointed to whole-brain networks as more sensitive to vigilance than RSNs had proven.

The clearest demonstration of this sensitivity came from a study combining whole-brain FC networks and a support-vector-machine-learning algorithm to separate wakefulness and the various stages of NREM sleep (Tagliazucchi et al., 2012b). Solely based on FC matrices belonging to the polysomnographically scored epochs of wakefulness, N1, N2, and SWS, Tagliazucchi and colleagues trained a classifier on one part of their large dataset, and was able to score, with high accuracy, the remaining data into vigilance stages (Tagliazucchi et al., 2012b). In a follow-up study Tagliazucchi and Laufs further proved the efficiency of this strategy by applying an equivalent classifier to non-sleep-scored resting-state recordings (Tagliazucchi and Laufs, 2014). They were able to show how the derived whole-brain FC networks became more similar to those of N1 sleep as the recording time progressed, essentially suggesting that participants were falling asleep while they were in fact supposed to be awake (... a rather common experience for those who have participated in many neuroimaging studies...).

The sleep-scored dataset from Spoormaker and colleagues Spoormaker et al. (2010), Spoormaker et al. (2012) was subsequently examined with a similar approach (Altmann et al., 2016). Again a support vector machine was able to effectively discriminate between epochs of wakefulness and individual stages of NREM sleep.

These studies confirm, that although it may not have been obvious from RSN studies, whole-brain FC networks do in fact undergo very specific re-configurations with changes in vigilance. It is therefore worth paying particular attention to the comparative network analyses that have been reported on the datasets from Spoormaker and colleagues (Spoormaker et al., 2010, Spoormaker et al., 2012, Altmann et al., 2016) and Tagliazucchi and colleagues Tagliazucchi et al. (2012b), Tagliazucchi and Laufs (2014). Among the most consistent of these are widespread increases in FC values during N1 and N2, an N1-specific dropout of thalamocortical FC, along with a generally decreased FC level in SWS (see Figure 1-10).

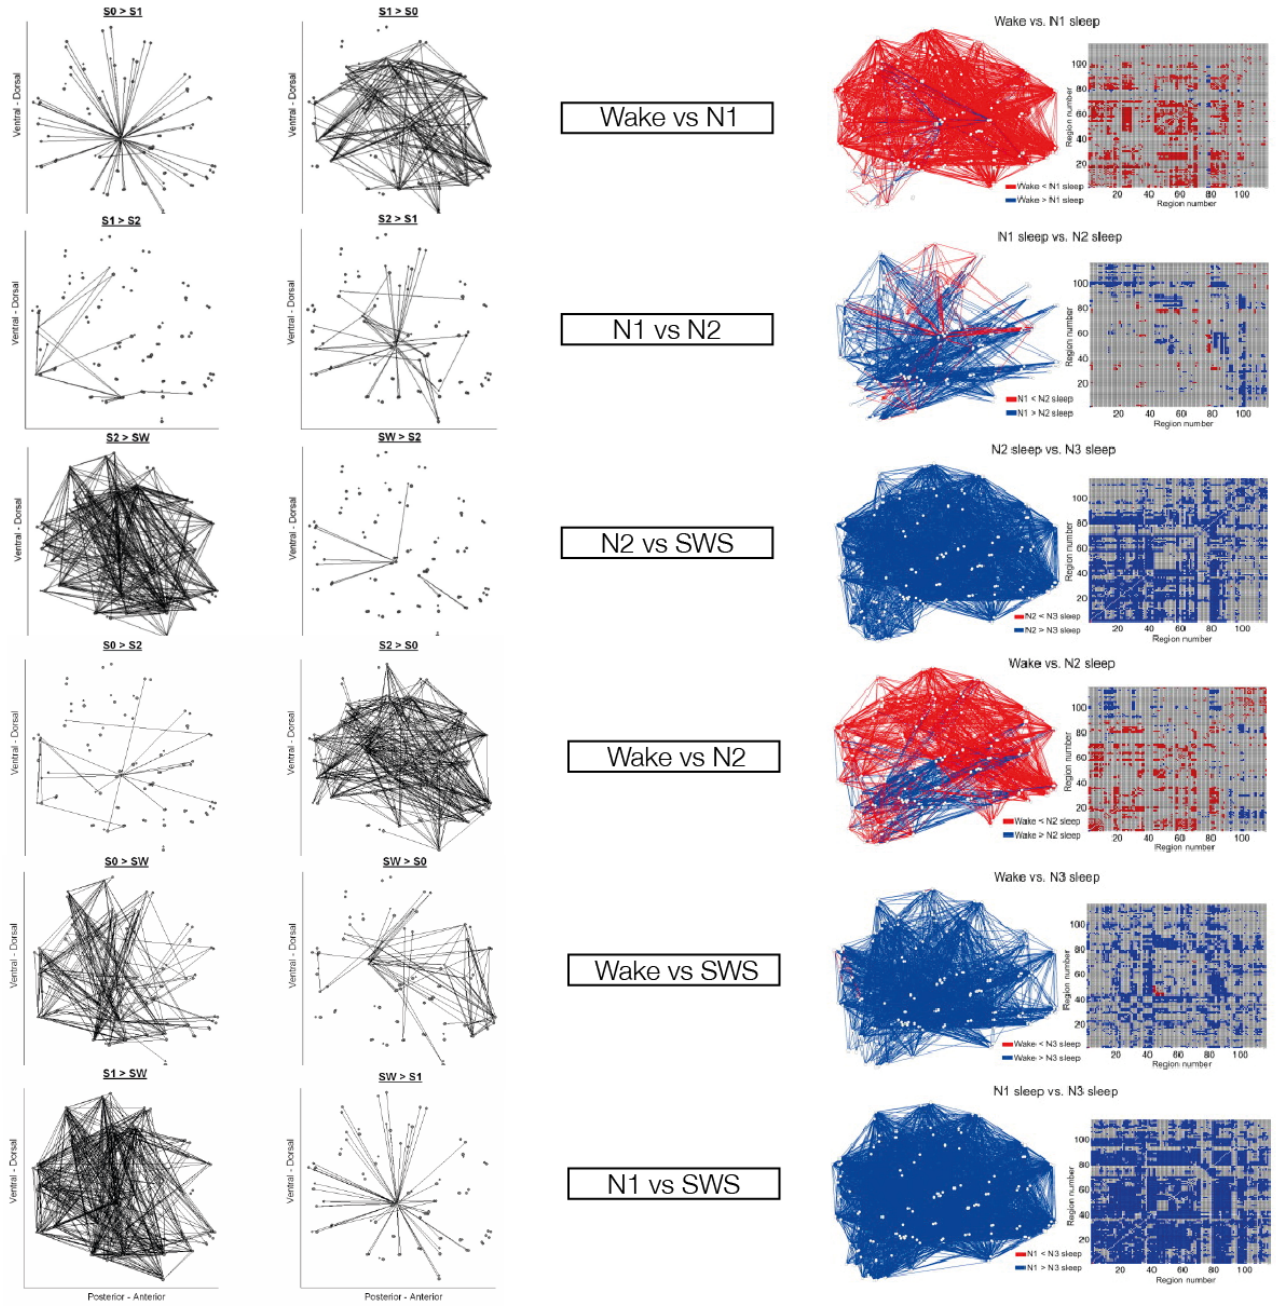

Figure 1-10 ***Whole-brain FC network changes from wakefulness through NREM sleep*** On the left are shown the results of the comparative network analyses from Spoormaker et al. (2010) (and supplementary material). No correction for multiple comparisons applied. On the right are the results from Tagliazucchi and Laufs (2014), where a Bonferroni method was used to correct for multiple comparisons. These differential networks are important because they comprise the information that made FC-based sleep classification possible, demonstrating the vigilance-sensitivity of whole-brain FC networks. Note the common N1-specific decreases in thalamocortical FC, as well as the large and widespread increases in N1 and N2. The clearest difference between the two sources is frontal increases in SWS compared to wake, reported by Spoormaker et al. (2010), but not by Tagliazucchi and Laufs (2014).

1.5.5 *Structural connectome basis of functional connectivity, and sleep*

Neural communication is facilitated through axonal connections and synapses, which makes information about the brain's structural wiring paramount to understand the basis of its activities. Determining structural connectivity between neurons at the micro and meso-scale requires elaborate invasive lesioning or tracing techniques, and can only be performed in animals. The careful mapping of the regulatory sleep-wake networks in the brainstem, hypothalamus, and basal forebrain, represent an example of meso-scale structural connectivity. In humans, equivalent mapping used to rely solely on post-mortem dissections, but the introduction of diffusion MR techniques, such as diffusion tensor imaging (DTI), has made it possible to resolve, in vivo and non-invasively, macro-scale structural connectivity (SC) (Hagmann, 2005, Hagmann et al., 2007).

Diffusion data from DTI can be combined with tractography analyses to get an estimate of how white-matter fibre tracts connect gray-matter areas throughout the brain. Given a parcellation of the brain, whole-brain SC networks can be constructed (see Figure 1-11).

Whole-brain SC networks can be represented in matrix form, and analysed in similar ways as whole-brain FC networks. It has become common to talk about such whole-brain networks as *connectomes* (Sporns et al., 2005), and the study of them as *connectomics* (Hagmann, 2005). SC and FC may then be regarded as different *modes* of the connectome. Please note that although the terms *connectome* and *connectomics* were originally defined from structural connectivity (Sporns, 2013), they are used rather ubiquitously today, including in reference to functional connectivity (Craddock et al., 2013).

When both of these modes are derived from the same parcellation, their patterns can be directly compared. Although FC networks are derived from constantly fluctuating activity timecourses, and SC networks are composed of much more static white-matter structures, it has become clear that these two connectivity modes show a high degree of overlap (Behrens and Sporns, 2012).

SC networks thus share many specific connections as well as graph-theoretical properties with FC networks. There are important subtleties to the relationship between SC and FC, however, and they are not to be thought of as mirror images of each other. In general, a structural connection is more predictive of a functional connection than the other way around (Honey et al., 2009). For example, within the areas of the DMN, strong structural connections can be found between the posterior cingulate cortex and the medial prefrontal cortex, and between the posterior cingulate cortex and the medial temporal lobe, but not between the latter and the medial prefrontal cortex (Greicius et al., 2009). Furthermore, while cross-hemispheric structural connectivity is generally sparse, functional connections between right and left homologues are among the strongest. Intriguingly, brain plasticity can also be measured at the macro-scale of SC, where persistent activity in certain areas, e.g. through behavioural training, or through invasive treatment with electrical stimulation, can induce changes in structural connectivity (Scholz et al., 2009, Hartevelt et al., 2014), adding further salience to the interdependence between brain activity and SC.

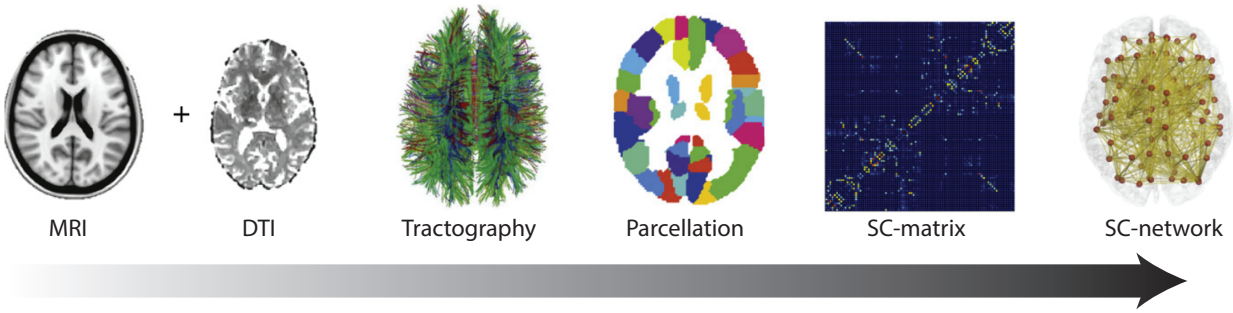

Figure 1-11 ***Construction of whole-brain SC networks*** Diffusion imaging is sensitive to the motion, or anisotropy, of water molecules. Since the diffusion of water becomes more anisotropic near white-matter fibres and their myelin-coated axons, it is possible to infer the distribution of such large fibres across the brain. By considering the voxel-wise diffusion map of the brain's white matter, continuous large-scale fibres can be tracked, representing bundles traversing the brain, linking voxels of gray matter. This step is called tractography (Hagmann et al., 2007, Le Bihan and Johansen-Berg, 2012). Within a given parcellation of the brain each ROI may therefore be assigned structural connections (SC) to all other areas. Repeating this for all ROIs yields a whole-brain SC network. DTI tractography is not as precise as invasive tracing, but have several shortcomings. Besides the technical challenges related to the correct estimation of voxel-wise tract-directionality, the modality itself has its fundamental limitations (Jones et al., 2013). First, tracts are estimates of large bundles of axons. Hence, it is not possible to determine the direction of a tract, but merely its presence or absence. Second, the estimation of the exact starting and termination points is challenging. Within a parcellation scheme these are assumed to correspond to the included ROIs, but parcellations are, as mentioned, tentative organisations of the brain. In principle, the strength of a non-directed structural connection should correspond to the number of axons it includes, but the current methods for tractography only allows for approximations, by assuming a relationship between the probability of the identified trajectories and fibre strength or count. (From Lord et al. (2017))

The intricate relationship between whole-brain SC and FC networks was recently explored in the context of sleep (Tagliazucchi et al., 2015). Intriguingly, it was found that sleep significantly influenced the correlation between SC and FC matrices, such that whole-brain FC patterns became more similar to the structural connectome in N2 and SWS compared to during wakefulness. Converging evidence had come from a slightly earlier study of general anaesthesia in monkeys, where the same increase in SC-FC similarity followed the increasing concentration of propofol dosage (Barttfeld et al., 2015). Both groups interpreted the increased SC-FC similarity as a sign of decreased exploration of the structural connectome during reduced levels of vigilance. Wakefulness was thus hypothesised to be associated with dynamics that would lead to a larger state repertoire diverging more from the structural backbone. Given the strong imprinting of RSNs in the structural connectome, these perspectives also offered a possible explanation for the consistent identification of RSNs during sleep and general anaesthesia (see 1.5.3).

1.5.6 *Whole-brain computational modelling, and sleep*

We have seen how SC appears to shape FC at the whole-brain scale, and we have seen how this relationship seems to change in meaningful ways in the context of sleep. These observations are valuable for our understanding of how large-scale brain activity is organised, but they remain descriptive in nature. As such, they do not address the mechanistic principles underlying *how* the structural connectome is being explored to produce the spatio-temporal patterns of FC. Determining the generative mechanisms of whole-brain activity represents a level of description that goes beyond the correlational. Even if the small-world structure of FC networks can be interpreted as a sign of scale-free or critical dynamics, this remains a qualified guess, since we only have access to the resulting patterns and not the cause of them. Moreover, the apparent similarity

between SC and FC networks may suggest that FC patterns are generated on a backbone of the SC. This too, however, remains a correlational observation and does not establish a causal relationship. These causal levels of description are hidden from observational analyses of neuroimaging data, and the only way to address them non-invasively is through modelling.

Whole-brain computational models explore the relationship between FC and SC patterns to simulate the observed characteristics of whole-brain activity (Deco et al., 2011, Cabral et al., 2014a). This is done by representing each ROI of a whole-brain network by a set of dynamical equations, which are then coupled through the structural connections given by an SC matrix. A general conclusion from a range of whole-brain computational modelling studies, exploring differing dynamical scenarios, is that whole-brain FC patterns can indeed emerge from the structural connectome (see Figure 1-12).

Local dynamics that have been tested include: chaos, noise, oscillations, multi-stability, criticality, as well as combinations of these. The fact that a range of different dynamics can be combined with the SC to fit the empirical FC patterns points to the challenge of validating such models that span a large number of parameters due to the many ROIs and connections (Woolrich and Stephan, 2013). Nevertheless, a common feature of most of the tested models is that the optimal working point is characterised by some kind of dynamic instability (Cabral et al., 2014a). As such this has been interpreted as a sign of the adaptation and flexibility through which the brain supports optimal behaviour and cognition (Deco et al., 2011).

An additional variable that is hidden from direct observation is effective connectivity. While functional connectivity describes *whether* information is shared between areas of the brain, *effective* connectivity includes the direction in which this information is flowing (Friston, 2011). Again this requires modelling to estimate non-invasively. Thus far effective connectivity has mainly been resolvable in rather simple networks involving a small number of nodes, as in dynamic causal modelling (DCM) (Friston et al., 2003). The large number of ROIs and connections in whole-brain networks complicates the necessary model inference (Woolrich and Stephan, 2013), although important progress is being made (Friston et al., 2014, Gilson et al., 2016).

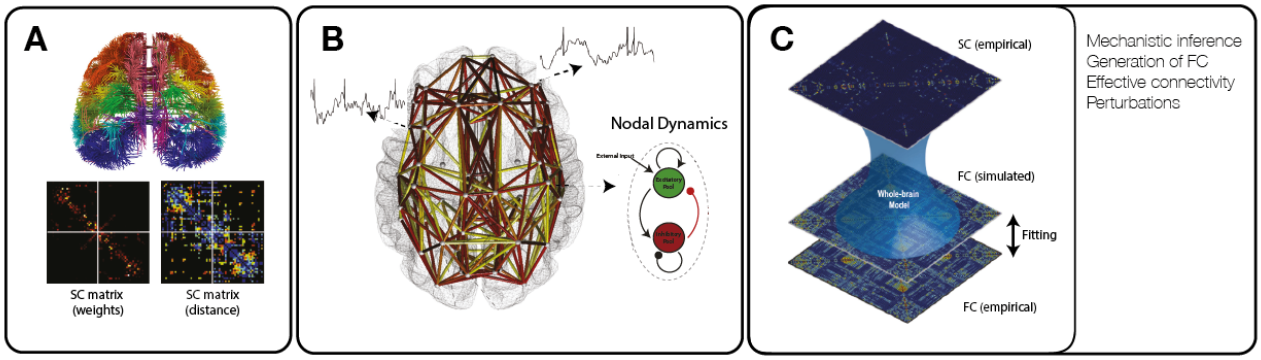

Figure 1-12 ***Whole-brain computational modelling***. **A)** In a whole-brain model (**B**) each area, or ROI, of a given parcellation is represented as a mathematically formulated unit of some specified internal activities. These units are then allowed to communicate with each other through the connections given by their structural connections derived from DTI tractography (**A)**. The computational element comes in when this network system of coupled units is simulated over time. By observing the timecourses created at each of the units, or ROIs, it is possible to analyse and derive patterns of FC in the same way as it is done for empirical neuroimaging data The first objective is then to establish whether the addition of local dynamics can explain the empirically recorded FC patterns over and beyond what may be inferred from the structural connections alone (**C**). In the search for local generative mechanisms of whole-brain FC networks it is important to note that a range of different model formulations have been tested and shown to yield rather comparable performances. Specifically these models differ in terms of the dynamic properties they allow for the individual local units (**B**). The mathematical implementations used to model these dynamics vary in their level of reduction, in the sense that some will try to model many components of ROI, such as different pools of excitatory and inhibitory neurons, while others will simply reduce the total group of neurons to one set of dynamics. Reduction is common to them all, since none of them are aiming to provide a full account of individual neurons. This is justified by observations from statistical physics, showing how macroscopic physical systems (brain areas in this case) obey laws that are independent from their microscopic constituents (here: individual neurons or circuits) (Haken, 1975). (From Lord et al. (2017))

Thus far, whole-brain computational modelling has primarily been used to explore generative mechanisms of whole-brain FC during wakefulness based on resting-state recordings. Only one study has addressed the changes occurring during sleep. Deco and colleagues modelled each cortical area of a 66 ROI parcellation (Hagmann et al., 2007) as a pool of mutually interconnected excitatory and inhibitory neurons, coupled through an SC matrix (Deco et al., 2014), in a spiking neurons attractor model (Deco et al., 2011). Compared to other more reduced models, the one used in Deco et al. (2014), includes a term for the ROI-specific concentration of the excitatory neurotransmitter acetylcholine. By tuning this parameter, Deco and colleagues showed how the local dynamics changed from fast desynchronised to slow synchronised spiking as the levels of acetylcholine were lowered (Deco et al., 2014). By taking this as an effective induction of slow oscillations in the cortical areas, and assuming this to represent sleep, the authors went on to investigate how simulated whole-brain FC patterns were affected by these changed local dynamics (Deco et al., 2014). It was thus shown that patterns of simulated FC maintained their resemblance to the empirical resting-state FC at moderate decreases of the acetylcholine parameter, and decreased at low levels of acetylcholine when slow oscillations dominated the model. The authors interpreted this in light of RSNs preservation during early stages of NREM sleep (Boly et al., 2008, Horovitz et al., 2008, Larson-Prior et al., 2009, Sämann et al., 2011) and the uncoupling of these during slow-wave sleep (Horovitz et al., 2009, Sämann et al., 2011) (see above 1.5.3). However as noted above, RSNs are not the most sensitive description of sleep (see 1.5.4). It would be particularly interesting to see if these cortical slow oscillations would explain some of the more

detailed, empirically derived changes in whole-brain FC networks associated with sleep (as reviewed in 1.5.4).

As for the increased similarity between SC and FC during sleep, reviewed above (Barttfeld et al., 2015, Tagliazucchi et al., 2015), this was first reported subsequently to the study of Deco and colleagues (2014). Meanwhile, Deco and colleagues also provided evidence in favour of the possible decrease in state exploration during SWS (see above 1.5.5). As such, they showed that the variance across the measured FC values was lowest during low levels of acetylcholine, when slow oscillations were present locally. This was taken as a proxy of decreased state exploration.

1.5.7 *Electrophysiological underpinnings of functional connectivity during sleep*

Until now, I have focussed my summary of whole-brain activity during sleep on findings obtained with fMRI. However, the fact that fMRI measures hemodynamics and not neural activity *per se* is a central limitation to fMRI as a modality to study brain function (Logothetis, 2008). The blood-oxygenation level dependent (BOLD) signal (Ogawa and Lee, 1990) is, as its name implies, sensitive to changes in blood oxygen. This contrast is driven by a cascade of physiological events in the vasculature of the brain, that results in a reduction in deoxyhemoglobin and an increase in the BOLD signal (see Figure 1-13).

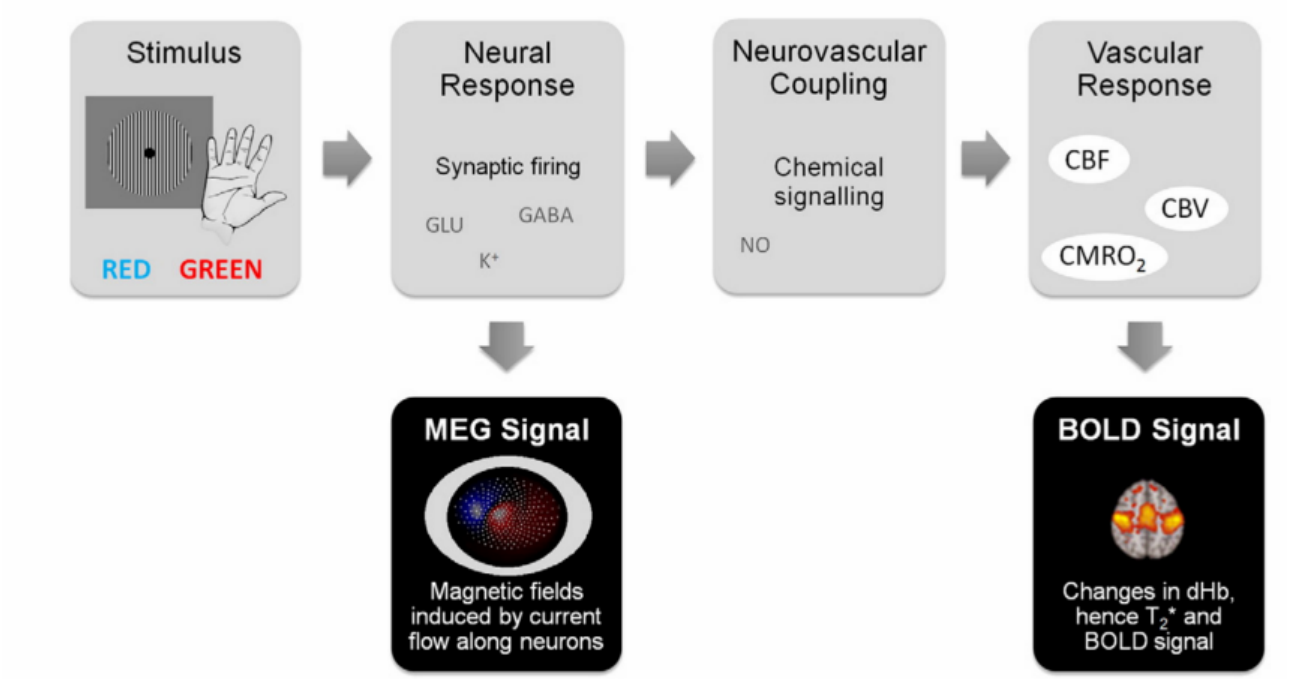

Figure 1-13 ***Signal origins of fMRI BOLD and MEG***. It is commonly accepted that a localised increase in the BOLD signal corresponds to an increase in brain activity, which can be sought uncovered through deconvolutions of the so-called hemodynamic response function (HRF) (Bandettini, 2016). Electrical activity in neurons will also produce extracranial signals directly. While EEG electrodes pick up the electrical potentials of these extracranial signals, MEG sensors measure the induced changes in the magnetic field. We know that mainly the current sources that are placed tangential to the scalp will generate extracranial magnetic signals, while both radial and tangential sources contribute to the electrical potentials, ultimately leading to differences in the anatomical profiles of EEG and MEG. (From Hall et al. (2014)).

From intracortical recordings in animals and human patients we know that a consistent relationship exists between locally generated local field potentials (LFPs) and local changes in BOLD activity. Specifically, different frequencies of the LFPs have been found to correlate differently with BOLD, such that lower frequencies in the delta to beta ranges are negatively correlated, while activity in the higher gamma range is positively correlated with local BOLD signals (Logothetis et al., 2001, Mukamel et al., 2005, Shmuel et al., 2006). Importantly, however, these local relationships do not necessarily explain the long-range co-variations observed in fMRI FC studies (Foster et al., 2016).

Intracortical recordings have shown that both very slow fluctuations in the LFP, so-called slow cortical potentials (SCP) (delta and sub-delta range) as well as slow modulations of gamma power can be expressed across long distances in the brain between areas that are functionally connected in the BOLD signal (He et al., 2008). Both of these electrophysiological phenomena have been proposed to underlie whole-brain FC patterns (He and Raichle, 2009).

A different approach to linking the fMRI-derived FC patterns with electrophysiology is through non-invasive techniques, such as EEG and MEG. We have a good understanding of the signal origins of these modalities and know that they are direct measures of neural activity (Baillet et al., 2001). As such, they are both sensitive to extracranial signals generated by extracellular electrical currents arising from synchronised post-synaptic potentials in the dendrites of thousands of neurons located mainly in cortical columns (see Figure 1-13).

Concomitant recordings of fMRI and EEG make it possible to relate specific FC patterns or RSNs derived from fMRI to oscillations measured with EEG. Such studies have shown that modulations in power (or amplitude) of different band-passed EEG signals correlate with varying degrees to FC patterns and RSNs of fMRI, albeit with no single frequency band showing a one-to-one match (Mantini et al., 2007, Laufs, 2008, Schölvinck et al., 2013, Tagliazucchi and Laufs, 2015). Similar approaches have revealed a temporal correlation between EEG microstates (Lehmann et al., 1998, Koenig et al., 2005) and fMRI-derived RSNs (Britz et al., 2010).

The obvious limitation of these EEG-fMRI studies is that the electrophysiological component is represented by oscillatory activity in electrodes on the scalp, and hence provides little spatial information. Both EEG and MEG are subject to a so-called inverse problem, denoting the fundamentally ill-posed relationship between the measured extracranial signals and the location of the intracortical sources (Helmholtz, 1853, Baillet et al., 2001). However, the fact that the signals measured by MEG get less distorted compared to the ones picked up by EEG has facilitated the development of efficient source-reconstruction strategies for MEG (Hansen et al., 2010). Based on the high number of sensors (~100 – 300) in a modern whole-head MEG system, and under various assumptions about the underlying sources, it is thus possible to estimate signal origins and significantly increase the spatial resolution of MEG.

Spontaneously co-varying power fluctuations of band-limited and source-reconstructed MEG signals, particularly in the alpha and beta range, have thus been observed to organise into large-scale distributed patterns of high resemblance to the FC networks known form fMRI resting-state studies (Pasquale et al., 2010, Brookes et al., 2011a, Brookes et al., 2011b, Mantini et al., 2011, Hipp et al., 2012, Luckhoo et al.,

2012, Hall et al., 2014). This makes MEG a particularly promising modality for linking the expansive fMRI literature on spontaneous brain activity to its electrophysiological basis in general, and to neural oscillations in particular. Whole-brain computational models have also shown promise to help elucidate the generative mechanisms of these MEG-derived FC networks (Cabral et al., 2014b). There are many hypotheses about, and evidence for, the functions related to specific oscillatory frequencies (Fries, 2005, Buzsaki, 2006, Wang, 2010, Palva and Palva, 2012, Engel et al., 2013), making such translations desirable.

To date the electrophysiological correlates of whole-brain FC changes in sleep have received little attention. I will now review the few studies related to this (see Figure 1-14).

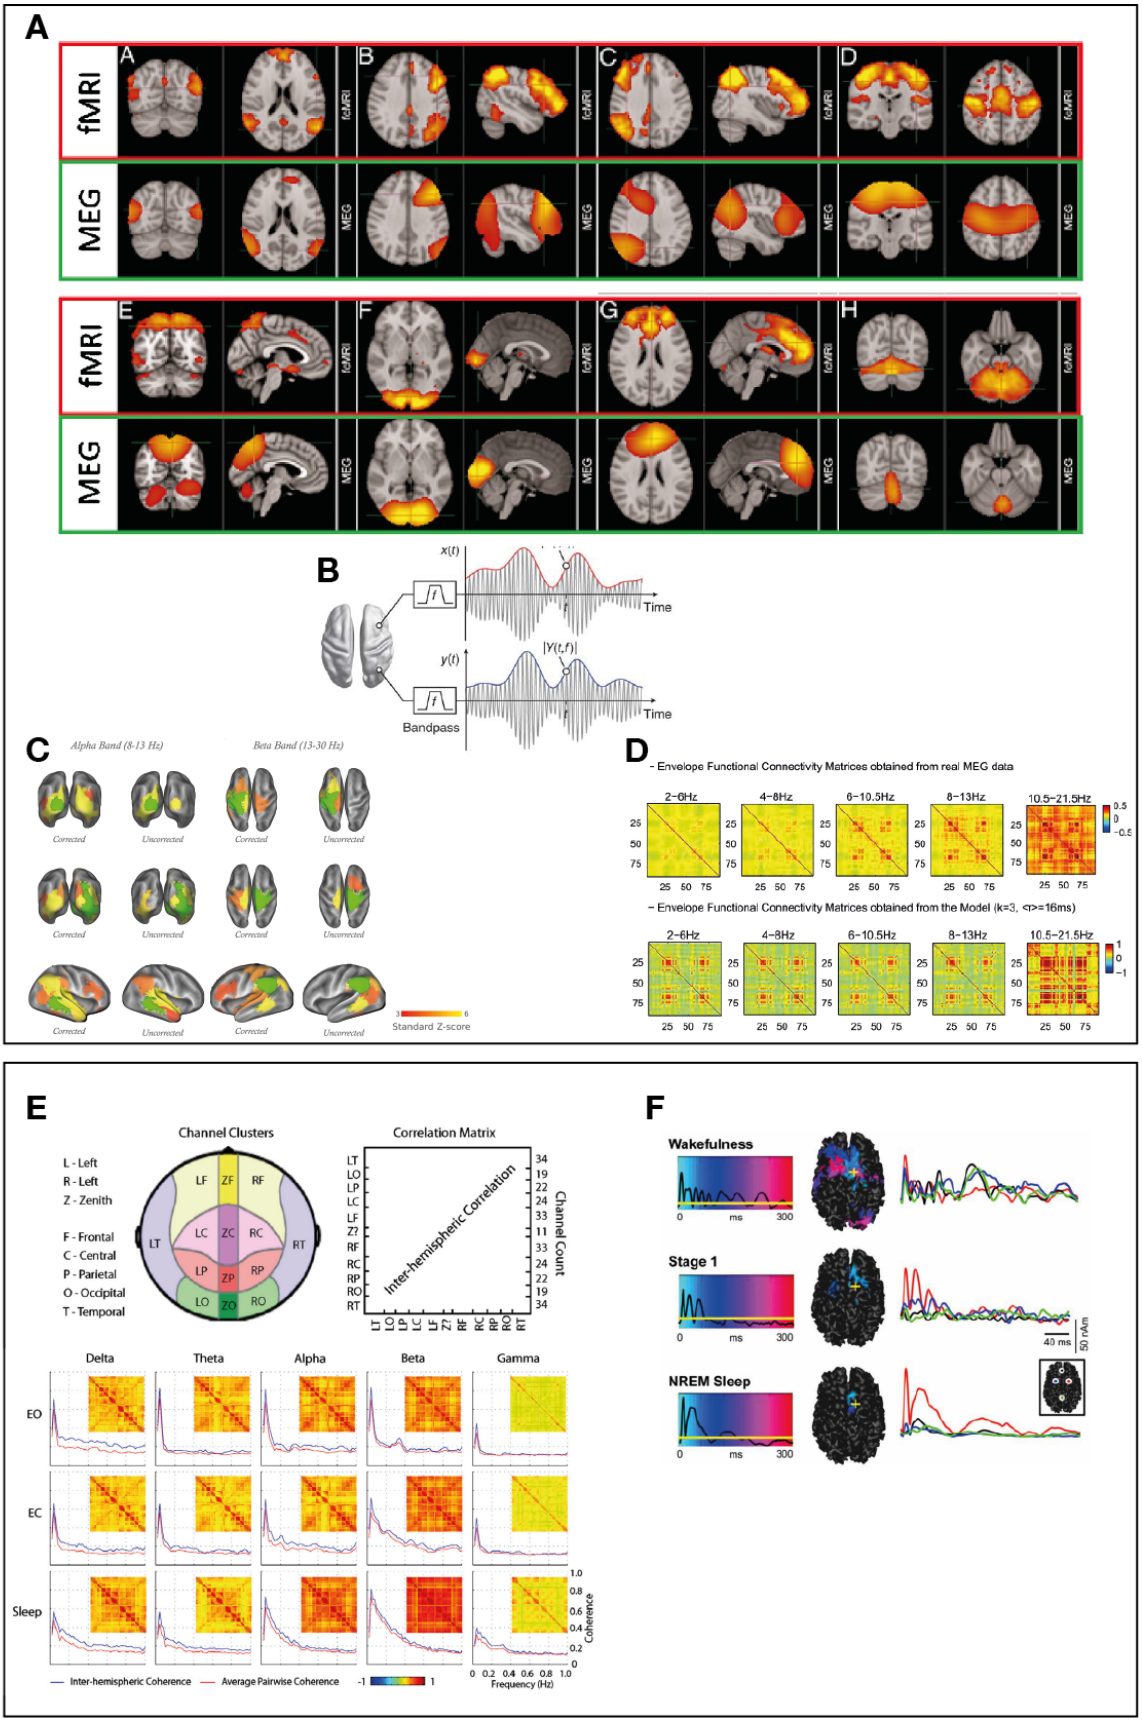

Figure 1-14 ***Electrophysiological underpinnings of whole-brain functional connectivity.*** **A**) Temporal ICA decomposition of source-reconstructed resting-state MEG data can yield whole-brain patterns with high resemblance to fMRI-derived RSNs. Adapted from Cabral et al. (2014a), but originally from Brookes et al. (2011b) **B**) Rather than the raw electrophysiological amplitudes whole-brain patterns are commonly derived as correlations between fluctuations in power of specific frequency bands. Adapted from Hipp et al. (2012) **C**: The estimation of FC from source-reconstructed MEG data is biased by the so-called source-leakage phenomenon. Careful correction for this improves the reproduction

of FC-maps known from fMRI (from Colclough et al. (2015)). **D**: A whole-brain computational model coupling Kuramoto oscillators through a structural connectome has been able to reproduce MEG-derived FC networks. Adapted from Cabral et al. (2014b) **E**: Coherence estimates of band-pass filtered power fluctuations in MEG sensor-space tested in wakefulness (eyes open and eyes closed) and early NREM sleep. Note the general peak in coherence for the very slow (~0.1 Hz) component and the increase in coherence across most frequency bands from wakefulness to sleep. Adapted from Liu et al. (2010). **F**: Investigations of changes in TMS-induced effective connectivity from wakefulness to sleep (from Massimini et al. (2005) and included supplementary material)

Using LFP and subsequent fMRI recordings on epilepsy patients, He and colleagues showed that during wakefulness the slow-cortical potential as well as gamma power modulations in the LFP showed correlation patterns within somato-motor areas that corresponded to FC extracted from the BOLD signal (He et al., 2008). The same was found during REM sleep, but in SWS it was only the SCP that showed correlations similar to BOLD, while the gamma power modulations were absent. Importantly, the fMRI recordings used to derive the BOLD FC were performed during wakefulness, and did thus not correspond to the sleep-staged LFP recordings, which makes it difficult to infer the specific impact of these findings for sleep FC.

Tagliazucchi and colleagues combined EEG and fMRI to relate whole-brain FC networks derived from BOLD to power in specific EEG frequency bands during wakefulness and N1 sleep (Tagliazucchi et al., 2012c). They were able to show that changes in vigilance significantly changed the EEG-FC correlations. For instance, while gamma power was positively correlated with FC patterns during wakefulness, power in the slower frequencies in the delta range took over during N1 sleep.

Earlier studies have assessed distributed coupling between EEG sensors during sleep (Ferri et al., 2007, 2008). Building networks between timecourses of 16 electrodes using the statistical measure synchronisation likelihood, Ferri and colleagues examined differences in small-worldness between wakefulness and NREM sleep (Ferri et al., 2008). Their results are however difficult to compare with the graph-theoretical fMRI studies reviewed above (see 1.5.4), because of the inherent differences in network nodes (electrodes vs. brain areas) and edges (synchronisation likelihood of raw EEG vs. linear correlations of slow BOLD). Unfortunately these EEG studies did not report changes in more translatable variables, such as overall coupling between electrodes.

Liu and colleagues used MEG recordings of both wakefulness and “light sleep” (where the sleep scoring was based on spectral content measures of the MEG signals and not traditional staging) and measured linear correlation and coherence between band-limited power fluctuations in the MEG sensors (Liu et al., 2010). They found evidence for globally increased correlation values, across frequency bands, in “light sleep” compared to wakefulness, which appear in line with the global FC increases during N1/N2 observed with fMRI (see 1.5.4). Liu and colleagues further performed a careful examination of the frequency components of the band-limited power fluctuations, and showed that for both vigilance conditions and for all frequency bands, cross-sensor coherence peaked for the very slow fluctuations (~ 0.1 Hz). This finding helped guide the later resting-state MEG FC studies to look for these slow modulations in power in source-reconstructed data, and earned the study much attention. Curiously few have appreciated the study’s implications for sleep.

A final line of evidence related to the electrophysiology of whole-brain networks during sleep has come from a unique methodology not mentioned so far. In a seminal study from 2005, Massimini and colleagues

combined transcranial magnetic stimulation (TMS) with high-density (hd)-EEG to compare cortical responsiveness between wakefulness and NREM sleep (Massimini et al., 2005). By stimulating cortical regions directly with magnetic pulses the authors were able to measure the responses in distant cortical regions through source-reconstructions of the simultaneous hd-EEG recordings. This elegant, and technically demanding, experimental setup allowed the authors to report a decrease in long-distance spread of the TMS-induced perturbation during N2 and SWS (pooled together as one condition). This was taken as evidence for decreased effective connectivity during NREM sleep (Massimini et al., 2005). It is tempting to relate this finding to the global decrease in whole-brain functional connectivity of SWS reported in fMRI studies, or even the increased graph-theoretical segregation and modularity (see 1.5.4). A number of reasons make this difficult, however. First, the dynamics created by the TMS perturbations are quite different from the spontaneous activity that the fMRI studies are based on. Second, if there is a direct relationship between the perturbation-based effective connectivity reported by Massimini and colleagues and the fMRI-based whole-brain FC, one would expect a significant increase in the cortical spread during N1 and N2 (see 1.5.4). However, in their supplementary material, Massimini and colleagues in fact report the opposite (Massimini et al., 2005).

1.5.8 *Dynamic properties of spontaneous brain activity during sleep*

An assumption implicit to the FC studies reviewed so far (throughout 1.5.3, 1.5.4, 1.5.5, 1.5.6, and 1.5.7) is that functional connections between brain areas can be described as stationary over time. Traditional analysis of RSNs and whole-brain FC networks are thus based on averaging temporal relationships over entire recording intervals, typically with durations of five minutes or longer. However, the idea of the brain as temporally stationary in its functional architecture is difficult to reconcile with its role in behaviour and cognition that requires constant adaptation. Indeed, some of the sleep FC studies, reviewed above, are themselves indicators of dynamic changes in the whole-brain network as a function of vigilance. Similar, albeit perhaps less pronounced, non-stationarities are to be expected within wakefulness and the resting-state as well.

Recent studies have started to explore resting-state recordings from this perspective, and are indeed finding empirical evidence for spontaneous, dynamic re-configurations of whole-brain networks in terms of RSNs (Kiviniemi et al., 2011), whole-brain FC patterns (Allen et al., 2014), graph-theoretical properties (Bassett et al., 2011), and the relationship to SC (Liégeois et al., 2015) (see Figure 1-15A–B). Dynamic functional connectivity (dFC) has become a prominent area of research, adding means and tools to the characterisation of whole-brain spontaneous activity (Chang and Glover, 2010, Hutchison et al., 2013b, Calhoun et al., 2014, Preti et al., 2016).

The most straightforward and intuitive way of exploring and defining dFC is through a sliding-window approach (Allen et al., 2014). By segmenting recordings into shorter windows with some degree of temporal overlap, the traditional metrics of FC can be applied to each of these windows, and as such tracked over time (see Figure 1-15A).

There are methodological challenges associated with choosing the duration used for both the window segments and their overlaps, since these will influence the temporal resolution of the observed dynamics. While such challenges, as well as potential alternative methods addressing these, are covered elsewhere (Hindriks et al., 2015, Leonardi and Van De Ville, 2015, Zalesky and Breakspear, 2015, Preti et al., 2016), it is worth noting that it is inherently difficult to determine which temporal scale of recorded activity that is in fact relevant to the brain. fMRI will only be sensitive to parts of this dynamic spectrum, since the sluggish HRF of the BOLD signal (see above 1.5.7) imposes a lower limit on temporal changes that can be recovered (on the order of tens of seconds). A recent study by Baker and colleagues highlighted the need for electrophysiological evidence, when they temporally decomposed source-reconstructed MEG recordings, using a data-driven hidden Markov model, and showed that RSNs are in fact dynamic on much faster timescales (on the order of hundreds of milliseconds) (Baker et al., 2014a, Vidaurre et al., 2016) (see Figure 1-15C).

As mentioned, the observation that the similarity between SC and FC can change with vigilance could point to a dynamic rather than a static relationship between these two modes of connectivity (see 1.5.5). We have seen how whole-brain computational modelling has been used to explore the SC-dependency of traditional static FC patterns to uncover generative mechanisms. The more recent empirical measures of dFC are now being used as additional validation criteria for the various dynamical scenarios (Hansen et al., 2015, Ponce-Alvarez et al., 2015, Deco and Kringelbach, 2016, Cabral et al., 2017).

While there appears to be no studies directly addressing dFC during sleep, a few studies have provided evidence related hereto.

Starting with a study mentioned previously, Tagliazucchi and colleagues (2012c) found changes from wake to N1 sleep in how fMRI FC patterns correlated with power in frequencies of simultaneously recorded EEG (see 1.5.7). This points to a vigilance-dependence in the dynamics underlying the observed FC.

Watanabe et al. measured the overall connectivity within the DMN and a frontoparietal network within separate epochs of SWS and REM sleep (Watanabe et al., 2014). They found a time-dependency in intra-network FC strength, such that networks were more strongly connected towards the end of a period of SWS, while the opposite development was observed in REM sleep.

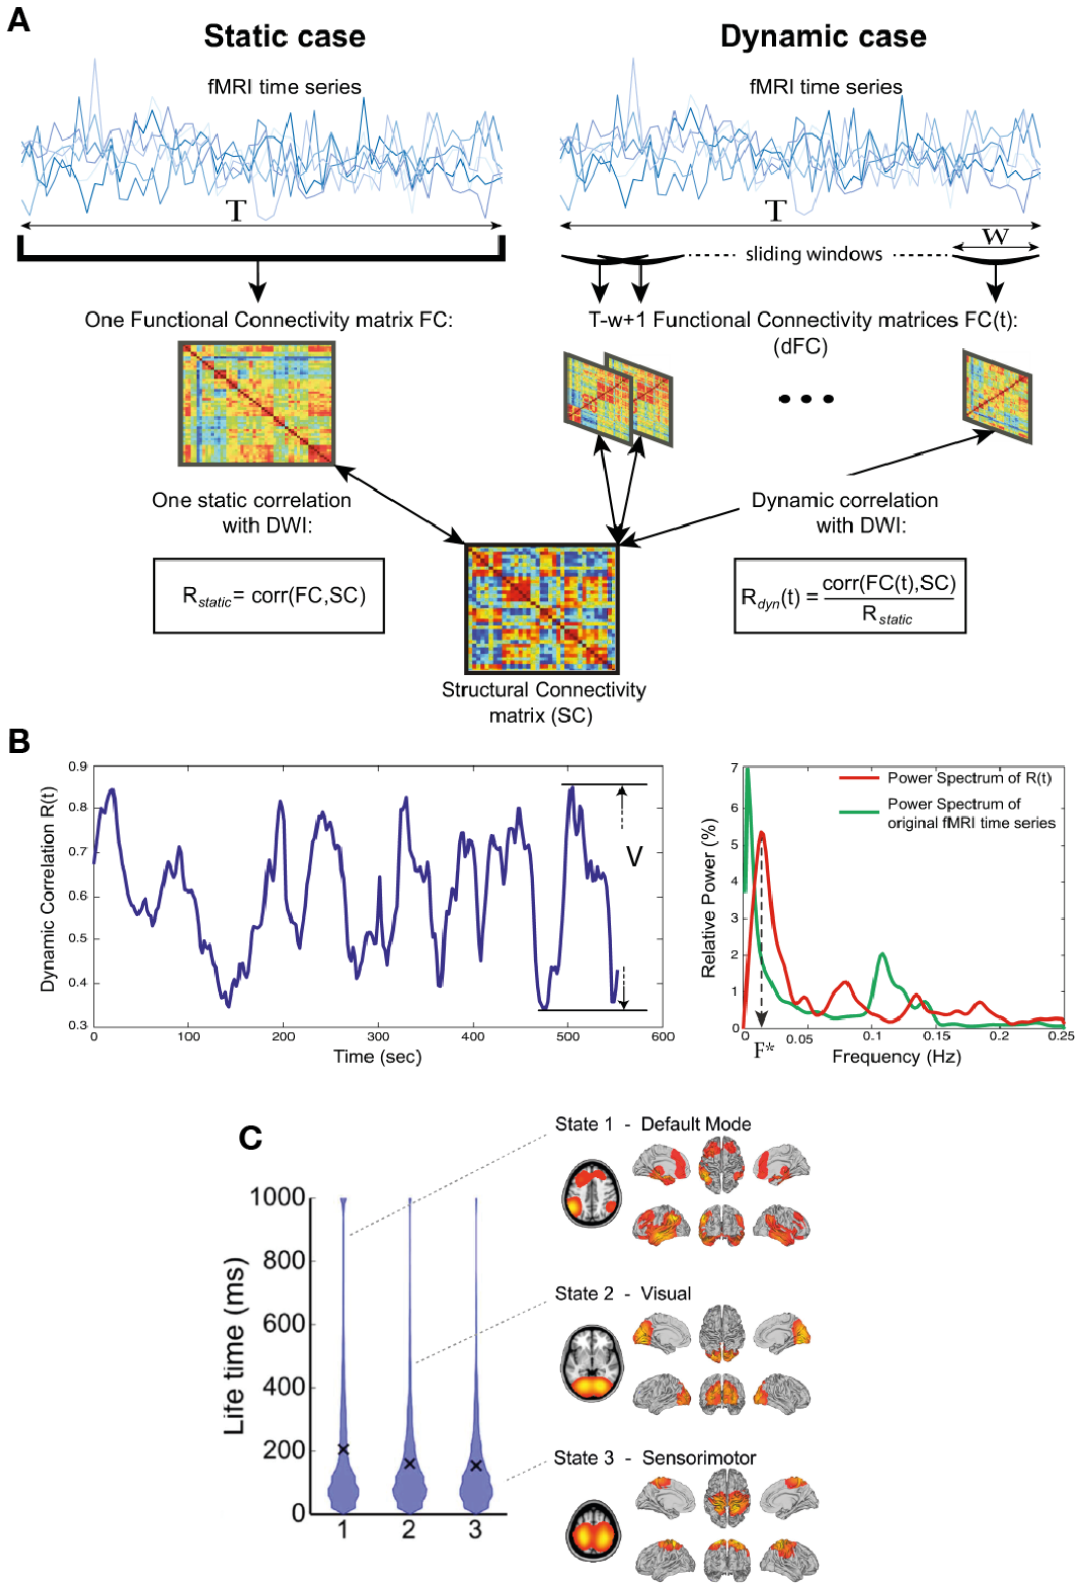

Figure 1-15 ***Dynamic (as opposed to static) functional connectivity.*** **A)** Workflow of a sliding-window FC analysis, where overlapping windows are moved through fMRI BOLD timecourses and an FC matrix for each window is derived (on the right). This contrasts the method for static FC where correlations are averaged over the full recording interval (on the left) (from Liégeois et al. (2015)). **B)** An example of what dFC may be used for is the temporal tracking of SC-FC similarity. On the left is shown how the correlation between dFC matrices and one SC matrix fluctuates over time. On the right the main frequency of this fluctuation is shown, and a peak is evident around 0.02 Hz. This corresponds to an average period of 50 seconds, indicating a rather slow dynamic reconfiguration of the whole-brain FC (from Liégeois et al. (2015)). **C)** With source-reconstructed MEG data it is possible to resolve much faster dynamics. Here a hidden Markov model was used to show how RSNs arise and dissolve with life times on the scale of hundreds of milliseconds (from Baker et al. (2014a)).

Tagliazucchi and colleagues (2013c) quantified the long-range temporal dependence of BOLD timecourses across the whole brain by examining the Hurst exponent of the auto-correlation. They found this dynamic measure of temporal self-similarity to decrease in NREM sleep compared to wake, and mapped these differences to areas belonging to the DMN. The authors emphasised how this could be a sign of decreased critical dynamics during NREM sleep, relating their finding to intracortical reports of sleep-dependent changes in distributions of neural avalanches (Priesemann et al., 2013). While the concept of critical dynamics and its associations with sleep remains a matter of debate (see e.g. Dehghani et al. (2012)), the findings in (Tagliazucchi et al., 2013c) at least show that BOLD dynamics within whole-brain networks change from wakefulness to sleep.

Brodbeck and colleagues took a different approach and explored EEG microstates during wakefulness, N1, N2, and SWS (Brodbeck et al., 2012). They were able to derive the same number (four) of microstates with matching spatial topographies (as defined in EEG sensor-space) across all vigilance conditions. Based on the putative temporal link between EEG microstates and fMRI-derived RSNs (Britz et al., 2010), this led the authors to speculate that the preservation of microstates across wakefulness and sleep could reflect the preservation of RSNs observed with fMRI (see 1.5.3). Keeping with this interpretation it is particularly interesting to note how these microstates significantly changed in their temporal characteristics. While the state lifetimes were 40-50 milliseconds on average during wakefulness, they increased significantly during N2 and almost doubled during SWS (Brodbeck et al., 2012).

1.6 ***Summary of the argument***

For almost a century our understanding of sleep has been closely linked to recordings of brain activity. Since early observations of the co-occurrence of clear visible changes in the EEG and the behavioural indications of sleep, major advances have been made in our description of the micro and mesoscopic determinants for many sleep-defining EEG signatures. Meanwhile our comprehension of the brain's spontaneous dynamics at the macroscopic scale has progressed. It is thus becoming increasingly clear that large-scale networks and patterns of brain activity are principal to a full characterisation of operations during wakefulness, such as cognition, in health as well as in disorders.

fMRI studies exploring functional connectivity and its interplay with structural connectivity at the whole-brain level are beginning to appreciate how these spatial networks are also at play during sleep. It is now clear that the brain is expressing patterns of spatial diversities beyond what a few EEG scalp electrodes are able to resolve. This has facilitated new approaches to fundamental questions about how the brain is able to support remarkably different states of consciousness like wakefulness and sleep. Theoretical concepts from information theory of networks and collective dynamics offer new and exciting interpretations of brain activity during sleep, when considered at this large scale.

From studying the effects of sleep deprivation, we know that sleep must serve a wide range of functions, but we do not currently have a clear understanding of how the brain is involved in this. A whole-brain view could be crucial for 'bridging the gap' between scientists and the realm of sleep, from where little information escapes due to a general lack of reportability and overt behaviour. However, in order to get a full

characterisation of the sleeping brain the whole-brain perspective has to be combined with an account of the many temporal scales involved in sleep. This remains to be done.

The main aim of this thesis is to provide a dynamic large-scale network perspective, by integrating spatial and temporal aspects of brain activity during sleep. Specifically I investigate what this level of description can tell us about the contrast between wakefulness and the various stages of sleep and the transitions between these.

Chapter 2 considers sleep-scored fMRI recordings of wakefulness and individual stages of NREM sleep in healthy adult participants to revisit previous findings from comparative network analyses of whole-brain FC – a level of description, which has proven particularly sensitive to changes in vigilance. Particular attention is paid to the participant-level of the datasets, as well as the means for controlling for multiple comparisons in the statistical inferences – issues that have largely been ignored in previous reports.

The chapter further explores the vigilance-dependence of the relationship between whole-brain SC and FC. While confirming a previously reported sleep-dependent increase in SC-FC similarity, the chapter focus on the robustness of this finding at the participant-level. A novel algorithm is introduced to identify the regional sub-networks that are driving the global differences in SC-FC similarity. Temporal fluctuations of the SC-FC relationship are considered, and the average strength of functional connectivity is identified as an important global covariate.

Chapter 3 explores the electrophysiological evidence of whole-brain FC networks during wakefulness and NREM sleep. Appreciating the indirect relationship between the BOLD signal of fMRI and neural activity as well as the limited temporal resolution of this imaging modality, analyses of MEG recordings of wakefulness and NREM sleep are performed in a group of healthy participants. Static FC networks are computed for six canonical frequency bands, covering the range from 0.25 – 48 Hz, in the cortical source space, using validated methods for both source-reconstruction of the MEG signals and FC analyses. The resulting FC networks from wakefulness, N2 sleep, and slow-wave sleep are compared, and used to point to the differential importance of specific frequencies for each of the included vigilance conditions.

Chapter 4 extends the within-frequency FC analyses, presented in *Chapter 3*, to a cross-frequency perspective. Evidence is provided for strong interactions between a number of frequency bands during N2 sleep in particular.

Chapter 5 relates the observed global increases in static FC during N2 sleep to the traditional spectral description of polysomnography. By correlating temporal changes in spectral power at the sensor level to fluctuations in average FC levels of the cortical networks, a specific type of transient events is found to underlie the strong N2-specific coupling in static FC values.

Chapter 6 returns to the fMRI data in order to explore the dynamic evolution of whole-brain FC networks across transitions from wakefulness through NREM sleep. Specifically, a data-driven method based on a hidden Markov model is implemented. This method is shown to temporally resolve a number of whole-brain FC networks, which in specific combinations express a high degree of temporal concordance with the polysomnographic sleep staging, with some states exclusive to wakefulness and others to sleep. The dynamic representation not only allows for a characterisation of the spatial FC patterns specific to certain vigilance conditions, but also facilitates an exploration of the temporal dynamics associated with wakefulness and sleep. Based on this it is discussed how these spatial and temporal aspects of whole-brain activity might be involved in the transitions to and from sleep, as well as sleep-dependent functions.

In sum, temporal dynamics of large-scale brain networks of functional connectivity is explored using fMRI, DTI, and MEG. The high spatial precision of fMRI is used to show meaningful fluctuations in functional connectivity in relation to sleep. When combined with structural information derived from DTI, these FC dynamics are evident as wider state explorations during wakefulness compared to slow-wave sleep. By virtue of its higher temporal resolution, and ability to link whole-brain FC patterns to electrophysiology, MEG is used to identify frequency bands with relevance to specific vigilance conditions at the level of whole-brain networks. As such, dynamics richer than those recordable with fMRI are highlighted. Finally, pushing the temporal resolution of fMRI a novel method for resolving dynamic FC is implemented, yielding a spatio-temporal perspective of whole-brain activity during wakefulness and sleep as well as transitions between them.

Finally, *Chapter 7* discusses the implications of the evidence provided in the empirical chapters in terms of our understanding of whole-brain activity during sleep, and in general. Some of the points in section 1.4, arguing for a whole-brain perspective of sleep, are revisited in light of the findings presented in Chapters 2 to 6. The integration of large-scale spatial networks and multiple timescales is highlighted as imperative to more informed formulations of theories seeking to explain specific functions and regulations of sleep. It is argued that MEG is a promising modality for providing evidence from the fastest end of the temporal scale, and it is suggested how this might be accomplished with regards to sleep.

Chapter 2

fMRI functional connectivity on top of structural connectivity during sleep

2.1 *Introduction*

Vigilance-dependent changes in brain activity can be readily identified as spectral re-arrangements of an EEG recording (see Introduction, 1.2.1), or at a higher spatial resolution from a map of regional decreases in brain metabolism derived with PET (see Introduction, 1.5.2). Meanwhile, such accounts only offer limited explanation for why the brain ceases to support responsiveness, awareness, and arousal in a state like sleep. Functional connectivity (FC) quantifies how brain regions share information, and has become increasingly appreciated as a more relevant substrate for decoding brain operations (see Introduction, 1.5.1). fMRI studies of sleep have made it clear, however, that measures of FC have to be refined to capture changes in vigilance. The preservation during sleep of RSNs, previously thought of as prerequisites of waking cognition, illustrated the dichotomic presence/non-presence of FC as an insufficient explanatory model (see Introduction, 1.5.3). Expanding the scope beyond the spatial constraints of RSNs, and paying more attention to the reconfigurations of FC strengths in whole-brain networks has proven a more sensitive strategy in this regard. This is most strikingly demonstrated by studies decoding stages of wakefulness and NREM sleep solely on the basis of whole-brain FC networks (see Introduction, 1.5.4).

Relative to global EEG spectra and localised amplitude changes, the vigilance-dependent re-organisations of FC allow for more complex functional interpretations from a network perspective. A flexible balance between integration and segregation of information in large-scale networks is highlighted as crucial for optimal brain function during wakefulness, in neural models of cognition and consciousness. In line with this, graph theory has been used to show that whole-brain FC networks become more segregated in NREM sleep (see Introduction 1.5.4). Another notion central to whole-brain connectomics is that the information processing between distant gray-matter regions is shaped by the white-matter fibres connecting them (see Introduction, 1.5.5).SC may thus serve as a blueprint for possible whole-brain patterns of FC interactions. The finding that the overall similarity between SC and FC increase in states of decreased vigilance, such as NREM sleep (see Introduction, 1.5.5), suggest that this blueprint becomes more doctrine than guide during such states, thus restricting flexible departures.

Previous studies have focussed their analyses at the group-level, pooling data of vigilance conditions in unbalanced ways from many participants, effectively ignoring the single-participant variance. This is a particularly convenient analysis strategy when dealing with neuroimaging datasets of sleep, since sleep represents an inherently ill-controlled experimental paradigm. It is thus impossible to guarantee that a given participant will fall asleep in the noisy MR scanner, let alone reach all stages of sleep. Simply collapsing all data available from each vigilance condition in a multi-participant cohort may be a sensible first step for data exploration, however statistically this leads to unbalanced designs. This makes it difficult to evaluate the extracted FC features in terms of their participant-specific sensitivity, and as a consequence to fully appreciate their functional relevance.

This chapter revisits some of the FC-derived markers of vigilance in spontaneous fMRI recordings. With a focus on the participant-level statistics, comparative FC-network analyses as well as SC-FC comparisons are performed on fMRI datasets of wakefulness and NREM sleep from two independent cohorts of healthy participants. A novel method is introduced to resolve regional aspects of SC-FC similarity offering new evidence for how structural connectivity is shaping brain activity during changes in vigilance from wakefulness to slow-wave sleep. Finally the relationship between SC and FC is tracked over time using a sliding-window analysis.

2.2 ***Methods***

2.2.1 *Data selection and pre-processing of sleep fMRI datasets*

For the work presented in this chapter, we analysed data from two independent fMRI datasets and one DTI dataset. Technical details of the recordings, including acquisition parameters and participant demographics are provided in the Appendix (see Appendix, Datasets included in thesis).

For the *Liege* fMRI dataset (see Appendix, fMRI dataset 1 (Liege)), we chose the subset of 12 participants who reached slow-wave sleep (SWS) during the fMRI recordings as confirmed by concomitant polysomnography. Wakefulness and SWS epochs were time-matched, such that each of the 12 participants was represented by equal amounts of wakefulness and SWS data. Note that this dataset was sleep scored according to the principles of Rechtschaffen and Kales (Rechtschaffen and Kales, 1968), operating with 4 stages of non-REM sleep. Periods scored as either NREM sleep stages 3 or 4 were here considered as SWS, in accordance with the more recent AASM guidelines (Iber et al., 2007) and the sleep-scoring of the other *Frankfurt* fMRI dataset (see below). The mean number of fMRI volumes of the wakefulness-SWS data pairs was 468 (± SE 47.79, N = 12), which, given the scanning repetition time (TR) of 2.46 seconds, corresponded to ~19 minutes (1,151.3 seconds (± SE 117.56, N = 12)).

Physiological interference in the fMRI data was reduced using CORSICA (McKeown et al., 1998), which separated the fMRI signals into components related to structured noise (such as respiration, heartbeat, and movements) and brain activity. Automatic selection of noise-related components was based on their spatial distribution (see supplementary material of Boly et al. (2012b)). The fMRI data were registered to the standard space of the Montreal Neurological Institute (MNI) (Maintz and Viergever, 1998).

Further pre-processing of the fMRI data was carried out using MELODIC (Multivariate Exploratory Linear Decomposition into Independent Components) Version 3.14 (Beckmann and Smith, 2004), part of FSL (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl). We used the default parameters of this imaging pre-processing pipeline on all participants: motion correction using MCFLIRT (Jenkinson et al., 2002); non-brain removal using BET (Smith, 2002); spatial smoothing using a Gaussian kernel of FWHM 5mm; grand-mean intensity normalisation of the entire 4D dataset by a single multiplicative factor and high-pass temporal filtering (Gaussian-weighted least-squares straight line fitting, with sigma = 50.0 seconds).

From the *Frankfurt* dataset, we used a subset of data from 18 participants. Since this dataset included periods from all stages of NREM sleep (and from wakefulness) we took the opportunity to include not only wakefulness and SWS, but also epochs of N1 and N2 sleep. The 18 participant included in our analyses

corresponded to the subgroup who went through all vigilance conditions during the recordings. The vigilance-specific epochs were time-matched within participants as had been done for the *Liege* dataset. Hence, for each participant the amount of data for any given vigilance condition was constrained by the condition with the shortest period of contiguous volumes. This yielded a mean number of fMRI volumes of 102.94 (± SE 11.14, N = 18), which, given the TR of 2.08 seconds, corresponded to a mean duration of ~4 minutes (214.12 seconds (± SE 23.16, N =18)). We took the participant-specific data from the centre of the longest continuing periods of each vigilance condition.

While the outlined selection process did exclude a substantial amount of data, it served to guarantee that each participant was represented by exactly equal amounts of data from each vigilance condition. We did this to avoid any vigilance-specific bias potentially introduced by duration-dependent uncertainty on the functional connectivity estimates.

Before parting the data into vigilance-specific epochs, the EPI data were pre-processed (Tagliazucchi et al., 2012b, Tagliazucchi and Laufs, 2014). MRI and pulse artefact correction were applied based on the average artefact subtraction method (Allen et al., 1998) in Vision Analyzer2 (Brain Products, Germany). The EPI data were realigned, normalised to MNI space, and spatially smoothed using a Gaussian kernel of 8 mm^3 FWHM in SPM8 (http://www.fil.ion.ucl.ac.uk/spm/). Spatial downsampling was then performed to a 4 x 4 x 4 mm resolution. From the simultaneously recorded ECG and respiration, cardiac-, and respiratory-related noise components were estimated using the RETROICOR method (Glover et al., 2000), and together with motion parameters these were regressed out of the signals. The data were temporally band-pass filtered in the range 0.01-0.1 Hz (Cordes et al., 2001) using a sixth-order Butterworth filter.

2.2.2 *Brain parcellation and computation of functional connectivity matrices*

For both fMRI datasets we used the Automated Anatomical Labeling (AAL) atlas (Tzourio-Mazoyer et al., 2002) as a parcellation scheme to part the brain into 90 cortical and sub-cortical regions-of interest (ROIs). An AAL template was transformed into the MNI space of the EPI data, and the FSL function *fslmeants* was used to extract the average timecourse over the voxels within each of the ROIs. Thus, each participant and vigilance-specific epoch was represented by 90 spatially averaged BOLD timecourses.

We used Pearson's linear correlation coefficient as measure of functional connectivity (FC), and computed this between each of the 90 ROIs for each vigilance-specific epoch. FC_{ij}, i = 1 … 90 , j = (1 … 90 | ≠ i) was estimated in MATLAB (The MathWorks Inc., Natick, MA, 2000) using the built-in function, *corrcoef*:

$$R_{ij} = \frac{1}{T-1}\sum_{t}^{T}\left(\frac{\overline{s_i(t)-\mu_{s_i}}}{\sigma_{s_i}}\right)\left(\frac{s_j(t)-\mu_{s_j}}{\sigma_{s_j}}\right) \qquad \textbf{Eq. 2.1}$$

where T is the duration of the epoch in question, s denotes the ROI timecourses, μ and σ are the mean and standard deviation over time.

Prior to further analyses the resulting correlation matrices were Gaussianated by application of a z-transformation:

$$FC = 0.5 \ln\big((1+R) \oslash (1 - R)\big) \quad \textbf{Eq. 2.2}$$

where ln(.) is the natural logarithm, and $\oslash$ denotes element-wise division.

2.2.3 *Paired t-tests for comparative network analysis between vigilance conditions*

The first part of this chapter considers changes in static FC between vigilance conditions. For this purpose we used the FC matrices computed from the *Frankfurt* dataset, as it included all stages of NREM sleep as well as wakefulness. To test for group-differences between each pair of the vigilance conditions, we implemented a paired t-test on each connection in the FC matrix. The paired t-test was possible due to the data selection, in which each participant contributed with one epoch (of matched durations) from each of the four vigilance conditions (see above, 2.2.1). Previous studies investigating similar contrasts of FC between wakefulness and sleep have typically ignored potential uneven distributions of data, such that certain participants would contribute with only a subset of vigilance conditions. Spoormaker and colleagues thus defined connection-specific group differences by applying an arbitrary threshold (of $\Delta R > 0.20$) (Spoormaker et al., 2010). Tagliazucchi and colleagues performed unpaired t-tests on each connection in the same dataset used here (*Frankfurt*, albeit using 116 regions of the AAL atlas as opposed to the 90 used here), correcting for multiple comparisons by the Bonferroni method.

There were two factors contributing to increases in the family-wise error rate: the four vigilance conditions, which we compared in pairs, and the multitude of connections in the FC matrices. To deal with multiple comparisons we employed non-parametric permutation testing of a general linear model (GLM). The design matrix was formulated as a paired t-test with pairs of vigilance conditions as the independent variable and the connections of the FC matrices as dependent variables. Permuting the design matrix a large number of times produced null-distributions, which could be used to define the probability of finding differences by chance; the probability of finding a null-statistic more extreme than the original test-statistic corresponded to the non-parametric p-value of a test.

This however left us with a choice regarding which test-statistic to use. We chose two complementary strategies; one considering local connections with a connection-wise correction scheme, and one searching for significant connected sub-components of the FC network.

For the connection-wise correction we used the FSL function *randomise* (Winkler et al., 2014) to perform 5000 permutations of the vigilance condition labels in the design matrix (constrained within paired observations) and calculate the corrected p-values by comparing the t-stats of the un-corrected difference matrix to the t-stats in the generated null-distributions.

To look for significant sub-components we used the network-based statistic (NBS) software (Zalesky et al., 2010). Here the test-statistic used for significance testing was computed as the sum of connection weights in supra-threshold connected components. In more detail, a threshold was applied to the difference

matrix of t-stats from the original, connection-by-connection paired t-tests, and potential, connected sub-networks (components) were identified amongst the connections surviving thresholding (supra-threshold connections). The t-stat magnitudes of the connections participating in such a supra-threshold component were summed, and the resulting value comprised the test-statistic. Null-values were now defined from similar components obtained in the permuted cases. Again we used 5000 permutations to build a distribution of null-values, and compared the original test-statistic to this null-distribution to get a corrected p-value. Importantly the corrected p-value was a feature of the supra-threshold component, and could thus not be assigned to any individual connections. In this regard, the method may be compared to the scenario in cluster-based correction strategies in voxel-based parametric mapping (Bullmore et al., 1999, Hayasaka and Nichols, 2004). As for the choice of t-stat threshold this represents a free parameter of the NBS method (Zalesky et al., 2010), which was experimentally varied, until one threshold was chosen for all contrasts (see Results section for more details).

For both the connection-by-connection and the NBS, corrected p-values below 0.05 were regarded as significant.

2.2.4 *Computation of group-averaged structural connectivity*

The second part of this chapter considers the relationship between functional connectivity and structural connectivity (SC) networks.

A group-averaged SC matrix was generated using the DTI data acquired from 16 healthy participants independent from the *Liege* and *Frankfurt* fMRI datasets (for more details, see Appendix, DTI dataset). For each participant data were obtained in two phase-encoding directions, which were processed together to get one estimate of SC for each participant. The construction of the SC matrices followed a three-step procedure described in more detail below. First, ROIs were defined as the 90 AAL areas (Tzourio-Mazoyer et al., 2002), used for the FC analysis above. Second, connections between ROIs were estimated using probabilistic tractography. Third, data were averaged across participants to get one group-estimate. The resulting SC network was constructed for and reported in previously published work (e.g. Deco et al. (2016b)).

The FLIRT tool from the FSL toolbox (www.fmrib.ox.ac.uk/fsl, FMRIB, Oxford) was used to co-register the b0 image in diffusion MRI space to the T1-weighted structural image and then to the T1 template of ICBM152 in MNI space (Collins et al., 1994). The two resulting transformation matrices were concatenated and inverted, and applied to warp the AAL templates (Tzourio-Mazoyer et al., 2002) from MNI space to the diffusion MRI native space.

Processing of the diffusion MRI data was carried out in the FSL diffusion toolbox *Fdt*, with the default parameters. Following this pre-processing, the local probability distributions of voxel-wise fibre directions were estimated (Behrens et al., 2003). The *probtrackx* tool in *Fdt* was used to provide automatic estimation of crossing fibres within each voxel (Behrens et al., 2007).

The connectivity probability from a seed voxel v_i to another voxel v_j was defined by the proportion of fibres passing through voxel v_i that reached voxel v_j using a sampling of 5000 streamlines per voxel (Behrens et al., 2007). This was extended from the voxel level to the region level, i.e. in an AAL region

consisting of n voxels, i.e. 5000 x n fibres were sampled. The connectivity probability P_{IJ} from ROI I to ROI J was calculated as the number of sampled fibres in region I that connected the two regions divided by 5000 x n, where n was the number of voxels in region I.

For each ROI, the connectivity probability was estimated to each of the other 89 ROIs. Due to the dependence of tractography on the seeding location, the probability from I to J is not necessarily equivalent to that from J to I. Previous studies have, however, shown a high correlation between these two directionalities (Cabral et al., 2013). As directionality of connections cannot be determined based on diffusion-weighted MRI (see 1.5.5), the undirectional connectivity probability P_{IJ} between regions I and J was defined by taking the mean of these two connectivity probabilities. This undirectional connectivity was considered a measure of the structural connectivity between the two areas, with $C_{IJ} = C_{JI}$, and the probability as a proxy of connectivity strength.

Finally, the 16 participant-specific SC matrices were averaged to get one group-averaged SC matrix.

2.2.5 *Relationship between functional connectivity and structural connectivity*

The objective of the following analyses was to estimate the similarity between the FC matrices, computed for both the *Liege* and *Frankfurt* datasets, and the group-averaged SC matrix.

We used Pearson's linear correlation between the vectorised SC matrix and the vigilance-specific FC matrices. Prior to calculating the correlations, we corrected for differences in the link densities between the FC and SC matrices, as in Tagliazucchi and colleagues (2015). The link density of a network is defined as the ratio between the number of present connections and the number of possible connections. Since an FC matrix is made up of values of correlations between timecourses, and correlations are always non-zero, non-thresholded FC matrices will be fully connected with link densities of $d_{FC} = 1$. In contrast, an SC network is sparsely connected with some connections present and some absent. As such, the group-averaged SC analysed here, expressed a link density of $d_{SC} = 0.6245$. We used the function *threshold_proportional* from the Brain Connectivity Toolbox (Rubinov and Sporns, 2010) to threshold all FC matrices to the link density of the SC. Unlike Tagliazucchi et al. (2015) we did not perform binarisation, but rather kept the SC values and the FC values after thresholding. Binarisation would have set all present connections to unity, which would have removed information regarding relative connection values in both the SC and the FC matrices, which we regarded as relevant.

The similarity between a given link-density-thresholded functional connectivity matrix, FC', and the group-averaged structural connectome, SC, was quantified as:

$$R_{SC,FC'} = \frac{1}{N_{connections} - 1} \sum_{n}^{N_{connections}} \left(\frac{\overline{SC_n - \mu_{SC}}}{\sigma_{SC}} \right) \left(\frac{FC'_n - \mu_{FC'}}{\sigma_{FC'}} \right) \qquad \textbf{Eq. 2.3}$$

where $N_{connections}$ is the total number of connections in the connectivity matrices (i.e. 90^2–90), μ is the mean connectivity value, and σ is the standard deviation of the connectivity values.

The similarity score thus represented a single summary value for the comparison of the structural connectome and each of the vigilance-specific FC' matrices.

2.2.6 *Group-comparisons of vigilance-specific differences in SC-FC similarity and in mean FC level*

We compared the SC-FC similarity, as well as mean levels of FC, between wakefulness and SWS in the *Liege* dataset and between all vigilance conditions for the *Frankfurt* dataset. The mean FC level was calculated by the average across all values in the FC matrices with the diagonals set to zeros. To account for participant-specific differences we employed paired t-tests between each pair of vigilance conditions and non-parametric permutations to establish significance. A null-distribution was build for each pairwise comparison by randomly shifting the condition label of the SC-FC (or the mean FC) observations and recording the resulting t-stats. 1000 permutations were performed. The t-stat from the paired t-test, with original condition labels, was then compared to the null-distribution, and this yielded the probability of finding a significant difference by chance. A significance level of $p < 0.05$ was employed.

2.2.7 *Regional differences in SC-FC similarity*

The SC-FC similarity is a global measure representing the overall correspondence between patterns in the SC and FC matrices. To resolve potential vigilance-dependent differences at a more regional level a novel strategy was introduced. The idea was to estimate the ROI-specific rank-proximities between the FC matrices and the SC matrix. Specifically, we used the thresholded FC matrices, FC', and sorted the included values by their magnitude, such that the strongest weight was given the index $90^2 = 8{,}100$, while the weakest was given the index 1. We named the resulting indexed matrices FC''. The same was done for the SC matrix, yielding SC''. We then defined the rank-error ratio (RER) between the SC'' and each of the FC'' matrices as:

$$RER_{SC'',FC''} = 1 - \left(abs(SC'' - FC'') \oslash \left(N_{ROI}{}^2 - N_{ROI}{}^2 d_{SC}\right)\right) \qquad \textbf{Eq. 2.4}$$

where abs(.) denotes the absolute operator, $\oslash$ elementary division. N_{ROI} is the total number of connections in each matrix, i.e. 90^2, and d_{SC} is the link density of the SC matrix (see above 2.3.5). The term $(SC'' - FC'')$ quantifies the element-wise differences in rank positions between the SC matrix and a given FC matrix, while the term $\left(N_{ROI}{}^2 - N_{ROI}{}^2 d_{SC}\right)$ corresponds to the largest possible difference, since we were interested only in rank values above the link-density threshold. Accordingly we ignored elements of the resulting RER matrices corresponding to zeros in the FC' matrices as well as the SC matrix:

$$RER_{SC'',FC''}(x) = 0 \begin{cases} SC(x) = 0 \\ FC'(x) = 0 \end{cases} \qquad \textbf{Eq. 2.5}$$

An $RER_{SC'',FC''}$ matrix thus represented the element-wise error in rank between a given FC'' matrix and the SC'' matrix. RER matrices were calculated for wakefulness and SWS for the *Lige* as well as the *Frankfurt*

dataset, and as for the FC matrices (see above 2.3.3) differences in sub-networks were estimated using the paired t-tests combined with the NBS method.

2.2.8 *Dynamic investigation of SC-FC similarity*

As a final level of analysis we used the *Frankfurt* dataset to track the dynamic development of the global SC-FC similarity. We did so by using a conventional sliding-window approach (Allen et al., 2014, Hansen et al., 2015, Liégeois et al., 2015) (see Introduction, 1.5.8). Since the *Frankfurt* dataset consisted of continuous fMRI recordings of participants cycling through wakefulness and stages of NREM sleep, we could use this to explore how vigilance transitions modulated the SC-FC similarity at a more detailed temporal scale. As such, for each of the participants reaching SWS (N = 18) we computed dynamic FC (dFC) matrices in windows of 60 seconds with 58 seconds overlap (window length = 30 TRs, moved in steps of 1 TR). As in Liégeois et al. (2015), we calculated the SC-FC similarity for each window, along with the mean FC level, following the procedure for the static FCs, including link-density thresholding of the dFC matrices.

2.3 ***Results***

2.3.1 *Vigilance-specific differences in static FC*

Figure 2-1A and C show the group-averaged FC matrices for the vigilance conditions included in the datasets from *Liege* and *Frankfurt*, respectively. The FC matrices were organised by the AAL regions as summarised in Table 2-1 and illustrated in the Appendix (Figure S 1), with indices 1 to 45 corresponding to left hemispheric ROIs, and indices 46 to 90 to right hemispheric ROIs. Left (I_{left}) and right (I_{right}) homologues were organised such that $I_{right} = 90 - I_{left}+1$. This means that the second and fourth quadrants of an FC matrix include intra-hemispheric connections, while the first and the third include cross-hemispheric connections (see Figure 2-1E).

From these representations it is evident that the overall structure of these static FC matrices remained fairly preserved across vigilance conditions. The mean levels of FC changed however, as shown in Figure 2-1B and D. As such, SWS expressed significantly lower mean FC levels than wakefulness in both datasets, and compared to N1 and N2 in the *Frankfurt* data. In addition, N1 was found significantly increased compared to N2.

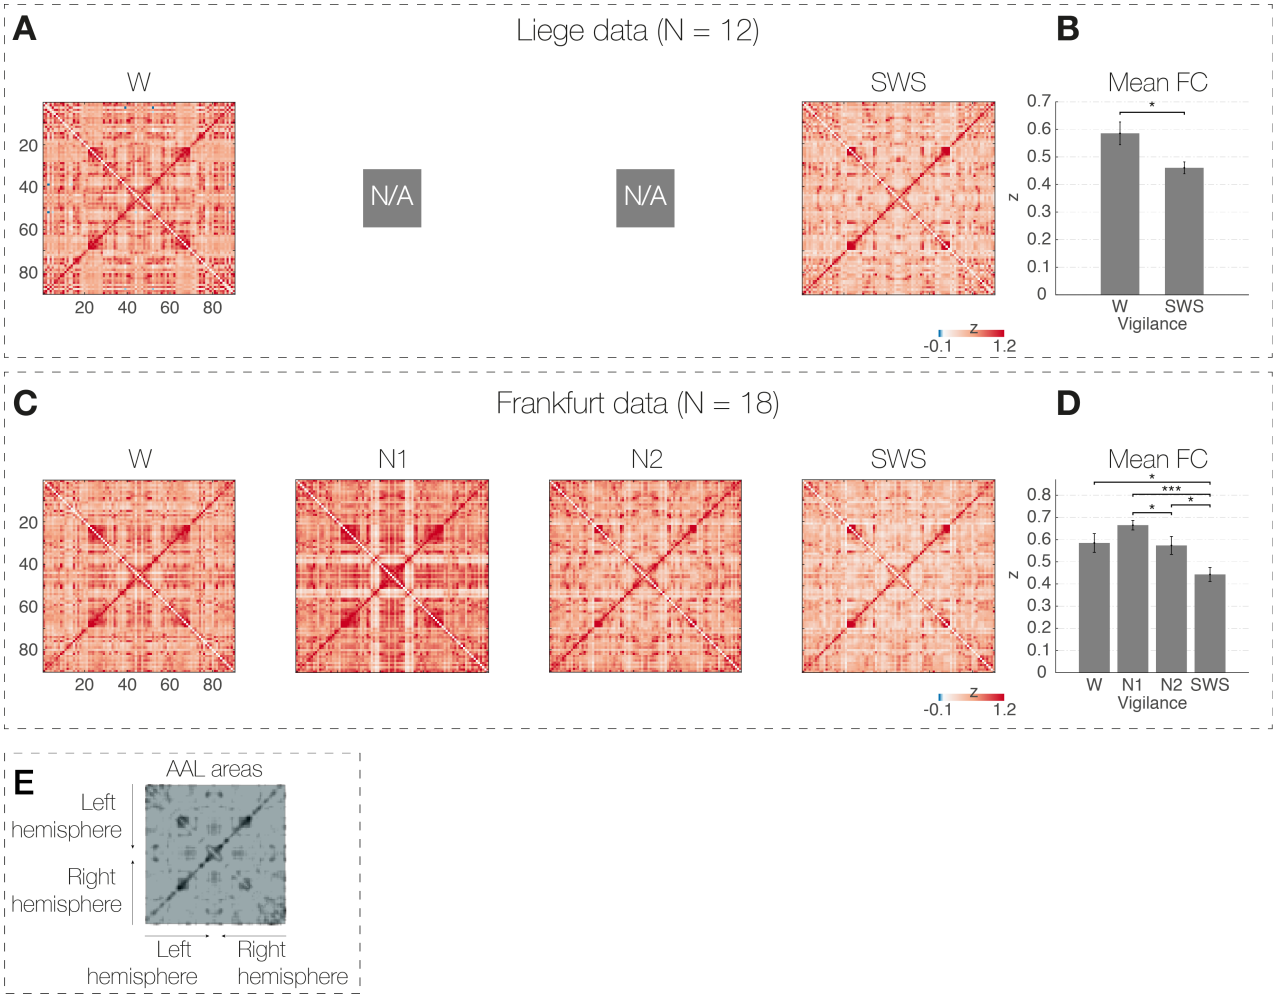

Figure 2-1 ***Group-averaged vigilance-specific FC matrices*** **A)** Group-averaged FC matrices for wakefulness (W) and slow-wave sleep (SWS) for the *Liege* dataset. N/A denotes that no data was available for other sleep stages. **C)** The same as **A** but for the *Frankfurt* dataset. **B)** and **D)** bar plots of the mean FC values for the included vigilance conditions. Cross-bars signal significant differences as estimated with non-parametric paired t-tests. *: $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$. **E)** schematic representation of the organisation of the AAL ROIs in the FC matrices.

AAL index	Hemishphere	Label	Hemishphere	AAL index
1	L	Precentral gyrus	R	90
2	L	Superior frontal gyrus	R	89
3	L	Superior frontal gyrus, orbital part	R	88
4	L	Middle frontal gyrus	R	87
5	L	Middle frontal gyrus, orbital part	R	86
6	L	Inferior frontal gyrus, opercular part	R	85
7	L	Inferior frontal gyrus, triangular part	R	84
8	L	Inferior frontal gyrus, orbital part	R	83
9	L	Rolandic operculum	R	82
10	L	Supplementary motor area	R	81
11	L	Olfactory Cortex	R	80
12	L	Superior frontal gyrus, medial	R	79
13	L	Superior frontal gyrus, medial orbital	R	78
14	L	Gyrus rectus	R	77
15	L	Insula	R	76
16	L	Anterior cingulate and paracingulate gyrus	R	75
17	L	Median cingulate and paracingulate gyrus	R	74
18	L	Posterior cingulate gyrus	R	73
19	L	Hippocampus	R	72
20	L	Parahippocampal gyrus	R	71
21	L	Amygdala	R	70
22	L	Calcarine fissure and surrounding cortex	R	69
23	L	Cuneus	R	68
24	L	Lingual gyrus	R	67
25	L	Superior occipital gyrus	R	66
26	L	Middle occipital gyrus	R	65
27	L	Inferior occipital gyrus	R	64
28	L	Fusiform gyrus	R	63
29	L	Postcentral gyrus	R	62
30	L	Superior parietal gyrus	R	61
31	L	Inferior parietal	R	60
32	L	Supramarginal gyrus	R	59
33	L	Angular gyrus	R	58
34	L	Precuneus	R	57
35	L	Paracentral lobule	R	56
36	L	Caudate	R	55
37	L	Putamen	R	54
38	L	Pallidum	R	53
39	L	Thalamus	R	52
40	L	Heschl gyrus	R	51
41	L	Superior temporal gyrus	R	50
42	L	Temporal pole: superior temporal gyrus	R	49
43	L	Middle temporal gyrus	R	48
44	L	Temporal pole: middle temporal gyrus	R	47
45	L	Inferior temporal gyrus	R	46

Table 2-1 ***Overview of indexing of SC and FC matrices based on the AAL parcellation*** (Tzourio-Mazoyer et al., 2002).

Local differences between the vigilance-specific FC matrices are shown in Figure 2-2. As explained in the Methods (see 2.2.3), two different correction schemes were employed for multiple comparisons; the connection-wise method, correcting at the single connection level, and the network-based statistic (NBS) correcting at the level of supra-threshold sub-networks. For the NBS a threshold was defined *a priori* on the t-stats, above which connected sub-networks were searched for. After testing a range of thresholds, we chose 4.0 as a good representative solution (see Zalesky et al. (2010) for a discussion on the choice of thresholds). As evident in Figure 2-2, the NBS produced more widespread differential networks than the connection-wise correction, indicating that the significance-threshold at the single connection level was higher than the one chosen for the NBS. This is to be expected, since the NBS assigns significance at the level of connected sub-networks and not single connections. The results for the connection-wise correction may be seen as a conservative reference for the NBS results. In line with this, please note that a majority of the significant networks resolved with the NBS overlapped with the connection-wise method.

As expected from the mean FC results above, SWS showed widespread decreases in connectivity relative to the other vigilance conditions. Compared to wakefulness, these decreases were particularly abundant in connections between temporal and occipital regions, and from those to regions in somatosensory and motor cortices. Relative to N1 sleep, significant decreases were found in SWS across the whole network with the noticeable exception of subcortical connections. Regional highlights of the comparison to N2 sleep included significant decreases during SWS in connections involving Heschl's gyrus, the paracentral lobule as well as the adjacent supplementary motor area, the precentral gyrus, and the anterior cingulate.

N1 showed significant increases compared to wakefulness in connections between medial orbitofrontal regions and the parietal as well as temporal cortices. The NBS identified a sub-network of significant decreases relative to wakefulness between subcortical regions, including the thalamus and the caudate, and the precuneus and visual cortices. Compared to N2 sleep, N1 expressed significant increases in a large network comprising the occipital, temporal and parietal cortices. Again, subcortical regions were involved in networks with decreased connectivity in N1, and in particular the thalamus showed stronger interregional connectivity during N2 sleep.

The comparison between wakefulness and N2 sleep yielded a sub-network of N2-specific increases with the anterior cingulate cortex as highly participating in connections to the temporal and somatosensory cortices. Wakefulness showed increases between temporal and occipital regions.

In sum, when wakefulness showed increased FC it was typically in connections between temporal and occipital regions. Cortico-cortical FC was commonly increased during N1 sleep, while subcortico-cortical connections were decreased. N2 sleep expressed strong FC in connections involving the somatosensory and motor system and the anterior cingulate cortices, while posterior parietal and occipital increases were generally absent. Finally, SWS was characterised by widespread decreases compared to all other conditions.

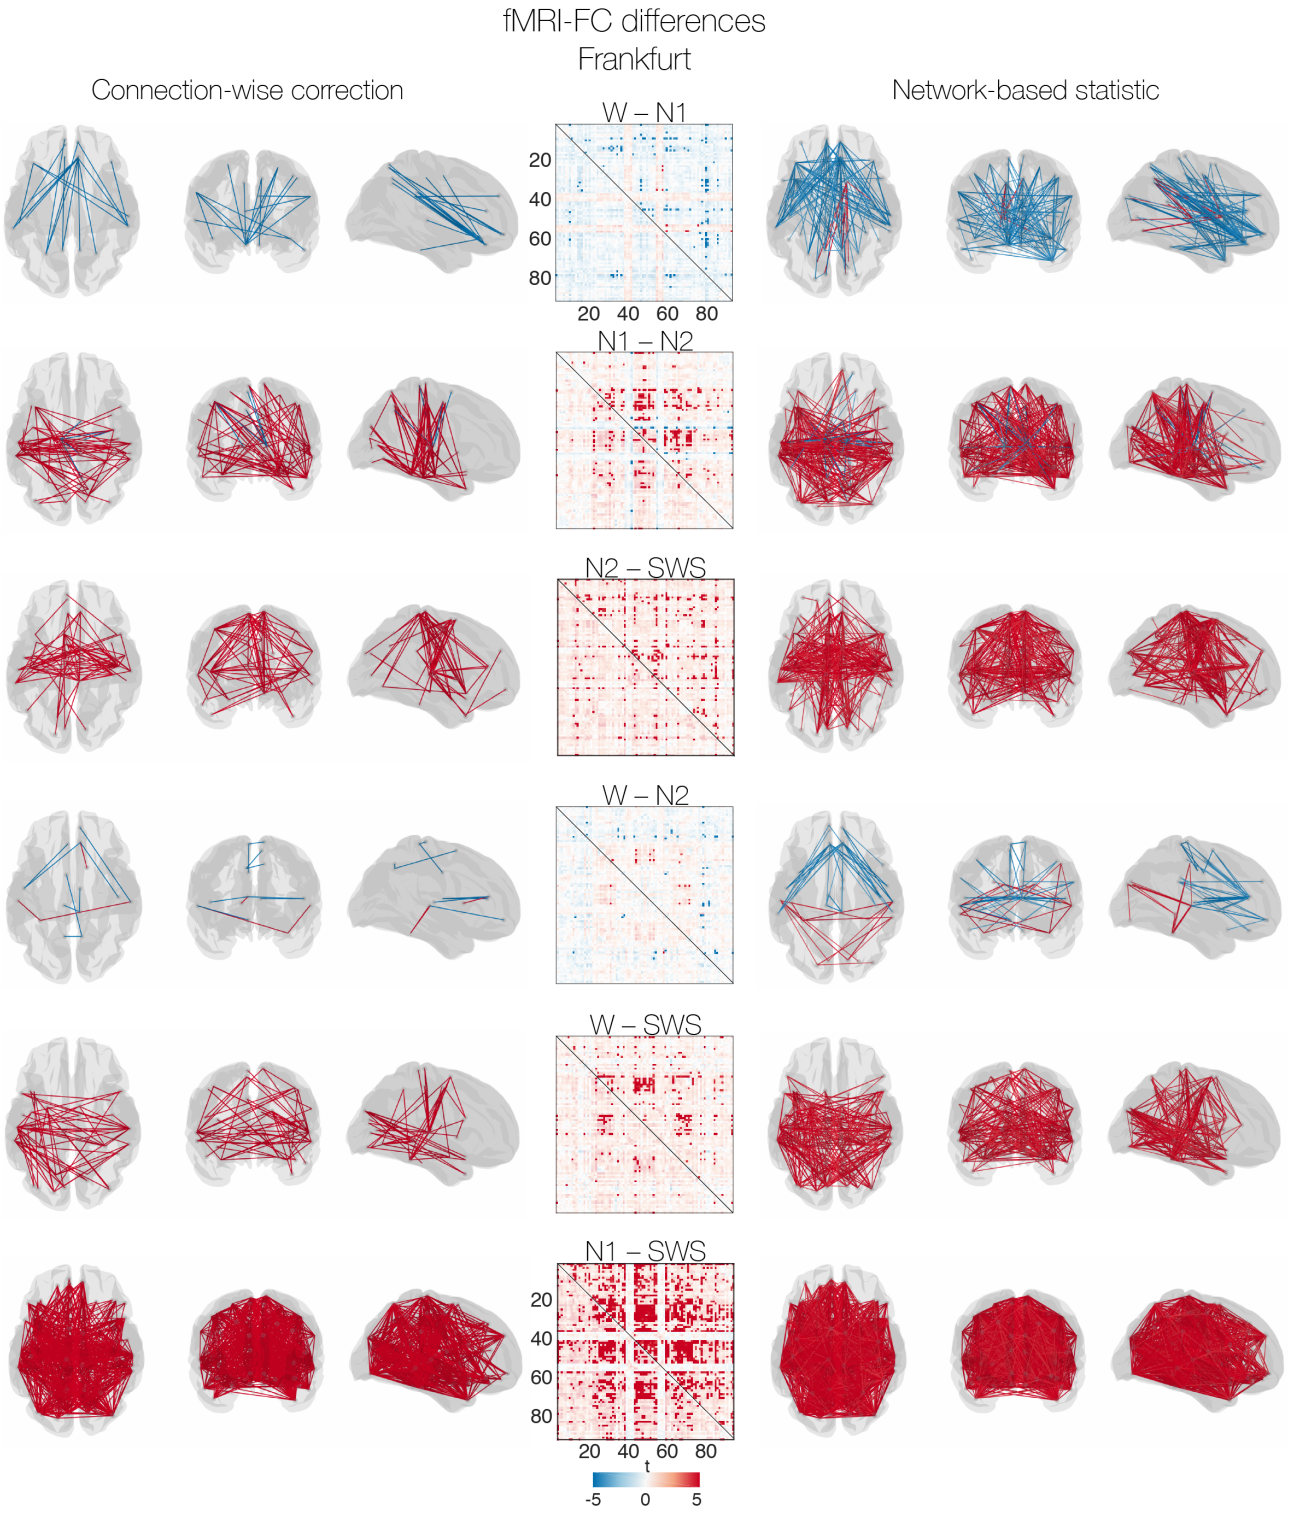

Figure 2-2 ***Comparative FC network analyses between vigilance conditions*** Paired t-tests were performed on the FC networks of 90 AAL areas derived from the 18 participants from the *Frankfurt* dataset that included all vigilance stages in their fMRI recordings. Multiple comparisons were corrected using two complementary methods; a connection-wise correction (left column) and a network-based statistic (right column). The results of each pairwise comparison are presented by a faded non-corrected t-stat matrix with corrected significance indicated by highlighted colours. Significant connections/sub-networks are overlaid on a glass-brain to the left/right. Red and blue correspond to increases and decreases, respectively, for the contrast given above the matrix. The sub-diagonal and the super-diagonal parts of a matrix show the results for the connection-wise correction and the network-based statistic, respectively. Note, for the connection-wise correction significance pertains to individual connections, while the network-based statistic considers the significance of sub-networks. The threshold chosen for the network-based statistic was $t > 4.0$.

2.3.2 *Vigilance-dependence of SC-FC similarity*

As presented in Figure 2-3 we observed a consistent increase in global SC-FC similarity during SWS relative to wakefulness. For the *Liege* dataset (Figure 2-3A) only 1 out of 12 participants showed the opposite relationship. An equivalent low proportion (3 out of 18) of participants did not show higher SC-FC similarity during SWS compared to wakefulness for the *Frankfurt* dataset (Figure 2-3C). In line with this, the non-parametric paired t-tests showed highly significant differences between these two vigilance conditions (Figure 2-3B and D). For the *Frankfurt* dataset we also tested the SC-FC similarity of N1 and N2 sleep. Similarly to wakefulness, values during N1 were found significantly lower than in SWS, and in addition significantly lower than N2.

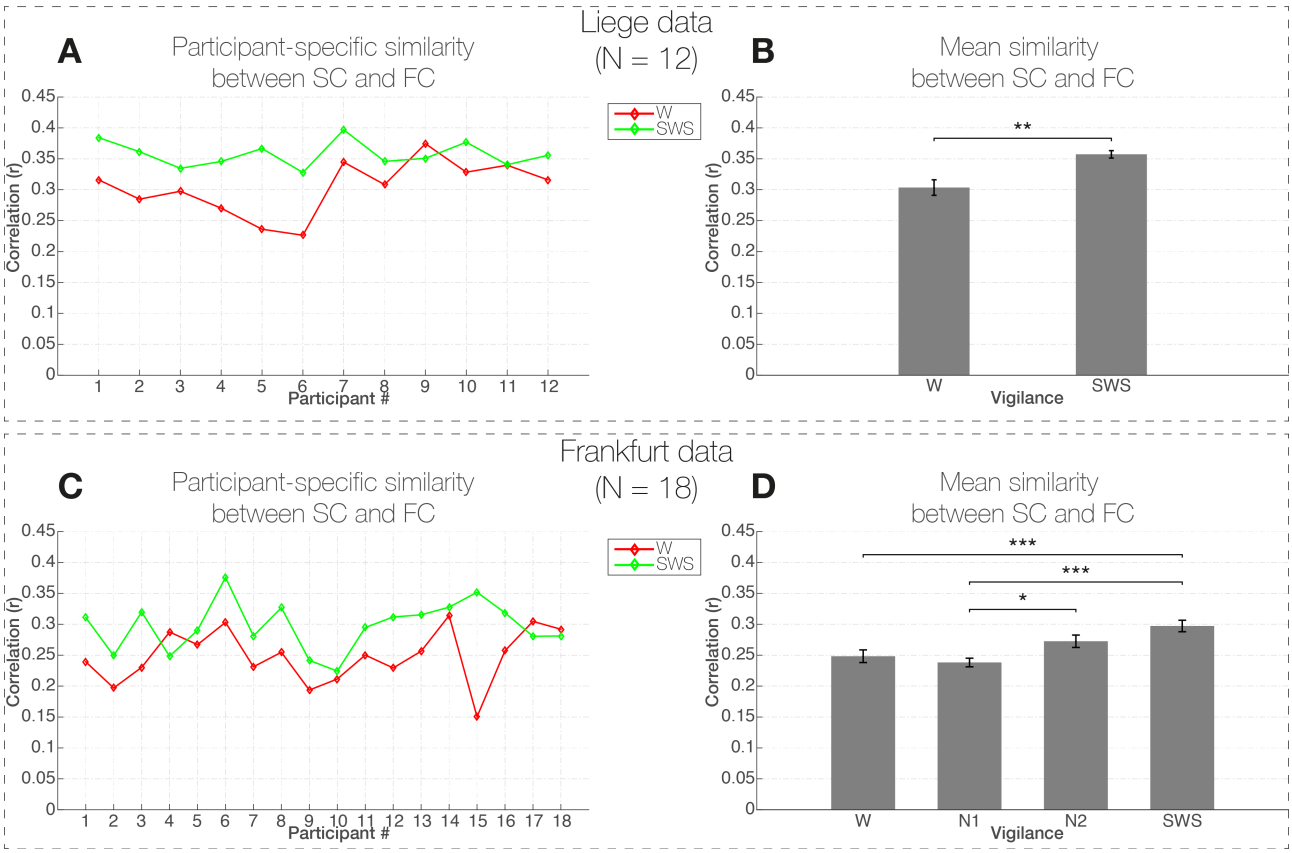

Figure 2-3 ***Global SC-FC similarities for two independent fMRI datasets*** **A)** and **C)** participant-specific SC-FC similarity, measured as Pearson's linear correlation (r) between the SC and link-density-thresholded FC matrices, for wakefulness (red) and SWS (green) for the *Liege* and *Frankfurt* dataset, respectively. Note the increase in SC-FC similarity from wakefulness to SWS across participants. The bar-plots in **B** and **D** represent the group-averages of SC-FC similarity across the available vigilance conditions for the *Liege* and *Frankfurt* dataset, respectively. Cross-bars signal significant differences as estimated with non-parametric paired t-tests. *: $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

We investigated the difference in SC-FC similarity between wakefulness and SWS in more regional detail for both the *Liege* and the *Frankfurt* dataset. To estimate the specific parts of the structural connectivity that the respective FC matrices followed, we computed rank-error ratio (RER) matrices (explained in Methods, 2.2.7) within wakefulness and SWS. The group-averaged RERs are shown in the sub-diagonal parts of the matrices in Figure 2-4A. An RER matrix may be seen as a mapping of each functional connection to its corresponding structural connection in terms of rank position. We went on to explore group differences with paired t-tests, correcting for multiple comparisons using the NBS. Figure 2-4B shows results for a t-stat

threshold of 2.8 for the NBS. The significant sub-networks identified by NBS are highlighted in the sub-diagonal and super-diagonal parts of the matrix for the *Liege* and the *Frankfurt* dataset, respectively.

From both the matrix and the brain representations in Figure 2-4B it is clear that consistent differences were found across the *Liege* and *Frankfurt* datasets. As such, during wakefulness functional connections were found to correspond more to structural connections between occipital and temporal regions, and between temporal and posterior parietal regions. On the other hand, in SWS functional connections expressed closer proximity to the structure in midline connections with a particular abundance between the posterior cingulate cortex and orbitofrontal regions.

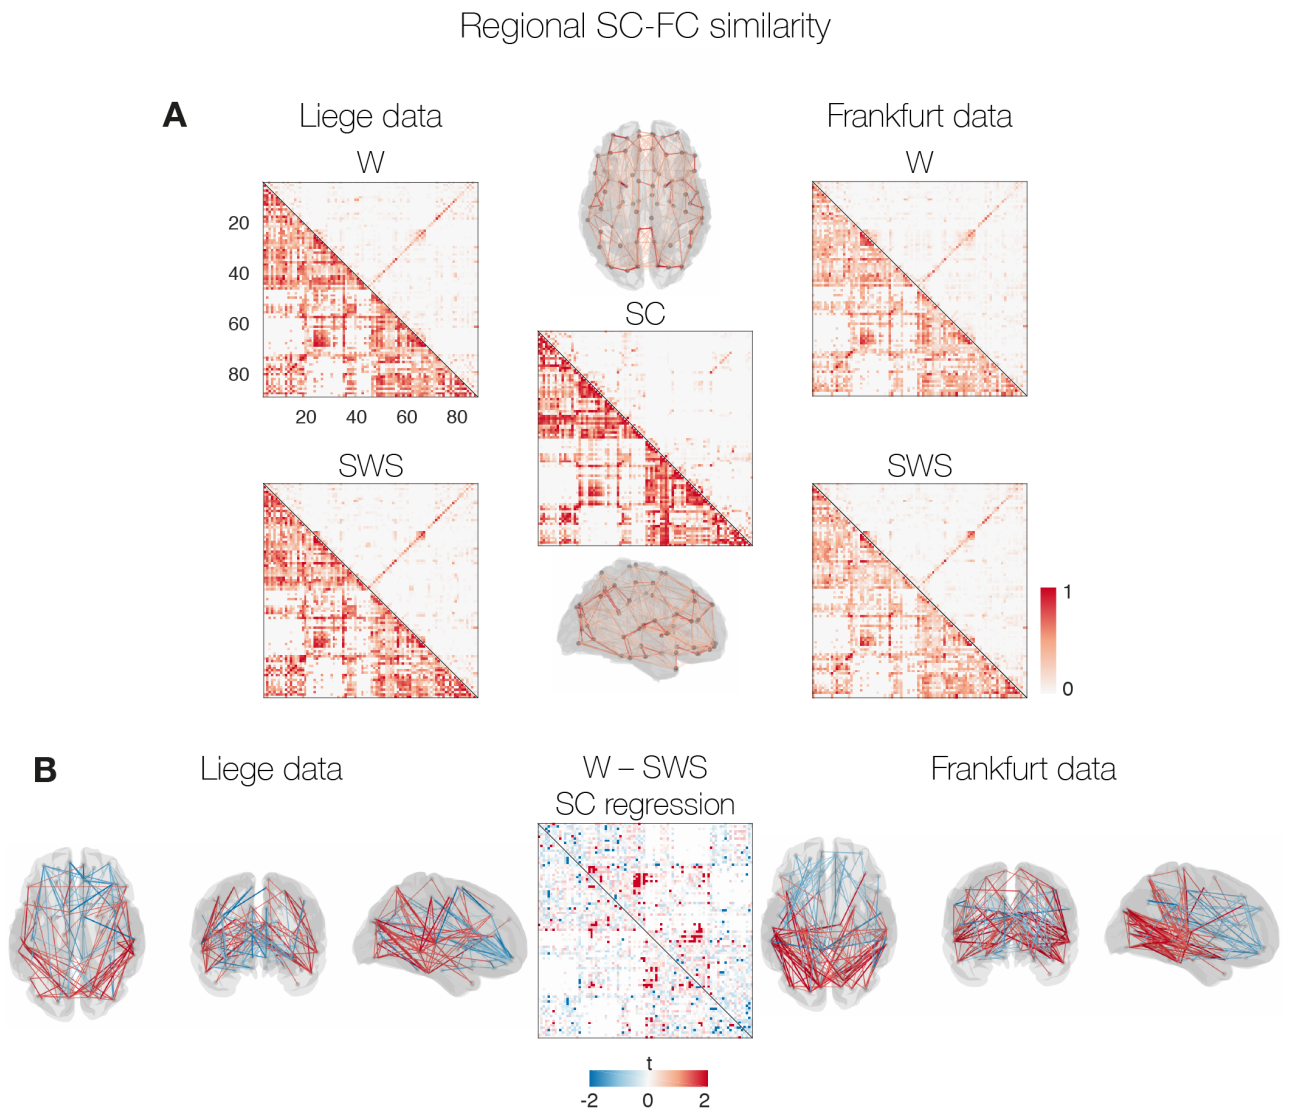

Figure 2-4 ***Regional determinants of SC-FC similarity for wakefulness and SWS*** **A)** The middle split matrix show the group-averaged structural connectome (SC). The super-diagonal part shows the original SC normalised by the value of its strongest connection, and is visualised on a glass brain above and below the matrix. The values of this original SC span several orders of magnitude, which makes it difficult to tell weaker connections. The sub-diagonal part shows the same structure but with values ranked according to strength, placing them on a common scale. The super-diagonal parts of the matrices for the *Liege* (left) and *Frankfurt* (right) datasets show the group-averaged FC networks for wakefulness and SWS. The FC networks have been link-density thresholded, and their values substituted by the original values from the SC to make them visually more comparable to the original SC matrix. The sub-diagonal parts show the group-averaged results of the rank-error ratio (RER), scaled from 0 to 1, where values close to 1 indicates that a functional connection had a ranked value close to the one of its corresponding structural connection. **B)** The RER matrices were compared between wakefulness and SWS using a paired t-test and NBS to correct for multiple comparisons. A t-stat threshold of 2.8 was used. The results are shown in the split t-stat matrix. The sub-diagonal and the super-diagonal parts show results for the *Liege* and the *Frankfurt* dataset respectively. The significantly different sub-networks are highlighted and visualised in glass-brains next to the matrix. Note the correspondence between the two datasets at this regional level. Notice also how FC during wakefulness followed the SC more closely in occipital, temporal and parietal connections, while SWS FC was more similar to SC in frontal regions.

2.3.3 *Dynamic tracking of SC-FC relationship*

To get a sense of the dynamic evolution of the SC-FC relationship across time, we computed dynamic FC (dFC) matrices by applying a sliding-window analysis to the recordings of the 18 participants that included wakefulness and SWS in the *Frankfurt* dataset. A window of length 30 TRs was moved in steps of 1 TR, and SC-FC similarity as well as mean FC level was computed for each window. From the static FC analyses it may be seen that the mean FC level and the SC-FC similarity followed opposite developments throughout the vigilance conditions (compare Figure 2-1D and Figure 2-3D). The dFC analysis should reveal whether this inversed relationship was also evident at shorter time scales.

The values of SC-dFC and mean FC from all time windows are plotted against each other in Figure 2-5A. As illustrated by the golden regression line, the two variables expressed a strong negative correlation. In Figure 2-5B this relationship is quantified at the participant level, illustrating that the negative correlation between SC-FC similarity and the mean FC level was consistent over time and participants. The group-average was r = –0.650 (± SE 0.035, N = 18). The full timecourses from an example recording (participant 5 with a value very close to the group average) are plotted in Figure 2-5C together with the sleep scoring obtained from the simultaneous polysomnography. From this representation it is obvious how the two variables fluctuated in opposing fashion.

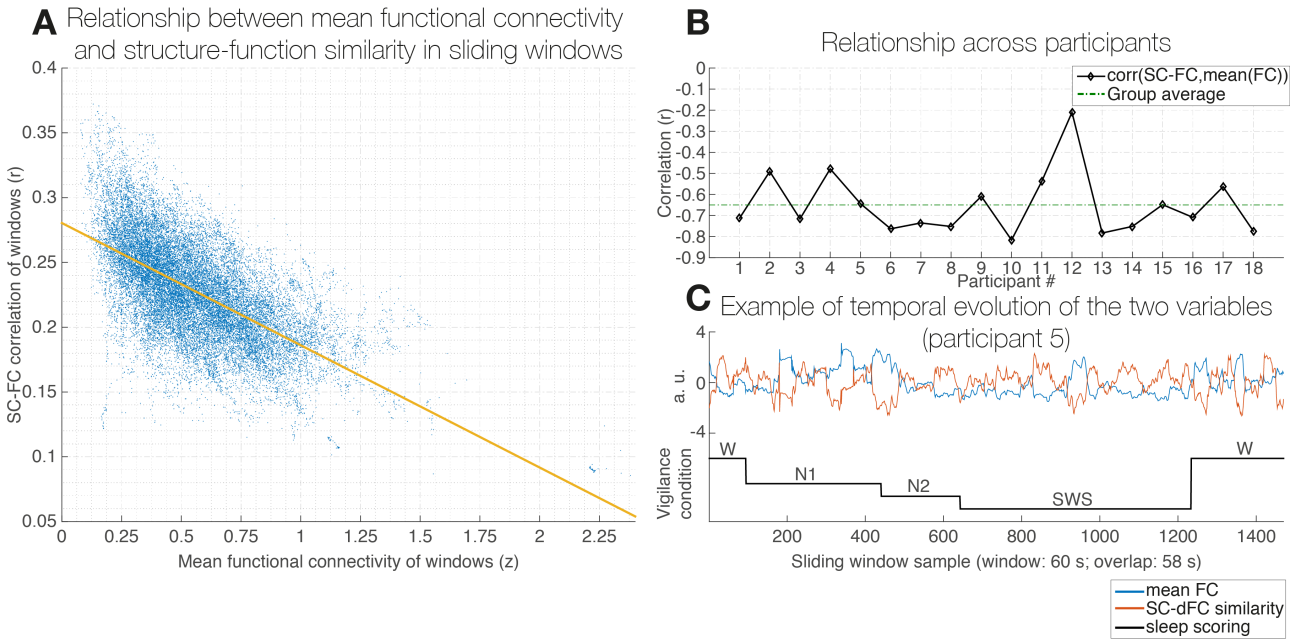

Figure 2-5 ***Dynamic tracking of SC-FC relationship and mean FC level*** **A)** A sliding-window approach (window length: 60 seconds, with 58 seconds overlap) was used to get a series of dynamic FC (dFC) matrices across the full recordings of 18 participants from the *Frankfurt* dataset. Similarity to SC as well as the mean FC level were calculated for each dFC matrix. The relationship between mean FC level and SC-FC similarity is presented as a scatter plot containing values from all participants and windows, and a significant negative correlation between these two variables is indicated by the golden regression line. **B)** Correlation between mean FC level and SC-dFC similarity, showing negative values consistently across the 18 participants. The green line marks the group-average. **C)** Participant 5 showed a relationship between mean FC and SC-dFC of a value very close to the group-average. The plot shows the timecourses of mean FC (blue) and SC-dFC (red). The two timecourses were demeaned and normalised by their standard deviations for visualisation purposes. Note, how the two variables fluctuate in opposite manners, expressing their negative dynamic correlation. The sleep scoring, from the simultaneously recorded EEG is shown in black below the two timecourses. When following the curve for the mean FC level it becomes clear that it fluctuates throughout the whole recording, which may be taken as a sign of dynamic functional connectivity. Specifically during SWS the mean FC level stabilises at a relatively low value, but shows intermittent periods of higher values.

2.4 ***Discussion***

Using data from two independent fMRI datasets of wakefulness and NREM sleep we have identified vigilance-dependent changes in whole-brain static FC networks in terms of global strengths and individual connections and sub-networks, while paying particular attention to the consistency across participants. We have shown that the global increase in SC-FC similarity during slow-wave sleep is robust across both cohorts. Moreover, using a dynamic FC analyses we have provided evidence of an overt negative relationship between SC-FC similarity and mean FC level.

2.4.1 *Comparative network analysis of static functional connectivity: frontal decreases in wakefulness, subcortico-cortical dropout in N1 sleep, and global suppression during slow-wave sleep*

From the whole-brain networks of static FC we found differences between wakefulness and individual stages of NREM sleep in global as well as regional connectivity strengths. We applied paired t-tests to the dataset, that was matched within participants in terms of time and vigilance conditions, and corrected for multiple comparisons using non-parametric statistics.

The mean FC level, obtained from averaging across all connections, was found highest during N1 sleep, and lowest during SWS, while wakefulness and N2 sleep took on intermediate values. This is largely in line with previous evidence (Spoormaker et al., 2010, Tagliazucchi and Laufs, 2014) (see Introduction, 1.5.4). We were unable to reproduce the global increase in FC during N2 sleep relative to wakefulness, reported in both Spoormaker et al. (2010) and Tagliazucchi and Laufs (2014). The methods used for computation of FC networks in these studies were very similar to ours; ROIs were defined using the AAL template and pairwise FC values were estimated in the static way. The latter used all 116 regions, while the former excluded the cerebellar regions, thus including the same 90 ROIs we used.

Our methods primarily diverged in the criteria for data inclusion, which aimed to guarantee that each participant was represented by the same amount of data from all vigilance conditions. In contrast, the two former studies analysed all artefact-free epochs, including data from participants with only a subset of the investigated vigilance conditions. While this allowed them to include data from participants and more samples for each vigilance condition, it also resulted in unbalanced dataset. Given that Tagliazucchi and Laufs (2014) used data from the same dataset as we did, which excludes cohort-driven differences, we argue that our results may be seen as representative of what can be statistically inferred when the participant-level of description is accounted for. The same may be said for our regional findings, discussed below.

From a functional perspective it is clear that global FC does not follow the level of vigilance in a linear way. Higher FC does not separate wakefulness from sleep in general, but only from SWS. This is corroborated by the findings within RSNs where FC values can be maintained or even increased during N1/N2 sleep (Horovitz et al., 2009, Larson-Prior et al., 2009). Assuming that a higher FC indicates a higher degree of information processing, it would seem that this is not a prerequisite of wakefulness. Graph-theoretical studies of whole-brain FC networks have pointed to topological differences in trying to explain why FC may be high during states of decreased vigilance. The most consistent finding from these studies is that information processing becomes more segregated during sleep, as quantified through the graph-

theoretical measure modularity (Boly et al., 2012b, Spoormaker et al., 2012, Tagliazucchi et al., 2013b). However this appears to mainly apply to N2 sleep and SWS, as Tagliazucchi and colleagues were unable to find differences in global modularity between N1 and wakefulness (Tagliazucchi et al., 2013b).

Still, from the classifier studies we know that whole-brain FC networks do change in very specific ways with vigilance (Tagliazucchi and Laufs, 2014, Altmann et al., 2016) (see Introduction, 1.5.4). It seems likely that these classifiers gain their sensitivity from regional FC information.

Accordingly, and in line with the previous reports (Spoormaker et al., 2010, Tagliazucchi and Laufs, 2014) (see Figure 1-10 in Introduction) our comparative network analyses identified an abundance of regional differences in FC strengths between wakefulness and the three NREM sleep stages. We were able to reproduce many aspects of the differential maps reported in Spoormaker et al. (2012) and Tagliazucchi and Laufs (2014), albeit often not to the same extent. Again, this may be attributed to our more conservatively designed datasets, and perhaps more importantly to our more rigorous, non-parametric corrections for multiple comparisons. While Spoormaker and colleagues (2012) ignored corrections altogether, Tagliazucchi and colleagues applied a parametric Bonferroni threshold (Tagliazucchi and Laufs, 2014).

Discrepancies included a less widespread increase in cortico-cortical FC for N1 compared to wakefulness, and much reduced differences between wakefulness and N2. Instead our analyses of these contrasts yielded more specific differential networks. Wakefulness was thus significantly decreased compared to N1 in fronto-temporal, fronto-parietal, and fronto-occipital FC, and compared to N2 in connections between the anterior cingulate cortex and superior temporal and somato-motor regions. In contrast, wakefulness was found increased relative to N2 in a network involving temporal and occipital areas.

Interestingly, the latest of the classifier studies by Altmann and colleagues imposed linearity upon their employed support vector machine, which allowed the authors to do *post hoc* analyses and identify the connections that represented the most relevant features for the discrimination between wakefulness and individual sleep stages (Altmann et al., 2016). Considering the resulting brain plots (provided in the supplementary material in Altmann et al. (2016)) it appears that in many cases they show more resemblance to the differential networks from our analyses than those from Spoormaker et al. (2010) and Tagliazucchi and Laufs (2014). This might suggest that our more careful statistical inferences increased the specificity towards the most discriminatory features of the whole-brain FC network.

Our analyses identified the N1-specific decrease in subcortico-cortical FC reported previously (Spoormaker et al., 2012, Tagliazucchi and Laufs, 2014). Compared to wakefulness, primarily the caudate, but also the thalamus, was involved in a sub-network with the occipital cortex of decreased FC in N1. Compared to N2, N1 showed examples of significant decreases in FC between the thalamus and regions in all cortical lobes. While cortico-cortical FC was widely increased in N1 compared to SWS the subcortical areas were largely spared. N1 represents a particularly interesting stage, given its placement between wakefulness and sleep. Meanwhile, RSN and graph-theoretical studies have collectively struggled to identify meaningful changes in whole-brain activity from wakefulness to N1 sleep (see Introduction, 1.5.3 and 1.5.4). The drop in subcortico-cortical FC revealed by our and previous whole-brain comparative FC analyses (Spoormaker et al., 2012, Tagliazucchi and Laufs, 2014) appear as one of the more specific signatures. The

thalamic component of this FC phenomenon is paralleled by intracortical evidence (Magnin et al., 2010), where a decrease in dynamic complexity was observed in the thalamus prior to in the cortex during sleep onset. It could be speculated that the thalamocortical disconnection is reflecting a decreased thalamic modulation of the cortex, which in turn could lead to the marked increase in global FC that we observe, and that might be incompatible with high responsiveness and wakefulness.

It is important to note that the thalamus is a diverse structure that can be divided by its anatomical and functional properties, which the AAL parcellation does not account for. Illustrating this point, Picchioni and colleagues divided the thalamus by a specific-relay (the lateral geniculate) and a non-specific nucleus (the centromedian), when comparing thalamocortical FC between wakefulness and NREM sleep (Picchioni et al., 2014). They found that in particular the non-specific nucleus, believed to have a more heterogeneous pattern of projections to the cortex than the sensory relay nuclei, showed higher connectivity to cortical areas in wakefulness than during N2 and SWS. Unfortunately the study did not report values for N1 sleep (Picchioni et al., 2014). Hale and colleagues used an even finer, functionally based, parcellation of the thalamus and also included N1 sleep, however the use of global signal regression make their FC results difficult to compare with ours (Hale et al., 2016).

For a more functional interpretation of the non-thalamic regional differences between wakefulness and N1 it may be instructive to consult the literature on neural correlates of insomnia. We found regions of the frontal cortex, mainly in the orbitofrontal part, to be involved in a range of increased connections to parietal, temporal as well as a few occipital regions during N1. In contrast, the caudate was the most engaged region of the sub-network where FC was increased in wakefulness relative to N1. The orbitofrontal cortex and the caudate have been reported to undergo structural (Altena et al., 2010, Stoffers et al., 2012, Joo et al., 2013) and functional (Stoffers et al., 2014) alterations, respectively, in patients suffering from insomnia. Although insomnia is a multifaceted disorder that does not reduce to any single symptom, issues related to the initiation of sleep does represent a common phenomenon and a central diagnostic criterion (Riemann et al., 2015). It is therefore interesting to find an overlap between areas that have been implicated in this disorder and areas that change their FC from wakefulness to N1 sleep in our results.

For N2 sleep we found evidence for increased FC in network patterns bearing high resemblance to activation maps from fMRI studies of sleep spindles and K-complexes (see Figure 1-5 in Introduction, but also Andrade et al. (2011)). As such, connections to the anterior cingulate, the superior temporal, motor, and sensory cortices as well as between these were abundantly increased. Given the established role of the thalamus in generating both of these N2-specfic oscillations (see Introduction, 1.3.1), it appears sensible that thalamocortical connections were stronger in N2 compared to N1. While the indirect relationship between the fMRI BOLD signal and electrophysiology makes it difficult to confirm, our findings indicate that FC patterns during N2 may be driven by transient phasic events. The static FC approach may thus be less sensitive to potential background activity unrelated to such transients, which emphasise a caveat of static FC modelling of spontaneous activity.

2.4.2 *Functional connectivity gets 'locked' in a frontoparietal network of the structural connectome during slow-wave sleep*

We found FC to be generally decreased during SWS, at the global as well as at the regional level. Previous reports (Spoormaker et al., 2010, Spoormaker et al., 2011, Tagliazucchi and Laufs, 2014) have linked this observation to evidence from combined TMS and EEG studies, showing decreased propagation across the cortex of TMS-induced perturbations during this sleep stage (Massimini et al., 2005, Sarasso et al., 2015). The theoretical proposition is that decreased information processing *per se* underlies the decreased level of consciousness in SWS. However, as noted in the Introduction (see 1.5.7), the parallel between fMRI-derived FC and TMS/EEG propagations does not hold for N1 sleep, which challenges this one-to-one interpretation. Furthermore, while FC is decreased in SWS it is not abolished, as illustrated by the fact that all major RSNs have been identified in this stage (see Introduction, 1.5.3). One way of characterising this remaining FC is by comparing it to the structural connectome.

In line with a previous report from human sleep (Tagliazucchi et al., 2015) and one from general anaesthesia in monkeys (Barttfeld et al., 2015), we found that global FC becomes more similar to a group-averaged SC during SWS. We showed that this is a very robust phenomenon across participants in the same dataset used in Tagliazucchi et al. (2015), and validated this in a second independent dataset. At a functional level, this offers some interesting perspectives. It suggests that FC may arise in SWS from a rather passive exploration of the underlying structural architecture, which may also explain why RSNs can be observed in this state (see Introduction, 1.5.5). As pointed to in the previous studies (Barttfeld et al., 2015, Tagliazucchi et al., 2015), the lower SC-FC similarity during wakefulness support the assumption that a larger repertoire of states are required for complex conscious operations during wakefulness, following a corollary of the Integrated Information Theory (Tononi, 2008, Marshall et al., 2016).

We further tested the consistent increase in SC-FC similarity during SWS at the regional level of sub-networks. In both of the independent datasets we found that the global increase was driven by particularly close resemblance between SC and FC in connections to the frontal cortices. As such, midline fronto-parietal and fronto-temporal functional connections were closer to their corresponding structural connections in SWS compared to wakefulness. In contrast wakefulness FC showed closer similarity to SC in occipital and temporal connections. Tagliazucchi and colleagues also investigated regional differences, but with a different approach based on graph-theoretical modularity detection (Tagliazucchi et al., 2015). Their findings seem in conflict with our results by showing a closer relationship between SC and FC in higher-order connections for SWS and primary sensory regions for wakefulness. Importantly, the analyses in Tagliazucchi et al. (2015) were based on binarised matrices, which implicated that relative FC and SC strengths were discounted. Our rank-error ratio measure explicitly aimed to model these strengths. If indeed a deviation of FC patterns from SC represents a higher exploration of states, and accepting this as a sign of complex computations required for conscious processing, we would argue that our results make more sense. As such, by our accounts wakefulness is mainly confined to the underlying structure in "low-level" sensory regions, while FC appears to get 'locked' in a fronto-parietal and fronto-temporal pattern during SWS, normally associated with higher-order functions.

It is important to note that SC-FC similarity did not follow vigilance linearly. In fact, N1 expressed even lower average values of this measure than wakefulness (see Figure 2-3D), albeit the two were not significantly different. Based on this it would be inaccurate to propose global SC-FC similarity as a direct marker of consciousness as such. Interestingly, however, a turn to electrophysiological evidence offers a different interpretation, unrelated to consciousness, of the increased SC-FC similarity during SWS. The electrophysiological hallmark of SWS is the slow oscillation (see Introduction, 1.3.1). Interestingly, EEG studies have suggested that slow oscillations propagate across the cortex in patterns resembling some of the largest white-matter connections of the brain (Massimini et al., 2004, Murphy et al., 2009b). Specifically, this propagation was described as occurring in an antero-posterior direction along the superior longitudinal fascicle (Kurth and Huber, 2012), fitting well with our regional analyses for SWS. A prominent hypothesis for SWS is that it serves specific roles in brain plasticity (Tononi and Cirelli, 2013), through the slow-wave activity. Although this is mainly based on evidence at the level of synapses (Tononi and Cirelli, 2013), dendrites (Yang et al., 2014) and gray-matter regions (Riedner et al., 2007), studies of cortical maturation in development point to a potential role of slow-wave activity for plasticity of white-matter at the macro-scale (Kurth and Huber, 2012). Assuming that whole-brain activity during SWS exerts a greater plastic effect upon the structural connectome offer a different explanation for the observed increase in SC-FC similarity during SWS.

2.4.3 *Increases in dynamic functional connectivity lead to departure from the structural connectome blueprint*

Noticing an opposite effect of vigilance on the global measures of mean FC and SC-FC similarity from our static analyses, we tested how these two variables correlated over time. We found a very robust and strong negative correlation, such that windows of high mean FC were characterised by low similarities between global patterns of SC and FC. A previous study, employing very similar methods, found that SC-FC similarity does indeed fluctuate dynamically within resting-state fMRI recordings (Liégeois et al., 2015). The authors proposed this to be empirical evidence of how the structural connectome represent a principled network pattern aiding the transition between various FC configurations of whole-brain activity, in line with theoretical evidence from whole-brain computational modelling (Deco et al., 2011, Cabral et al., 2014a) (see Introduction, 1.5.6). Owing to the insight provided by static FC analyses of not only wakefulness but also NREM sleep, we were able to show: 1) that such fluctuations persist during sleep, and 2) that the structural connectome mainly exerts its constraints when the general level of FC is low. Assuming that the global strength of FC reflects a higher functional engagement and information processing, the latter observation suggests that patterns included in the structural connectome are insufficient for the specific operations carried out by the brain.

2.4.4 *Methodological considerations and limitations*

The lack of available data from REM sleep hampers the interpretation of the results presented in this chapter. The markedly different range of mental content commonly associated with REM would have been valuable

in the attempt to assign functional meaning to the FC measures, their vigilance-dependence, and their relationship to SC. As noted in the Introduction (see 1.5.3) sleep restriction is a commonly used experimental means to facilitate sleep in participants placed in the noisy and uncomfortable environment of an MR scanner, and was also included for the datasets analysed here. Sleep restriction has a tendency to augment the natural predominance of NREM sleep in the beginning of a sleep bout, and thus effectively restrict the amount of REM. An alternative approach would be to selectively deprive participants REM sleep, which would, at least theoretically, increase the proportion of REM in subsequent recordings, although this is quite an extensive and laborious intervention.

Further, it is important to note that sleep deprivation does not reflect regular wakefulness (van Someren et al., 2015). Recently, fMRI studies have addressed the effect of sleep-deprivation on spontaneous whole-brain activity, reporting robust alterations of FC networks, at the level of RSNs (Gujar et al., 2010, Sämann et al., 2010, Havas et al., 2012, Bosch et al., 2013, Kaufmann et al., 2015) and graph-theoretical measures (Liu et al., 2014). Future studies should focus on the effect of sleep deprivation on the FC differences reported here.

Despite the well-established dynamic nature of normal sleep architecture, most clearly illustrated by the decrease in power of slow-wave activity from early to late sleep cycles, most existing datasets, including the ones analysed here, are limited to the first sleep cycle. It is highly conceivable that whole-brain FC networks vary within the defined sleep stages over full-night recordings. Characterising such changes could yield a whole-brain account of sleep homeostasis and provide valuable functional information of this key concept of sleep and its regulation (see Introduction, 1.3.3).

Our analyses yielded some qualitatively different results from previous studies (Spoormaker et al., 2010, Tagliazucchi and Laufs, 2014). We have already highlighted how our explicit modelling of the participant-level as well as our more careful corrections for multiple comparisons most likely account for these differences. When creating balanced datasets we made sure to match the number of samples within participants, such that each participant was represented by equal number of time samples from each vigilance condition. The number of samples included for each participant was determined by the vigilance condition with the shortest contiguous epoch, and data from vigilance conditions including more than this minimum were taken from the centre of these epochs. However, our sliding-window analyses indicated a fluctuating FC across all vigilance conditions. It is therefore conceivable that different results would have been obtained had we chosen different segments of the vigilance epochs. This warrants a more dynamic FC investigation.

The negative relationship between mean FC level and SC-FC similarity, which we have identified across participants, should be sought validated in other datasets. The sliding-window approach is associated with challenges regarding statistical inference (Hindriks et al., 2015) and a somewhat arbitrary choice of window length (Preti et al., 2016), and may not be the optimal way of resolving FC dynamics. Further, it may be criticised that the group-averaged structural connectome used for SC-FC similarity estimations was derived from DTI data independent from the fMRI recordings. A similar approach was taken by Tagliazucchi et al. (2015), when investigating the same question, and structural connectivity matrices are known to show very high consistency across a healthy population. Thus for the current dataset of structural connectomes, we quantified the consistency across participants by correlating the 16 matrices in an all-to-all manner, and

found a very high average correlation of r = 0.947, with a minimum value of r = 0.909. As such, we would argue that the current group-average was a good representation of the healthy structural connectome. Moreover, our use of weighted, as opposed to binarised, SC values in the computation of global as well as regional SC-FC similarities requires a consideration of the interpretation issues related to deriving structural connectivity strengths from probabilistic tractography (see Introduction, 1.5.5 and Figure 1-11).

Finally, we used the AAL parcellation for both SC and FC computations. As mentioned earlier, this represents merely one possible large-scale partitioning of the brain (see Introduction, 1.5.4), and testing the methods and findings presented in this chapter would be another important level of validation.

Chapter 3
Frequency-resolved functional connectivity of sleep using MEG

3.1 *Introduction*

While still outnumbered by studies of wakefulness and the resting-state, investigations of functional connectivity (FC) networks have provided new insights regarding the multifaceted nature of whole-brain activity during sleep (see Introduction, 1.5). Some examples hereof were given in Chapter 2, where fMRI recordings of sleep were used to derive whole-brain FC networks. fMRI is generally the modality-of-choice for studying large-scale functional connectivity. Meanwhile, the interpretation of fMRI results is obscured by the indirect relationship between the BOLD signal and neural activity (see Introduction, 1.5.7).

Recent refinement of tools for source-reconstructing electrophysiological signals picked up at sensor arrays outside the head, offering solutions to the ill-posed inverse problem, have enriched the use of MEG. The source-reconstruction of neuronal oscillations measured with MEG, offers a unique opportunity to probe large-scale brain activity directly. Notably, MEG has been used to corroborate the fMRI evidence of large-scale patterns of functional connectivity in spontaneous brain activity during wakeful rest. As such, resting-state networks (RSNs) have been identified as co-variations of power fluctuations in specific frequency bands (Pasquale et al., 2010, Brookes et al., 2011b, Hipp et al., 2012, Luckhoo et al., 2012) (see Introduction, 1.5.7).

Linking large-scale network patterns to specific frequency bands not only suggests their neural origin but also facilitates more detailed interpretations of them. Much in parallel to the traditional characterisation of brain activity during sleep (see Introduction, 1.2.1, 1.3.1), evidence from spatially non-specific EEG measurements have been used to allocate functional aspects of wakefulness to specific frequency bands (Buzsaki, 2006, Engel and Fries, 2010, Wang, 2010, Palva and Palva, 2012, Lopes da Silva, 2013). The findings from MEG source-reconstruction of RSN-like structures in the alpha and beta bands may thus be seen in light of the assumed roles of these oscillations in complex cognitive phenomena like attention, memory, and perception. The study of spontaneous brain activity is exploratory in nature, and a multi-modal integration between fMRI and electrophysiology is principal to its success (Schoffelen and Gross, 2009, Schölvinck et al., 2013). MEG represents a promising mediator in this respect.

The majority of resting-state evidence from MEG studies have pointed to co-variations of power envelopes in the alpha and beta bands as the main contributor of large-scale FC, showing spatial correspondence to fMRI FC networks (Hillebrand et al., 2012, Hall et al., 2014, Tewarie et al., 2014, Colclough et al., 2015, Hipp and Siegel, 2015). This has promoted the idea of such brain-wide power modulations in alpha and beta as substrates for long-distance integration of more segregated and spatially confined information processing in higher frequencies like gamma (Von Stein and Sarnthein, 2000, Siegel et al., 2012). This view is generally in line with the overall composition of the EEG spectrum during wakefulness, where alpha and beta oscillations are prominent. In sleep, however, and NREM sleep in particular, these higher frequencies are diminished compared to wake (see Introduction, 1.2.1). Yet, fMRI

reports have consistently shown a rich presence of FC across the brain during all stages of sleep (see Introduction, 1.5.4 and Chapter 2). The question is what the electrophysiological basis might be for the preserved fMRI-derived FC patterns during sleep. To date, this remains scarcely explored (see Introduction, 1.5.7), and there have been no attempts at expanding the MEG and source-reconstruction-based FC framework to sleep. Besides linking fMRI results to underlying electrophysiology, such evidence could provide a new large-scale perspective on the established spectral characterisation of sleep.

In this chapter we report on power-amplitude couplings between 39 cortical regions, applying beamforming on whole-head MEG recordings of 8 healthy participants undergoing changes in vigilance from wakefulness to NREM sleep. A recently validated pipeline was employed to correct the source-leakage bias prior to the derivation of FC measures. We test the vigilance-specificity of the resulting FC patterns derived from 6 canonical frequency bands, spanning the range from 0.25 to 48 Hz, for the three vigilance conditions: wakefulness, N2 sleep and SWS. Performing comparative network analyses we compare MEG FC strengths between vigilance conditions.

3.2 ***Methods***

3.2.1 *Pre-processing of MEG sleep data*

Whole-head MEG recordings from 11 healthy participants were sleep-scored according to the principles of Rechtschaffen and Kales (Rechtschaffen and Kales, 1968) (for a full description of acquisition and sleep scoring, see Appendix, MEG dataset).

All pre-processing steps were performed in MATLAB (The MathWorks, Inc.) employing in-house built software, borrowing functions from other software packages, such as SPM8 (Litvak et al., 2011), FieldTrip (Oostenveld et al., 2011), and FSL (Jenkinson et al., 2012).

Each sleep recording (~40 minutes) was preprocessed as one continuous piece. The raw sensor data were high-pass filtered at 0.1 Hz, and channels with excessive noise levels were removed from further analyses, based on visual inspection.

Data were then submitted to Blind Source Separation by temporal Independent Component Analysis (ICA) using the FastICA algorithm (Hyvarinen and Oja, 2000, Vigario et al., 2000) as implemented in the FastICA toolbox for MATLAB (http://research.ics.aalto.fi/ica/fastica). Each independent component (IC) was represented by an independent timecourse and a topographic map describing each of the MEG sensors' contribution to the given IC. All ICs were reviewed manually and categorised as either "neural" or "artefact". Artefactual ICs included cardiac, ocular (Vigario et al., 2000, Mantini et al., 2008) or facial muscle-interference (Muthukumaraswamy, 2013b) and were removed from the full-rank data (Mantini et al., 2011).

Following ICA de-noising the continuous data segments were epoched according to vigilance conditions as given by the sleep scoring. Potentially remaining artefacts were removed from these epochs after visual inspection.

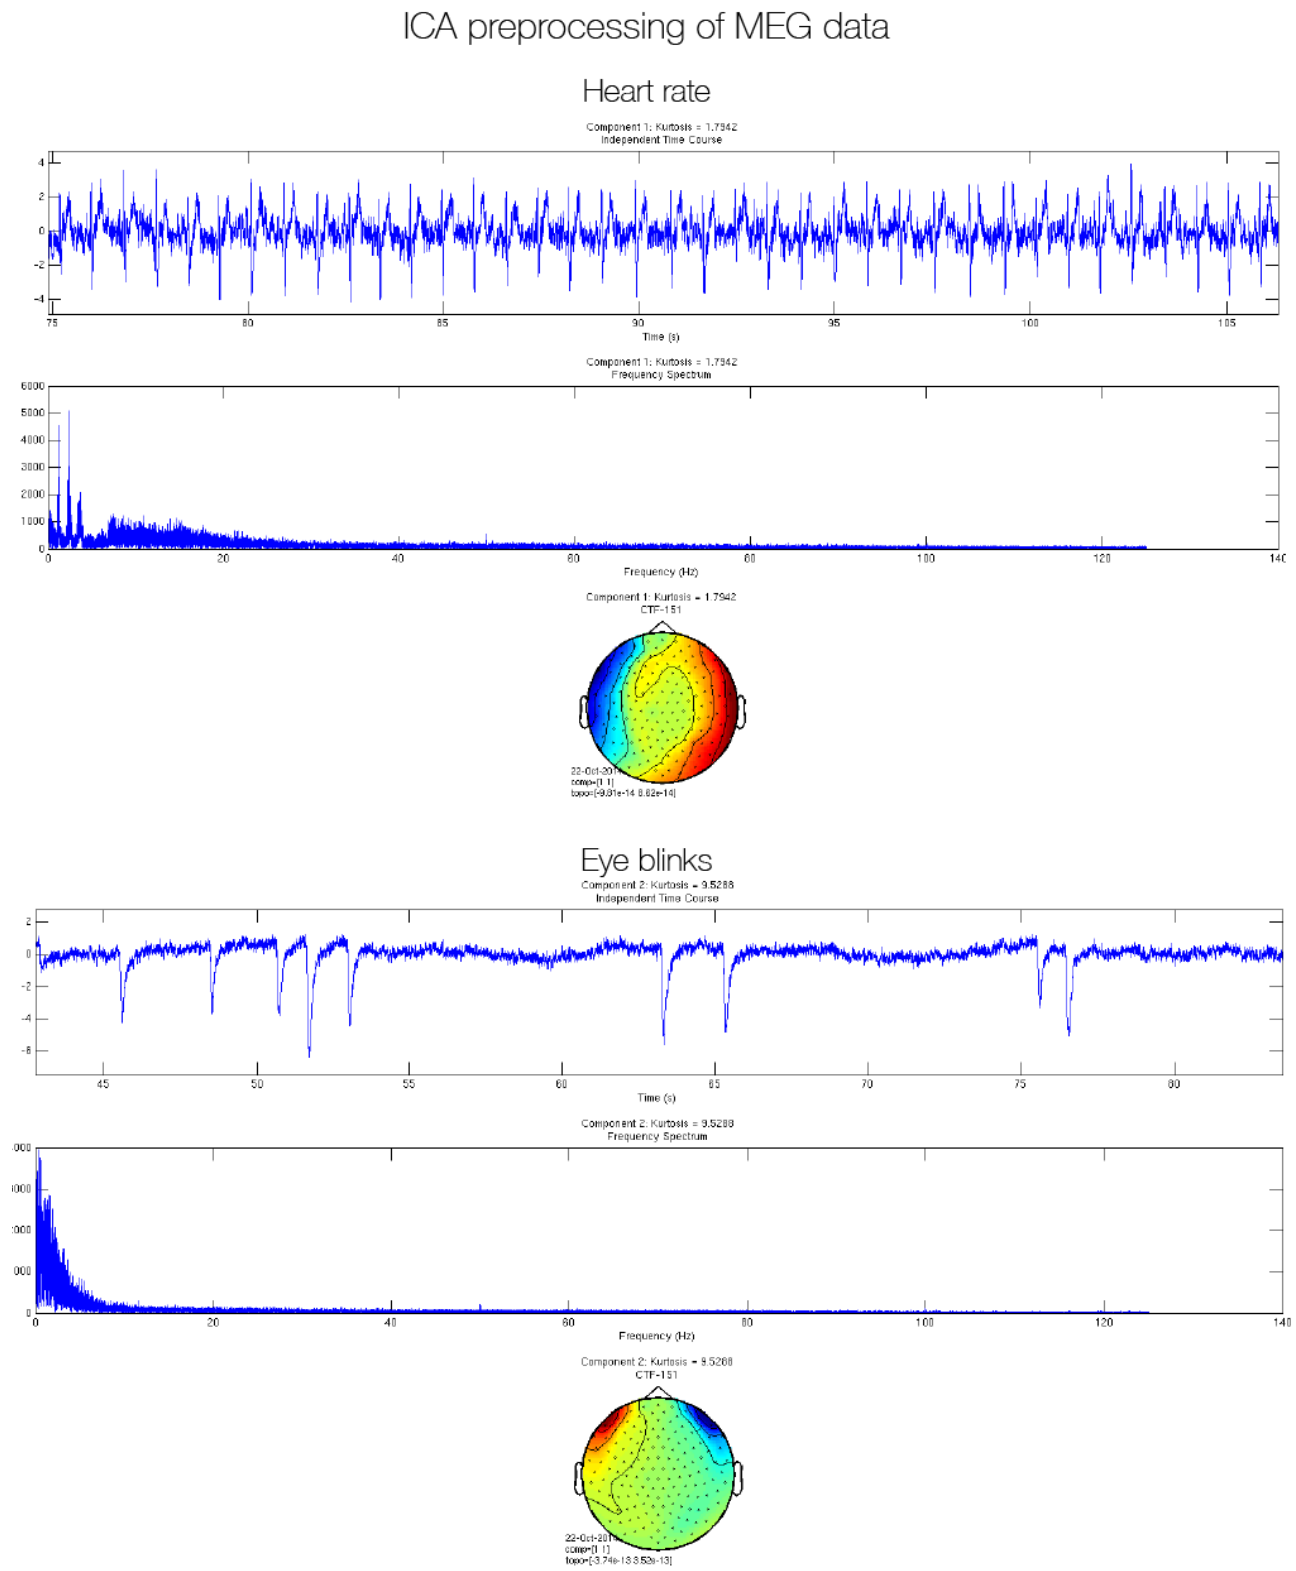

Figure 3-1 ***Examples of artefactual ICs removed by pre-processing*** Top figure shows the timecourse, spectral decomposition, and topographical map of an heart beat related IC. Bottom figure shows the same for an IC related to eye blinks.

3.2.2 *Choice of vigilance condition epochs*

Similarly to the fMRI datasets analysed in Chapter 2, the vigilance conditions were distributed unevenly across as well as within participants in the present MEG data. We based our data selection criteria on a number of considerations. As in Chapter 2 we aimed to get a balanced dataset with each participant contributing with equal amounts of data from each vigilance condition. Moreover, in order to get reliable estimates of MEG FC, we required a minimum duration of data from each vigilance condition. Finally, we wanted segments of data from non-interrupted vigilance stages, taking this as a sign of more reliable sleep scoring. Only three vigilance conditions fulfilled these criteria, and we thus focussed our investigations on wakefulness, N2 sleep, and SWS.

Choosing epochs of durations between 270 and 300 seconds (corresponding to between 9 and 10 sleep scoring windows), allowed for at data set consisting of 11 triplets of vigilance conditions from 8 different participants (3 participants contributed with 2 triplets from different recording days). Results shown in this and the next two chapters are thus based on data from these 8 participants.

3.2.3 *Source-reconstruction of MEG sensor signals*

An overlapping spheres forward model was implemented to allow for source reconstruction of the MEG sensor signals (Huang et al., 1999), in which individual spheres were fitted to each of the sensors of the MEG array. Sources were defined as current dipoles distributed on a 7 mm grid representation of the MNI152 template brain. No structural scans were available for the participants, and thus the template brain was co-registered to each participant's head position estimated at the beginning of each recording session through the use of head position coils. More specifically, the head position coils were used to determine the positions of three fiducial points; the nasion, and the left and right pre-auricular. Through the signals emitted by the head position coils the position of the anatomical point could be related to the MEG sensors. This relationship was then used to co-register, through an affine transformation, the fiducial points in MNI152 space to the space of the MEG sensors. Using a template brain to represent the anatomy of a given participant is an approximation, the accuracy of which will depend on how 'average' the anatomy of the given participant is. For one particular participant, this approximation was inaccurate, leading to an improbable positioning of the co-registered template brain inside the sensor array. In order to avoid the loss of valuable data, we chose to use a structural scan from a different participant from an unrelated study, in the form of a T1-weighted MR image, which was found to fit the dimensions of the participant-specific fiducial positions. A similar strategy was shown to be an efficient and sensible strategy to deal with MEG datasets where structural information is missing or corrupted (Holliday et al., 2003). The chosen T1 image was first registered to the MNI152 template by affine transformation and the resulting mesh was then coregistered to the space of the MEG sensors following the procedure outlined above.

A linearly constrained minimum variance (LCMV) beamformer was used to reconstruct into source space the continuous MEG signals (Van Veen et al., 1997, Hillebrand et al., 2005, Woolrich et al., 2011). The data from each vigilance-specific epoch were filtered into 6 frequency bands: delta (0.25–4 Hz), theta (4–8 Hz), alpha (8–13 Hz), sigma (13–16 Hz), beta (16–30 Hz), and gamma (30–48 Hz), and separate

covariance matrices were estimated for each band and vigilance-specific epoch. The sigma band is commonly disregarded in resting-state studies, where it is included as part of a wide beta band, however it was included here to increase specificity of sleep-spindle-related activities. While slightly higher than some estimations in the literature (Warby et al., 2014), our choice of 13 – 16 Hz for the sigma band was based on visual inspections of frequency spectra of the data epochs (see Figure 3-2). When taking the mean of the spectrograms over all channels included in the N2-epochs, a peak was evident between 13 and 16 Hz. This peak was absent for wakefulness and less pronounced in SWS, as one would expect, under the assumption of spindling activity. The dominant frequencies of sleep spindles are known to vary with time of day (Aeschbach et al., 1997), and the daytime recordings in healthy young adults of the present study is in line with a slightly increased spindle-specific frequency range.

As indicated above, estimation of sensor covariance is necessary for beamforming (Hillebrand and Barnes, 2005), and the precision of the latter is contingent on the reliability of the former. MEG resting-state studies have commonly estimated sensor covariance over the full extent of the recordings, under the assumption that this is relatively stable over time (Brookes et al., 2011a). It has been shown, however, that subtle non-stationarities in the sensor covariance, e.g. related to task, can produce spurious beamforming results (Woolrich et al., 2013). The significant spectral reconfiguration occurring in sensor signals, upon which traditional sleep staging is based (see Introduction, 1.2.1), are likely to produce bigger differences in sensor covariances from wakefulness and across the specific sleep stages. This was an important reason for creating individual spatial filters for each vigilance-specific epoch.

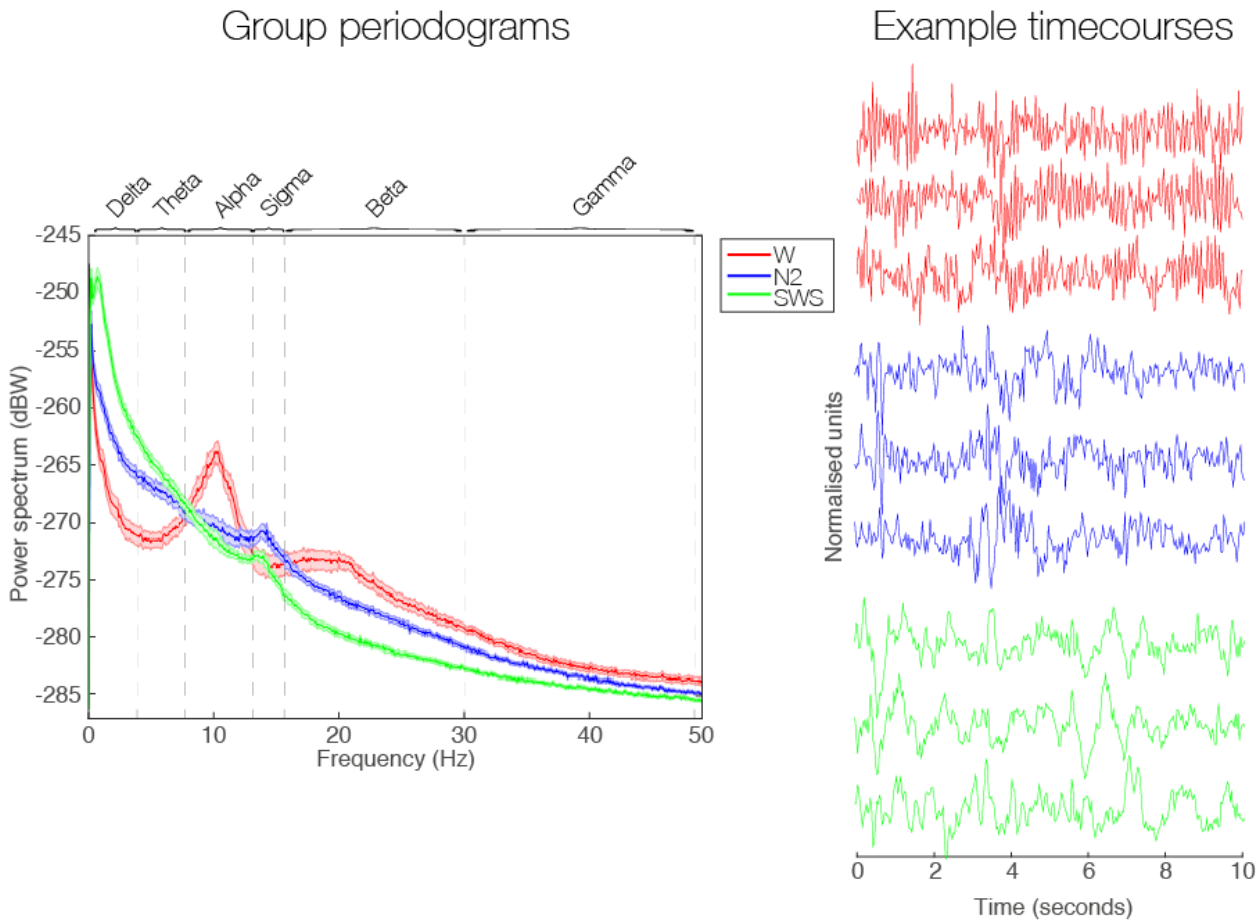

Figure 3-2 ***Spectrogram of sensor recordings included in functional connectivity analysis.*** *Left*: Power values were averaged across all sensors and sessions within vigilance conditions. Mean values are shown in solid, while error zones represent standard error of the mean. Vertical dashed lines delineate the frequency bands employed to construct the FC networks. *Right*: 10 seconds of example data from each of the vigilance conditions (red: wakefulness, blue: N2 sleep, green SWS)

3.2.4 *Brain parcellation and functional connectivity*

We followed the methods outlined in Colclough and colleagues (2015, 2016) for computing static functional connectivity (FC) from source reconstructed MEG signals. This included the choice of brain parcellation. As such, 39 non-overlapping, and spatially contiguous cortical ROIs were derived based on a clustering analysis on an ICA decomposition of fMRI resting-state data from the first 200 subjects of the Human Connectome Project (HCP) database (Van Essen et al., 2013). Sub-cortical regions were omitted from the analyses since the sensitivity of MEG is limited in deeper areas of the brain. Each ROI spanned a high number of voxels and associated beamformed timecourses, and hence we used Principal Component Analysis (PCA) to estimate one representative timecourse per ROI. The first principal component, explaining the majority of the variance was used.

When estimating FC from source-reconstructed MEG signals one has to pay particular attention to the so-called source-leakage bias (Brookes et al., 2012, Hipp et al., 2012, Colclough et al., 2015). Source-leakage describes the phenomenon that ambiguities in the source reconstruction (i.e. beamforming) result in a spatially varying and anisotropic point spread function of the source reconstructed timecourses, which may

cause part of a signal from one source, to 'leak' into the estimated signal from a different source. This can cause statistical interdependencies, such as linear correlations, between ROIs to become spuriously high. The method devised by Colclough and colleagues (2015) was designed specifically to deal with this potential bias and has proven its performance in simulations and cross-validations (Colclough et al., 2015, Colclough et al., 2016).

We applied said method to each vigilance-specific set of 39 PCA-derived timecourses in each of the six frequency bands. This resulted in 39 mutually orthogonal ROI timecourses for each epoch and frequency band. A number of recent studies have found the low-frequency component of power fluctuations in band-limited oscillatory activity optimal for capturing functional connectivity in MEG source space (Liu et al., 2010, Brookes et al., 2011a, Smith et al., 2011a, Hipp et al., 2012, Luckhoo et al., 2012, Cabral et al., 2014c). Thus, functional connectivity information has been found to be concentrated around fluctuations in power envelopes between 0.1 and 1 Hz. Accordingly, we used the Hilbert transform to extract power envelopes from the complex analytic signal, S, for each time point, t, vigilance-state-specific epoch, v, and frequency band, f:

$$S(r_i,t)^{v,f} = d(r_i,t)^{v,f} + jH(r_i,t)^{v,f} \quad \textbf{Eq. 3.1}$$

where $d(r_i,t)^{v,f}$ is the beamformed, PCA-derived, and orthogonalised timecourse at location r_i, i = 1 … 39 cortical ROIs, j represents the complex unit, and $H(r_i,t)$ is the Hilbert transform. Expressing the complex-valued analytic signal in polar form using Euler's formula yielded:

$$S(r_i,t)^{v,f} = A(r_i,t)^{v,f} e^{j\psi(r_i,t)^{v,f}} \quad \textbf{Eq. 3.2}$$

where $A(r_i,t)^{v,f}$ is the amplitude of the analytic signal, which we will refer to as the envelope, and the phase of the analytic signal is denoted by $\psi(r_i,t)^{v,f}$.

To obtain the low-frequency component of the power fluctuations in the six frequency bands of interest, we low-pass filtered the absolute values of the envelopes at 0.5 Hz and downsampled to 1 Hz. We chose two different, albeit related, measures of FC between the corrected ROI timecourses: i) full linear correlations and ii) partial correlations. For the full linear correlation the Pearson's correlation was used (as in Eq. 2.1). Partial correlations were calculated from the inverse of the covariance matrix between the ROI timecourses, also known as the precision matrix (Marrelec et al., 2006b). The covariance matrix was computed using the built-in MATLAB function, *cov*, and inverted using the Cholesky decomposition. Estimations of partial correlations can be improved by regularising the estimate of the precision matrix and suppress small values, which promotes sparsity. We regularised the partial correlations as in Colclough and colleagues (2015), using the graphical lasso (Friedman et al., 2008). Full correlations and partial regularised correlations were z-transformed using Fisher's transformation. An additional scaling was applied to the z-scores according to the standard deviation of an empirical null distribution, computed through the simulation of null networks of Gaussian uncorrelated timecourses with the same number of nodes as ROIs and temporal smoothness

derived from the original data. Simulations were iterated 8 times for each vigilance-state-specific network and the temporal smoothness were modelled using the first auto-regressive coefficient (Colclough et al., 2015).

The use of full linear correlations for FC estimation by far exceeds the use of partial correlations in fMRI as well as MEG (Smith et al., 2011b, Woolrich and Stephan, 2013). Meanwhile, full correlations do not distinguish between direct and indirect connections. Given a network of ROIs: A, B, and C, in which A and C, and B and C are functionally connected, the timecourses between A and B might show a high correlation, even if they are only indirectly connected through C. Partial correlation adjusts the correlation between the network elements, say A and B, by regressing out contributions from other parts of the network, and as such resolves the unique and direct connections.

Assuming a successful and realistic parcellation of the brain into its functional subdivisions, partial correlations arguably represent a more sensible choice of FC measure (Marrelec et al., 2006a, Smith et al., 2011b, Duff et al., 2013). Colclough and colleagues recently showed that partial correlations produce more realistic FC networks in source-reconstructed and leakage-corrected MEG data, compared to full correlations (Colclough et al., 2015). Moreover, evidence from the same group suggest that full as well as partial correlations between band-limited power envelopes yield networks of high intra-participant repeatability compared to a range of other FC metrics that have been suggested for MEG (Colclough et al., 2016).

To get a group estimate of static FC networks for each vigilance condition and each frequency band, we treated each vigilance-specific epoch as independent observations in a fixed-effects analysis testing the significance of the mean of the standardised z-scores between each pair of ROIs (Colclough et al., 2015). For an overview of the analysis pipeline from sensor recordings to FC networks, see Figure 3-3.

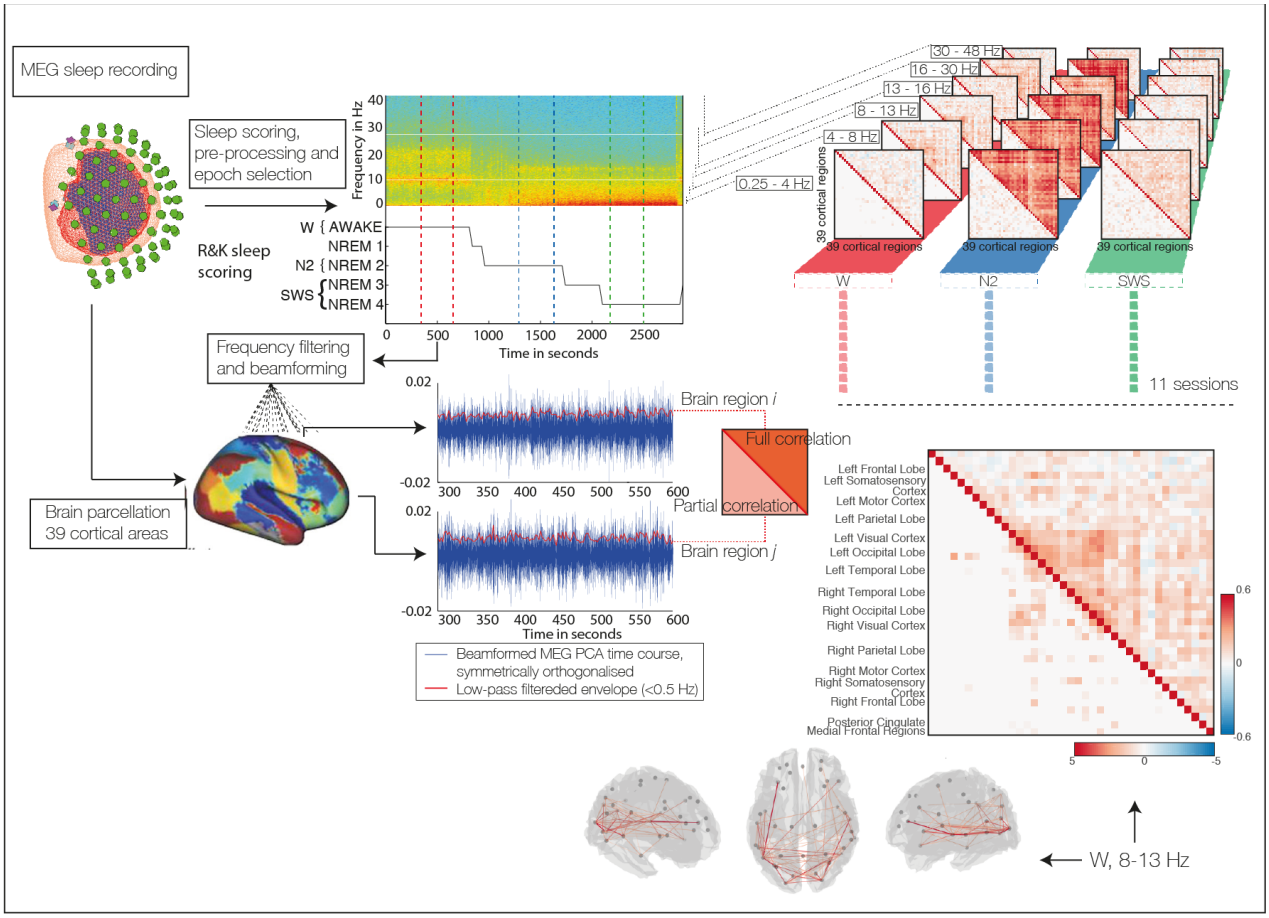

Figure 3-3 ***Overview of analysis workflow for computation of vigilance-specific MEG FC.*** *Top left*: MEG recordings of healthy participants were sleep scored in accordance with standard principles (Rechtschaffen and Kales, 1968). Sensor data were pre-processed in continuous segments (with no regard to sleep scoring) using ICA-based removal of heart rate, ocular, and muscular interference. *Middle left*: Pre-processed sensor data were epoched into ~300 second segments of wakefulness (W), NREM sleep stage 2 (N2), and slow wave sleep (SWS), band-pass filtered into six canonical frequency bands (0.25-4 Hz, 4-8 Hz, 8-13 Hz, 13-16 Hz, 16-30 Hz, and 30-48 Hz). An LCMV beamformer was used to source-reconstruct each epoch onto a 7mm grid in MNI space with covariance matrices estimated separately for each frequency band. Voxel-data was extracted corresponding to 39 cortical ROIs of a functional fMRI-derived atlas, and PCA was used to summarise brain activity in each ROI as the principal component explaining the majority of the variance. A symmetric multivariate orthogonalisation procedure was used to correct for source leakage, and the Hilbert transform was used to extract power envelopes from the corrected ROI timecourses. Envelopes were low-pass filtered at 0.5 Hz and downsampled to 1 Hz. The 39 low-pass filtered power parcel envelopes were associated with each other within frequency bands by computing full and partial correlations. *Top right*: The resulting functional connectomes from an example participant within the six frequency bands and three vigilance conditions. Super-diagonal parts of the matrices show the full correlations while the sub-diagonal parts show the partial correlations. In total, functional connectomes were computed for 11 triplets of the three vigilance conditions. *Bottom right*: Example of a functional connectome from the alpha band (8-13 Hz) for wakefulness showing a dominance of visual connectivity. The corresponding brain plot shows the 10 % strongest connections, taken from the partial correlations in the sub-diagonal part of the matrix.

3.2.5 *Vigilance-specificity of functional connectivity patterns*

As discussed in the Introduction as well as in Chapter 2 fMRI-derived FC networks have been shown to change with vigilance in very specific ways, allowing for the classification of wakefulness and sleep solely based on FC patterns (Tagliazucchi et al., 2012b, Tagliazucchi and Laufs, 2014, Altmann et al., 2016). A sensible question to ask is whether the MEG FC networks derived here show a similar level of vigilance-specificity. Our limited amount of observations, i.e. data epochs, prohibited a similar machine learning classification approach. However, we were able to test whether network patterns within each of the three included vigilance conditions (W, N2, and SWS) were more alike themselves than any of the two other states. A similar approach was recently used by Colclough et al. (2016) to quantify the reproducibility of intra-participant FC patterns in a large sample of resting-state recordings from the HCP database (Larson-Prior et al., 2013).

For each frequency band we paired FC matrices within and across vigilance conditions in all combinations. Given the 11 sessions from each vigilance condition, discounting comparisons within recording sessions, this yielded $\frac{11!}{2!(11-2)!} = 55$ intra-condition comparisons, and $11^2 - 11 = 110$ across-condition comparisons. The symmetric FC matrices were converted to column vectors using only their sub-diagonal parts, and Pearson's linear correlation was used to estimate the similarity between these column vectors. In Chapter 2, we compared FC matrices from fMRI to a group-averaged SC matrix (see Eq. 2.3) in terms of their patterns similarity. Here, instead, we compared MEG FC networks to each other. Pattern similarities were computed for both the full and partial correlation-based FC matrices.

For each vigilance condition the distribution of intra-condition similarities, R_{intra}, was represented by a pseudo-statistic, represented by the ratio between its mean value and standard deviation:

$$\mu_{intra} = \frac{1}{N}\sum_{i=1}^{N} r_i \qquad \textbf{Eq. 3.3}$$

$$\sigma_{intra} = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N} \left| r_i - \mu_{intra} \right|^2} \qquad \textbf{Eq. 3.4}$$

$$pseudostat_{intra} = \frac{\mu_{intra}}{\sigma_{intra}} \qquad \textbf{Eq. 3.5}$$

where N = 1…55.

In order to quantify the level of intra-condition specificity, we embedded the intra and across-condition comparisons in a non-parametric permutation-based testing. Random distributions were thus created from a mix of the original intra-condition comparisons and each of the two across-condition similarity distributions, $R_{across1}$ and $R_{across2}$. For each vigilance condition this yielded two super-distributions of 110+55 = 165 similarity values each: $R_{intra+across1}$ and $R_{intra+across2}$. From each of these two random super-distributions we picked 55 values at random, and calculated $pseudostat_{rand1}$ and $pseudostat_{rand2}$, 5000 times. In the end this yielded two null-distributions, R_{rand1} and R_{rand2} consisting of 5000 pseudostat values each. The significance of

the difference between the original $pseudostat_{intra}$ and each of the two null-distributions could then be estimated as:

$$p_{|intra|vs|across|1} = P\left(pseudostat_{intra} \leq R_{rand1}\right) = \left[\sum_{R_{rand1} \geq pseudostat_{intra}} R_{rand1}\right] / N_{permutations} \quad \textbf{Eq. 3.6}$$

$$p_{|intra|vs|across|2} = P\left(pseudostat_{intra} \leq R_{rand2}\right) = \left[\sum_{R_{rand2} \geq pseudostat_{intra}} R_{rand2}\right] / N_{permutations} \quad \textbf{Eq. 3.7}$$

where P(.) denotes a cumulative distribution function, and $N_{permutations} = 5000$. p-values below 0.05 were considered significant. As an example, for W-epochs in the 8-13 Hz band, $p_{|intra|vs|across|1}$ could correspond to the comparison between the intra-W similarity values and the W-to-N2 similarity values. A significant p-value would then signal that FC matrices were more similar to other W matrices than to N2 matrices.

The analysis described above was devised to quantify the relative consistency and specificity of the derived MEG FC networks within vigilance conditions. Having established this, we were further interested in whether this FC information was particularly sensitive to changes in vigilance. As a reference, we therefore carried out the same analysis on an alternative representation of the 39 ROIs, considering the mean power values as opposed to the FCs between them. We computed the 39 mean values of the band-specific envelopes, which yielded a vector, which could readily be correlated with equivalent distributions of power from other sessions. We followed the steps outlined above, and for each frequency band and each vigilance condition thus obtained a p-value for higher similarity within versus across vigilance conditions.

In order to make sure that any potential differences in terms of vigilance-specificity between FC networks and power distributions were not due to unequal observations in the two cases (741 versus 39), we performed a third specificity analyses based on the degree distributions of the FC networks (Rubinov and Sporns, 2010). As such, we summarised each of the FC matrices in 39 values by summing over their columns, and correlated these within and across vigilance conditions, again using the steps described above. Ultimately, these analyses address how consistent MEG FC networks are within the included vigilance stages, and whether FC networks are more discriminative than power distributions, which do not take into account inter-regional interactions.

3.2.6 *Sleep-dependent changes in functional connectivity*

As in Chapter 2, we tested for differences between vigilance conditions at the global level of mean FC as well as at the regional level of individual connections and sub-networks.

Paired t-tests were performed on the mean full correlations and mean partial regularised correlations within each frequency band. We used a sign-flipping approach to estimate the significance of each test, by randomly shifting around the condition labels (within recording sessions). We performed the sign-shifting 1000 times, and for each random configuration of condition labels we performed a paired t-test and used the resulting t-stat to build a null-distribution. We then compared the null-distribution and the t-stat of the original test (e.g. WvsN2 in the 8-13 Hz band), by calculating a p-value as the probability of finding absolute

values from the null distribution exceeding the absolute value of the original t-stat. p-values below 0.05 were considered statistically significant.

At the regional level, we performed comparative network analyses on the FC strengths as measured with full and partial correlations, combining paired t-tests with both the connection-wise and the NBS correction for multiple comparisons (see Chapter 2, 2.3.3). As in Chapter 2, a t-stat threshold had to be chosen for the NBS analysis. It was clear that the t-stats fell into two different domains for the full and the partial regularised correlations respectively. As a result we searched for significant differences in different ranges for the two measures; an interval around t = 5 for the full correlations, and an interval around t = 2 for the partial regularised correlations.

For the connection-wise correction a connection was considered significantly different at corrected p-values < 0.05. The same level was employed for the NBS, where significance pertained to sub-networks and not individual connections.

3.3 ***Results***

3.3.1 *Group averaged, vigilance specific functional connectomes*

The resulting group averages of the FC networks are shown in Figure 3-4. Using full and partial correlations between the beamformed, low-pass filtered power envelopes, we were able to summarise the functional connectivity profiles within each vigilance condition and frequency band. There are several interesting things to note from these group averages: i) The mean level of FC was clearly increased in N2 sleep compared to both wakefulness and SWS across all frequency bands except gamma. This effect is quantified in 3.3.3. ii) The patterns of the functional FC networks of wakefulness followed previous results from resting-state recordings, with FC concentrated in the alpha and beta bands. These included strong connectivity in visual areas extending towards parietal and temporal areas in the beta band. Notably, only wakefulness showed clear signs of FC in the gamma band, especially in fronto-temporal connections iii) The use of partial correlations had a particularly strong impact on the N2 condition, where relative FC strengths and patterns changed markedly. The most robust direct connections involved frontal, temporal, and parietal regions in the alpha, sigma, and beta bands. iv) The most distinct connectivity pattern for SWS was found in the alpha band between frontal and temporal regions, whereas the lower frequency bands of delta and theta (0.25–4 Hz and 4–8 Hz), showed low levels of FC and no clear signs of structured patterns.

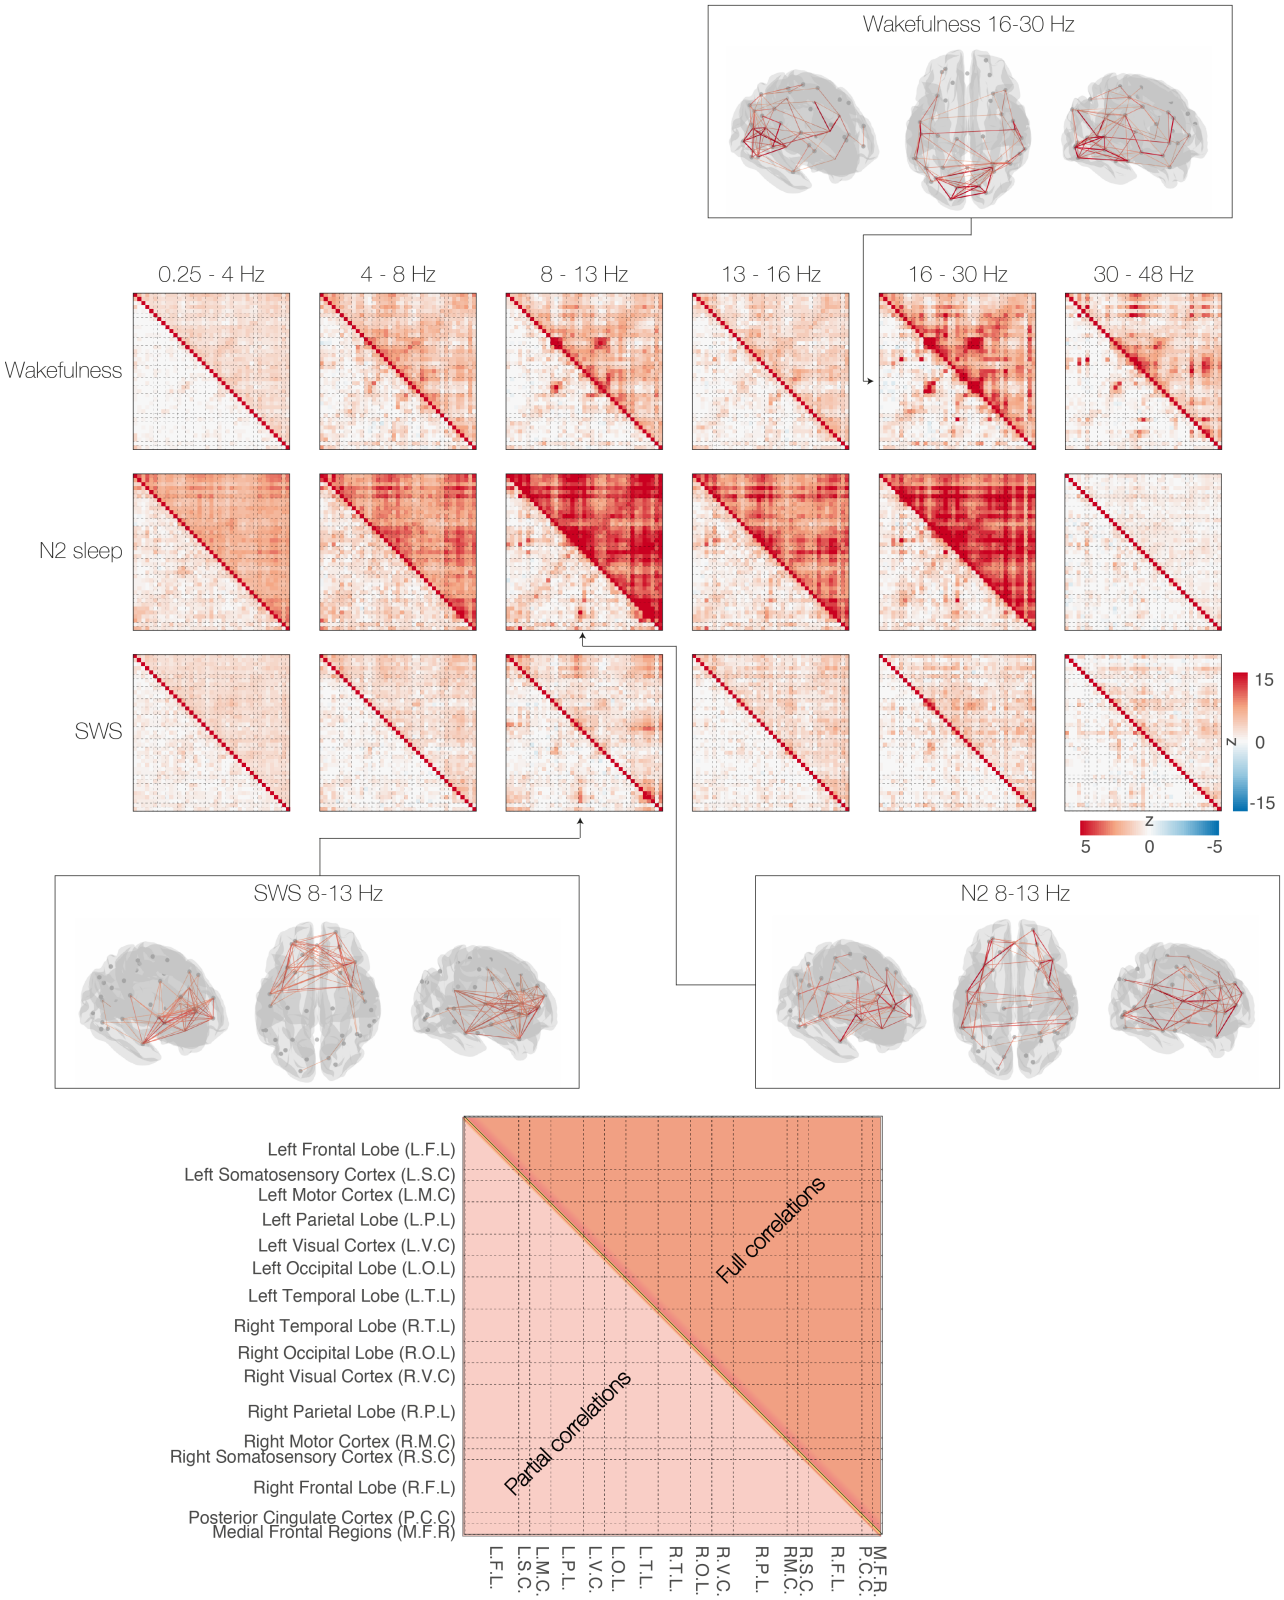

Figure 3-4 ***Group-averaged functional connectomes from wakefulness, N2, and SWS***: A fixed-effects analysis was performed within each vigilance condition and frequency band for every connection between the 39 cortical ROIs Resulting z-statistics are shown for full correlations and partial correlations in upper diagonal and lower diagonal matrix parts, respectively. Note the difference in scaling between full and partial correlations (full: vertical colour-bar, partial: horizontal colour-bar). The brain plots show the 10 % strongest connections in the beta band for wakefulness and in the alpha band for N2 and SWS.

3.3.2 *Specificity of FC networks to changes in vigilance*

The FC networks showed visible differences in their patterns from wakefulness to sleep (see Figure 3-4). We wanted to test how specific these patterns were to each of the included vigilance conditions, and whether certain frequency bands would prove more important than others. By correlating FC matrices within and across conditions, we could probe the similarity of network patterns (as explained in 3.2.5). The results of this analysis are shown in Figure 3-5A and B for FC networks of both full and partial correlations. Plotted here are, for each frequency band, the mean values of the linear Pearson's correlation between FC matrices; within conditions (solid lines), and from a condition to the two other conditions (dashed lines). When a solid line shows a significantly higher value than a dashed line (represented by a black star), this signals that the FC patterns within the condition in question are more similar to themselves than to patterns from the condition represented by the dashed line.

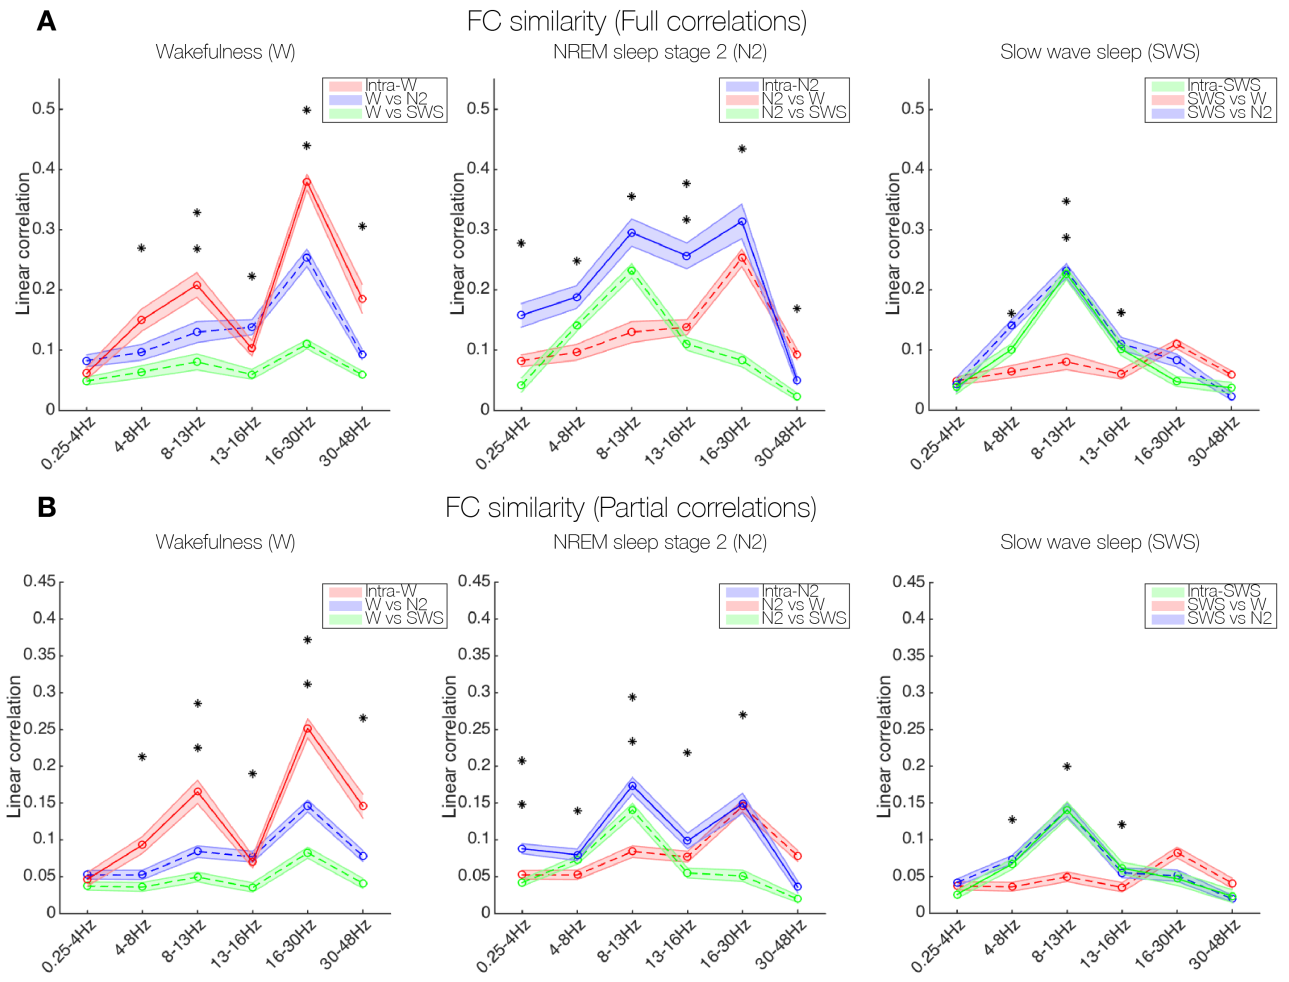

Figure 3-5 ***Similarity of FC patterns within versus across vigilance conditions:*** **A)** FC matrices computed using full correlations were compared within each of the three vigilance conditions by correlating sub-diagonal parts of the matrices. The means of intra-condition similarities are plotted against frequency bands in solid and the standard errors of the means shown as error zones. Within each vigilance condition matrices were also compared to matrices from the remaining two conditions. The means of the across-condition similarities are plotted as dashed lines and the standard errors of the means are shown as error zones. Black stars denote significant differences between intra-condition correlations and across-condition correlations. The significance of these differences was estimated using non-parametric permutation tests. **B)** Same as **A** but for partial correlations**.**

The significance of differences between a given condition and any of the two other conditions were calculated using permutation testing on the distributions making up intra-condition correlations and across-condition correlations (for more details, see 3.2.5). Relative specificity was fairly consistent from full to partial correlations (compare A and B in Figure 3-5), albeit the general level of similarity (correlations between FC matrices) was generally lower for partial correlations.

For wakefulness, FC patterns showed high specificity in the alpha and beta bands. For N2 sleep, specificity was highest in the sigma band when full correlations were used. This changed to the delta and alpha bands when considering partial correlations. For both measures of connectivity, N2 sleep was more similar to SWS in the lower frequency bands, and more similar to wakefulness in the higher frequency bands. The curve of intra-SWS correlations followed the curve of 'SWS vs N2' quite closely for both full and partial correlations. A clear peak was evident for the alpha band; here SWS was easily separated from wakefulness, and from N2 sleep when full correlations were considered.

To illustrate the discriminatory importance of including FC information, we tested whether a similar separation could be obtained from the power values of the beamformed, low-pass filtered envelopes (for more details, see 3.2.5). Results are shown in Figure 3-6A. With the exception of the delta band for N2 sleep, the envelope power values could not be used to separate between vigilance conditions. Solid and dashed lines, representing intra- and across-condition correlations respectively, followed each other closely at a high level of similarity/correlation, and were generally not specific to any vigilance condition.

We performed the same analysis on degree derivations of the FC networks, thus transforming the latter to the same dimensionality as the envelope power distributions (see Figure 3-7 for examples of degree distributions from the three vigilance conditions). We did this for both full and partial correlations, and the results are shown in Figure 3-6B and C. At this level of description, separation between vigilance conditions was once again evident.

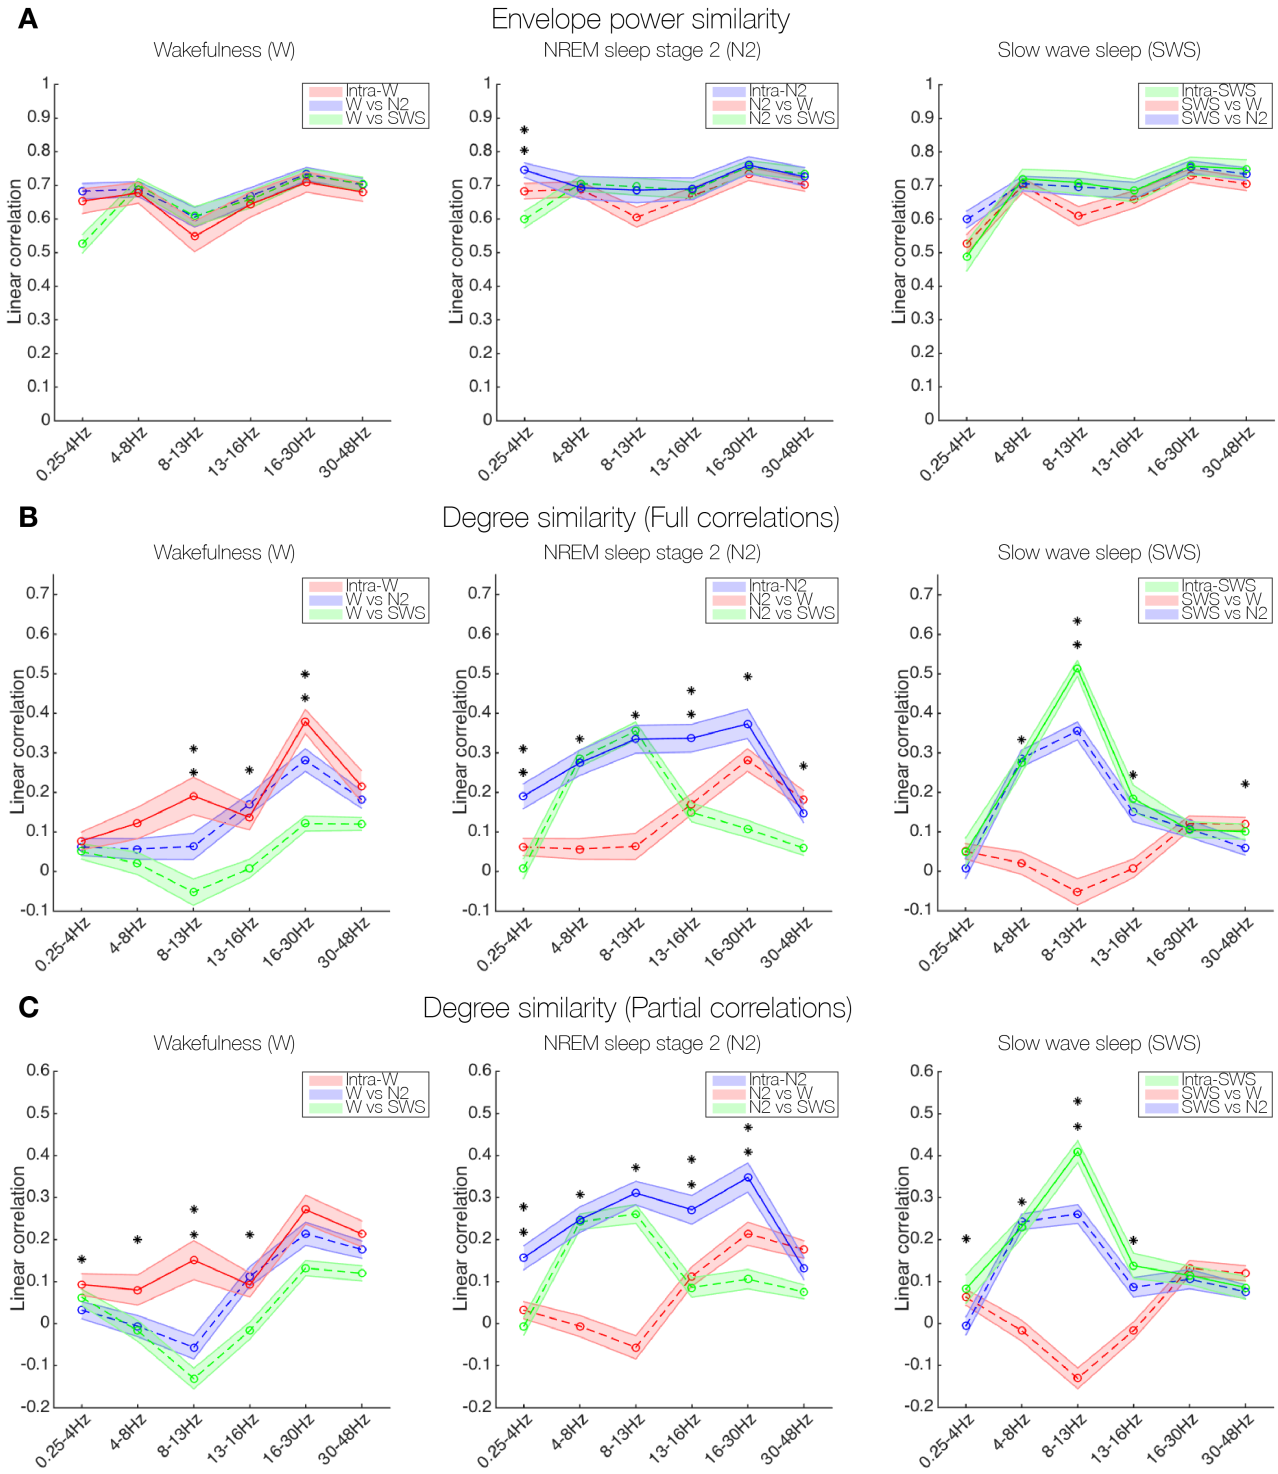

Figure 3-6 ***Similarity of envelope power distributions and functional connectivity degree distributions within versus across vigilance conditions***: **A)** Similar to results shown in **Figure 3-5** but instead of comparing functional connectivity matrices, distributions of mean envelope power values from the 39 ROIs were correlated to estimate similarity within and across vigilance conditions. Mean intra-condition similarity is plotted in solid against frequency bands with error zones of standard errors of the means. **B)** Summing over columns of the FC matrices, computed using full correlations, degree distributions were derived, and similarities of these distributions were estimated, as previously, within and across vigilance conditions. **C)** Same as B, but for degree distributions based on partial correlations.

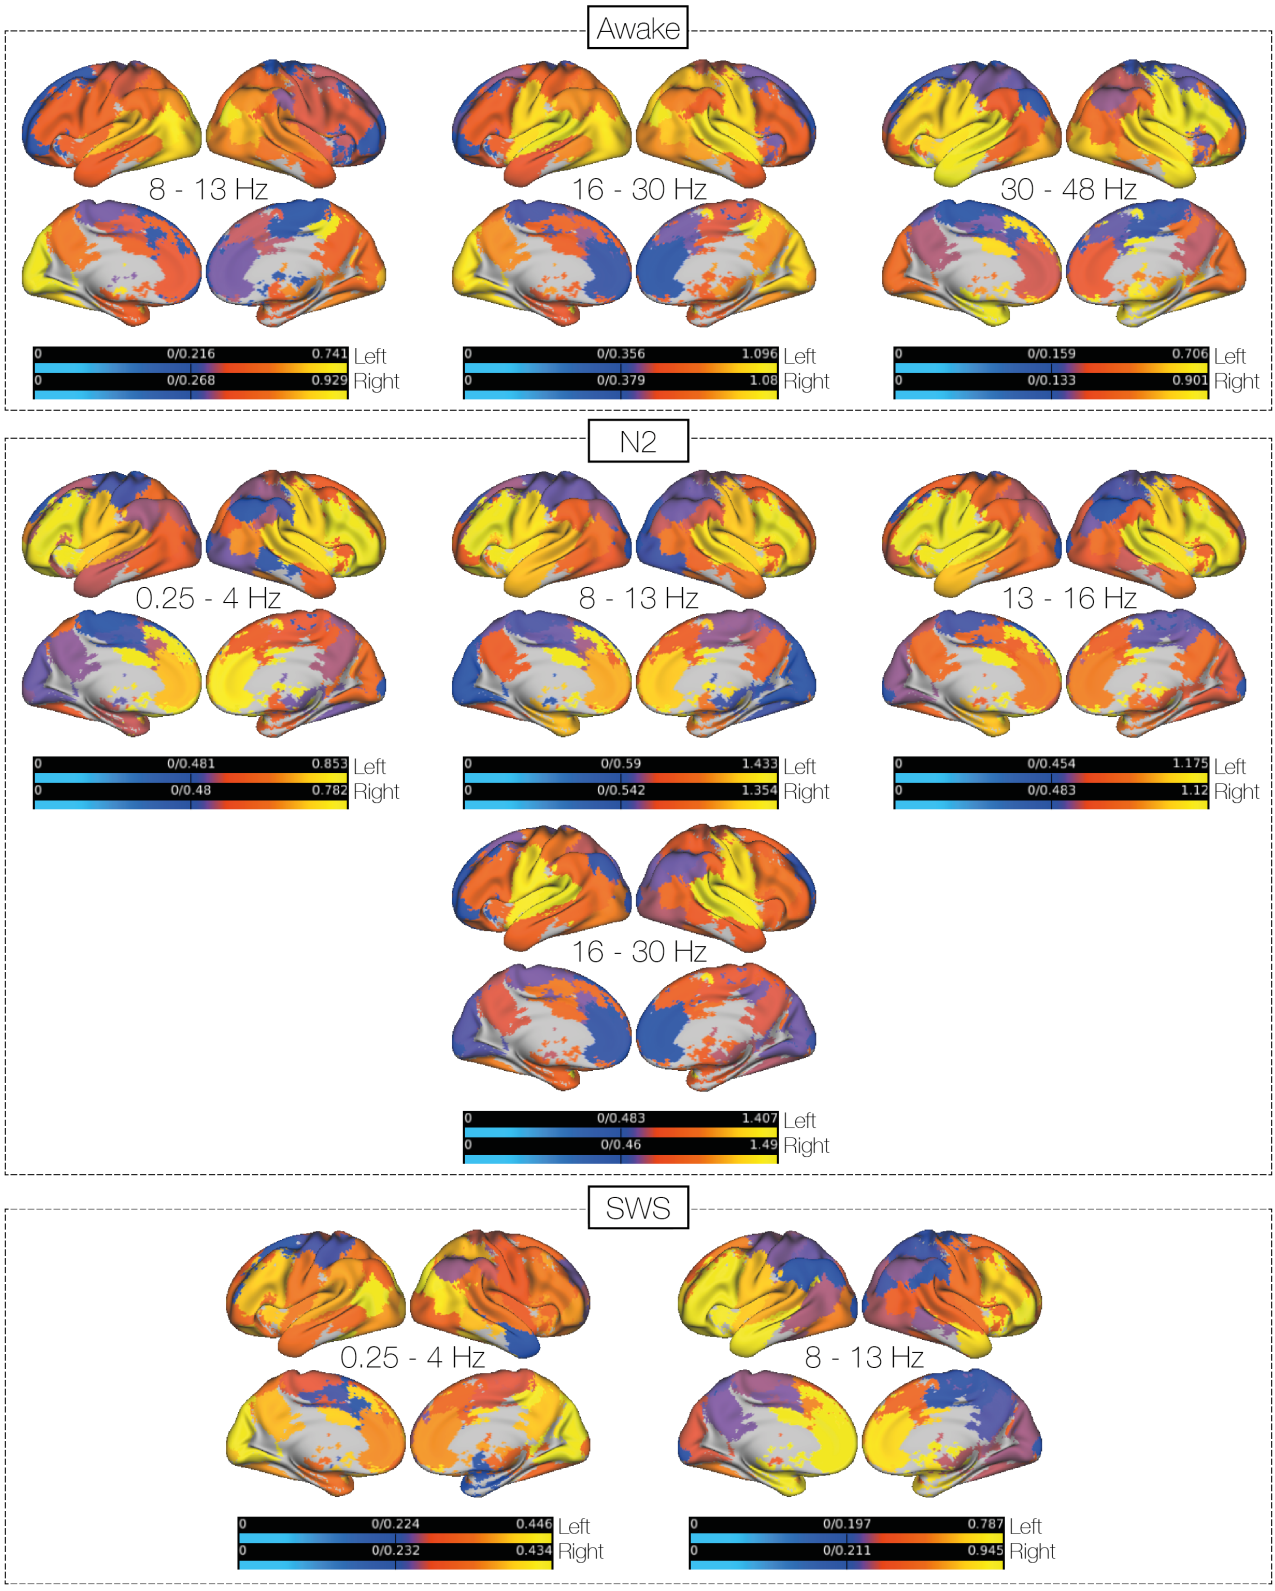

Figure 3-7 ***Degree maps from each of the three vigilance conditions***: By summing over the columns of the FC matrices each ROI can be represented by its total amount of connectivity, which is equivalent to its non-thresholded degree. In this case the partial correlation matrices were used to illustrate how FC changes across vigilance conditions in some prominent frequency bands. The maps show the group-averaged degree. Degree distributions were used for the calculation of the within- and cross-vigilance similarity estimation.

3.3.3 *Changes in mean functional connectivity*

The vigilance-specific mean values of functional connectivity are shown as bar plots in Figure 3-8A and B for full and partial correlations, respectively.

Elevated levels of FC in N2 were already indicated in Figure 3-4, an effect that was particularly evident for full correlations. Using paired t-tests and sign-flipping, mean levels of FC of N2 were indeed found significantly higher than both of the other two conditions in all frequency bands except gamma, for full correlations. For partial correlations, N2 was only found higher than wakefulness in the delta band, while SWS was still lower in all bands but gamma. Wakefulness and SWS only differed significantly in the beta band, where wakefulness showed higher mean values of connectivity for both full and partial correlations.

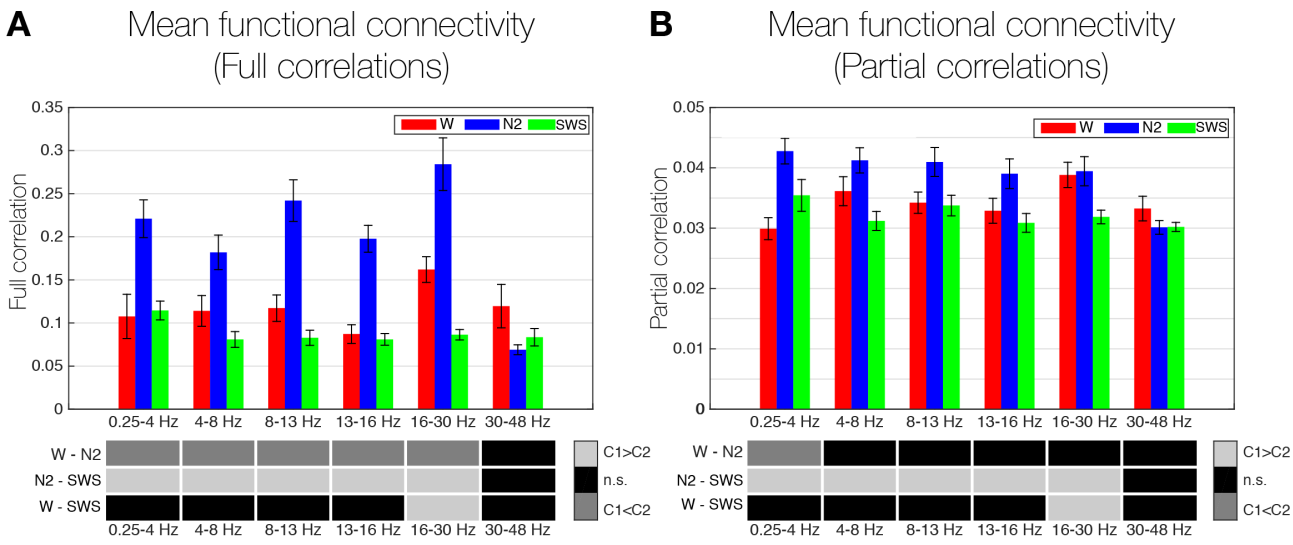

Figure 3-8 ***Vigilance-dependent differences in mean FC*****: A)** Bar plot shows the mean of values in the sub-diagonal parts of functional connectivity matrices, based on full correlations, within vigilance conditions, error bars represent standard error of the mean; red: wakefulness, blue: N2 sleep, green: SWS. Greyscale matrix indicate direction of potential significant differences between vigilance conditions: light grey: significant positive difference of contrast given by matrix row, dark grey: significant negative difference of contrast given by matrix row, black: no significant difference. Significance of differences were evaluated using sign-flipping of paired t-tests. p-values < 0.05 were considered significant. **B)** Same as A, but mean FC values calculated from partial correlations.

3.3.4 *Changes in local connections and sub-networks*

The similarity and mean FC analyses presented above indicate that consistent, frequency-specific re-configurations take place in FC networks following changes in vigilance from wakefulness to sleep. To characterise these changes in more detail we tested vigilance-dependent changes at the regional level, using connection-wise and NBS corrections for multiple comparisons. The resulting difference matrices are shown in Figure 3-9.

For the full correlations, the initial observation of large increases in N2 sleep manifested as dense and widespread networks of increased FC relative to both wakefulness and SWS. This was true for both the connection-wise and the NBS method, which showed high correspondence for the full correlations. The NBS was thresholded at $t > 5.0$ for the results reported here.

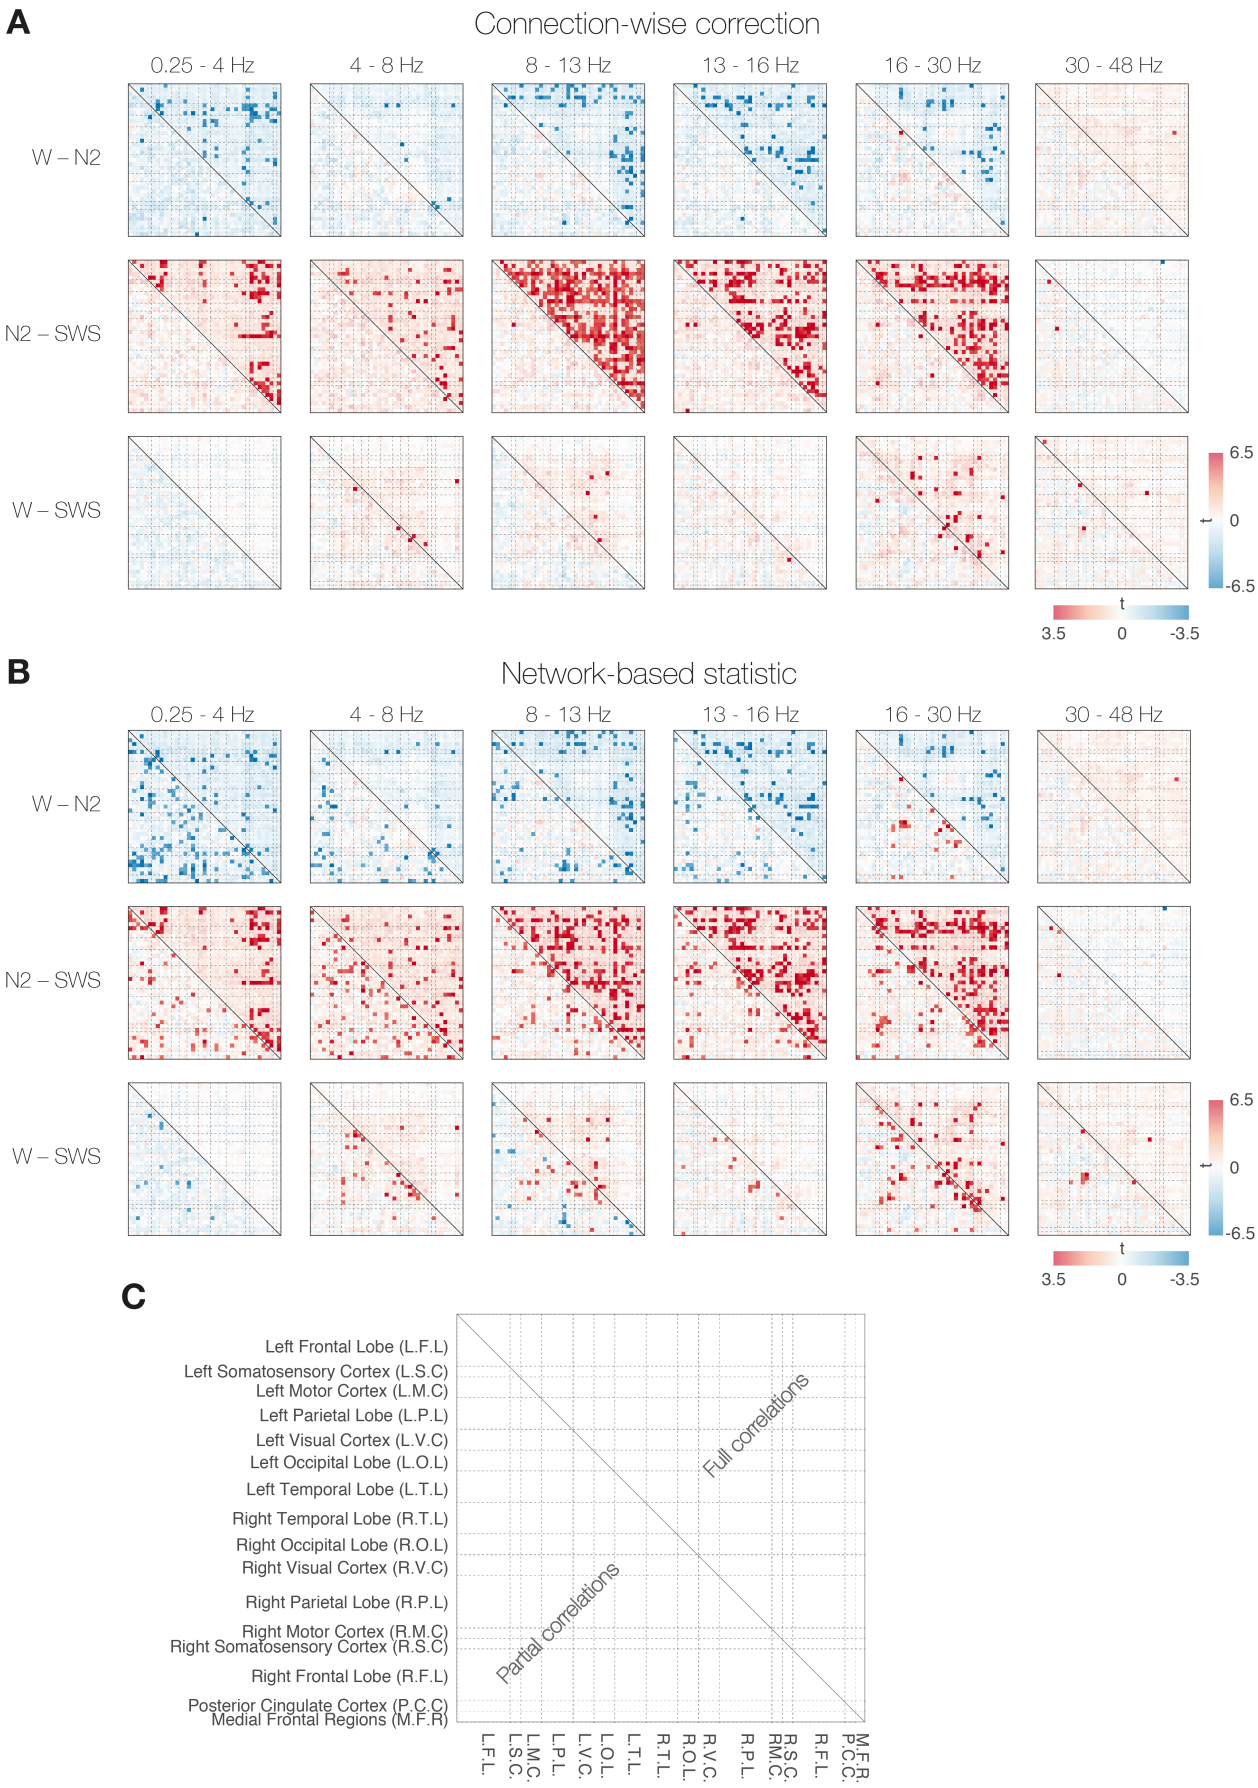

Figure 3-9 ***Difference matrices from paired t-tests between vigilance conditions*****: A)** Results of comparative network analyses between the three vigilance conditions with connection-wise correction for multiple comparisons. Paired t-tests were performed within each frequency band, and the resulting matrices of t-stats are shown with the significant connections highlighted. Red/blue entries correspond to connections that were larger/smaller in the first condition of a

contrast. The contrasts are shown on the left of the matrices. The super-diagonal parts are results based on full correlations, while the sub-diagonal parts are based on partial correlations. Significant connections had corrected p-values < 0.05. **B)** Same as **A,** but using NBS to correct for multiple comparisons. Note that significance here is a feature of sub-networks and not individual connections. Sub-networks of corrected p-values < 0.05 were regarded significant. **C)** Schematic representation of the FC matrix layout.

The contrast between wakefulness and N2 sleep yielded a consistent pattern of N2-specific increases in the delta, alpha, sigma, and beta bands involving frontal, temporal, and somatosensory cortices. Compared to SWS, N2 showed widespread increases in the alpha, sigma, and beta bands. For the delta band N2 was increased mainly in connections involving the frontal cortices and midline areas, while the gamma band showed little evidence of differences. Significant increases were found in wakefulness compared to SWS most notably in the alpha and beta bands for connections involving posterior parietal and occipital areas.

Differences became re-scaled when partial correlations were considered. This effect was most striking for contrasts including N2, which showed fewer increased connections in the alpha, sigma, and beta bands, compared to the results for full correlations. The connection-wise correction generally found few significant differences, however at a threshold of $t > 2.2$ the NBS was able to identify several differential sub-networks across frequencies and contrasts. Highlighting some of these (shown in Figure 3-10), N2 still showed higher connectivity than wakefulness in and between temporal, frontal and somatosensory cortices, particularly in the alpha band, whereas the significant sub-network in the sigma band was dominated by long-range connections between frontal and occipital regions together with intra- and inter-hemispheric temporal connections. Interestingly, in the beta band, wakefulness now showed a sub-network of stronger connectivity than N2. This comprised occipital and posterior parietal regions, the posterior cingulate cortex, and the right temporal cortex. Similar regions in the posterior parts of the cortex were involved in the differential beta band sub-network between wakefulness and SWS. In the alpha band, the contrast between wakefulness and SWS presented a distinct sub-network of stronger fronto-temporal connectivity in SWS, akin to the one observed between wakefulness and N2 in the same band. In the alpha band, wakefulness was more strongly connected than SWS in an occipito-parieto-temporal sub-network. The contrast between N2 and SWS still showed rather global networks of increased connectivity in N2, albeit mostly sparser than for the full correlations. For the alpha, sigma, and beta bands; inter-hemispheric connections between temporal and parietal regions were consistently increased.

In summary, when using full correlations between frequency-specific and low-pass filtered envelopes, N2 showed large global FC increases in brain-wide networks compared to both wakefulness and SWS in all frequency bands but gamma. These differences were less pronounced when partial regularised correlations were used, especially in the alpha, sigma, and beta bands. The connections that remained higher in N2 mainly involved temporal, frontal, and somatosensory areas. In the beta band, using partial correlations, wakefulness showed a network of increased FC compared to N2 in and between posterior parietal and occipital regions. Compared to SWS, FC was increased in wakefulness for occipital and posterior parietal regions, and these differences were relatively resilient to whether full or partial correlations were used. For partial correlations, SWS showed increased FC compared to wakefulness in a sub-network comprising fronto-temporal connections in the alpha band.

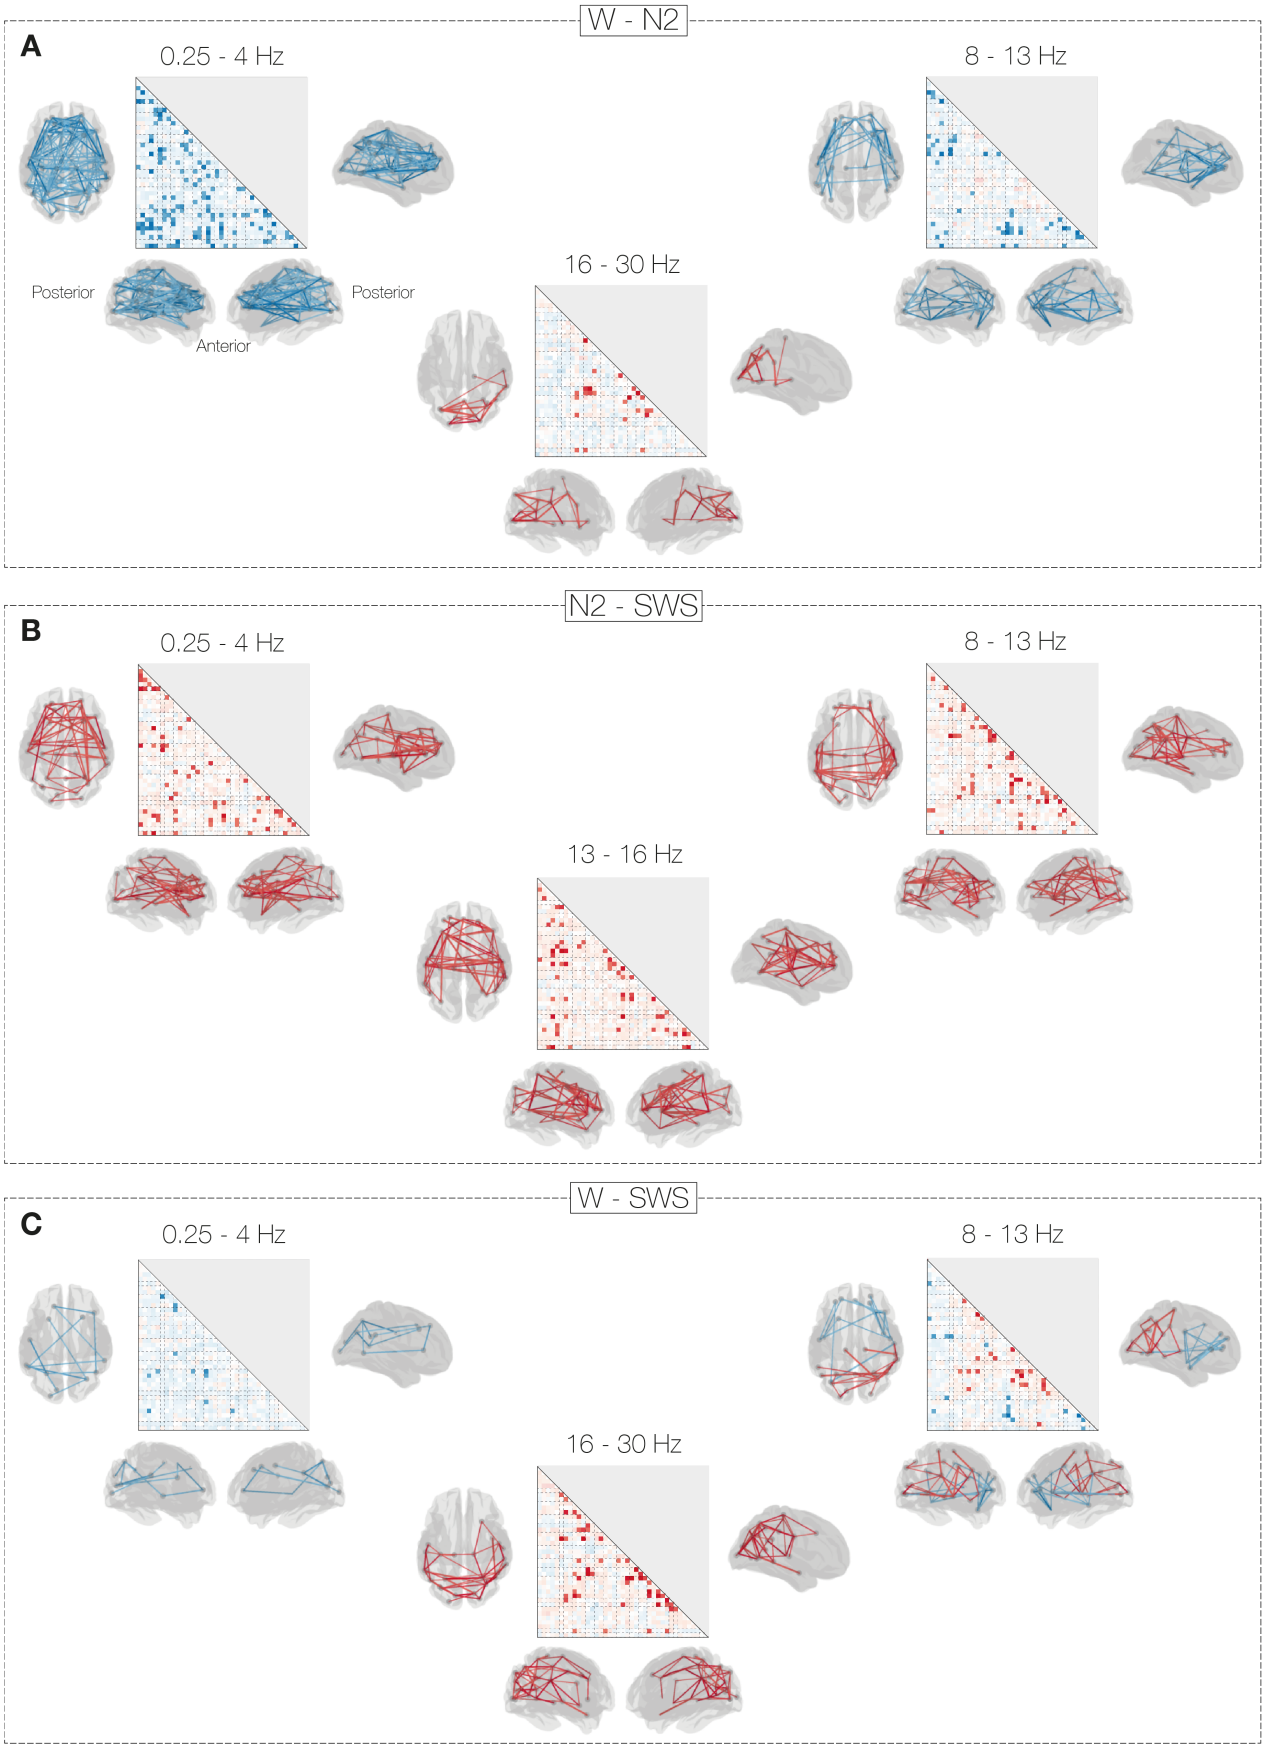

Figure 3-10 ***Highlights of significantly different sub-network identified with NBS*** **A)** Prominent differential networks of partial correlations for the contrast between wakefulness and N2 sleep. Note the widespread N2-specific increases (blue) in the delta frequency band. In the alpha band, N2-related increases were confined to a fronto-temporal network. Wakefulness-specific increases (red) were found in a posterior network in the beta band. **B)** Selected differential networks of partial correlations for the contrast between N2 sleep and SWS. FC was increased for N2 (red) in widespread networks, particularly involving the frontal and temporal cortices, but also more posterior regions in the

alpha band. **C)** Differential networks of partial correlations for the contrast between wakefulness and SWS. Note the posterior dominance of wakefulness-specific increases (red) in the alpha and beta bands. Note also the SWS-specific increases (blue) in a network connecting temporal and frontal regions.

3.4 ***Discussion***

To our knowledge, this chapter presents the first attempt to measure FC in sleep from source-reconstructed MEG recordings. Compared to fMRI, the use of MEG in the estimation of FC remains a less matured enterprise, both in terms of interpretations and applications. Employing analysis tools, recently validated for resting-state recordings, we have tested whether full and partial correlations between beamformed band-limited power amplitudes can be used to discriminate between periods of wakefulness, N2 sleep, and SWS. Finding evidence in favour of this for specific frequency bands, we went on to quantify the vigilance-dependent re-configurations in terms of global mean levels of FC and regional differences in connections and sub-networks.

3.4.1 *Vigilance-specificity of MEG FC networks; frequency contributions and sleep graphoelements*

FC patterns were particularly consistent in the alpha and beta bands during wakefulness, in the delta, alpha, and sigma bands during N2 sleep, and in the alpha band during SWS. This suggests that FC can indeed be detected with MEG as correlations between band-limited power fluctuations, not only during the wakeful resting-state as reported in a number of previous studies (Brookes et al., 2011a, Brookes et al., 2011c, Hipp et al., 2012, Wens et al., 2014, Colclough et al., 2015) (see Introduction, 1.5.7), but also in sleep. This is a non-trivial point, since MEG FC first and foremost has been devised to capture patterns in spontaneous resting-state activity.

Correlations between beamformed band-limited power envelopes are believed to reflect exchange of neuronal information through oscillatory coherence (Siegel et al., 2012, Engel et al., 2013). The spatial overlap with RSNs and large-scale FC networks known from fMRI stands as an important validation for the use of this MEG metric (Brookes et al., 2011a, Brookes et al., 2011c, Hipp et al., 2012, Luckhoo et al., 2012, Marzetti et al., 2013, Hipp and Siegel, 2015). A recent study compared band-limited power envelopes correlations (full and partial) to a range of alternative measures of resting-state MEG FC, and identified the former as outperforming the latter in terms of intra- and inter-participant repeatability and consistency (Colclough et al., 2016). Our observation that FC patterns change in manners specific to certain vigilance conditions adds an important element of verification to the use of this particular FC metric.

The higher specificity of network patterns in the alpha and beta bands during wakefulness, reported here, resonates well with previous MEG FC resting-state studies. As mentioned, spatial overlap has been found between fMRI-derived RSNs and power envelope correlations in these bands (Brookes et al., 2011b). Since then, studies have consistently reported MEG FC to be concentrated in these frequencies (Hillebrand et al., 2012, Hipp et al., 2012, Luckhoo et al., 2012, Schölvinck et al., 2013, Hall et al., 2014, Colclough et al., 2015, Hipp and Siegel, 2015). A particularly important reference for our specificity analyses are the results reported in Colclough et al. (2016) for inter-participant consistency of the FC measures used in our analyses, i.e. leakage-corrected band-limited power amplitude correlations, applied to resting-state data. In line with

our results for within-wakefulness FC specificity, Colclough and colleagues reported high consistency for the alpha and beta bands, followed by theta and gamma, while the delta band showed the lowest levels.

Hillebrand and colleagues investigated the influence of source power on MEG FC values, as estimated with the Phase Lag Index (Stam et al., 2007) between beamformed timecourses (Hillebrand et al., 2012). While a positive linear relationship was identified across most of the included frequency bands (delta, theta, alpha, beta), source power only explained a minor proportion of the estimated FC (Hillebrand et al., 2012). The potential bias that MEG FC distributions across the brain could be driven solely by a heterogeneous signal-to-noise ratio (SNR) of a given oscillation was thus rejected (Hillebrand et al., 2012). This is a relevant point for our purpose, given the marked changes in spectral content commonly observed across frequencies from wakefulness to sleep (see Figure 3-2). To make sure that the vigilance-specificities observed in our FC estimations were not merely effects of inter-condition differences in spectral compositions and source-SNR, we re-performed our specificity analyses on source-power distributions. Our results clearly show that source-power distributions were quite stable across vigilance conditions and, unlike the FC network patterns, held little discriminatory information.

3.4.2 *The role of sleep graphoelements in FC patterns*

The fact that N2 sleep expressed highest FC specificity in the delta, alpha, and sigma bands could point to an involvement of phasic sleep oscillations, such as K-complexes and sleep spindles, known as hallmarks of this sleep stage (see Introduction, 1.2, 1.3). Unlike most resting-state MEG FC studies, we included a separate frequency band, sigma, for the purpose of capturing spindle-related activity. Yet, spindles are known to vary in terms of frequency and are commonly grouped into 'fast' (13 – 16 Hz), and 'slow' (11 – 13 Hz) (Werth et al., 1997, Mölle et al., 2002, De Gennaro and Ferrara, 2003). Even if the group-averaged sensor-periodograms showed a distinct N2-specific peak between 13 and 16 Hz (see Figure 3-2), it is conceivable that latent slower spindles could have driven envelope correlations in the alpha band.

The N2-specific FC networks and degree maps for the alpha, sigma, and beta further support the notion that spindles constitute a main contributor to the estimated envelope correlation patterns. Based on scalp topographical EEG evidence 'fast' and 'slow' spindles are traditionally described as expressing posterior and anterior predominance, respectively (Werth et al., 1997, De Gennaro and Ferrara, 2003). The higher spatial detail that has been provided by fMRI studies of BOLD activity time-locked to EEG-derived spindles (see Introduction, 1.5.2) have implicated the anterior cingulate, the thalamus, and the superior temporal cortex as generators common to both 'fast' and 'slow' spindles (Schabus et al., 2007), with more participation of parietal areas in the 'fast' ones (Schabus et al., 2007, Andrade et al., 2011). In line with this, the degree maps for N2 sleep in Figure 3-7 show a trend of increased mean FC in parietal regions, and decreased mean FC in medial frontal regions from the alpha, over the sigma, to the beta band.

While topographical EEG and hemodynamic fMRI studies have indicated widespread co-occurrences of spindles, intra-cortical electrophysiology have established that spindles can be expressed locally with little or no inter-regional cross-talk (Andrillon et al., 2011, Nir et al., 2011, Johnson et al., 2012, Nobili et al., 2012). Simultaneous EEG/MEG studies employing source-reconstruction of spindles have shown that these

modalities vary in their sensitivity to spindles, and that MEG more readily picks up spatially asynchronous and focal spindles (Dehghani et al., 2010b, Dehghani et al., 2011). It has been hypothesised and tested with thalamocortical modelling that this modality-dependent variability and the dichotomy between local and widespread spindles can be attributed to two different pathways of generation, involving the core (specific) and matrix (non-specific) thalamic nuclei, respectively (Bonjean et al., 2012, Piantoni et al., 2016a). Apart from the 'fast' versus 'slow' categorisation, a distinction between local and widespread properties is thus important for a full characterisation of sleep spindles. Interestingly, a recent study used intra-cortical recordings with high spatial electrode-coverage of eight epileptic patients to show that local spindles could be identified over most of the cortical mantle, while widespread co-occurrences of spindles were most likely to be found between the lateral prefrontal and the superior temporal cortices (Piantoni et al., 2016b). The latter observation is very much in line with the involvement of these regions in the FC patterns we report for N2.

K-complexes, the other main graphoelement of N2 (see Introduction, 1.3.1), might have been driving the specificity of the FC networks for N2 the delta band. Unlike for sleep spindles, little agreement exists between the fMRI and the EEG/MEG literature in terms of the time-locked correlates of K-complexes (for a discussion, see Wennberg and Cheyne (2013) and Jahnke et al. (2012)). As such, K-complex-related BOLD activity has been found in many regions, including the auditory, and visual cortices, portions of the cingulate, posterior parietal regions, prefrontal cortex, and the thalamus (Caporro et al., 2012, Jahnke et al., 2012). In contrast, previous electrophysiological evidence from source-localisation of EEG and MEG as well as intra-cortical recordings have mainly implicated frontal regions (Wennberg, 2010, Wennberg and Cheyne, 2013). Our analyses revealed MEG FC in the delta band as focussed towards lateral, medial frontal and superior temporal regions in a pattern quite similar to the one observed for the alpha and sigma bands. Once again, we may refer to a more recent intra-cortical study showing that K-complexes can be identified in most cortical regions, while a particular preponderance of co-occurrence was reported between frontal and temporal regions (Mak-McCully et al., 2015).

K-complexes have been likened to isolated slow oscillations (Cash et al., 2009) (see Introduction, 1.3.1). Specifically, K-complexes are thought to reflect the DOWN-state of the slow oscillation (Cash et al., 2009). In this respect, it is interesting that FC networks from the delta band in SWS showed little specificity. This suggests that the continuous slow oscillations and waves, which dominated during SWS, were less able to produce consistent correlations between envelope fluctuations compared to K-complexes in N2 sleep. This may be due to the 'travelling-wave' character of slow oscillations, which has been established through scalp topography and source-reconstruction of EEG data (Massimini et al., 2004, Murphy et al., 2009b). It is likely that envelope correlations may be a sub-optimal measure for capturing this kind of lagged propagation.

The most distinct and specific FC for SWS was found in the alpha band. This may appear surprising, given the general view that higher frequencies are suppressed in this sleep stage (see Introduction, 1.2.1 and Figure 3-2). Meanwhile, latent higher frequency activities during SWS have been a known phenomenon for many years. Soon after the discovery of the slow oscillation, it was shown that neuronal UP-states are associated with bursts of fast spiking in frequencies including the alpha range (for review see Steriade

(2006)). At the extra-cranial EEG level, alpha activity during SWS was previously thought to reflect pathology under the common term 'alpha-delta' sleep (Hauri and Hawkins, 1973, Moldofsky, 1989). Since then, alpha activity during SWS has been described as a common characteristic of the human sleep EEG (Scheuler et al., 1983, Pivik and Harman, 1995), while mechanisms of its generation remain poorly illuminated compared to those of the other sleep-specific EEG signatures (Olbrich et al., 2011). Once again, sleep spindles might play a role, since slow spindles, in particular, have been identified in SWS (Rasch and Born, 2013). Of particular interest for our FC findings, are a number of reports on sensor-level coherence of EEG signals during SWS (Achermann and Borbély, 1998, Duckrow and Zaveri, 2005). Common to these studies are indications of high levels of coherence in the alpha frequency range, in particular over frontal derivations. This fits well with our FC estimations for SWS with high values between frontal and temporal regions.

3.4.3 *Comparative analyses of wake and sleep FC; global and regional differences*

The vigilance-specificities of the FC networks, discussed in the previous section, were based on similarities between overall pattern structures. This level of analysis tells us *that* FC values are distributed differently between vigilance conditions across the whole network, but not *how* the networks differed between conditions, or whether FC strengths changed. Importantly, the vigilance-specificity test was devised to estimate the consistency of the full FC networks, and could therefore not be used to rule out potentially consistent sub-patterns. Using global and regional comparative network analyses we were able to characterise such differences.

Beginning with the global differences, our results clearly state that mean levels of FC varied across vigilance conditions (see Figure 3-8). N2 sleep showed significantly higher values than SWS over all frequency bands except gamma. This was also true for the contrast between wakefulness and N2, however only when full correlations were used. Interestingly, when partial correlations were considered, N2 was only higher than wakefulness in the delta band. In contrast, the significantly stronger FC of wakefulness compared to SWS in the beta band was found for both full and partial correlations. This suggests an N2-specific re-scaling of FC values in the theta, alpha, sigma, and beta band, when indirect connections were removed. Keeping with the interpretation that sleep graphoelements as main determinants of the FC networks derived from N2, it would seem that sleep spindles in the alpha to beta range produce more indirect envelope correlations than K-complexes in the delta band. An explanation for this phenomenon may be found in how these transient sleep oscillations are generated. As such, spindles are known to arise from thalamocortical loop interactions (Luthi, 2014, Piantoni et al., 2016a), while K-complexes and slow oscillations are believed to depend on the thalamus in a less strict sense (Steriade, 2006, Cash et al., 2009) (although see David and Schmiedt (2014), Lemieux et al. (2014), Crunelli et al. (2015)). In particular the widespread spindle volleys orchestrated by the non-specific matrix nuclei of the thalamus could well have given rise to indirect connections between the cortical ROIs. Conversely, K-complex-dependent connections may have been less shared between ROIs due to a relatively smaller thalamic contribution.

Our findings of frequency-wide N2-specific increases in envelope correlations are in line with the sensor-space study by Liu et al. (2010) (see Introduction, 1.5.7, and Figure 1-14E). Here, significant increases in coherence and correlation between power fluctuations were reported for 'light sleep' compared to wakefulness. 'Light sleep' was defined on the basis of the spectral ratio between alpha and delta power, and most likely included a combination of N1 and N2 sleep. Although the authors did not discuss the involvement of sleep-specific oscillations, the perspectives offered above apply to their results as well. Including separate data from N1 sleep should give more insight in this respect, and this represents an important next step for MEG FC exploration.

The extent to which power envelope correlations in MEG parallel slow BOLD correlations in fMRI remains an unresolved matter. As mentioned, out of the frequencies investigated with MEG FC, the alpha and beta bands have shown the highest correspondence with resting-state fMRI FC (see e.g. Hipp and Siegel (2015)). It is not immediately obvious, however, that the same frequency-translation would apply to the case of sleep data, since the amount of alpha and beta oscillations are generally decreased during sleep (see Figure 3-2). In line with this, Tagliazucchi and colleagues showed that the electrophysiological basis of fMRI FC networks varies with vigilance, such that the relative contributions of EEG frequencies change from wakefulness to early N1 sleep (Tagliazucchi et al., 2012c). Comparing the vigilance-specific MEG FC networks of the present chapter (see Figure 3-4) to the equivalent fMRI FC networks from Chapter 2 (see Figure 2-1) it is clear that in particular SWS showed much more preserved FC in fMRI. In general, FC appeared more diverse in MEG across vigilance conditions, while the fMRI-derived networks seemed to suggest a high degree of shared, and thus vigilance-independent connectivity. While it would require fMRI and MEG networks derived from equivalent parcellations to confirm, such as in Hipp and Siegel (2015) and Tewarie et al. (2016a), it seems unlikely from our results that the employed MEG FC methods would be able to capture the same robust and shared FC as observed in fMRI.

Yet, what is interesting to note is that our regional comparative MEG FC analyses yielded differential networks (Figure 3-10), which appeared to share some of the characteristics that we and others have observed when deriving equivalent differences in fMRI (compare to Figure 2-2 in Chapter 2, and Figure 1-10 in the Introduction). As such, when partial correlations were used, we found wakefulness to be associated with increased FC in a sub-network involving occipital, parietal, and temporal ROIs, relative to SWS in the alpha and beta bands, and relative to N2 in the beta band. This is in line with the wakefulness-specific increases we reported in Chapter 2 (compare to Figure 2-2, the 'W–SWS' and 'W–N2' contrasts). The N2-specific results appear similarly cross-modal, with an abundance of fronto-temporal and fronto-parietal connections increased relative to wakefulness (compare to Figure 2-2, the 'W-N2' contrast) in the delta, alpha, and sigma bands, and more widespread increases relative to SWS (compare to Figure 2-2, the 'N2–SWS' contrast). While not evident in the fMRI dataset analysed in Chapter 2, the fronto-temporal network of increased FC in SWS compared to wakefulness bears resemblance to other fMRI reports for this contrast (Spoormaker et al., 2010, Altmann et al., 2016).

3.4.4 *MEG power envelope correlations as a use-dependent signature of the functional connectome*

Based on the above it would seem that while our MEG-derived FC networks might have failed to reproduce the vigilance-robust part of the fMRI-FC, we were in fact able to capture many of the inter-condition differences. One interpretation of this could be that MEG is able to resolve the sources of differences between fMRI FC networks, which may unfold 'on top' of a more robust baseline of connectivity. The frequency-specific MEG FCs may thus represent some form of use-dependent activity, which drive the more subtle differences observed in the fMRI-derived FC. This would imply some underlying BOLD FC generator capable of producing a backbone network, on top of which the power envelope correlations become incorporated. In the Introduction it was mentioned how slow cortical potential (SCP) have been suggested as a potential electrophysiological basis for large-scale FC patterns of fMRI (see 1.5.7 and He and Raichle (2009)). Given that these very slow dynamics are difficult to extract from extra-cranial recordings such as MEG, it is possible that they account for a baseline network of BOLD correlations, insensitive to changes in vigilance, and concealed from our band-limited power envelope analyses. An alternative scenario, supported by a recent empirical study (Tewarie et al., 2016a) and a recent whole-brain computational modelling study (Deco et al., 2016a), is that the individual frequency bands, i.e. delta, theta, alpha, sigma, beta, and gamma, have to be modelled in combination to explain the fMRI FC patterns (Tewarie et al., 2016a). A multimodal dataset of both fMRI and MEG from sleeping participants, as well as an inclusion of data from N1 sleep could aid in the further testing of these different scenarios.

3.4.5 *Functional significance of vigilance-dependent FC networks*

We have argued for the involvement of K-complexes and sleep spindles in the generation of N2-specific FC increases. Hence, we may refer to the many proposed roles of these sleep-waveforms when interpreting the FC differences that we have observed. As it stands, most sleep oscillations as well as individual sleep stages have to some extent been implicated in sleep-dependent learning and memory consolidation (see Introduction, 1.4.3). Both spindles and SWS have received considerable attention in this respect, and shown to correlate with overnight improvement of declarative memory (for review, see Rasch and Born (2013)). The fronto-temporal sub-network of increased FC that we observed for N2 and SWS in the alpha band could represent large-scale substrate of such memory-related processes. Indeed this would be in line with the view of NREM sleep as particularly beneficial for the transfer of short-term, hippocampus-dependent memories to long-term storage in neocortical areas (Diekelmann and Born, 2010), where the coupling between medial temporal and prefrontal regions are believed to play a crucial part (Simons and Spiers, 2003) (see the discussion of Piantoni et al. (2016b)). The finding of this network in the alpha band for SWS is particularly interesting in light of recent spectral EEG evidence pointing to increased alpha activity during SWS as a specific marker of overnight performance-improvement on word-pairing (Marshall et al., 2006) and gain of explicit knowledge of hidden rules in number sequences (Yordanova et al., 2012).

The increased inter-hemispheric connectivity between temporal regions observed in N2 compared to SWS could be related to the putative gating mechanisms of spindles for sensory information processing that

has been demonstrated for N2 specifically using auditory stimuli and fMRI (Dang-Vu et al., 2011, Schabus et al., 2012).

The functional role of the K-complex has been subject to a long history of ideas (Colrain, 2005). Two opposing views regard K-complexes respectively as sleep protective (De Gennaro et al., 2000, Colrain, 2005, Jahnke et al., 2012) and as signs of arousal (Halász, 2005). The on-going debate is rooted in the observation that this large graphoelement can occur spontaneously as well as in response to external stimulation (Colrain, 2005, Halász, 2005). Our finding of wide-spanning FC increases in the delta band of N2 suggests that K-complexes were particularly effective at bringing the cortical network into a synchronised state. This would appear to favour the sleep protective view, since a synchronised network would be expected to limit incorporation of external information, and as such be non-conducive to arousal.

The sub-networks of wakefulness-specific FC increases, in the alpha and beta bands compared to SWS and in the beta band compared to N2, consistently involved areas in the posterior part of the cortex. Interestingly, encompassing occipital, temporal, and posterior parietal regions, this network shares considerable overlap with a 'posterior cortical hot zone', which through a theoretical synthesis of content-specific and state-based evidence was recently proposed as a neural correlate of consciousness (Koch et al., 2016).

3.4.6 *Methodological considerations*

The lack of N1 sleep in the present dataset represents a limitation to the interpretations of the presented results. Evidence from this sleep stage would have been useful in formulating the link to the fMRI results in Chapter 2. In particular, it would be interesting to test if the N1-specific subcortico-cortical dropout can be detected with MEG. The SNR of subcortical regions is low in MEG, which was why we chose a cortical parcellation. A number of studies are however suggesting that it is in fact possible to extract sensible information from deeper structures (Attal et al., 2012, Roux et al., 2013). A finding of decreased subcortico-cortical FC with MEG in N1 would represent important affirmation in this regard. We limited our analyses to wakefulness, N2, and SWS due to our inclusion criteria (see 3.3.2). A careful concatenation of non-contiguous segments from the current dataset may yield enough time-points from N1 to extract sensible FC measures.

We used approximately 300 seconds of each vigilance condition for each participant. This is in line with most MEG resting-state studies. Yet, the inclusion of more data, through concatenation, could potentially make FC estimates more robust, especially for the lower frequency bands.

The lack of individual structural information for the participants of the current dataset will most likely have limited the accuracy of the applied source-reconstruction. The same may be said of the minimal number of anatomical reference points (only three fiducials compared to the hundreds of digitized head position points used in more contemporary acquisitions). It is important to note, however, that the potential localisation bias would not have been the same had the FC analyses been performed in the sensor space. In the end we would argue that source-reconstruction, together with the employed leakage-correction, was a

sensible strategy, as this is recognised as an important un-mixing step promoting interpretation of MEG FC (Schoffelen and Gross, 2009, Palva and Palva, 2012).

Chapter 4
Large-scale cross-frequency coupling in sleep using MEG

4.1 *Introduction*

Source-reconstruction of MEG data offers the opportunity to resolve frequency-specific information regarding the large-scale networks of inter-regional interactions in the brain. In Chapter 3 it was shown how the extraction and correlation of power amplitude envelopes could be used to complement sleep-related findings from fMRI, and to label vigilance-dependent changes in FC in the electrophysiological frequency domain. While promoting a more comprehensive account with a closer link to neural activity (see Introduction, 1.5.7), the increased frequency resolution also presents the challenge of handling an expanded space of information. Specifically, one has to decide on a strategy to structure the temporal richness of the electrophysiological data. The parting of recorded oscillations into separate bands of frequencies dates back to Berger, who introduced the use of the Greek letters alpha and beta to denote recorded activities (Berger, 1929, Buzsáki and Wang, 2012). More bands have been added since then, such as the delta, theta, sigma, and gamma, however the exact demarcations and definitions of these vary greatly throughout past and present literature. A common consensus is that these bands should reflect the generative mechanisms of their oscillations (Buzsáki and Wang, 2012). Even if these mechanisms are yet to be fully uncovered (see e.g. Introduction, 1.3.1), there are empirical grounds for asserting that the suggested frequency bands serve separate computational roles in the brain (Lopes da Silva, 2013). Chapter 3 may be said to follow in this tradition, along with most previous reports on MEG FC, by restricting analyses to within-frequency interactions.

Meanwhile, it is becoming increasingly recognised that oscillations from separate frequency bands can affect each other, through cross-frequency coupling. This inter-frequency communication has been shown to correlate with various cognitive processes (Palva et al., 2005, Canolty et al., 2006, Osipova et al., 2006, Jensen and Colgin, 2007, Canolty and Knight, 2010). A similar phenomenon is well-described in the specific case of NREM sleep, where intra-cortical as well as scalp EEG and MEG evidence speak for the grouping of faster oscillations by the ubiquitous slow oscillation (Steriade et al., 1993b, Contreras and Steriade, 1995, Clemens et al., 2007, Le Van Quyen et al., 2010, Cox et al., 2012, Cox et al., 2014, Klinzing et al., 2016). In a similar vein, a leading hypothesis for sleep-dependent memory consolidation proposes an important role for interactions between sleep spindles and fast hippocampal ripples (Diekelmann and Born, 2010, Rasch and Born, 2013).

There is thus good reason to assume that a within-frequency view is too restrictive to capture how sleep modifies FC networks. It is only very recently, however, that MEG studies have started to explore the role of cross-frequency interactions at the large-scale network level during the resting-state (Florin and Baillet, 2015, Brookes et al., 2016, Tewarie et al., 2016b). With the specific aim of expanding the use of amplitude envelope correlations for MEG FC estimation, Brookes and colleagues introduced the concept of a multi-layer network (Brookes et al., 2016, Tewarie et al., 2016b). Considering the traditional within-frequency ROI

networks of envelope correlations as intra-frequency layers (e.g. delta-to-delta, alpha-to-alpha, and beta-to-beta), so-called inter-layers were incorporated, consisting of correlations between ROI envelopes from different frequency bands (e.g. delta-to-alpha, alpha-to-beta, beta-to-delta).

This multi-layer network approach thus offers a comprehensive overview of the multi-dimensional MEG data of cortical regions and their spatio-temporal power envelope interactions across frequency bands. Theoretically it is motivated by insights from physics and network theory (for details, see Tewarie et al. (2016b) and Brookes et al. (2016)). It has thus been demonstrated that a network of units (here: cortical regions), connected by a number of qualitatively different layers of communication (here: power fluctuations of discrete frequency bands) can be fully modelled only when possible interactions between these layers are taken into account. An analogy to this scenario is found in radio broadcasting, where various channels will be transmitting different content to the listeners. Some channels will be conveying debates or cultural magazines, others will be playing music ranging from the classical to the popular, such that each channel delivers material independent of each other. However, if some unexpected event of significant public interest takes place, many of the channels may be interrupted by a news update, whereby the contents of the channels become overlapping, thus illustrating their interdependence.

In Brookes et al (2016) the multi-layer network approach was shown sensitive to task as well as differences in resting-state dynamics between a group of healthy participants and a population of schizophrenic patients. Tewarie and colleagues provided evidence for a common mode of connectivity spanning both intra- and inter-layers of frequencies, hypothesising that this shared pattern reflect efficient integration between the frequency-specific networks and computations (Tewarie et al., 2016b).

A multi-layer approach represents a natural next step for analysing the source-reconstructed MEG power envelopes extracted in Chapter 3. Employing this pan-spectral interrogation we are able to ask whether the identified FC differences are restricted to within frequency layers, or if inter-layer cross talk is also vigilance-sensitive. We test for the existence of a common connectivity mode and probe the multi-layer integration and segregation across vigilance conditions.

4.2 ***Methods***

4.2.1 *Cross-frequency correlations and the super-adjacency matrix*

In order to explore cross-frequency relationships in the source-reconstructed MEG data, we followed a method introduced in two recent studies by Brookes et al. (2016) and by Tewarie et al. (2016b). Instead of restricting the correlation estimates within frequency bands, as in Chapter 3, we correlated envelopes from all bands in an all-to-all manner.

We re-used the 39 ROI envelopes extracted for the analyses in Chapter 3, in the same six frequency bands: delta (0.25-4 Hz), theta (4-8 Hz), alpha (8-13 Hz), sigma (13-16 Hz), beta (16-30 Hz), and gamma (30-48 Hz). Each downsampled ROI envelope from a given frequency band was thus correlated to 39 other envelopes from the remaining five frequency bands. For each vigilance-specific session this produced 39x5 correlation values, or 5+4+3+2+1 = 15 unique cross-frequency connectivity matrices. An overview of the workflow is given in Figure 4-1, and the cross-frequency FC estimation is illustrated in Figure 4-1A.

Following the notation in Brookes et al. (2016), and by tiling the cross-frequency matrices, a so-called super-adjacency matrix was produced for each vigilance-specific session (see Figure 4-1B). The super-adjacency matrix comprised 6x6 block matrices or tiles. Tiles along the diagonal of the super-adjacency matrix included within-frequency correlations, and may be referred to as intra-layer FC matrices (Tewarie et al., 2016b), and corresponded to the full correlation FC matrices presented in Chapter 3 (see Figure 3-4). The off-diagonal tiles of the super-adjacency matrix comprised cross-frequency correlations, and each off-diagonal tile may be referred to as an inter-layer FC matrix (Tewarie et al., 2016b). As an example, an (alpha, beta)-tile (i.e. an inter-layer) was composed of correlations between downsampled envelopes from the alpha and the beta bands. Unlike the within-frequency correlation matrices (i.e. intra-layers) the inter-layer matrices were not symmetrical around their diagonals, since a correlation between ROI *a* from the alpha band and ROI *b* in the beta band, does not necessarily correspond to the correlation between ROI *a* from the beta band and ROI *b* from the alpha band. Further, the diagonal of an inter-layer matrix would not be a vector of ones, but include the within-ROI correlations between the two frequencies. The super-adjacency matrix itself was symmetrical around its own super-diagonal, since for instance the (alpha, beta)-tile was a mirror image of the (beta, alpha)-tile.

4.2.2 *Vigilance-dependent changes in cross-frequency correlations*

As in Chapter 2 and Chapter 3, we tested for both global and regional differences between vigilance conditions. Once again we used paired t-tests in a GLM framework and corrected for multiple comparisons with FSL randomise (see Chapter 2, 2.3.3). For the connection-wise differences, comparisons were made within the cross-frequency tiles of the super-adjacency matrix. The correlations were Fisher z-transformed prior to the tests. For the differences in mean layer-FC we simply computed the mean value of each tile and used these as observations. The design matrix was permuted 5000 times for each test, and differences with a corrected p-value below 0.05 were regarded as significant.

It was clear that negative correlations were much more common between envelopes from differing frequency bands, compared to what we had observed in our within-frequency analyses. In order to determine whether negative correlations were driving potential differences between vigilance conditions, we repeated the comparative analyses, using the absolute of the correlation values as observations.

4.2.3 *Interactions between intra- and inter-layer FC*

Tewarie and colleagues proposed a measure of interactions between intra- and inter-layer FC matrices (Tewarie et al., 2016b). By correlating the patterns of within-frequency and cross-frequency matrices in resting-state recordings it was shown that the pattern of a given inter-layer was highly dependent on the FC patterns within each of the frequency bands in this inter-layer combined. It was hypothesised that this similarity between network patterns from intra-layers and inter-layers might reflect a process of integration occurring throughout the multi-layer network (Tewarie et al., 2016b). As such, while each layer (intra as well as inter) may express individual and specific signature patterns, with potential relevance to diverse functions, a common mode of connectivity might aid the integration between the layers and their functions.

We wanted to test if such a common mode was present in our data, and whether interactions between layers would be affected by changes in vigilance. For each intra-layer matrix, along the diagonal of the super-adjacency matrix, we thus computed the Pearson's correlation to each inter-layer matrix of the same tile column. The procedure of layer pairings is sketched in Figure 4-1C. As in Tewarie et al. (2016b) we computed correlations after taking the absolute of the intra- and inter-layer elements. This was done for super-adjacency matrices from each participant for each vigilance condition. The interactions between intra- and inter-layer matrices were compared across vigilance conditions using paired t-tests, in the same way as for the mean cross-frequency correlations above.

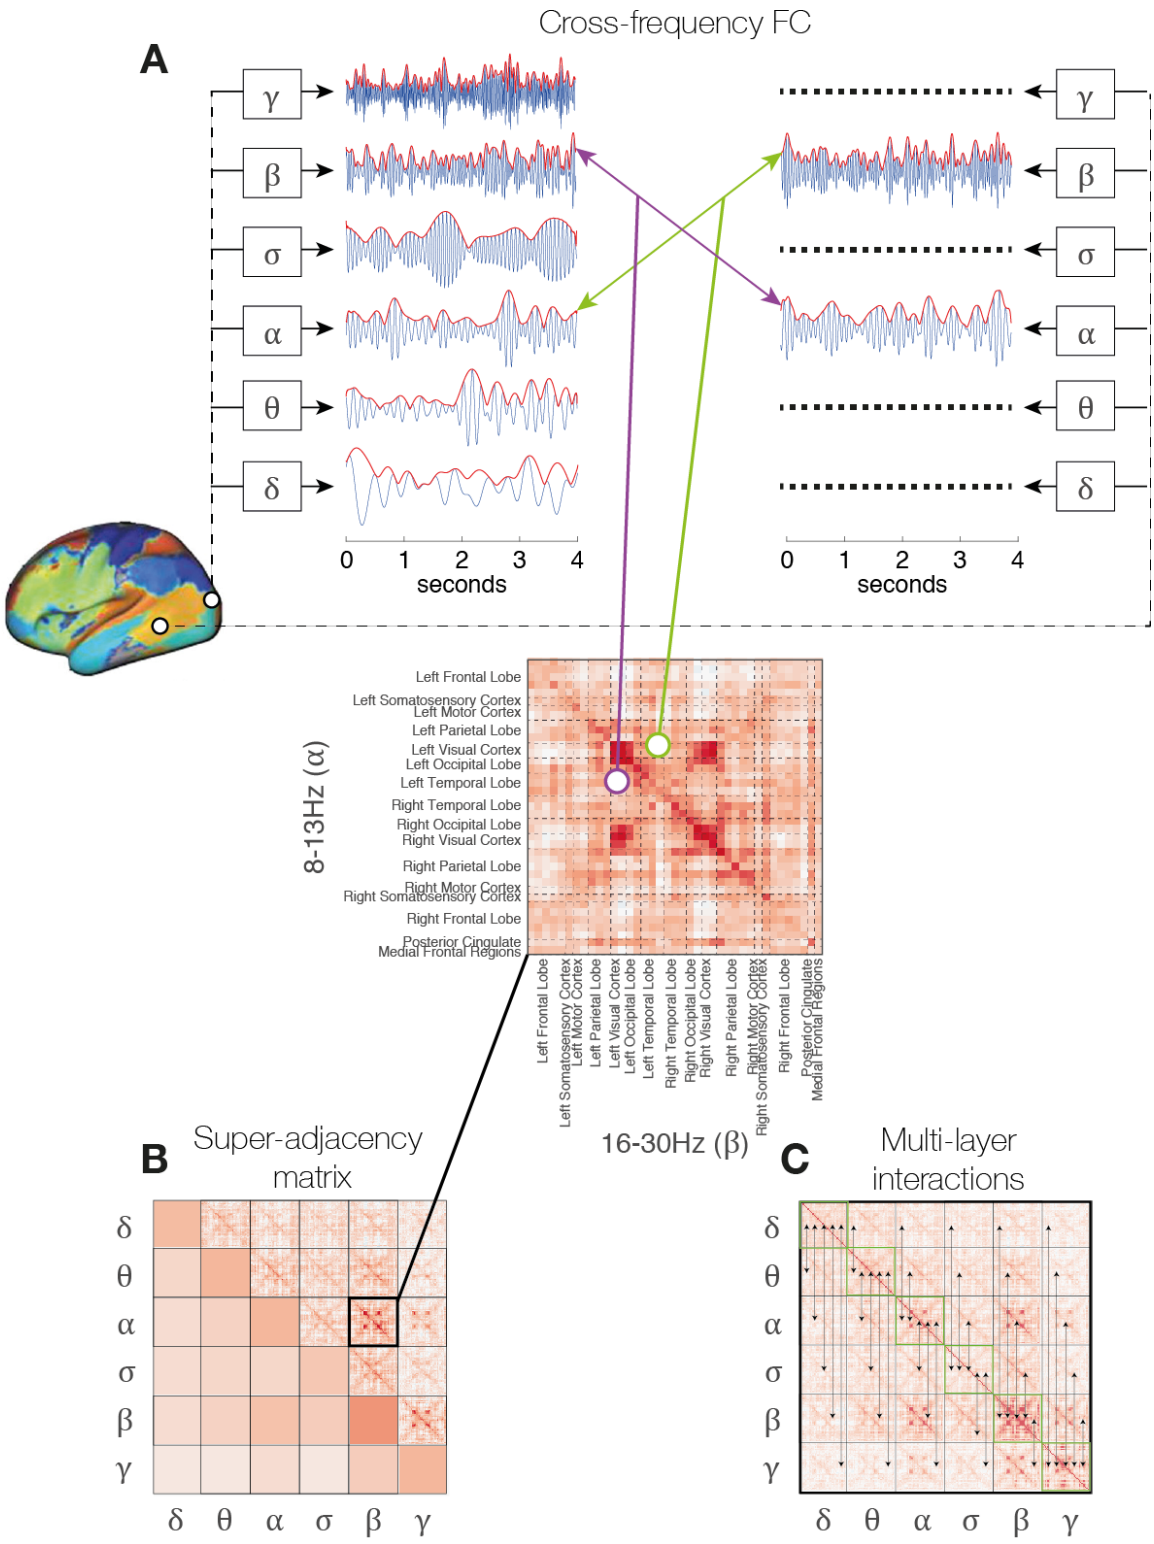

Figure 4-1 ***Deriving cross-frequency FC and multi-layer interactions in the super-adjacency matrix*** **A)** The beamformed timecourses (plotted in blue) of each of the six frequency bands were used to correlate amplitude envelopes (plotted in red) across frequencies. Cross-frequency correlations (FC) made up the inter-layer matrix. An example is shown for a visual and a temporal ROI and correlations between the alpha and beta band. Note how the two ROIs correspond to entries on each side of the diagonal of the inter-layer FC matrix. **B)** The inter-layer FC matrix makes up a tile of the super-adjacency matrix, shown here. **C)** Multi-layer interactions were estimated as the Pearson's correlation between intra-layers and inter-layers, as depicted by the arrows.

4.3 ***Results***

4.3.1 *Group-averaged super-adjacency matrices*

The super-adjacency matrices resulting from our cross-frequency analyses within vigilance conditions are shown in Figure 4-2A. As explained in Figure 4-2B, the super-adjacency matrices are presented with cross-frequency tiles, i.e. inter-layer networks, in the upper triangular part, and mean values of the inter-layers in the corresponding tiles of the lower triangular part. Since the super-adjacency matrices were symmetric around their diagonals, no information was lost in this representation. For reference, we included the mean of the within-frequency correlations, i.e. intra-layers, in the diagonal tiles of the super-adjacency matrix. Note, that these values correspond to the ones presented in Figure 3-8A. In Figure 4-2C are highlighted some of the prominent inter-layer networks for each of the three vigilance conditions.

For wakefulness we observed strong cross-frequency correlations between envelopes of the alpha and beta bands, and this inter-layer also showed the highest mean value of cross-frequency correlations. Strong interactions were also seen between the beta and gamma bands, and between the alpha and gamma bands, albeit to a lesser extent. A common network structure was evident in most of the inter-layers of the super-adjacency matrix, and these cross-frequency tiles were close to symmetrical around their diagonals. This close-to-symmetric matrix structure is seen in the highlighted (alpha, beta)-layer and (alpha, gamma)-layer for wakefulness in Figure 4-2C; the common structure included relatively high correlations between visual and parietal regions.

The super-adjacency matrix of N2 showed generally elevated levels of cross-frequency correlations, except in tiles involving the gamma band. Signs of a common structure in the inter-layer networks were mainly evident between neighbouring frequency bands, and scarcely so for inter-layers of the gamma band, as exemplified in the highlighted cross-frequency-tiles for N2 in Figure 4-2C. The common structure was of a different configuration than the one in wakefulness and more focussed around somato-sensory and frontal cortices.

Compared to the two other vigilance conditions, SWS showed less distinct cross-frequency interactions in general throughout the super-adjacency matrix. Interestingly, we found an abundance of negative cross-frequency correlations between the gamma band and all other frequency bands except beta.

4.3.2 *Differences between vigilance-specific super-adjacency matrices*

The results for the vigilance-dependent differences in cross-frequency interactions are presented in Figure 4-3**.** We investigated both the raw (Figure 4-3A) and the absolute values (Figure 4-3B) of the cross-frequency correlations. The regional, connection-specific differences are shown in the upper triangular part of each differential super-adjacency matrix, while differences between the global mean values are shown in their lower triangular parts.

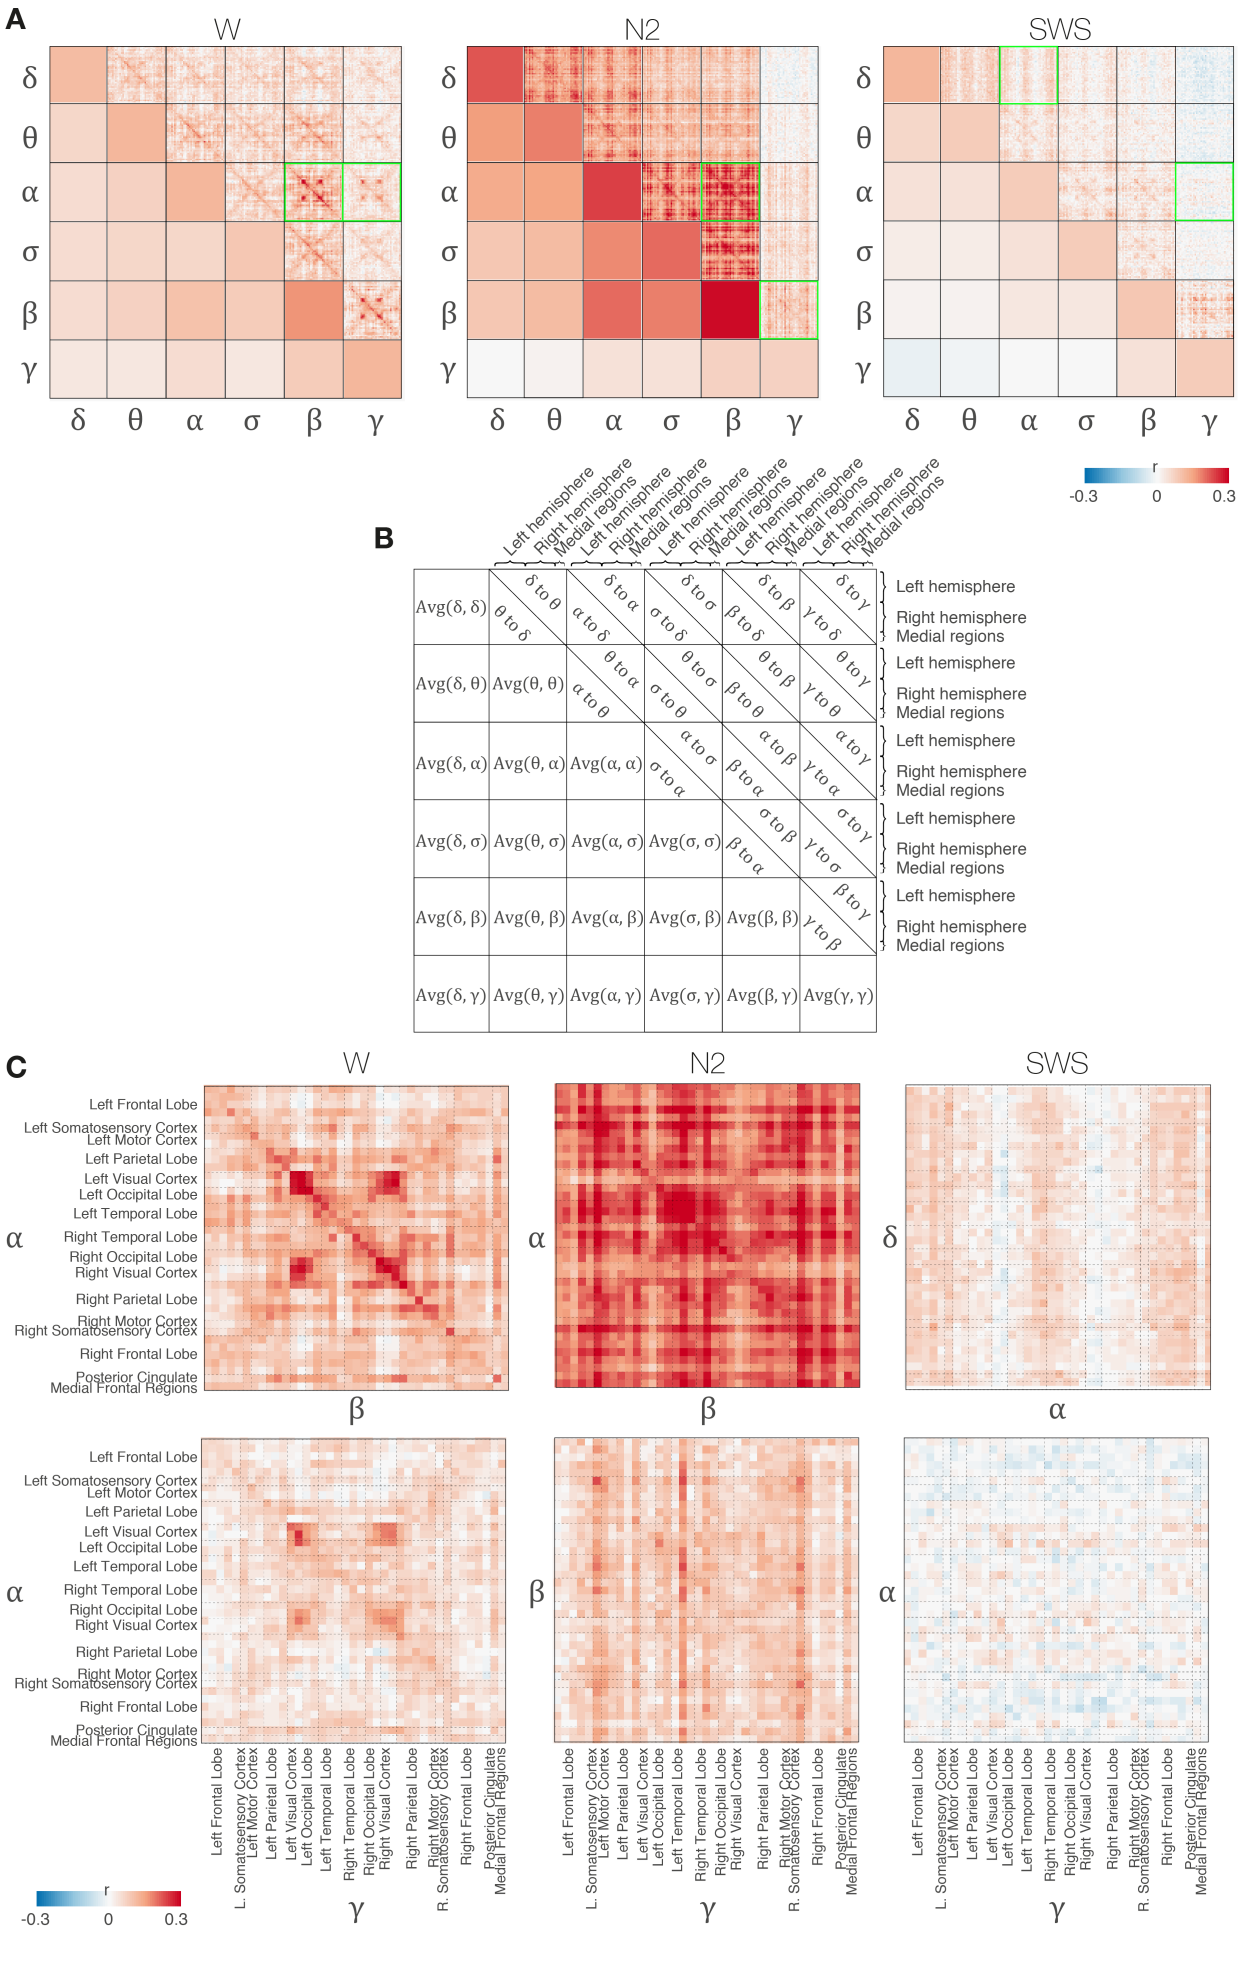
A
W
N2
SWS
δ
θ
α
σ
β
γ
r
-0.3
0
0.3
B
Left hemisphere
Right hemisphere
Medial regions
Avg(δ, δ)
Avg(θ, θ)
Avg(α, α)
Avg(σ, σ)
Avg(β, β)
Avg(γ, γ)
C
Left Frontal Lobe
Left Somatosensory Cortex
Left Motor Cortex
Left Parietal Lobe
Left Visual Cortex
Left Occipital Lobe
Left Temporal Lobe
Right Temporal Lobe
Right Occipital Lobe
Right Visual Cortex
Right Parietal Lobe
Right Motor Cortex
Right Somatosensory Cortex
Right Frontal Lobe
Posterior Cingulate
Medial Frontal Regions
L. Somatosensory Cortex

Figure 4-2 ***Within-condition cross-frequency super-adjacency matrices (previous page)***: **A)** Group-averaged super-adjacency matrix for wakefulness (W), N2 sleep (N2), and slow-wave sleep (SWS). Each super-diagonal tile corresponds to an inter-frequency layer, with the individual cross-frequency power envelope correlations. The tiles below the diagonal show the mean value of the cross-frequency correlations. For reference, the diagonal tiles show the mean value of the within-frequency power-envelope correlations, which correspond to the mean FC values presented in Chapter 2. **B)** Schematic representation showing the organisation of the super-adjacency matrices in **A**. **C)** Selected inter-frequency layers, corresponding to tiles highlighted by green squares in **A**. Note how a preserved structure is present for the inter-layers of wakefulness. δ (delta): 0.25 – 4 Hz; θ (theta): 4 – 8 Hz; α (alpha): 8 – 13 Hz; σ (sigma): 13 – 16 Hz; β (beta): 13 – 30 Hz; γ (gamma): 30 – 48 Hz.

As expected from the within-condition group averages presented in Figure 4-2, N2 showed higher cross-frequency correlations than wakefulness for most inter-layers except the ones including gamma. Particularly strong differences were found in a 'cluster' of inter-layers involving the alpha, sigma, and beta bands, and another one involving delta, theta, alpha, and sigma. This did not change when considering the absolute correlations. Wakefulness showed signs of higher values in 'long-range' interactions, i.e. between frequency bands distant from each other. As such, the mean (delta, gamma)-layer reached statistical significance, however only for the raw correlations. When using the absolute correlations this significance disappeared, suggesting that this difference was driven by negative cross-frequency correlations in N2.

The contrast N2–SWS was also dominated by higher values for N2 across most frequency-pairings. This time the main effect of using absolute as opposed to raw correlation was a diminished difference in inter-layers involving the gamma band, suggesting that SWS included even more negative correlations than N2. Only the (sigma, gamma)-layer remained significantly larger in N2 than SWS in terms of mean values, when absolute correlations were used.

As for the comparison between wakefulness and SWS, the beta band was involved in most of the significantly different inter-layers at the mean level, when considering absolute correlations. Once again, significant increases for wakefulness involving the gamma band disappeared for absolute correlations, pointing to the non-negligible negative cross-frequency correlations in SWS.

In terms of differences in regional cross-frequency connections it is helpful to consider these from a degree-perspective. For N2 particularly, increases across many of the inter-layers involved the same regions, which accordingly had high differential degrees. In Figure 4-3C are highlighted some of the differential inter-layer matrices that capture this trend. For N2, significant involvement of temporal, somatosensory, and frontal regions were evident, in particular for inter-layers involving the sigma band. These regions corresponded to the rows or columns in the inter-layers that included most of the N2-specific increases, against both wakefulness and SWS.

Wakefulness showed increased cross-frequency interactions mainly in parietal and occipital regions; in the (beta, gamma)-layer compared to N2, and in most inter-layers of the beta band compared to SWS.

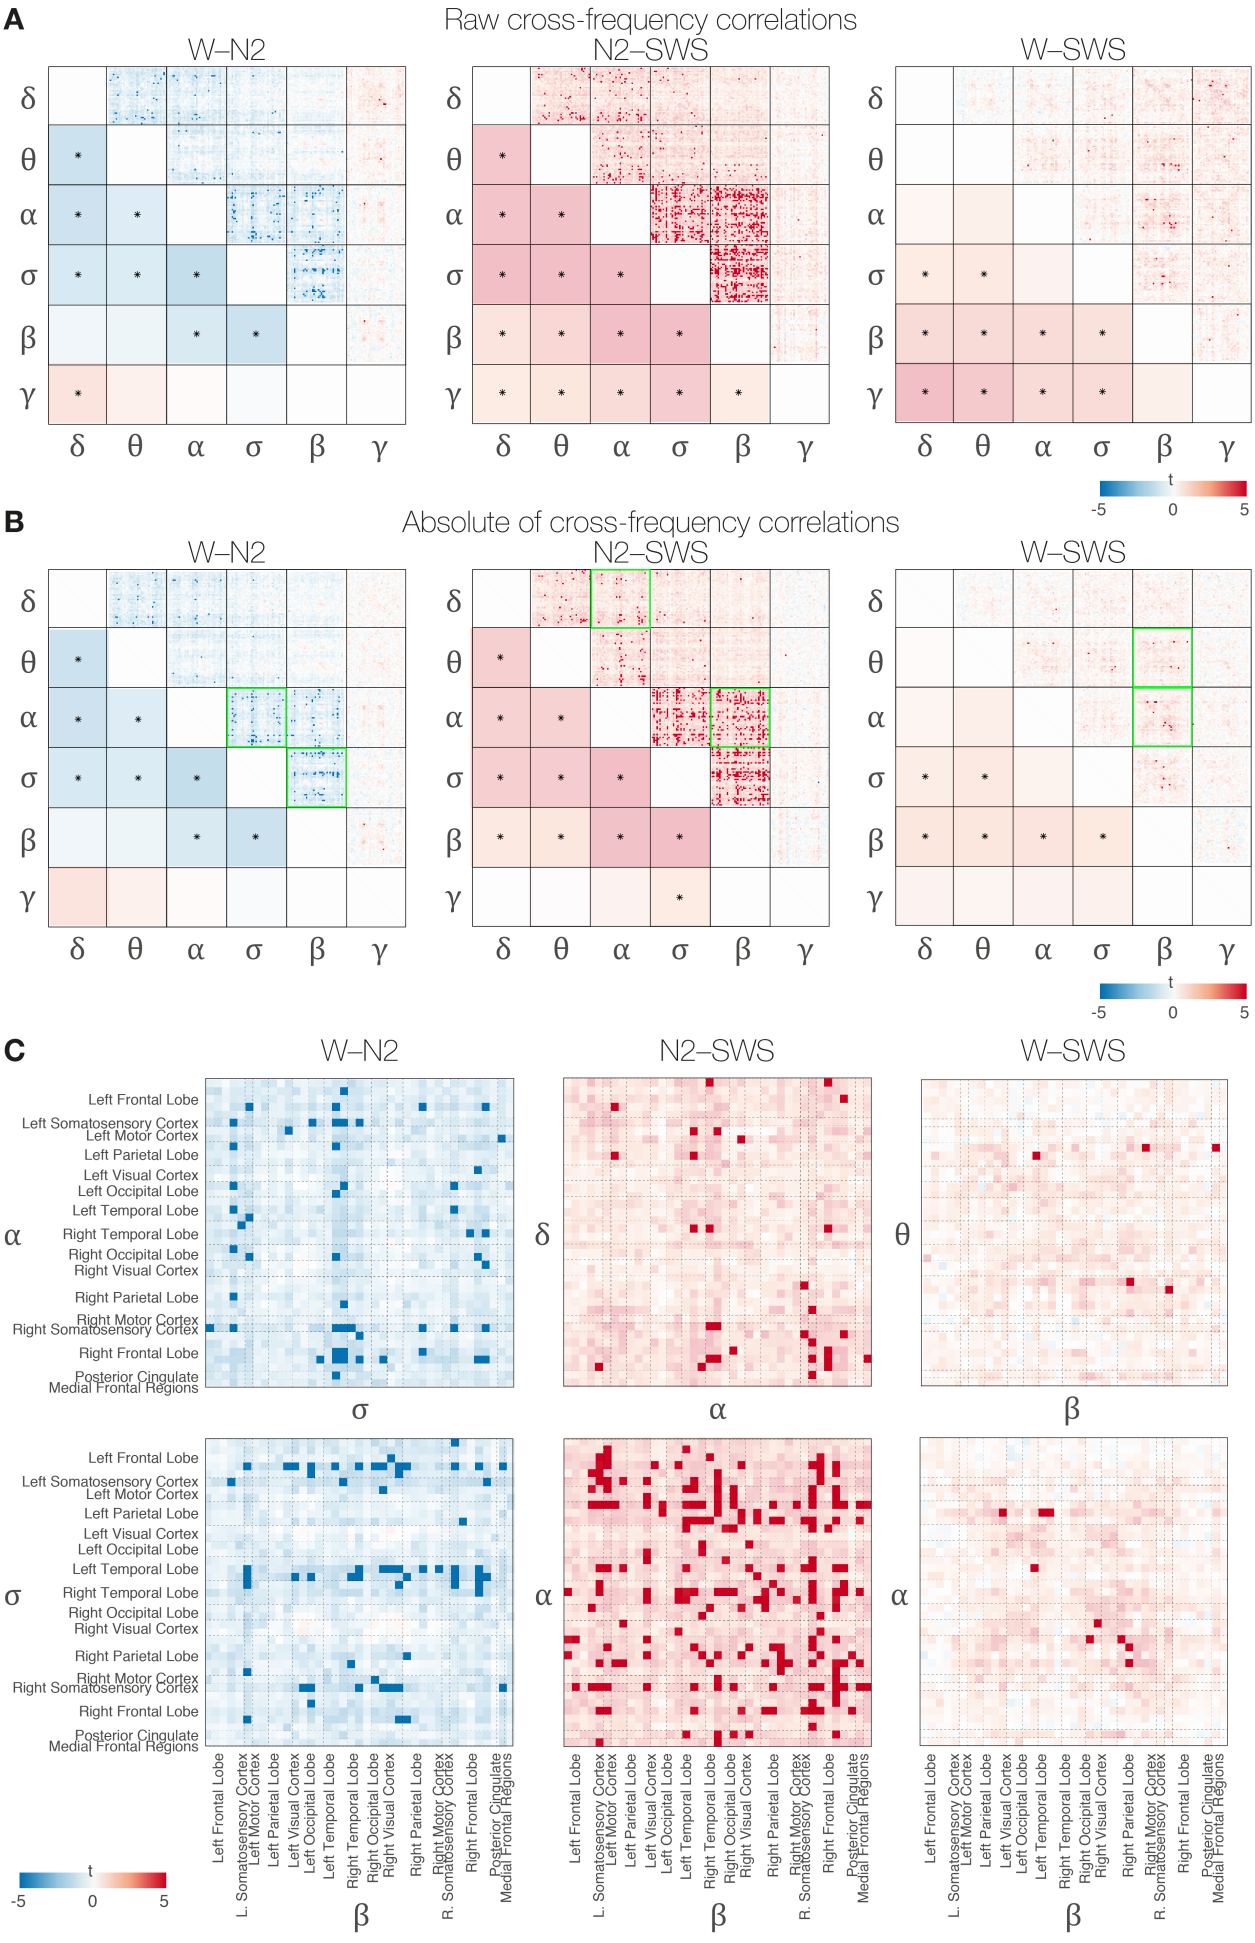
A
Raw cross-frequency correlations
W–N2
N2–SWS
W–SWS
δ θ α σ β γ
t
-5 0 5
B
Absolute of cross-frequency correlations
W–N2
N2–SWS
W–SWS
δ θ α σ β γ
t
-5 0 5
C
W–N2
N2–SWS
W–SWS
Left Frontal Lobe
Left Somatosensory Cortex
Left Motor Cortex
Left Parietal Lobe
Left Visual Cortex
Left Occipital Lobe
Left Temporal Lobe
Right Temporal Lobe
Right Occipital Lobe
Right Visual Cortex
Right Parietal Lobe
Right Motor Cortex
Right Somatosensory Cortex
Right Frontal Lobe
Posterior Cingulate
Medial Frontal Regions
α
δ
θ
σ
β
L. Somatosensory Cortex
R. Somatosensory Cortex
t
-5 0 5

Figure 4-3 ***Paired t-tests between vigilance-specific cross-frequency correlations (previous page)***: **A)** Differential t-stat super-adjacency matrices based on raw cross-frequency power envelope correlations. *Super-diagonal tiles*: Significant cross-frequency connections are highlighted. Significance was evaluated using connection-wise permutations within each tile with corrected p-values < 0.05. *Sub-diagonal tiles*: Paired t-tests performed on the mean of the super-diagonal tiles. * denotes a significant difference at a corrected p-value < 0.05. **B)** Due to the presence of negative cross-frequency correlations paired t-tests were also performed after taking the absolute of the cross-frequency correlations. **C)** Selected tiles from **B** (marked by green squares). δ (delta): 0.25 – 4 Hz; θ (theta): 4 – 8 Hz; α (alpha): 8 – 13 Hz; σ (sigma): 13 – 16 Hz; β (beta): 13 – 30 Hz; γ (gamma): 30 – 48 Hz.

4.3.3 *Vigilance-effect on multi-layer connectivity modes*

In Figure 4-4 are shown the results of the interactions between intra- and inter-layers of the super-adjacency matrices. The vigilance-specific group-averages are shown in Figure 4-4B, computed according to the tile-pairing depicted in Figure 4-4A. For wakefulness the strongest intra-to-inter-layer interaction was found between the alpha and beta bands, followed by the one between the beta and gamma bands. For N2, we observed the same dual-clustered behaviour as described for the super-adjacency matrix above (4.4.1). Hence, the alpha, beta, and sigma bands formed one cluster of strong intra-to-inter-layer interactions, while the delta, theta, and alpha bands formed another. For SWS interactions were generally decreased, and the highest level was found between the beta and gamma bands.

The comparisons between vigilance conditions revealed that for N2 the common mode of connectivity expressed through intra-to-inter-layer interactions was particularly strong in the lower frequency bands, i.e. from delta to sigma. For these frequency bands N2 was significantly increased compared to both wakefulness and SWS. Conversely, interactions between bands more distant apart were elevated for wakefulness. As such, most interactions to the gamma band were significantly increased in wakefulness compared to N2, with the exception of the sigma band. Relative to SWS it was mainly the interactions from the gamma band that were increased in wakefulness. SWS was generally decreased in most interactions compared to both wakefulness and SWS, apart from for interactions to the gamma band. Theta-to-gamma and beta-to-gamma interactions were thus significantly higher for SWS relative to N2, and not significantly different from wakefulness.

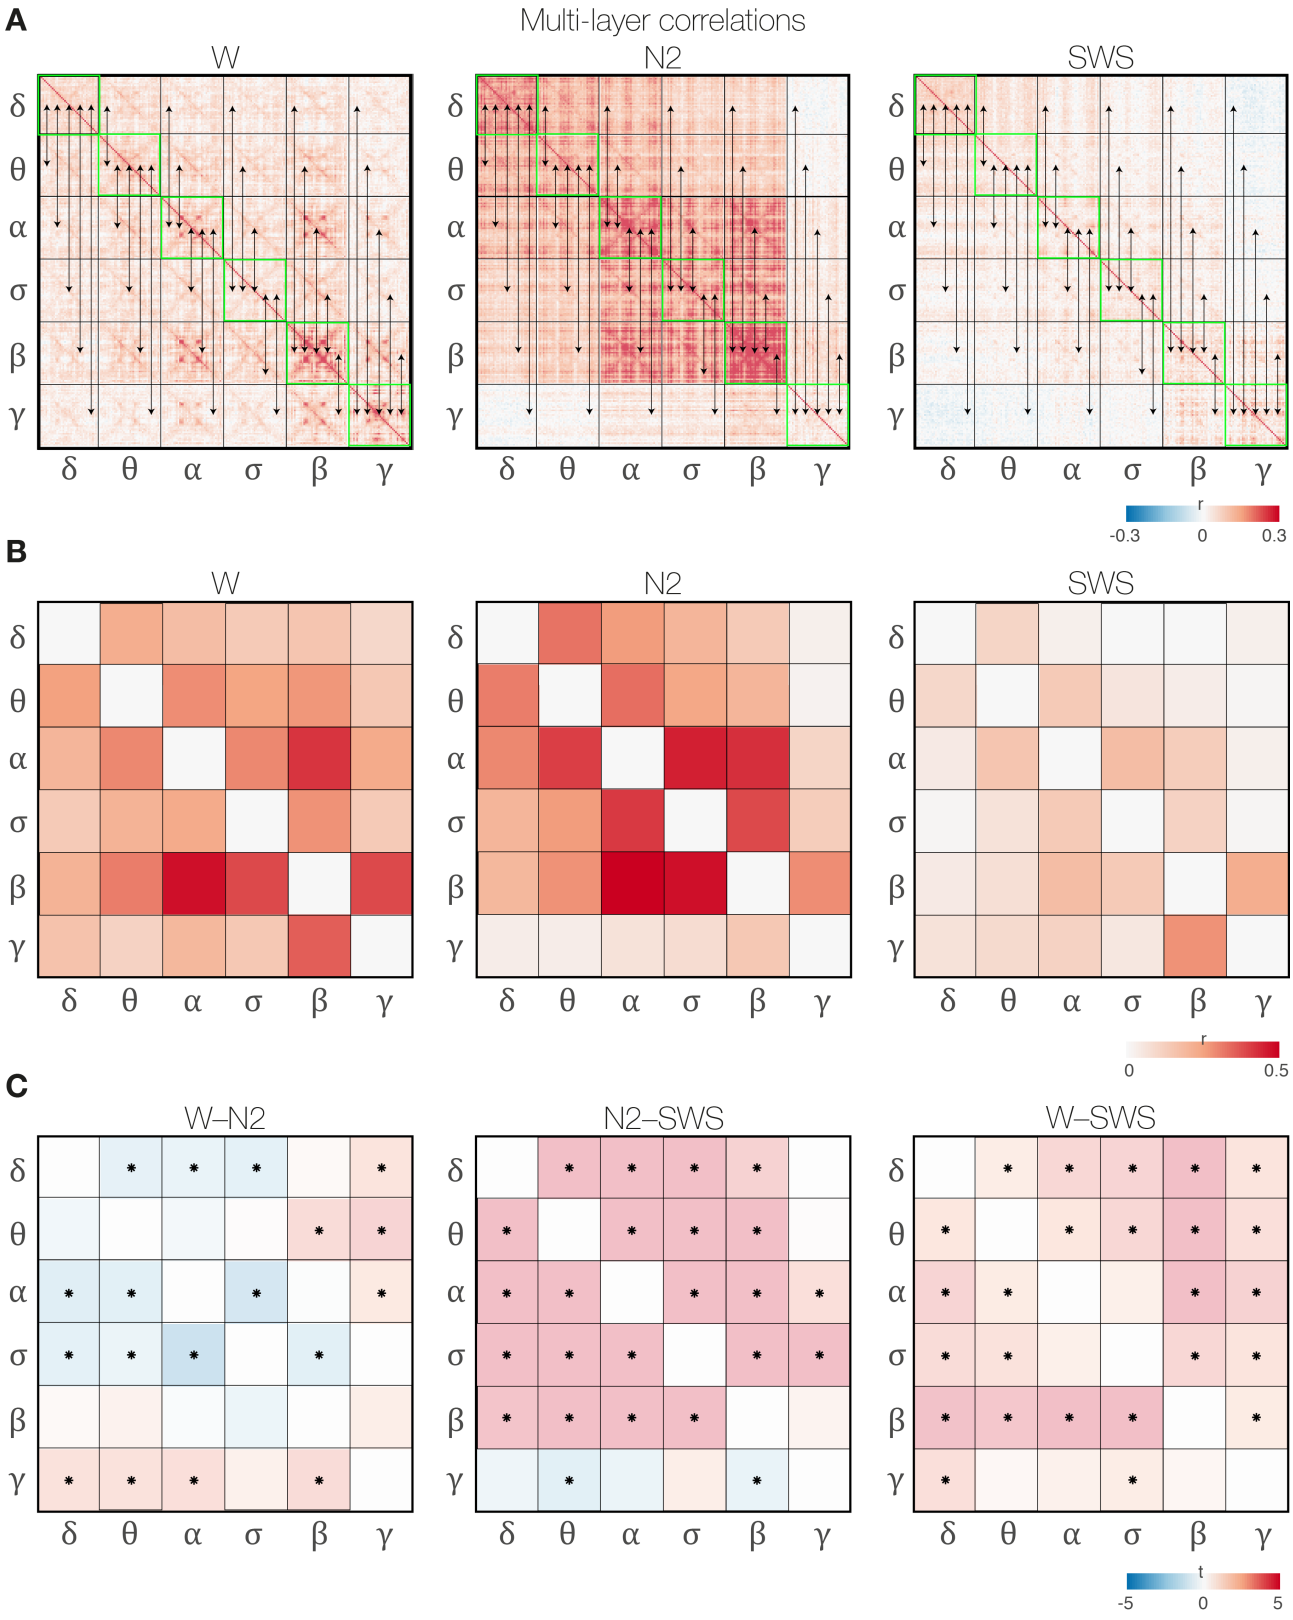

Figure 4-4 ***Interactions between intra- and inter-layer FC***: **A)** Schematic representation of the principle of estimating pattern-similarities between layers of the super-adjacency matrices. Arrows show the tile-pairings leading to the group-averaged multi-layer matrices in **B**. **C)** Paired t-tests between pairs of vigilance conditions for each tile of the multi-layer matrices. Elements show resulting t-stats and * denotes significant differences at a corrected p-value < 0.05.

4.4 ***Discussion***

With the aim of applying a more comprehensive pan-spectral view of the MEG envelope correlations, we have explored cross-frequency interactions between six canonical frequency bands in 39 cortical areas. We have provided group-averages within three vigilance conditions (wakefulness, N2 sleep, and SWS), and tested for significant differences between these in terms of magnitudes of the cross-frequency correlations. Finally, we have tested whether a recently proposed common mode of multi-layer connectivity is affected by changes in vigilance.

4.4.1 *Cross-frequency coupling changes from wakefulness to sleep*

The group-averaged super-adjacency matrix for wakefulness showed considerable overlap with the ones reported for resting-state data by Brookes and colleagues and Tewarie and colleagues (Brookes et al., 2016, Tewarie et al., 2016b). As such, cross-frequency correlations were high between the alpha and beta bands, as well as between the beta and gamma bands. The high cross-correlations between envelopes of occipital and parietal regions were also consistent with these previous reports (Brookes et al., 2016, Tewarie et al., 2016b). These inter-layers and regions were also involved in the significant differences between wakefulness and SWS. The beta band, in particular, was generally more connected to other frequencies in wakefulness than in SWS. At the same time our results show that this is not a feature specific to wakefulness, since N2 also showed many increases in these inter-layers compared to SWS. Some of these inter-layers of N2 even exceeded wakefulness in terms of mean values.

When differences were observed in cross-frequency correlations between vigilance conditions these typically included the same regional connections identified in the within-frequency differences, reported in Chapter 3. The wakefulness-specific increases in visual and parietal regions compared to SWS were thus also found within the alpha and beta bands (see Figure 3-9A and B). Likewise, N2 showed particularly high cross-frequency coupling between frontal, somatosensory, motor, and temporal regions, in patterns corresponding to those within the delta, alpha, sigma, and beta bands (see Figure 3-9A and B). This suggests that strong connections within a frequency band tended to be predictive of strong connections from that band to other frequencies. An equivalent conclusion was reached for resting-state data by Tewarie and colleagues (2016b). Cross-frequency relationships between the band-limited power envelopes are therefore best understood as highly dependent on their within-frequency connectivity. This explains why we observed minimal cross-frequency coupling in SWS, since within-frequency FC is generally low for this vigilance condition, as shown in Chapter 3 (see e.g. Figure 3-8). Further, cross-frequency coupling to and from the gamma band was negligible for N2, in line with the low levels of FC within this band.

Many of the cross-frequency findings could thus be explained by within-frequency FC. Meanwhile, the pan-spectral view did identify some unique vigilance-specific multi-layer characteristics. The contrast between wakefulness and N2 sleep was particularly illustrative of how frequency bands express significant interactions during N2 sleep. N2-specific increases were thus observed in two clusters of inter-layers; one between the delta, theta, and alpha bands, and one between the alpha, sigma, and beta bands. This could be interpreted as a sign of two distinct FC generators in N2, giving rise to two fundamental types of

explorations of the multi-layer network. As discussed in Chapter 3, K-complexes and sleep spindles are likely drivers of within-frequency FC, and these graphoelements may well be causing the two clusters of cross-frequency interaction. Overall, it is clear that N2-specific increases do not reduce to within frequency FC, but rather involve high degrees of cross-frequency interactions, involving all bands except gamma.

The interdependence of oscillations during sleep is a well-established phenomenon (Steriade, 2006). The slow oscillation, in particular, is believed to orchestrate the activity of faster activities, such as spindles and transients in the gamma range (see 1.3.1). This grouping of frequencies, originally identified in intra-cortical and intra-cellular recordings (Steriade et al., 1993b, Contreras and Steriade, 1995, Clemens et al., 2007, Le Van Quyen et al., 2010), has been demonstrated in the human EEG (Mölle et al., 2002, Piantoni et al., 2013), MEG (Klinzing et al., 2016), and proposed as a neural correlate of sleep-dependent memory consolidation (Diekelmann and Born, 2010, Cox et al., 2012, Rasch and Born, 2013, Cox et al., 2014). Here it is important to note that such cross-frequency relations have been described mainly as interdependencies between the *phase* of a low-frequency oscillation, i.e. the UP-phase of the slow oscillation, and the *amplitude* of higher-frequency activities, such as spindles, gamma activity, and hippocampal ripples. Phase-amplitude coupling (Canolty and Knight, 2010) differs from the power-amplitude coupling (Engel et al., 2013) used to compute the multi-layer networks in this chapter. Specifically, coupling through power-amplitude correlations rely on fluctuations in the oscillatory power of both bands in question, whereas phase-amplitude coupling can arise even if the low-frequency band maintains constant power. Our results may therefore not be directly translatable to the observations referenced above, but rather address whether power fluctuations express coupling across frequencies during sleep.

Power-amplitude correlations of the delta band change markedly from N2 sleep to SWS, with the two cluster-configurations of N2 sleep engaging and connecting all frequency bands, except gamma, while few interactions were evident in SWS. This stands in contrast to phase-amplitude coupling, which is generally reported higher for SWS (Amiri et al., 2016). Meanwhile, our results did produce an interesting overlap between the two vigilance conditions in terms of the negative association between the gamma band and the lower frequency bands, delta and theta. Through the comparisons to wakefulness, and by investigating the effect of taking the absolute of the correlations, we were able to show that, albeit small in magnitude, these negative relationships are non-negligible. Furthermore, the comparison between N2 and SWS suggests that this phenomenon was particularly consistent for SWS. A more careful examination of this phenomenon could provide important evidence, complementing the existing hypotheses regarding phase-amplitude interactions of the slow oscillation and faster frequencies in sleep (Rasch and Born, 2013).

4.4.2 *Multi-layer integration and segregation are affected by changes in vigilance*

The multi-layer network approach offered a comprehensive overview of the multi-dimensional MEG data of 39 cortical regions and their spatio-temporal power envelope interactions in six frequency bands. Assuming that each frequency layer holds the potential of conveying unique information throughout the cortical network, it is clear from our results that this potential is poorly explored in both N2 sleep and SWS. During SWS few layers of frequencies produce consistent and distinct patterns of connectivity (see also Chapter 3),

while in N2 all frequency bands, except gamma, appear to be transmitting highly similar content. Returning to the radio analogy in the introduction of this chapter, this would corresponds to SWS transmitting on a restricted number of channels, while N2 may be likened to a hijacking of most of the available channels. The apparent negative relationship between the gamma band and other bands in N2 and SWS may be paralleled to a channel that goes off-line whenever other channels are 'on air'.

The multi-layer character of wakefulness is found somewhere between those of N2 and SWS. Tewarie and colleagues addressed the balance between intra-layer (within-frequency) uniqueness and inter-layer (cross-frequency) sharing of information in resting-state data (Tewarie et al., 2016b). By comparing intra- and inter-layers in terms of their patterns, they were able to quantify how cross-frequency networks share a common mode of connectivity, on top of which specialised intra-frequency interactions are added (Tewarie et al., 2016b). Applying an equivalent analysis of interactions between intra- and inter-layers, and comparing these across the vigilance conditions, our results provide potential clues to the functional importance of such multi-layer integration and segregation (Figure 4-4). As such, cross-frequency networks were more similar in wakefulness compared to SWS across most layers, pointing to more integration in wakefulness. In contrast, the significantly higher similarity in N2 relative to wakefulness between most layers and inter-layers involving bands from delta to sigma could reflect a decreased level of frequency-specific specialisation for N2. In the pan-spectral framework, SWS and N2 may thus represent two opposite extremes of segregation and integration, respectively, in between which wakefulness strikes a balance.

The higher degree of multi-layer integration of the gamma band during wakefulness, compared to both N2 and SWS, is intriguing (Figure 4-4). In Chapter 3 we were unable to identify clear wakefulness-specific increases in FC within the gamma band. This may seem surprising, given theories linking gamma oscillations to high-order cognition and perception (Singer and Gray, 1995, Jensen et al., 2007, Fries, 2009). However, these theories are mostly based on evidence of localised synchronisation of gamma activity rather than long-range interactions between brain regions. Gamma oscillations thus primarily seem to represent important substrates of local neural computations, which may rely on slower frequencies for their inter-regional integration (Von Stein and Sarnthein, 2000, Siegel et al., 2012, Engel et al., 2013). The common mode of connectivity that we find to be particularly preserved during wakefulness, compared to both N2 and SWS (see Figure 4-4), could be reflecting a wakefulness-specific basis for incorporating these high-order computations into the large-scale, multi-layer network.

4.4.3 *Methodological considerations*

As alluded to in preceding sections, amplitude envelope coupling is but one way of probing inter-regional interactions in electrophysiological data (Engel et al., 2013). This measure has gained popularity from its ability to identify interdependencies among cortical regions within frequency bands, in networks resembling those known from fMRI FC (see Introduction, 1.5.7 and Colclough et al. (2016)), and in Chapter 3 we showed that it may be used to extract vigilance-specific information. Specifically, for the purpose of capturing cross-frequency interactions, phase-amplitude coupling has nevertheless been the measure-of-choice (Canolty and Knight, 2010). The interaction between the phase of a low-frequency oscillation, such as

delta or theta, and the power amplitude of higher frequencies, such as beta or gamma, has mainly been demonstrated in rather localised regions with intra-cortical recordings or in topographical sensor space in EEG/MEG (Canolty and Knight, 2010). A very recent study showed that local low-to-high frequency phase-amplitude coupling co-vary across the cortex in patterns resembling RSNs, suggesting that this kind of oscillatory interaction may complement amplitude envelope coupling as a principle for large-scale networks (Florin and Baillet, 2015). To our knowledge there has been no attempts at mapping large-scale interactions of phases and amplitudes between distant regions. We opted for an amplitude-amplitude cross-frequency approach, since this yielded immediately interpretable results in relation to the within-frequency analyses presented in Chapter 3. As noted in Tewarie et al. (2016b), however, the concept of a multi-layer network could in principle incorporate a range of coupling measures both within and across frequencies. An exciting future avenue for MEG sleep data would be to incorporate phase-amplitude coupling, as this might provide a large-scale view on the grouping of sleep oscillations, as discussed in section 4.4.1.

Chapter 5
Transient events in sleep functional connectivity using MEG

5.1 ***Introduction***

In the preceding chapters it has been shown how large-scale FC can offer new perspectives on differences in brain activity between wakefulness and sleep, the understanding of which has traditionally been firmly rooted in spectral analyses of a limited number of EEG channels. While fMRI offers a significantly increased spatial specificity (see Chapter 2), MEG is endowed with the temporal resolution to provide frequency-specific information of the large-scale networks (see Chapter 3 and Chapter 4). Meanwhile, it is important to get a handle on how the more derived FC measures relate to the traditional polysomnographic description of spectral changes.

Many of the findings presented in the foregoing chapters have been discussed in relation to sleep graphoelements, like K-complexes and sleep spindles. These sleep-defining electrophysiological signatures are likely to have some form of impact on the FC estimations, however the relationship between these two levels of description remains unexplored. Sleep graphoelements are essentially transient events in different frequency bands of the spectrogram of sleep recordings. As such, we employed a sliding window approach to the MEG data, calculating for each window the FC matrices in each of the six frequency bands as well as the mean spectral power of the corresponding sensor signals. Tracking frequency-specific FC and sensor power over time in this way allowed for an exploration of the relationship between transient spectral events and MEG FC.

5.2 ***Methods***

5.2.1 *Temporal relationship between sensor power and FC*

We wanted to investigate the relationship between our estimated MEG FC measure, i.e. band-passed power envelope correlations, and signal power. Given the transient nature of sleep graphoelements, such as K-complexes and spindles, it was important to resolve this relationship in time. Considering, once more, the beamformed ROI timecourses derived in Chapter 3 (see 3.2.3), within the six chosen frequency bands (delta, theta, alpha, sigma, beta, and gamma), we employed a sliding window approach to calculate FC matrices for smaller segments of time. We estimated FC between the 39 ROIs, as described in section 3.2.4, within time windows of 30 seconds, moving the window 5 seconds at a time, such that consecutive windows would overlap by 25 seconds. The mean FC level was calculated for each window, using full correlations between the power envelopes. This resulted in a timecourse of mean FC within each frequency band, for each participant, for each vigilance condition, v. For the same windows, W_T, we then turned to the signals in sensor space, and computed the power within each frequency band, f. For each MEG sensor, s, used for the beamforming, power was estimated using the built-in MATLAB function *bandpower*:

$$Pow_{f,v,W_t,s} = \frac{1}{N_{samples}} \sum_{t=t_1}^{t_2} s^2(t) \qquad \textbf{Eq. 5.1}$$

where t_1 and t_2 demarcates the time window of 30 seconds, as defined by W_t, while $N_{samples}$ denotes the number of samples between t_1 and t_2 at a sampling frequency of 250 Hz.

Averaging across the single-sensor-values for each window yielded the mean power value timecourse, Pow_{f,v,W_T}, on the same time scale as the mean FC timecourse. We were thus able to correlate the mean FC and mean power over time, using Pearson's correlation. Doing this across all frequency pairs yielded an $N_{frequency}$ x $N_{frequency}$ matrix for each participant and vigilance condition, representing a measure of the relationships between sensor signal power and the derived FC. We tested for differences between vigilance conditions for each pairing of mean FC and sensor signal power, using paired t-tests corrected for multiple comparisons using non-parametric permutations. Corrected p-values < 0.05 were considered significant.

5.2.2 *Temporal variability of MEG FC*

The basic question of this chapter was whether fluctuations in power modulate the derived FC measure. Such temporal modulation presupposes that FC itself is fluctuating over time. We quantified the amount of variability in the mean FC level by the ratio between its standard deviation and its mean over time. Once again we tested for cross-vigilance differences using paired t-tests. This was done within frequency bands, and non-parametric permutations were used to correct for multiple comparisons. Corrected p-values < 0.05 were considered significant.

5.3 ***Results***

An overview of the relationships between average sensor power and mean levels of FC is presented for each vigilance condition in Figure 5-1A. The relationship was tracked with overlapping 30-seconds-windows, with the aim of resolving potential influence of relatively short-lived events in the sensor data upon the derived source-space FC. Each vigilance condition is represented by a matrix of correlations between the development of sensor spectral power in each of the six frequency bands and the development of mean levels of FC throughout the 30-seconds sliding windows. For wakefulness, the most striking observation was the apparent negative association between sensor power in the alpha band and mean FC in the beta band. A similar negative association was seen in N2 for sensor power in the theta band and mean FC in the higher frequency bands. Unlike wakefulness, N2 expressed a high positive correlation between sensor power and mean FC in the higher frequency bands – beta in particular. This relationship was found significant compared to both wakefulness and SWS (see Figure 5-1B). SWS showed some signs of positive correlations between sensor power in the beta band and mean FC in lower bands, albeit not at the same magnitudes as during N2. Figure 5-1C shows the amount of variability of mean FC for each of the vigilance conditions in each frequency band. When quantified as the ratio between the standard deviation and the mean over time, it is clear that both N2 and wakefulness expressed significantly more fluctuations in FC compared to SWS. Further, N2 showed higher mean values than wakefulness, reaching significance in the delta and beta bands.

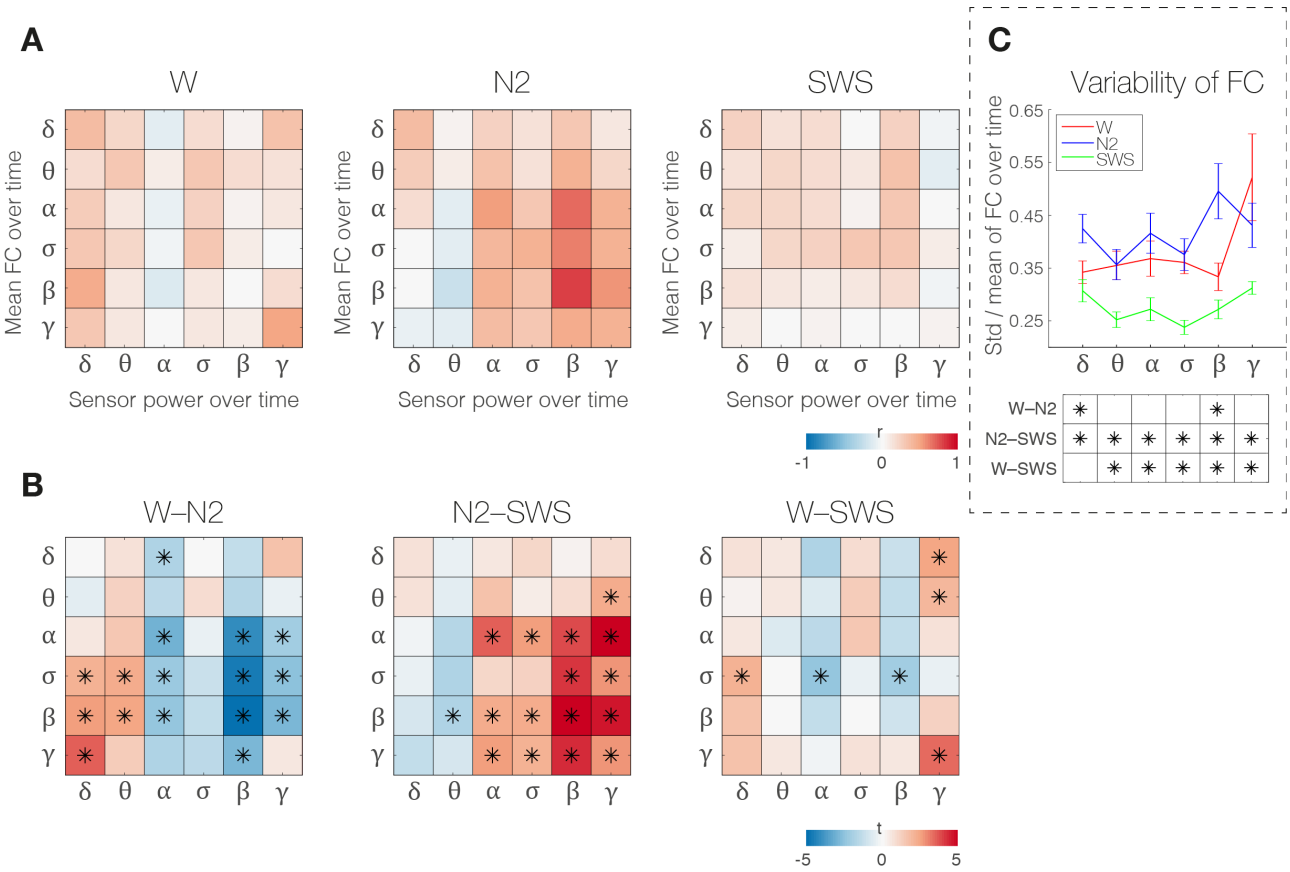

Figure 5-1 ***Overview of relationships between sensor spectral power and mean levels of FC in six frequency bands***: **A)** Group-averaged Pearson's correlation between sensor power (x-axis) and mean level of FC (y-axis) across the sliding windows, for wakefulness (W), N2 sleep (N2), and slow wave sleep (SWS). **B)** results of paired t-tests between vigilance conditions for each pairing of spectral power and mean FC level. **C)** *Top***:** Group-averaged variability of mean FC level, represented as the ratio between the standard deviation and the mean across time windows. Error bars represent the standard error (N = 11). *Bottom*: Results of paired t-tests between vigilance conditions. * indicates a significant difference at a corrected p-value < 0.05.

The results in Figure 5-1 point to vigilance-dependent interactions between relatively short spectral events and the mean FC level of the MEG data. For illustration and to investigate these phenomena further, we plotted time-frequency representations of the sensor data together with the temporal development of mean FC levels in source-space for each vigilance-specific epoch. Epochs from a representative participant are shown in Figure 5-2. Similar representations are given for the remaining participants in the Appendix (see Figure S 2 to S 7). The negative correlation between sensor power in the alpha band and mean FC for wakefulness can be appreciated visually for most participants. Hence, a fading of power in the alpha range of the spectrogram was often associated with an increase in mean FC, particularly in the beta band. Similarly visible is the high variability in mean FC during N2. Furthermore, it is particularly noticeable how specific events in the spectrogram of N2 were temporally closely related to increases in mean FC, particularly in the alpha, sigma and beta bands. These events were consistent across participants, with a distinct spectral signature of increases in the delta, alpha and beta bands and decreases in the theta band, thus reflecting the plot for N2 in Figure 5-1A. These rather broadband events expressed durations of about 5 – 10 seconds and occurred between 1 and 4 times within the 300 seconds epochs analysed for N2. Unlike, the two other vigilance conditions, SWS showed no clear associations between spectral events and changes in mean FC.

For a closer look at the spectral events with clear association to FC in N2, Figure 5-3 shows extracts of MEG sensor timecourses from the three relevant time windows of the representative participant. Considering these timecourses more carefully they would appear to share a common sequence of a slow and rather large amplitude deflection for ~ 1 second followed by a fast burst-like segment of lower amplitude from 2 to 4 seconds. Again this was a consistent feature across participants (see Appendix, Supplementary figures, Figure S 2 to S 7).

Relationship between sensor signal power and mean functional connectivity

W

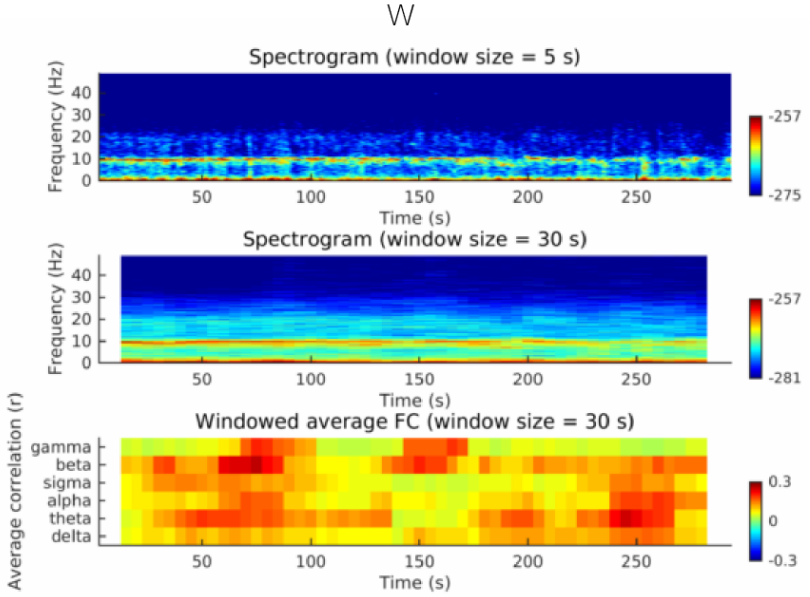

N2

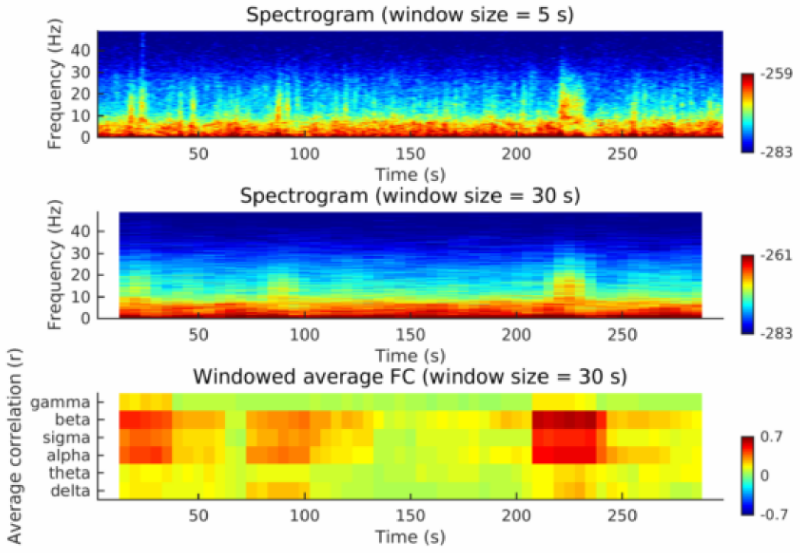

SWS

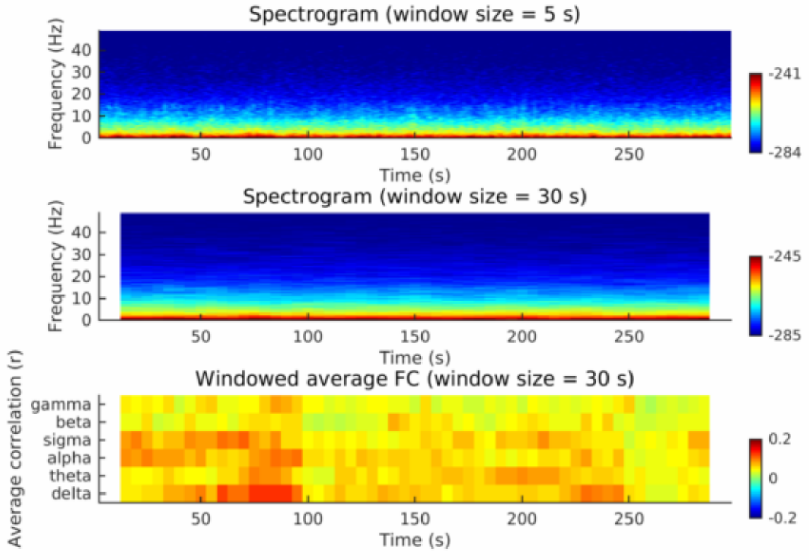

Figure 5-2 ***Time-frequency representations of relationship between sensor power and mean level of FC in representative participant (previous page)***: Two spectrograms were computed from the broadband MEG sensor timecourses for each vigilance condition; both using the short-time Fourier transform in MATLAB, one with a window size of 5 seconds and 4 seconds overlap, and one with a window size of 30 seconds and 25 seconds overlap. The latter corresponded to the window size and overlap of the sliding window mean FC computation. Note how the epoch from wakefulness (W) is characterised by increases in mean FC for time windows where the spectral power of the alpha band is low. Note also, how specific transient spectral power increases appear to be driving mean FC in N2 sleep. These transients occurred on the background of a presumed spindle activity in the 10 – 20 Hz band. For SWS, no clear associations can be found between spectral events and mean FC levels. Note that colour scales have been adjusted for each vigilance condition.

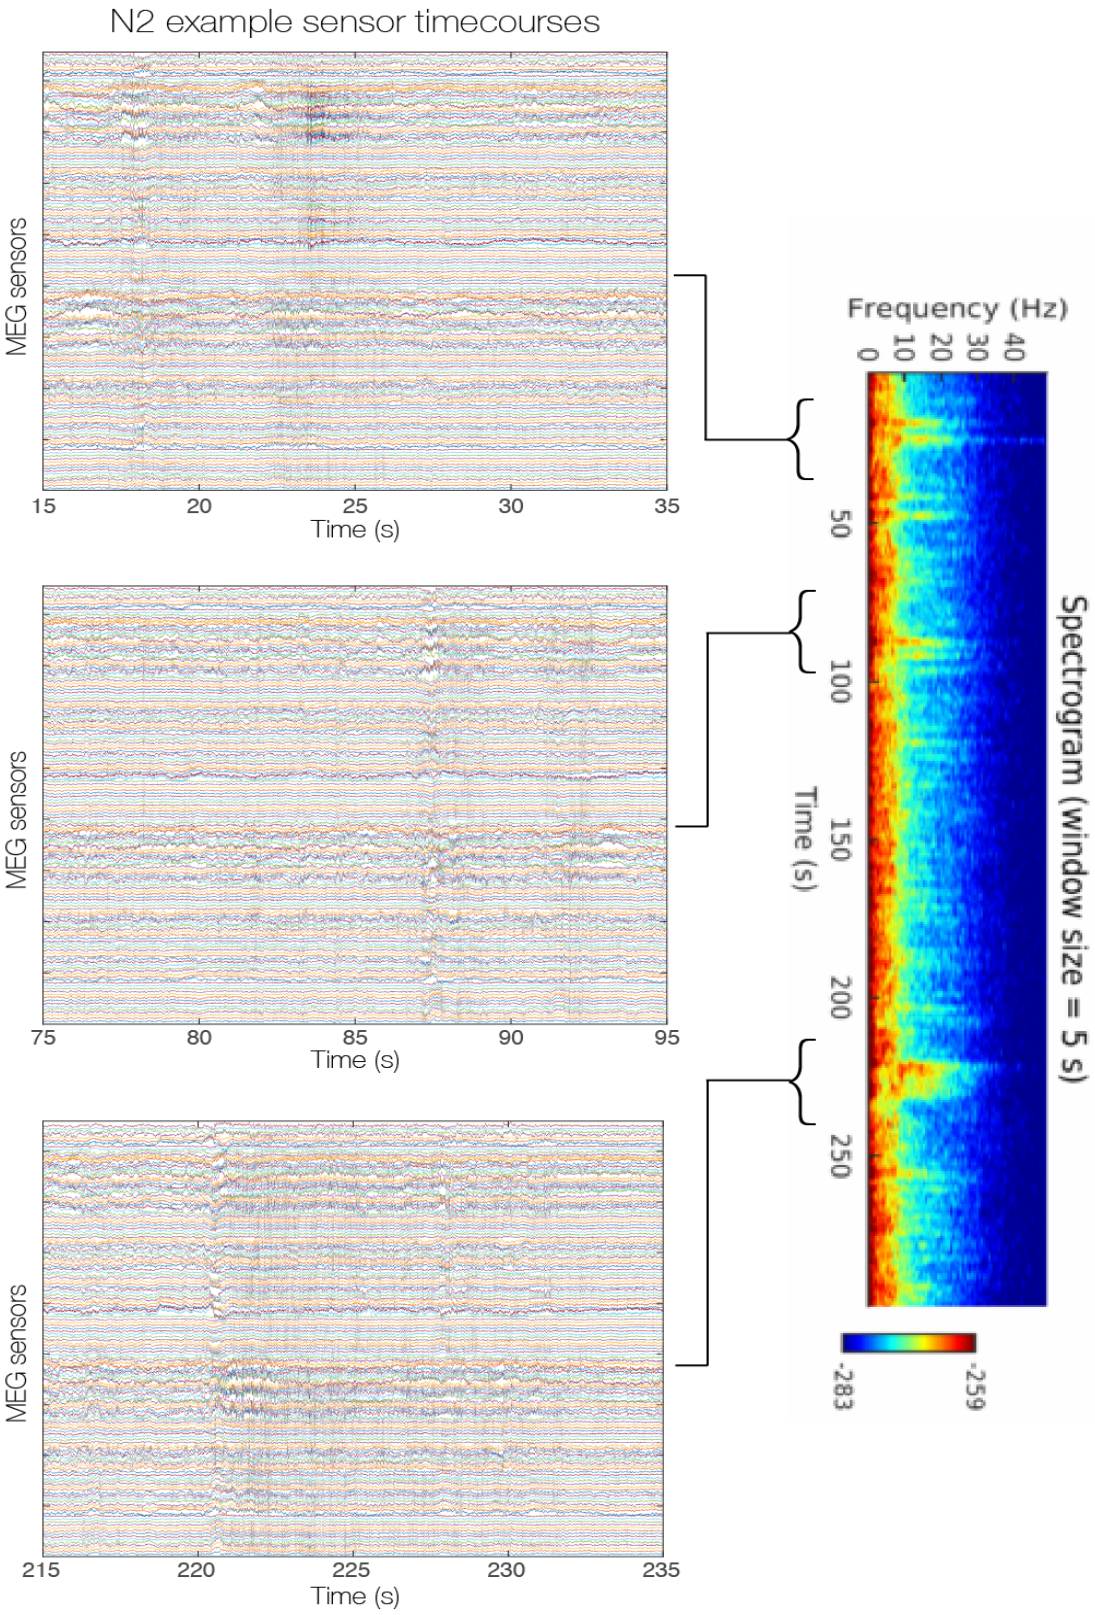

Figure 5-3 ***Sensor timecourses of N2-specific transients from representative participant***: The spectrogram shown vertically on the right is the 5 seconds window version from **Figure 5-2**. The three epochs surround the transient spectral events that significantly increased mean FC during N2 sleep, and show the sensor timecourses of the broadband MEG signal.

5.4 ***Discussion***

Following up on the within- and cross-frequency MEG FC findings reported in Chapter 3 and Chapter 4, we set out to investigate the influence of transient changes in the MEG power spectrum and the employed measure of FC, power envelope correlations between beamformed timecourses. We have shown that the interplay between spectral power and overall FC varies with vigilance. As such, the temporal development of alpha power was negatively associated with mean FC in higher frequencies, such as beta, for wakefulness. A similar inhibiting influence of spectral power appeared in the theta band in N2 sleep. Unlike in wakefulness and SWS, mean FC in N2 further expressed a clear positive contingency on short-lived spectral events of rather broad frequency profiles with peaks in the beta range. These transients were found consistently across participants and were characterised by a common sequence in the MEG sensor signals.

5.4.1 *Transient spectral events linked to mean FC in N2 sleep*

In Chapter 3 as well as in Chapter 4, sleep spindles and K-complexes were discussed as conceivable drivers of the within- and cross-frequency MEG FC patterns identified for N2. The results of the present chapter provide direct evidence that N2 is characterised by an association of specific transient events in the spectrogram and mean FC. This positive correlation was found particularly strong between spectral power increases in the beta band and mean FC in the frequency ranges of alpha, sigma, beta and to a lesser extent gamma. This grouping of frequencies is in line with the clustering of cross-frequency correlations described in Chapter 4.

The spectrogram representations in Figure 5-2, Figure S 4, and Figure S 5, identifying the N2-specific transient events, showed a primary involvement of increased power from ~ [10 – 25] Hz. As such, these events seemed to engage frequencies above the normal upper range of sleep spindles (~16 Hz). Further, it is worth noting that they often occurred against a background of more confined activity in what may be regarded ‘classical’ spindle territory. This could be a manifestation of the distinction between local and global spindling, discussed in Chapter 3, since these more broadband transients were particularly efficient at increasing mean FC levels.

The timecourse representations in Figure 5-3 identified a common timecourse pattern for the transient events, according to which a large-amplitude slow deflection was followed by faster activity of lower amplitude. An obvious candidate for the early part of this pattern is the K-complex. The temporally successive relationship between K-complexes and spindles is a well-described phenomenon in the EEG literature (Colrain, 2005). Drawing on the results of the regionally more specific FC analyses carried out in Chapter 2 and Chapter 3, we propose that the identified transient events are rather global synchronisations of functional networks across regions as well as across a large group of frequencies during N2 sleep.

We note that the employed strategy of using large-scale FC networks as the primary observation was key to the identification of these specific events. The present analyses are thus an example of approaching sleep data from a large-scale perspective resulting in the identification of specific graphoelements. This stands in contrast to the tradition of focussing on historic definitions of sleep graphoelements by their own right where

less attention is paid to their large-scale properties. Future work should investigate the role of these broadband transients of N2 in independent datasets.

A curious observation is that SWS showed few signs of associations between spectral events and FC. In Chapter 3 it was established that the majority of FC during SWS is found in the alpha band (see Figure 3-5), and it was discussed whether the envelope correlations were reflecting spindle activities or another less well-defined source of alpha oscillations. The results of the present chapter would seem to indicate that alpha FC in SWS is rather independent of transient spectral events, whereby a prominent role of spindles appears unlikely.

5.4.2 *Non-static FC in N2 sleep and wakefulness*

The significant influence of spectral events on MEG FC during N2 was naturally reflected in the temporal variability of mean FC, which was highest during this vigilance condition (Figure 5-1C). As such, N2 showed significantly more FC variability than SWS in all frequency bands, and more than wakefulness in the delta and beta bands. In turn, wakefulness showed higher mean FC variability than SWS in all bands but delta. Even if tracking the temporal evolution of mean FC levels represents a rather non-specific representation of the dynamic properties of FC, the presented results give strong indication that, for N2 and wakefulness, FC depends noticeably on the time window chosen for its estimation. This observation poses a conflict with the temporal stationarity that is assumed when FC is averaged over time (see Introduction, 1.5.8). This is in line with more recent attempts in resting-state studies at resolving dynamic as opposed to static FC in both fMRI (Hutchison et al., 2013a, Allen et al., 2014, Calhoun et al., 2014, Preti et al., 2016) and MEG (Baker et al., 2014a, O'Neill et al., 2015, Vidaurre et al., 2016). Our results suggest that similar approaches are warranted for the case of sleep data, and N2 in particular. In contrast, the developments of FC in SWS were more in line with the concept of stationarity.

5.4.3 *An inhibiting role of alpha upon FC during wakefulness*

An underlying motivation for using large-scale FC to characterise spontaneous brain activity is that this level of description is assumed to provide information above and beyond increases or decreases in spectral power (Hipp et al., 2012). This has been shown to be the case for MEG FC when averaging over time, i.e. for the static approaches (Hillebrand et al., 2012, Hipp et al., 2012). Indeed in Chapter 3 we showed that the vigilance-specificity of static MEG FC networks cannot be explained by changes in source power (see Figure 3-6). To our knowledge, however, there have been no accounts on how the interplay between spectral power and MEG FC evolve over time. The results of the present chapter indicate that MEG FC during wakefulness is not static, but varies depending on the time window used for its estimation (see Figure 5-1C and Figure 5-2). When tracked over time for each of the six included frequency bands, it was clear that mean FC level did in fact express an association with changes in spectral power. The negative correlation between power in the alpha band and mean FC, in the beta band in particular, was a specific characteristic of wakefulness. This suggests an antagonistic relationship between power in the alpha band and interregional communication, which interestingly aligns well with the so-called inhibition-timing hypothesis (Klimesch et

al., 2007) and related ideas about alpha oscillations (Mazaheri and Jensen, 2010). According to these views, the power of on-going alpha oscillations dictates the brain's capacity to process information. Specifically, increased power in alpha is believed to inhibit cognitive processing. While this has mainly been tested in task-related studies, the negative influence that we observe of alpha power on mean FC of beta, and to a lesser extent of alpha itself, could by deduction be interpreted as important evidence that FC does in fact mark cognitive processes in spontaneous brain activity.

It is worth noting that resting-state studies combining EEG and fMRI have provided concurrent evidence of the negative association between EEG alpha power and FC derived from fMRI (Scheeringa et al., 2012, Tagliazucchi et al., 2012c, Chang et al., 2013). The study by Tagliazucchi et al. (2012c) further reported a positive correlation between spectral power in the gamma band and FC, a finding which may be reflected in the positive correlation between spectral power and MEG FC that we observed in the gamma band for wakefulness (Figure 5-1A).

5.4.4 *Methodological considerations*

A primary finding of the present chapter is that specific transient increases in spectral power from ~ 10 to 25 Hz is associated with marked increases in mean level of FC during N2 sleep. An obvious concern in this regard is whether potential muscular artefacts may have escaped our preprocessing procedures and caused the transient power increases. The current dataset did not allow us to control for this, since no additional information was collected about muscle activity or head movements. It is worth stressing, however, that the most significant effects of the power increases on FC were found in the frequency ranges of alpha, sigma, and beta, and not in the gamma band, which would be expected in the case of muscular interference (Muthukumaraswamy, 2013a). The transient events identified in the spectrogram of N2 were largely contained in the alpha to beta ranges. Furthermore, these events were seen solely for N2. Muscular tone is known to be highly depressed during SWS, however in terms of muscular interference we would have expected at least similar amounts during wakefulness, which was not found. Future datasets with better reference recordings of e.g. scalp, facial, and neck muscle activity as well as continuous head tracking should be used to address the possible influence of muscular contractions and movement.

As alluded to above, measuring the mean FC level over time is a rather crude strategy for representing FC dynamics. While a marked change in mean FC is indicative of a network change, smaller network reconfiguration can take place under reasonable stable mean FC levels. Furthermore, even though we used a sliding window to capture temporal non-stationarities in FC, we still downsampled the power envelopes as in Chapter 3 and Chapter 4 before calculating the correlations between them. As such, we kept the focus on the slow part of the envelope signals. The perhaps most significant insight offered by recent dynamic MEG FC studies of the resting-state is that FC states undergo switching at a much more rapid timescale, on the order of hundreds of milliseconds (Baker et al., 2014a, O'Neill et al., 2015, Vidaurre et al., 2016). Future work should apply recently developed methods for dynamic MEG FC to sleep data with an aim of resolving these timescales.

Chapter 6
Dynamic functional connectivity of sleep using a hidden Markov model of fMRI

6.1 *Introduction*

The data reviewed and presented thus far have shown how brain activity during wakefulness and sleep, as well as the contrast between them, can be characterised in light of large-scale functional networks. It has been argued and demonstrated that investigating sleep data with the tools and framework developed for the exploration of spontaneous brain activity of the waking resting-state may provide a more comprehensive and interpretable account of sleep-related brain processes than what is offered by a few EEG electrodes. Yet, studies carried out so far with respect to large-scale functional connectivity and its relation to sleep, in this thesis as well as in the wider literature, have taken as a common and natural starting point the traditional EEG-based staging of polysomnography (see 1.2.1). Analyses have thus been constrained by the data structure already given by the traditionally defined vigilance stages, and FC networks have been computed and compared from these data segments (see Chapter 2, Chapter 3, Chapter 4, and Chapter 5). While this approach can be used to derive large-scale correlates of already defined brain states, it fails to evade the constraints of the spatially non-specific EEG. If large-scale FC networks are truly functionally relevant for the brain in its transitions in vigilance, one would expect sleep to be identifiable from these networks alone. The machine-learning studies, reviewed previously (see 1.5.4), strongly suggest that the large-scale functional connectome does change in specific ways with vigilance (Tagliazucchi and Laufs, 2014, Altmann et al., 2016). However the fact that the EEG-based sleep scoring is used to train the employed pattern-recognising algorithms imposes a fundamental restriction on the resulting vigilance features extracted from the fMRI BOLD data.

It thus remains to be seen whether sleep can be identified solely based on large-scale functional networks. The question is whether a completely data-driven temporal decomposition of whole-brain recordings of vigilance transitions will reflect polysomnographic staging. This chapter aims to address this question. Importantly, this is not possible within the framework of static FC. Vigilance obviously unfolds in time, and hence does not align with the assumption of temporal stationarity that is inherent to static FC analyses (see 1.5.8). In order to unequivocally establish a relationship between vigilance and FC it is necessary to resolve FC networks dynamically, that is, to decompose the large-scale networks in time. The primary motivation for such an examination is to get a better handle on the functional meaning of these networks, in terms of their relevance to vigilance and sleep. The task may be likened to the scientific situation that Loomis and his colleagues were in when they made the first observations and descriptions of vigilance-related EEG changes (see 1.1). Here, however, single-channel EEG is replaced by whole-head coverage brain recordings, in the form of fMRI, and instead of relying of patterns distinguishable by eye, we aim to decompose the data in a data-driven way based on FC information. The pressing question is: will whole-head recordings and

networks temporally concur with the EEG-based staging, or will the increased spatial resolution and FC information possibly suggest an alternative categorisation of vigilance?

Two main challenges arise. First, the optimal way of resolving dynamics represents an outstanding issue within the field of large-scale connectomics (see 1.5.8). As such, a consensus has yet to be reached with respect to which method that provides the most meaningful temporal decomposition of large-scale brain networks. The common sliding-window approach comes with the obvious challenge and limitation of defining the size of the applied windows, which will affect the method's temporal sensitivity (Preti et al., 2016). This is in conflict with the aim of resolving FC dynamics in a completely data-driven way. We have thus chosen to turn to an alternative approach based on a specific implementation of a hidden Markov model (HMM), which has shown great promise at resolving large-scale networks in resting-state recordings with MEG (Baker et al., 2014a) and more recently with fMRI (Ryali et al., 2016, Vidaurre et al., 2017). The HMM is able to identify periods of stable patterns in multivariate sets of data based on probability distributions, and, crucially, can do this with no prior assumptions regarding temporal scales. Each stable period is modelled as one out of a finite number of HMM states, each of which represent a unique configuration of the large-scale network.

As required from a dynamic FC representation an HMM solution characterises fMRI data as a set of FC states, each with an associated timecourse describing when it is active and when it is inactive. This leads to the second challenge of the chapter: how to speak of a 'relationship' between FC dynamics and vigilance. Based on what is known about vigilance-dependent changes in static large-scale FC networks, there appears to be two likely scenarios for the outcome of a dynamic FC analysis of fMRI sleep data. i) evidence from the machine-learning studies (see 1.5.4) suggest that the large-scale network re-configures in quite specific ways as a function of vigilance stages (Tagliazucchi and Laufs, 2014, Altmann et al., 2016). This predicts that a dynamic decomposition should identify states of high specificity to each vigilance stage. ii) On the other hand, the demonstration that RSNs, such as the DMN and attention networks, can be found in all stages of sleep (see Introduction, 1.5.3) might suggest that instead of changing the patterns themselves, vigilance changes would manifest as changes to the dynamics, such as life times and switching, of these networks, as suggested in Brodbeck et al. (2012) and (Tagliazucchi et al., 2013c).

Applying a completely data-driven HMM with variational Bayes (VB) inference on the *Frankfurt* fMRI dataset, with no use of the sleep-scoring information, we hypothesised that the resulting temporal decomposition would show some form of relationship to the EEG-based, traditional sleep staging.

6.2 ***Methods***

6.2.1 *Data selection and pre-processing*

We considered fMRI recordings from 57 participants of the Frankfurt dataset (see Appendix, fMRI dataset 2 (Frankfurt)). All of these recordings included wakefulness and N1 sleep, 40 of these included N2 sleep, while a subset of 18 included all of the vigilance conditions: wake (W), N1, N2, and slow wave sleep (SWS). We included the full recordings of 1500 samples each (TR = 2.08 seconds). The simultaneously recorded EEG was scored by a sleep-scoring expert according to the AASM guidelines (Iber et al., 2007). Each 30-

second window of fMRI data was thus labelled as either wakefulness (W), N1, N2, or SWS. The EPI data were pre-processed as explained for the Frankfurt dataset in Chapter 2 (see 2.3.1).

6.2.2 *Brain parcellation and extraction of ROI timecourses*

As in Chapter 2 we chose the AAL atlas (Tzourio-Mazoyer et al., 2002) as a parcellation scheme to part the brain into 90 cortical and sub-cortical areas. An AAL template was transformed into the MNI space of the EPI data, and the FSL function *fslmeants* was used to extract the average timecourse over the voxels within each of the regions of interest (ROIs). Thus, each participant was represented by 90 averaged BOLD timecourses of 1500 samples.

6.2.3 *Fitting of hidden Markov model (HMM) to ROI timecourses*

Prior to fitting the hidden Markov model (HMM), we applied some additional pre-processing steps to the AAL timecourses. First, for each individual AAL timecourse within each participant, we removed the mean and divided by the standard deviation. This step served two purposes; 1) to reduce the effect of differences in mean level of BOLD activity between participants, and 2) to prepare the data for the following decomposition by principal component analysis (PCA). The demeaned and standardised timecourses were then concatenated across participants, yielding a data matrix of dimensions 90 by [1500 x 57], corresponding to the number of ROIs times the total number of samples. PCA was applied to the concatenated data to make the HMM-inference more robust (due to the lower dimensional space) and to remove noise in the data. By considering the amount of variance of the concatenated BOLD timecourses explained by each of the principal component, we chose to retain just above 90 per cent of the variance, which corresponded to D = 25 principal components (see Figure 6-1). Hence, in PCA space the data was now represented as a 25 by [1500 x 57] data matrix. Importantly, we kept the M x 90 mixing matrix of the PCA decomposition for subsequent back-projection of the HMM output to the original AAL space.

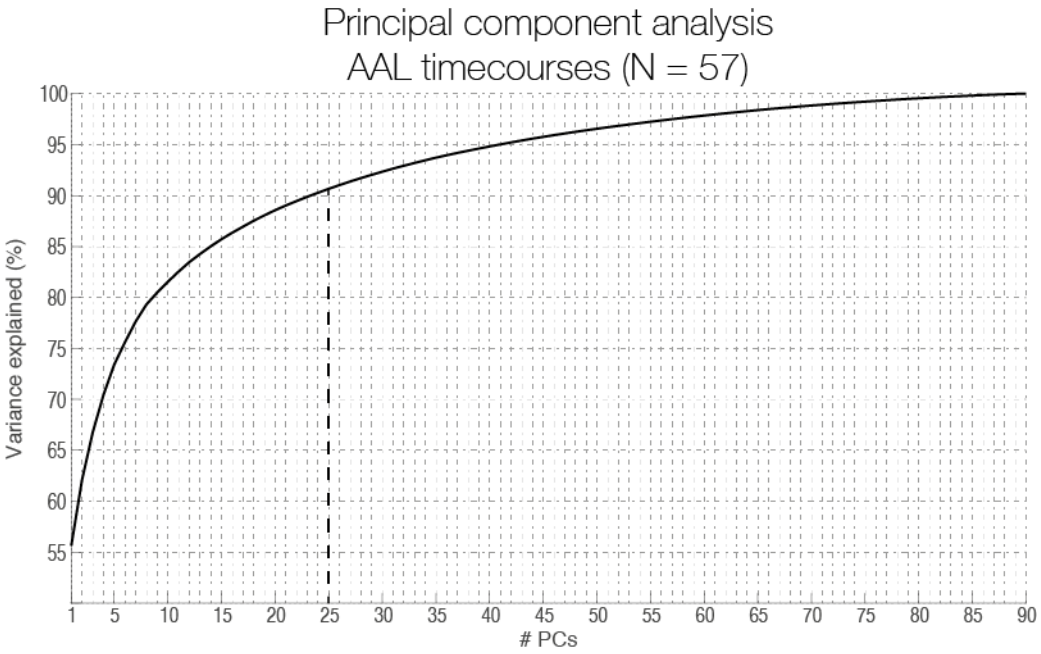

Figure 6-1 ***Variance explained in PCA decomposition***: The top 25 principal components, corresponding to just above 90 % of the variance of the demeaned and standardised BOLD timecourses, were retained and input to the HMM (marked by the dashed line).

Hidden Markov models are widely used as analysis tools for biomedical signals (Rezek and Roberts, 2005, Vidaurre et al., 2016). A main strength of this family of models is their ability to resolve non-stationary properties in multivariate datasets in an unsupervised fashion. The HMM works through two core assumptions: 1) that the observed data, **y**, may be explained as generated from a finite number, K, of *hidden* states, and 2) that the transition into a given of these states, s_t, depend only on the state that is active at the immediately preceding time point, s_{t-1}, i.e. that the system is *Markovian*. In a Bayesian framework this second property may be formulated as:

$$P(s_t|s_1 \ldots s_{t-1}) = P(s_t|s_{t-1}) = \pi_t \qquad \textbf{Eq. 6.1}$$

where π is the K x K state transition matrix.

Assume that the observations $\mathbf{y}_t$ are conditional only on the state at time t:

$$P(\mathbf{y}_t|s_1 \ldots s_t) = P(\mathbf{y}_t|s_{t-1}) \qquad \textbf{Eq. 6.2}$$

The joint distribution of the state sequence and the observations may then be factored as the full true posterior probability of the model:

$$P(\mathbf{y},\mathbf{s},\Theta) = P(s_0|\pi_0)P(\pi_0)\prod_t^T P(s_t|s_{t-1},\pi_t)\,P(y_t|s_t,\theta)P(\pi_t)P(\theta) \qquad \textbf{Eq. 6.3}$$

where $P(\pi_t)$, $P(\pi_0)$, and $P(\theta)$ may be set as non-informative priors. Θ contains the HMM parameters π_0, π_t, and θ. π_0 parameterises the initial state probability, $P(s_0)$, π_t is the aforementioned transition probability matrix, which determines the transition probabilities $P(s_t|s_{t-1})$, while θ parameterises the observation probabilities $P(y_t|s_t)$. The data, $\mathbf{y}$, was in our case represented by the 25 top principal components of the concatenated BOLD timecourses from 57 participants, which summed to a total length of 85,500 samples with each sample corresponding to the TR of 2.08 seconds.

The observation model, i.e. $P(y_t|s_t\theta)$, is a key component of an HMM implementation. In essence this dictates how the HMM should relate the hidden states to the observed variables, and as such which features in the data that are being modelled. The application of the HMM to fMRI BOLD timecourses presented here was inspired by the work of Baker et al. (2014a), in which an HMM was used to temporally resolve functional connectivity in source-reconstructed power envelopes from MEG. As in Baker et al. (2014a) we used a multivariate normal distribution as the observation model, such that each state, k, inferred by the HMM was represented by a distribution, θ_k, and its first and second order statistics, μ_k and Σ_k:

$$P(y_t|s_t = k, \theta) \sim \mathcal{N}(\mu_k,\Sigma_k) \qquad \textbf{Eq. 6.4}$$

This observation model is particularly interesting in the perspective of functional connectivity, since the estimation of covariance (or correlation) between observed variables represents the most widely used means of quantifying statistical dependencies between activity timecourses in neuroimaging, and hence is fundamental to our current understanding of functional connectivity (see Introduction, 1.5.3, 1.5.4). Compared to static approaches where covariance structures are estimated by averaging over long recording intervals, the HMM offers a non-stationary alternative, through which dynamic functional connectivity may be resolved. To date most attempts at estimating dynamic functional connectivity have relied on the application of sliding window averages. Duration and overlap of windows must be specified as free parameters in these approaches, and this tends to complicate their application, besides introducing assumptions regarding the life times of the dynamic states identified. A particularly noteworthy advantage of the HMM is the fact that the observation model and the transition model are fitted simultaneously. Due to the conditional dependence between these two terms (as given by Eq. 6.3) the temporal dynamics represented by the transition model helps constraining the spatial structures resolved by the observation model, by favouring those that repeat in a predictable way. As a consequence, the fitting of an HMM includes no prior assumptions regarding the temporal scale, e.g. life times, of the inferred states.

Following Baker et al. (2014a) and Vidaurre et al. (2016) variational Bayes (VB) inference was used on the HMM (Rezek and Roberts, 2005). The parameter outputs of the model used for further investigations included the multivariate normal distribution inferred for each state, i.e. an D x D covariance matrix, and a 1 x D vector of relative mean values, where D was the dimensionality of the PCA-decomposition (in our case D = 25). In order to make these state representations tractable to interpretation we used the mixing matrix, M, from the PCA decomposition to project back to the space of the AAL through the use of the D x 90 mixing matrix, M of the PCA decomposition:

$$\mu_k^{AAL} = \mu_k M^T \quad \textbf{Eq. 6.5}$$

$$\Sigma_k^{AAL} = M\Sigma_k M^T \quad \textbf{Eq. 6.6}$$

where superscript T denotes the matrix transpose.

Another important output from the HMM is the transition probability matrix, which is a K x K matrix describing the probability of transitioning from each state to each of the other states. This was used to analyse the transition dynamics between the inferred states of a given HMM solution.

In order to resolve the timing of the inferred states, we used Viterbi decoding to get the marginal posterior inference on the state variables $P(s_t|\mathbf{y})$. This gave us K state timecourses, $\mathbf{s}_k = s_t \ldots s_T$, describing the probability of each state, k = 1…K, at each time point, t = 1…T. Depending on the application of the information included in the state variables we used either a soft or a hard classification on the states. Soft decoding corresponded to the raw output of the HMM, whereby each state was represented as a probability of being active at each time point, while the hard classification, obtained through the Viterbi decoding, assigned *one* state per time point, t, by taking the most probable a posteriori state, u_t:

$$u_t = \arg_k \max P(s_t = k|\mathbf{y}) \quad \textbf{Eq. 6.7}$$

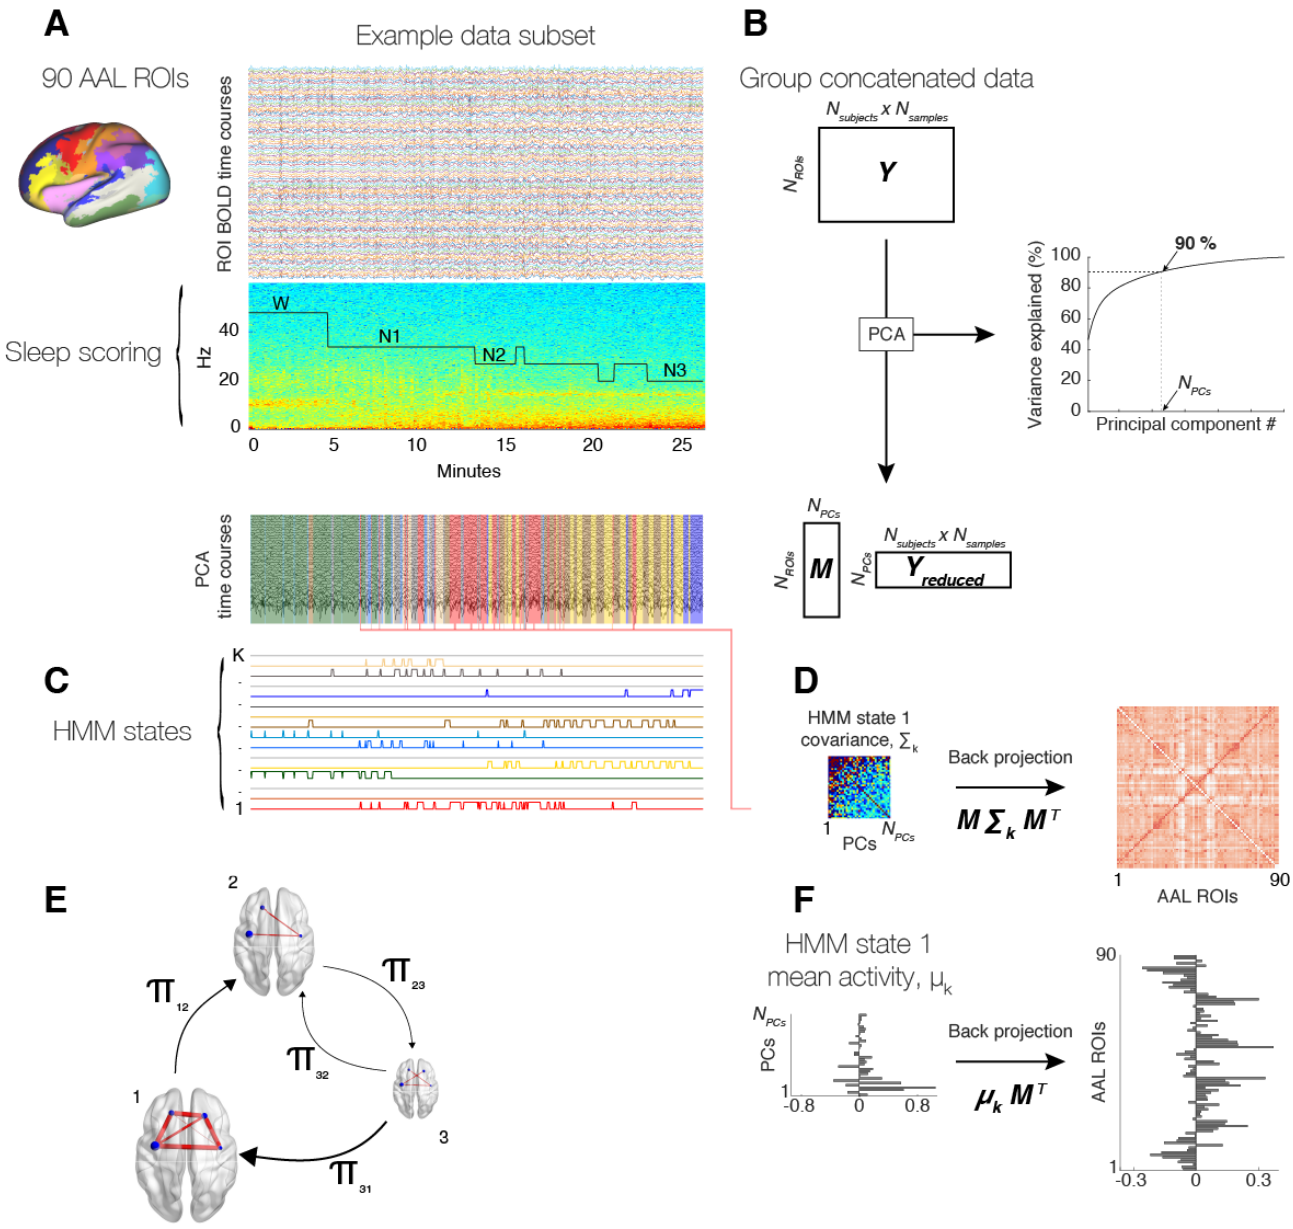

Figure 6-2 ***Overview of HMM analysis pipeline***: **A)** ROI timecourses were extracted by averaging BOLD signals across voxels within each of the 90 cortical and subcortical AAL areas. Each ROI timecourse was demeaned and normalised by its standard deviation. An example subset is shown and below it the corresponding sleep scoring, which was performed on concomitantly collected EEG. **B)** The data dimensionality was reduced using PCA, and ~90 % of the variance of the ROI timecourses was retained. **C)** The HMM was run on the PCA timecourses, resulting in K number of states, each characterised by: i) a specific multivariate Gaussian distribution of a mean distribution, μ_K (shown in **F**), and a covariance matrix, Σ_K (shown in **D**), ii) a timecourse describing its probability of being active at any point of the recording (see **C**), and iii) a transition probability, π, to any other state (shown in **E**). The state-specific mean distributions and covariance matrices could be back-projected to the space of the AAL by the use of the mixing matrix from the PCA decomposition.

6.2.4 *Estimating the optimal number of HMM states and relationship to sleep structure*

A main motivation for implementing the HMM in the current study was to limit the number of assumptions about the data we were aiming to characterise. As mentioned, the HMM is particularly appealing since its use for resolving functional connectivity dynamics does not rely on specifications of temporal scales through sliding windows. Our implementation of the HMM did however share with clustering-based methods (Allen

et al., 2014, Hansen et al., 2015), and mixture models (such as spatial ICA) the challenge of having to specify a number of states, K, to be inferred from the data. The VB framework used for the HMM inference equipped us with a measure of model evidence, estimated as the free energy. It is important to note, however, that a trade-off exists between simplicity/interpretability and explanatory power/redundancy. Thus, while an HMM with a high number of K might show higher model evidence, this improved fit might not reflect a gain in relevance nor interpretability. This caveat is generally true for the characterisation of spontaneous brain activity. It remains inherently difficult to know how many different brain states/networks one is to expect from a given resting-state recording. This is particularly clear in studies of the ill-monitored and poorly controlled resting-state, where in contrast to task-based paradigms, the researcher is left with few clues to the number of states to expect.

Sleep fMRI represents an exciting object of study in this respect, since the EEG-based organisation of the data into vigilance stages should be somehow reflected in the state decomposition. While the cognitive status within each vigilance stage remains inaccessible and uncontrolled, the fMRI data is characterised at a very fundamental level by the underlying sleep structure. As such we probed the relationship between the states output by the HMM and the EEG-based sleep scoring, in order to get a physiologically oriented measure of how well a given model order explained the data.

For a given model order, K, the HMM output K number of state timecourses, $\mathbf{s}_k$ (soft decoding). Using multivariate analysis of variance (MANOVA) we were able to relate the K x T matrix of $\mathbf{S}_K$ to the 1 x T sleep scoring vector, labelling each time point as one of the four stages of vigilance (W, N1, N2, or SWS). We used the built-in MATLAB function *manova1*, grouping $\mathbf{S}_K$ by the categorical sleep-scoring vector. This works by comparing the multivariate means of $\mathbf{S}_K$ across time with the sleep organisation. In order to test for significant multivariate relationships between the state timecourses and the sleep scoring, we embedded the MANOVA in a non-parametric permutation-based testing. The extent to which the MANOVA could relate the sleep-scoring grouping to the state timecourses may be quantified by the statistic Wilk's lambda, which is output by *manova1*. To define a chance-level for the said relationship we permuted the labels of the sleep-scoring vector within each participant. The number of occurrences of each vigilance condition as well as the sleep scoring time segments were preserved. For each participant we thus shuffled the labels of the existing segments of vigilance conditions. We did this a total of 1000 times, and a separate MANOVA was run between the shuffled sleep scoring vector and the original HMM state timecourses for each permutation and each model order of the HMM. Finally, we collected the Wilk's lambda for each of the permuted MANOVA runs.

Besides the free-energy and the MANOVA results the search for an optimal model order solution could also be helped by quantifying the level of redundancy in the HMM states for each model order. For this we chose to consider the covariance matrices associated with the HMM states. Specifically, we estimated the level of pairwise similarity between all state covariance matrices from each model solution. This was accomplished by means of the pairwise Pearson's correlation between the K number of matrices, similarly to the matrix comparisons used in Chapter 2 (see 2.3.5) and Chapter 3 (see 3.3.5). The mean across all

combinations of states (nchoosek = $\frac{K!}{2!(K-2)!}$) was taken as a proxy for the state redundancy of the given model order of the HMM.

To summarise, we consulted a number of metrics for the choice of optimal model order for the HMM: a purely data-driven summary represented by the Bayesian free energy; a physiology- and hypothesis-reflecting summary given by the MANOVA between the sleep staging and the HMM state timecourses; and a summary of the functional connectivity redundancy from the mean similarity of the HMM state covariances.

6.2.5 *Brain representation of HMM states*

Once back-projected from PCA to AAL space, each state inferred by the HMM was represented by a 1 x 90 mean vector and a 90 x 90 covariance matrix, as given by Eq. 6.4.

The mean distributions can be understood as specific configurations of activity increases and/or decreases in the given states. Note that these changes in activity were relative as they were inferred from the demeaned and standardised BOLD timecourses. For representation, the mean distributions were visualised on the cortical surface using the open-source Human Connectome Project software, *Connectome Workbench* and the function *wb_volume_to_surface_mapping*.

The covariance matrices expressed the inferred pairwise relationships between the 90 AAL timecourses for each of the HMM states. For matrix representations we converted the covariance matrices C_k, k = 1…K, into correlation matrices, R_k, using the built-in MATLAB function *corrcov*:

$$R_k = C_k \oslash (\sigma_k \sigma_k^T) \qquad \textbf{Eq. 6.8}$$

$$\sigma_k = \sqrt{C_k I} \qquad \textbf{Eq. 6.9}$$

where $\oslash$, T, I denoted element-wise division, matrix transpose, and the identity matrix respectively.

As will be clear from the results, the inferred covariance/correlation structures had a high degree of similarity across HMM states. Hence, many traditional representations of functional connectivity networks, such as degree or strongest connections would not necessarily be the most efficient way of characterising differences between the states. To ameliorate the dominance of the connectivity shared between states with the aim of extracting the information that was unique to each state-specific FCs, we computed a differential FC matrix, FC_{diff}, for each HMM state. For a given HMM state, k, FC_{diff} was defined as the mean FC matrix, averaged across all other states, subtracted from the state-specific FC matrix, R_k:

$$FC_{diff,k} = R_k - \frac{1}{K-1} \sum_{j \neq k}^{K} R_j \qquad \textbf{Eq. 6.10}$$

Positive values in FC_{diff} of a given HMM state would correspond to connections estimated as stronger than the average over the remaining state FCs, while a negative value would correspond to connections of

lower magnitude. Of note, we checked that the averaged FC matrices, i.e. the second term of Eq. 6.10, included positive values only regardless of the value of k. This ruled out the possibility of finding positive values in an FC_{diff} due to negative connection in the averaged FC matrices. The differential FC matrices were visualised by taking the 1% most positive and 1% most negative connections of each FC_{diff} and plot these in a glass-brain.

6.2.6 *Summary statistics*

The dynamic nature of the HMM allowed us to review a number of interesting features in the temporal domain of the inferred brain states. As such, following Baker et al. (2014a), we defined a number of summary statistics on the state timecourses. These included for each state: mean life time, fractional occupancy, and mean interval length:

$$\text{mean life time (k)} = \frac{\sum_t(u_t == k)}{N_{occurrences}(k)} \quad \textbf{Eq. 6.11}$$

$$\text{fractional occupancy (k)} = \frac{1}{T}\sum_t(u_t == k) \quad \textbf{Eq. 6.12}$$

$$\text{mean interval length (k)} = \frac{T - \sum_t(u_t == k)}{N_{occurrences}(k)} \quad \textbf{Eq. 6.13}$$

where $N_{occurrences}(k) = \sum_t\Big(\big((u_t == k) - (u_{t-1} == k)\big) == 1\Big)$, and $[u_t == k]$ was one if $u_t = k$ and zero otherwise. For this purpose we used the hard classification of the Viterbi decoding (see Eq. 6.7). Importantly, we computed these measure within participants in two ways. First, across the complete recording, and second, restricted within individual vigilance stages, as defined by the sleep scoring. For the full-recording estimates T and t would be indexed according to length of the participant-specific recordings, i.e. 1500 samples. For the vigilance-specific estimates T_V would be defined (within each participant) for each vigilance condition, V, according to the sleep-scoring vector, Ss:

$$T_V = T(Ss == V),\ V = \{W, N1, N2, SWS\} \quad \textbf{Eq. 6.14}$$

In addition to the measures described above (in Eq. 6.11, Eq. 6.12, and Eq. 6.13) we were interested in quantifying some general dynamic features within the vigilance conditions. First, by collapsing over HMM states, and using the hard classification on the Viterbi decoding, we were able to quantify the amount of state switching for each vigilance condition, V, as the number of state transitions per vigilance-specific duration (within each participant):

$$\text{switching(V)} = \frac{\sum_{t_V}\big(u_{t_V} - u_{t_V-1} \neq 0\big)}{T_V} \quad \textbf{Eq. 6.15}$$

The term $\left[u_{t_V} - u_{t_V-1} \neq 0\right]$ defines a 1 x T_V vector of ones and zeros, trans(V) where ones correspond to time points at which a HMM-state transition occurs within vigilance conditions. We used trans(V) to index the HMM-state transitions within each vigilance condition. For each transition time point, t_{Vtrans}, we defined a measure of state exploration by computing the Pearson's correlation coefficient of the mean activity distribution for the state preceding the transition and the state following the transition:

$$\text{state exploration (V)} = \frac{1}{\sum \text{trans(V)}} \sum_{i}^{\text{trans(V)}==1} \left(\frac{1}{N_{ROIs}} \sum_{j}^{N_{ROIs}} \left(\frac{\overline{A_j - \mu_A}}{\sigma_A} \right) \left(\frac{B_j - \mu_B}{\sigma_B} \right) \right) \quad \textbf{Eq. 6.16}$$

where A and B are the absolute of the mean activity distributions $\mu_{u(i-1)}^{AAL}$ and $\mu_{u(i)}^{AAL}$ respectively, and μ and σ are the mean and standard deviation of these distributions. N_{ROIs} is the number of AAL regions. The state exploration measure was devised to capture the amount of diversity involved in state transitions within vigilance conditions.

As a final level of description, we re-visited the analysis of Chapter 2 of SC-FC similarity (see 2.3.5). By means of Pearson's correlation we compared each of the K inferred HMM state FC matrices to the group-averaged SC matrix derived in Chapter 2:

$$\text{SC-FC (k)} = \frac{1}{N_{connections}} \sum_{n}^{N_{connections}} \left(\frac{\overline{SC_n - \mu_{SC}}}{\sigma_{SC}} \right) \left(\frac{\Sigma_k{}^{AAL}{}_n - \mu_{\Sigma_k{}^{AAL}}}{\sigma_{\Sigma_k{}^{AAL}}} \right) \quad \textbf{Eq. 6.17}$$

where $N_{connections}$ is the number of connections in the AAL networks, i.e. $90^2 - 90$.

6.3 ***Results***

6.3.1 *HMM performance and MANOVA*

In this section we describe our evaluation of the HMM's performance on the concatenated, sleep-scored fMRI data from 57 participants. We do so from three different perspectives; first, by considering how well the HMM described the data in general in terms of the free energy on the Bayesian inference; second, by asking if any specific HMM solution (i.e. model order K) showed a particularly close correspondence with the sleep scoring as quantified by the MANOVA; and third, by examining the similarity/redundancy of the HMM state FC matrices. Based on these, we will argue that a model order of K = 19 represents a good trade-off between richness and redundancy.

Figure 6-3 shows the various summaries of the overall HMM performance. The minimum free energy is plotted as a function of model order K, in Figure 6-3A, and quantifies the error of the model. A lower value of minimum free energy signals a better fit to the data. As evident from the plot, we ran the HMM with model orders from 4 to 45 in steps of one, and we observed a monotonically decreasing minimum free energy, suggesting that even higher model orders might in fact perform even better on this strictly data-driven measure.

To aid the choice of which model order K to further investigate, we added a hypothesis-driven measure to each HMM run, by quantifying the relationship between the state timecourses of the HMM and the EEG-based sleep scoring using a permutation-embedded MANOVA. Although, the HMM was run in a way that was naïve to any sleep scoring, we hypothesised that the output would somehow reflect the sleep structure. The plots in Figure 6-3B and C respectively show the results of this MANOVA analysis for the whole data (N = 57), and for a subset of the HMM solution, corresponding to the participants that included all 4 vigilance conditions (N = 18). For both cases the MANOVA showed higher correspondence when the HMM state timecourses were grouped by the original sleep scoring (solid blue line), than by the permuted sleep scoring (dashed blue line, with error zones corresponding to the standard error across permutations). The measure Wilk's lambda, Λ, tests whether the means compared by the MANOVA (in this case means of the HMM state timecourses) align in a certain dimension. Due to the four possible values of the sleep scoring (W, N1, N2, SWS) we tested whether the means would group in a three-dimensional space. The quantity $(1 - \Lambda)$ is often interpreted as the proportion of variance of the dependent variables (the HMM state timecourses in this case) explained by the model effects (the sleep scoring). As such, low values of Λ signals a better MANOVA fit. As evident for both the N = 57 and the N = 18 case, the MANOVA describes a closer relationship for higher model orders, which seemed to stabilize from around K = 19.

A third summary of the HMM solutions is provided in Figure 6-3D, where the state redundancy is plotted as a function of model order K. The state redundancy was defined as the overall correlation between the FC matrices of the HMM states for a given model order. A decreasing trend can be seen from K = 4 to around K = 15, above which it fluctuates about a correlation of $r \sim 0.65$.

As a final level of overall description of the HMM solutions, Figure 6-3E plots the average state life time as a function of model order. The average state life time decreased with the number of states from a value around 30 seconds for K = 4. Interestingly, however, it appears to stabilise somewhat from around K = 14 at a level between 15 and 20 seconds. The fact that the mean life time of the HMM states stagnates would suggest that an increase in the model order did not result in splitting of states into shorter lived ones, but rather the sporadic occurrence of new states replacing states that had been re-occurring at lower model orders.

In an attempt to synthesise these previously presented summaries of the model order parameter space, and in the absence of any obvious, optimal solution we chose a model order of 19 as a good trade-off between richness and redundancy. For K = 19, the MANOVA showed high difference between the original and permuted sleep scoring, the state redundancy was in the lower end, and the life time curve was only just beginning to stagnate.

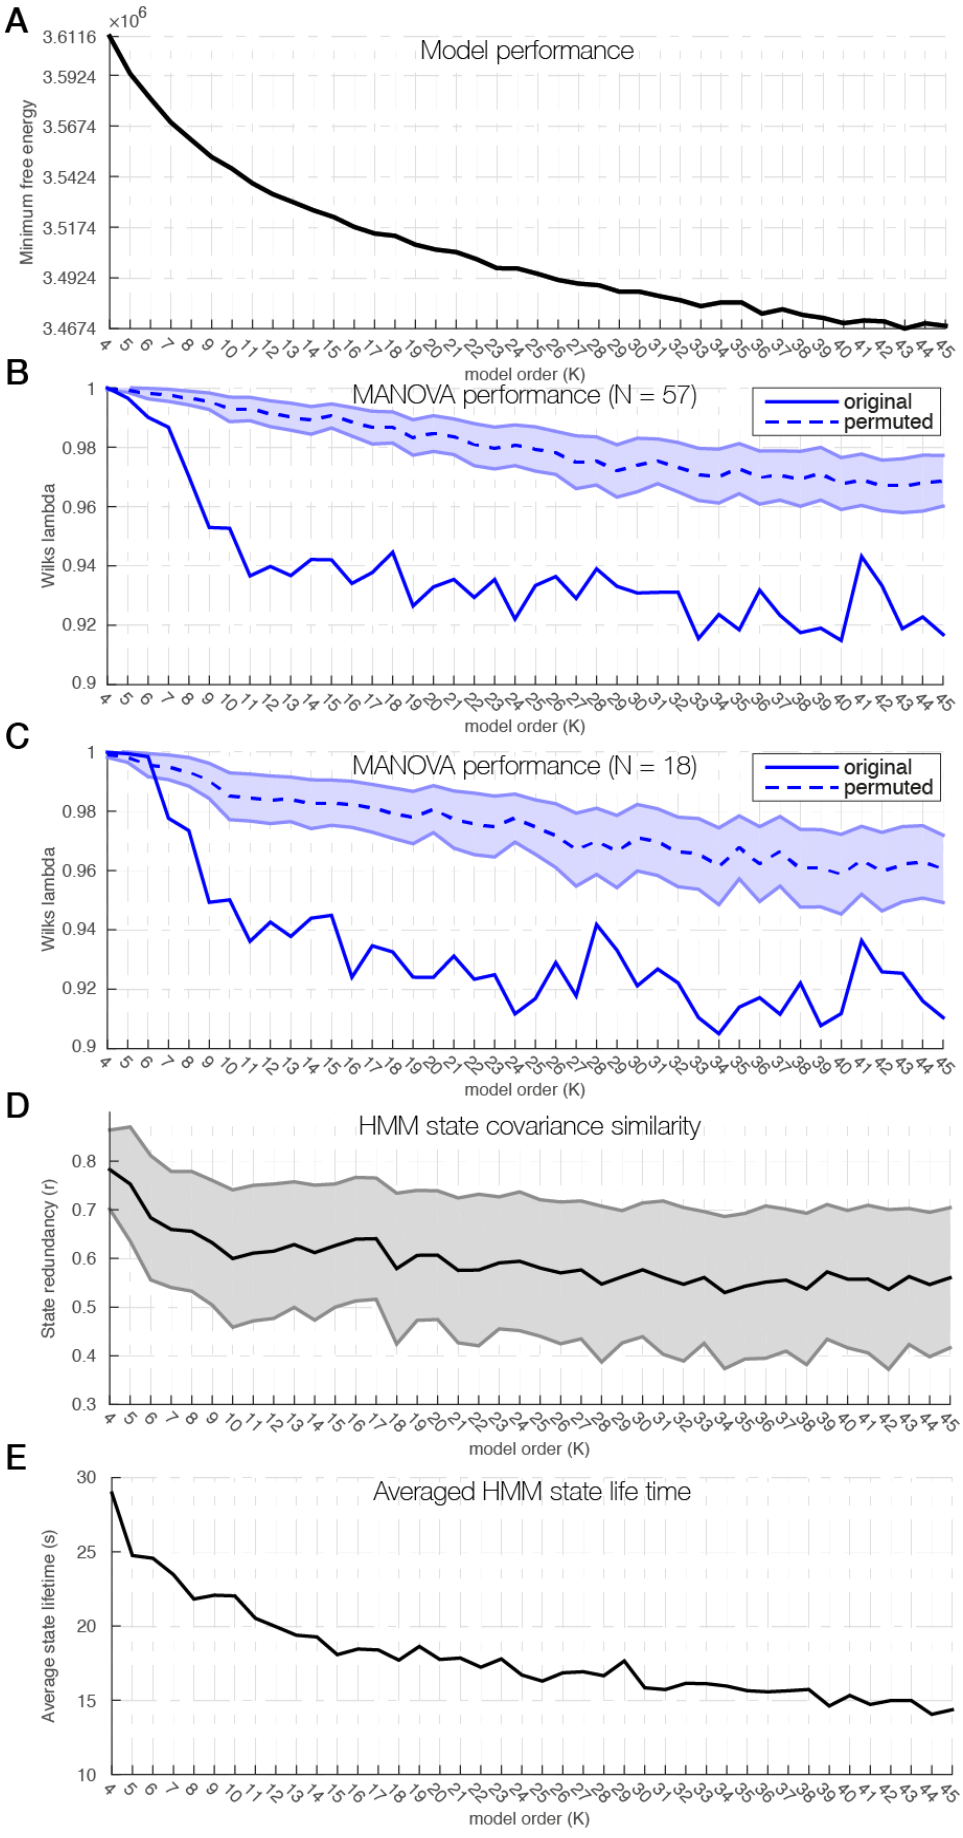

Figure 6-3 ***Summary measures of HMM solutions across a range of model orders***: **A)** The minimum free-energy on the VB inference decreased monotonically for increasing model orders. **B)** Wilk's lambda is an output of the MATLAB

function *manova1*, and describes how well the HMM state timecourses could be grouped according to the EEG-based sleep scoring. Lower values correspond to a better fit. Here Wilk's lambda is shown for the full HMM solution of all 57 participants and across a range of model orders. The solid blue line shows depicts the output of the MANOVA run with the original sleep scoring, while the dashed blue line corresponds to the average Wilk's lambda from 1000 MANOVAs where the sleep scoring was randomly permuted. Error zones correspond to the standard deviation of the permuted cases. **C)** Equivalent to **B** but where the MANOVAs were run only on the subset of participant that included all four vigilance stages. Note how the original sleep scoring expressed significantly better fitting with the HMM state timecourses in both **B** and **C**, and how Wilk's lambda appear to stagnate somewhere between $10 < K > 20$. **D)** Similarity of HMM state covariances, computed as the Pearson's correlation between all pairs of states. The solid line shows the average while error zones correspond to the standard error. **E)** The development of state life times averaged across all HMM states within model orders.

6.3.2 *HMM state timecourses and their relationship to EEG-based sleep scoring*

The following sections present the results of the HMM of model order K = 19. In Figure 6-4A the HMM state timecourses are plotted for the 18 participants that included all vigilance conditions. Immediately below the HMM state timecourses are included the independent EEG-based sleep scoring for the same data. Each HMM state timecourse represent the probability of that state being active at any given point in time (soft decoding). Even from this simple presentation a clear interdependence between the HMM states and the vigilance stages is evident, such that certain HMM states are found more predominantly during certain stages of vigilance. Some clear examples include HMM states 8, 1, 6, and 16, the probability of which appear to follow periods of wakefulness, N1, N2, and SWS respectively. The MANOVA results, presented above, can be seen as a quantification of this visible vigilance-dependence of the HMM states. Notably, this vigilance-dependence appeared consistent across participants. With a few exceptions (see later section, 6.4.6), the HMM states showed little participant-specificity.

The FC structures of each of the 19 HMM states are shown in Figure 6-4B in the form of correlation matrices. Note that these are direct outputs of the HMM. In line with the state redundancy plotted in Figure 6-3D, a high degree of shared connectivity was evident in the FC matrices, which expressed patterns of fairly high similarity. The mean state redundancy for K = 19 was r = 0.606 (± SD 0.133, $N_{nchoosek}$ = 171). Eight of the nineteen states (2, 7, 8, 9, 10, 14, 16, and 18) included negative correlations. These were however small in magnitude compared to the positive correlations, as reflected in the colour scaling of the matrices.

Of relevance to the SC-FC similarity results, reviewed below, the group-averaged SC matrix is included at the bottom of Figure 6-3B.

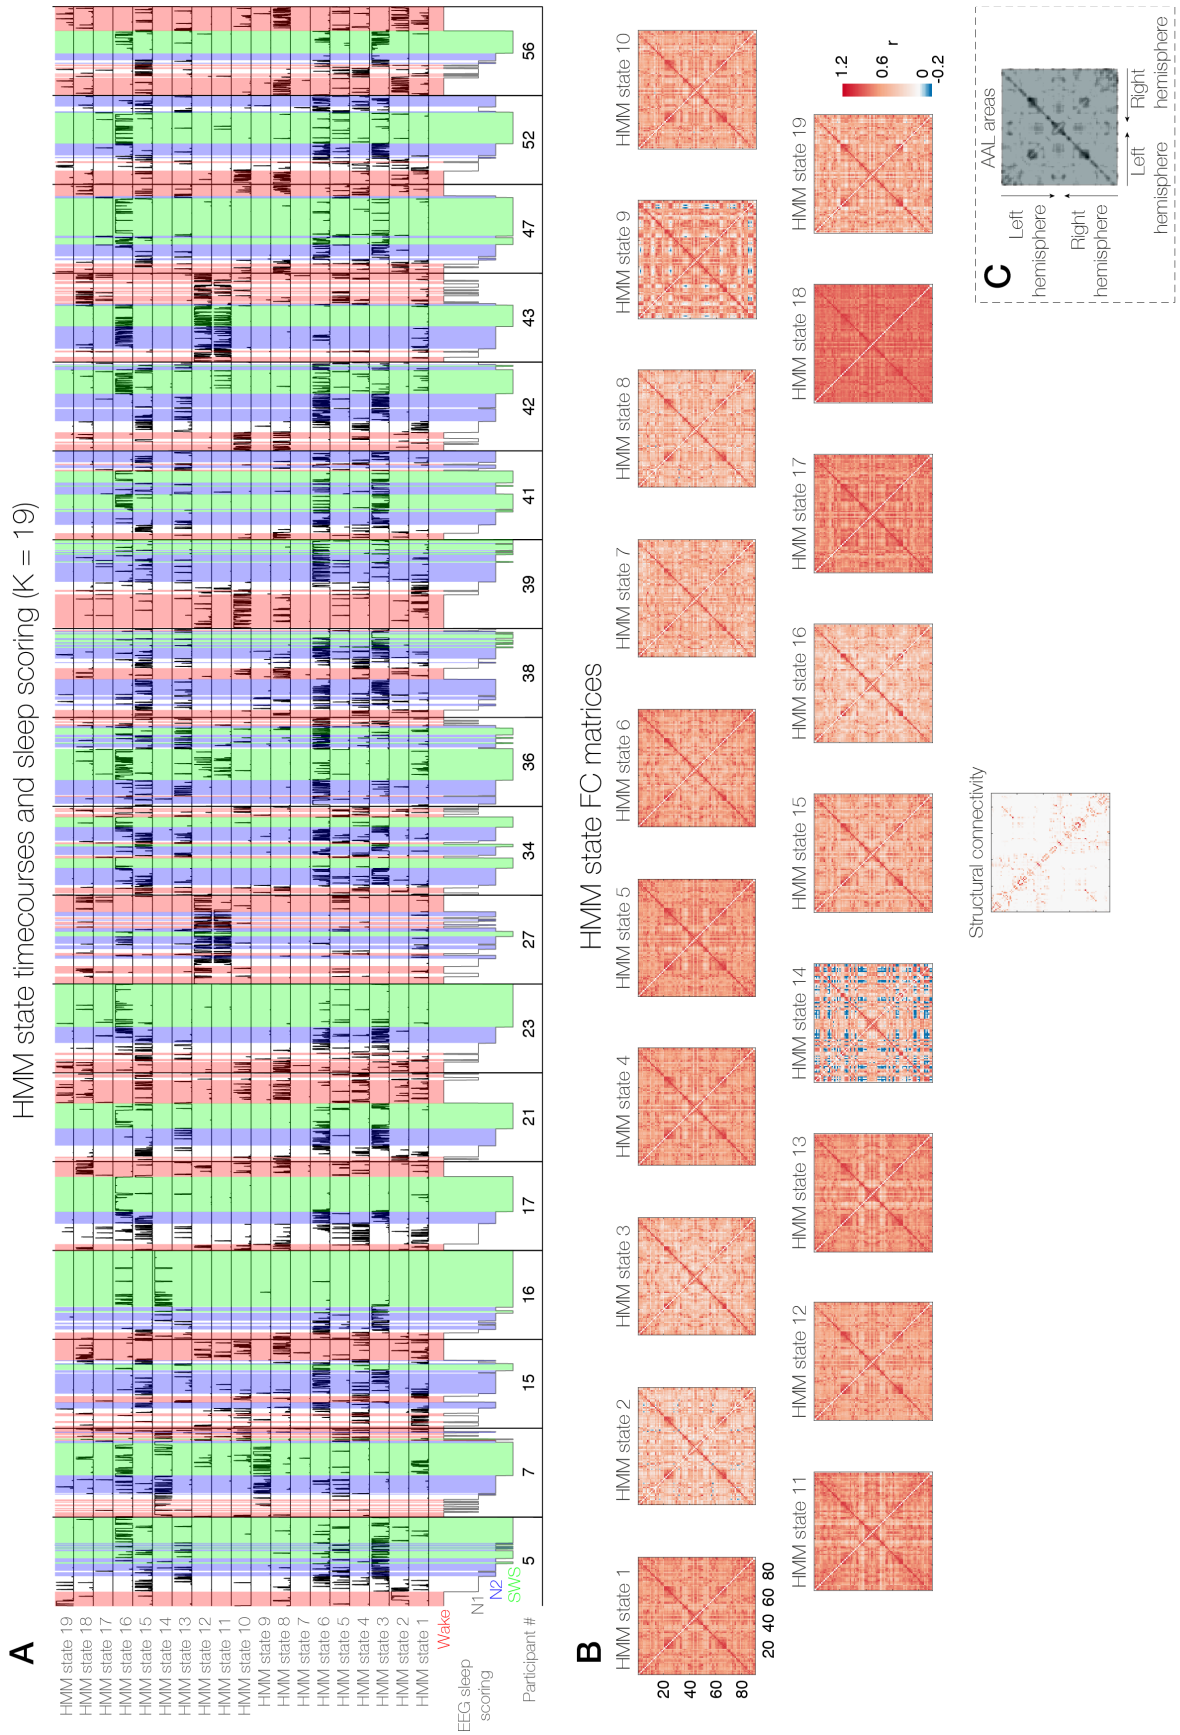

Figure 6-4 ***State timecourses and FC matrices for HMM model order K = 19***: **A)** 19 HMM state timecourses describing the probability of activity at each sample point of the fMRI sessions of the 18 participants that reached all four stages of vigilance. Below the HMM state timecourses are shown the corresponding EEG-based sleep scoring. The coloured overlay show periods scored as wakefulness (red), N1 (white), N2 (blue), and SWS (green). Note how many of the HMM state timecourses express a strong temporal correspondence with specific stages of vigilance. **B)** FC matrices for each of the 19 HMM states. At the bottom is included the group-averaged SC matrix from an independent dataset. **C)** Schematic representation of the organisation of the AAL ROIs in the FC matrices.

6.3.3 *Mean life time, fractional occupancy, and mean interval length*

Thus far, we have presented summaries of the HMM solutions, attempting to quantify their relationship with the EEG-based sleep scoring through a MANOVA analysis, settled at K = 19 as our model order of choice, and shown the reliability of the HMM state timecourses across participants. In order to further investigate the vigilance-dependence of the HMM solution, we will now consider a number of summary statistics of the HMM, calculated within EEG-defined vigilance conditions.

Figure 6-5 shows the summary statistics; mean life time, fractional occupancy, and mean interval length, within vigilance conditions for the 18 participants that included all four vigilance conditions. The fractional occupancy plots in the middle column of Figure 6-5 illustrate how different HMM states dominated during different vigilance stages. HMM states 8 and 10 thus showed high fractional occupancy during wakefulness; HMM state 15 during N1 sleep; HMM states 3 and 6 during N2 sleep; and HMM state 16 during SWS. Further it is worth noting how HMM states with high fractional occupancy in N2/SWS were generally low in wakefulness, and vice versa. This kind of mutual exclusiveness was less apparent between 'neighbouring' vigilance conditions, i.e. W/N1, N1/N2, N2/SWS, suggesting that the HMM might not just have captured the individual vigilance conditions, but also the sequencing between them.

The vigilance-sensitivity of the HMM states suggested by the fractional occupancy plots can help structure the information of mean life time (leftmost column of Figure 6-5) and mean interval length (rightmost column of Figure 6-5). Since the fractional occupancy highlights certain HMM states for each vigilance condition, we focus on the mean life time, and mean interval length of these states. Thus, within wakefulness, HMM states 8 and 10 had mean life times of 16.0 seconds (± SE 1.9, N = 18) and 11.6 seconds (± 2.3, N = 18) respectively; within N1 sleep, HMM state 15 had a mean life time of 12.5 seconds (± SE 1.0, N = 18); within N2 sleep states 3 and 6: 28.6 seconds (± SE 4.8, N = 18), and 16.7 seconds (± SE 1.9, N = 18); and within SWS sleep state 16: 64.7 seconds (± SE 16.3, N = 18). Although this is not an exhaustive description of the states present within each of the vigilance conditions, it would at least seem that the dominant states (as measured by fractional occupancy) for the vigilance conditions tended to have longer life times with the deepening of sleep.

An equivalent, fractional occupancy-guided overview may be given of the mean interval lengths of the HMM states. Thus, within wakefulness states 8 and 10 had mean interval lengths of 134.3 seconds (± SE 37.2, N = 18), and 154.2 seconds (± SE 29.0, N = 18); within N1 sleep state 15: 146.0 seconds (± SE 48.5, N = 18); within N2 sleep states 3 and 6: 183.9 seconds (± 70.9, N = 18), and 91.4 seconds (± SE 23.0, N = 18); and within SWS sleep state 16: 79.2 seconds (± 19.2, N = 18). As would be expected these HMM states with highest fractional occupancy within the given vigilance conditions had low mean interval lengths compared to HMM states with low fractional occupancy in the same condition, meaning that the former states were revisited with shorter delays than the latter.

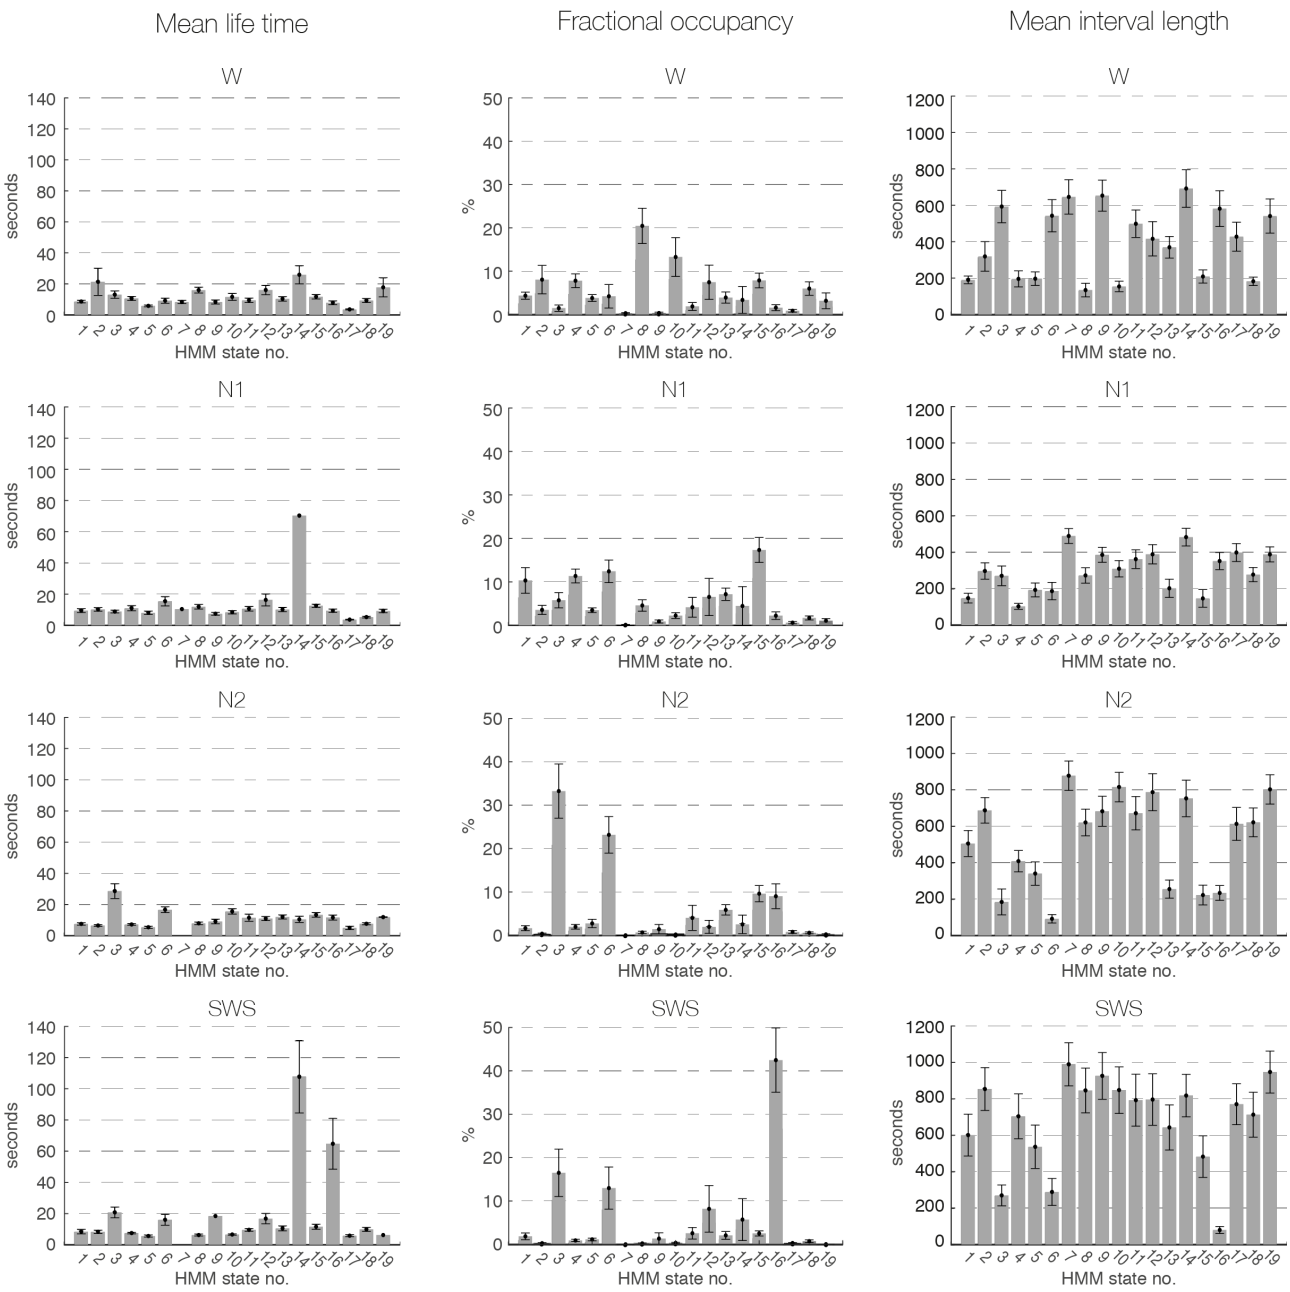

Figure 6-5 ***Summary statistics of HMM states for K = 19 grouped within vigilance conditions***: From left to the right are shown mean life time, fractional occupancy, and mean interval length of each of the 19 HMM states within each of the four vigilance conditions (first row: wakefulness, second row: N1, third row: N2, fourth row: SWS). Bars show the mean across the 18 participants that included all four vigilance conditions and error bars correspond to the standard error. Note how certain HMM states dominated, in terms of high fractional occupancy, during each of the vigilance conditions. Notice also how the fewer dominant states of N2 and SWS compared to wakefulness and N1, and their longer mean life times.

6.3.4 *Specificity of HMM states to vigilance*

While the fractional occupancies, reviewed above, suggest that each of the vigilance conditions were dominated by different HMM states, this measure does not answer the reverse question, namely how likely a given HMM state is to occur during a specific vigilance condition. A high fractional occupancy during a given vigilance condition may thus be seen as reflecting vigilance-sensitivity of an HMM state, however this does not imply vigilance-specificity. Equally, some HMM states may be very specific to a given vigilance stage, while only occupying small proportions of it. Figure 6-6A gives an alternative representation of the

mean life times of the HMM states. Here, the mean life time estimations were not restricted within vigilance conditions, but rather computed across the entirety of the recordings. Each HMM-state-specific bar is colour coded according to the probability of finding this state in each of the four vigilance conditions. A bar dominated by one colour would correspond to a state with high specificity towards the vigilance condition represented by that colour. Conversely, a bar coloured by an equal mix of all four colours would imply that a state had low specificity for any one vigilance condition. The *specificity* is an essential complementary measure to the fractional occupancy when evaluating the importance of the HMM states for the vigilance conditions. Hence, whereas the states highlighted for their fractional occupancy in wakefulness (i.e. states 8 and 10), in N2 (i.e. states 3 and 6), and in SWS (i.e. state 16) also showed high specificities for the same vigilance conditions, this was not unequivocally the case for HMM state 15, which showed high fractional occupancy in N1. In fact, state 15 was equally likely found during N2 sleep as during N1 sleep. The HMM state with highest probability for N1 sleep was state 1, even though this state did not show up on the fractional occupancy plots. In a similar way, HMM states 2 and 18, showed high specificity for wakefulness.

Another interesting thing to note from the plot in Figure 6-6A is how HMM states with high probability of belonging to either N2 or SWS sleep clearly expressed longer mean life times.

Figure 6-6B shows a clustering dendrogram derived from the MANOVA run for the HMM at K = 19. From this it is evident that the HMM did indeed capture the natural sequence of sleep. Hence, in order to get from wakefulness or N1 sleep to SWS in the 'HMM state space', one would have to pass through N2 sleep.

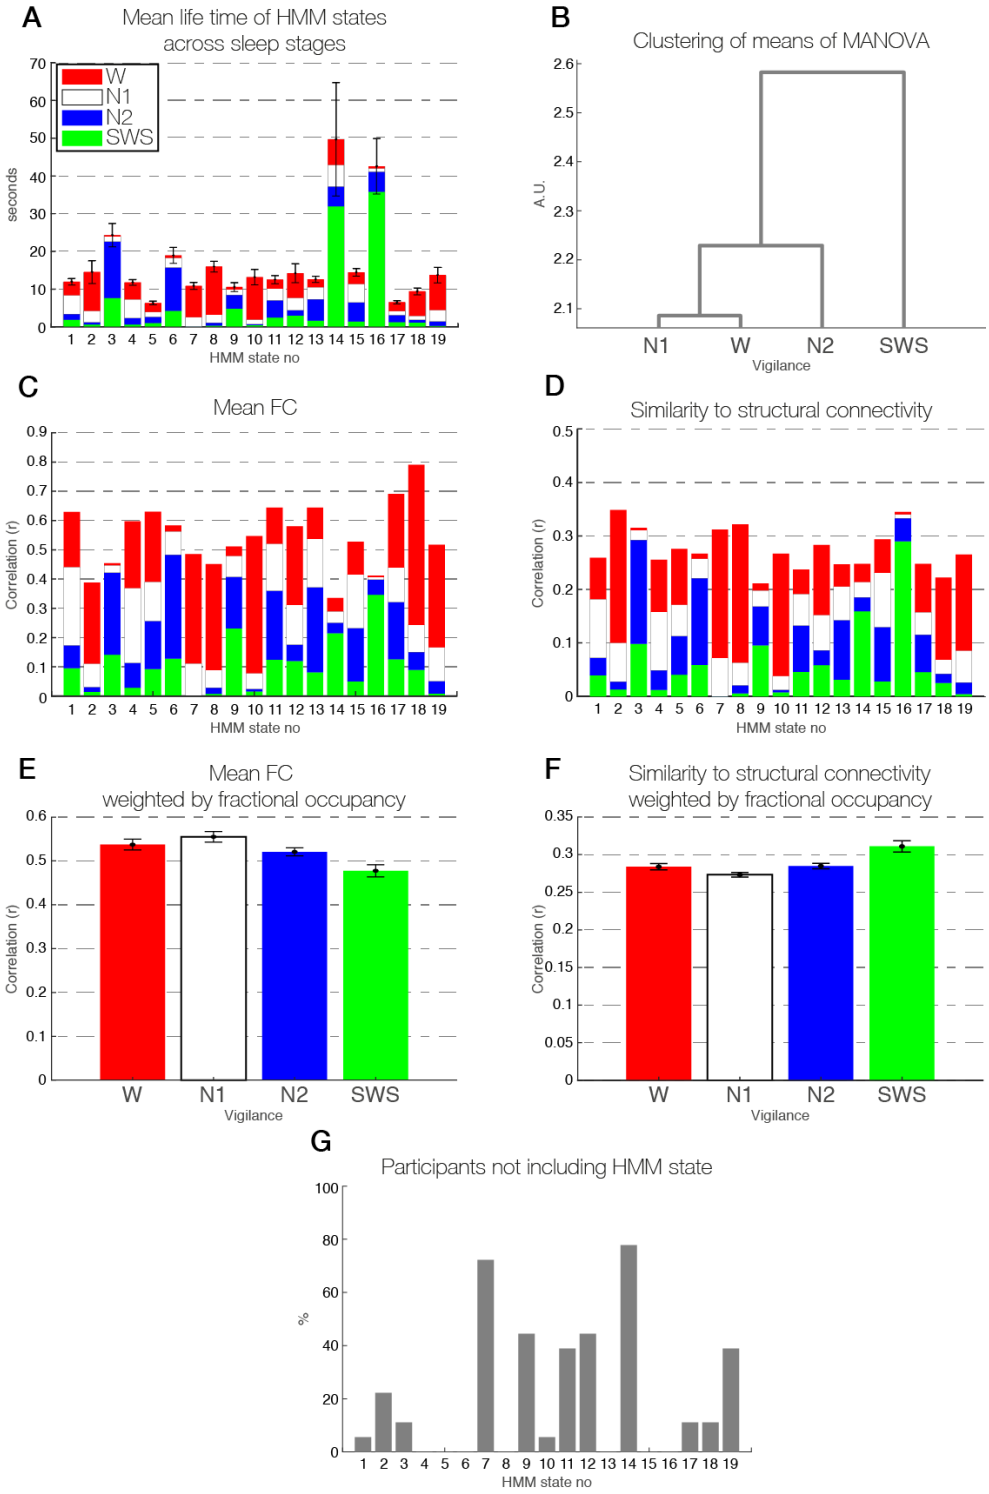

Figure 6-6 ***Summary of mean life times, mean FC, and SC-FC similarity*** **A)** Mean life time of HMM states across vigilance conditions. Each bar is colour-coded according to the probability of finding it in each of the four vigilance conditions. Note that states with high specificity of N2 and SWS showed higher mean life times than those with high specificity for wakefulness and N1. **B)** Clustering dendrogram output by the MANOVA analysis for K = 19 on the 18 participants that included all vigilance conditions. This shows the relative distance of the EEG-based vigilance stages in the space of the HMM states. Wakefulness and N1 are grouped closely together, while they are far away from SWS, and N2 are situated in between. This is in agreement with the natural sequence of a hypnogram and reflects the good temporal correspondence between the 19 HMM states and the categorical sleep staging. **C)** Mean FC of each HMM state. Note how HMM states with high mean FC tended to be less similar to SC. **D)** SC-FC similarity of each HMM state. **E)** Mean FC within vigilance conditions computed by taking the mean of the state-specific mean FC values weighted by the vigilance-specific fractional occupancies. **G)** Equivalent to **F** but showing the vigilance-specific SC-FC similarity values. **G)** Overview of how common each HMM state was across the 18 participants. Note that a number of HMM states, 11 and 14 in particular, were absent in many participants, suggesting that they modelled participant-specific features.

6.3.5 *Dynamic mean FC and SC-FC similarity*

Each HMM state was associated with a covariance matrix, describing its FC pattern and profile (see Figure 6-4B). Taking the mean across the non-diagonal elements of each of these matrices yielded the state-specific mean FC values, plotted in Figure 6-6C. Again each bar has been coloured according to the probability of finding the given HMM state during each of the four vigilance stages. In a similar way, each HMM state FC matrix could be compared to the group-averaged SC matrix (see bottom of Figure 6-4B) to get the state-specific SC-FC similarity. The results of this are shown in Figure 6-6D. As a sanity check, and for the sake of relating the HMM results to those of the static FC analyses carried out on the same *Frankfurt* dataset in Chapter 2, we summed up the mean FC values and SC-FC similarities within each of the four vigilance stages. This was done by weighting each of the HMM state values by the vigilance-specific fractional occupancies (as given by the plots in the middle column of Figure 6-5) and adding these together. The results are shown in Figure 6-6E–F for mean FC and SC-FC similarity respectively. Comparing Figure 6-6E to Figure 2-1D, and Figure 6-6F to Figure 2-3D it becomes clear that the HMM solution corresponds well with a static FC analyses, when averaging over time. Please note that the mean FC results shown in Figure 2-1D of Chapter 2 are z-transformed, whereas the results in Figure 6-6E are not. This means that the comparison of mean FC values only applies in a relative sense across vigilance stages, and not numerically.

6.3.6 *"Outlier" states*

As alluded to in section 6.4.2, and as observed in Figure 6-4A, the HMM states mixed well across participants. Exceptions were present, however, and a few states appeared to model participant-specific features. In Figure 6-6G is provided an overview of the proportion of participants that did not include each of the 19 HMM states. From this it is clear that in particular HMM states 7 and 14 occurred only in a small subset (less than 30%) of the participants. Tracking its timecourse in Figure 6-4A reveals that HMM state 7 occurred very rarely, and when it did occur it was brief (see also fractional occupancy and mean life time in Figure 6-5). In contrast, HMM state 14 occurred very specifically in two participants in Figure 6-4A. Similarly, HMM states 11 and 12 showed high affinity for the same two participants, which were occupied almost exclusively by these two states. Returning, to the issue of model order selection, it was our experience that the main effect of increasing the number of states, K, was to induce more of these, in our context unwanted, cases of participant-specific states. This could also explain why the average mean life time did not change markedly when increasing the model order above K~14.

6.3.7 *Transition matrix and dynamic interactions between HMM states*

The HMM not only assigns a state to each time point of the fMRI data input to the model, it further models the probability of transitioning between any of these HMM states. The transition probability matrix quantifies the pairwise probability of switching from any two HMM states. Ultimately, for our purpose, this offers yet another way of representing the space of the HMM states in light of the vigilance conditions under investigation. In Figure 6-7 we present an account of the HMM state transitions. Figure 6-7A shows the 19 x 19 transition matrix output by the HMM, while Figure 6-7C visualises the 8% most likely transitions. The

transition map in Figure 6-7B was constructed using the open-source software *MapEquation* (http://www.mapequation.org). Each state is represented by a circle plot coloured according to the proportion of time it is active in each of the four vigilance conditions. Transitions between states are depicted as directed arrows scaled according to the probability of the transitions. Notably, the transitions formed a general path moving from wakefulness in the top-right over N1, N2, towards SWS in the lower-left of the figure. Furthermore, wakefulness appeared to fall into two divisions; one including HMM states 2, 10, and 8, and another comprising HMM states 5, 17 and 18. Notably these two divisions were not connected through transitions, but parted by states with high activity in N1 and N2 sleep, such as HMM states 15, 4, and 13.

It is also worth noting that only a subset of the total 19 HMM states are shown here. Hence states 7, 9, 14, and 19 did not survive the thresholding (top 8%) of the transition matrix. These were all states, which occurred only in a small proportion of the participants as summarised in Figure 6-6G. States 11 and 12 were also characterised by a low mixing across participants, but a strong transition probability between each other. This is why they appear as disconnected from the rest of the main network in the transition map.

The transition map representation supports what was suggested by the vigilance sensitivity and specificity, and indeed already by the MANOVA, namely that the HMM was able to pick up the sleep structure inherent to the BOLD timecourses. At the same time, it offers a novel account, at the group-level, of how whole-brain BOLD activity transitions through states, as characterised by mean and covariance, from wakefulness and through the stages of non-REM sleep.

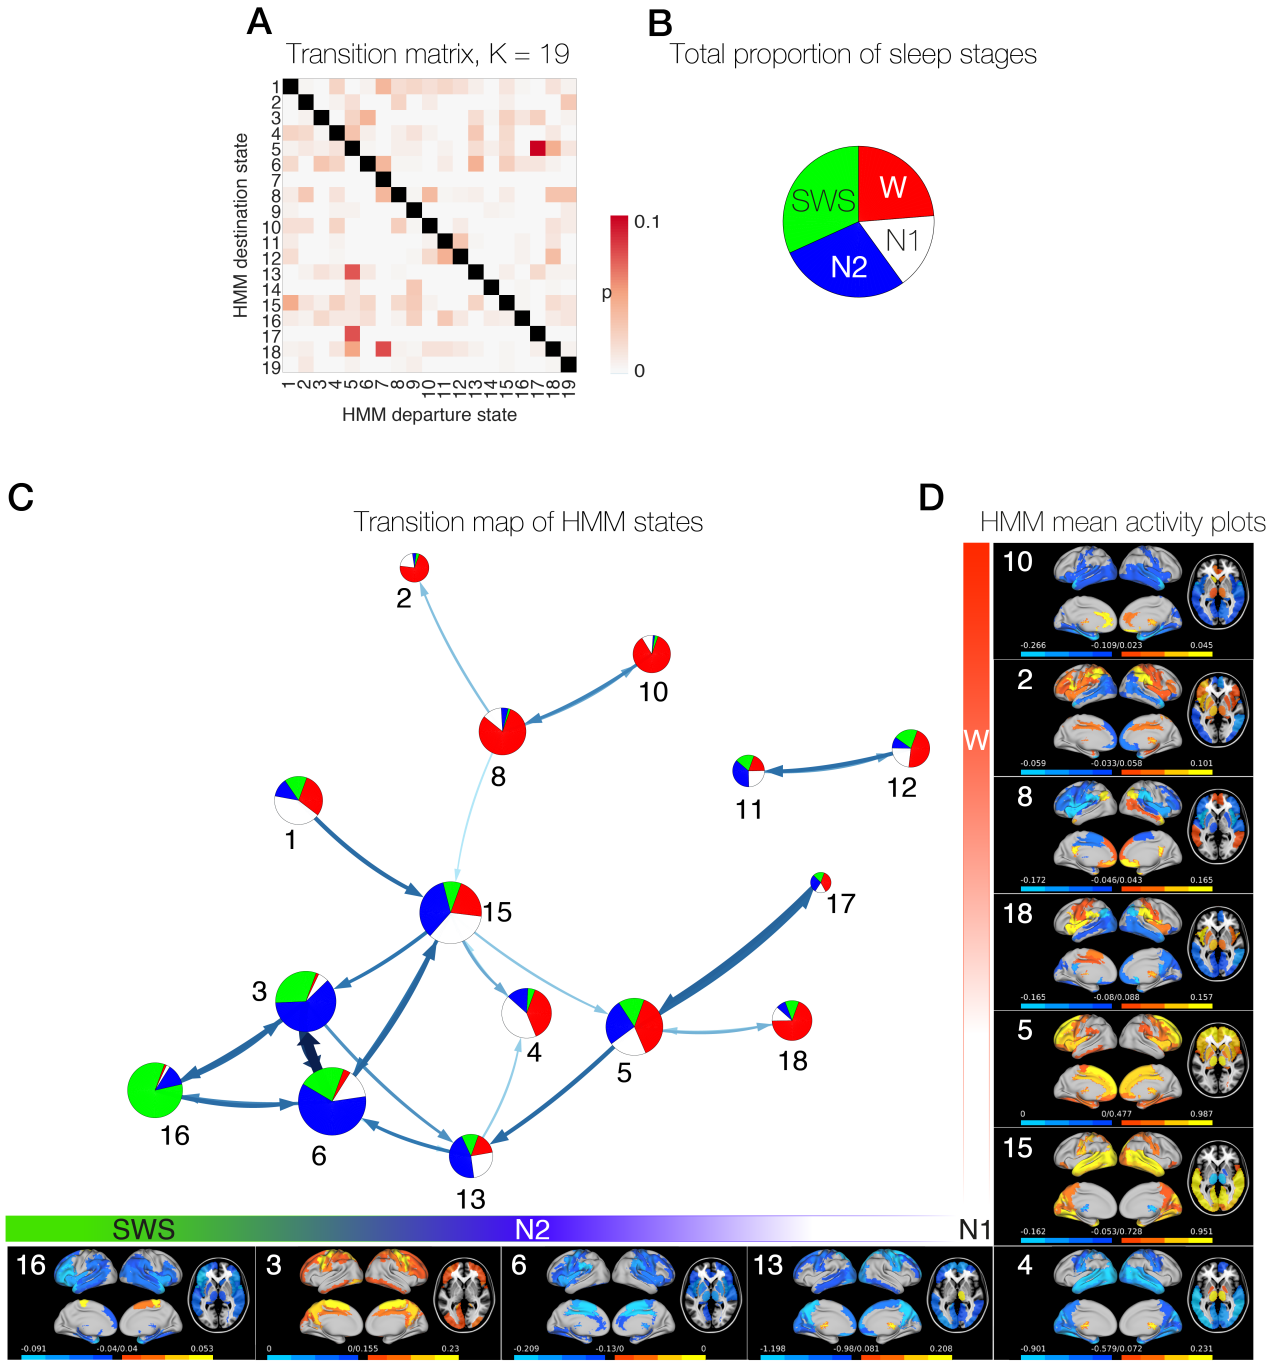

Figure 6-7 ***HMM state transitions***: **A)** Transition matrix of the HMM states calculated for the 18 participants that included all four vigilance conditions. This quantifies the likelihood of transitioning from any given (departure-) state to any other (destination-) state. **B)** Pie plot showing the total proportion vigilance stages within the 18 participants. **C)** Transition map showing the 8 % most frequent transitions. Each state is depicted as a pie plot expressing its likelihood of occurrence within each of the four vigilance conditions. Arrows show the direction of the transitions with thickness proportional to the transition probability. The transitions describe a passage from states with high activity during wakefulness in the top right, towards the left through states including more N1, and down to states specific to N2 and finally SWS. Note how wakefulness appear to be represented as two divisions, separated by states with higher N1-affinity **D)** Mean activity distributions of some of the more prominent states of **C**.

6.3.8 *Brain representations of HMM states: mean activity*

A logical next step is to ask which features in the BOLD activity that the HMM was able to pick up to achieve its high level of sensitivity and specificity to vigilance. To this end, we turn to the brain representations of the HMM states. In Figure 6-7D the transition diagram is accompanied by surface

renderings and mid-horizontal slices of the mean activity changes modelled by the HMM for a number of highlighted states of the transition map, and a closer look at these states is offered by Figure 6-8 and Figure 6-9.

When reviewing these representations it is interesting to note, how the patterns of mean activity changes most likely to be found in wakefulness coincided with distributions known from ICA- and seed-based connectivity resting-state studies. Thus, state 8 showed increased mean activity in areas closely overlapping the default-mode network, such as the bilateral posterior cingulate cortex, bilateral angular gyrus, bilateral superior medial frontal cortex, bilateral medial orbitofrontal cortex, and bilateral middle temporal cortex. These increases were accompanied by mean activity decreases in the bilateral insula, in particular, but also the bilateral supramarginal gyrus, and bilateral inferior frontal gyrus – areas often described as anti-correlated with the default-mode network in seed-based and ICA studies (see 1.5.3). State 2, another likely state during wakefulness, showed highest increases in the bilateral supramarginal gyrus, the bilateral thalamus, and right frontal cortex, with decreases in many of the previously referred-to default-mode areas, apart from the posterior cingulate cortex. State 10, the final of the first wakefulness division, showed large decreases in mean activity in the primary sensory areas, including bilateral visual, bilateral temporal, as well as bilateral somatosensory, together with increases in subcortical areas, such as the bilateral thalamus, and caudate, and the cortical anterior cingulate. From the second division of wakefulness, state 5 (and 17) was dominated by increases in frontal cortices, while state 18 resembled an inversed version of state 8 (with increases and decreases swapped around).

The HMM states with significant contributions to N1 sleep were characterised by reversed activities between the bilateral subcortical (thalamus and parts of the basal ganglia) and primary sensory areas. This was true for states 4 and 15 (and state 1 although its decreases were not confined to subcortical areas, but supplemented by decreases in the anterior and middle cingulate cortex). State 13 had highest specificity for N2 and N1, and showed thalamic increases combined with decreases across the rest of the brain.

N2 sleep was dominated by HMM states 3 and 6. The supplementary motor area played a significant role in both of these states; in state 3 as increases in concert with the bilateral precuneus and primary motor cortices; and in state 6 as decreases together with the bilateral thalamus, middle cingulate, supramarginal cortex, and the rolandic operculum.

The prevailing HMM state during SWS was state 16. This was characterised by decreases in mean activity in the bilateral middle and superior temporal pole, the orbital part and the operculum of the inferior frontal cortex, bilateral insula, as well as medial temporal areas. Some very localised increases were seen in the paracentral lobule and adjacent supplementary motor area.

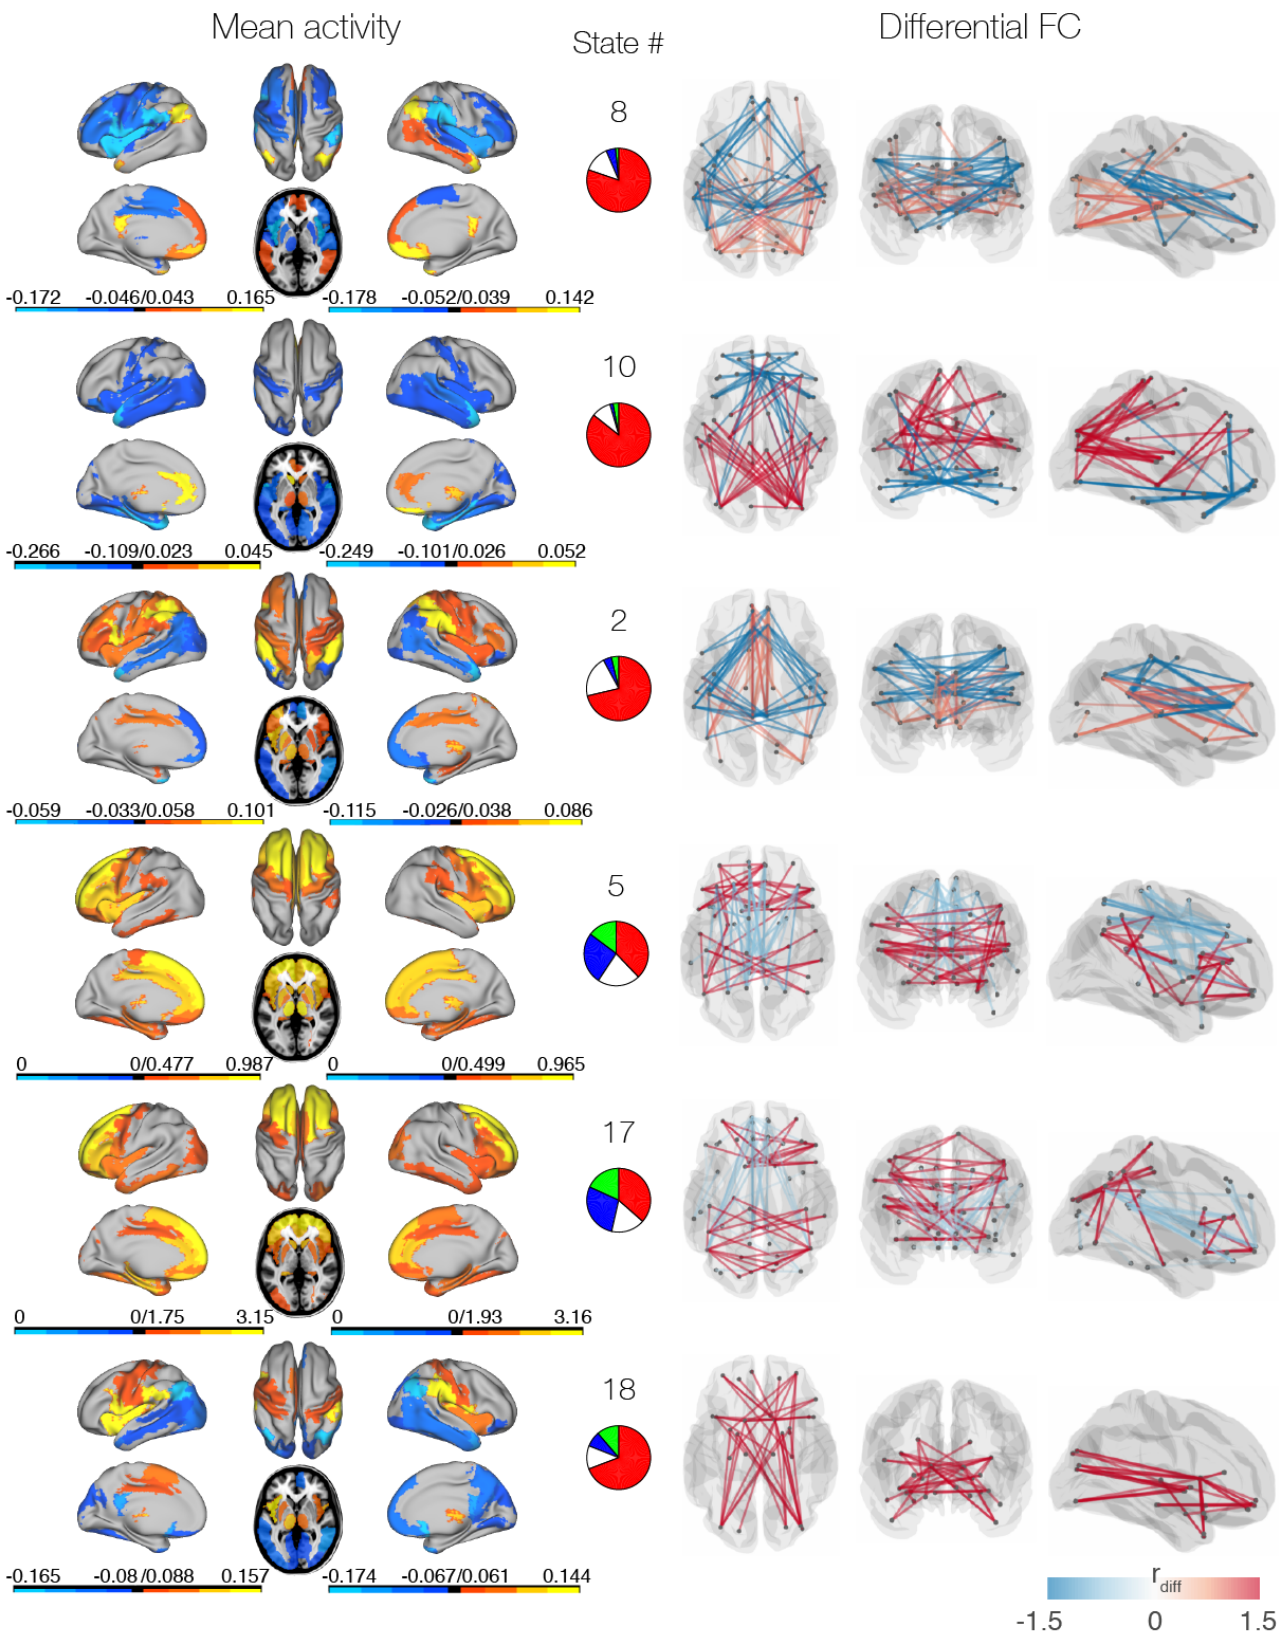

Figure 6-8 ***First and second division of wakefulness represented as largest mean activity changes and differential FCs***: *Left*: mean activity distributions of the first (states 8, 10, and 2) and second (states 5, 17, and 18) divisions of wakefulness as suggested by the transition map in Figure 6-7C. Maps were thresholded relative to each state, such that the 50% most negative and 50 percent most positive values are shown with red/yellow and blue/light blue respectively. Note the DMN appearance of the increases in state 8 and the anti-correlated network resemblance of the decreases (Fox et al., 2005). *Middle*: Pie plots showing the likelihood of finding each state within each of the four vigilance conditions. *Right*: Corresponding differential FC networks (FC_{diff}), computed by taking the state-dependent FC matrices and subtracting the average of the FC matrices of all other states. The maps were thresholded relative to each state, such that the 1% most positive and 1% most negative connections are shown in red and blue respectively.

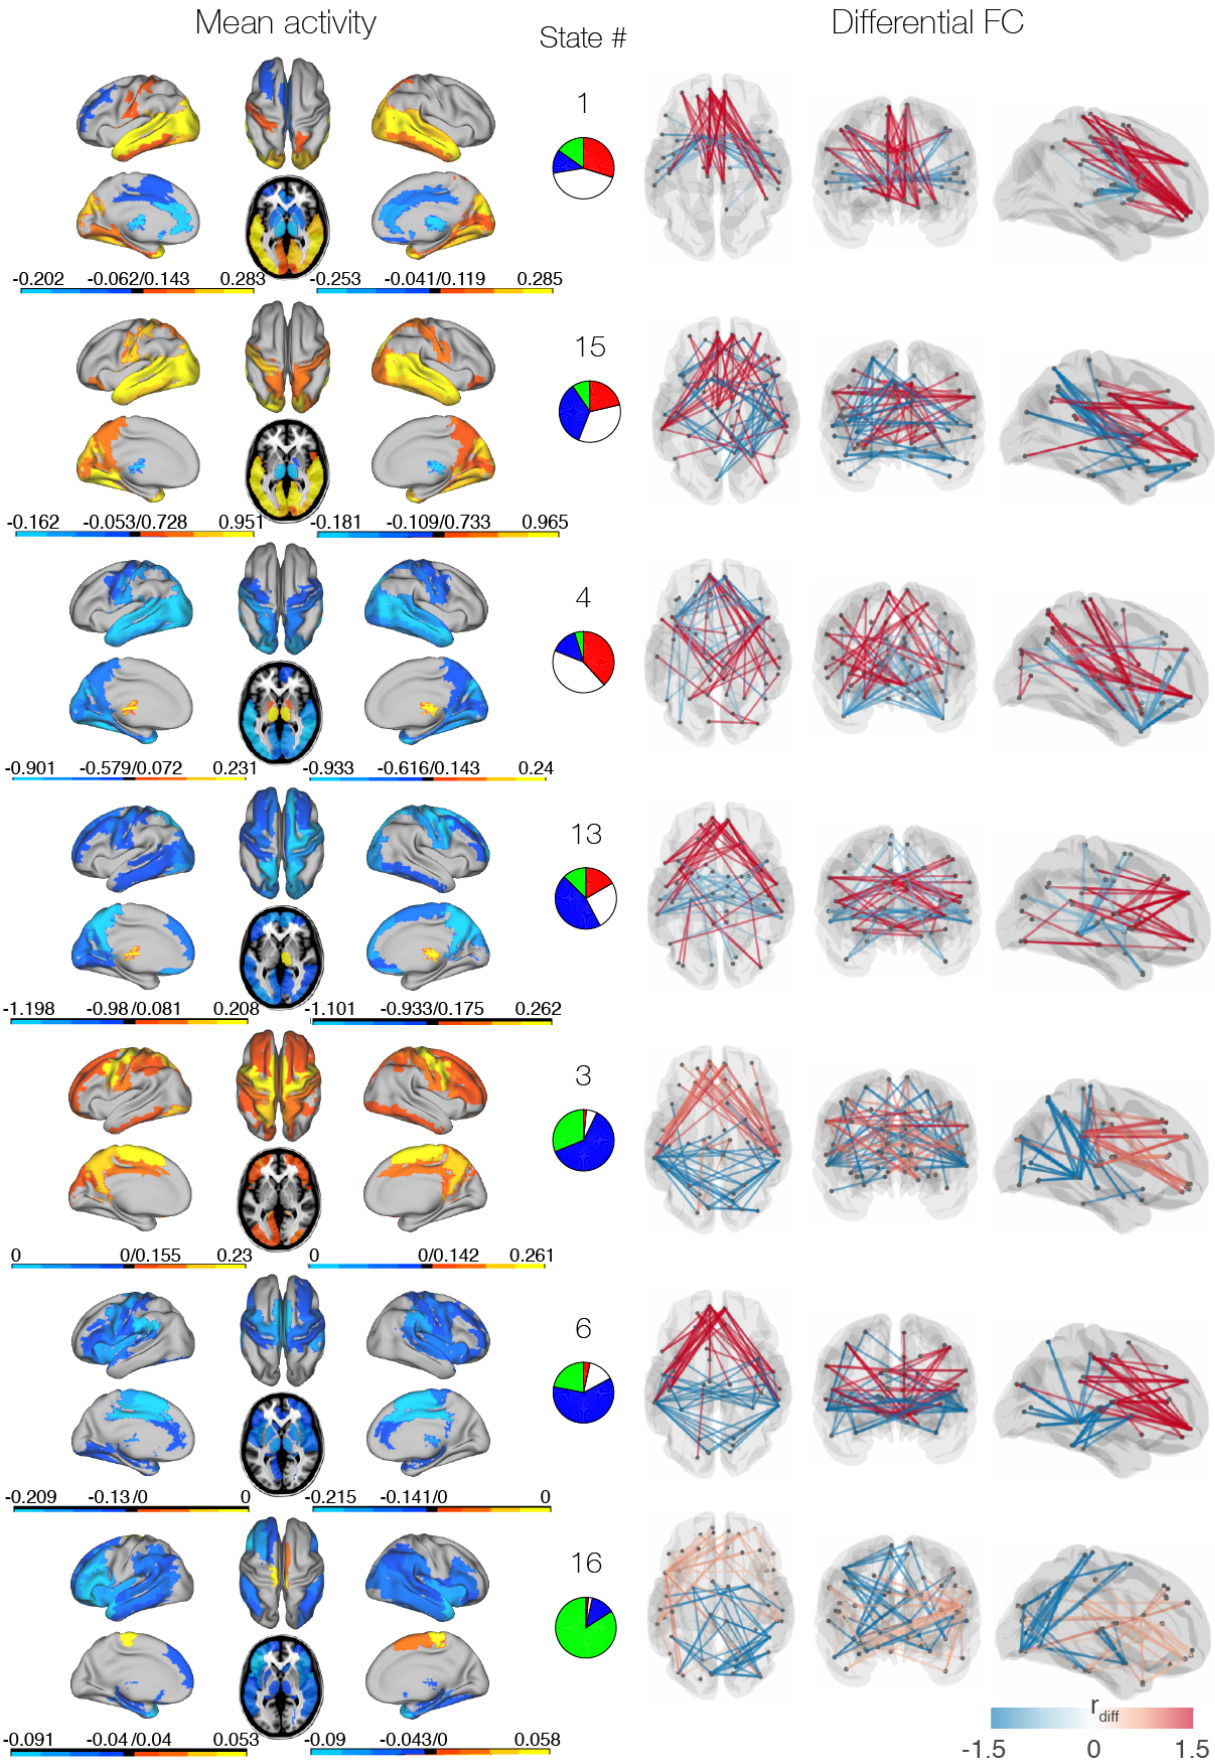

Figure 6-9 ***N1-, N2-, and SWS-related states represented as mean activity distributions and differential FCs***: Equivalent to Figure 6-8, but for the most prominent states of N1, N2, and SWS.

6.3.9 *Brain representations of HMM states: differential FC*

In addition to the mean activity changes, the HMM modelled the covariance between the 90 AAL ROI-timecourses when deriving the 19 states (see Figure 6-4B). As mentioned in the Methods, the covariance of an HMM state is equal to the temporally resolved FC. From the state redundancy plot in Figure 6-3D it is clear that the HMM state FC matrices correlated at quite high levels, implying that the states generally shared much of the connectivity. Focussing on the more subtle differences in FC patterns between the HMM states, the differential FCs (see Eq. 6.10) for a number of highlighted HMM states are presented in the right-hand columns of Figure 6-8 and Figure 6-9. The 1% largest positive and 1% largest negative differences are plotted on glass-brains as red and blue connections, respectively. When comparing these representations to the mean activity plots it becomes clear that the FC signatures of the HMM states did not necessarily follow the mean activity changes in trivial ways.

In state 8, where mean activity increases were seen in a DMN-like pattern, FC was mainly increased, compared to the average, in connections between the occipital and temporal cortices. The inferior occipital gyrus was particularly involved, as was the superior temporal gyrus. The largest decreases in FC formed a network involving the posterior cingulate cortex as a key area, and connections from here to and around the supramarginal gyrus, which had shown large decreases in mean activity in this state. The supramarginal gyrus itself showed decreased connections to the angular gyrus, insula, anterior cingulate, and medial frontal areas. State 2, which had shown mean activity increases in many areas where state 8 was decreased, also expressed a differential FC with an interesting involvement of the posterior cingulate cortex. An abundance of increased connections were thus evident from here towards the frontal cortices, while, as in state 8, connectivity was decreased from the posterior cingulate to the supramarginal gyrus and from the latter to the angular gyrus and the anterior cingulate. The last of the three states from the first division of wakefulness, state 10, had relative increases in connectivity from the middle and superior occipital gyri to the postcentral and precentral gyri, as well as the middle and superior temporal gyri. A distinct network of connectivity decreases was centred on the gyrus rectus and adjacent medial orbital parts of the frontal cortices. From here connectivity decreases were evident within the frontal cortices as well as to inferior parts of the temporal and occipital cortices.

In the second division of wakefulness, which included HMM states 5, 18 and arguably 17, the relationship between mean activity changes and differential FC was no less complex. For states 5 and 17 some concordance was suggested, as both of these had shown strong mean activity increases in the frontal cortices, which seemed reflected in increased within-frontal connectivity. At the same time, however, they both showed decreased FC between frontal areas and the posterior cingulate cortex. HMM state 18 was distinguishable by the fact that it had no FC decreases relative to the average. The strongest of its increases gravitated towards the caudate, from the occipital cortices, the temporal pole, and the orbital parts of the frontal cortices.

In terms of the reversals in mean activity between subcortical and cortical areas seen in the states related to N1 sleep, this was reflected in the FC for some of these states. HMM state 1 thus showed very specific decreases in connections involving the caudate, as did state 15, while the thalamus was involved in many

decreased connections in state 13. State 4 had thalamic FC decreased to the temporal pole, and increased to the medial orbitofrontal cortex. The subcortical FC decreases commonly involved the temporal cortices and the precentral areas. Despite rarely reflected in the mean activity changes, frontal increases in FC were common to all of the N1-related HMM states. In particular the medial parts of the frontal cortices were well represented, and increased connectivity between these and the supramarginal gyrus was a common trait. For state 1, projections to the frontal cortices were further noticeable from the median cingulate cortex, and the paracentral lobule; from the insula in state 13 and 4, the latter of which also included connections from the superior parietal gyrus; and the median cingulate cortex in state 15.

HMM states 3 and 6 were signatures of N2 sleep, and the FC of these states differentiated from the average in comparable ways. In terms of relatively decreased FC they both showed anchors in the temporal cortices, and the middle temporal gyrus was consistently involved. From here decreased connections were seen to other temporal areas and the occipital cortices, including the inferior occipital gyrus in state 3 and the fusiform gyrus in state 6. In state 3 the parietal cortices were also part of this network, represented by the inferior and superior parietal gyri. The insulae were specific to the network of decreases in state 6. The relative increases were reminiscent of the increases seen in the N1-related states, with an abundance of connections between the supramarginal gyrus and medial frontal areas. The participation of the anterior cingulate cortex appeared as specific to the N2 states. HMM state 3 further expressed increased FC between the thalamus and the paracentral lobule.

HMM state 16 dominated during SWS. The very specific mean activity increases seen in the supplementary motor area and the paracentral lobule were somewhat reflected in FC decreases between these areas and occipital areas, including the inferior occipital gyri, as well as the right inferior parietal gyrus. Decreased connectivity was further seen between these inferior parietal and occipital gyri, as well as between the latter and the amygdalae. Interestingly, the main hub of the network of increased FC was the opercular part of the left inferior frontal gyrus, which had shown one of the strongest decreases in mean activity in this state. Increases were also seen from the temporal cortices and the supramarginal gyri towards the frontal cortices.

Representations of both mean activity changes and differential FCs for all HMM states can be found in the Appendix (see Figures S 8–9), along with matrix representations of FC_{diff}.

6.3.10 *Switching within and differences across vigilance conditions*

The fact that the HMM solution assigned each data time point a specific HMM state made it possible to ask whether dynamic features of the large-scale brain activity and networks were vigilance-dependent. By parting the HMM state timecourses according to the EEG-based sleep scoring, we were able to quantify the vigilance-specific: i) amount of switching, and ii) state exploration.

The amount of switching was defined as the number of HMM state transitions occurring within each of the four vigilance conditions. As shown in Figure 6-10A, there was significantly more switching taking place during wakefulness and N1 than during both N2 and SWS, and significantly less in SWS than in N2.

While higher switching suggests a more dynamic exploration of HMM states, this measure does not take into account the diversity of the states that are being switched between. To capture this relative state diversity/similarity, we characterised each transition taking place within each of the four vigilance stages in terms of the correlation between the HMM states switched from and to. This was done using both the HMM state FC matrices, shown in Figure 6-10C, and the absolute of the HMM state mean activity profiles, shown in Figure 6-10B.

As seen in Figure 6-10C, the similarity of the HMM states switched between was high, when FC matrices were considered, for all vigilance stages, all of which showed mean values above $r = 0.7$. It is worth noting that this was higher than the all-to-all correlation level, referred to as the state redundancy in section 6.4.2 ($r = 0.606$, see Figure 6-3D), thus suggesting that the HMM states that were transitioned between within vigilance stages were more similar in terms of their FC matrices than the average similarity across the HMM states. The highest mean value of FC similarity was found for N2, which was significantly increased compared to wakefulness.

The mean activity of the HMM states varied markedly more than the FC patterns, as illustrated by the relatively low correlation values in Figure 6-10B. The most striking observation was that the mean activity similarity was significantly increased for N1 compared to all other vigilance stages, implying that absolute activity changes were distributed more similarly across ROIs for the states transitioned between in N1.

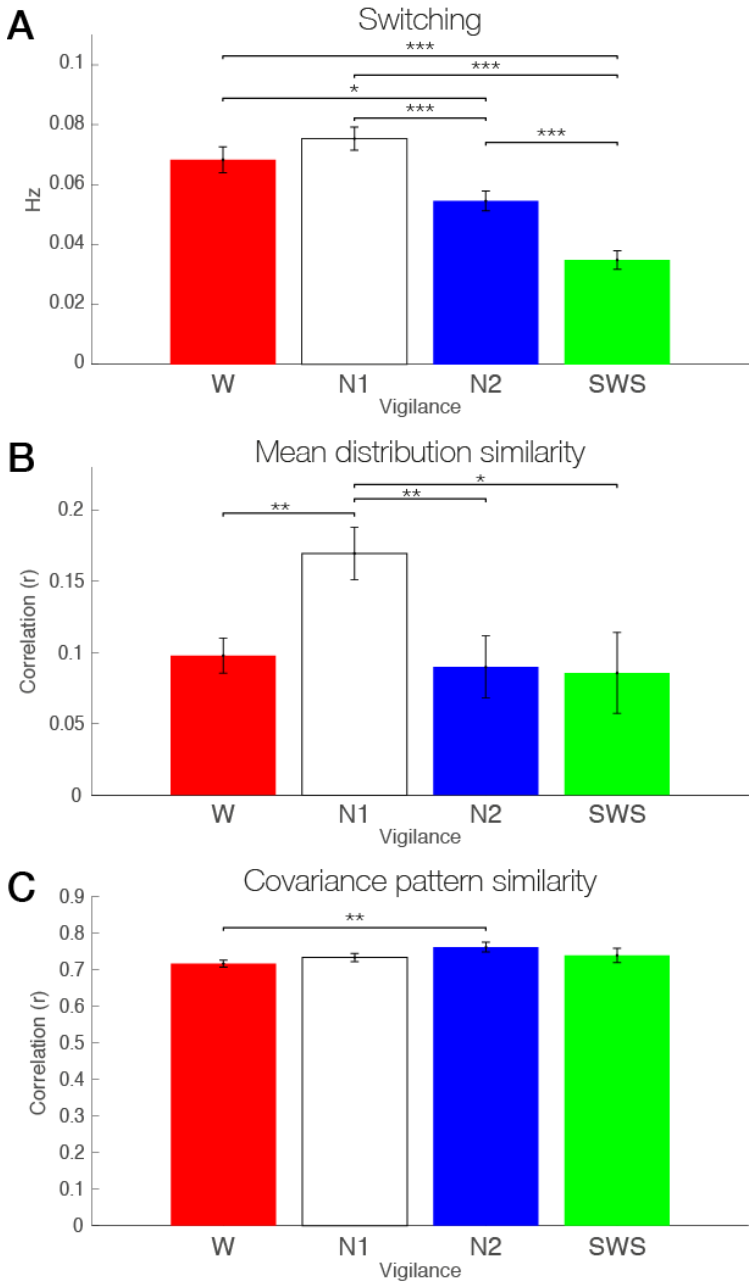

Figure 6-10 ***Large-scale dynamics of wakefulness and sleep stages in the HMM state space***: **A)** Switching frequency averaged within each stage of vigilance, estimated as the number of HMM state transitions per time within each of the 18 participants. **B)** Mean distribution similarity within vigilance stages, estimated as the Pearson's correlation of the absolute of the mean distribution just prior to and just after following individual transitions. **C)** Covariance pattern similarity within vigilance stages, estimated as the Pearson's correlation between the upper triangular parts of the covariances of the HMM states just prior to and just following individual transitions. Significance was evaluated using paired t-tests and permutations. * $p < 0.05$, ** $p < 10^{-2}$; *** $p < 10^{-3}$.

6.4 ***Discussion***

By applying an HMM with VB inference on concatenated fMRI data from 57 participants undergoing transitions between wakefulness and NREM sleep, we have produced a temporal decomposition of large-scale networks, driven solely by the statistical attributes inherent to the data. The HMM parted the data into segments, where the data fit certain configurations of mean activity and covariance (FC) in and between the 90 AAL ROIs, and each of these configurations corresponded to one out of K HMM states. Subsequent to the HMM derivation we quantified the multivariate relationship, using MANOVA, between the set of HMM states and the sequence of vigilance stages given by the EEG-based sleep scoring. For the majority of model orders between 4 and 45, the HMM states expressed a non-random dependence on vigilance. Evaluating each model order solution, in terms of model performance (minimum free energy), relationship to vigilance (MANOVA), state redundancy (covariance similarity), and average state life time, we found the model order K = 19 to represent a good trade-off between richness and redundancy. Examining this HMM solution, it was shown how many of the HMM states showed high specificity for individual vigilance stages, and were distributed in a temporal sequence emulating the EEG-based sleep scoring. Furthermore, our results strongly suggest that a higher number of HMM states were specific to wakefulness compared to the NREM sleep stages.

The fact that each HMM state was defined in time allowed us to probe the nature of the HMM-state transitions within and across the traditionally defined stages of vigilance. From this it was clear that the HMM states grouped, such that states with high probability of occurring within the same vigilance stage were more likely to transition between each other. Considering the transition dynamics, it was shown that the overall switching was markedly higher in wakefulness and N1 sleep compared to the more consolidated stages of NREM sleep, N2 and SWS. While the switching frequency was comparable between wakefulness and N1, we have shown that the states switched between in N1 tended to be more similar than the states switched between during wakefulness.

6.4.1 *HMM sensitivity and specificity to vigilance in fMRI data*

A main result of the present chapter is that a strictly data-driven dynamic representation of large-scale networks in fMRI data can yield states of high temporal concordance with the traditional EEG-based sleep scoring. The fact that this can be realised simply by visually tracking the HMM states together with the sleep scoring illustrates the robustness of this finding (see Figure 6-4A). From a quantitative perspective, this was confirmed by the MANOVA results (see Figure 6-3B–C). To our knowledge, this provides the first evidence that wakefulness and sleep can be extracted directly from continuous fMRI recordings without any explicit use of polysomnographic information. In terms of the two scenarios listed in the Introduction to this chapter, regarding the possible ways vigilance changes could be envisaged *a priori* to manifest in a dynamic decomposition of large-scale networks, the HMM solution suggests that wakefulness and each stage of NREM sleep are in fact represented as distinct states, rather than different temporal modes of a shared set of states. As such, the HMM identified states that were both sensitive and specific to each of the included, traditionally defined, stages of vigilance. This is in line with the successful machine-learning vigilance-

classification reported by Tagliazucchi et al. (2012b), Tagliazucchi and Laufs (2014) and Altmann et al. (2016). Of note, these studies trained their respective classifiers on windows of correlation matrices, i.e. FC information only, while the HMM implemented in the present chapter, in addition to FC, i.e. covariance of the fitted Gaussian distributions, was sensitive also to changes in mean activity of the BOLD timecourses. In fact, the high cross-HMM-state covariance similarities (see Figure 6-3D) may suggest that the HMM to a significant extent relied on mean activity information to resolve the vigilance organisation of the data. It is important to note, however, that the HMM modelled the mean and covariance of each state simultaneously, and that these first and second order statistics in concert characterise the unique configuration of a state. Furthermore, although it is generally assumed that an FC network contain more, or at least different, information than what can be explained by changes in mean activities of the connected ROIs, the two measures should not be viewed as completely independent of each other. Indeed, comparing the largest mean activity changes to the largest changes in FC of the HMM states (see Figure 6-8 and Figure 6-9) revealed that the relationship between the two pieces of information is rather complex. We argue that the mean activity changes and FC represent complementary measures characterising the behaviour of the large-scale network of fMRI signals.

While seed-based correlation estimates, employed in both static (see Introduction, 1.5.3, and Chapter 2, 2.3.2) and sliding-window approaches (see Introduction, 1.5.8, and Chapter 2, 2.3.8) focus on second order statistics and do not explicitly model changes in mean activity or signal amplitude, the latter first-order statistics do play increasingly important roles in recently proposed directions for studying large-scale dynamics in fMRI (Tagliazucchi et al., 2012a, Karahanoğlu and Ville, 2015, Preti et al., 2016, Tagliazucchi et al., 2016). It has thus been shown that RSNs and large-scale FC networks can be dynamically resolved by considering how ROIs co-activate transiently as given by the amplitudes of whole-brain BOLD signals (Tagliazucchi et al., 2012a, Karahanoğlu and Ville, 2015, Tagliazucchi et al., 2016). The fact that the HMM picked up on changes in mean BOLD activity makes sense from the perspective of earlier neuroimaging activation studies, contrasting stages of wakefulness and NREM sleep (see Introduction, 1.5.2). As such, PET studies have shown that metabolic changes take place in heterogeneous patterns throughout the brain volume, often with the biggest sleep-related decreases in regions of higher-order functions (Maquet, 2000, Dang-Vu et al., 2010b, Meyer et al., 2012, Nir et al., 2013). As pointed out earlier, investigations of the fMRI in sleep have traditionally focussed on BOLD-correlates of sleep-specific oscillations, such as spindles, K-complexes, and slow oscillations (see Introduction, 1.5.2). The few studies that have specifically investigated overall stage-dependent activity changes in the continuous BOLD signal from wakefulness to NREM sleep are more or less in line with the PET findings (Kaufmann et al., 2006)(although, see Laufs et al. (2007)).

It remains an open question whether the HMM reached its high level of sensitivity and specificity to vigilance by modelling mean activity, FC information, or indeed a crucial combination of both. Interestingly, this would be testable within the HMM framework. As shown in a recent study using an HMM to probe dynamics of predefined RSNs during the resting-state, the HMM can be tweaked to consider only the covariance when deriving its states (Ryali et al., 2016). The reverse can also be accomplished by fixing the

covariance structure, such that the HMM only searches for changes in mean activity distributions. Comparing the vigilance-sensitivity of both of these HMM implementations in turn with the model of the present chapter could, at least in theory, provide insight into the respective relevance of the mean activity and FC information.

6.4.2 *Classification error or a matter of more spatiotemporal detail?*

While we have thus far focussed on the correspondence between the HMM state solution and the EEG-based sleep scoring, it is worth examining in more detail the nature of this correspondence. The MANOVA results essentially expressed how well the HMM state timecourses aligned with a grouping given by the EEG-based sleep staging. Performing the same analysis with a shuffled version of the sleep staging revealed that the vigilance sequence was a non-random trait of the HMM state timecourses. Another salient and, perhaps even more, noteworthy implication of the MANOVA results was that the fit improved as the number of HMM states were increased. This means that the HMM did a better a job at capturing vigilance changes when it was allowed to use more than the four states employed in polysomnography. If, by analogy, the HMM is thought of as a data-driven sleep scorer, given nothing but the fMRI signals, it is telling us that four stages are too limited to summarise the data. As discussed in the Introduction (see 1.4.4), the validity of current polysomnographic staging protocols has since their inceptions been subject to prolific criticism, and the number of stages and their criteria for identification remain at the core of sceptical views. We highlight the limited spatial resolution of a few EEG electrodes as another shortcoming of polysomnography, and regard the higher number of states identified by the HMM as illustrating that a richer picture is called for, of wakefulness as well as of NREM sleep, when the whole-brain perspective of fMRI is adopted.

Wakefulness was modelled with the highest number of HMM states, which seems in line with the intuitive idea that heightened vigilance requires more diverse brain activity. Compared to the other stages of vigilance, the HMM struggled to identify any single states unequivocally specific to N1 sleep. Rather the HMM states with highest activity during polysomnographically defined N1, such as HMM states 1, 4, and 15 also showed significant occurrences during wakefulness and N2 sleep. This issue will be further addressed below. Meanwhile, two HMM states (3 and 6) were found highly specific to N2 sleep. Finally, SWS appeared as the vigilance stage where the HMM agreed most with polysomnography in modelling the far majority of this stage as one state.

Although the MANOVA revealed a significant correspondence between sets of HMM states and the EEG-based sleep scoring, it is also important to note that, even if we accept that vigilance stages are composed of more than one HMM state each, the fit was not perfect. HMM states that occurred with high specificity to a given vigilance stage could always be found, albeit much less often, during other vigilance stages as well. As an example, HMM state 16 dominated the fractional occupancy of SWS (see Figure 6-5) and showed high specificity for SWS (see Figure 6-7), however on occasion was also found during the polysomnographically-defined stage N2. Similarly, HMM states that modelled wakefulness, such as 8, 10, and 2, showed some activity during N1 sleep, and those modelling N2, such as 3 and 6, had some prevalence in N1 and SWS. In particular the HMM states dominating N1 sleep in terms of fractional occupancy did not

reach the same level of specificity for N1 as other HMM states showed for their respective vigilance stages, but rather showed considerable activity in wakefulness and N2.

Notably the HMM consistently 'misclassified' states in a sensible way within 'neighbouring' vigilance stages, i.e. wakefulness to N1; N1 to wakefulness and N2; N2 to N1 and SWS; and SWS to N2. A reasonable question is therefore whether the imperfect specificities of the HMM states were caused by classification errors of the HMM, or by potential inconsistencies in the EEG-based sleep scoring? While we are merely able to speculate about this, there are a couple of factors that speak in favour of the latter option. First, visual sleep scoring relies on segmenting EEG data into 30-seconds windows, which are then categorically labelled as one of the possible vigilance stages (see Introduction, 1.2.1). This effectively corresponds to a substantial down-sampling of vigilance, and any potential transitioning or mix of stages within a pre-defined 30-seconds window will be ignored. Since the HMM was not constrained by such windowing, it is conceivable that some of the mismatches between the HMM and the sleep soring may be attributed to the latter being too coarse. Secondly, it is worth pointing out that visual polysomnographic scoring is associated with a non-negligible inter-observer variability. A recent report by the American Academy of Sleep Medicine (AASM) thus estimated the reliability of sleep scoring performed according to the AASM criteria (Iber et al., 2007) on the same data by different sleep experts to 83% (Rosenberg and Van Hout, 2013), and even lower levels have been reported for the manual of Rechtschaffen & Kales (Norman et al., 2000). In the AASM report it is particularly interesting to note that the reliability was found lowest for N1 sleep (Rosenberg and Van Hout, 2013). This may be unsurprising since the EEG of N1 is less characteristic than N2, which is defined from spindles and K-complexes, and from SWS with its associated slow waves. Meanwhile, this low reliability of N1 staging could potentially explain the lowered level of specificity of the HMM states. Its data-driven nature, we argue, favours the HMM in this respect.

6.4.3 *The HMM reveals time-resolved large-scale networks with RSN resemblance and strong vigilance-dependence*

The crucial difference between the approach presented in the present chapter and the machine-learning studies (Tagliazucchi et al., 2012b, Tagliazucchi et al., 2015, Altmann et al., 2016) lies in the fact that while the latter aim to identify features in the covariance structure of the fMRI signals fitting the vigilance-specific segments given by EEG-based sleep scoring, the HMM extracts state features in a strictly data-driven fashion. The convincing temporal overlap between the sleep scoring and the HMM decomposition was thus not driven by a classification procedure, but rather established post-hoc. If the traditional sleep scoring is taken as a proxy for the level of vigilance rather than as a gold standard set in stone, it may be argued that the HMM-derived large-scale network patterns offer a view of how vigilance is represented inherently in the fMRI signals. These patterns will now be discussed.

Beginning with the HMM states with high activity during wakefulness, it is interesting to note that these expressed mean activity and FC changes in patterns resembling known RSNs (see Introduction, 1.5.3). As such, HMM state 8 showed increased mean activity in core areas of the DMN (Raichle et al., 2001, Raichle, 2015), including the angular gyrus, the posterior cingulate cortex, and the medial prefrontal cortex. Notably,

concomitant mean activity decreases were seen in a set of regions resembling the anti-correlated network known from the study by Fox and colleagues (2005). A number of networks have been defined as comprising these areas. One is the attention network, where a distinction is often made between a dorsal (DAN) and a ventral (VAN) part, (Fox et al., 2006). Another is the central executive network (CEN), which is often described as involving the supramarginal gyrus and the dorsolateral part of the frontal cortex (Koechlin and Summerfield, Ryali et al., 2016). The antagonistic relationship between the DMN and these task-positive networks has received much attention since the discovery of their anti-correlation in static FC analyses during the resting-state (Fox et al., 2005) (see Introduction, 1.5.3). While the identification of this anti-correlation has traditionally relied on global signal regression, which remains a controversial pre-processing step of the BOLD timecourses (Fox et al., 2009, Murphy et al., 2009a, Murphy and Fox, 2016), our results suggest that this antagonism may be found as opposite signs in mean activity from fMRI data without global signal regression. Apart from HMM state 8, this was also found for state 2 where mean activity increases delineated a DAN-like pattern, while decreases were found in many of the areas of the DMN, and for state 18, the mean activity changes of which were very similar to state 8 but with the signs of the changes reversed.

Finding higher-order RSN-like structures in HMM states with high specificity to wakefulness is significant, given the theories and ideas linking these networks to cognition and complex mental phenomena (Buckner et al., 2008, Andrews-Hanna et al., 2010, Bressler and Menon, 2010). Meanwhile, this finding would appear to contradict previous static FC studies that have struggled to show clear differences in the presence of these networks between wakefulness and NREM sleep (Horovitz et al., 2008, Horovitz et al., 2009, Larson-Prior et al., 2009, Sämann et al., 2011), even on the very same dataset analysed here (Tagliazucchi et al., 2013c) (see Introduction, 1.5.3). The question is; what makes the HMM capable of assigning these higher-order RSN structures so specifically to wakefulness, when seed-based and ICA approaches do not? A straightforward suggestion is that the HMM models mean activity while the static FC methods focus on the correlation between timecourses. Hence it may be instructive to consult the patterns of covariance, i.e. FC, of the wakefulness-related states as modelled by the HMM. As mentioned before, the FC patterns output by the HMM correlated quite highly across vigilance conditions, which is indeed in line with the previous reports of a generally preserved functional connectome with high vigilance-resilience (Horovitz et al., 2008, Horovitz et al., 2009, Larson-Prior et al., 2009, Sämann et al., 2011, Tagliazucchi et al., 2013c). However, when focussing on the way each of the HMM states deviated from the average FC structure, by considering the differential FC, it becomes clear that the state-specific FC and mean activity distribution did in fact complement each other, albeit in complex ways. The antagonism between the DMN and DAN/VAN/CEN was thus reflected in the FC_{diff} of state 8 and state 2 by decreased connectivity between areas of these networks, such as the posterior cingulate cortex, medial prefrontal areas, and angular gyrus of the DMN and the supramarginal gyrus of the DAN/VAN/CEN. In contrast, increases in connectivity were seen between DMN core areas in state 2, i.e. posterior cingulate cortex to medial prefrontal areas, and medial prefrontal areas to the angular gyrus. The largest relative FC increases of state 8 and state 10 were found in connections between the occipital and temporal/parietal cortices, forming networks overlapping with those

identified as strongest during wakefulness by static comparative network analyses in Chapter 2 (with fMRI, see Figure 2-2) and in Chapter 3 (with MEG, see Figure 3-10A and C).

The states with high activity during N1 sleep commonly showed reversed mean activities between subcortical and cortical areas. This observation is consistent with electrophysiological evidence from intracortical recordings in rats (Baker et al., 2014b) and in humans (Magnin et al., 2010) showing that thalamic changes in dynamics precede those of cortical areas near the onset of NREM sleep. The same phenomenon has been used to explain the drop in thalamocortical FC that is consistently observed for N1 sleep, by others (see e.g. Spoormaker et al. (2011)) and in Chapter 2. To our knowledge, however, the very specific inverse relationship between activity changes in subcortical and cortical areas, suggested by the HMM states 1, 4, 13, and 15, is the first evidence that this can be observed directly from the BOLD timecourses, providing what would seem like a close correspondence between electrophysiology- and fMRI-derived activities around the onset of NREM sleep. The increased spatial resolution of fMRI offers the novel insight that this inverse relationship was particularly clear between subcortical and primary sensory areas of the cortex, including the visual, auditory and somatosensory systems. It is tempting to speculate that this distribution of activity is involved in the decreased, and less reliable, responsiveness to external stimuli that is defining for this early process of falling asleep (Ogilvie, 2001).

In terms of the relative FC changes taking place during the N1-related HMM states, only a subset of them expressed markedly decreased connectivity between cortical and subcortical areas. The caudate was thus found de-connected from temporal areas in HMM states 1 and 15, while state 13 expressed decreased temporal connectivity to the thalamus. Increases in FC were common for the N1-related states to medial frontal areas, from the supramarginal gyrus, the insula, and the median cingulate cortex. Once again, this is in line with the comparative network analyses of Chapter 2 (see Figure 2-2).

N2 sleep was dominated by HMM states 3 and 6. Interestingly, these two states respectively expressed solely increases and solely decreases in areas, which have been consistently identified in a number of studies, as fMRI-correlates of sleep spindles (Caporro et al., Laufs et al., 2007, Schabus et al., 2007, Andrade et al., 2011) (see Introduction, 1.5.2). These included the supplementary motor area, pre- and post-central gyri, paracentral lobule, precuneus, and lateral parts of the frontal cortices. It is likely that the HMM was sensitive to the signal changes caused by the spindles that constitute the most salient graphoelement of N2. In further support of this, the differential FC of both HMM state 3 and 6 included increases in thalamocortical connections. As mentioned in the Introduction (see 1.3.1) of this thesis and the Discussion of Chapter 3 (see 3.5.2), the interplay between the cortex and thalamus is essential to the generation of sleep spindles (Luthi, 2014, Piantoni et al., 2016a). In HMM state 3, specific increases were thus seen between the thalamus and the paracentral lobule, while state 6 expressed increased connectivity between the thalamus and the superior occipital gyrus as well as the orbital part of the superior frontal gyrus. Future studies should investigate the relationship between N2-related HMM states and electrophysiologically defined spindles, as well as other sleep oscillations.

The overall appearance of the differential FC was quite consistent across both of the N2-related states and was once again in line with the comparative network analyses in Chapter 2 (see Figure 2-2). This included

decreased connectivity from the middle temporal gyri to the occipital cortices in both states, between the temporal cortices in state 6, and from the middle temporal cortex to the posterior parts of the parietal cortex in state 3. Apart from the thalamocortical connections, discussed above, the relative increases in connectivity followed the trends of the N1-related states, and were especially localised between the supramarginal gyrus and medial areas of the frontal cortices.

The majority of SWS was comprised of HMM state 16. The corresponding mean activity distribution was dominated by decreases, the largest of which were found in the insulae, the temporal pole, and the lateral parts of the frontal cortices. In particular the frontal decreases could be interpreted in light of the earlier PET findings of SWS-related decreased metabolism in the frontal cortices, which is generally believed to reflect the higher concentration of slow-wave activity in these areas (Werth et al., 1997, Meyer et al., 2012) (see Introduction, 1.5.2). The only observed increases for state 16 were found in the paracentral lobule and the adjacent supplementary motor area. These might reflect activity in a system involved in SWS-dependent consolidation of motor-skill learning (Tamaki et al., 2013). From the FC perspective, the paracentral lobule and the supplementary motor area were also involved, by virtue of their decreased connectivities to occipital areas and the inferior parietal gyrus. Further, it is interesting to note that the relative FC increases of state 16 formed a fronto-temporal network, not too dissimilar from the strongest SWS-related connectivity identified in Chapter 3 with MEG in the alpha frequency band (see Figure 3-4 and Figure 3-10C).

6.4.4 *Large-scale network transitions from wakefulness to sleep*

Having consulted and discussed the large-scale networks of the individual HMM states, and establishing the close relationship between the temporal occurrence of these states and the EEG-based sleep scoring, it is now possible to review and interpret the state transitions modelled by the HMM.

Figure 6-7C illustrates how the most significant state transitions describe the natural succession of vigilance stages descending from wakefulness, through N1 and N2 sleep to SWS. This complements the MANOVA results (see Figure 6-6B) in supporting the claim that the HMM successfully identified the vigilance organisation of large-scale networks inherent to the fMRI data. However, the transition map in Figure 6-7C further offers some interesting and novel perspectives on the general nature of switching between wakefulness and sleep. For instance, the transitions appear to suggest that wakefulness may be parted into two subdivisions, separated by HMM states mainly occurring during sleep. A likely explanation for this is that one division corresponds to wakefulness-periods prior to sleep, while wakefulness following sleep is captured in the second division. While it would take more specific testing to confirm, a visual inspection of the HMM state timecourses from Figure 6-4A do indeed appear to show some support for this claim. As such, states 5, 18, and arguably 17 from the second division visibly seemed to increase their probabilities in wakefulness-epochs that occurred after sleep. If this is a true finding, it might very well reflect the different nature of being awake prior to sleep and waking up from sleep. While these two types of wakefulness are regarded equivalent from the perspective of polysomnography, it is well known that they differ markedly in terms of behaviour and cognition. Little is known, however, about the neural processes underlying the distinctive sensation of awakening characterised by subjective feeling of grogginess and

decreased cognitive performance, also referred to as sleep inertia (Tassi and Muzet, 2000, Nir et al., 2013). Interestingly, findings of scalp EEG studies of the waking-up condition have been interpreted as signs that the posterior parts of the cortex are particularly slow at returning to levels of activity seen prior to sleep (Ferrara et al., 2006, Marzano et al., 2011). Consistent with this, a recent fMRI study showed frontal BOLD activity to be increased relative to that of posterior areas upon awakening (Tsai et al., 2014). Two of the HMM states found in, what may be tentatively referred to as, the 'wake-up' division, states 5 and 17, were indeed characterised by marked increases in the frontal cortices, but notably also by decreased connectivity between some of these frontal areas and the posterior cingulate cortex, as may be seen from their differential FCs.

If it is accepted that states 5, 17, and 18 primarily described a condition of awakening, it follows that states 2, 8, and 10 likely belonged to wakefulness prior to sleep. This means that the process of falling asleep may be queried through a closer look at the transitions taking place between this particular division of wakefulness and HMM states with high prevalence in sleep. When doing so, it becomes clear that the transition from HMM state 8 to HMM state 15 was the best candidate for a representation of the descent from wakefulness to N1 sleep, i.e. of falling asleep. There is currently little consensus of how to identify the exact onset of sleep, both in terms of behaviour and brain activity (Ogilvie, 2001, Nir et al., 2013). Responsiveness to stimuli as well as subjective experience can vary quite substantially during this state of drowsiness and sleep induction, without any definite correlates in the EEG (Ogilvie, 2001). Polysomnographic staging avoids addressing this issue by simply classifying 30-seconds windows according to their 'average' appearance (see Introduction, 1.2.1). From this perspective, the actual transition from wake to sleep could take place anywhere within a 30-second window scored as N1. The scoring system introduced by Hori and colleagues provides a significantly richer picture of N1 sleep and the sleep onset, with as many as 9 stages, defined based on subtle and detailed signatures of the EEG (Hori et al., 1994). The large-scale perspective offered by the HMM, potentially used in combination with the Hori-system, could prove valuable for the development of a more consistent way of talking about the sleep onset period. It has already been discussed how N1-related HMM states consistently expressed inversed mean activities between the subcortical and cortical areas. A specific hypothesis could be that this particular distribution of brain activity implies that sleep has in fact been entered. Testing this would require a dataset that included additional behavioural measures, such as a continuous tracking of responsiveness or reaction times (Ogilvie and Wilkinson, 1984, Hori et al., 1994), which could be compared to the states derived by the HMM.

The fact that transitions into N1-related HMM states were most likely to happen from HMM state 8 may imply that the latter describes a large-scale configuration of brain activity from where sleep is facilitated. HMM state 8 expressed increased mean activity in areas corresponding to the DMN. This seems sensible, if, as often proposed, the DMN does in fact play a role in cognitive states of inwardly directed mentation, such as self-awareness, mind-wandering, and autobiographical memory (Buckner et al., 2008, Fox et al., 2015). The position of the DMN as a "gate" into the N1-related HMM states may thus be reflecting the mental retraction from the external world that commonly precedes sleep onset. One immediate and exciting extension of this finding would be to test if the particular configuration of brain activity, described by HMM

state 8, is less present in populations experiencing troubles initiating sleep. As an example, applying the HMM to data from a cohort of patients suffering from insomnia (Riemann et al., 2010, Riemann et al., 2015) and/or hypersomnia (Larson-Prior et al., 2014) could represent a promising approach to studying the involvement of large-scale networks in these invalidating disorders (see Introduction, 1.4.1).

6.4.5 *Vigilance-specific signatures of large-scale dynamics*

The HMM resolves large-scale brain activity in time, and hence makes it possible to study the dynamic features of the networks extracted from the ROI BOLD timecourses. This stands in contrast to static FC investigations carried out in Chapter 2, where networks are assumed to be stationary over time (see also Introduction, 1.5.8). Importantly, however, comparing the dynamic FC modelled by the HMM to the static FC results presented in Chapter 2 demonstrates that the two converge quite well when the HMM states are summed over time (as noted in the Results, 6.4.5). Weighting each HMM state by its vigilance-specific fractional occupancy thus yielded values of mean FC and SC-FC similarity corresponding to the static cases (compare Figure 6-6E to Figure 2-1D, and Figure 6-6F to Figure 2-3D). With the HMM it thus becomes clear that neither mean FC nor SC-FC similarity were stable within vigilance conditions. As an example, wakefulness included HMM states of relatively low (states 2, 8, and 10) and relatively high (states 5 and 18) mean FC values. Interestingly the HMM states of low mean FC seemed to belong to the 'pre-sleep' division of wakefulness, while those of high mean FC were mainly found in the 'wake-up' division. Furthermore, the negative association between mean FC and SC-FC similarity, which was demonstrated in Chapter 2, using a sliding-window approach (see Figure 2-5), was also captured by the HMM. As such, the states with weakest mean FC tended to show higher similarity to SC than the states with high mean FC.

While the HMM thus showed good correspondence with the static FC analyses, when summed over time, it further allowed for a direct investigation of vigilance-specific dynamic features, such as state switching and the related dynamic repertoire. These are concepts that have previously been speculated to change with vigilance. As discussed earlier, a decreased dynamic repertoire during sleep is a direct prediction of the Integrated Information Theory (Marshall et al., 2016), and used to explain the increased SC-FC similarity observed for SWS (Tagliazucchi et al., 2015) and general anaesthesia (Barttfeld et al., 2015) (see Chapter 2, 2.5.2), however, so far, direct empirical evidence for this has been scarce. From the present results it is quite clear that the HMM generally required a higher number of states to model wakefulness and N1 sleep, than to model the more consolidated stages of NREM sleep, N2 and SWS (see Figure 6-7). Furthermore, the state visits during wakefulness and N1 tended to have shorter duration, which led to significantly higher switching frequency compared to that of both N2 and SWS (see Figure 6-10A). Both of these observations align well with the ideas of the Integrated Information Theory that information-processing demands are higher during wakefulness than in consolidated stages of NREM sleep, and thus require a wider repertoire of states and more transitions. Importantly, however, switching levels could not distinguish between wakefulness and N1. In fact, N1 expressed the highest average switching frequency, albeit it was not significantly different from wakefulness. Fast switching does therefore not seem to be reserved for wakeful consciousness and mentation. As alluded to before, however, mental content is far from obliterated during N1 sleep, but rather

takes on a less controlled and often hypnagogic or even hallucinogenic character. An intriguing hypothesis could be that the common N1-related feeling of loosing control over thought processes is associated with a hyperactive switching that the HMM assigns to this condition (Goupil and Bekinschtein, 2012, Bareham et al., 2014).

Still, if the dynamic network features are in fact involved in mentation and spontaneous thought processes, one would expect them to change in some meaningful way from wakefulness to N1 sleep. Interestingly, when we estimated the similarity of the HMM states that were switched between within each vigilance condition, N1 stood out quite clearly (see Figure 6-10A). As such, the switched-between HMM states overlapped significantly more for N1 in terms of the areas involved, suggesting that the fast switching occurred between HMM states of little diversity. Taking all of this together, wakefulness could be distinguished from consolidated NREM sleep by its higher amount of switching, and from N1 sleep due to its higher diversity of activity distributions.

Chapter 7
Conclusion

In the preceding chapters I have provided a large-scale perspective of brain activity during wakefulness and sleep. Regardless of the recording modality, the measurement of spontaneous brain activity results in a plethora of signals, which can be analysed and decoded in numerous ways. The validity of any of these views of brain activity will ultimately depend on how well it captures and reflects operations of the brain, observed as behaviour and cognition. Wakefulness and sleep are fundamentally different, behaviourally as well as cognitively, and while it is beyond doubt that the brain is involved in these changes, it is less clear which aspects of brain activity that are responsible for them. In particular at the largest spatial scale of whole-brain networks, where the dimensionality and amount of possible interactions are close to innumerable, it remains challenging and unresolved how to best structure and interpret the recordings of brain activity. Identifying the large-scale signatures that differ between wakefulness and sleep is an important and fundamental first step towards understanding how the brain employs this level to support these but also less extreme changes in behaviour and cognition that we experience while awake, and processes that we are unaware of during sleep.

The main aim of my thesis has been to provide a novel dynamic large-scale network perspective on human sleep based on the analysis of unique human neuroimaging data. Specifically, as shown in the chapters of the thesis and summarised here, I have been able to provide new information based on integrating spatial and temporal aspects of brain activity both *during* and in the *transitions between* wakefulness and various stages of non-REM sleep. Overall, my findings are providing new and important insights for our current understanding of sleep, which has primarily been informed by findings from behaviour and spatially limited EEG recordings.

In the following sections, I will briefly summarise the main empirical findings of my thesis and, in light of these, revisit some of the outstanding questions regarding sleep, which were highlighted in the Introduction as contingent on evidence from the whole-brain perspective.

7.1 ***New insights***

7.1.1 *Robust changes in static functional connectivity between wakefulness and sleep*

When reviewing the literature on studies investigating large-scale network characteristics of brain activity during sleep, through concepts such as functional connectivity, resting-state networks, and graph theory it became clear that surprisingly few observations have in fact resulted in consistent models and explanations for the marked differences in behaviour, cognition and consciousness between wakefulness and sleep. Many potentially interpretable findings seemed matched by contradicting observations of their own or of other studies. The seeming lack of vigilance-sensitivity has led some researchers to altogether defy the role of large-scale FC in cognition and higher-order mentation (see e.g. Snyder (2016)), perhaps challenging the hope that such measures would be able to provide the whole-brain perspective of sleep that is called for.

However, seminal findings by Tagliazucchi and colleagues (Tagliazucchi et al., 2012b, Tagliazucchi and Laufs, 2014), later repeated by Altmann and colleagues (2016), represented a turning point, showing that vigilance can indeed be reliably decoded from fMRI FC networks.

The power of these studies compared to previous attempts were linked to their use of whole-brain FC networks, which were not constrained by assumptions such as resting-state network configurations. Taking this as a starting point for my own investigations, in Chapter 2 I performed a closer examination of comparative network differences between wakefulness and three NREM sleep stages (N1 sleep, N2 sleep, and slow-wave sleep (SWS)), using polysomnographically scored fMRI data from 18 healthy participants (from Prof Laufs' lab in Frankfurt, see Appendix). Paying particular attention to balancing the data at the participant level, I identified the global and regional vigilance-dependent differences in fMRI FC that may be regarded robust within a statistically rigorous framework. Compared to previously reported findings, obtained with less careful statistical inferences, differential FC networks were, according to my analyses, less pronounced and widespread with higher regional specificity. Still I was able to reproduce some previously highlighted characteristics of the vigilance stages, including mean FC increases along with subcortico-cortical connectivity decreases for N1, as well as generally decreased connectivity for SWS. Of note, the improved regional specificity allowed me to relate my findings to sleep-specific oscillations, such as K-complexes and spindles in case of the N2-spcific networks, and to the neurobiological literature of insomnia for the difference between wakefulness and N1.

From a more theoretical connectome-view, Chapter 2 further revisited recent findings suggesting that SWS is associated with an increased similarity between structural and functional connectivity. Probing this global measure in two independent fMRI datasets (from Prof Maquet's lab in Liege and Prof Laufs' lab in Frankfurt, see Appendix), I showed that this is indeed a highly consistent phenomenon across participants. The potential interpretation that a closer FC-dependence on the underlying structural backbone of SC signals a lowered state repertoire has previously been suggested to reflect decreased levels of consciousness associated with SWS and general anaesthesia (Barttfeld et al., 2015, Tagliazucchi et al., 2015). A novel method for deriving the regional determinants for this increased SC-FC similarity found FC of SWS to be particularly constrained within higher-order connections between fontal cortices and midline posterior regions, which seems in line with such an interpretation. Importantly, however, the SC-FC similarity did not express a linear relationship with vigilance, and N1 actually expressed values even lower than wakefulness, which obscures the view that a decrease in this measure should reflect consciousness *per se*. Furthermore, through a sliding-window analysis of the dynamic evolution of SC-FC similarity I found the mean level of FC as an important co-variate, which should be included in theoretical considerations of the relationship between structure and function.

7.1.2 *Extraction of objective large-scale spatiotemporal network signatures of vigilance*

While Chapter 2 focussed mainly on the static aspects of FC in the fMRI sleep datasets, the included sliding-window analyses did point to the important and well-established fact; that FC networks should not be regarded as stationary, but rather dynamic over time. This notion is easily reconciled with common

subjective experience during wakefulness, where several different mental and behavioural states are in constant flux, but sleep of course is also a temporally variable phenomenon, including NREM and REM stages, as well as stage-specific oscillations and graphoelements. Our understanding of sleep dynamics has in the meantime always relied on evidence from EEG and its related signatures – often identified through visual inspection. Even the highlighted machine-learning-based demonstrations of vigilance-specificity in large-scale FC networks were constrained by their training on polysomnographic EEG features of sleep (Tagliazucchi et al., 2012b, Tagliazucchi and Laufs, 2014, Altmann et al., 2016). Driven by the nagging conviction that the superior spatial resolution of fMRI ought to hold potential for more than merely provide correlates of sleep stages already defined from the crude whole-brain representation of EEG, Chapter 6 set out to investigate whether a completely data-driven dynamic decomposition of the sleep-containing fMRI dataset (from Prof Laufs' lab in Frankfurt, see Appendix) would reflect vigilance in a meaningful way. Using a hidden Markov model (HMM) to decompose the BOLD timecourses into a set of temporally resolved large-scale network states, and subsequently comparing these to the EEG-based sleep-scoring showed an unequivocal correspondence. As such, vigilance was indeed found objectively extractable from the fMRI data without explicit use of the EEG. Importantly, the fit between the HMM states and polysomnography got better as the number of HMM states were increased above 4, suggesting that more than the original four traditionally defined vigilance stages were present in the data.

The HMM states organised in such a way that a high number were specific for wakefulness and N1, while N2 and SWS were modelled by fewer states. Interestingly, I found the states with high activity during wakefulness to express large-scale patterns resembling known higher-order RSNs. This may suggest that these networks, previously found rather resilient to changes in vigilance in static FC analyses, are in fact wakefulness-specific. A specific state resembling the default-mode network (DMN) turned out to represent the most common route from wakefulness to N1 sleep – a finding I proposed could be valuable for future neuroimaging investigations of sleep disorders, such as insomnia. N1-related states were characterised by desynchronised activities between subcortical areas, including the thalamus, and the primary sensory cortices, in parallel with the subcortico-cortical dropout in FC and intracortical evidence (Magnin et al., 2010). Activity distributions of the states modelling N2 showed overlap with reported activity maps for sleep spindles, while the main state of SWS showed widespread decreases in activity in frontal cortices, and very localised increases in and around the supplementary motor areas.

The temporal resolution of the HMM decomposition further allowed me to measure the vigilance-specific levels of FC dynamics. Wakefulness and N1 clearly differentiated from N2 and SWS by showing markedly more switching between the derived HMM states. Intriguingly, while switching was found comparable between wakefulness and N1 sleep, the latter tended to include switches between states of higher similarity, suggesting that wakefulness was unique in terms of its high switching combined with high state diversity.

7.1.3 *Using MEG to increase spatiotemporal information about sleep FC*

Given the indirect relationship between fMRI and neural activity it was important to investigate how large-scale aspects of sleep may unfold in electrophysiological data. Recent advances in the source-reconstruction

of and FC derivation from whole-head MEG data, makes this modality particularly interesting for such purposes. In Chapters 3, 4 and 5 I presented a comprehensive examination of large-scale networks of power envelope correlations in a unique set of MEG recordings from 11 sessions in 8 healthy participants in wakefulness, N2 sleep and SWS (from Prof Van Someren's lab in Amsterdam, see Appendix).

In Chapter 3, I focussed on resolving large-scale FC networks in six canonical, and sleep-relevant, frequency bands. I used recently developed and validated methods to source-reconstruct the MEG signals and construct FC networks comprising 39 cortical areas for each frequency band, participant and vigilance condition, while importantly correcting for the non-trivial source-leakage bias (Colclough et al., 2015). To my knowledge Chapter 3 represents the first example of source-space MEG FC of sleep. Functional connections were estimated as correlations (full as well as partial) between power envelopes, band-pass filtered into six frequency bands, ranging from 0.25 – 48 Hz.

Comparing these networks in terms of global and regional FC across and within the included vigilance conditions, I was able to show that consistent MEG FC patterns can be derived, not only from wakefulness as shown previously, but indeed also from NREM sleep. As expected, intra-vigilance FC consistency was high in the alpha and beta bands for wakefulness. N2 showed high consistency across most of the frequency bands, and uniquely in the sigma spindle band, but not in the gamma band. Interestingly, and somewhat surprisingly SWS only showed signs of consistent FC patterns in the alpha band – a frequency range normally regarded as suppressed during this sleep stage. Of importance I showed that FC information was more specific to vigilance than averaged power values of the included cortical areas.

In parallel to the fMRI analysis of Chapter 2, I resolved the pairwise MEG network differences between the vigilance conditions. In particular, when using partial correlations, with the aim of discounting indirect connectivity, it was clear that the differential networks expressed some overlap with the ones identified in Chapter 2, using equivalent contrasts, for fMRI. Given the observation that FC patterns appeared much more preserved from wakefulness to sleep in fMRI relative to MEG, I proposed that MEG FC might be capturing the use-dependent electrophysiological signatures of each of the vigilance stages, which get filtered into fMRI FC networks. Providing evidence not just from the resting-state, but also from sleep, my findings are relevant to the overall aim of finding a translation between the widely used fMRI FC networks and electrophysiology.

The frequency distribution and spatial patterns of FC in N2 suggested that the use-dependent signatures of this stage might be related to sleep oscillations, such as K-complexes and spindles. The cross-frequency interactions investigated in Chapter 4 provided evidence in line with this interpretation. As such two distinct clusters of cross-frequency interactions were formed for N2. Probing the link between transient sleep oscillations and MEG FC even more directly in Chapter 5, I found that N2 was particularly dependent on frequency-specific changes in power. When tacked over time through a sliding-window analysis, mean FC was highly variable in N2 and appeared driven by increases in spectral power in the higher frequency bands – beta in particular.

Although MEG FC of wakefulness has been investigated before in a number of resting-state studies, the comparison to sleep-derived FC revealed some novel interesting aspects. In the cross-frequency

investigations of Chapter 4 I was thus able to show that integration of the gamma band within a multi-layer FC network was particularly strong for wakefulness (Tewarie et al., 2016b). This is interesting, given the role that gamma oscillations are thought to play in higher-order cognition (Singer and Gray, 1995, Jensen et al., 2007, Fries, 2009). Furthermore, in Chapter 5 I identified a negative correlation between sensor power in the alpha band and mean FC of the beta band as a specific trait of wakefulness. I interpreted this in light of the so-called inhibition-timing hypothesis (Klimesch et al., 2007) and related ideas regarding cognitive functions of alpha oscillations (Mazaheri and Jensen, 2010).

The strongest FC observed for SWS was found between temporal and frontal cortices in the alpha band. This is exciting given the recent spectral EEG evidence pointing to increased alpha activity during SWS as a specific marker of overnight performance-improvement on word-pairing (Marshall et al., 2006) and gain of explicit knowledge of hidden rules in number sequences (Yordanova et al., 2012).

Overall I have strived to provide an overview of what the additional spectral information offered by MEG compared to fMRI FC networks can tell us about large-scale aspects of wakefulness and sleep, as well as the contrast between them. Notably, the added temporal resolution of MEG comes at the expense of an expansive data dimensionality, which requires patience and rigour to navigate through. With this in mind it can seem almost unfathomable that the analyses presented in this thesis are only scraping the surface of the vast amount of temporal information MEG and electrophysiology provides. By focussing on the slowly fluctuating component of MEG power envelopes I have presented findings of merely a part of the temporal scale. Recent work has demonstrated that much faster large-scale interactions can be resolved with MEG. The application of the HMM has proven particularly promising in this respect (Baker et al., 2014a, Vidaurre et al., 2016). Given my successful use of an HMM on fMRI in relation to sleep, and obvious next step is to explore its performance on MEG sleep data.

7.2 ***Perspectives***

7.2.1 *The role of large-scale networks in the transitioning between wakefulness and sleep*

As discussed in section 1.4.1 the subjective experience of falling asleep tells us that sleep is more than a fully automated behaviour. We are able to voluntarily defy sleep, even in the face of immense tiredness. This is difficult to explain from the perspective of low-level sleep regulation alone, and suggests an involvement of higher cortical areas. In Chapter 6, the HMM-based dynamic decomposition of large-scale networks near the onset of sleep suggested that a specific configuration of brain activity may be required for sleep initiation. Specifically, the transition into N1 sleep was found most likely to occur from a pattern of brain activity resembling the default-mode network (DMN) in wakefulness. Findings of this nature could guide new hypothesis regarding the involvement of large-scale networks in the process of falling asleep, as well as in conditions where sleep initiation is disturbed, such as insomnia or indeed hypersomnia disorders.

7.2.2 *Theories of local sleep at large?*

Sleep has traditionally been regarded a global phenomenon of the brain. However, as mentioned in section 1.4.2, recent prominent theories posit that sleep may in fact be fundamentally local, in terms of its

evolutionary origin (Krueger et al., 2016) and functions (Krueger and Tononi, 2011, Tononi and Cirelli, 2013). The slow oscillation and slow-wave activities are thus thought to exert local plastic effects, e.g. at the level of synapses. The observation that slow-wave activity is not homogeneously distributed across the brain, but rather expressed in areas in proportion to the labour put on these prior to sleep is key to such theories (Murphy et al., 2011). The high density of slow-wave activity in frontal cortices is believed to reflect this phenomenon. In Chapter 6, the marked decreases in mean BOLD activity and concomitant increases in FC in the frontal cortices, found during SWS, are in line with these ideas. Furthermore, in Chapter 2 I speculated that the plasticity-promoting effects of SWS might also be reflected at the macro-scale of SC networks. If the brain is generally more prone to plastic change during SWS, it is conceivable that the structural connectome gets shaped by functional connectivity patterns expressed in this particular stage. Although direct evidence of such an effect of sleep do not currently exist, it offers a theoretical explanation for the observed SWS-specific increase in SC-FC similarity.

7.2.3 *Whole-brain signatures of sleep-dependent memory and learning*

Sleep is repeatedly being highlighted for its role in memory and learning, which has been consistently demonstrated in behavioural studies (Stickgold, 2005, Diekelmann and Born, 2010). A number of findings in this thesis may be interpreted as potential large-scale correlates of the neural processes underlying these phenomena. Chapter 3 thus showed that the most marked MEG FC patterns of SWS is found in the alpha frequency band in a network between temporal and frontal cortices. Fronto-temporal connectivity is generally believed to play a key role in the long-term consolidation of memories (Simons and Spiers, 2003), and the alpha band has specifically been shown to correlate with SWS-dependent improvement of word-pairing (Marshall et al., 2006) and gain of explicit knowledge of hidden number rules (Yordanova et al., 2012). In Chapter 6 mean BOLD activity was shown to increase very specifically in the supplementary motor area and paracentral lobule, areas that have been implicated in SWS-dependent consolidation of motor-skill learning (Tamaki et al., 2013).

7.2.4 *Refined spatio-temporal features for sleep staging*

Since the 1930's the understanding of sleep has been intimately linked to the EEG-based definitions of sleep stages (Loomis et al., 1935a, b, Davis et al., 1937, Loomis et al., 1937). Whether EEG patterns are derived through visual inspection, as in standard polysomnography (Rechtschaffen and Kales, 1968, Iber et al., 2007), or by automated spectral algorithms (Olbrich and Achermann, 2005), the low spatial resolution of the EEG ultimately ignores much of the spatial complexity of whole-brain activity. In Chapter 6 it was shown how the inclusion of precise spatial information, offered by fMRI, lead to an alternative data-driven state decomposition with significant overlap with traditional polysomnography, but importantly also with considerable differences. This approach could aid in future improvements of objective criteria for defining vigilance-level and for identifying transitions.

7.2.5 *Potential implications for consciousness?*

As alluded to in section 1.4.6, sleep and equivalent alterations of vigilance have been instrumental in the search for neural correlates of consciousness (Koch et al., 2016). The approach of contrasting brain activity during wakefulness with that of either sleep or general anaesthesia specifically aims at identifying the substrates controlling the *state* of consciousness as opposed to the *content* of consciousness (Dehaene et al., 2014). Since each of the empirical chapters of this thesis report on exactly such contrasts, it is worth asking if the identified differences may be related to consciousness. From the perspective of static FC, comparative network analyses in Chapter 2 and 3 showed, consistently across fMRI and MEG, that wakefulness was differentiated from NREM sleep by strong connectivity in and between occipital, temporal and posterior parietal cortices. This constellation of connections overlap with a so-called 'posterior cortical hot zone', which was recently proposed as a neural correlate of consciousness (Koch et al., 2016). The strongly increased similarity to SC of FC patterns in SWS compared to wakefulness has been interpreted by others (Barttfeld et al., 2015, Tagliazucchi et al., 2015) and in Chapter 2 as a sign that wakefulness is associated with more exploration of the state repertoire, in line with the Integrated Information Theory of consciousness (Tononi, 2008, Tononi et al., 2016). Chapter 6 provided, what is, to the best of my knowledge, the first direct evidence of an increased large-scale network repertoire in wakefulness compared to NREM sleep, by estimating the amount of switching between HMM states and the diversity of the states switched between. Findings of Chapter 6 further suggest that resting-state network configurations believed to reflect higher-order cognition and spontaneous thought processes were specific to wakefulness and rarely found in sleep stages deeper than N1. This would seem in line with ideas of the Global Neuronal Workspace Theory (Baars, 1993, Dehaene et al., 2003, Dehaene and Changeux, 2011), that consciousness results from a distributed network of interacting brain regions with some regions being able to broadcast information globally.

The importance of coming up with new ways of understanding the brain mechanisms underlying different states of consciousness is clearly seen in relation to general anaesthesia. A primary concern of an anaesthesiologist is to make sure that a patient is in fact fully unconscious, when "put under" for surgery. It is becoming increasingly clear, however, that standard EEG-based criteria do not always suffice to rule out 'leftover' consciousness and awareness in seemingly anaesthetised patients, as the so-called isolated forearm test can vividly demonstrate (Sanders et al., 2012). A recently proposed measure of slow-wave saturation could be one way of objectively identifying loss of consciousness (Ní Mhuircheartaigh et al., 2013), along with the so-called perturbational complexity index (Casali et al., 2013, Sarasso et al., 2015). Meanwhile, the large-scale network view, promoted in my thesis could also turn out to be informative for these purposes.

Appendix

Datasets included in thesis

fMRI dataset 1 (Liege)

Acquisition of fMRI and simultaneous EEG

fMRI time series were acquired using a 3-Tesla MR scanner (Allegra; Siemens). Multislice T2*-weighted fMRI images were obtained with a gradient echo-planar sequence using axial slice orientation [32 slices; voxel size, 3 × 3.4 × 3mm^3; matrix size, 64 × 64 × 32; repetition time (TR) = 2,460 ms; echo time (TE) = 40 ms; flip angle = 908; field of view (FOV) = 220 mm; delay = 0]. The three initial scans were discarded to avoid magnetic saturation effects. A structural T1-weighted 3D Magnetization Prepared RApid Gradient Echo (MP-RAGE) sequence (TR = 1,960 ms; TE = 4.43 ms; inversion time, 1,100 ms; FOV = 230 × 173 mm2; matrix size, 256 × 192 × 176; voxel size, 0.9 × 0.9 × 0.9 mm^3) was also acquired in all participants.

EEG was recorded simultaneously with the fMRI, using two MR-compatible 32-channel amplifiers (Brainamp MR plus; Brain Products) and an MR-compatible EEG cap (Braincap MR; Falk Minow Services) with 64 ring-type electrodes. The EEG cap included 62 scalp electrodes that were online referenced to a frontocentral electrode (FCz), as well as one electrooculogram (EOG) and one electro-cardiogram (ECG) channel. Using abrasive electrode paste (ABRALYT 2000; FMS), electrode–skin impedance was kept below 5 kOhm in addition to the 5-kOhm resistor built into the electrodes. To avoid movement-related EEG artifacts, participants' heads were immobilized in the head coil by a vacuum pad. EEG was digitized at a 5,000-Hz sampling rate with a 500-nV resolution. Data were analog filtered by a band-limiter low-pass filter at 250Hz (30 dB per octave) and a high-pass filter with 10-s time constant corresponding to a high-pass frequency of 0.0159 Hz. Data were transferred outside the scanner room through fiber-optic cables to a personal computer where the EEG system running Vision Recorder Software v1.03 (Brain Products) was synchronized to the scanner clock. Sleep EEG was monitored online with BrainProducts RecView Software. For analysis, EEG data were low-pass filtered (finite impulse response filter, −36 dB at 70 Hz), down-sampled to 250Hz, and re-referenced to linked mastoids. Scanner artifacts were removed in Vision Analyzer, using an adaptive average subtraction. Ballistocardiographic artifacts were removed using an algorithm based on independent component analysis.

Sleep staging of scanning artefacts-corrected data followed the standard criteria of Rechtschaffen and Kales (Rechtschaffen and Kales, 1968).

Participants

Twenty-five young healthy participants (eleven females, mean age of 22.0 years; age range 18-25 years) were recruited by advertisement. From this group, 13 participants maintained periods of steady NREM sleep. Demographic and sleep parameters for these subjects have been reported in a previous study. 12 out of the 13 participants reached slow-wave sleep. None were on medication or sleep deprived. Written informed consent was obtained from all participants. The study was approved by the Ethics Committee of the Medical School at the University of Liege.

fMRI dataset 2 (Frankfurt)

Acquisition of fMRI and polysomnography

fMRI was acquired on a 3 T system (Siemens Trio, Erlangen, Germany) with the following settings: 1505 volumes of T2*-weighted echo planar images with a repetition time (TR) of 2.08 seconds, and an echo time of 30 ms; matrix 64 x 64, voxel size 3 x 3 x 2 mm^3, distance factor 50 %, FOV 192 mm^2.

Simultaneous polysomnography was performed through the recording of EEG, EMG, ECG, EOG, pulse oximetry, and respiration. EEG was recorded using a cap (modified BrainCapMR, Easycap, Herrsching, Germany) with 30 channels, of which the FCz electrode was used as reference. The sampling rate of the EEG was 5 kHz, and a low-pass filter was applied at 250 Hz. EMG was collected with chin and tibial derivations, and as the ECG and EOG recorded bipolarly at a sampling rate of 5 kHz with a low-pass filter at 1 kHz. Pulse oximetry was collected using the Trio scanner, and respiration with MR-compatible devices (BrainAmp MR+, BrainAmp ExG; Brain Products, Gilching, Germany).

Sleep classification was performed by a sleep expert based on the EEG recordings in accordance with the AASM criteria (2007).

Participants

71 participants were selected from a larger database. Exclusion criteria focussed on the quality of the concomitant acquisition of EEG, EMG, fMRI, and physiological recordings. Participants were instructed to lie still in the scanner with their eyes closed and relax. Written informed consent was obtained, and the study was approved by the local ethics committee.

MEG dataset

MEG recordings were acquired with 151 axial gradiometers of a CTF-151 system (CTF Systems Inc., Port Coquitlam, Canada) from 11 healthy participants (6 males, 19-27 years of age) with no history of neurological disorders. Participants gave written informed consent, and the study protocol was approved by the medical ethics committee of the VU University Medical Center, Amsterdam, The Netherlands. Each participant was recorded on two separate days. Both days included ~80 minutes of attempted daytime sleep with participants placed in a supine position inside the scanner with head movements limited by foam pads (for a full description of recording procedure and experimental conditions, see Piantoni et al. (2015)). In order to facilitate their ability to fall asleep, participants were instructed to restrict their sleep to a maximum of 5 hours the night prior to each recording day.

The MEG signals were recorded in a magnetically shielded room at a sampling frequency of 250 Hz. Based on the axial gradiometer traces, the MEG sleep recordings were qualitatively scored by an expert sleep technician according to the criteria of Rechtschaffen and Kales (Rechtschaffen and Kales, 1968, Simon et al., 2000). As such, each 30-second window of data was scored as belonging to one of the following six categories: wakefulness, rapid eye movement sleep (REM), non-REM sleep stage 1 (NREM-1), non-REM sleep stage 2 (NREM-2), non-REM sleep stage 3 (NREM-3), non-REM sleep stage 4 (NREM-4).

DTI dataset

Acquisition of Diffusion Tensor Imaging data

MRI data (structural MRI and DTI) were collected in one session on a 3T Siemens Skyra scanner at CFIN, Aarhus University, Denmark. The parameters for the structural MRI T1 scan were as follows: voxel size of 1 mm^3; reconstructed matrix size 256x256; echo time (TE) of 3.8 ms and repetition time (TR) of 2300 ms. The Diffusion Tensor Imaging (DTI) data were collected using TR = 9000 ms, TE = 84 ms, flip angle = 90˚, reconstructed matrix size of 106x106, voxel size of 1.98x1.98 mm with slice thickness of 2 mm and a bandwidth of 1745 Hz/Px. Furthermore, the data were collected with 62 optimal nonlinear diffusion gradient directions at b=1500 s/mm^2. Approximately one non-diffusion weighted image (b=0) per 10 diffusion weighted images was acquired. Additionally, the DTI images were collected with different phase encoding directions. One set was collected using anterior to posterior phase encoding direction and the second acquisition was performed in the opposite direction. The matrices of both phase encoding directions were thresholded, such that the top 99% connections were retained.

Participants

All participants were recruited through the online recruitment system at Aarhus University. Data were collected at CFIN, Aarhus University, Denmark, from 16 healthy right-handed participants (11 men and 5 women, mean age: 24.75 ± 2.54). Participants with psychiatric or neurological disorders (or a history thereof) were excluded from participation in this study. This study was approved by the internal research board at CFIN, Aarhus University, Denmark. Ethics approval was granted by the Research Ethics Committee of the Central Denmark Region (De Videnskabsetiske Komitéer for Region Midtjylland). Written informed consent was obtained from all participants prior to participation.

Supplementary figures

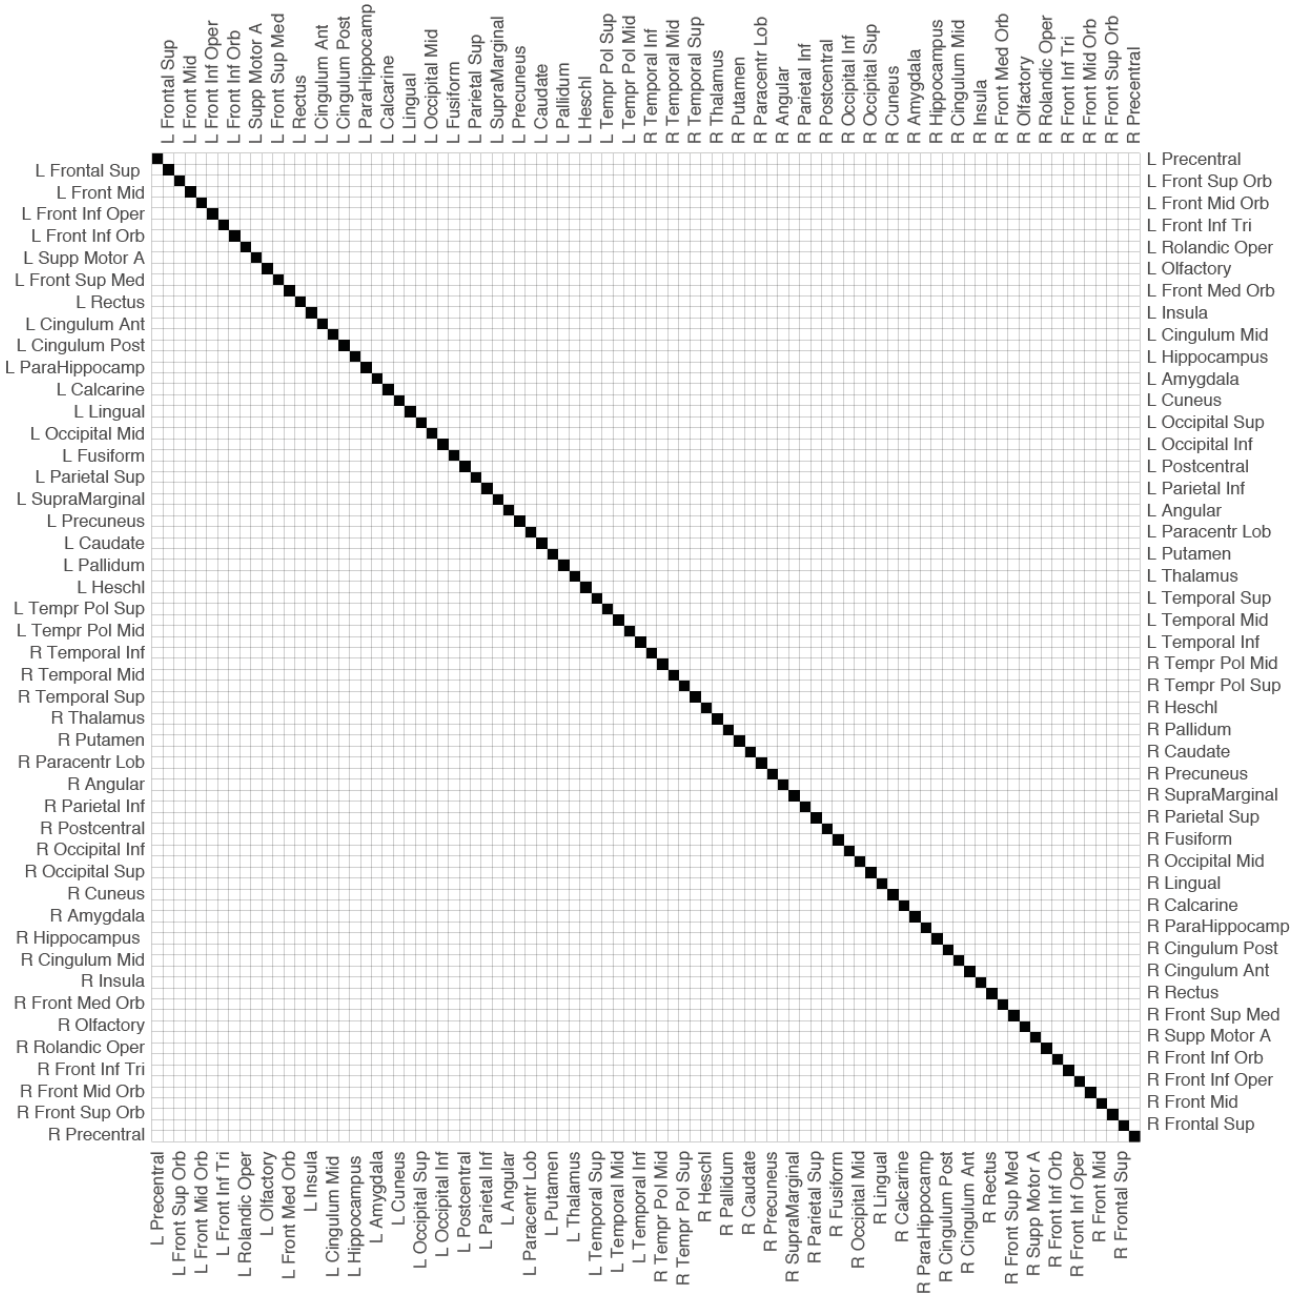

Figure S 1 ***Matrix organisation of AAL regions-of-interest*** AAL ROIs were organised into left (matrix entries 1 to 45) and right (matrix entries 46 to 90) hemisphere. (Supplement to Chapters 2 and 6).

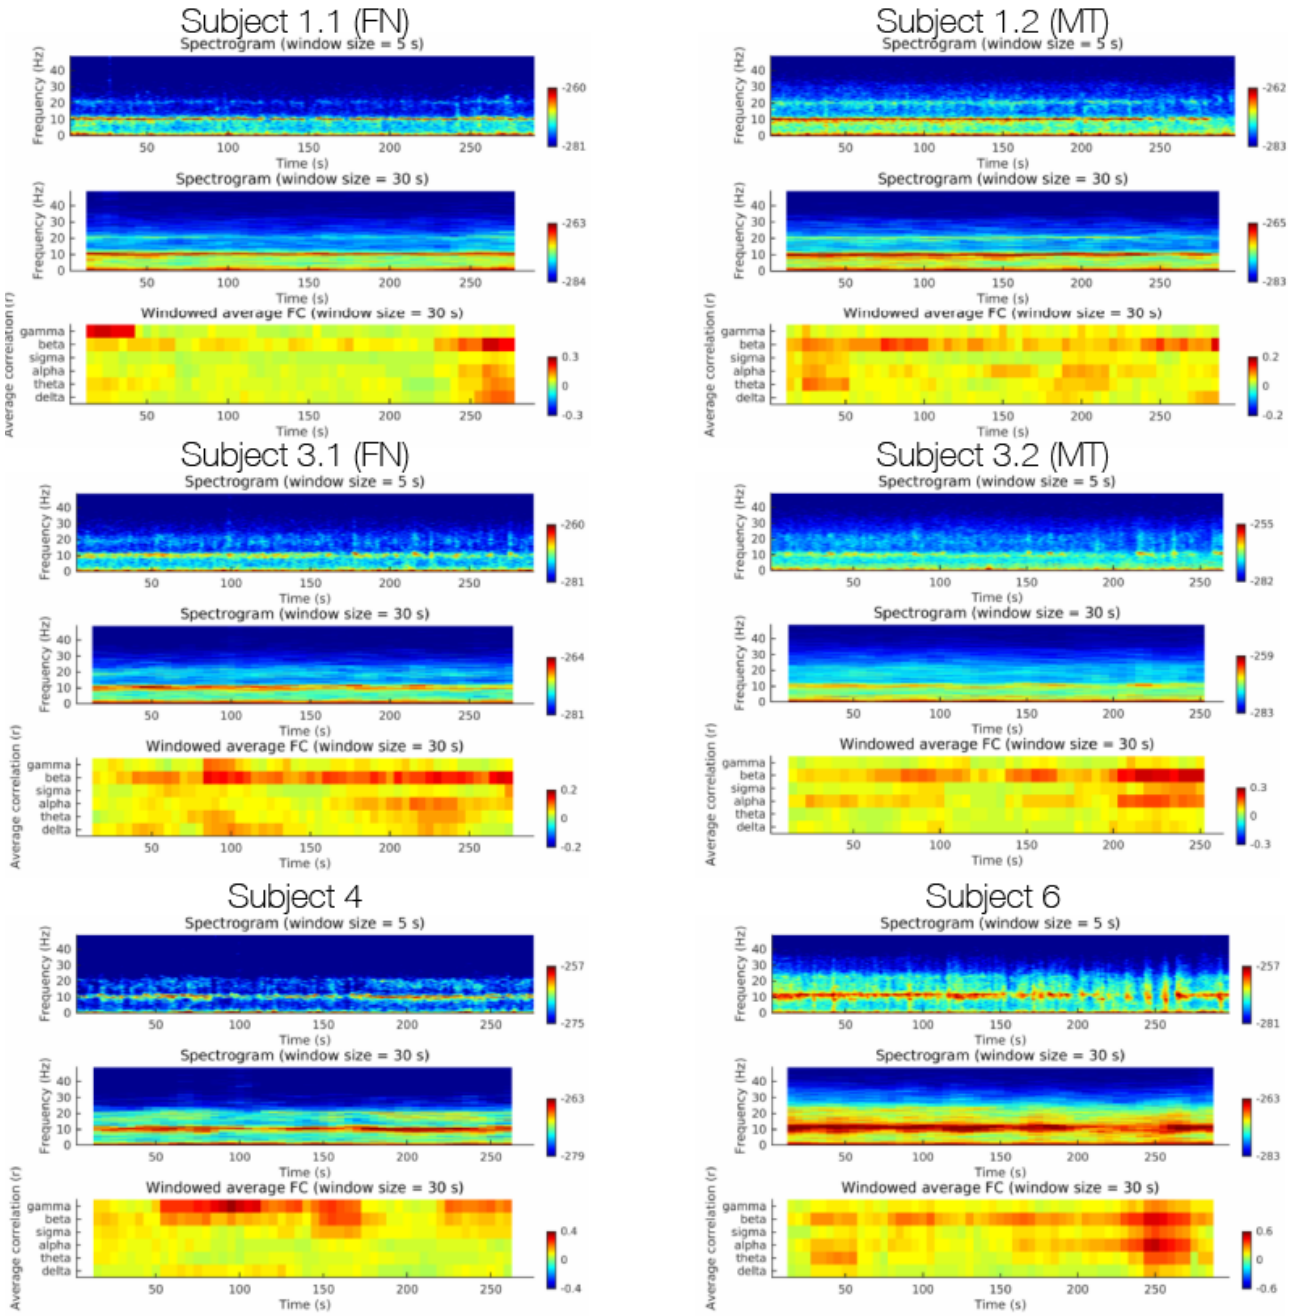

Figure S 2 ***Time-frequency representations of relationship between sensor power and mean level of FC in wakefulness from participant 1 to 4 (***Supplement to **Figure 5-2*).*** Spectrograms were computed from the broadband MEG sensor timecourses using the short-time Fourier transform in MATLAB, one with a window size of 5 seconds and 4 seconds overlap, and one with a window size of 30 seconds and 25 seconds overlap. The latter corresponded to the window size and overlap of the sliding window mean FC computation. 'FN' and 'MT' denotes the two different recording days. Note how decreases in power in the alpha frequency band often co-occurred with an increase in mean FC, particularly in the beta band.

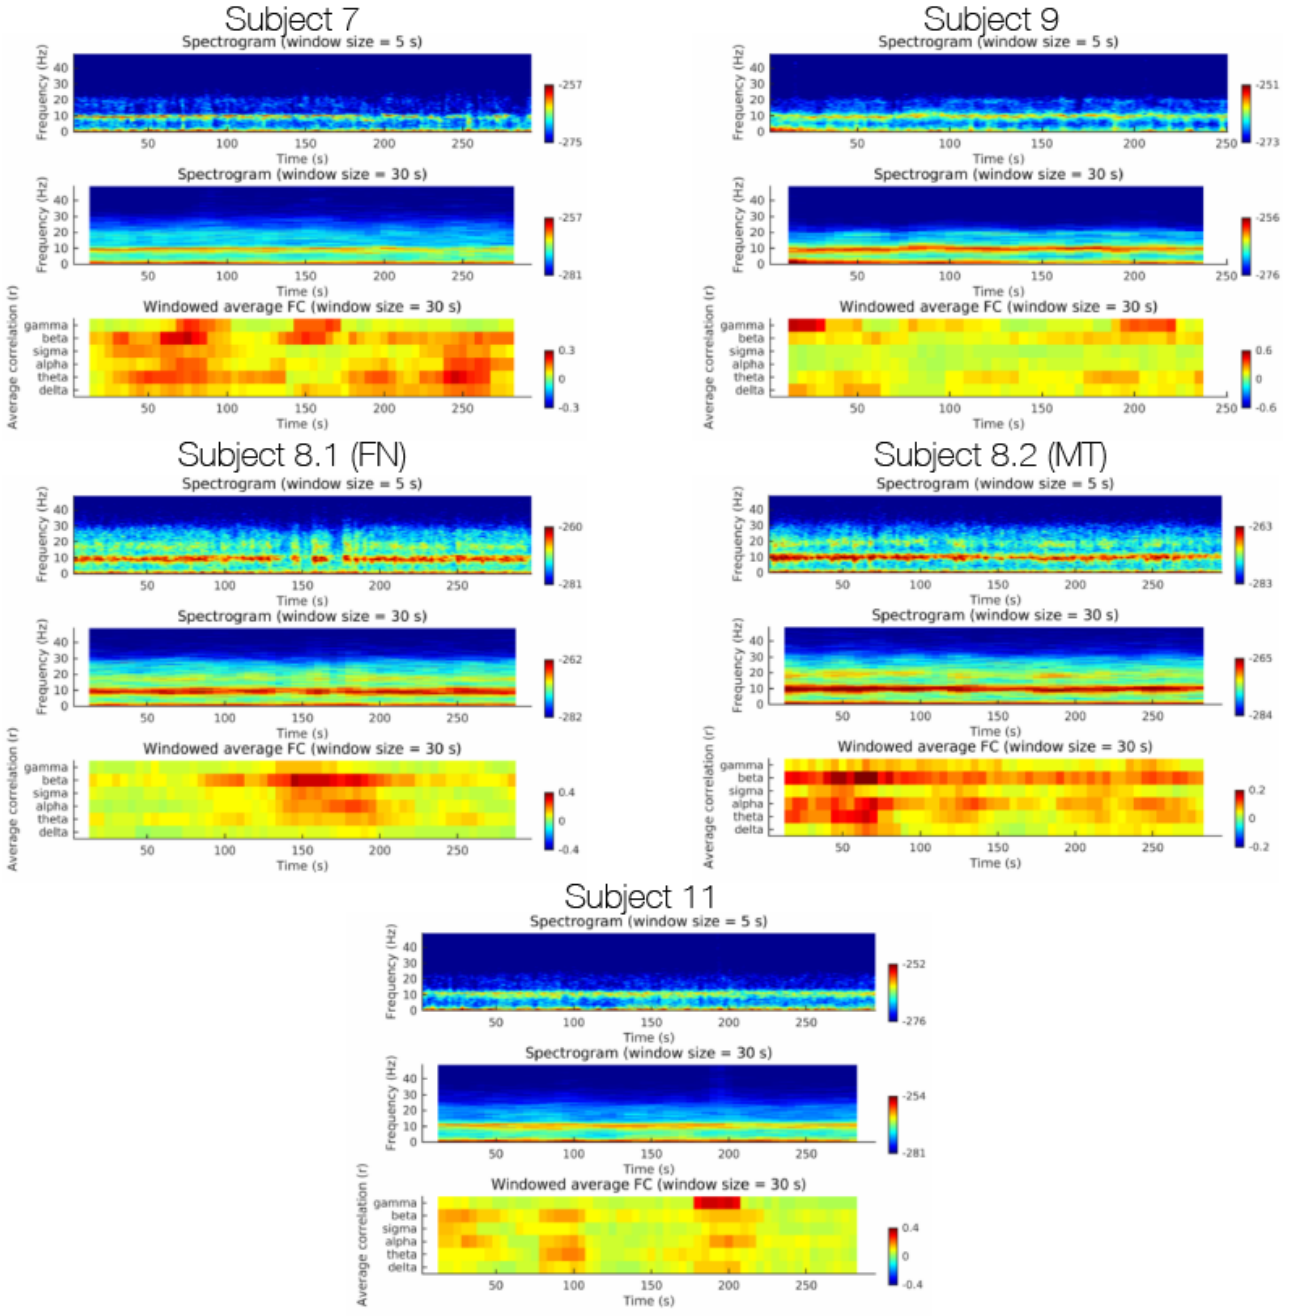

Figure S 3 ***Time-frequency representations of relationship between sensor power and mean level of FC in wakefulness from participant 5 to 8 (***Supplement to **Figure 5-2*).*** Spectrograms were computed from the broadband MEG sensor timecourses using the short-time Fourier transform in MATLAB, one with a window size of 5 seconds and 4 seconds overlap, and one with a window size of 30 seconds and 25 seconds overlap. The latter corresponded to the window size and overlap of the sliding window mean FC computation. 'FN' and 'MT' denotes the two different recording days. Note how decreases in power in the alpha frequency band often co-occurred with an increase in mean FC, particularly in the beta band.

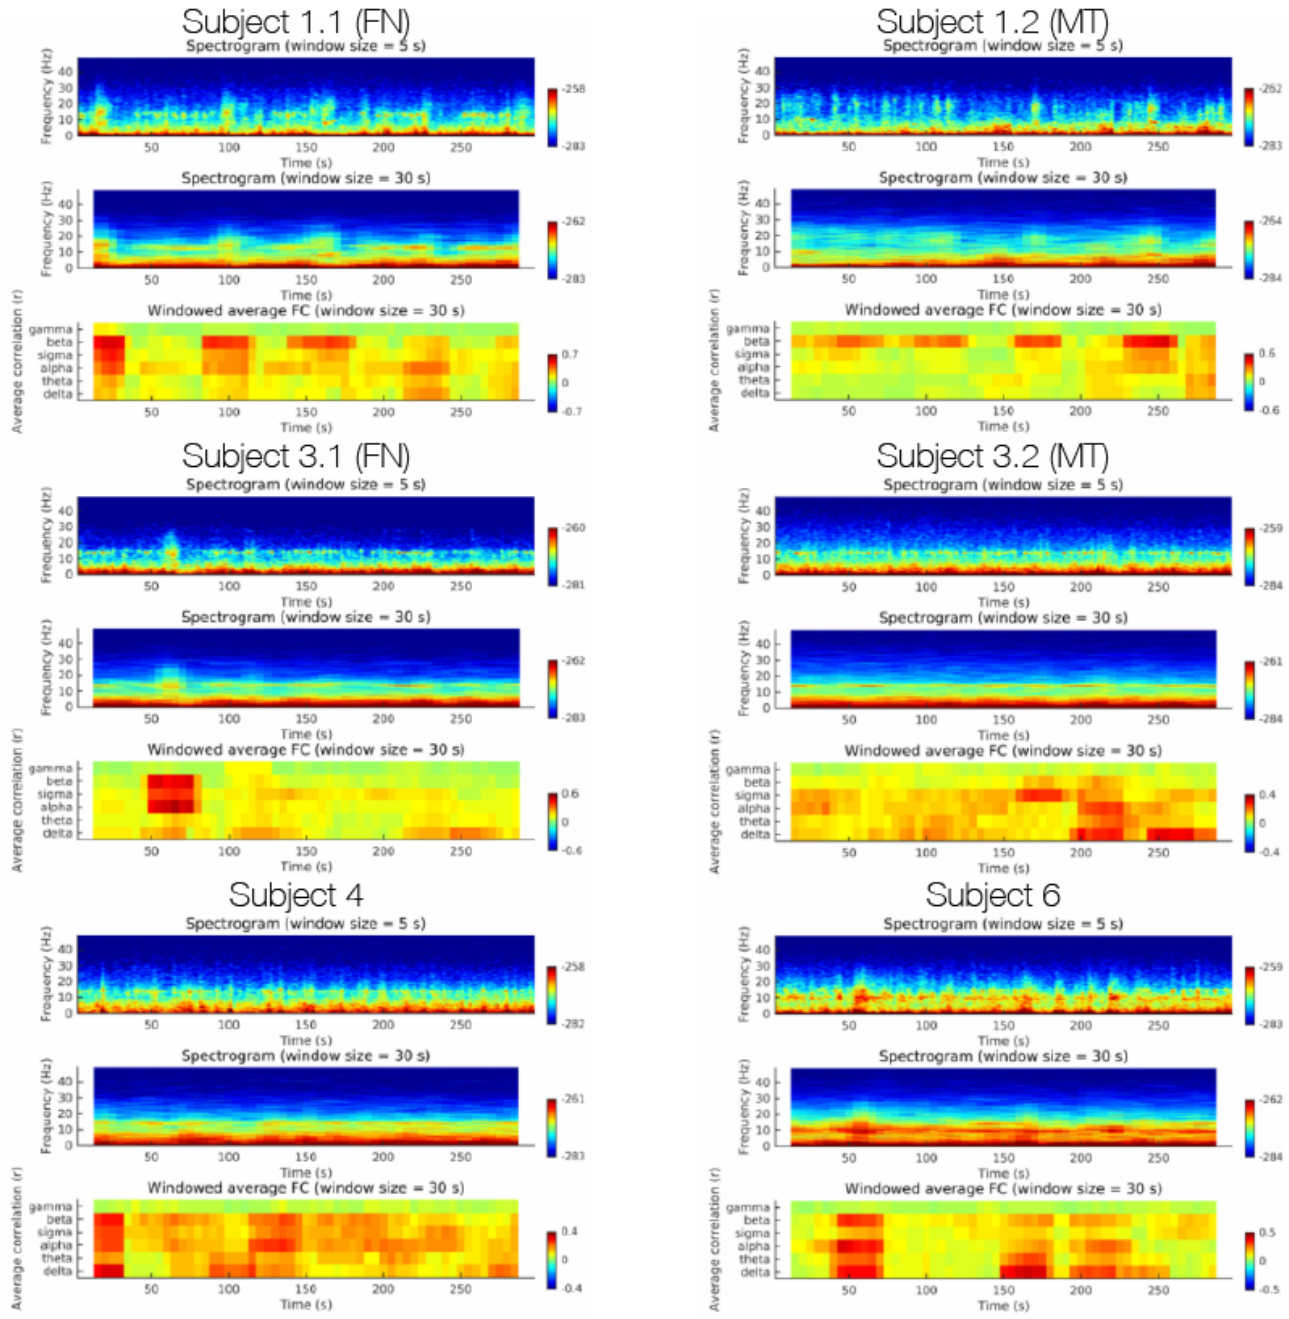

Figure S 4 ***Time-frequency representations of relationship between sensor power and mean level of FC in N2 sleep from participant 1 to 4 (***Supplement to **Figure 5-2*).*** Spectrograms were computed from the broadband MEG sensor timecourses using the short-time Fourier transform in MATLAB, one with a window size of 5 seconds and 4 seconds overlap, and one with a window size of 30 seconds and 25 seconds overlap. The latter corresponded to the window size and overlap of the sliding window mean FC computation. 'FN' and 'MT' denotes the two different recording days. Note how specific transient spectral power increases appear to be driving mean FC across a range of frequencies.

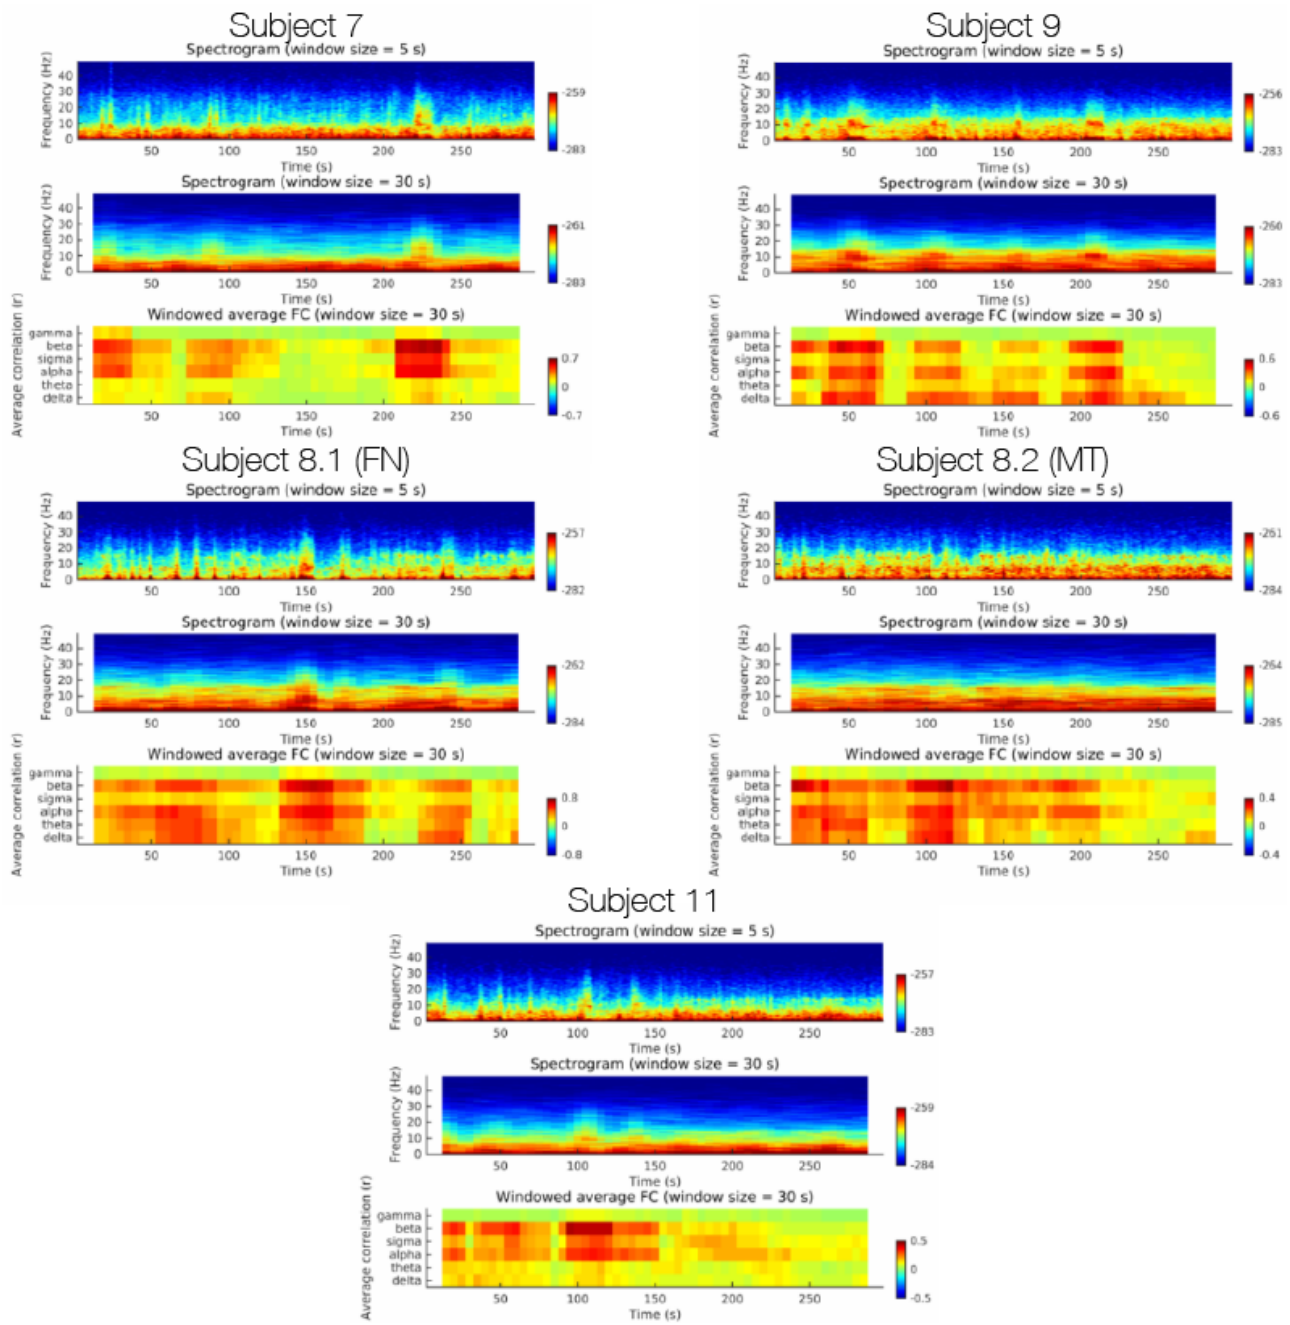

Figure S 5 ***Time-frequency representations of relationship between sensor power and mean level of FC in N2 sleep from participant 5 to 8 (***Supplement to **Figure 5-2*).*** Spectrograms were computed from the broadband MEG sensor timecourses using the short-time Fourier transform in MATLAB, one with a window size of 5 seconds and 4 seconds overlap, and one with a window size of 30 seconds and 25 seconds overlap. The latter corresponded to the window size and overlap of the sliding window mean FC computation. 'FN' and 'MT' denotes the two different recording days. Note how specific transient spectral power increases appear to be driving mean FC across a range of frequencies.

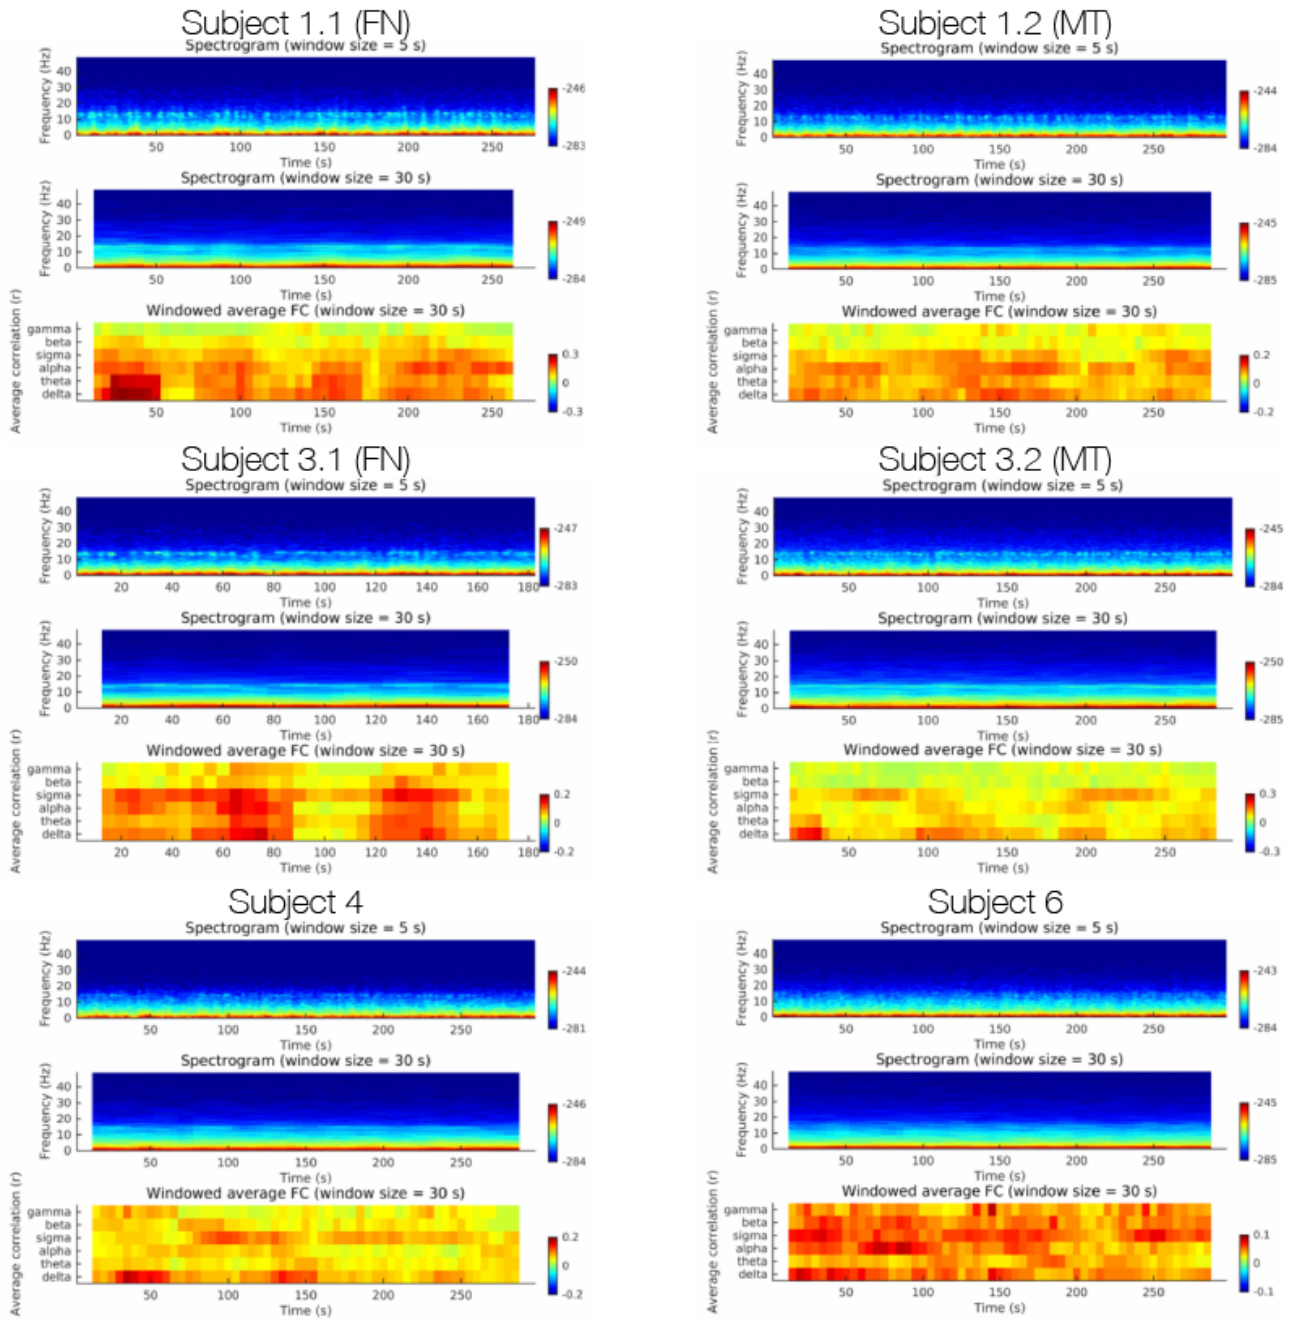

Figure S 6 ***Time-frequency representations of relationship between sensor power and mean level of FC in SWS from participant 1 to 4 (***Supplement to **Figure 5-2*).** Spectrograms were computed from the broadband MEG sensor timecourses using the short-time Fourier transform in MATLAB, one with a window size of 5 seconds and 4 seconds overlap, and one with a window size of 30 seconds and 25 seconds overlap. The latter corresponded to the window size and overlap of the sliding window mean FC computation. 'FN' and 'MT' denotes the two different recording days. No clear associations between sensor power and mean FC are visible.

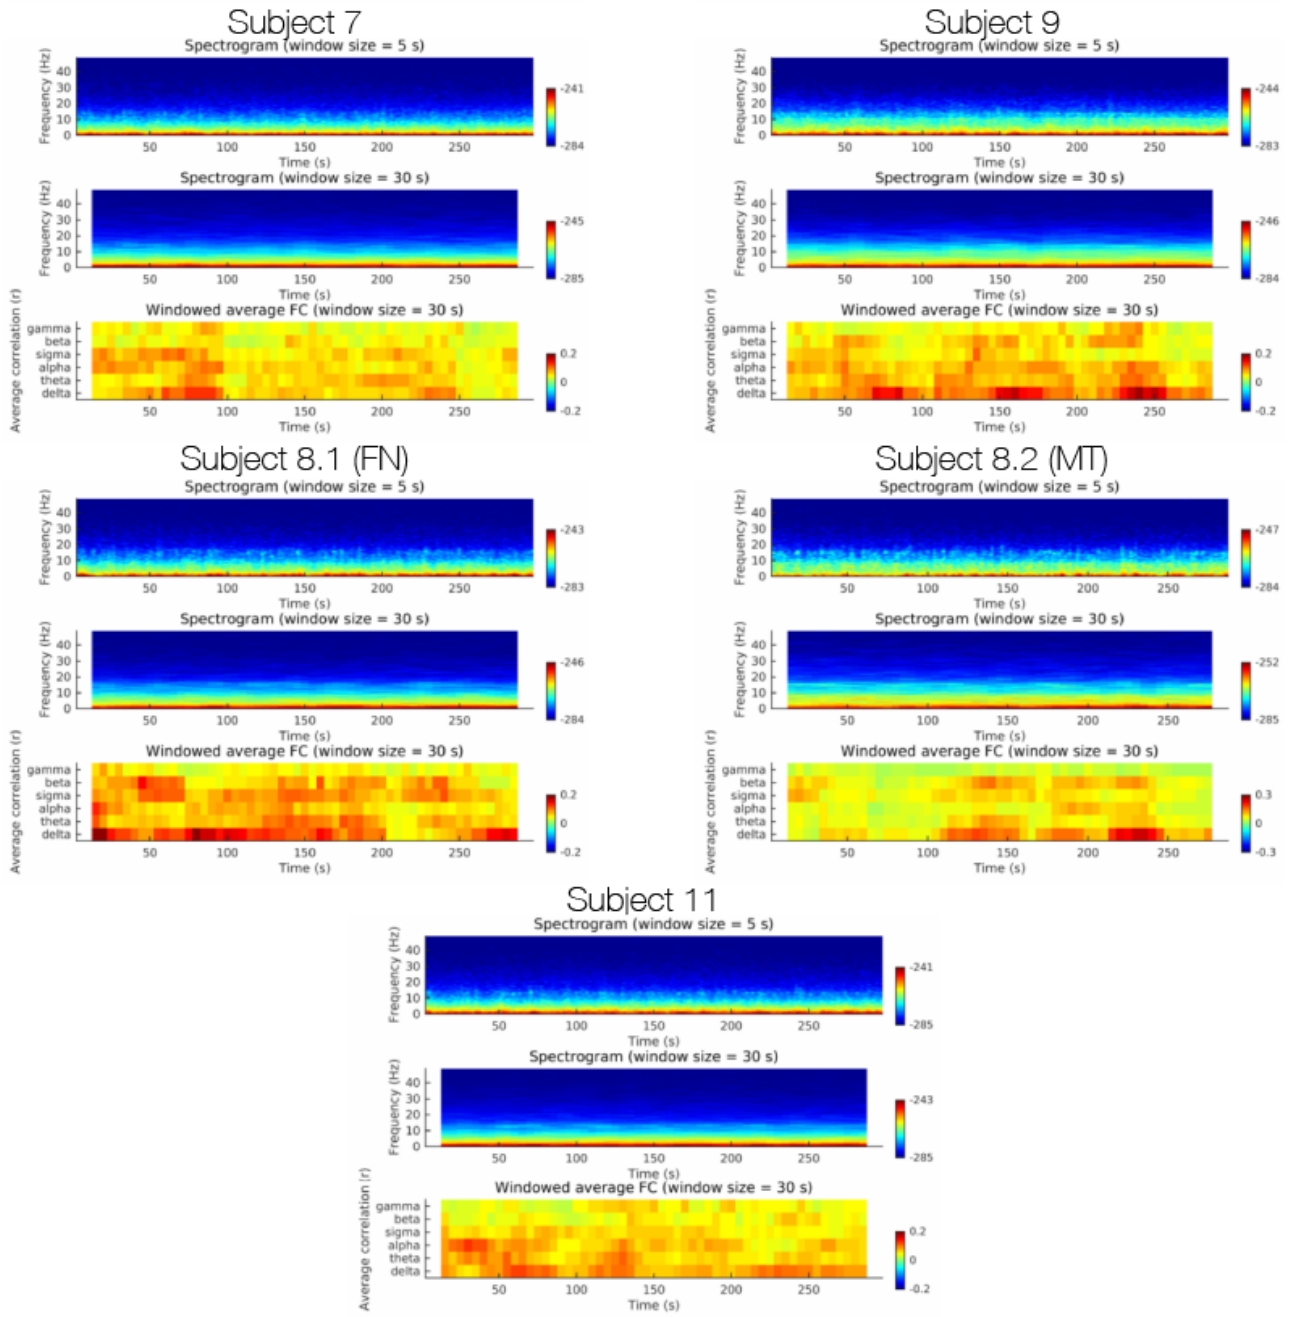

Figure S 7 ***Time-frequency representations of relationship between sensor power and mean level of FC in SWS from participant 5 to 8 (***Supplement to **Figure 5-2*).*** Spectrograms were computed from the broadband MEG sensor timecourses using the short-time Fourier transform in MATLAB, one with a window size of 5 seconds and 4 seconds overlap, and one with a window size of 30 seconds and 25 seconds overlap. The latter corresponded to the window size and overlap of the sliding window mean FC computation. 'FN' and 'MT' denotes the two different recording days. No clear associations between sensor power and mean FC are visible.

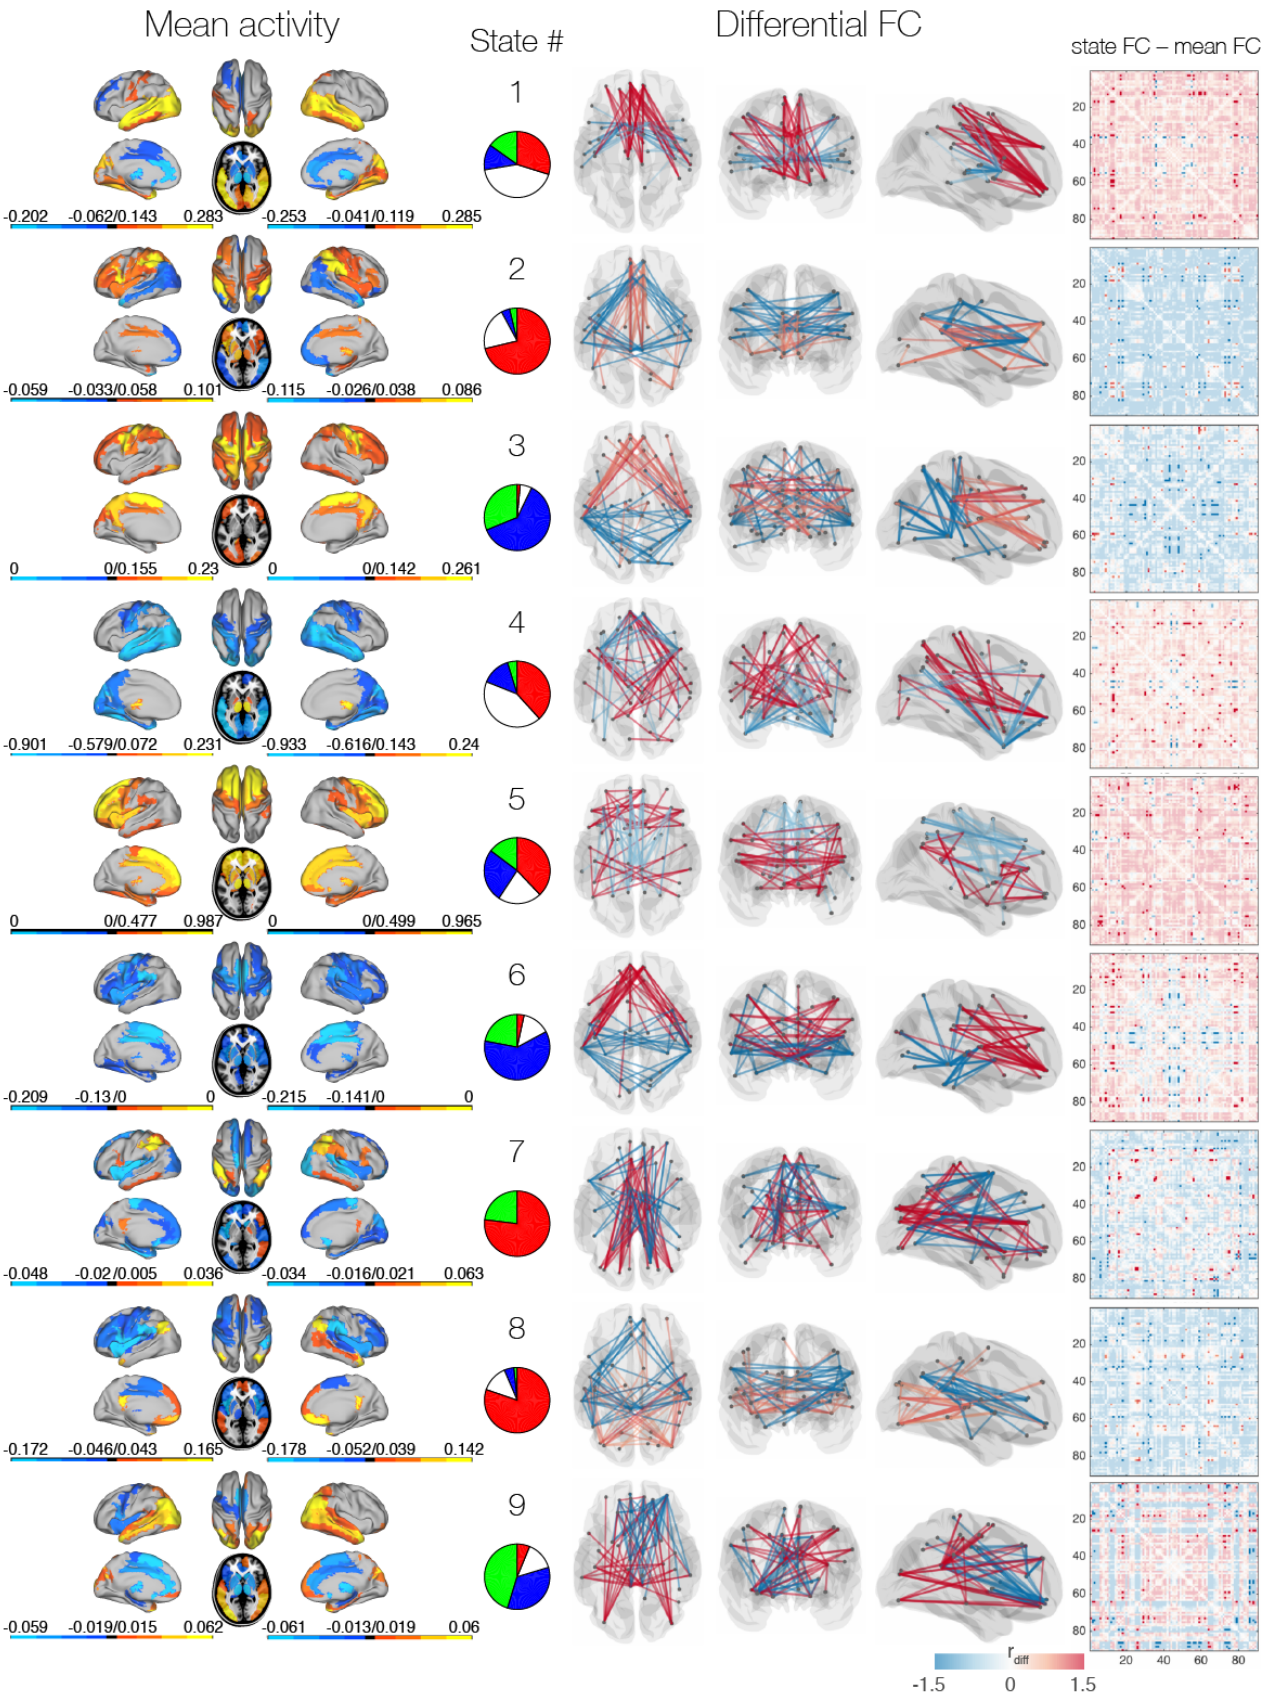

Figure S 8 ***Mean activity maps and differential FC networks of HMM states 1 to 9 (***Supplement to **Figure 6-8** and **Figure 6-9*).** *Left*: Mean activity maps were thresholded relative to each state, such that the 50% most negative and 50 percent most positive values are shown with red/yellow and blue/light blue respectively. *Middle*: Pie plots show the likelihood of finding each state within each of the four vigilance conditions. *Right*: Corresponding differential FC networks (FC_{diff}), computed by taking the state-dependent FC matrices and subtracting the average of the FC matrices of all other states. The brain plots and matrices were thresholded relative to each state, such that the 1% most positive and 1% most negative connections are shown in red and blue respectively.

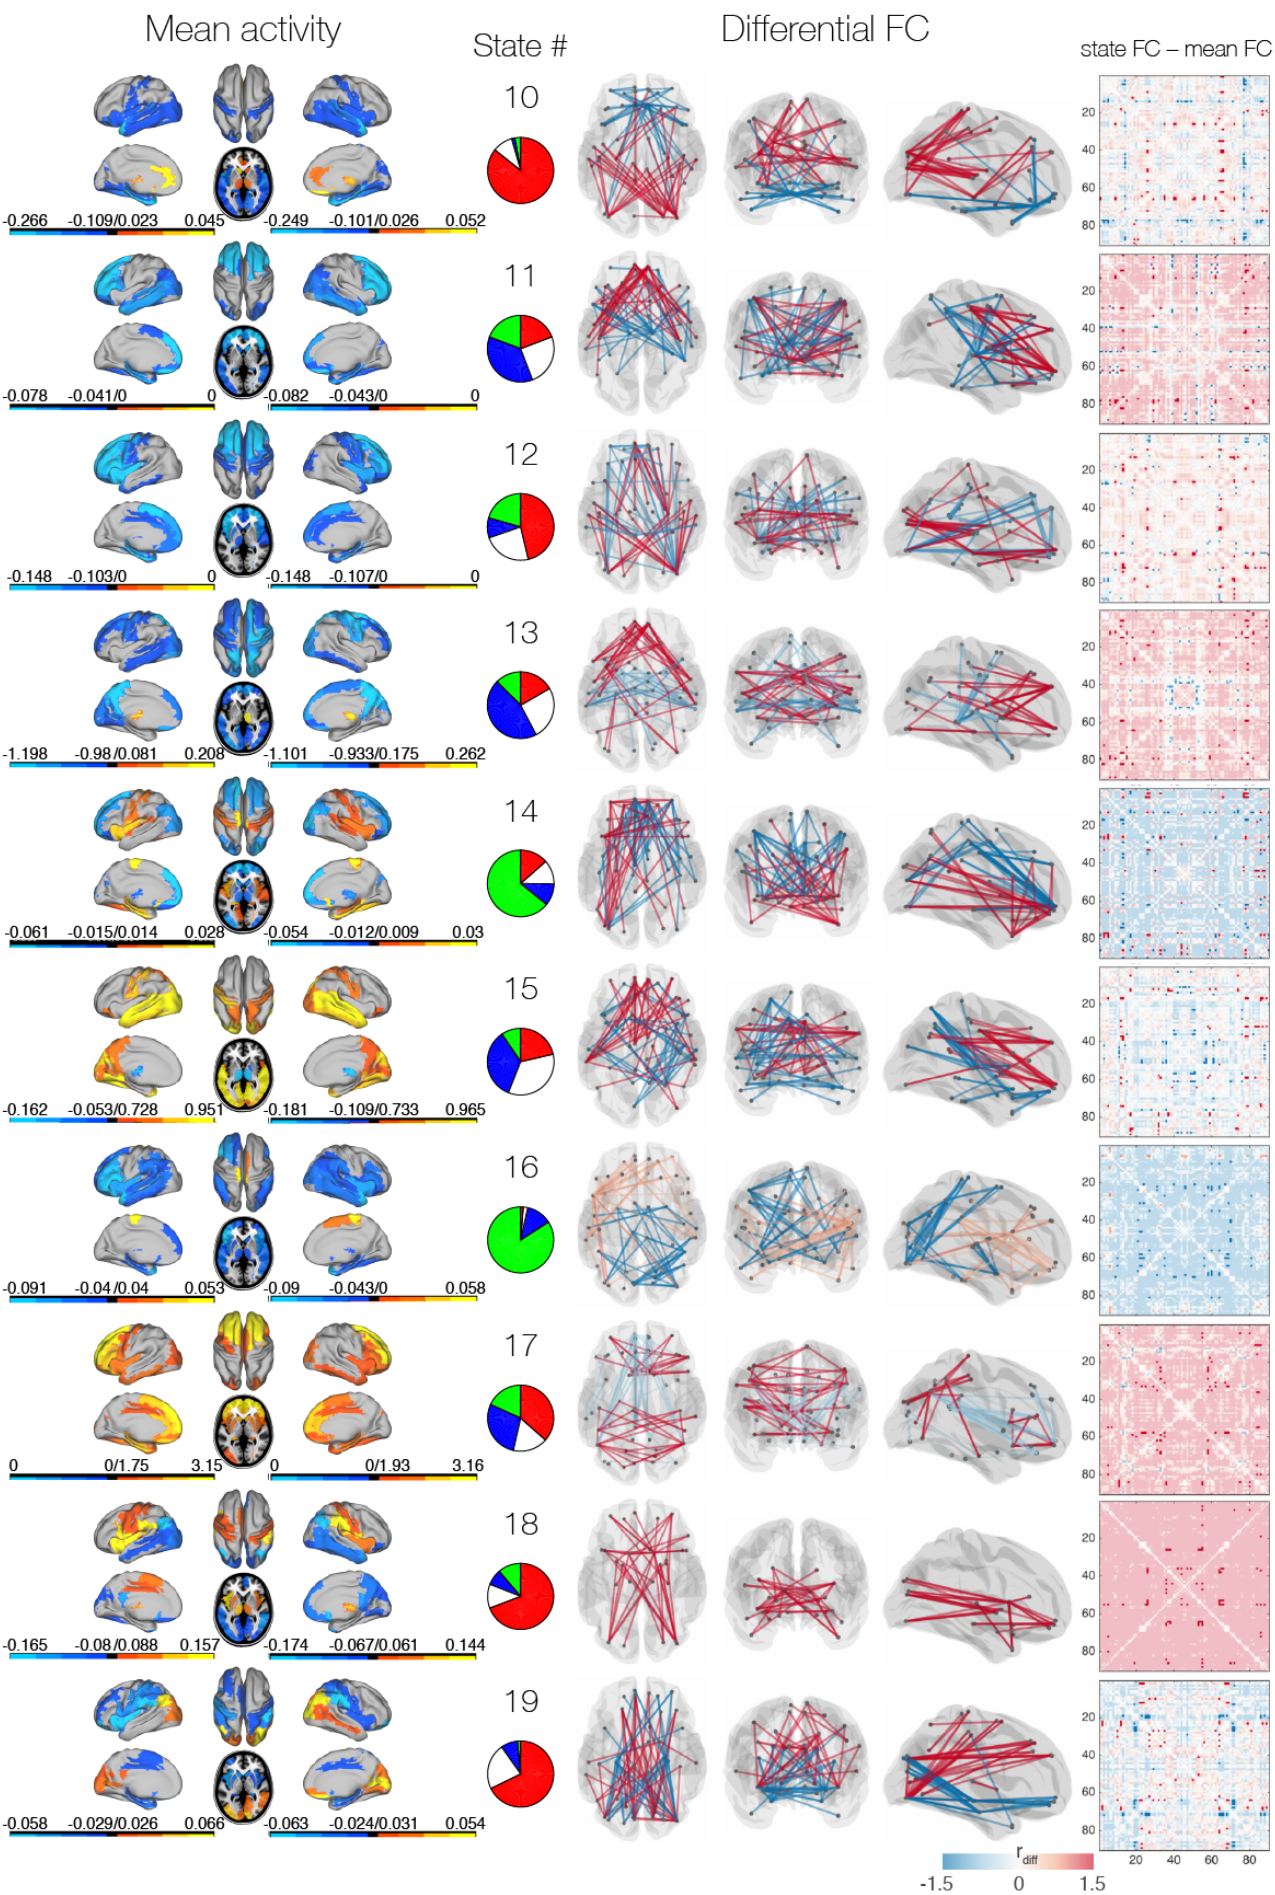
Mean activity
State #
Differential FC
state FC – mean FC
10
-0.266 -0.109/0.023 0.045
-0.249 -0.101/0.026 0.052
11
-0.078 -0.041/0 0
-0.082 -0.043/0 0
12
-0.148 -0.103/0 0
-0.148 -0.107/0 0
13
-1.198 -0.98/0.081 0.208
-1.101 -0.933/0.175 0.262
14
-0.061 -0.015/0.014 0.028
-0.054 -0.012/0.009 0.03
15
-0.162 -0.053/0.728 0.951
-0.181 -0.109/0.733 0.965
16
-0.091 -0.04/0.04 0.053
-0.09 -0.043/0 0.058
17
0 0/1.75 3.15
0 0/1.93 3.16
18
-0.165 -0.08/0.088 0.157
-0.174 -0.067/0.061 0.144
19
-0.058 -0.029/0.026 0.066
-0.063 -0.024/0.031 0.054
20 40 60 80
r diff
-1.5 0 1.5

Figure S 9 ***Mean activity maps and differential FC networks of HMM states 10 to 19 (***Supplement to **Figure 6-8** and **Figure 6-9*****).*** *Left*: Mean activity maps were thresholded relative to each state, such that the 50% most negative and 50 percent most positive values are shown with red/yellow and blue/light blue respectively. *Middle*: Pie plots show the likelihood of finding each state within each of the four vigilance conditions. *Right*: Corresponding differential FC networks (FC_{diff}), computed by taking the state-dependent FC matrices and subtracting the average of the FC matrices of all other states. The brain plots and matrices were thresholded relative to each state, such that the 1% most positive and 1% most negative connections are shown in red and blue respectively.

References

Abeysuriya RG, Rennie CJ, Robinson PA (2015) Physiologically based arousal state estimation and dynamics. Journal of Neuroscience Methods 253:55-69.

Abeysuriya RG, Robinson PA (2016) Real-time automated EEG tracking of brain states using neural field theory. Journal of Neuroscience Methods 258:28-45.

Achermann P, Borbély AA (1997) Low-frequency (< 1 Hz) oscillations in the human sleep electroencephalogram. Neuroscience 81:213-222.

Achermann P, Borbély AA (1998) Coherence analysis of the human sleep electroencephalogram. Neuroscience 85:1195-1208.

Achermann P, Tarokh L (2014) Human sleep and its regulation. Kosmos 63:173-180.

Aeschbach D, Dijk DJ, Borbely AA (1997) Dynamics of EEG spindle frequency activity during extended sleep in humans: relationship to slow-wave activity and time of day. Brain research 748:131-136.

Allada R, Siegel JM (2008) Unearthing the phylogenetic roots of sleep. Current biology : CB 18:R670-R679.

Allen EA, Damaraju E, Plis SM, Erhardt EB, Eichele T, Calhoun VD (2014) Tracking whole-brain connectivity dynamics in the resting state. Cerebral cortex (New York, NY : 1991) 24:663-676.

Allen PJ, Polizzi G, Krakow K, Fish DR, Lemieux L (1998) Identification of EEG events in the MR scanner: the problem of pulse artifact and a method for its subtraction. NeuroImage 8:229-239.

Altena E, Vrenken H, Van Der Werf YD, van den Heuvel OA, Van Someren EJ (2010) Reduced orbitofrontal and parietal gray matter in chronic insomnia: a voxel-based morphometric study. Biological psychiatry 67:182-185.

Altmann A, Schröter MS, Spoormaker VI, Kiem SA, Jordan D, Ilg R, Bullmore ET, Greicius MD, Czisch M, Sämann PG (2016) Validation of non-REM sleep stage decoding from resting state fMRI using linear support vector machines. NeuroImage 125:544-555.

Amiri M, Frauscher B, Gotman J (2016) Phase-amplitude coupling is elevated in deep sleep and in the onset zone of focal epileptic seizures. Frontiers in human neuroscience 10.

Andrade KC, Spoormaker VI, Dresler M, Wehrle R, Holsboer F, Sämann PG, Czisch M (2011) Sleep Spindles and Hippocampal Functional Connectivity in Human NREM Sleep. The Journal of Neuroscience 31:10331-10339.

Andrews-Hanna JR, Reidler JS, Huang C, Buckner RL (2010) Evidence for the default network's role in spontaneous cognition. Journal of neurophysiology 104:322-335.

Andrillon T, Nir Y, Staba RJ, Ferrarelli F, Cirelli C, Tononi G, Fried I (2011) Sleep spindles in humans: insights from intracranial EEG and unit recordings. The Journal of neuroscience : the official journal of the Society for Neuroscience 31:17821-17834.

Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. Science 273:1868-1871.

Attal Y, Maess B, Friederici A, David O (2012) Head models and dynamic causal modeling of subcortical activity using magnetoencephalographic/electroencephalographic data. Reviews in the Neurosciences 23:85-95.

Baars BJ (1993) A cognitive theory of consciousness. Cambridge University Press.

Baillet S, Mosher JC, Leahy RM (2001) Electromagnetic brain mapping. IEEE Signal Processing Magazine 18:14-30.

Baker AP, Brookes MJ, Rezek IA, Smith SM, Behrens T, Smith PJP, Woolrich M (2014a) Fast transient networks in spontaneous human brain activity. eLife 3:e01867.

Baker R, Gent TC, Yang Q, Parker S, Vyssotski AL, Wisden W, Brickley SG, Franks NP (2014b) Altered activity in the central medial thalamus precedes changes in the neocortex during

transitions into both sleep and propofol anesthesia. The Journal of neuroscience : the official journal of the Society for Neuroscience 34:13326-13335.
Bandettini PA (2016) Functional Imaging: Magnetic Resonance Imaging. In: Neuroscience in the 21st Century (Pfaff, D. W. and Volkow, N. D., eds), pp 1-28 New York, NY: Springer New York.
Banks S, Dinges DF (2007) Behavioral and physiological consequences of sleep restriction. Journal of clinical sleep medicine : JCSM : official publication of the American Academy of Sleep Medicine 3:519-528.
Bareham CA, Manly T, Pustovaya OV, Scott SK, Bekinschtein TA (2014) Losing the left side of the world: Rightward shift in human spatial attention with sleep onset. Scientific Reports 4.
Barttfeld P, Uhrig L, Sitt JD, Sigman M, Jarraya B, Dehaene S (2015) Signature of consciousness in the dynamics of resting-state brain activity. Proceedings of the National Academy of Sciences of the United States of America 112:887-892.
Bassett DS, Bullmore E (2006) Small-world brain networks. The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry 12:512-523.
Bassett DS, Bullmore ET (2009) Human brain networks in health and disease. Current opinion in neurology 22:340-347.
Bassett DS, Wymbs NF, Porter MA, Mucha PJ, Carlson JM, Grafton ST (2011) Dynamic reconfiguration of human brain networks during learning. Proceedings of the National Academy of Sciences 108:7641-7646.
Beckmann CF, DeLuca M, Devlin JT, Smith SM (2005) Investigations into resting-state connectivity using independent component analysis. Philosophical transactions of the Royal Society of London Series B, Biological sciences 360:1001-1013.
Beckmann CF, Smith SM (2004) Probabilistic independent component analysis for functional magnetic resonance imaging. IEEE transactions on medical imaging 23:137-152.
Beggs JM (2008) The criticality hypothesis: how local cortical networks might optimize information processing. Philosophical transactions Series A, Mathematical, physical, and engineering sciences 366:329-343.
Beggs JM, Timme N (2012) Being critical of criticality in the brain. Frontiers in physiology 3:163.
Behrens TE, Berg HJ, Jbabdi S, Rushworth MF, Woolrich MW (2007) Probabilistic diffusion tractography with multiple fibre orientations: What can we gain? NeuroImage 34:144-155.
Behrens TE, Woolrich MW, Jenkinson M, Johansen-Berg H, Nunes RG, Clare S, Matthews PM, Brady JM, Smith SM (2003) Characterization and propagation of uncertainty in diffusion-weighted MR imaging. Magnetic resonance in medicine : official journal of the Society of Magnetic Resonance in Medicine / Society of Magnetic Resonance in Medicine 50:1077-1088.
Behrens TEJ, Sporns O (2012) Human connectomics. Current opinion in neurobiology 22:144-153.
Berger H (1929) Über das elektrenkephalogramm des menschen. European Archives of Psychiatry and Clinical Neuroscience 87:527-570.
Biswal B, Yetkin FZ, Haughton VM, Hyde JS (1995) Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. Magnetic resonance in medicine 34:537-541.
Biswal BB, Mennes M, Zuo X-N, Gohel S, Kelly C, Smith SM, Beckmann CF, Adelstein JS, Buckner RL, Colcombe S, Dogonowski A-M, Ernst M, Fair D, Hampson M, Hoptman MJ, Hyde JS, Kiviniemi VJ, Kötter R, Li S-J, Lin C-P, Lowe MJ, Mackay C, Madden DJ, Madsen KH, Margulies DS, Mayberg HS, McMahon K, Monk CS, Mostofsky SH, Nagel BJ, Pekar JJ, Peltier SJ, Petersen SE, Riedl V, Rombouts SARB, Rypma B, Schlaggar BL, Schmidt S, Seidler RD, Siegle GJ, Sorg C, Teng G-J, Veijola J, Villringer A, Walter M, Wang L, Weng X-C, Whitfield-Gabrieli S, Williamson P, Windischberger C, Zang Y-F, Zhang H-Y, Castellanos FX, Milham MP (2010) Toward discovery science of human brain function. Proceedings of the National Academy of Sciences 107:4734-4739.
Boly M, Massimini M, Garrido MI, Gosseries O, Noirhomme Q, Laureys S, Soddu A (2012a) Brain connectivity in disorders of consciousness. Brain connectivity 2:1-10.

Boly M, Perlbarg V, Marrelec G, Schabus M, Laureys S, Doyon J, Pelegrini-Issac M, Maquet P, Benali H (2012b) Hierarchical clustering of brain activity during human nonrapid eye movement sleep. Proceedings of the National Academy of Sciences of the United States of America 109:5856-5861.

Boly M, Phillips C, Tshibanda L, Vanhaudenhuyse A, Schabus M, Dang-Vu TT, Moonen G, Hustinx R, Maquet P, Laureys S (2008) Intrinsic brain activity in altered states of consciousness: how conscious is the default mode of brain function? Annals of the New York Academy of Sciences 1129:119-129.

Bonhomme V, Boveroux P, Brichant JF, Laureys S, Boly M (2012) Neural correlates of consciousness during general anesthesia using functional magnetic resonance imaging (fMRI). Archives italiennes de biologie 150:155-163.

Bonjean M, Baker T, Bazhenov M, Cash S, Halgren E, Sejnowski T (2012) Interactions between core and matrix thalamocortical projections in human sleep spindle synchronization. The Journal of Neuroscience 32:5250-5263.

Borbely AA (1982) A two process model of sleep regulation. Human neurobiology 1:195-204.

Borbely AA, Daan S, Wirz-Justice A, Deboer T (2016) The two-process model of sleep regulation: a reappraisal. Journal of sleep research 25:131-143.

Born AP, Law I, Lund TE, Rostrup E, Hanson LG, Wildschiodtz G, Lou HC, Paulson OB (2002) Cortical deactivation induced by visual stimulation in human slow-wave sleep. NeuroImage 17:1325-1335.

Bosch OG, Rihm JS, Scheidegger M, Landolt H-P, Stämpfli P, Brakowski J, Esposito F, Rasch B, Seifritz E (2013) Sleep deprivation increases dorsal nexus connectivity to the dorsolateral prefrontal cortex in humans. Proceedings of the National Academy of Sciences of the United States of America 110:19597-19602.

Boveroux P, Vanhaudenhuyse A, Bruno MA, Noirhomme Q, Lauwick S, Luxen A, Degueldre C, Plenevaux A, Schnakers C, Phillips C, Brichant JF, Bonhomme V, Maquet P, Greicius MD, Laureys S, Boly M (2010) Breakdown of within- and between-network resting state functional magnetic resonance imaging connectivity during propofol-induced loss of consciousness. Anesthesiology 113:1038-1053.

Breakspear M (2017) Dynamic models of large-scale brain activity. Nature neuroscience 20:340-352.

Bressler SL, Menon V (2010) Large-scale brain networks in cognition: emerging methods and principles. Trends in cognitive sciences 14:277-290.

Britz J, Ville DVD, Michel CM (2010) BOLD correlates of EEG topography reveal rapid resting-state network dynamics. NeuroImage 52:1162-1170.

Brodbeck V, Kuhn A, Wegner Fv, Morzelewski A, Tagliazucchi E, Borisov S, Michel CM, Laufs H (2012) EEG microstates of wakefulness and NREM sleep. NeuroImage 62:2129-2139.

Brookes MJ, Hale JR, Zumer JM, Stevenson CM, Francis ST, Barnes GR, Owen JP, Morris PG, Nagarajan SS (2011a) Measuring functional connectivity using MEG: methodology and comparison with fcMRI. NeuroImage 56:1082-1104.

Brookes MJ, Tewarie PK, Hunt BAE, Robson SE, Gascoyne LE, Liddle EB, Liddle PF, Morris PG (2016) A multi-layer network approach to MEG connectivity analysis. NeuroImage 132:425-438.

Brookes MJ, Woolrich M, Luckhoo H, Price D, Hale JR, Stephenson MC, Barnes GR, Smith SM, Morris PG (2011b) Investigating the electrophysiological basis of resting state networks using magnetoencephalography. Proceedings of the National Academy of Sciences of the United States of America 108:16783-16788.

Brookes MJ, Woolrich M, Luckhoo H, Price D, Hale JR, Stephenson MC, Barnes GR, Smith SM, Morris PG (2011c) Investigating the electrophysiological basis of resting state networks using magnetoencephalography. Proceedings of the National Academy of Sciences 108:16783-16788.

Brookes MJ, Woolrich MW, Barnes GR (2012) Measuring functional connectivity in MEG: a multivariate approach insensitive to linear source leakage. Neuroimage 63:910-920.
Brown EN, Lydic R, Schiff ND (2010) General Anesthesia, Sleep, and Coma. The New England journal of medicine 363:2638-2650.
Brown RE, Basheer R, McKenna JT, Strecker RE, McCarley RW (2012) Control of sleep and wakefulness. Physiological reviews 92:1087-1187.
Buckner RL, Andrews-Hanna JR, Schacter DL (2008) The brain's default network. Annals of the New York Academy of Sciences 1124:1-38.
Bullmore E, Sporns O (2009) Complex brain networks: graph theoretical analysis of structural and functional systems. Nature reviews Neuroscience 10:186-198.
Bullmore E, Sporns O (2012) The economy of brain network organization. Nature reviews Neuroscience 13:336-349.
Bullmore ET, Suckling J, Overmeyer S, Rabe-Hesketh S, Taylor E, Brammer MJ (1999) Global, voxel, and cluster tests, by theory and permutation, for a difference between two groups of structural MR images of the brain. IEEE transactions on medical imaging 18:32-42.
Buzsaki G (2006) Rhythms of the brain: Oxford University Press.
Buzsáki G, Wang X-J (2012) Mechanisms of gamma oscillations. Annual review of neuroscience 35:203-225.
Cabral J, Fernandes HM, Van Hartevelt TJ, James AC, Kringelbach ML, Deco G (2013) Structural connectivity in schizophrenia and its impact on the dynamics of spontaneous functional networks. Chaos: An Interdisciplinary Journal of Nonlinear Science 23:046111.
Cabral J, Kringelbach ML, Deco G (2014a) Exploring the network dynamics underlying brain activity during rest. Progress in neurobiology 114:102-131.
Cabral J, Kringelbach ML, Deco G (2017) Functional connectivity dynamically evolves on multiple time-scales over a static structural connectome: Models and mechanisms. NeuroImage.
Cabral J, Luckhoo H, Woolrich M, Joensson M, Mohseni H, Baker A, Kringelbach ML, Deco G (2014b) Exploring mechanisms of spontaneous functional connectivity in MEG: how delayed network interactions lead to structured amplitude envelopes of band-pass filtered oscillations. NeuroImage 90:423-435.
Cabral J, Luckhoo H, Woolrich M, Joensson M, Mohseni H, Baker A, Kringelbach ML, Deco G (2014c) Exploring mechanisms of spontaneous functional connectivity in MEG: How delayed network interactions lead to structured amplitude envelopes of band-pass filtered oscillations. Neuroimage 90:423-435.
Calhoun Vince D, Miller R, Pearlson G, Adalı T (2014) The Chronnectome: Time-Varying Connectivity Networks as the Next Frontier in fMRI Data Discovery. Neuron 84:262-274.
Campbell SS, Tobler I (1984) Animal sleep: a review of sleep duration across phylogeny. Neuroscience and biobehavioral reviews 8:269-300.
Canolty RT, Edwards E, Dalal SS, Soltani M, Nagarajan SS, Kirsch HE, Berger MS, Barbaro NM, Knight RT (2006) High gamma power is phase-locked to theta oscillations in human neocortex. Science 313:1626-1628.
Canolty RT, Knight RT (2010) The functional role of cross-frequency coupling. Trends in cognitive sciences 14:506-515.
Caporro M, Haneef Z, Yeh HJ, Lenartowicz A, Buttinelli C, Parvizi J, Stern JM Functional MRI of sleep spindles and K-complexes. Clinical Neurophysiology 123:303-309.
Caporro M, Haneef Z, Yeh HJ, Lenartowicz A, Buttinelli C, Parvizi J, Stern JM (2012) Functional MRI of sleep spindles and K-complexes. Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology 123:303-309.
Carskadon M, Dement W (2005) Normal human sleep: an overview. In: Principles and Practice of Sleep Medicine (Kryger, M. et al., eds), pp 13-23: Elsevier Saunders.
Casali AG, Gosseries O, Rosanova M, Boly M, Sarasso S, Casali KR, Casarotto S, Bruno M-A, Laureys S, Tononi G, Massimini M (2013) A theoretically based index of consciousness

independent of sensory processing and behavior. Science Translational Medicine 5:198ra105-198ra105.
Cash SS, Halgren E, Dehghani N, Rossetti AO, Thesen T, Wang C, Devinsky O, Kuzniecky R, Doyle W, Madsen JR, Bromfield E, Eross L, Halasz P, Karmos G, Csercsa R, Wittner L, Ulbert I (2009) The human K-complex represents an isolated cortical down-state. Science 324:1084-1087.
Chang C, Glover GH (2010) Time-frequency dynamics of resting-state brain connectivity measured with fMRI. NeuroImage 50:81-98.
Chang C, Liu Z, Chen MC, Liu X, Duyn JH (2013) EEG correlates of time-varying BOLD functional connectivity. NeuroImage 72:227-236.
Chialvo DR (2010) Emergent complex neural dynamics. Nature Physics 6:744-750.
Chow HM, Horovitz SG, Carr WS, Picchioni D, Coddington N, Fukunaga M, Xu Y, Balkin TJ, Duyn JH, Braun AR (2013) Rhythmic alternating patterns of brain activity distinguish rapid eye movement sleep from other states of consciousness. Proceedings of the National Academy of Sciences of the United States of America 110:10300-10305.
Cirelli C, Tononi G (2008) Is sleep essential? PLoS biology 6:e216.
Clemens Z, Mölle M, Erőss L, Barsi P, Halász P, Born J (2007) Temporal coupling of parahippocampal ripples, sleep spindles and slow oscillations in humans. Brain : a journal of neurology 130:2868-2878.
Colclough GL, Brookes MJ, Smith SM, Woolrich MW (2015) A symmetric multivariate leakage correction for MEG connectomes. NeuroImage 117:439-448.
Colclough GL, Woolrich MW, Tewarie PK, Brookes MJ, Quinn AJ, Smith SM (2016) How reliable are MEG resting-state connectivity metrics? NeuroImage 138:284-293.
Cole DM, Smith SM, Beckmann CF (2010) Advances and pitfalls in the analysis and interpretation of resting-state FMRI data. Frontiers in systems neuroscience 4:8.
Collins D, Neelin P, Peters T, Evans AC (1994) Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space. Journal of Computer Assisted Tomography 18:192-205.
Colrain IM (2005) The K-complex: a 7-decade history. Sleep 28:255-273.
Colrain IM, Campbell KB (2007) The use of evoked potentials in sleep research. Sleep medicine reviews 11:277-293.
Contreras D, Steriade M (1995) Cellular basis of EEG slow rhythms: a study of dynamic corticothalamic relationships. The Journal of Neuroscience 15:604-622.
Cordes D, Haughton VM, Arfanakis K, Carew JD, Turski PA, Moritz CH, Quigley MA, Meyerand ME (2001) Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data. AJNR American journal of neuroradiology 22:1326-1333.
Cox R, Driel Jv, Boer Md, Talamini LM (2014) Slow oscillations during sleep coordinate interregional communication in cortical networks. The Journal of neuroscience : the official journal of the Society for Neuroscience 34:16890-16901.
Cox R, Hofman WF, Talamini LM (2012) Involvement of spindles in memory consolidation is slow wave sleep-specific. Learning & Memory 19:264-267.
Craddock RC, Jbabdi S, Yan CG, Vogelstein JT, Castellanos FX, Di Martino A, Kelly C, Heberlein K, Colcombe S, Milham MP (2013) Imaging human connectomes at the macroscale. Nature methods 10:524-539.
Crunelli V, David F, Lőrincz ML, Hughes SW (2015) The thalamocortical network as a single slow wave-generating unit. Current opinion in neurobiology 31:72-80.
Czisch M, Wetter TC, Kaufmann C, Pollmacher T, Holsboer F, Auer DP (2002) Altered processing of acoustic stimuli during sleep: reduced auditory activation and visual deactivation detected by a combined fMRI/EEG study. NeuroImage 16:251-258.
Damoiseaux JS, Rombouts SA, Barkhof F, Scheltens P, Stam CJ, Smith SM, Beckmann CF (2006) Consistent resting-state networks across healthy subjects. Proceedings of the National Academy of Sciences of the United States of America 103:13848-13853.

Dang-Vu TT (2012) Neuroimaging of phasic and non-phasic NREM activities. In: Neuroimaging of Sleep and Sleep Disorders: (Nofzinger, E. et al., eds) Cambridge: Cambridge University Press.

Dang-Vu TT, Bonjean M, Schabus M, Boly M, Darsaud A, Desseilles M, Degueldre C, Balteau E, Phillips C, Luxen A, Sejnowski TJ, Maquet P (2011) Interplay between spontaneous and induced brain activity during human non-rapid eye movement sleep. Proceedings of the National Academy of Sciences of the United States of America 108:15438-15443.

Dang-Vu TT, Schabus M, Desseilles M, Albouy G, Boly M, Darsaud A, Gais S, Rauchs G, Sterpenich V, Vandewalle G, Carrier J, Moonen G, Balteau E, Degueldre C, Luxen A, Phillips C, Maquet P (2008a) Spontaneous neural activity during human slow wave sleep. Proceedings of the National Academy of Sciences of the United States of America 105:15160-15165.

Dang-Vu TT, Schabus M, Desseilles M, Albouy G, Boly M, Darsaud A, Gais S, Rauchs G, Sterpenich V, Vandewalle G, Carrier J, Moonen G, Balteau E, Degueldre C, Luxen A, Phillips C, Maquet P (2008b) Spontaneous neural activity during human slow wave sleep. Proceedings of the National Academy of Sciences of the United States of America 105:15160-15165.

Dang-Vu TT, Schabus M, Desseilles M, Sterpenich V, Bonjean M, Maquet P (2010a) Functional Neuroimaging Insights into the Physiology of Human Sleep. Sleep 33:1589-1603.

Dang-Vu TT, Schabus M, Desseilles M, Sterpenich V, Bonjean M, Maquet P (2010b) Functional neuroimaging insights into the physiology of human sleep. Sleep 33:1589-1603.

Datta S (2012) Chapter 7 - Phasic pontine-wave (P-wave) generation: cellular-molecular-network mechanism and functional significance A2 - Frank, Marcos G. In: Sleep and Brain Activity, pp 147-164 San Diego: Academic Press.

David F, Schmiedt JT (2014) Thalamus and cortex: inseparable partners in shaping sleep slow waves? The Journal of neuroscience : the official journal of the Society for Neuroscience 34:11517-11518.

Davis H, Davis PA, Loomis AL, Harvey EN, Hobart G (1937) Changes in human brain potentials during the onset of sleep. Science 86:448-450.

De Gennaro L, Ferrara M (2003) Sleep spindles: an overview. Sleep medicine reviews 7:423-440.

De Gennaro L, Ferrara M, Bertini M (2000) The spontaneous K-complex during stage 2 sleep: is it the 'forerunner' of delta waves? Neuroscience letters 291:41-43.

Deco G, Cabral J, Woolrich M, Stevner A, van Hartevelt T, Kringelbach M (2016a) Single or multi-frequency generators in on-going brain activity: a mechanistic whole-brain model following empirical MEG evidences. bioRxiv.

Deco G, Hagmann P, Hudetz AG, Tononi G (2014) Modeling resting-state functional networks when the cortex falls asleep: local and global changes. Cerebral cortex (New York, NY : 1991) 24:3180-3194.

Deco G, Hartevelt TJV, Fernandes HM, Stevner A, Kringelbach ML (2016b) The most relevant human brain regions for functional connectivity: Evidence for a dynamical workspace of binding nodes from whole-brain computational modelling. NeuroImage 146:197-210.

Deco G, Jirsa VK, McIntosh AR (2011) Emerging concepts for the dynamical organization of resting-state activity in the brain. Nature reviews Neuroscience 12:43-56.

Deco G, Kringelbach M (2016) Metastability and Coherence: Extending the Communication through Coherence Hypothesis Using a Whole-Brain Computational Perspective. Trends in neurosciences 39:432.

Deco G, Kringelbach ML (2014) Great expectations: using whole-brain computational connectomics for understanding neuropsychiatric disorders. Neuron 84:892-905.

Dehaene S, Changeux J-P (2011) Experimental and theoretical approaches to conscious processing. Neuron 70:200-227.

Dehaene S, Charles L, King J-R, Marti S (2014) Toward a computational theory of conscious processing. Current opinion in neurobiology 25:76-84.

Dehaene S, Sergent C, Changeux JP (2003) A neuronal network model linking subjective reports and objective physiological data during conscious perception. Proceedings of the National Academy of Sciences of the United States of America 100:8520-8525.
Dehghani N, Cash SS, Chen CC, Hagler DJ, Jr., Huang M, Dale AM, Halgren E (2010a) Divergent cortical generators of MEG and EEG during human sleep spindles suggested by distributed source modeling. PloS one 5:e11454.
Dehghani N, Cash SS, Chen CC, Hagler Jr DJ, Huang M, Dale AM, Halgren E (2010b) Divergent cortical generators of MEG and EEG during human sleep spindles suggested by distributed source modeling. PloS one 5:e11454.
Dehghani N, Cash SS, Halgren E (2011) Emergence of synchronous EEG spindles from asynchronous MEG spindles. Human brain mapping 32:2217-2227.
Dehghani N, Cash SS, Rossetti AO, Chen CC, Halgren E (2010c) Magnetoencephalography demonstrates multiple asynchronous generators during human sleep spindles. Journal of neurophysiology 104:179-188.
Dehghani N, Hatsopoulos NG, Haga ZD, Parker RA, Greger B, Halgren E, Cash SS, Destexhe A (2012) Avalanche analysis from multielectrode ensemble recordings in cat, monkey, and human cerebral cortex during wakefulness and sleep. Frontiers in physiology 3.
Diekelmann S, Born J (2010) The memory function of sleep. Nat Rev Neurosci 11:114-126.
Duckrow RB, Zaveri HP (2005) Coherence of the electroencephalogram during the first sleep cycle. Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology 116:1088-1095.
Duff E, Makin T, Madugula S, Smith SM, Woolrich MW (2013) Utility of Partial Correlation for Characterising Brain Dynamics: MVPA-based Assessment of Regularisation and Network Selection. In: Pattern Recognition in Neuroimaging (PRNI), 2013 International Workshop on, pp 58-61.
Engel AK, Fries P (2010) Beta-band oscillations—signalling the status quo? Current opinion in neurobiology 20:156-165.
Engel AK, Gerloff C, Hilgetag CC, Nolte G (2013) Intrinsic coupling modes: multiscale interactions in ongoing brain activity. Neuron 80:867-886.
Ferrara M, Curcio G, Fratello F, Moroni F, Marzano C, Pellicciari MC, Gennaro LD (2006) The electroencephalographic substratum of the awakening. Behavioural brain research 167:237-244.
Ferri R, Rundo F, Bruni O, Terzano MG, Stam CJ (2007) Small-world network organization of functional connectivity of EEG slow-wave activity during sleep. Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology 118:449-456.
Ferri R, Rundo F, Bruni O, Terzano MG, Stam CJ (2008) The functional connectivity of different EEG bands moves towards small-world network organization during sleep. Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology 119:2026-2036.
Finelli LA, Borbely AA, Achermann P (2001) Functional topography of the human nonREM sleep electroencephalogram. The European journal of neuroscience 13:2282-2290.
Florin E, Baillet S (2015) The brain's resting-state activity is shaped by synchronized cross-frequency coupling of oscillatory neural activity. NeuroImage 111.
Fornito A, Zalesky A, Breakspear M (2013) Graph analysis of the human connectome: promise, progress, and pitfalls. NeuroImage 80:426-444.
Fornito A, Zalesky A, Breakspear M (2015) The connectomics of brain disorders. Nature reviews Neuroscience 16:159-172.
Foster BL, He BJ, Honey CJ, Jerbi K, Maier A, Saalmann YB (2016) Spontaneous neural dynamics and multi-scale network organization. Frontiers in systems neuroscience 10.
Fox KCR, Spreng RN, Ellamil M, Andrews-Hanna JR, Christoff K (2015) The wandering brain: meta-analysis of functional neuroimaging studies of mind-wandering and related spontaneous thought processes. NeuroImage 111:611-621.

Fox MD, Corbetta M, Snyder AZ, Vincent JL, Raichle ME (2006) Spontaneous neuronal activity distinguishes human dorsal and ventral attention systems. Proceedings of the National Academy of Sciences of the United States of America 103:10046-10051.
Fox MD, Raichle ME (2007) Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. Nat Rev Neurosci 8:700-711.
Fox MD, Snyder AZ, Vincent JL, Corbetta M, Van Essen DC, Raichle ME (2005) The human brain is intrinsically organized into dynamic, anticorrelated functional networks. Proceedings of the National Academy of Sciences of the United States of America 102:9673-9678.
Fox MD, Zhang D, Snyder AZ, Raichle ME (2009) The global signal and observed anticorrelated resting state brain networks. Journal of neurophysiology 101:3270-3283.
Franken P (2002) Long-term vs. short-term processes regulating REM sleep. Journal of sleep research 11:17-28.
Friedman J, Hastie T, Tibshirani R (2008) Sparse inverse covariance estimation with the graphical lasso. Biostatistics 9:432-441.
Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. Trends in cognitive sciences 9:474-480.
Fries P (2009) Neuronal gamma-band synchronization as a fundamental process in cortical computation. Annual review of neuroscience 32:209-224.
Friston KJ (1994) Functional and effective connectivity in neuroimaging: a synthesis. Human brain mapping 2:56-78.
Friston KJ (2011) Functional and effective connectivity: a review. Brain connectivity 1:13-36.
Friston KJ, Harrison L, Penny W (2003) Dynamic causal modelling. NeuroImage 19:1273-1302.
Friston KJ, Kahan J, Biswal B, Razi A (2014) A DCM for resting state fMRI. NeuroImage 94:396-407.
Fukunaga M, Horovitz SG, van Gelderen P, de Zwart JA, Jansma JM, Ikonomidou VN, Chu R, Deckers RH, Leopold DA, Duyn JH (2006) Large-amplitude, spatially correlated fluctuations in BOLD fMRI signals during extended rest and early sleep stages. Magnetic resonance imaging 24:979-992.
Gilson M, Moreno-Bote R, Ponce-Alvarez A, Ritter P, Deco G (2016) Estimation of directed effective connectivity from fMRI functional connectivity hints at asymmetries of cortical connectome. PLoS computational biology 12.
Glasser MF, Coalson TS, Robinson EC, Hacker CD, Harwell J, Yacoub E, Ugurbil K, Andersson J, Beckmann CF, Jenkinson M, Smith SM, Van Essen DC (2016) A multi-modal parcellation of human cerebral cortex. Nature 536:171-178.
Glover GH, Li TQ, Ress D (2000) Image‐based method for retrospective correction of physiological motion effects in fMRI: RETROICOR. Magnetic resonance in medicine 44:162-167.
Goupil L, Bekinschtein TA (2012) Cognitive processing during the transition to sleep. Archives italiennes de biologie 150:140-154.
Greicius MD, Kiviniemi V, Tervonen O, Vainionpaa V, Alahuhta S, Reiss AL, Menon V (2008) Persistent default-mode network connectivity during light sedation. Human brain mapping 29:839-847.
Greicius MD, Supekar K, Menon V, Dougherty RF (2009) Resting-state functional connectivity reflects structural connectivity in the default mode network. Cerebral cortex (New York, NY : 1991) 19:72-78.
Guillery RW (2000) Brodmann's'Localisation in the Cerebral Cortex'. Journal of anatomy 196:493-496.
Gujar N, Yoo S-S, Hu P, Walker MP (2010) The unrested resting brain: sleep deprivation alters activity within the default-mode network. Journal of cognitive neuroscience 22:1637-1648.
Gulevich G, Dement W, Johnson L (1966) Psychiatric and eeg observations on a case of prolonged (264 hours) wakefulness. Archives of general psychiatry 15:29-35.

Hagmann P (2005) From diffusion MRI to brain connectomics. Institut de traitement des signaux Programme Doctoral en informatique et communications pour l'obtention du grade docteur es sciences par Docteur en médecine, Université de Lausanne.
Hagmann P, Kurant M, Gigandet X, Thiran P, Wedeen VJ, Meuli R, Thiran J-P (2007) Mapping Human Whole-Brain Structural Networks with Diffusion MRI. PloS one 2:e597.
Haken H (1975) Cooperative phenomena in systems far from thermal equilibrium and in nonphysical systems. Reviews of Modern Physics 47:67-121.
Halász P (2005) K-complex, a reactive EEG graphoelement of NREM sleep: an old chap in a new garment. Sleep medicine reviews 9:391-412.
Hale JR, White TP, Mayhew SD, Wilson RS, Rollings DT, Khalsa S, Arvanitis TN, Bagshaw AP (2016) Altered thalamocortical and intra-thalamic functional connectivity during light sleep compared with wake. NeuroImage 125:657-667.
Hall EL, Robson SE, Morris PG, Brookes MJ (2014) The relationship between MEG and fMRI. NeuroImage 102:80-91.
Hansen ECA, Battaglia D, Spiegler A, Deco G, Jirsa VK (2015) Functional connectivity dynamics: modeling the switching behavior of the resting state. NeuroImage 105:525-535.
Hansen PC, Kringelbach ML, Salmelin R (2010) MEG an introduction to methods. New York: Oxford University Press.
Hartevelt TJv, Cabral J, Deco G, Møller A, Green AL, Aziz TZ, Kringelbach ML (2014) Neural plasticity in human brain connectivity: the effects of long term deep brain stimulation of the subthalamic nucleus in Parkinson's disease. PloS one 9:e86496.
Hauri P, Hawkins DR (1973) Alpha-delta sleep. Electroencephalography and clinical neurophysiology 34:233-237.
Havas JAD, Parimal S, Soon CS, Chee MWL (2012) Sleep deprivation reduces default mode network connectivity and anti-correlation during rest and task performance. NeuroImage 59:1745-1751.
Hayasaka S, Nichols TE (2004) Combining voxel intensity and cluster extent with permutation test framework. Neuroimage 23:54-63.
He BJ, Raichle ME (2009) The fMRI signal, slow cortical potential and consciousness. Trends in cognitive sciences 13:302-309.
He BJ, Snyder AZ, Zempel JM, Smyth MD, Raichle ME (2008) Electrophysiological correlates of the brain's intrinsic large-scale functional architecture. Proceedings of the National Academy of Sciences of the United States of America 105:16039-16044.
Heine L, Soddu A, Gomez F, Vanhaudenhuyse A, Tshibanda L, Thonnard M, Charland-Verville V, Kirsch M, Laureys S, Demertzi A (2012) Resting state networks and consciousness: alterations of multiple resting state network connectivity in physiological, pharmacological, and pathological consciousness States. Frontiers in psychology 3:295.
Helmholtz H (1853) Ueber einige Gesetze der Vertheilung elektrischer Ströme in körperlichen Leitern, mit Anwendung auf die thierisch-elektrischen Versuche (Schluss.). Annalen der Physik 165:353-377.
Hillebrand A, Barnes GR (2005) Beamformer Analysis of MEG Data. In: International review of neurobiology, vol. Volume 68, pp 149-171: Academic Press.
Hillebrand A, Barnes GR, Bosboom JL, Berendse HW, Stam CJ (2012) Frequency-dependent functional connectivity within resting-state networks: an atlas-based MEG beamformer solution. NeuroImage 59:3909-3921.
Hillebrand A, Singh KD, Holliday IE, Furlong PL, Barnes GR (2005) A new approach to neuroimaging with magnetoencephalography. Human brain mapping 25:199-211.
Himanen S-L, Hasan J (2000) Limitations of Rechtschaffen and Kales. Sleep medicine reviews 4:149-167.
Hindriks R, Adhikari MH, Murayama Y, Ganzetti M, Mantini D, Logothetis NK, Deco G (2015) Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? NeuroImage 127:242-256.

Hipp JF, Hawellek DJ, Corbetta M, Siegel M, Engel AK (2012) Large-scale cortical correlation structure of spontaneous oscillatory activity. Nature neuroscience 15:10.1038/nn.3101.
Hipp Joerg F, Siegel M (2015) BOLD fMRI correlation reflects frequency-specific neuronal correlation. Current Biology 25:1368-1374.
Hobson JA (1989) Sleep / J. Allan Hobson. New York: Scientific American Library : Distributed by W.H. Freeman.
Hobson JA (2005) Sleep is of the brain, by the brain and for the brain. Nature 437:1254-1256.
Hobson JA, Pace-Schott EF (2002) The cognitive neuroscience of sleep: neuronal systems, consciousness and learning. Nature reviews Neuroscience 3:679-693.
Holliday IE, Barnes GR, Hillebrand A, Singh KD (2003) Accuracy and applications of group MEG studies using cortical source locations estimated from participants' scalp surfaces. Hum Brain Mapp 20:142-147.
Honey CJ, Sporns O, Cammoun L, Gigandet X, Thiran JP, Meuli R, Hagmann P (2009) Predicting human resting-state functional connectivity from structural connectivity. Proceedings of the National Academy of Sciences of the United States of America 106:2035-2040.
Hori T, Hayashi M, Morikawa T (1994) Topographical EEG changes and the hypnagogic experience.
Horovitz SG, Braun AR, Carr WS, Picchioni D, Balkin TJ, Fukunaga M, Duyn JH (2009) Decoupling of the brain's default mode network during deep sleep. Proceedings of the National Academy of Sciences of the United States of America 106:11376-11381.
Horovitz SG, Fukunaga M, de Zwart JA, van Gelderen P, Fulton SC, Balkin TJ, Duyn JH (2008) Low frequency BOLD fluctuations during resting wakefulness and light sleep: a simultaneous EEG-fMRI study. Hum Brain Mapp 29:671-682.
Huang MX, Mosher JC, Leahy RM (1999) A sensor-weighted overlapping-sphere head model and exhaustive head model comparison for MEG. Physics in medicine and biology 44:423-440.
Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, Della Penna S, Duyn JH, Glover GH, Gonzalez-Castillo J (2013a) Dynamic functional connectivity: promise, issues, and interpretations. NeuroImage 80:360-378.
Hutchison RM, Womelsdorf T, Allen EA, Bandettini PA, Calhoun VD, Corbetta M, Penna SD, Duyn JH, Glover GH, Gonzalez-Castillo J, Handwerker DA, Keilholz S, Kiviniemi V, Leopold DA, Pasquale Fd, Sporns O, Walter M, Chang C (2013b) Dynamic functional connectivity: promise, issues, and interpretations. NeuroImage 80:360-378.
Hyvarinen A, Oja E (2000) Independent component analysis: algorithms and applications. Neural networks : the official journal of the International Neural Network Society 13:411-430.
Iber C, Ancoli-Israel S, Chesson A, Quan SF (2007) The AASM manual for the scoring of sleep and associated events: rules, terminology and technical specifications. American Academy of Sleep Medicine Westchester, IL.
Jagannath A, Peirson SN, Foster RG (2013) Sleep and circadian rhythm disruption in neuropsychiatric illness. Current opinion in neurobiology 23:888-894.
Jahnke K, Wegner Fv, Morzelewski A, Borisov S, Maischein M, Steinmetz H, Laufs H (2012) To wake or not to wake? The two-sided nature of the human K-complex. NeuroImage 59:1631-1638.
Jbabdi S, Johansen-Berg H (2011) Tractography: where do we go from here? Brain connectivity 1:169-183.
Jenkinson M, Bannister P, Brady M, Smith S (2002) Improved optimization for the robust and accurate linear registration and motion correction of brain images. NeuroImage 17:825-841.
Jenkinson M, Beckmann CF, Behrens TE, Woolrich MW, Smith SM (2012) Fsl. Neuroimage 62:782-790.
Jensen O, Colgin LL (2007) Cross-frequency coupling between neuronal oscillations. Trends in cognitive sciences 11:267-269.
Jensen O, Kaiser J, Lachaux J-P (2007) Human gamma-frequency oscillations associated with attention and memory. Trends in neurosciences 30:317-324.

Johnson LA, Blakely T, Hermes D, Hakimian S, Ramsey NF, Ojemann JG (2012) Sleep spindles are locally modulated by training on a brain–computer interface. Proceedings of the National Academy of Sciences 109:18583-18588.
Joiner WJ (2016) Unraveling the evolutionary determinants of sleep. Current biology : CB 26:R1073-R1087.
Jones DK, Knösche TR, Turner R (2013) White matter integrity, fiber count, and other fallacies: The do's and don'ts of diffusion MRI. NeuroImage 73:239-254.
Joo EY, Noh HJ, Kim JS, Koo DL, Kim D, Hwang KJ, Kim JY, Kim ST, Kim MR, Hong SB (2013) Brain gray matter deficits in patients with chronic primary insomnia. Sleep 36:999-1007.
Karahanoğlu FI, Ville DVD (2015) Transient brain activity disentangles fMRI resting-state dynamics in terms of spatially and temporally overlapping networks. Nature communications 6:7751.
Kaufmann C, Wehrle R, Wetter TC, Holsboer F, Auer DP, Pollmächer T, Czisch M (2006) Brain activation and hypothalamic functional connectivity during human non-rapid eye movement sleep: an EEG/fMRI study. Brain : a journal of neurology 129:655-667.
Kaufmann T, Elvsåshagen T, Alnæs D, Zak N, Pedersen PØ, Norbom LB, Quraishi SH, Tagliazucchi E, Laufs H, Bjørnerud A, Malt UF, Andreassen OA, Roussos E, Duff EP, Smith SM, Groote IR, Westlye LT (2015) The brain functional connectome is robustly altered by lack of sleep. NeuroImage 127:324-332.
Kerkhofs M, Lavie P (2000) HISTORICAL NOTE: Frédéric Bremer 1892–1982: a pioneer in sleep research. Sleep medicine reviews 4:505-514.
Kiviniemi V, Kantola JH, Jauhiainen J, Hyvarinen A, Tervonen O (2003) Independent component analysis of nondeterministic fMRI signal sources. NeuroImage 19:253-260.
Kiviniemi V, Vire T, Remes J, Elseoud AA, Starck T, Tervonen O, Nikkinen J (2011) A sliding time-window ICA reveals spatial variability of the default mode network in time. Brain connectivity 1:339-347.
Kiviniemi VJ, Haanpaa H, Kantola JH, Jauhiainen J, Vainionpaa V, Alahuhta S, Tervonen O (2005) Midazolam sedation increases fluctuation and synchrony of the resting brain BOLD signal. Magnetic resonance imaging 23:531-537.
Klimesch W, Sauseng P, Hanslmayr S (2007) EEG alpha oscillations: The inhibition–timing hypothesis. Brain Research Reviews 53:63-88.
Klinzing JG, Mölle M, Weber F, Supp G, Hipp JF, Engel AK, Born J (2016) Spindle activity phase-locked to sleep slow oscillations. NeuroImage 134:607-616.
Koch C, Massimini M, Boly M, Tononi G (2016) Neural correlates of consciousness: progress and problems. Nature reviews Neuroscience 17:307-321.
Koechlin E, Summerfield C An information theoretical approach to prefrontal executive function. Trends in cognitive sciences 11:229-235.
Koenig T, Studer D, Hubl D, Melie L, Strik WK (2005) Brain connectivity at different time-scales measured with EEG. Philosophical Transactions of the Royal Society B: Biological Sciences 360:1015-1024.
Koike T, Kan S, Misaki M, Miyauchi S (2011) Connectivity pattern changes in default-mode network with deep non-REM and REM sleep. Neuroscience research 69:322-330.
Krueger JM, Frank MG, Wisor JP, Roy S (2016) Sleep function: Toward elucidating an enigma. Sleep medicine reviews 28:46-54.
Krueger JM, Rector DM, Roy S, Van Dongen HP, Belenky G, Panksepp J (2008) Sleep as a fundamental property of neuronal assemblies. Nat Rev Neurosci 9:910-919.
Krueger JM, Tononi G (2011) Local Use-Dependent Sleep; Synthesis of the New Paradigm. Current topics in medicinal chemistry 11:2490-2492.
Kurth S, Huber R (2012) Chapter 10 - Sleep slow oscillations and cortical maturation A2 - Frank, Marcos G. In: Sleep and Brain Activity, pp 227-261 San Diego: Academic Press.

Larson-Prior LJ, Ju YE, Galvin JE (2014) Cortical-subcortical interactions in hypersomnia disorders: mechanisms underlying cognitive and behavioral aspects of the sleep-wake cycle. Frontiers in neurology 5:165.

Larson-Prior LJ, Oostenveld R, Della Penna S, Michalareas G, Prior F, Babajani-Feremi A, Schoffelen JM, Marzetti L, de Pasquale F, Di Pompeo F, Stout J, Woolrich M, Luo Q, Bucholz R, Fries P, Pizzella V, Romani GL, Corbetta M, Snyder AZ, Consortium WU-MH (2013) Adding dynamics to the Human Connectome Project with MEG. Neuroimage 80:190-201.

Larson-Prior LJ, Power JD, Vincent JL, Nolan TS, Coalson RS, Zempel J, Snyder AZ, Schlaggar BL, Raichle ME, Petersen SE (2011) Modulation of the brain's functional network architecture in the transition from wake to sleep. Progress in brain research 193:277-294.

Larson-Prior LJ, Zempel JM, Nolan TS, Prior FW, Snyder AZ, Raichle ME (2009) Cortical network functional connectivity in the descent to sleep. Proceedings of the National Academy of Sciences of the United States of America 106:4489-4494.

Laufs H (2008) Endogenous brain oscillations and related networks detected by surface EEG-combined fMRI. Human brain mapping 29:762-769.

Laufs H, Walker MC, Lund TE (2007) 'Brain activation and hypothalamic functional connectivity during human non-rapid eye movement sleep: an EEG/fMRI study'--its limitations and an alternative approach. Brain : a journal of neurology 130:e75; author reply e76.

Laumann Timothy O, Gordon Evan M, Adeyemo B, Snyder Abraham Z, Joo Sung J, Chen M-Y, Gilmore Adrian W, McDermott Kathleen B, Nelson Steven M, Dosenbach Nico UF, Schlaggar Bradley L, Mumford Jeanette A, Poldrack Russell A, Petersen Steven E (2015) Functional system and areal organization of a highly sampled Individual human brain. Neuron 87:657-670.

Le Bihan D, Johansen-Berg H (2012) Diffusion MRI at 25: Exploring brain tissue structure and function. NeuroImage 61:324-341.

Le Van Quyen M, Staba R, Bragin A, Dickson C, Valderrama M, Fried I, Engel J (2010) Large-scale microelectrode recordings of high-frequency gamma oscillations in human cortex during sleep. Journal of Neuroscience 30:7770-7782.

Lehmann D, Strik WK, Henggeler B, Koenig T, Koukkou M (1998) Brain electric microstates and momentary conscious mind states as building blocks of spontaneous thinking: I. Visual imagery and abstract thoughts. International Journal of Psychophysiology 29:1-11.

Lemieux M, Chen JY, Lonjers P, Bazhenov M, Timofeev I (2014) The impact of cortical deafferentation on the neocortical slow oscillation. The Journal of neuroscience : the official journal of the Society for Neuroscience 34:5689-5703.

Leonardi N, Van De Ville D (2015) On spurious and real fluctuations of dynamic functional connectivity during rest. NeuroImage 104:430-436.

Liégeois R, Ziegler E, Phillips C, Geurts P, Gómez F, Bahri MA, Yeo BTT, Soddu A, Vanhaudenhuyse A, Laureys S, Sepulchre R (2015) Cerebral functional connectivity periodically (de)synchronizes with anatomical constraints. Brain structure & function 221:2985-2997.

Litvak V, Mattout J, Kiebel S, Phillips C, Henson R, Kilner J, Barnes G, Oostenveld R, Daunizeau J, Flandin G, Penny W, Friston K (2011) EEG and MEG data analysis in SPM8. Computational intelligence and neuroscience 2011:852961.

Liu H, Li H, Wang Y, Lei X (2014) Enhanced brain small-worldness after sleep deprivation: a compensatory effect. Journal of sleep research 23:554-563.

Liu X, Lauer KK, Ward BD, Li SJ, Hudetz AG (2013) Differential effects of deep sedation with propofol on the specific and nonspecific thalamocortical systems: a functional magnetic resonance imaging study. Anesthesiology 118:59-69.

Liu Z, Fukunaga M, Zwart JAd, Duyn JH (2010) Large-scale spontaneous fluctuations and correlations in brain electrical activity observed with magnetoencephalography. NeuroImage 51:102-111.

Logothetis NK (2008) What we can do and what we cannot do with fMRI. Nature 453:869-878.
Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. Nature 412:150-157.
Loomis AL, Harvey EN, Hobart G (1935a) Further observations on the potential rhythms of the cerebral cortex during sleep. Science 82:198-200.
Loomis AL, Harvey EN, Hobart G (1935b) Potential rhythms of the cerebral cortex during sleep. Science 81:597-598.
Loomis AL, Harvey EN, Hobart GA (1937) Cerebral states during sleep, as studied by human brain potentials. Journal of Experimental Psychology 21:127.
Lopes da Silva F (2013) EEG and MEG: relevance to neuroscience. Neuron 80:1112-1128.
Lord L-D, Stevner AB, Deco G, Kringelbach ML (2017) Understanding principles of integration and segregation using whole-brain computational connectomics: implications for neuropsychiatric disorders. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 375.
Luckhoo H, Hale JR, Stokes MG, Nobre AC, Morris PG, Brookes MJ, Woolrich MW (2012) Inferring task-related networks using independent component analysis in magnetoencephalography. NeuroImage 62:530-541.
Luthi A (2014) Sleep spindles: where they come from, what they do. The Neuroscientist : a review journal bringing neurobiology, neurology and psychiatry 20:243-256.
Lyamin OI, Manger PR, Ridgway SH, Mukhametov LM, Siegel JM (2008) Cetacean sleep: An unusual form of mammalian sleep. Neuroscience & Biobehavioral Reviews 32:1451-1484.
MacDonald AA, Naci L, MacDonald PA, Owen AM (2015) Anesthesia and neuroimaging: investigating the neural correlates of unconsciousness. Trends in cognitive sciences 19:100-107.
Magnin M, Rey M, Bastuji H, Guillemant P, Mauguiere F, Garcia-Larrea L (2010) Thalamic deactivation at sleep onset precedes that of the cerebral cortex in humans. Proceedings of the National Academy of Sciences of the United States of America 107:3829-3833.
Maintz JA, Viergever MA (1998) A survey of medical image registration. Medical image analysis 2:1-36.
Mak-McCully RA, Rosen BQ, Rolland M, Régis J, Bartolomei F, Rey M, Chauvel P, Cash SS, Halgren E (2015) Distribution, Amplitude, Incidence, Co-Occurrence, and Propagation of Human K-Complexes in Focal Transcortical Recordings. eneuro 2:ENEURO.0028-0015.2015.
Mantini D, Della Penna S, Marzetti L, de Pasquale F, Pizzella V, Corbetta M, Romani GL (2011) A signal-processing pipeline for magnetoencephalography resting-state networks. Brain Connect 1:49-59.
Mantini D, Franciotti R, Romani GL, Pizzella V (2008) Improving MEG source localizations: an automated method for complete artifact removal based on independent component analysis. Neuroimage 40:160-173.
Mantini D, Perrucci MG, Gratta CD, Romani GL, Corbetta M (2007) Electrophysiological signatures of resting state networks in the human brain. Proceedings of the National Academy of Sciences of the United States of America 104:13170-13175.
Maquet P (2000) Functional neuroimaging of normal human sleep by positron emission tomography. Journal of sleep research 9:207-231.
Marrelec G, Krainik A, Duffau H, Pelegrini-Issac M, Lehericy S, Doyon J, Benali H (2006a) Partial correlation for functional brain interactivity investigation in functional MRI. Neuroimage 32:228-237.
Marrelec G, Krainik A, Duffau H, Pélégrini-Issac M, Lehéricy S, Doyon J, Benali H (2006b) Partial correlation for functional brain interactivity investigation in functional MRI. NeuroImage 32:228-237.
Marshall L, Helgadottir H, Molle M, Born J (2006) Boosting slow oscillations during sleep potentiates memory. Nature 444:610-613.

Marshall W, Gomez-Ramirez J, Tononi G (2016) Integrated information and state differentiation. Frontiers in psychology 7:926.
Martuzzi R, Ramani R, Qiu M, Rajeevan N, Constable RT (2010) Functional connectivity and alterations in baseline brain state in humans. NeuroImage 49:823-834.
Marzano C, Ferrara M, Moroni F, De Gennaro L (2011) Electroencephalographic sleep inertia of the awakening brain. Neuroscience 176:308-317.
Marzetti L, Della Penna S, Snyder AZ, Pizzella V, Nolte G, de Pasquale F, Romani GL, Corbetta M (2013) Frequency specific interactions of MEG resting state activity within and across brain networks as revealed by the multivariate interaction measure. NeuroImage 79:172-183.
Mashour GA (2014) Top-down mechanisms of anesthetic-induced unconsciousness. Frontiers in systems neuroscience 8:115.
Massimini M, Ferrarelli F, Huber R, Esser SK, Singh H, Tononi G (2005) Breakdown of cortical effective connectivity during sleep. Science (New York, NY) 309:2228-2232.
Massimini M, Huber R, Ferrarelli F, Hill S, Tononi G (2004) The sleep slow oscillation as a traveling wave. The Journal of Neuroscience 24:6862.
Mazaheri A, Jensen O (2010) Rhythmic pulsing: linking ongoing brain activity with evoked responses. Frontiers in human neuroscience 4.
McKeown MJ, Makeig S, Brown GG, Jung TP, Kindermann SS, Bell AJ, Sejnowski TJ (1998) Analysis of fMRI data by blind separation into independent spatial components. Human brain mapping 6:160-188.
Meisel C, Olbrich E, Shriki O, Achermann P (2013) Fading signatures of critical brain dynamics during sustained wakefulness in humans. The Journal of Neuroscience 33:17363-17372.
Meunier D, Lambiotte R, Bullmore E (2010) Modular and hierarchically modular organization of brain networks. Frontiers in Neuroscience 4.
Meyer C, Muto V, Jaspar M, Kussé C, Foret A, Mascetti L, Maquet P (2012) Chapter 8 - Neural correlates of human sleep and sleep-dependent memory processing A2 - Frank, Marcos G. In: Sleep and Brain Activity, pp 165-186 San Diego: Academic Press.
Mhuircheartaigh RN, Rosenorn-Lanng D, Wise R, Jbabdi S, Rogers R, Tracey I (2010) Cortical and subcortical connectivity changes during decreasing levels of consciousness in humans: a functional magnetic resonance imaging study using propofol. The Journal of neuroscience : the official journal of the Society for Neuroscience 30:9095-9102.
Mignot E (2008) Why we sleep: the temporal organization of recovery. PLoS biology 6:e106.
Moldofsky H (1989) Sleep and fibrositis syndrome. Rheumatic diseases clinics of North America 15:91-103.
Mölle M, Marshall L, Gais S, Born J (2002) Grouping of spindle activity during slow oscillations in human non-rapid eye movement sleep. Journal of Neuroscience 22:10941-10947.
Moruzzi G, Magoun HW (1949) Brain stem reticular formation and activation of the EEG. Electroencephalography and clinical neurophysiology 1:455-473.
Mukamel R, Gelbard H, Arieli A, Hasson U, Fried I, Malach R (2005) Coupling between neuronal firing, field potentials, and FMRI in human auditory cortex. Science 309:951-954.
Murphy K, Birn RM, Handwerker DA, Jones TB, Bandettini PA (2009a) The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? NeuroImage 44:893-905.
Murphy K, Fox MD (2016) Towards a consensus regarding global signal regression for resting state functional connectivity MRI. NeuroImage.
Murphy M, Huber R, Esser S, Riedner BA, Massimini M, Ferrarelli F, Ghilardi MF, Tononi G (2011) The cortical topography of local sleep. Current topics in medicinal chemistry 11:2438-2446.
Murphy M, Riedner BA, Huber R, Massimini M, Ferrarelli F, Tononi G (2009b) Source modeling sleep slow waves. Proceedings of the National Academy of Sciences 106:1608-1613.
Muthukumaraswamy S (2013a) High-frequency brain activity and muscle artifacts in MEG/EEG: A review and recommendations. Frontiers in human neuroscience 7.

Muthukumaraswamy SD (2013b) High-frequency brain activity and muscle artifacts in MEG/EEG: a review and recommendations. Front Hum Neurosci 7:138.
Newman ME (2006) Modularity and community structure in networks. Proceedings of the National Academy of Sciences of the United States of America 103:8577-8582.
Ní Mhuircheartaigh R, Warnaby C, Rogers R, Jbabdi S, Tracey I (2013) Slow-wave activity saturation and thalamocortical isolation during propofol anesthesia in humans. Science Translational Medicine 5:208ra148-208ra148.
Nir Y, Massimini M, Boly M, Tononi G (2013) Sleep and consciousness. In: Neuroimaging of consciousness, pp 133-182: Springer.
Nir Y, Staba RJ, Andrillon T, Vyazovskiy VV, Cirelli C, Fried I, Tononi G (2011) Regional slow waves and spindles in human sleep. Neuron 70:153-169.
Nobili L, De Gennaro L, Proserpio P, Moroni F, Sarasso S, Pigorini A, De Carli F, Ferrara M (2012) Chapter 13 - Local aspects of sleep: Observations from intracerebral recordings in humans. In: Progress in brain research, vol. Volume 199 (Andries Kalsbeek, M. M. T. R. and Russell, G. F., eds), pp 219-232: Elsevier.
Norman RG, Pal I, Stewart C, Walsleben JA, Rapoport DM (2000) Interobserver agreement among sleep scorers from different centers in a large dataset. Sleep 23:901-908.
O'Neill GC, Bauer M, Woolrich MW, Morris PG, Barnes GR, Brookes MJ (2015) Dynamic recruitment of resting state sub-networks. NeuroImage 115:85-95.
Ogawa S, Lee T-M (1990) Magnetic resonance imaging of blood vessels at high fields: In vivo and in vitro measurements and image simulation. Magnetic resonance in medicine 16:9-18.
Ogilvie RD (2001) The process of falling asleep. Sleep medicine reviews 5:247-270.
Ogilvie RD, Wilkinson RT (1984) The detection of sleep onset: behavioral and physiological convergence. Psychophysiology 21:510-520.
Olbrich E, Achermann P (2005) Analysis of oscillatory patterns in the human sleep EEG using a novel detection algorithm. Journal of sleep research 14:337-346.
Olbrich E, Claussen JC, Achermann P (2011) The multiple time scales of sleep dynamics as a challenge for modelling the sleeping brain. Philosophical transactions Series A, Mathematical, physical, and engineering sciences 369:3884-3901.
Oostenveld R, Fries P, Maris E, Schoffelen JM (2011) FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. Computational intelligence and neuroscience 2011:156869.
Osipova D, Takashima A, Oostenveld R, Fernández G, Maris E, Jensen O (2006) Theta and gamma oscillations predict encoding and retrieval of declarative memory. Journal of neuroscience 26:7523-7531.
Pace-Schott EF, Hobson JA (2002) The neurobiology of sleep: genetics, cellular physiology and subcortical networks. Nature reviews Neuroscience 3:591-605.
Palva JM, Palva S, Kaila K (2005) Phase synchrony among neuronal oscillations in the human cortex. Journal of Neuroscience 25:3962-3972.
Palva S, Palva JM (2012) Discovering oscillatory interaction networks with M/EEG: challenges and breakthroughs. Trends in cognitive sciences 16:219-230.
Pasquale Fd, Penna SD, Snyder AZ, Lewis C, Mantini D, Marzetti L, Belardinelli P, Ciancetta L, Pizzella V, Romani GL, Corbetta M (2010) Temporal dynamics of spontaneous MEG activity in brain networks. Proceedings of the National Academy of Sciences of the United States of America 107:6040-6045.
Pearlmutter BA, Houghton CJ (2009) A new hypothesis for sleep: tuning for criticality. Neural computation 21:1622-1641.
Piantoni G, Astill RG, Raymann RJEM, Vis JC, Coppens JE, Van Someren EJW (2013) Modulation of gamma and spindle-range power by slow oscillations in scalp sleep EEG of children. International Journal of Psychophysiology 89:252-258.
Piantoni G, Halgren E, Cash SS (2016a) The Contribution of Thalamocortical Core and Matrix Pathways to Sleep Spindles. Neural plasticity 2016:3024342.

Piantoni G, Halgren E, Cash SS (2016b) Spatiotemporal characteristics of sleep spindles depend on cortical location. Neuroimage 146:236-245.
Piantoni G, Van Der Werf YD, Jensen O, Van Someren EJ (2015) Memory traces of long-range coordinated oscillations in the sleeping human brain. Hum Brain Mapp 36:67-84.
Picchioni D, Duyn JH, Horovitz SG (2013) Sleep and the functional connectome. NeuroImage 80:387-396.
Picchioni D, Pixa ML, Fukunaga M, Carr WS, Horovitz SG, Braun AR, Duyn JH (2014) Decreased connectivity between the thalamus and the neocortex during human nonrapid eye movement sleep. Sleep 37:387-397.
Pivik RT, Harman K (1995) A reconceptualization of EEG alpha activity as an index of arousal during sleep: all alpha activity is not equal. Journal of sleep research 4:131-137.
Ponce-Alvarez A, Deco G, Hagmann P, Romani GL, Mantini D, Corbetta M (2015) Resting-state temporal synchronization networks emerge from connectivity topology and heterogeneity. PLoS computational biology 11:e1004100.
Portas CM, Krakow K, Allen P, Josephs O, Armony JL, Frith CD (2000) Auditory processing across the sleep-wake cycle: simultaneous EEG and fMRI monitoring in humans. Neuron 28:991-999.
Preti MG, Bolton TAW, Ville DVD (2016) The dynamic functional connectome: State-of-the-art and perspectives. NeuroImage.
Priesemann V, Valderrama M, Wibral M, Le Van Quyen M (2013) Neuronal avalanches differ from wakefulness to deep sleep – evidence from intracranial depth recordings in humans. PLoS computational biology 9:e1002985.
Rabinovich MI, Friston KJ, Varona P (2012) Principles of brain dynamics: global state interactions: MIT Press.
Raichle ME (2010) Two views of brain function. Trends in cognitive sciences 14:180-190.
Raichle ME (2015) The brain's default mode network. Annual review of neuroscience 38:433-447.
Raichle ME, MacLeod AM, Snyder AZ, Powers WJ, Gusnard DA, Shulman GL (2001) A default mode of brain function. Proceedings of the National Academy of Sciences 98:676-682.
Raichle ME, Mintun MA (2006) Brain work and brain imaging. Annual review of neuroscience 29:449-476.
Rasch B, Born J (2013) About sleep's role in memory. Physiological reviews 93:681-766.
Rechtschaffen A, Kales A (1968) A manual of standardized terminology, techniques and scoring system for sleep stages of human subjects. Bethesda, Md.: U.S. Dept. of Health, Education, and Welfare, Public Health Services-National Institutes of Health, National Institute of Neurological Diseases and Blindness, Neurological Information Network.
Rezek I, Roberts S (2005) Ensemble hidden Markov models with extended observation densities for biosignal analysis. In: Probabilistic modeling in bioinformatics and medical informatics, pp 419-450: Springer.
Riedner BA, Vyazovskiy VV, Huber R, Massimini M, Esser S, Murphy M, Tononi G (2007) Sleep homeostasis and cortical synchronization: III. A high-density EEG study of sleep slow waves in humans. Sleep 30:1643-1657.
Riemann D, Nissen C, Palagini L, Otte A, Perlis ML, Spiegelhalder K (2015) The neurobiology, investigation, and treatment of chronic insomnia. The Lancet Neurology 14:547-558.
Riemann D, Spiegelhalder K, Feige B, Voderholzer U, Berger M, Perlis M, Nissen C (2010) The hyperarousal model of insomnia: A review of the concept and its evidence. Sleep medicine reviews 14:19-31.
Robinson PA, Phillips AJK, Fulcher BD, Puckeridge M, Roberts JA (2011) Quantitative modelling of sleep dynamics. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 369:3840-3854.
Rosenberg RS, Van Hout S (2013) The American Academy of Sleep Medicine inter-scorer reliability program: sleep stage scoring. Journal of clinical sleep medicine : JCSM : official publication of the American Academy of Sleep Medicine 9:81-87.

Roux F, Wibral M, Singer W, Aru J, Uhlhaas PJ (2013) The phase of thalamic alpha activity modulates cortical gamma-band activity: evidence from resting-state MEG recordings. The Journal of Neuroscience 33:17827-17835.
Rubinov M, Sporns O (2010) Complex network measures of brain connectivity: uses and interpretations. NeuroImage 52:1059-1069.
Ryali S, Supekar K, Chen T, Kochalka J, Cai W, Nicholas J, Padmanabhan A, Menon V (2016) Temporal dynamics and developmental maturation of salience, default and central-executive network interactions revealed by variational bayes hidden markov modeling. PLoS computational biology 12:e1005138.
Sämann PG, Tully C, Spoormaker VI, Wetter TC, Holsboer F, Wehrle R, Czisch M (2010) Increased sleep pressure reduces resting state functional connectivity. Magma (New York, NY) 23:375-389.
Sämann PG, Wehrle R, Hoehn D, Spoormaker VI, Peters H, Tully C, Holsboer F, Czisch M (2011) Development of the brain's default mode network from wakefulness to slow wave sleep. Cerebral cortex 21:2082-2093.
Sanders RD, Tononi G, Laureys S, Sleigh JW (2012) Unresponsiveness≠ unconsciousness. The Journal of the American Society of Anesthesiologists 116:946-959.
Saper CB, Fuller PM, Pedersen NP, Lu J, Scammell TE (2010) Sleep state switching. Neuron 68:1023-1042.
Sarasso S, Boly M, Napolitani M, Gosseries O, Charland-Verville V, Casarotto S, Rosanova M, Casali AG, Brichant JF, Boveroux P, Rex S, Tononi G, Laureys S, Massimini M (2015) Consciousness and complexity during unresponsiveness induced by propofol, xenon, and ketamine. Current biology : CB 25:3099-3105.
Schabus M, Dang-Vu TT, Albouy G, Balteau E, Boly M, Carrier J, Darsaud A, Degueldre C, Desseilles M, Gais S, Phillips C, Rauchs G, Schnakers C, Sterpenich V, Vandewalle G, Luxen A, Maquet P (2007) Hemodynamic cerebral correlates of sleep spindles during human non-rapid eye movement sleep. Proceedings of the National Academy of Sciences of the United States of America 104:13164-13169.
Schabus M, Dang-Vu TT, Heib DP, Boly M, Desseilles M, Vandewalle G, Schmidt C, Albouy G, Darsaud A, Gais S, Degueldre C, Balteau E, Phillips C, Luxen A, Maquet P (2012) The fate of incoming stimuli during NREM sleep is determined by spindles and the phase of the slow oscillation. Frontiers in neurology 3:40.
Scheeringa R, Petersson KM, Kleinschmidt A, Jensen O, Bastiaansen MCM (2012) EEG α power modulation of fMRI resting-state connectivity. Brain connectivity 2:254-264.
Scheuler W, Stinshoff D, Kubicki S (1983) The Alpha-Sleep Pattern. Neuropsychobiology 10:183-189.
Schoffelen JM, Gross J (2009) Source connectivity analysis with MEG and EEG. Human brain mapping 30:1857-1865.
Schölvinck ML, Leopold DA, Brookes MJ, Khader PH (2013) The contribution of electrophysiology to functional connectivity mapping. Neuroimage 80:297-306.
Scholvinck ML, Maier A, Ye FQ, Duyn JH, Leopold DA (2010) Neural basis of global resting-state fMRI activity. Proceedings of the National Academy of Sciences of the United States of America 107:10238-10243.
Scholz J, Klein MC, Behrens TE, Johansen-Berg H (2009) Training induces changes in white-matter architecture. Nature neuroscience 12:1370-1371.
Schulz P, Steimer T (2009) Neurobiology of circadian systems. CNS drugs 23 Suppl 2:3-13.
Searle JR (1984) Minds, brains and science. Harvard University Press.
Sela Y, Vyazovskiy VV, Cirelli C, Tononi G, Nir Y (2016) Responses in rat core auditory cortex are preserved during sleep spindle oscillations. Sleep 39:1069-1082.
Shmuel A, Augath M, Oeltermann A, Logothetis NK (2006) Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. Nature neuroscience 9:569-577.

Shulman GL, Fiez JA, Corbetta M, Buckner RL, Miezin FM, Raichle ME, Petersen SE (1997) Common blood flow changes across visual tasks: II. decreases in cerebral cortex. Journal of cognitive neuroscience 9:648-663.
Siegel JM (2005) Clues to the functions of mammalian sleep. Nature 437:1264-1271.
Siegel M, Donner TH, Engel AK (2012) Spectral fingerprints of large-scale neuronal interactions. Nature reviews Neuroscience 13:121-134.
Simon NR, Manshanden I, Lopes da Silva FH (2000) A MEG study of sleep. Brain research 860:64-76.
Simons JS, Spiers HJ (2003) Prefrontal and medial temporal lobe interactions in long-term memory. Nature Reviews Neuroscience 4:637-648.
Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. Annual review of neuroscience 18:555-586.
Smith JK, Humes DJ, Head KE, Bush D, White TP, Stevenson CM, Brookes MJ, Marciani L, Spiller RC, Gowland PA, Francis ST (2011a) fMRI and MEG analysis of visceral pain in healthy volunteers. Neurogastroenterology and motility : the official journal of the European Gastrointestinal Motility Society 23:648-e260.
Smith SM (2002) Fast robust automated brain extraction. Human brain mapping 17:143-155.
Smith SM, Fox PT, Miller KL, Glahn DC, Fox PM, Mackay CE, Filippini N, Watkins KE, Toro R, Laird AR, Beckmann CF (2009) Correspondence of the brain's functional architecture during activation and rest. Proceedings of the National Academy of Sciences of the United States of America 106:13040-13045.
Smith SM, Miller KL, Salimi-Khorshidi G, Webster M, Beckmann CF, Nichols TE, Ramsey JD, Woolrich MW (2011b) Network modelling methods for FMRI. Neuroimage 54:875-891.
Snyder AZ (2016) Intrinsic Brain Activity and Resting State Networks. In: Neuroscience in the 21st Century (Pfaff, D. W. and Volkow, N. D., eds), pp 1-52 New York, NY: Springer New York.
Spoormaker VI, Czisch M, Maquet P, Jancke L (2011) Large-scale functional brain networks in human non-rapid eye movement sleep: insights from combined electroencephalographic/functional magnetic resonance imaging studies. Philosophical transactions Series A, Mathematical, physical, and engineering sciences 369:3708-3729.
Spoormaker VI, Gleiser PM, Czisch M (2012) Frontoparietal connectivity and hierarchical structure of the brain's functional network during sleep. Frontiers in neurology 3:80.
Spoormaker VI, Schroter MS, Gleiser PM, Andrade KC, Dresler M, Wehrle R, Samann PG, Czisch M (2010) Development of a large-scale functional brain network during human non-rapid eye movement sleep. The Journal of neuroscience : the official journal of the Society for Neuroscience 30:11379-11387.
Sporns O (2011) The non-random brain: efficiency, economy, and complex dynamics. Frontiers in computational neuroscience 5:5.
Sporns O (2013) The human connectome: origins and challenges. NeuroImage 80:53-61.
Sporns O (2015) Cerebral cartography and connectomics. Philosophical transactions of the Royal Society of London Series B, Biological sciences 370.
Sporns O, Tononi G, Edelman GM (2000) Connectivity and complexity: the relationship between neuroanatomy and brain dynamics. Neural networks : the official journal of the International Neural Network Society 13:909-922.
Sporns O, Tononi G, Kötter R (2005) The human connectome: a structural description of the human brain. PLoS computational biology 1:e42.
Spreng RN, Mar RA, Kim AS (2009) The common neural basis of autobiographical memory, prospection, navigation, theory of mind, and the default mode: a quantitative meta-analysis. Journal of cognitive neuroscience 21:489-510.
Stam CJ, Nolte G, Daffertshofer A (2007) Phase lag index: assessment of functional connectivity from multi channel EEG and MEG with diminished bias from common sources. Human brain mapping 28:1178-1193.

Steenland H (2014) Staying awake: Top-down systems control of sleep. OA Neurosciences 2(1):1.
Steriade M (2006) Grouping of brain rhythms in corticothalamic systems. Neuroscience 137:1087-1106.
Steriade M, Contreras D, Curro Dossi R, Nunez A (1993a) The slow (< 1 Hz) oscillation in reticular thalamic and thalamocortical neurons: scenario of sleep rhythm generation in interacting thalamic and neocortical networks. The Journal of neuroscience : the official journal of the Society for Neuroscience 13:3284-3299.
Steriade M, Nunez A, Amzica F (1993b) Intracellular analysis of relations between the slow (< 1 Hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. The Journal of neuroscience : the official journal of the Society for Neuroscience 13:3266-3283.
Steriade M, Nunez A, Amzica F (1993c) A novel slow (< 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. The Journal of neuroscience : the official journal of the Society for Neuroscience 13:3252-3265.
Stern JM, Caporro M, Haneef Z, Yeh HJ, Buttinelli C, Lenartowicz A, Mumford JA, Parvizi J, Poldrack RA (2011) Functional imaging of sleep vertex sharp transients. Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology 122:1382-1386.
Stickgold R (2005) Sleep-dependent memory consolidation. Nature 437:1272-1278.
Stickgold R (2013) Parsing the role of sleep in memory processing. Current opinion in neurobiology 23:847-853.
Stickgold R, Walker MP (2009) The neuroscience of sleep. Amsterdam: Academic Press, Elsevier.
Stoffers D, Altena E, Werf YDvd, Sanz-Arigita EJ, Voorn TA, Astill RG, Strijers RLM, Waterman D, Someren EJWV (2014) The caudate: a key node in the neuronal network imbalance of insomnia? Brain : a journal of neurology 137:610-620.
Stoffers D, Moens S, Benjamins J, van Tol MJ, Penninx BW, Veltman DJ, Van der Wee NJ, Van Someren EJ (2012) Orbitofrontal gray matter relates to early morning awakening: a neural correlate of insomnia complaints? Frontiers in neurology 3:105.
Tagliazucchi E, Balenzuela P, Fraiman D, Chialvo DR (2012a) Criticality in large-scale brain fMRI dynamics unveiled by a novel point process analysis. Frontiers in physiology 3:15.
Tagliazucchi E, Behrens M, Laufs H (2013a) Sleep neuroimaging and models of consciousness. Frontiers in psychology 4:256.
Tagliazucchi E, Crossley N, Bullmore ET, Laufs H (2015) Deep sleep divides the cortex into opposite modes of anatomical-functional coupling. Brain structure & function.
Tagliazucchi E, Laufs H (2014) Decoding wakefulness levels from typical fMRI resting-state data reveals reliable drifts between wakefulness and sleep. Neuron 82:695-708.
Tagliazucchi E, Laufs H (2015) Multimodal imaging of dynamic functional connectivity. Frontiers in neurology 6:10.
Tagliazucchi E, Siniatchkin M, Laufs H, Chialvo DR (2016) The voxel-wise functional connectome can be efficiently derived from co-activations in a sparse spatio-temporal point-process. Frontiers in Neuroscience 10:381.
Tagliazucchi E, von Wegner F, Morzelewski A, Borisov S, Jahnke K, Laufs H (2012b) Automatic sleep staging using fMRI functional connectivity data. NeuroImage 63:63-72.
Tagliazucchi E, von Wegner F, Morzelewski A, Brodbeck V, Borisov S, Jahnke K, Laufs H (2013b) Large-scale brain functional modularity is reflected in slow electroencephalographic rhythms across the human non-rapid eye movement sleep cycle. NeuroImage 70:327-339.
Tagliazucchi E, Wegner Fv, Morzelewski A, Brodbeck V, Jahnke K, Laufs H (2013c) Breakdown of long-range temporal dependence in default mode and attention networks during deep sleep. Proceedings of the National Academy of Sciences of the United States of America 110:15419-15424.

Tagliazucchi E, Wegner Fv, Morzelewski A, Brodbeck V, Laufs H (2012c) Dynamic BOLD functional connectivity in humans and its electrophysiological correlates. Frontiers in human neuroscience 6:339.

Tamaki M, Huang T-R, Yotsumoto Y, Hämäläinen M, Lin F-H, Náñez JE, Watanabe T, Sasaki Y (2013) Enhanced spontaneous oscillations in the supplementary motor area are associated with sleep-dependent offline learning of finger-tapping motor-sequence task. The Journal of Neuroscience 33:13894-13902.

Tassi P, Muzet A (2000) Sleep inertia. Sleep medicine reviews 4:341-353.

Terzano MG, Parrino L, Sherieri A, Chervin R, Chokroverty S, Guilleminault C, Hirshkowitz M, Mahowald M, Moldofsky H, Rosa A, Thomas R, Walters A (2001) Atlas, rules, and recording techniques for the scoring of cyclic alternating pattern (CAP) in human sleep. Sleep medicine 2:537-553.

Tewarie P, Bright MG, Hillebrand A, Robson SE, Gascoyne LE, Morris PG, Meier J, Mieghem PV, Brookes MJ (2016a) Predicting haemodynamic networks using electrophysiology: The role of non-linear and cross-frequency interactions. NeuroImage 130:273-292.

Tewarie P, Hillebrand A, Dellen Ev, Schoonheim MM, Barkhof F, Polman CH, Beaulieu C, Gong G, Dijk BWv, Stam CJ (2014) Structural degree predicts functional network connectivity: a multimodal resting-state fMRI and MEG study. NeuroImage 97:296-307.

Tewarie P, Hillebrand A, Dijk BWv, Stam CJ, O'Neill GC, Mieghem PV, Meier JM, Woolrich MW, Morris PG, Brookes MJ (2016b) Integrating cross-frequency and within band functional networks in resting-state MEG: A multi-layer network approach. NeuroImage 142:324-336.

Tognoli E, Kelso JA (2014) The metastable brain. Neuron 81:35-48.

Tononi G (2008) Consciousness as integrated information: a provisional manifesto. The Biological Bulletin 215:216-242.

Tononi G, Boly M, Massimini M, Koch C (2016) Integrated information theory: from consciousness to its physical substrate. Nature reviews Neuroscience 17:450-461.

Tononi G, Cirelli C (2003) Sleep and synaptic homeostasis: a hypothesis. Brain research bulletin 62:143-150.

Tononi G, Cirelli C (2013) Sleep and the price of plasticity: from synaptic and cellular homeostasis to memory consolidation and integration. Neuron 81:12-34.

Tononi G, Edelman GM, Sporns O (1998) Complexity and coherency: integrating information in the brain. Trends in cognitive sciences 2:474-484.

Tsai P-J, Chen SC-J, Hsu C-Y, Wu CW, Wu Y-C, Hung C-S, Yang AC, Liu P-Y, Biswal B, Lin C-P (2014) Local awakening: regional reorganizations of brain oscillations after sleep. NeuroImage 102 Pt 2:894-903.

Tyvaert L, LeVan P, Grova C, Dubeau F, Gotman J (2008) Effects of fluctuating physiological rhythms during prolonged EEG-fMRI studies. Clinical neurophysiology 119:2762-2774.

Tzourio-Mazoyer N, Landeau B, Papathanassiou D, Crivello F, Etard O, Delcroix N, Mazoyer B, Joliot M (2002) Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. NeuroImage 15:273-289.

Uehara T, Yamasaki T, Okamoto T, Koike T, Kan S, Miyauchi S, Kira J, Tobimatsu S (2014) Efficiency of a "small-world" brain network depends on consciousness level: a resting-state FMRI study. Cerebral cortex 24:1529-1539.

Uhrig L, Dehaene S, Jarraya B (2014) Cerebral mechanisms of general anesthesia. Annales francaises d'anesthesie et de reanimation 33:72-82.

Valera FJ, Toro A, John ER, Schwartz EL (1981) Perceptual framing and cortical alpha rhythm. Neuropsychologia 19:675-686.

Van Essen DC, Smith SM, Barch DM, Behrens TE, Yacoub E, Ugurbil K, Consortium WU-MH (2013) The WU-Minn Human Connectome Project: an overview. Neuroimage 80:62-79.

van Someren E, Cluydts R (2013) Sleep regulation and insomnia. In: Neuroscience in the 21st Century: From Basic to Clinical (Pfaff, D. W., ed), pp 1889-1916 New York, NY: Springer New York.
van Someren EJW, Cirelli C, Dijk D-J, Cauter EV, Schwartz S, Chee MWL (2015) Disrupted sleep: from molecules to cognition. Journal of Neuroscience 35:13889-13895.
Van Veen BD, van Drongelen W, Yuchtman M, Suzuki A (1997) Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. IEEE transactions on bio-medical engineering 44:867-880.
van Wijk BCM, Stam CJ, Daffertshofer A (2010) Comparing brain networks of different size and connectivity density using graph theory. PloS one 5:e13701.
Velly LJ, Rey MF, Bruder NJ, Gouvitsos FA, Witjas T, Regis JM, Peragut JC, Gouin FM (2007) Differential dynamic of action on cortical and subcortical structures of anesthetic agents during induction of anesthesia. Anesthesiology 107:202-212.
Vidaurre D, Abeysuriya R, Becker R, Quinn AJ, Alfaro-Almagro F, Smith SM, Woolrich MW (2017) Discovering dynamic brain networks from big data in rest and task. NeuroImage.
Vidaurre D, Quinn AJ, Baker AP, Dupret D, Tejero-Cantero A, Woolrich MW (2016) Spectrally resolved fast transient brain states in electrophysiological data. NeuroImage 126:81-95.
Vigario R, Sarela J, Jousmaki V, Hamalainen M, Oja E (2000) Independent component approach to the analysis of EEG and MEG recordings. IEEE transactions on bio-medical engineering 47:589-593.
Vincent JL, Patel GH, Fox MD, Snyder AZ, Baker JT, Van Essen DC, Zempel JM, Snyder LH, Corbetta M, Raichle ME (2007) Intrinsic functional architecture in the anaesthetized monkey brain. Nature 447:83-86.
Von Economo C (1930) Sleep as a problem of localization. The Journal of nervous and mental disease 71.
Von Stein A, Sarnthein J (2000) Different frequencies for different scales of cortical integration: from local gamma to long range alpha/theta synchronization. International journal of psychophysiology 38:301-313.
Vyazovskiy VV, Faraguna U, Cirelli C, Tononi G (2009) Triggering slow waves during NREM sleep in the rat by intracortical electrical stimulation: effects of sleep/wake history and background activity. Journal of neurophysiology 101:1921-1931.
Vyazovskiy VV, Harris KD (2013) Sleep and the single neuron: the role of global slow oscillations in individual cell rest. Nature Reviews Neuroscience 14:443-451.
Walker MP (2009a) The Role of Sleep in Cognition and Emotion. Annals of the New York Academy of Sciences 1156:168-197.
Walker MP (2009b) The role of slow wave sleep in memory processing. Journal of clinical sleep medicine : JCSM : official publication of the American Academy of Sleep Medicine 5:S20-26.
Wang XJ (2010) Neurophysiological and computational principles of cortical rhythms in cognition. Physiological reviews 90:1195-1268.
Warby SC, Wendt SL, Welinder P, Munk EG, Carrillo O, Sorensen HB, Jennum P, Peppard PE, Perona P, Mignot E (2014) Sleep-spindle detection: crowdsourcing and evaluating performance of experts, non-experts and automated methods. Nature methods 11:385-392.
Watanabe T, Kan S, Koike T, Misaki M, Konishi S, Miyauchi S, Miyahsita Y, Masuda N (2014) Network-dependent modulation of brain activity during sleep. NeuroImage 98:1-10.
Watts DJ, Strogatz SH (1998) Collective dynamics of 'small-world' networks. Nature 393:440-442.
Weber F, Dan Y (2016) Circuit-based interrogation of sleep control. Nature 538:51-59.
Wehrle R, Kaufmann C, Wetter TC, Holsboer F, Auer DP, Pollmacher T, Czisch M (2007) Functional microstates within human REM sleep: first evidence from fMRI of a thalamocortical network specific for phasic REM periods. The European journal of neuroscience 25:863-871.

Wennberg R (2010) Intracranial cortical localization of the human K-complex. Clinical Neurophysiology 121:1176-1186.
Wennberg R, Cheyne D (2013) On noninvasive source imaging of the human K-complex. Clinical Neurophysiology 124:941-955.
Wens V, Bourguignon M, Goldman S, Marty B, Beeck MOd, Clumeck C, Mary A, Peigneux P, Bogaert PV, Brookes MJ, Tiège XD (2014) Inter- and intra-subject variability of neuromagnetic resting state networks. Brain topography 27:620-634.
Werth E, Achermann P, BorbÉLy A (1997) Fronto‐occipital EEG power gradients in human sleep. Journal of sleep research 6:102-112.
Winkler AM, Ridgway GR, Webster MA, Smith SM, Nichols TE (2014) Permutation inference for the general linear model. NeuroImage 92:381-397.
Wong CW, Olafsson V, Tal O, Liu TT (2013) The amplitude of the resting-state fMRI global signal is related to EEG vigilance measures. NeuroImage 83:983-990.
Woolrich M, Hunt L, Groves A, Barnes G (2011) MEG beamforming using Bayesian PCA for adaptive data covariance matrix regularization. Neuroimage 57:1466-1479.
Woolrich MW, Baker A, Luckhoo H, Mohseni H, Barnes G, Brookes M, Rezek I (2013) Dynamic state allocation for MEG source reconstruction. NeuroImage 77:77-92.
Woolrich MW, Stephan KE (2013) Biophysical network models and the human connectome. NeuroImage 80:330-338.
Wulff K, Gatti S, Wettstein JG, Foster RG (2010) Sleep and circadian rhythm disruption in psychiatric and neurodegenerative disease. Nature reviews Neuroscience 11:589-599.
Yang G, Lai CSW, Cichon J, Ma L, Li W, Gan W-B (2014) Sleep promotes branch-specific formation of dendritic spines after learning. Science (New York, NY) 344:1173-1178.
Yordanova J, Kolev V, Wagner U, Born J, Verleger R (2012) Increased alpha (8–12 Hz) activity during slow wave sleep as a marker for the transition from implicit knowledge to explicit insight. Journal of cognitive neuroscience 24:119-132.
Zalesky A, Breakspear M (2015) Towards a statistical test for functional connectivity dynamics. NeuroImage 114:466-470.
Zalesky A, Fornito A, Bullmore E (2012) On the use of correlation as a measure of network connectivity. NeuroImage 60:2096-2106.
Zalesky A, Fornito A, Bullmore ET (2010) Network-based statistic: identifying differences in brain networks. NeuroImage 53:1197-1207.
Zepelin H, Siegel JM, Tobler I (2005) Mammalian sleep. In: Principles and Practice of Sleep Medicine, pp 91-100.
Zhang D, Raichle ME (2010) Disease and the brain's dark energy. Nature reviews Neurology 6:15-28.