

Mathematical Modelling and Inference for Alzheimer's Disease



Pavan Chaggar
Exeter College
University of Oxford

A thesis submitted in partial fulfillment
of the requirements for the degree

Doctor of Philosophy

Hilary 2024

To Mama, Papa and Gaggun

Acknowledgements

Working on this thesis has been a privilege that would not have been possible without my colleagues, friends, and family. I often fear that the people around me are unaware of how grateful I am for their presence in my life. I hope I can now convey the depth of appreciation I feel for them.

First, I would like to thank my supervisor, Alain Goriely. Alain has been extraordinary in his supervision, not only as a towering intellectual role model but as an inspiring example of personal compassion and balance. Working with and learning from Alain has been among the great experiences of my life thus far, and I am profoundly grateful for his guidance, wisdom, and tireless support throughout my thesis.

I thank my additional supervisors, mentors and collaborators, Saad Jbabdi, Stefano Magon and Gregory Klein, Bryn Roberts and Oskar Hansson, for your continued encouragement and engaging conversations throughout my doctoral journey. I would like to extend my thanks to the Doctoral Training Centre, specifically to Dave, Martin, Lorraine, Samantha and Melanie, for providing me the opportunity to study at Oxford and for your support throughout my DPhil.

I have had the good fortune of working with some exceptional early career researchers. In particular, I would like to thank Travis Thompson and Jake Vogel. Travis was instrumental in my early mathematical development, ensuring I made progress and felt supported in what was, at the time, a new field of study for me. My DPhil journey would have been substantially less fulfilling without his influence. Jake has been my closest collaborator during my DPhil and his influence on my DPhil is vast. Jake has always enthusiastically championed my work, provided me with countless opportunities to shine and constantly pushed me to be a better researcher. Without Jake's friendship and mentorship, enduring this DPhil would have been far less enjoyable than it has been.

I thank my friends (listed alphabetically to avoid squabbles): Andrew, Annika, Bea, Bisi, Caro, David, Eloise, Farbod, Geoff, George, Johnny, Joseph, Kris, Leo, Lex, Lionel, Michael, Sian, Sophie, and Soraya. With your enduring love and affection, I have grown to be a content and confident individual, capable of enjoying all that life has offered me. I look forward to enjoying more dinners and desserts with you all for many years to come.

To Bennie and Louis, thank you for being the best doggies.

Finally, to Mama, Papa, Gaggun and Pip. Thank you. Thank you for always allowing me to pursue my goals and for your unconditional love. Any and all of my accomplishments would not have been possible without you. Your steadfast belief in me has been my backbone throughout this journey, and I am endlessly grateful for the sacrifices you have made to see me succeed. I am blessed to have such incredible role models, and I will always strive to make you proud for all the love and kindness you have shown me.

Abstract

Alzheimer's disease is a devastating neurological condition characterised by the accumulation of two toxic proteins, amyloid-beta and tau, resulting in neurodegeneration and cognitive decline. The aetiology of Alzheimer's disease is multifaceted and complex, and there remain a limited number of treatment options available for patients. The mechanisms of Alzheimer's are particularly challenging to probe due to the numerous difficulties in observing disease mechanisms *in vivo* in the human brain. Instead, researchers must investigate from afar, relying on macro-scale neuroimaging methods and animal models of disease mechanisms. In this work, we address these obstacles by employing physics-based models of Alzheimer's pathology that can be paired with neuroimaging data to investigate disease mechanisms.

In Chapter 2, we introduce the two components necessary for modelling tau pathology: transport through axonal connections and production via prion-like replication. We show that a simple network-based reaction-diffusion model of tau and neurodegeneration can be usefully applied to human neuroimaging data to identify mechanisms of disease propagation, namely accelerated production of tau. In Chapter 3, we introduce a novel model of tau progression that accounts for the regional heterogeneity in vulnerability to tau pathology. We show that regionally specific production rates of tau are necessary for the accurate prediction of regional tau data and that the model can identify differences in transport and production dynamics across the Alzheimer's disease continuum. In Chapter 4, we use results from previous chapters to develop the first mechanistic model of the amyloid-tau-neurodegeneration pathway and apply it to multimodal longitudinal neuroimaging data. The model explains and predicts several aspects of Alzheimer's disease, including the necessity of amyloid for the acceleration of tau pathology, amyloid-induced regional heterogeneity of tau deposition, and tau-induced atrophy progression. Furthermore, we show that the spatial deposition of amyloid imparts a particular vulnerability to the entorhinal cortex for the deposition of tau seeds, providing a novel theory of amyloid-induced tau seeding.

Overall, we provide a framework for how mechanism-based modelling of Alzheimer's disease can aid investigation into Alzheimer's disease pathology and provide valuable tools for application in clinical research.

Contents

1	Introduction	1
1.1	Motivation and Thesis Aims	1
1.2	Biological Foundations of Alzheimer’s Disease	2
1.2.1	The Culprits: $A\beta$ and tau	2
1.2.2	The Amyloid Cascade Hypothesis	4
1.2.3	Staging Alzheimer’s and the ATN Framework	5
1.3	Brain Imaging for Alzheimer’s Disease	7
1.3.1	Structural MRI	7
1.3.2	Diffusion Weighted Imaging	9
1.3.3	PET	10
1.4	Dynamical Modelling of Alzheimer’s Disease	13
1.5	Thesis Outline and Contributions	15
2	A Modelling and Inference Pipeline for Alzheimer’s Disease	18
2.1	Introduction	18
2.2	Mathematical Modelling of Toxic Proteins	18
2.2.1	Protein Transport and the Graph Laplacian	19
2.2.2	Prion-like Proliferation	21
2.3	A Bayesian Workflow	27
2.4	Inference, Identifiability and Uncertainty	28
2.5	Applications to PET and sMRI Data	41
2.5.1	Methods	41
2.5.2	Results	46
2.6	Discussion	53
3	Regional Modelling of Tau Transport and Production	56
3.1	Introduction	56
3.2	Results	58

3.2.1	Deriving A Generative Model Of Tau Dynamics	58
3.2.2	Regional Heterogeneity Is Necessary for Longitudinal Forecasting	63
3.2.3	Local Model Accurately Forecasts Regional Tau Progression .	70
3.2.4	Early AD progression is driven by tau transport	76
3.3	Discussion and Conclusions	84
4	Dynamical Modelling of the ATN Pathway	88
4.1	Introduction	88
4.2	Modelling the ATN Pathway	90
4.2.1	Linear Stability and Bifurcation Analysis	92
4.3	A Measurement Model for ATN Biomarkers	95
4.3.1	Estimating Fixed Parameters	97
4.3.2	Simulating ATN Dynamics	100
4.4	Identifiability Analysis	103
4.5	Inference with the Coupled Model	108
4.6	Epicentre of Amyloid-Tau Colocalisation	116
4.7	Amyloid Induced Symmetry Breaking Promotes EC Vulnerability . .	123
4.8	Discussion	126
5	Concluding Thoughts	131
	Bibliography	136

List of Figures

1.1	Biomarker abnormalities predicted by Jack’s curves.	6
1.2	T1 weighted images from subjects across the AD continuum.	8
1.3	The DKT parcellation	9
1.4	An example connectome	11
1.5	Stages of τ -protein progression from SUVR.	12
2.1	Simulation from the diffusion model	21
2.2	Reaction kinetics of the heterodimer model.	22
2.3	Dynamics of the heterodimer model.	23
2.4	Simulation from the global FKPP model	26
2.5	Local structural identifiability with all regions observable.	31
2.6	Local structural identifiability with temporal regions observable. . . .	32
2.7	Local structural identifiability with entorhinal cortex observable. . . .	33
2.8	Local structural identifiability with inferior temporal lobe observable.	34
2.9	Probabilistic modelling framework	37
2.10	Practical identifiability analysis for the network FKPP model.	40
2.11	Regional tau and atrophy for A^+ and A^- groups.	47
2.12	Posterior distributions for model parameters	48
2.13	Posterior predictive simulations for tau.	51
2.14	Posterior predictive simulations for atrophy.	52
2.15	Posterior predictive residual analysis.	53
3.1	Simulated transport and production dynamics in the local FKPP model.	62
3.2	Model fit for in-sample and out-of-sample data.	67
3.3	Residual analysis.	69
3.4	Regional Residual analysis.	70
3.5	Posterior predictive plots for the left inferior temporal region.	72
3.6	Posterior predictive plots for the left entorhinal cortex.	73
3.7	Posterior predictive plots for the left entorhinal cortex.	74

3.8	Posterior predictive plots for the left entorhinal cortex.	75
3.9	Inferred population level parameters using ADNI data.	79
3.10	Inferred population level parameters using BF2 data.	80
3.11	Bivariate posterior distributions for production and transport.	81
3.12	Inferred population level parameters using ADNI data with changes to population-level prior distributions.	82
3.13	Inferred population level parameters using BF2 data with changes to population-level prior distributions.	83
4.1	Phase diagrams for the coupled $A\beta$ -tau model.	93
4.2	Comparison of atrophy with tau.	96
4.3	Population analysis for estimated fixed parameters of the ATN model.	99
4.4	Simulation from the ATN model.	102
4.5	Simulation from the ATN model with no $A\beta$ -tau coupling.	103
4.6	Sensitivity analysis for ATN model.	105
4.7	Identifiability the ATN model.	107
4.8	Posterior distributions from individualised model.	113
4.9	Posterior distributions from pooled model.	114
4.10	Posterior predictions from the ATN model calibrated to A^+T^+ subjects from ADNI.	115
4.11	Quantifying $A\beta$ -tau colocalisation sites.	119
4.12	Variable arrival time thresholds for $A\beta$ -tau colocalisation.	122
4.13	Posterior distributions from the ATN model calibrated to A^+T^+ sub- jects from ADNI.	125

Chapter 1

Introduction

1.1 Motivation and Thesis Aims

Alzheimer's disease (AD) is a complex neurological disorder characterised by the accumulation of toxic proteins, progressive structural brain damage and cognitive decline. AD remains incurable and difficult to treat, placing a significant strain on those affected by the disease and society as a whole. Despite the plentiful resources devoted to the treatment and understanding of the diseases, little progress has been made toward the development of effective pharmacological interventions. Research on the aetiology of AD predominantly comprises investigation into two toxic proteins, amyloid-beta, $A\beta$, and tau, through theories such as the amyloid hypothesis and the prion-like propagation hypothesis, among others [1, 2]. The challenges in making progressed toward developing effective disease altering treatments has motivated interdisciplinary investigation to help unify our knowledge of how $A\beta$ and tau promote AD progression. As path of this interdisciplinary topic, mathematical and computational modelling has proved to be a valuable tool for probing disease mechanisms in-silico and interpreting in-vivo observations. In this thesis, we hope to demonstrate how mathematical modelling and probabilistic inference can facilitate meaningful progress into understanding how AD unfolds in the human brain.

A significant obstacle hindering the scientific investigation into AD is the difficulty in making and understanding observations of the disease in-vivo in humans. In recent decades, sophisticated technologies have been developed to image human brains, notably positron emission tomography (PET) scanning and magnetic resonance imaging (MRI), allowing scientists to build upon theories of AD developed using in-vitro and in-vivo animal studies. Neuroimaging studies have validated post-mortem studies into spatiotemporal pattern of protein staging and brain atrophy [3, 4, 5, 6, 7]. Researchers have recognised the value of these data and several large neuroimaging databases now

exist to accelerate investigation, including the Alzheimer’s Disease Neuroimaging Initiative (ADNI), the Human Connectome Project, the Swedish BioFINDER study and more, making the process of obtaining and analysing observations less cumbersome and ultimately making the study of AD easier for scientists.

However, progress is still slow, in part because of the unyielding challenge of modelling and understanding high dimensional neuroimaging data. In this thesis I aim to address this problem through the marriage of dynamical systems modelling and probabilistic inference. The former allows for interpretability of high dimensional data through parsimonious mechanistic models; the latter allows us to quantify the uncertainty present throughout the stages of modelling, from parameter identification and uncertainty quantification to model selection. Together, these allow for the assessment of longitudinal patient data and the prediction of potential outcomes. Interrogating AD in this way requires at least three components: (1) biological understanding, (2) data, and (3) models. Biological understanding is needed to formulate models and data is needed to test them. In the following sections, I will briefly overview key features of $A\beta$ and tau and their role in the amyloid cascade theory (the leading theory of AD development), how stages of AD are characterised, current methods used for making in-vivo observations in humans, and finally the modelling framework employed and the challenges present in the modelling literature.

1.2 Biological Foundations of Alzheimer’s Disease

Two leading theories for the mechanisms of Alzheimer’s are the amyloid hypothesis [1, 8] and the prion-like propagation hypothesis [9, 2], each of which have undergone stages of redevelopment and now constitute complementary theories of $A\beta$ and tau progression. These theories, along with evidence from longitudinal biomarker progression, have become encompassed within the newly formulated amyloid-tau-neurodegeneration (ATN) framework for classifying stages of AD progression [10]. In this section, I will review the function and pathogenesis of $A\beta$ and tau, along with their roles in the amyloid cascade hypothesis and the ATN framework.

1.2.1 The Culprits: $A\beta$ and tau

$A\beta$ has been implicated in AD for decades, following the discovery of extracellular $A\beta$ plaques in the brains of AD patients post-mortem [11]. Since then, a wealth of evidence has come from genome wide association studies, in which researchers have found mutations related to $A\beta$ are associated with early onset AD. Mutations in

the amyloid precursor protein (APP), causing AD in people with Down’s syndrome, mutations in the APP processing protein presenilin cause early onset AD [12] and mutations in apolipoprotein E-4 (APOE- ϵ 4), associated with the breakdown of amyloid peptides, are a strong risk factor for A β pathology, plaque deposition and the development of AD [13, 14]. These findings have motivated extensive research into the role of A β in AD pathogenesis.

APP is a membrane bound protein believed to have a number of roles, including in neuroimmune responses, neuron formation and repair, and neuronal transport [15, 16]. APP undergoes significant post-translational modifications, including cleavage into A β . APP is cleaved through two pathways, the amyloidogenic pathway, driven by β -secretase, and the non-amyloidogenic pathway, driven by α -secretase, both are then cleaved by γ -secretase [16, 17]. Through the amyloidogenic pathway γ -secretase produces peptides of either length 40 or 42, A β -40 and A β -42, respectively. A β -42 is the more toxic of the two, with the ratio of A β -42 and A β -40 being indicative of A β pathology [18, 19], but both peptides contribute to the formation of oligomeric A β , fibrils and plaque formation [16, 20]. While the larger fibrils and plaques are insoluble, oligomeric A β is soluble and able to diffuse around in extracellular space. A β has also been shown to exhibit prion-like behaviour, meaning toxic oligomers of A β are able to induce conformational changes to healthy A β and turn them into toxic forms of A β that allow toxic A β to proliferate throughout the brain [21, 22]. The numerous findings of A β plaque involvement in post-mortem AD patients and the genetic risk factors that predispose AD have lead to the development of the amyloid cascade hypothesis [1, 23, 8], which we will discuss in the following section.

While there has been decades of research into A β , leading to drug development and prolonged clinical trials, tau has only recently gained research traction as a promising clinical target since it shows stable spatiotemporal staging that correlates well with brain atrophy and symptom onset, in contrast with A β whose accumulation largely precedes atrophy and is more homogeneously deposited throughout the brain [24, 25, 6, 10, 26].

Tau is derived from the microtubule-associated protein tau (MAPT) gene and is broadly associated with the structure of neurons and neuronal transport through its role in stabilising microtubules, scaffold-like structures that support neuronal structure and across which neuronal motor proteins travel [27, 28]. The role of tau in modulating microtubule function is mediated by phosphorylation. Tau binds to microtubule when dephosphorylated and unbinds when phosphorylated, each of which cause conformational changes to tau. The effective control of microtubule-mediated

neuronal structure and transport is highly dependent on a constant cycle phosphorylation and dephosphorylation of tau [27, 29]. Given its important role in neuronal stability, tau is ubiquitous across the brain.

In AD, tau aggregates and forms tangles, resulting in neuronal dysfunction and cell death. Tau aggregation follows a process of *hyperphosphorylation*, leading to excessive amounts of unbound tau. The hyperphosphorylation-driven increase in unbound tau is believed to be the first stage of tau misfolding and aggregation, where misfolded tau proceeds to aggregate into pretangles, followed by paired-helical-filaments (PHFs) and finally neurofibrillary tangles [30, 31, 32]. Like $A\beta$, tau exhibits prion-like behaviour, with experiments showing that hyperphosphorylated tau is capable of inducing conformational changes in healthy tau [33, 34, 35, 36, 2]. This process greatly accelerates the production of tau and results in widespread tauopathy. During normal ageing, there is natural accumulation of tau in the medial temporal lobe, termed primary age related tauopathy (PART), that is commonly found in healthy individuals during autopsy [4, 37]. PART is a *primary* tauopathy, in which no other agents are involved in the proliferation of tau. AD is a *secondary* tauopathy, in which another agent, namely $A\beta$, causes abhorrent proliferation of tau and is necessary for the expansion of tau pathology into the neocortex [38, 39, 40]. Therefore, the combination of both $A\beta$ and tau pathology are necessary for AD to develop.

Another key component of tau biology that will be frequently discussed throughout this thesis is the transport through axonal connections, leading to marked and conserved staging of tau during the course of AD. The particular spreading pattern of tau was documented by Braak and Braak, and show the progression of NFTs from the transentorhinal cortex, to limbic regions and then to the isocortex [41]. More recently, experimental work has shown that overexpression of tau seeds in the entorhinal cortex results in tau aggregation and spread to axonally connected regions [33, 42, 43]. Together with work showing the prion-like accumulation of tau, these results suggest the stereotypical spread of tau is a result of local accumulation and transsynaptic transport.

1.2.2 The Amyloid Cascade Hypothesis

The amyloid cascade hypothesis has been the leading theory of AD progression since its inception [44]. The theory makes the following predictions:

- $A\beta$ is the primary causative agent in the development of AD.
- $A\beta$ leads to the development of tau neurofibrillary tangles and neuronal death.

This theory has elicited continued controversy and sustained criticism, resulting has been regularly updates. There have been two main criticisms of the amyloid hypothesis. First, amyloid plaque burden does not correlate well with cognitive impairment, unlike neurofibrillary tau tangles, or similarly, patients can exhibit significant plaque burden but not be noticeably demented [44]. Second, there are cases where people have neurofibrillary tau tangles with no $A\beta$ burden, known as primary age-related tauopathy (PART), contradicting the causal role of $A\beta$ in initiating tau pathology as proposed by the amyloid hypothesis [37, 32].

These criticisms have been addressed by recent findings that $A\beta$ fibrils and plaques accelerate the formation of toxic tau and neurofibrillary tau tangles, demonstrating a causal role for $A\beta$ in the precipitation of tau pathology [38, 39, 40]. Newer incarnations of the amyloid hypothesis propose that a role of amyloid in AD as an agent of secondary tauopathy, that is as a catalyst for tau to spread further than it otherwise might during primary age related tauopathy [45, 32, 46, 10]. AD progression given by the updated amyloid hypothesis suggests AD starts with $A\beta$ pathology, followed by secondary tauopathy and finally neurodegeneration as a result of tau toxicity [44]. This staging is supported by clinical data is reflected in Jack's curves, a popular (qualitative) model of biomarker abnormality during Alzheimer's disease [46], shown in Figure 1.1. The amyloid hypothesis has recently received support following successful clinical trials of $A\beta$ clearing drugs, such as donanemab, lecanemab and aducanemab [47, 48, 49]. However, the clinical trial results demonstrate only mild slowing of disease progression. This could either be because administration occurred too late in the AD continuum [50] or because $A\beta$ does not directly result in AD pathology but instead acts in concert with tau to drive AD progression. In sum, the current version of the amyloid cascade hypothesis summarises and highlights the dual roles of $A\beta$ and tau in the pathogenesis of AD.

1.2.3 Staging Alzheimer's and the ATN Framework

The stereotypical progression of $A\beta$ and tau described by the amyloid cascade hypothesis has recently been used to propose a new framework for defining the stages of AD, the amyloid-tau-neurodegeneration (ATN) framework, now used for the classification of AD stages by the National Institute of Ageing [10]. Through this framework the authors propose to move away from a solely phenotypic definition of AD toward one that is informed by biomarker progression, in particular biomarkers related to amyloid, tau and neurodegeneration. The framework defines an AD *continuum* characterised by the levels of ATN biomarkers, such as $A\beta$ or tau fluid biomarkers, PET




Figure 1.1: **Biomarker abnormalities predicted by Jack's curves.** Figure adapted from [46]

and structural MRI. The framework defines eight categories of ATN profiles and diagnostic labels, shown in table Table 1.1. The positivity of each biomarker is denoted with a superscript, for example, amyloid positive is denoted A^+ . The Alzheimer's continuum comprised four sets of biomarker results, $A^+T^-N^-$, $A^+T^+N^-$, $A^+T^+N^+$, and $A^+T^-N^+$, with other profiles being either non-pathological or relating to some other pathology not related to AD. The ATN framework requires $A\beta$ positivity for the definition of AD. Similarly to the amyloid hypothesis, the ATN framework suggests that $A\beta$ pathology is followed by tau pathology and finally neurodegeneration. While the ATN framework indicates a staging of Alzheimer's that is predicted by the amyloid hypothesis, it does not make claims as to the causal structure dictating the interaction among the ATN biomarkers. In Chapter 4 we will address how these $A\beta$ and tau interact to cause the expected AD staging.

ATN profile	Biomarker category	
$A^-T^-N^-$	Normal AD biomarkers	
$A^+T^-N^-$	Alzheimer's pathological change	
$A^+T^+N^-$	Alzheimer's disease	
$A^+T^+N^+$	Alzheimer's disease	Alzheimer's disease continuum
$A^+T^-N^+$	Alzheimer's and concomitant suspected non Alzheimer's pathological change	
$A^-T^+N^-$	Non-AD pathological change	
$A^-T^-N^+$	Non-AD pathological change	
$A^-T^+N^+$	Non-AD pathological change	

Table 1.1: ATN biomarker profiles. (Adapted from [10].)

1.3 Brain Imaging for Alzheimer's Disease

In this section, I will review the brain imaging methods that are used throughout this thesis to generate the data used to assess models. In particular, I will discuss three imaging modalities: (1) structural MRI (sMRI), used to measure brain atrophy; (2) diffusion weighted MRI (dwMRI) used to track axonal connections; and (3) PET, used to quantify $A\beta$ and tau deposition in-vivo. These data are widely used by the AD research community and represent a valuable tool through which we are able to probe the progression of AD and calibrate models.

1.3.1 Structural MRI

Structural magnetic resonance imaging (sMRI) uses magnetic fields to detect differences in brain tissue composition and allows us to image the structure of the brain in-vivo. Images are obtained by inducing a static magnetic field around the tissue, causing protons in the tissue to align with the static field. Next a brief radiofrequency (RF) pulse is applied to perturb the magnetisation of proton, tilting them transversely to the static magnetic field and inducing synchronised precession. The precession creates a current in measurement coils around the MRI machine, generating the measured signal. Following the RF pulse, protons undergo two processes,




Figure 1.2: **T1 weighted images from subjects across the AD continuum. Figure adapted from [51]**, showing three stages of AD progression, from cognitively normal (CN), amnesic mild cognitive impairment (MCI) and finally Alzheimer’s Disease (AD).

first, realigning with the static magnetic field in the longitudinal plane, and second, losing precession in the transverse plane, each occurring over a different time period, T1 and T2, respectively. Different tissue types have different T1 and T2 times depending on their composition, allowing us to differentiate between them. Different RF pulse sequences are used to vary measurement sensitivity to either T1 or T2 signal, producing either T1- or T2- weighted images. Typically, T1-weighted images are used to monitor structural changes to brain matter in AD patients. Fat has a faster T1 time compared to water and therefore produces more signal in a T1-weighted image, making fatty tissues such as axons appear brighter. An example of a T1-weighted image from people across the AD spectrum is shown in Figure 1.2.

In AD research, sMRI is used as an accurate and non-invasive measure of how the structure of brain regions changes over time, for example as a measure of tissue atrophy due to disease. An example of sMRI scans from subjects across the AD continuum is shown in Figure 1.2. To do this, sMRI images (such as those in Figure 1.2) must first be *segmented* into different regions of interest (ROIs), together comprising a *parcellation*. There are many types of parcellation, each of which characterise different properties of the brain, for example, brain anatomy, functional connectivity, gyrification, among others [52, 53]. For consistency with the AD research community, we adopt the commonly used Desikan-Killiany-Tourville (DKT) parcellation which delineates regions based on anatomical landmarks and is supplied as standard in the FreeSurfer software package [54, 55], with the addition of deep grey nuclei, bilateral amygdala and hippocampus [56]. The regions of the parcellation are shown in Figure 1.3. In AD research, subjects are scanned over a period of years to measure brain




Figure 1.3: **The DKT parcellation.** The standard FreeSurfer DKT parcellation shown on a MNI brain. Top: 68 cortical regions for left and right hemispheres. Bottom: 15 subcortical regions, including the brainstem

atrophy. For example, the hippocampus in the medial temporal lobe is among the first regions to undergo significant degeneration in AD and can predict disease subtype and progression [57, 26]. In Chapters 2 and 4 we will use sMRI to model AD progression along with other quantitative biomarkers such as tau and $A\beta$ PET.

1.3.2 Diffusion Weighted Imaging

Diffusion weighted imaging is an MRI technique that measures the anisotropic diffusion of hydrogen atoms present in water across axonal connections and can be used to construct networks of connectivity between brain regions. Such a network can be used to model the intracellular spread of tau through axons in the brain.

In unrestricted spaces, such as cell bodies, water atoms undergo isotropic diffusion, moving randomly in any direction. In axons, the diffusion of atoms is restricted and it undergoes anisotropic diffusion along the axon. Diffusion weighted imaging relies on the sequential application of opposing magnetic gradients at a given angle to dephase and rephase the precession of protons in a particular axis. The successful rephasing of protons relies on minimal movement of protons between the applications of the opposing magnetic gradients and therefore the signal loss between magneti-

sations can be used to measure the movement of protons. There is low signal loss in directions in which there is restricted movement (e.g. transverse to an axon) and high signal loss in directions of unrestricted movement (e.g. along an axon). To remove the effect of signal loss due to T2 relaxation, the diffusion weighted images are interpreted relative to a baseline T2 image. Regions of high signal loss relative to the baseline T2 correspond to high rates of diffusion in the direction of the magnetic gradient. By applying magnetic gradients at multiple angles, it is possible to characterise the directions of proton movement, with greater angular resolution providing more detailed characterisations of proton diffusion, and therefore axon direction.

Once diffusion weighted images are acquired, it is possible to estimate the trajectory of axonal fibres between regions. Fibre reconstruction is achieved by *tractography*, which uses local information about diffusive anisotropy to piece together trajectories of diffusion and compile a synthetic axonal fibre, or *streamline*. There are multiple ways to perform tractography, throughout this work we use probtrackx, as provided by the FMRIB Software Library (FSL) [58, 59]. Probtrackx samples multiple streamlines from each voxel and each streamlines follows a particular trajectory with a probability given by the orientation distribution function, which is estimated voxel-wise from the diffusion weighted images. Streamlines terminate according to set criteria, including thresholds relating to fractional anisotropy, curvature or path length. There are also exclusion criteria based on anatomical priors, for example, tracts are deemed invalid if they stop outside the brain or in the ventricles. The streamlines from voxel to voxel are then summarised over a given parcellation. Since we use the DKT parcellation to process and understand sMRI and PET data, we also use it for summarising streamlines. The final output is a connectome, which is a dense graph with edge weights corresponding to the density of streamlines between each pair of regions. An example connectome generated from healthy subjects in the human connectome project is shown in Figure 1.4.

1.3.3 PET

Positron Emission Tomography (PET) is a method for brain imaging that relies on the activity of radioactive tracer. Radiotracers, such as flortaucipir, used for tracking tau in-vivo, are injected into the subject’s bloodstream and allowed to perfuse to relevant tissue. The radiotracers undergo beta decay to produce *positrons* that quickly interact with nearby electrons to release a pair of photons in opposite directions. The photons are detected by detectors around the subject. The presence of the radioactive tracer therefore depends on the simultaneous detection of photons travelling in opposite




Figure 1.4: **An example connectome.** A connectome generated using 150 healthy subjects from the Human Connectome Project, processed using FSL and summarised over the DKT atlas. Left: adjacency matrix of the network. Right: network visualised on the brain. Nodes sizes are weighted by degree and edges are coloured by relative density.

directions. The detection of these coincidence events allow for the reconstruction of three-dimensional brain images showing the relative density of tracer activity.

The recent advent of PET tracers that bind to AD biomarkers has allowed for observations to be made in-vivo [60, 61, 62, 63]. Two of the most popular radiotracers are the first generation AV1451, used in ADNI [56], and second generation RO948 used in the second Swedish BioFINDER study (BF2) [64]. These tracers primarily bind to tau lesions comprised of paired helical filaments, which allows for the quantification of tau related AD pathology [65]. However, radiotracers also exhibit off-target binding, attaching to molecules other than paired-helical filaments of tau and resulting in false signal. This can significantly limit the utility of tracers in the effected brain regions. AV1451 signal exhibits particularly strong contamination in the basal ganglia and choroid plexus [66, 67, 68]. While these are not areas that typically show particularly strong invasion in AD, the affect of off-target binding interferes with signal in important surrounding areas, namely the hippocampus [69, 66]. These issues are partially circumvented with next generation tracers, such as RO498, that show negligible binding in the choroid plexus and reduced binding the basal ganglia [65, 70, 71].

An example of how tau PET can be used to monitor disease progression is shown in Figure 1.5, adapted from [4]. Here, the authors use a combination of an $A\beta$ PET tracer (PiB) and a tau tracer (AV1451) to track disease progression from young adults (YA), to older adults without $A\beta$ (PiB- OA), older adults with $A\beta$ (PiB+ OA) and finally patients diagnosed with AD. Figure 1.5 shows the mean tau PET standardised

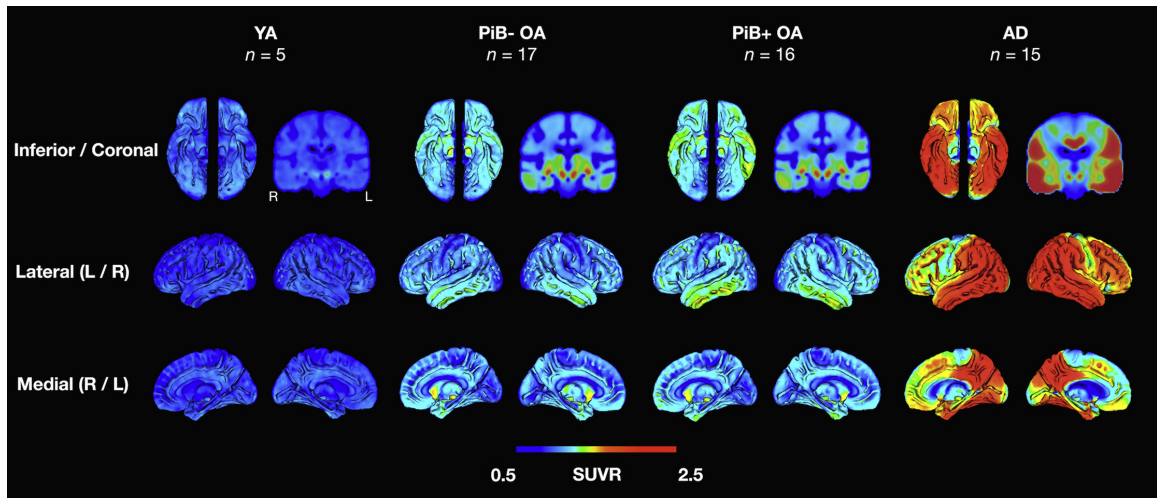


Figure 1.5: **Stages of τ -protein progression from SUVR.** Figure adapted from [4], showing the mean tau SUVR in a young adult population (YA), older adults (OA) with and without amyloid (determined by PiB PET), PiB- and PiB+, respectively, and subjects diagnosed with AD.

uptake value ratio (SUVR) in each of these groups. The figure demonstrates two important features of tau PET. First, tau PET is able to capture the pronounced staging of tau through the brain during AD progression. Second, tau staging observed through PET differs from that observed through post-mortem histological analysis [41], in particular through strong tau signal involvement in the occipital lobe. The differences in staging likely arise from regional variations in tracer dynamics, such as uptake and binding affinity. Tracer staging profiles have now been extensively validated [5, 6, 72, 73, 74].

These data have been used extensively by modellers for parameter calibration and validation [72, 75, 76]. Similarly, studies using structural MRI have shown how atrophy progresses during AD [7], highlighting the relationship between atrophy and tau, but not $A\beta$ [3, 4, 5]. As yet, few studies, have used structural MRI as observations to fit models, however, their relative abundance compared to PET data make them a valuable resource for modelling, especially if they can be usefully combined for multimodal inference [77, 78, 79].

Overall, neuroimaging has provided a wealth of knowledge and data that has advanced AD research in its own right, and has also facilitated modelling studies probing disease mechanism.

1.4 Dynamical Modelling of Alzheimer’s Disease

Network models of neurodegeneration have been used extensively to study the progression of toxic proteins during AD but come with a number of challenges, some inherent with dynamical systems models and some particular to the modelling domain. The use of models allows for a mechanistic understanding of high-dimensional neuroimaging data that would otherwise be only qualitatively understood. There have been some attempts at continuum models of spatiotemporal prion progression [80], reviewed further in Chapter 2. However, simulating continuum reaction-diffusion across a fine brain mesh is expensive and prohibits comparisons against data. Therefore, modelling work has largely focused on connectome-based models of disease progression, in particular of tau. This simplifying assumption is well justified, since tau predominantly spreads via axonal connections (see Section 1.2.1). Network models of AD typically describe two processes: (1) transport across axons; and (2) growth via an autocatalytic prion-like process (see Section 1.2.1). The first model, describing only transport of tau through axons, was provided by Raj et al. 2012, who show that a network diffusion model is capable of predicting tau concentration and brain atrophy [77, 78], however the model does not account for the non-linear production processes underlying tau dynamics, limiting its application to real-world longitudinal data. Another popular model of protein abnormality is the epidemic spreading model, a non-linear probabilistic model describing transport and production, developed by Iturria-Medina et al., 2014. which was later shown to fit the progression of tau more accurately than the network diffusion model [81, 72]. However, the model lacks interpretability and is unable to disentangle transport and production dynamics. Numerous groups have built on this work to expand on these models with more expressive, interpretable, physics-based descriptions of toxic protein production [82, 83, 80, 84], however there has been limited application of physics-based models to longitudinal data and there are many challenges preventing the widespread application of modelling in clinical research.

Models of toxic protein evolution typically take the form of a system of ordinary differential equations on a graph. Let $G = (V, E)$ be a graph with R vertices $v_i \in V$, corresponding to regions in the brain atlas, and edge set $e_{ij} \in E$, for all $i, j \in 1, \dots, R$, representing the axonal connections between regions. Then, for a single species of toxic protein with variables $\mathbf{p} = \mathbf{p}(t) = (p_1(t), \dots, p_R(t))$ associated with each node

i, the general network reaction-diffusion system is given by

$$\frac{d\mathbf{p}}{dt} = \underbrace{-\mathbf{L}\mathbf{p}}_{\text{Transport}} + \underbrace{\mathbf{R}(\mathbf{p}, \boldsymbol{\theta})}_{\text{Production}}. \quad (1.1)$$

The first term represents transport across the network, with transport coefficient $\rho \in \mathbb{R}_{\geq 0}$, and graph Laplacian, $\mathbf{L} \in \mathbb{R}^{R \times R}$ which acts similarly to the continuous Laplacian operator to describe transport of toxic protein across the network. The second term is a function \mathbf{R} with parameters $\boldsymbol{\theta} \in \mathbb{R}^\theta$ dictating local processes such as the local production or clearance of toxic protein. Representing the system in this way highlights the challenges faced by the disease progression modelling community, since each element of the system brings their own challenges that should be addressed.

First, there is parametric uncertainty associated with the parameter vector $\boldsymbol{\theta}$. For any given dynamical system, variations in parameters can lead to different model behaviours. In general, it is difficult to infer from observations alone the parameters of the dynamical system, a problem that is made more challenging in the presence of observation noise. Popular methods for inferring parameter values include least squares regression, maximum likelihood estimation and maximum a-posteriori estimation and have been used in the network neurodegeneration literature for model validation [77, 78, 72]. However, for ill-posed problems with sparse and noisy observations, and where data are assumed to be generated from a non-linear process, these methods are unsuitable since they do not account for potentially significant parametric uncertainty. Being able to account for this uncertainty is a vital step in bridging the application of disease progression models to clinical research practice. A framework for handling different sources of uncertainty will be discussed in Chapter 2 and Chapter 3.

Second, there is uncertainty about the nature of the function \mathbf{R} , principally due to our limited understanding of AD pathology. Several models have been presented in the literature, including the simple network diffusion model, [77, 78], prion-like proliferation, [82], Smoluchowski-type aggregation-fragmentation [83, 84]. Each of these models imbue different processes and assumptions into a model of AD pathology and at present there is little evidence to choose one over another. Thus it is important to be able to identify which of these models is most appropriate given the available data. There are numerous methods available for this, including a host of information criteria, such as the Akaike information criteria, Bayesian leave-one-out cross validation or comparison on out-of-sample data [85]. Developing suitable functions to describe

the local dynamics of toxic proteins in AD is a common theme throughout this thesis and will be addressed in later research chapters.

Third, there is considerable uncertainty associated with the graph used to define the dynamical system. The graph not only defines the state vector, \mathbf{p} , but the transport between its elements, and therefore its properties can have a large effect on the dynamics of the system. As discussed in Section 1.3.2, the process of generating a connectome can be broadly separated into two processes: (1) parcellating the brain; and (2) performing tractography. Each of these two processes are active areas of research in imaging neuroscience and there is no canonical choice for either. For the current work, we generate connectomes using the DKT [54, 55] with the addition of deep grey nuclei, bilateral hippocampus and amygdala, and brainstem. The parcellation is visualised in Figure 1.3. We use FSL to perform tractography, however, it should be noted that different tractography pipelines can result in different connectomes. Furthermore, different edge weightings and Laplacian normalisations can affect dynamics. In [86] we address how different connectomes generation pipelines and weighting schemes effect protein staging in [86]. We show that different connectome construction pipelines and sizes yield significantly different tau staging dynamics. We also conclude that length-free edge weightings are most likely to yield Braak staging (assuming simple reaction-diffusion dynamics). Connectomes averaged over groups, especially those taken from healthy populations such as from the Human Connectome Project, have limited utility in disease progression modelling since they do not account for how disease processes affect the structural brain network. To address this deficiency, researchers have employed null model testing [87] to test the validity of network-based models and are working towards personalised connectomes for precision medicine [88].

1.5 Thesis Outline and Contributions

This thesis will in large part be devoted to addressing the challenges describes in the previous section, and is organised as follows.

1. In Chapter 2, I develop the mathematical and inference framework allowing application to patient data and show that a simple model of tau propagation and atrophy can be calibrated to patient data to gain insight into disease dynamics.
2. In Chapter 3, I develop a more accurate model of tau propagation that accounts for regional variations in tau pathology and PET sensitivity, and show that

the model can be used for accurate, regionally specific predictions, as well as understanding differences in disease dynamics across the AD continuum.

3. In Chapter 4, I develop a model of the interactions between $A\beta$, tau and neurodegeneration to gain greater understanding on how these two proteins conspire to cause AD pathology.
4. In Chapter 5, I will offer some concluding remarks and proposals for future research goals.

The results of this thesis are the result of a several collaborative efforts and I would like to acknowledge the specific contributions of my coauthors. In Chapter 2, Amelie Schafer was primarily responsible for the analysis on data processed by myself and ADNI, in particular results in Figure 2.11b. In Chapter 3, we use data provided by BioFINDER-2, acquired and processed by the BioFINDER team and kindly provided by Oskar Hansson.

The following is a list of publications during my thesis with my specific contributions specified.

1. Travis B. Thompson, **Pavanjit Chaggar**, Ellen Kuhl, Alain Goriely, and Alzheimer's Disease Neuroimaging Initiative. "Protein-protein interactions in neurodegenerative diseases: A conspiracy theory." *PLoS computational biology* 16, no. 10 (20 20): e1008267.

I conducted the data analysis for this project and helped with interpretation, writing and revisions. The model developed in this paper is further analysed in Chapter 4 of this thesis.

2. Prama Putra, Travis B. Thompson, **Pavanjit Chaggar** and Alain Goriely. "Braiding Braak and Braak: Staging patterns and model selection in network neurodegeneration." *Network Neuroscience* 5 (2021): 929 - 956.

I analysed dMRI data from the Human Connectome Project to create the probabilistic connectomes used in the paper and throughout this thesis.

3. Amelie Schäfer, **Pavanjit Chaggar**, Travis B. Thompson, Alain Goriely, Ellen Kuhl, and Alzheimer's Disease Neuroimaging Initiative. "Predicting brain atrophy from tau pathology: A summary of clinical findings and their translation into personalized models." *Brain Multiphysics* 2 (2021): 100039. **Joint first author.**

I helped develop models and visualisation in this project. I conducted the analysis of sMRI data, coding and running simulation and inference analysis, and manuscript writing and revisions for this project.

4. Amelie Schäfer, **Pavanjit Chaggar**, Alain Goriely and Ellen Kuhl. "Correlating tau pathology to brain atrophy using a physics-based Bayesian model." *Engineering with Computers* 38 (2022): 3867 - 3877. **Joint first author.**

I helped with visualisation, manuscript writing and revisions. I conducted the analysis of sMRI data from a larger cohort of subjects in ADNI, and programmed and ran then simulation and inference analysis for this project. This work is included in Section 2.5.

5. **Pavanjit Chaggar**. Jacob Vogel, Alexa Pichet Binette, Travis B. Thompson, Olof Strandberg, Niklas Mattsson-Carlsson, Linda Karlsson, Erik Stomrud, Saâd Jbabdi, Stefano Magon, Gregory Klein, Oskar Hansson and Alain Goriely. "Personalised regional modelling predicts tau progression in the human brain." *bioRxiv* (2023)

I developed the models, conducted the relevant data analysis, programmed and ran the simulation and inference studies, created visualisations and drafted the initial manuscript. This manuscript forms Chapter 3 of this thesis.

Chapter 2

A Modelling and Inference Pipeline for Alzheimer's Disease

2.1 Introduction

Mathematical modelling of AD has becoming a fertile research programme in recent years, producing models that describe disease mechanisms and making valuable predictions about patient trajectories [78, 82, 89, 79]. In particular, network models of neurodegeneration have become increasingly popular due to their desirable balance of simplicity and explanatory power. However, for models to have practical applications they must be informative about relevant data (for example, sMRI, PET). To assess the suitability of models to explain data, we look at the *structural* and *practical* identifiability of models. Structural identifiability refers to models having unique input-output relations (they are injective functions), whereas practical identifiability refers to the whether inputs (initial conditions, parameters) can be inferred from observed data. Assessing whether models are appropriately expressive and identifiable from data will be the aim of this chapter. First, I will describe the the mathematical basis for network models of proteopathy. Second, I explore identifiability a simple non-linear network reaction-diffusion model of proteopathy. Third, I will apply this pipeline to sMRI and PET data from ADNI and show how models can be used to gain understanding about disease processes.

2.2 Mathematical Modelling of Toxic Proteins

Mathematical models of neurodegeneration have focused on modelling the spread, accumulation and aggregation of two toxic protein species, tau and $A\beta$. While there

have been contributions focusing on continuum dynamics [80, 83], as well as probabilistic network models [72], we will focus on dynamical models on structural brain networks, formulated as a system ODEs. In this chapter, we will explore models of tau, since it displays richer spatiotemporal dynamics that are more tightly coupled with atrophy and symptom onset (see Section 1.2 for review). The use of network models is motivated by experimental results demonstrating that tau preferentially travels through axonal fibres [42, 43, 90]. An additional benefit is that solutions to ODEs on graphs are computationally less expensive to obtain than their partial differential equation counterparts and therefore facilitate inference, which in itself is a computationally demanding task. As mentioned in Section 1.4, network models of AD protein pathology should aim to describe at least transport across axons and production via an autocatalytic process akin to prion-like templating. In the following section I will review the mathematical tools used to describe these two processes.

2.2.1 Protein Transport and the Graph Laplacian

The graph Laplacian is the central object used to describe transport along axonal fibres. The topology of axonal pathways can be obtained through the analysis of diffusion-weighted MRI data, such as those available from the Human Connectome Project. Here and throughout this thesis we use a connectome generated from 150 subjects healthy subjects in the Human Connectome Project [91, 92] and tractography was performed using probtrackx in FSL with default parameters [59]. The output of this analysis is a graph $G = (V, E)$, comprised of R nodes $v_i \in V$ (corresponding to brain regions) and edge set $e_{ij} \in E$ for $i, j = 1, \dots, R$. The transport of protein concentration between regions of the brains can be modelled using the graph Laplacian, defined by:

$$\mathbf{L} = \mathbf{D} - \mathbf{A}, \quad (2.1)$$

where \mathbf{A} is an adjacency matrix encoding the connectivity between regions.

$$A_{ij} = \begin{cases} 1 & \text{if an edge connects } v_i \text{ to } v_j \\ 0 & \text{otherwise} \end{cases}, \quad (2.2)$$

and degree matrix $\mathbf{D} = \text{diag}(\mathbf{A}\mathbf{1})$, where $\mathbf{1} = (1, \dots, 1)$ is the one vector, a diagonal matrix containing the total number of edges associated with a given node.

Using the graph Laplacian, the transport of proteins by diffusion on a graph is given by the network heat equation:

$$\frac{d\mathbf{p}}{dt} = -\rho\mathbf{L}\mathbf{p}, \quad \mathbf{p}(0) = \mathbf{p}_0, \quad (2.3)$$

where $\mathbf{p} \in \mathbb{R}_{\geq 0}^R$ is a dimensionless protein concentration at each node in the brain network, $\rho \in \mathbb{R}_{\geq 0}$ is the transport coefficient, and $\mathbf{L} \in \mathbb{R}^{R \times R}$ is the graph Laplacian.

We can include more information about the connectivity of the brain regions by weighting the adjacency matrix by the number and length of streamlines estimated between regions from tractography. In Putra et al., [86], we discuss three weighting systems based on the number and length of streamlines: length-free, $A_{ij} = n_{ij}$, ballistic, $A_{ij} = n_{ij}/\ell_{ij}$, and diffusive $A_{ij} = n_{ij}/\ell_{ij}^2$, where A_{ij} is the weighted edge connecting v_i and v_j , n_{ij} is the number of streamlines between v_i and v_j , and ℓ_{ij} is the average length of those streamlines for all $i, j = 1, \dots, R$. Each weighting confers different physical assumptions. The transport process should respect two conditions. 1) mass conservation ($\mathbf{1L} = \mathbf{0}$), ensuring that there is no clearance or production through the transport process; 2) Fick's condition ($\mathbf{L1} = \mathbf{0}$), ensuring transport is driven by concentration differences between regions. All of the above weightings satisfy both conditions, however, we show in [86] that length-free weighting provides greatest likelihood of observing Braak staging (see Figure 9 in [86]). Furthermore, length is naturally penalised in geometric graphs since there is a lower probability of long-range connections existing [93]. Furthermore, each weighting confers specific units to the transport parameters, ρ . Length-free weights entail ρ has dimension $1/t$, ballistic weights provide units of speed ℓ/t for ρ , and diffusive weights have units of diffusion, ℓ^2/t for ρ . Throughout this work, we will use graph Laplacian with length-free weightings and therefore, unless stated otherwise, ρ will have units of $1/t$ with t measured in years. streamlines. Each weightings confer different physical assumptions. All of the above weightings satisfy both mass conservation and Fick's condition, however, we show in Putra et al., 2020 [86] that length-free weighting provides the greatest likelihood of observing Braak staging (see Figure 9 [86]). Furthermore, length is naturally penalised in geometric graphs since there is a lower probability of long-range connections existing [93].

The dynamics given by the network diffusion model with length-free weighted graph Laplacian with $R = 83$ nodes is shown in Figure 2.1, using a seeding concentration of $p_i(0) = 0.5$ in the bilateral entorhinal cortex, $p_i(0) = 0$ everywhere else, and $\rho = 1.0$. These initials conditions are chosen to be representative of the spreading process that occurs during AD, which typically starts in the entorhinal cortex. The network diffusion model has been analysed against data using regression studies [77, 78] and Bayesian inference [75], however, there is a wealth of evidence to suggest that this model is not an accurate description of the prion-like process underlying AD, since it does not account for production coming from the prion-like templating

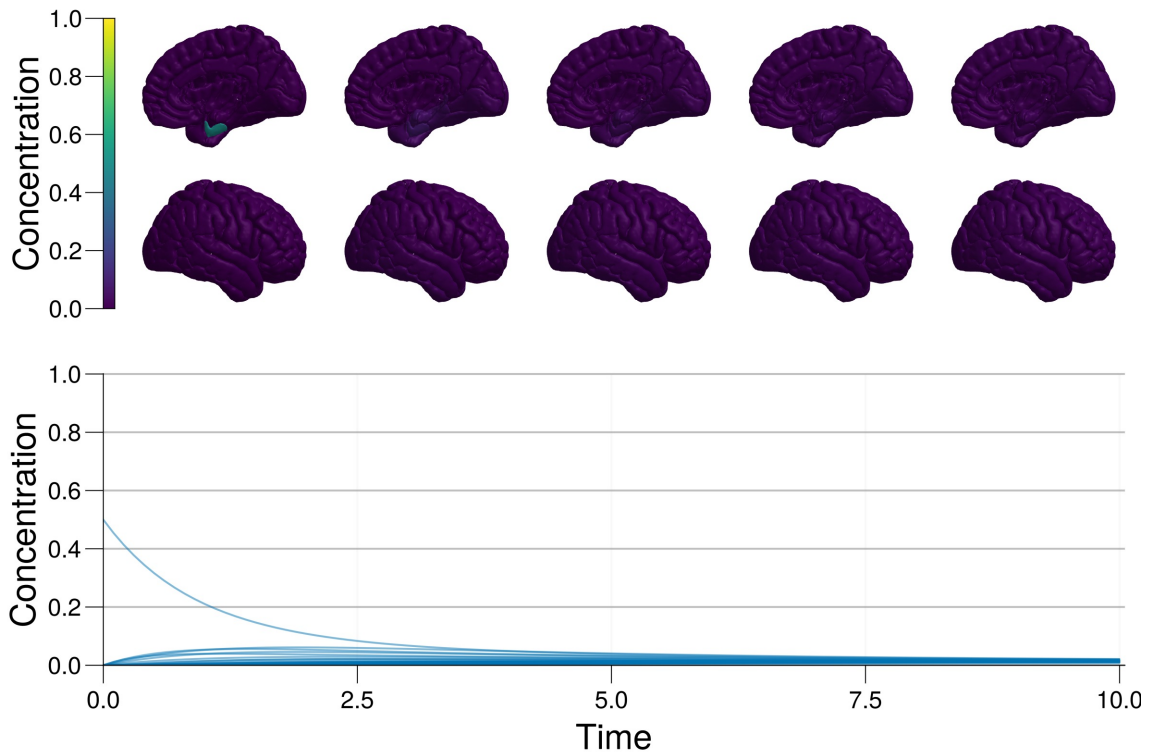


Figure 2.1: **Simulation from the diffusion model.** Equation (3.11), The diffusion model is initialised with bilateral seeding in the entorhinal cortex, $p_i(0) = 0.5$ for $i \in \{27, 63\}$, $p_i(0) = 0.0$ elsewhere, and $\rho = 1.0$.

process [36, 82]. This can be clearly seen in Figure 2.1, where the concentration uniformly approaches a steady state concentration of $(\sum_i p_i(0))/R$, the geometric mean of the initial conditions on R nodes. In the next section we will describe different models that account for protein production.

2.2.2 Prion-like Proliferation

To augment the network heat equation with appropriate production and clearance term of proteins, consider the kinetic diagram shown in Figure 2.2, describing Pruisner’s model of prion-like propagation through templating [94, 95].




Figure 2.2: **Reaction kinetics of the heterodimer model.** Adapted from [82].

There is a healthy protein concentration, $p \in \mathbb{R}_{\geq 0}$, and toxic protein concentration, $\tilde{p} \in \mathbb{R}_{\geq 0}$. The toxic protein binds the healthy protein with rate $k_{11'} \in \mathbb{R}_{\geq 0}$ and induces a conformational change into the healthy protein, turning it into a toxic protein at rate $k_{1'2'} \in \mathbb{R}_{\geq 0}$, before separating into two toxic monomers with rate $k_{2'2} \in \mathbb{R}_{\geq 0}$. This is collectively summarised as rate $k_{12} \in \mathbb{R}_{\geq 0}$. Both healthy and toxic protein species are cleared with a natural clearance rate of $k_1 \in \mathbb{R}_{\geq 0}$ and $\tilde{k}_1 \in \mathbb{R}_{\geq 0}$, respectively. Since we are only concerned with relative amounts of p and \tilde{p} and in the absence of relevant data, we take the concentration of each species to be unitless. Therefore, rate parameters have dimensions $1/t$, where t is assumed to be in years. This can be expressed for each node $i = 1, \dots, R$ on the brain network and coupled with the diffusion model:

$$\frac{dp_i}{dt} = -\rho \sum_{j=1}^R L_{ij} p_i + k_0 - k_1 p_i - k_{12} p_i \tilde{p}_i, \quad p_i(0) = p_{i,0} \quad (2.4a)$$

$$\frac{d\tilde{p}_i}{dt} = -\rho \sum_{j=1}^R L_{ij} \tilde{p}_i - \tilde{k}_1 \tilde{p}_i + k_{12} p_i \tilde{p}_i, \quad \tilde{p}_i(0) = \tilde{p}_{i,0}. \quad (2.4b)$$

The physical insight about the prion process comes at the cost of an increased number of parameters. For simulation purposes, this does not pose any issues, however, it will prove problematic for inference. The major obstacle we face during inference is that we only have observations of toxic protein species, $\tilde{\mathbf{p}}$, and therefore the system is only partially observed. In conjunction with the temporal sparsity, typically 3-5 observations, and noise of data, it will be challenging to identify the rate parameters. Additionally, in the absence of sufficient data about healthy proteins, we are unable to fix rate parameters to aid inference. Therefore, while the heterodimer model provides a detailed physical model of the prion-like process, it is unsuitable for inference with sparse, partially observed data.

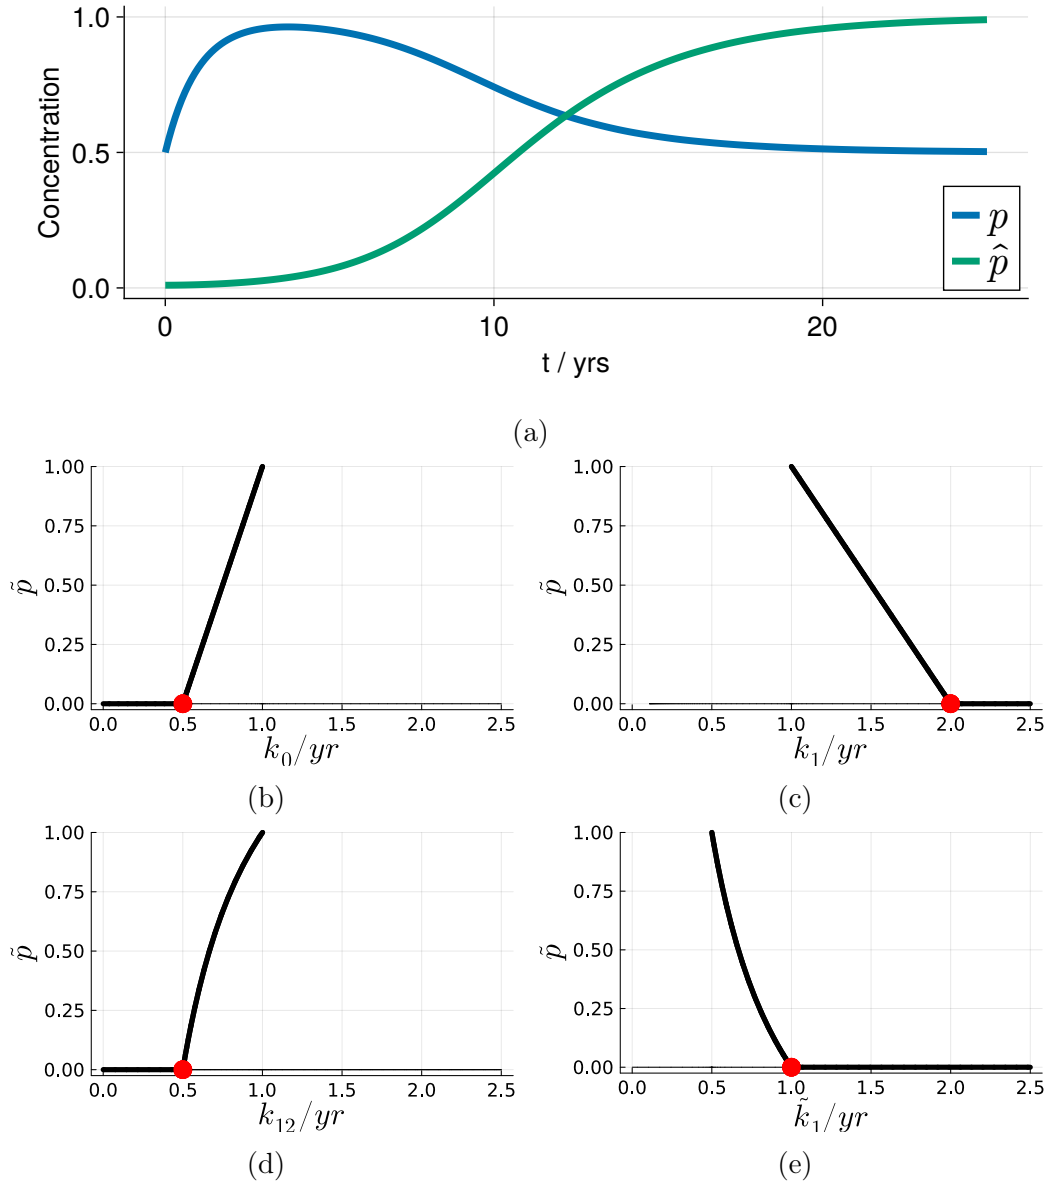


Figure 2.3: **Dynamics of the heterodimer model.** **2.3a** Simulated trajectories from the heterodimer model on a single node and in the absence of transport, using synthetic parameters, $p(0) = 0.5$, $\tilde{p} = 0.01/yr$, $k_0 = 1.0/yr$, $k_1 = 1.0/yr$, $\tilde{k}_1 = 0.5/yr$, $k_{12} = 1.0/yr$. **2.3b–2.3e**, bifurcation diagrams for each parameter with respect to toxic protein \tilde{p} . Thick lines represent stable manifolds, thin lines represent unstable manifolds and bifurcation points are represented by red points.

To address this, we seek to derive a simpler model that captures the same qualitative dynamics. To do so, let us further investigate the dynamics of the heterodimer so as to ensure a reduced order model displays similar dynamics. The heterodimer model on a single node i with health protein p and toxic protein \tilde{p} has two fixed points.

$$p_i = \frac{k_0}{k_1}, \quad \tilde{p}_i = 0, \quad \text{or} \quad p_i = \frac{\tilde{k}_1}{k_{12}}, \quad \tilde{p}_i = \frac{k_0 k_{12} - k_1 \tilde{k}_1}{k_{12} \tilde{k}_1}. \quad (2.5)$$

The first set of fixed points corresponds to a healthy state, with no toxic protein concentration, and the second set of fixed points corresponds to a diseased state. Time series dynamics from the diseased steady state are shown in figure Figure 2.3a where synthetic rate parameters are used and $p_i(0) \gg \tilde{p}_i$ for $i = 1, \dots, R$. The trajectory of the toxic protein concentration is sigmoidal, saturating a fixed point depending the balance of clearance and production rates. Additionally, each of the system parameters acts a bifurcation parameter in a transcritical bifurcation, switching stability between each fixed point, shown in Figures 2.3b to 2.3e. This analysis tells us that if there is sufficient *clearance* of healthy (k_1) or toxic (\tilde{k}_1) protein, the system is stable in a healthy state. In contrast, if there is a sufficient *source* of healthy protein (k_0) or conversion rate (\tilde{k}_1) to toxic protein, the system becomes stable in the disease state. In a reduced order model, we should therefore hope to derive a model that produces sigmoidal dynamics and contains a transcritical bifurcation parameter to control the stability of a healthy and diseased steady state.

To simplify the heterodimer model Equation (2.4) we first assume a healthy, homogeneous state with $\tilde{p}_i \ll p_i$, implying $\frac{dp_i}{dt} \approx 0$ and $-\sum_{j=1}^R L_{ij} p_j \approx 0$ for $i = 1, \dots, R$. Then, using Equation (2.4a), we can rewrite Equation (2.4a)

$$p_i(\tilde{p}_i) = \frac{k_0}{k_1 + k_{12} \tilde{p}_i}, \quad i = 1, \dots, R. \quad (2.6)$$

Linearising around $\tilde{p}_i = 0$, we have

$$p_i(\tilde{p}_i) \approx \frac{k_0}{k_1} \left(1 - \frac{k_{12}}{k_1} \tilde{p}_i \right) \quad i = 1, \dots, R. \quad (2.7)$$

Substituting this expression for p_i into Equation (2.4b) we obtain,

$$\frac{d\tilde{p}_i}{dt} = -\rho \sum_{j=1}^R L_{ij} \tilde{p}_j + \alpha \tilde{p}_i - \beta \tilde{p}_i^2, \quad i = 1, \dots, R, \quad (2.8)$$

where

$$\alpha = \frac{k_0}{k_1} k_{12} - \tilde{k}_1 \quad \text{and} \quad \beta = \frac{k_0 k_{12}^2}{k_1^2}. \quad (2.9)$$

To derive the canonical Fisher-Kolmogorov-Petrovsky-Piskunov (FKPP) model, we make a change of variables $\tilde{p}_i = c_i \alpha \beta^{-1}$, for $i = 1, \dots, R$ giving:

$$\frac{dc_i}{dt} = -\rho \sum_{j=1}^R L_{ij} c_j + \alpha c_i (1 - c_i), \quad c_i(0) = c_{i,0}, \quad i = 1, \dots, R. \quad (2.10)$$

since the variables $c_i(t)$ are unitless concentrations, ρ and α have dimensions of $1/t$ where we take t to be years. The simulated dynamics of the FKPP model, Equation (2.10), are shown in Figure 2.4. The simulation showed has tau seed in the bilateral entorhinal cortex, $c_{i,0} = 0.5$ for $i \in \{27, 63\}$ and $c_{i,0} = 0.0$ zero elsewhere, with $\rho = 0.1/yr$ and $\alpha = 1.2/yr$. These are synthetic parameters to used to illustrate the dynamics in a growth dominated regime, $\rho \ll \alpha$, using initial conditions that are representative of tau deposition in initial AD [41, 6]. During the early stages of AD, the staging of tau across the connectome shares qualitative similarity with observed staging in AD patients and we show in [86] that it is capable of producing Braak staging of tau. The reduced order model captures the same qualitative behaviour as the heterodimer model, showing sigmoidal growth on each node. Additionally, the production term $\alpha c_i(1 - c_i)$ is the normal form of the transcritical bifurcation, with α as the bifurcation parameter. From Equation (2.9), we can further see that the sign of α , and therefore the stability of the system depends on a balance the clearance and production terms of p and \tilde{p} ; $\alpha < 0$ implies $\tilde{k}_1 > k_0 k_{12}/k_1$ and vice-versa, $\alpha > 0$ implies $\tilde{k}_1 < k_0 k_{12}/k_1$. Therefore, the FKPP model is a satisfactory approximation to the heterodimer model that provides the same qualitative features we wished to preserve. In the following section we will test the identifiability of the network FKPP model to confirm that it is suitable for calibration using neuroimaging data.

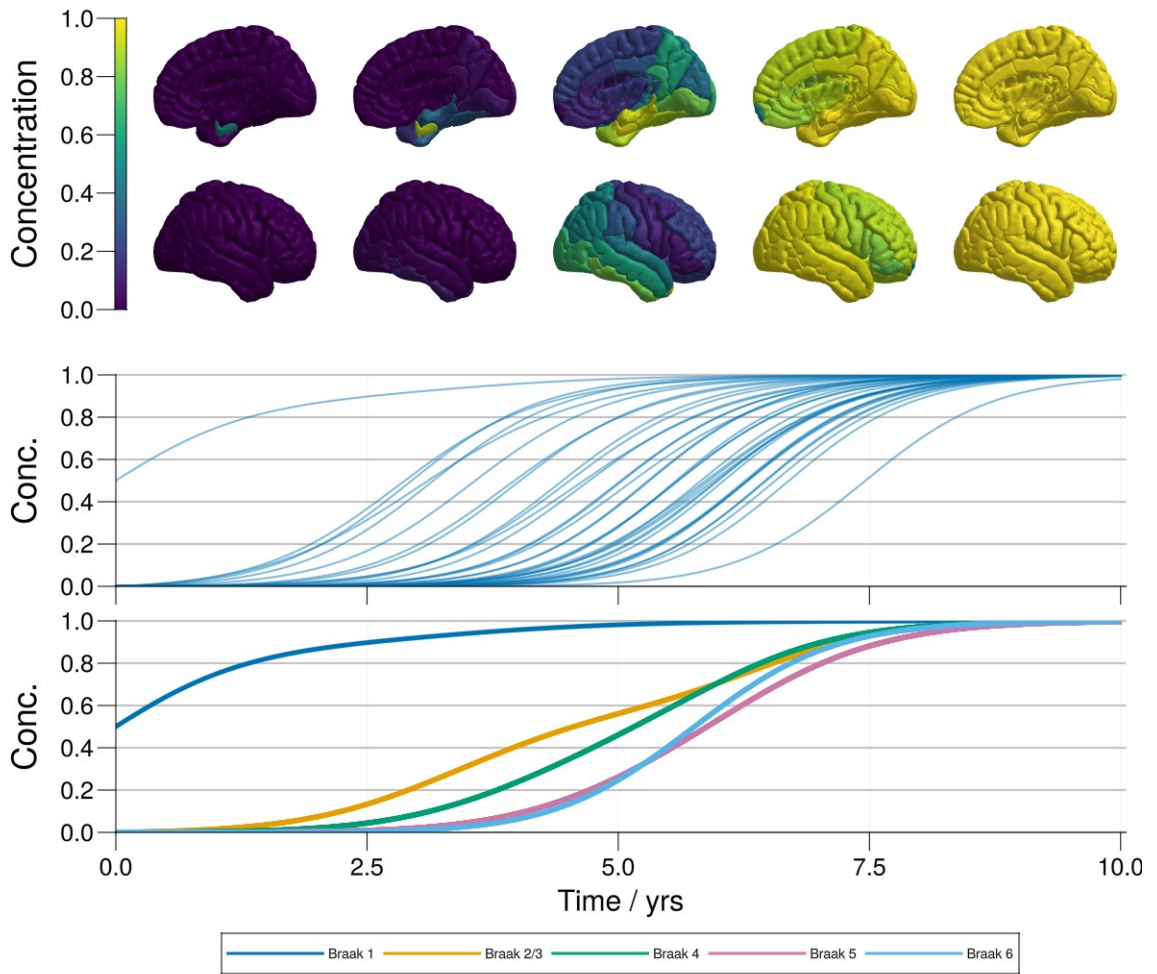


Figure 2.4: **Simulation from the global FKPP model.** Simulation using Equation (2.10) with an initial tau seed in the entorhinal cortex of $c_{i,0} = 0.5$ for $i \in 27, 63$ and $c_{i,0} = 0.0$ zero elsewhere, $\rho = 0.1/yr$ and $\alpha = 1.2/yr$.

2.3 A Bayesian Workflow

So far, we have discussed how different models may account for AD. In practice, we want to fit the models to clinically relevant neuroimaging data and learn about disease dynamics and enable longitudinal prediction. In a system of ODEs, the trajectory is uniquely determined by the initial conditions and the parameters. The process of solving this trajectory is the forward problem, typically involving numerical integration of the ODE system. The inverse problem asks whether the initial conditions and parameters can be determined from partially observed trajectories. While traditional frequentist methods have been used in the literature [78, 72], here, we will explore a Bayesian approach since it offers a better account of identifiability and uncertainty.

After making some observations \mathbf{Y} , we hope to use inference to identify the distributions for unobserved parts of the system, namely model parameters, $\boldsymbol{\theta}$, and parameters of measurement error, $\boldsymbol{\epsilon}$. We do this using Bayesian inference, where the problem corresponds to learning a posterior density $p(\boldsymbol{\theta} \mid \mathbf{Y})$ from a joint probability distribution of the model and data $p(\boldsymbol{\theta}, \mathbf{Y})$. This is achieved by making use of the Bayes' theorem:

$$\underbrace{p(\boldsymbol{\theta} \mid \mathbf{Y})}_{\text{posterior}} = \frac{\overbrace{p(\mathbf{Y} \mid \boldsymbol{\theta})}^{\text{likelihood}} \overbrace{p(\boldsymbol{\theta})}^{\text{prior}}}{\underbrace{p(\mathbf{Y})}_{\text{evidence}}}. \quad (2.11)$$

For even mildly high-dimensional problems, solving for $p(\boldsymbol{\theta} \mid \mathbf{Y})$ becomes intractable using analytic techniques since it relies on performing integration over increasingly large domains to evaluate the evidence term. The evidence term can be reformulated as:

$$p(\mathbf{Y}) = \int p(\mathbf{Y}, \boldsymbol{\theta}) d\boldsymbol{\theta}, \quad (2.12)$$

which shows explicitly that it is necessary to integrate over all of the dimensions of $\boldsymbol{\theta}$ to compute $p(\mathbf{Y})$. Markov chain Monte Carlo (MCMC) methods are a popular class of algorithms used for sampling from the unnormalised posterior distribution, the most common among them being the Metropolis-Hastings (MH) algorithm. In this work, we use Hamiltonian Monte Carlo (HMC), a variation of MH that is particularly useful for fast and accurate inference of high dimensional models with covarying parameters [96, 97]. MCMC methods work by sampling from the unnormalised posterior, denoted $\pi(\boldsymbol{\theta})$, which is proportional to the full posterior:

$$\pi(\boldsymbol{\theta}) = p(\mathbf{Y} \mid \boldsymbol{\theta})p(\boldsymbol{\theta}) \propto \frac{p(\mathbf{Y} \mid \boldsymbol{\theta})p(\boldsymbol{\theta})}{p(\mathbf{Y})} \quad (2.13)$$

To draw samples that capture the posterior, we construct a Markov process with a transition function $Q(\boldsymbol{\theta}' | \boldsymbol{\theta}^t)$, where $\boldsymbol{\theta}^t$ are the accepted parameters at step t , and $\boldsymbol{\theta}'$ are the proposed parameters for $t + 1$. Given certain conditions, this transition function defines a Markov process whose stationary distribution equals the unnormalised posterior. In random walk MH, $Q(\boldsymbol{\theta}' | \boldsymbol{\theta}^t)$ is chosen to be a diagonal multivariate Normal distribution centred around $\boldsymbol{\theta}^t$, $\mathcal{N}(\boldsymbol{\theta}^t, \boldsymbol{\Sigma})$. Proposals parameters are accepted based on the following if

$$u \leq \min \left(1, \frac{\pi(\boldsymbol{\theta}')Q(\boldsymbol{\theta}' | \boldsymbol{\theta}^t)}{\pi(\boldsymbol{\theta}^t)Q(\boldsymbol{\theta}^t | \boldsymbol{\theta}')} \right), \quad (2.14)$$

where u is sampled from a uniform distribution, $u \sim \mathcal{U}(0, 1)$, and rejected otherwise. This Markov process asymptotically converges to the true posterior distribution as the number of samples tends to infinity. The transition function used in random walk MH imbues some notable deficiencies that limit their utility in tackling high-dimensional problems. In high-dimensional spaces, the probability that random proposals will be accepted decreases substantially and therefore the sampler cannot efficiently explore posterior space. This can lead to unfeasibly slow convergence times. HMC obviates this problem by using a more sophisticated transition function to make proposals. In HMC the posterior is represented as a Hamiltonian function and proposals are generated by simulating the dynamics associated with the Hamiltonian. By doing so the transition function utilises the geometry of posterior space and the gradients therein to quickly find areas of high probability mass and generate samples that have a high probability of being accepted [98, 97]. This is true even in high-dimensional spaces with highly covarying parameters, making it an attractive alternative to MH. Throughout this work we use a No-U-Turn-Sampler (NUTS), a HMC sampler which adaptively finds the optimal step size with which to simulate the Hamiltonian [96]. In the next section we will provide an example of the Bayesian framework can be used to assess parameter identifiability and uncertainty.

2.4 Inference, Identifiability and Uncertainty

To make this more concrete, let us consider an example of probabilistic modelling process using the network FKPP model to generate data. The aims of this workflow are:

1. Ensure the model is *structurally* and *practically* identifiable.

2. Given some partially observed longitudinal data, infer the distribution of model parameters θ .

First, we assess whether the model is structurally identifiable. Structural identifiability can be categorised in two ways, *global* and *local* identifiability. If a system is locally identifiable, it means parameters can be identified from observable outputs up to *finitely many* values. Whereas global identifiability means parameters can be *uniquely* determined from observable outputs. To compute identifiability we use the software package `StructuralIdentifiability.jl`, implemented in the Julia programming language [99]. The package relies on computing polynomial approximations of model inputs in terms of observable outputs. For global identifiability, this amounts to computing a field of rational polynomial generators for input-output relations [99] and for local identifiability it requires power series approximations to calculate the observability matrix [100, 101]. These methods are not exact and both utilise randomisation for efficient computation. In global identifiability analysis, randomisation is used to determine field membership for rational polynomial functions, whereas in local identifiability, it is used in the context of numerical linear algebra to approximate the rank of the observability matrix and therefore the uniqueness of input-output relationships. Both randomisation methods allow for uncertainty quantification for correctness and for all tests here, we set a probability of correctness threshold of $p = 0.99$. To ensure models are structurally and practically identifiable it is important to analyse the system in a variety of scenarios, ranging from idealised cases of observability and data availability, to real-world examples, in which limited data may be available. To this end, we perform a number of tests for structural and practical identifiability, varying the observability of nodes, network structure, the availability of data and the level of observation noise. If model parameters are found to be either structurally or practically non-identifiable, we may seek to define reparameterised models in which non-identifiable parameters are either eliminated, fixed or combined to make an identifiable model [101].

Using `StructuralIdentifiability.jl` we can assess the local and global identifiability of our model Equation (2.10). The determination of global identifiability scales poorly with model complexity and we are therefore prohibited from calculating global identifiability on the full brain network. We therefore check the global identifiability of parameters on a subnetwork containing only nodes in the temporal lobe, $R = 22$. This particular subnetwork is chosen since it is the most common location from which Alzheimer’s pathology originates and spreads, with many who experience AD having tau pathology that is largely restricted to the temporal lobe. The analysis shows that

ρ and α are globally identifiable on the restricted temporal network with a correctness probability of 0.99. The same computational constraints do not apply to local identifiability analysis and we therefore conduct a more thorough investigation of local identifiability. We once again restrict our attention to a subnetwork of the full brain network, this time comprising regions in the right hemisphere with $R = 36$ regions, since inter-hemispheric connections are sparse (for example, see Figure 1.4) and AD pathology is bilateral. We seek to determine how network topology and partial observability effect local identifiability. We test four cases of observability: 1) all regions are observable, $R = 36$; 2) regions in the right temporal lobe are observable, $R = 11$; 3) only the right entorhinal cortex is observable, $R = 1$; 4) only the right inferior temporal lobe is observable, $R = 1$. For each of these four cases, we vary the cutoff level with which edges in the connectome are filtered. For example, a cutoff of 0.01 means edges with a normalised density below 0.01 are set to 0. As the cutoff level increases, as does the sparsity of the graph. Therefore, varying this parameter allows us to vary the connectedness of the connectome. The results for of these experiments are shown in Figures 2.5 to 2.8. With all nodes observable, the initial conditions and model parameters are all locally identifiable invariably of filter cutoff level. For only temporal regions observable we see that several initial conditions become *unidentifiable* as the filter cutoff increases, however, the model parameters remain identifiable. We note that the initial conditions for the temporal regions remain identifiable and regions that become unidentifiable are those which become disconnected from the observed temporal subnetwork. When only the entorhinal cortex is observed, all but the initial condition in the entorhinal cortex and the production parameter, α , remain identifiable with increasing filter cutoff, while the initial conditions for unobserved nodes and the transport parameter ρ are unidentifiable. It is particularly noticeable that majority of regions and the transport parameters become unidentifiable at filter cutoff value of 0.28, since this is the level at which the observed node, the entorhinal cortex, becomes disconnected from the network. To test whether the identifiability of the transport parameter depends on connectedness of the observable node, we test the identifiability with varying filter cutoffs for the right inferior temporal lobe which remains connected across the filter cutoff range. In this case, shown in Figure 2.8, the initial conditions for the inferior temporal lobe, connected regions and both production and transport parameters are locally identifiable.

From this analysis we can see how network topology effects the local structural identifiability of model state variables, comprising the initial conditions and model

parameters. In particular, we can see that initial condition for node i remains identifiable if it observable or is connected to an observable node. The production parameter, α , is identifiable in all cases and the transport parameter remains identifiable if the observable nodes are not disconnected. In Section 2.5 and in later chapters, where models are applied to data, we will assume fixed initial conditions based on initial observations to reduce the the computational cost of inference, therefore the identifiability of initial conditions does not pose a problem for us. Furthermore, we have access to observations across all regions of the cortex, therefore, the transport parameter will be structurally identifiable.

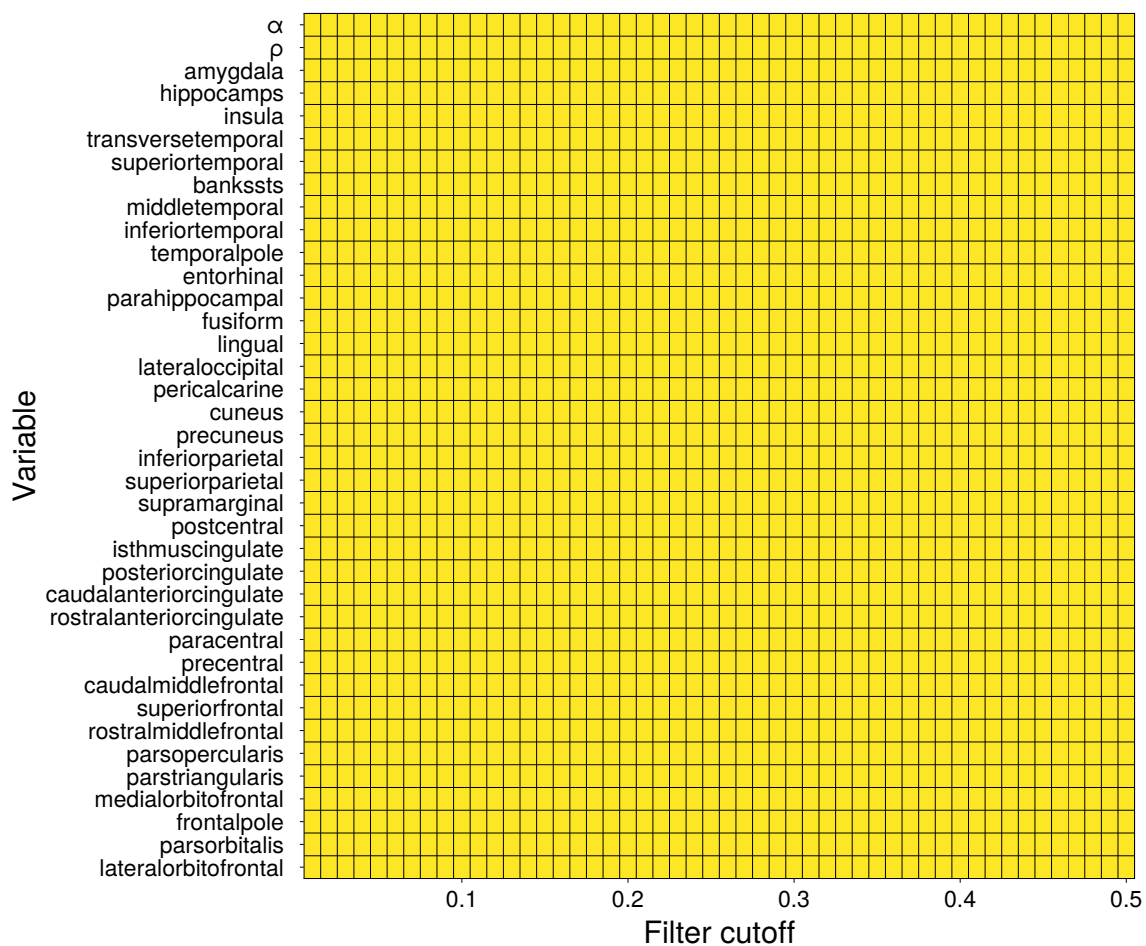


Figure 2.5: **Local structural identifiability with all regions observable.** Local identifiability is calculating using StructuralIdentifiability.jl on the right cortical hemisphere with $R = 36$ regions, for initial conditions $c_i(0)$ for $i = 1, \dots, R$ and model parameters ρ and α .




Figure 2.6: **Local structural identifiability with temporal regions observable.** Local identifiability is calculated using StructuralIdentifiability.jl on the right cortical hemisphere with $R = 11$ regions, for initial conditions $c_i(0)$ for $i = 1, \dots, R$ and model parameters ρ and α .




Figure 2.7: **Local structural identifiability with entorhinal cortex observable.** Local identifiability is calculating using StructuralIdentifiability.jl on the right cortical hemisphere with $R = 36$ regions, for initial conditions $c_i(0)$ for $i = 27$ and model parameters ρ and α .




Figure 2.8: **Local structural identifiability with inferior temporal lobe observable.** Local identifiability is calculated using StructuralIdentifiability.jl on the right cortical hemisphere with $R = 1$ regions, for initial conditions $c_i(0)$ for $i = 29$ and model parameters ρ and α .

We next examine the practical identifiability of the model. Practical identifiability has been defined in many ways. Practical identifiability is typically defined in terms of the model likelihood, where a model is said to be practically identifiable if the confidence intervals of the likelihood are finite [102, 103]. Similarly, in the Bayesian case, a parameter can be said to be practically identifiable if its posterior distribution has finite variance and there is significant movement away from the prior distribution, representing information gained from data [102, 104, 105]. In Section 2.5, we will apply the network FKPP model to patient tau PET data which is scaled to the $[0, 1]$ interval to correspond to the FKPP model output, therefore, in this section, we will probe specifically the practical identifiability of this model without a specific measurement model.

To assess the practical identifiability of the model we generate synthetic data and use Bayesian inference to assess whether parameters can be inferred from ideal data i.e. from data with low noise and samples at many time points capturing the range of dynamics in the model. We assume the following generative process:

$$\mathbf{Y} = \mathbf{f}(\mathbf{c}_0, \mathbf{t}, \boldsymbol{\theta}) + \boldsymbol{\epsilon}, \quad (2.15)$$

where $\mathbf{Y} \in \mathbb{R}^{R \times T}$ are the data for R nodes and T time points, generated by the function \mathbf{f} , with initial conditions \mathbf{c}_0 , at times \mathbf{t} , parameters $\boldsymbol{\theta}$, and observation noise $\boldsymbol{\epsilon}$. We generate synthetic data we choose \mathbf{f} to be solutions to the network FKPP model Equation (2.10), with an initial protein concentration seeded in the entorhinal cortex, $c_i(0) = 0.5$, and $c_i(0) = 0.0$ elsewhere, $\mathbf{t} = \{1.5n \mid 0 \leq n \leq 10\}$, $\rho = 0.1$ and $\alpha = 1.25$. We add identically distributed Gaussian noise with $\sigma = 0.025$. These parameter values are chosen to represent production dominated spread, $\alpha \gg \rho$, which produces sequential Braak-like staging. We begin with a low level of noise to assess the effect of partial observability on the identifiability of parameters. We focus on varying noise levels in later experiments, see Figure 2.10c. The time series of the solution and the generated data are shown in Figure 2.9a. Our data generating distribution is then:

$$\mathbf{Y} \sim \mathcal{N}(\mathbf{f}(\mathbf{c}_0, \mathbf{t}, \boldsymbol{\theta}, \mathbf{L}), \sigma^2 \mathbf{I}), \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2 \mathbf{I}). \quad (2.16)$$

The likelihood for observing data \mathbf{Y} , assuming the generative model is given by


$$p(\mathbf{Y} \mid \theta, \sigma, \mathbf{c}_0, \mathbf{t}) = \prod_t \prod_i p(Y_{ij} \mid \theta, \sigma, \mathbf{c}_0, t_j), \quad (2.17)$$

and the (unnormalised) posterior is

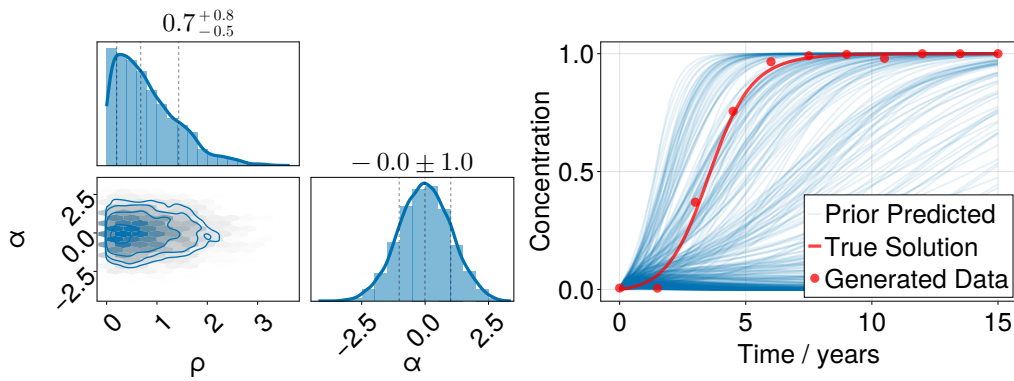
$$p(\theta, \sigma \mid \mathbf{Y}, \mathbf{u}_0, \mathbf{t}) \sim p(\mathbf{Y} \mid \theta, \sigma, \mathbf{c}_0, \mathbf{t}) p(\theta, \sigma), \quad (2.18)$$

The prior model encapsulates our initial beliefs about the data and should generate synthetic data that covers the observed data. The prior distributions and predictive trajectories can be seen in Figure 2.9b, using a standard half-normal distribution for ρ and a standard normal prior for α . These priors are chosen for two reasons. First, they are informative about the bounds particular parameters can take, namely that ρ must be positive while α can be either positive or negative, relating to production or clearance of toxic protein. Second, they produce trajectories that are consistent with domain knowledge, specifically, they produce dynamics over the scale of years. Trajectories of the FKPP model from the prior are shown in blue and sufficiently


cover the data points, here shown only for the right hippocampus. With the data and model in place, we use NUTS to sample from the posterior. The posterior densities can be seen in Figure 2.9c. The true values of the parameters are recovered with very little variation, despite a slight negative correlation between the transport and production parameters. The negative correlation here likely arises from correlations in early transport based seeding of non-seed nodes and production. If transport is faster, non-seed nodes will have a great initial transport driven deposition and therefore will grow faster with a smaller production coefficient. Overall, the results indicate that the FKPP model is identifiable in this idealised case.



(a) FKPP solution and synthetic data



(b) Prior distribution with sample trajectories



(c) Posterior distributions with sample trajectories

Figure 2.9: **Probabilistic modelling pipeline.** **2.9a** Solution to the FKPP model (grey) with initial seeding concentration of $c_{i,0} = 0.1$ in the bilateral entorhinal cortex, $i \in \{27, 63\}$ and $c_{i,0} = 0$ elsewhere, $\rho = 0.1/\text{yr}$ and $\alpha = 1.25/\text{yr}$, and synthetic data (red) generated by taking solution values at every 5 years. **2.9b** Left: Prior distribution for parameters: $\rho \sim \mathcal{N}^+(0, 1)$, $\alpha \sim \mathcal{N}(0, 1)$. Right: FKPP solution trajectories for the right hippocampus using parameters drawn from the prior distributions. **2.9c** Left: Posterior distributions obtained using NUTS. Right: FKPP solutions for the right hippocampus using parameters drawn from the posterior distributions.

In practice, we will not have ideal data since longitudinal PET data is sparse. Therefore, we seek to assess if and how identifiability changes across the disease timeline. To do this, we use two methods, (1) local sensitivity analysis and (2) inference on synthetic data saved at successive intervals of 3 time points, the median number of longitudinal tau PET scans available from ADNI. The results of this analysis are shown in Figure 2.10. Local sensitivity analysis is used to describe the effect of small perturbations to parameters on the output of the function. In this case, the output is the solution to our ODE system, the network FKPP model, Equation (2.10) and the parameters of interest are the transport coefficient, ρ , and the production rate, α . The sensitivities of interest are the partial derivatives of the output with respect to these parameters, $\partial_\rho c_i(t)$ and $\partial_\alpha c_i(t)$ for $i = 1, \dots, R$, respectively. The sensitivities are shown in Figure 2.10a, given the solution in Figure 2.9a. Higher sensitivities tell us that the output is more sensitive to the input parameter, or conversely that regions of high sensitivity convey more information about the parameter. The sensitivity behaviour shows the solutions are increasingly affected by input parameters as tau spreads across the network. Sensitivities decrease as regional concentration saturates, at which point the solution conveys little to no information about the input parameters, ρ and α . Therefore, we expect to have good practical identifiability of parameters up until most seeds are saturated. We test this by generating five synthetic data sets given by solutions to the ODE model saved at successive time windows $\mathbf{t} \in \{(3n, 3n + 1.5, 3n + 3)\}_{n=0}^4$. These time intervals are chosen to be representative of the intervals between longitudinal tau PET scans in ADNI. For instance, the first dataset are the solutions of the simulation in Figure 2.9a saved at $\mathbf{t} = \{0, 1.5, 3\}$. The results of this inference shown in Figure 2.10b. We see time windows for which the parameters are determined with the least uncertainty are those at the mid-point of the solution. There is higher uncertainty for early time periods, but the mean of the posterior still corresponds to the true parameter value for both ρ and α . However, at later time periods, the parameters become practically unidentifiable, displaying especially broad posteriors, consistent with the sensitivity analysis. We next examine the effect of increasing the level of noise in the synthetic dataset. Taking the solution of the simulation in Figure 2.9a at $\mathbf{t} = 3, 4.5, 6$ years, we generate five synthetic datasets by adding independent and identically distributed Gaussian noise for $\sigma \in 0.025, 0.05, 0.1, 0.25, 0.5$. The results of this analysis are shown in Figure 2.10c. The analysis shows that the practical identifiability of the transport parameter is more sensitive than the production parameter to increasing levels of noise. The transport parameter is non-identifiable for $\sigma \in 0.25, 0.5$, however the production parameter

remains identifiable, but with increased uncertainty. Therefore, for data with high levels of noise, the transport parameter is unlikely to be identifiable. This analysis shows that the network FKPP model is practically identifiable unless the regions are already well saturated or the data contains high levels of noise. In practice, we are unlikely to make tau PET measurements when the brain is near total saturation of tau so it does not present a problem for parameter inference. We will further examine the identifiability of the network FKPP model in the next section by applying the modelling and inference framework to real-world longitudinal patient data from ADNI.

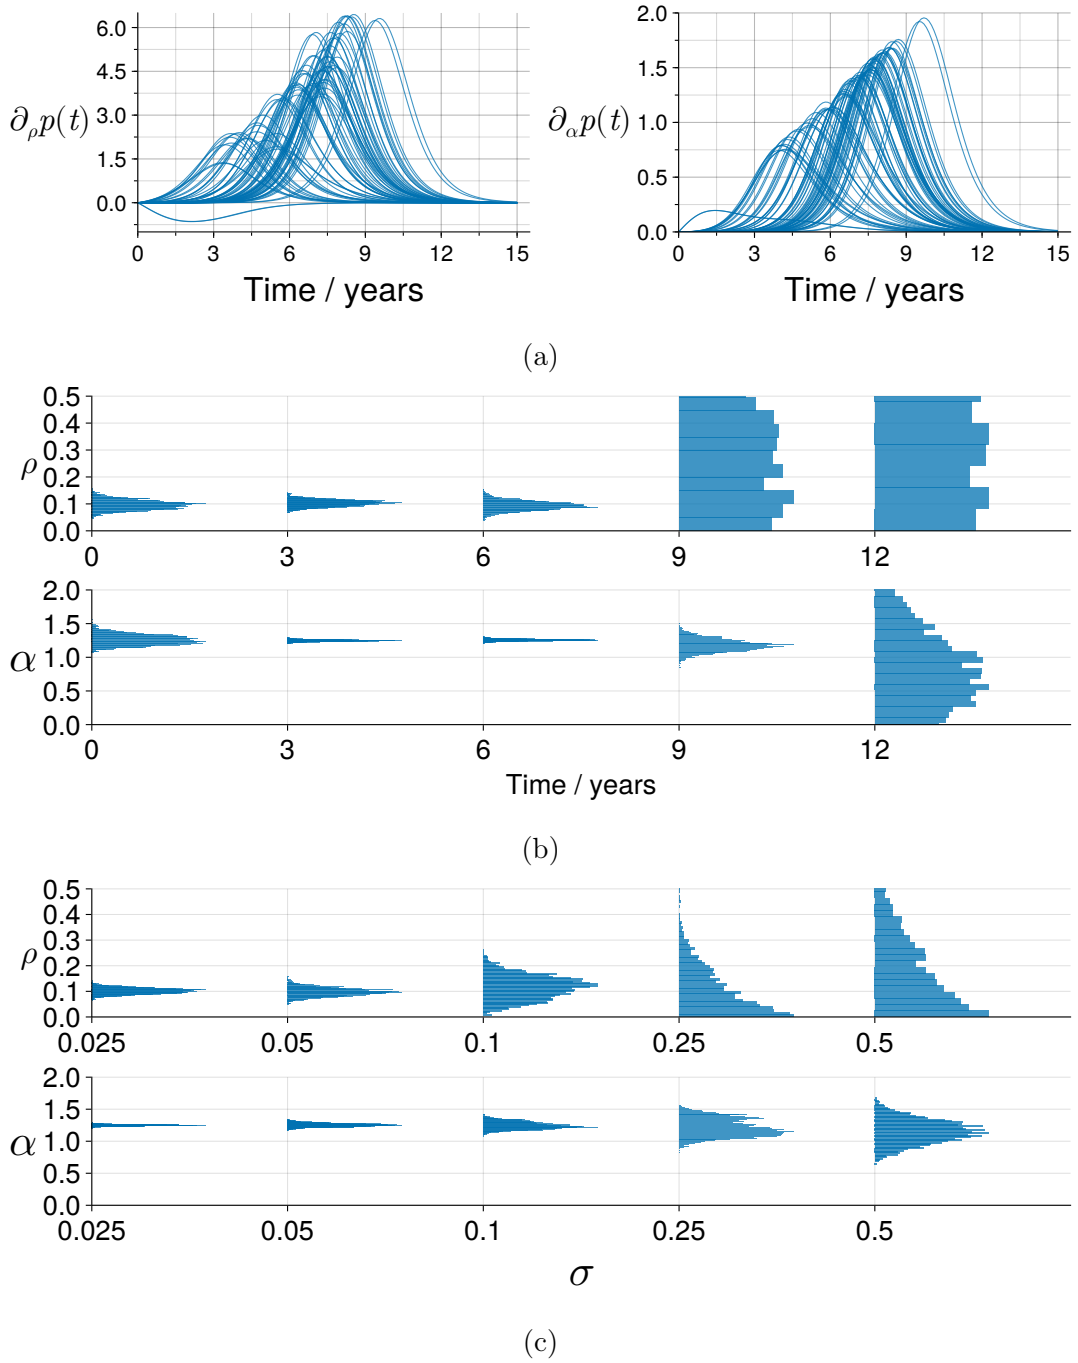


Figure 2.10: **Practical identifiability analysis for the network FKPP model.** **2.10a** Local sensitivities of parameters ρ (top) and α (bottom) from the simulation Figure 2.9a. **2.10b** Posterior densities for ρ and α at each time interval $\mathbf{t} \in \{(3n, 3n + 1.5, 3n + 3)\}_{n=0}^4$ of the synthetic data in Figure 2.9a, aligned with the initial time point for each data set. **2.10c**, Posterior densities for ρ and α for synthetic dataset in Figure 2.9a at $\mathbf{t} = 3, 4.5, 6$ years with varying noise levels of $\sigma \in \{0.025, 0.05, 0.1, 0.25, 0.5\}$. For both Figure 2.10b and Figure 2.10c the true values for transport and production are $\rho = 0.1/\text{year}$ and $\alpha = 1.25/\text{year}$.

2.5 Applications to PET and sMRI Data

In this section, we will demonstrate an application of the modelling and inference framework to tau PET and sMRI data from subjects from ADNI.

Macroscopically, the degeneration and loss of neurons manifests themselves in a loss of gray matter tissue, thinning of the cortex, widening of cortical sulci, and expansion of the lateral ventricles [26, 106, 107, 108]. A multitude of clinical studies have investigated the characteristics of brain atrophy in comparison to healthy aging and longitudinally across advancing stages of Alzheimer’s disease. Three main findings are ubiquitous in the literature: brain atrophy rates are higher in Alzheimer’s patients than in healthy age-matched controls [109, 110, 111], atrophy rates increase over time and with advancing disease [112, 113, 114, 115] but may decrease at late disease stages [109, 116, 114], and regional atrophy is strongly related to regional intensity of tau pathology as indicated by autopsy or on tau PET images [117, 118, 24]. In fact, tau pathology and cortical atrophy seem to follow the same stereotypical spatiotemporal progression [119, 107, 120, 121]: changes are first observed in the medial temporal lobe, with the hippocampus and entorhinal cortex representing the first affected regions [122]. With advancing disease, an increasing number of neocortical regions are affected by tau neurofibrillary tangles and atrophy, initially the lateral temporal lobe, followed by frontal and parietal lobes. The sensorimotor and visual cortices are typically the only areas spared from tau pathology and atrophy up until late disease stages [120].

Here we apply a model of tau and atrophy to a cohort of $N = 61$ subjects from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database, of which $N = 38$ subjects have previously been identified as amyloid positive, A^+ , and $N = 23$ as amyloid negative, A^- [123]. We use a hierarchical model structure for parameter inference to gain insight into group differences depending on amyloid status. Since amyloid is a known precursor for tau pathology and an indicator for progression to Alzheimer’s disease, we hypothesize that tau and atrophy dynamics differ between A^+ and A^- individuals.

2.5.1 Methods

A coupled network model for tau and atrophy

We describe the spatiotemporal dynamics of tau protein misfolding and propagating across the brain using the network FKPP model on a graph with $R = 83$ nodes Equation (2.10). We define the graph Laplacian using an adjacency matrix with

length-free weights, $A_{ij} = n_{ij}$, where n_{ij} denotes the average number of streamlines between regions i and j . To describe the relationship between tau and atrophy, we use a simple linear coupling that saturates as tissue undergoes atrophy,

$$\frac{da_i}{dt} = \eta c_i [1 - a_i], \quad a_i(0) = a_{i,0}, \quad i = 1, \dots, R, \quad (2.19)$$

where $a_i \in [0, 1]$ denotes the local amount of tissue atrophy depending on the local concentration of misfolded tau, $c_i \in [0, 1]$, and a global tau-induced atrophy coefficient $\eta \in \mathbb{R}_{\geq 0}$ [124].

Subject data

We calibrate our coupled tau-atrophy model using longitudinal tau PET and sMRI data from $N = 61$ subjects of the ADNI database [125]. All subjects have undergone between three and five consecutive tau PET scans and have corresponding structural imaging within six months of the tomography scans. On average, longitudinal scans were separated by 1.15 years. Out of the full set, $N = 38$ subjects were previously identified as A^+ , and $N = 23$ as A^- [123]. Table 2.1 summarizes the composition of cognitive diagnoses for the full cohort and each amyloid group. Since reliable tau PET measurements cannot be made from subcortical regions [68, 67], we exclude these from the data. Therefore, for each tau PET and MRI scan for a subject, we have observations for $R_{obs} = 68$ regions. Note the connectome used for the forward model has $R = 83$ regions, inclusive of subcortical regions. Following the analysis in Section 2.4, we expect the model parameters, ρ and α to remain identifiable despite observing only a sub-network of the connectome.

In total, we have data $\mathbf{Y} = \{(\mathbf{C}^n, \mathbf{A}^n)\}_{n=1}^N$ where $\mathbf{C}^n \in \mathbb{R}^{R_{obs} \times T_n}$ are tau PET data and $\mathbf{A}^n \in \mathbb{R}^{R_{obs} \times T_n}$ are atrophy data for the n -th subject over R_{obs} regions and T_n total scans for the n -th subject. We enumerate the observation times for a subject as $\mathbf{t}^n = \{t_1^n, \dots, t_{T_n}^n\}$ for $n = 1 \dots N$. \mathbf{C}_0^n and \mathbf{A}_0^n are the initial scans for the n -th subjects tau PET and atrophy measurements, respectively.

Tau data preparation

Preprocessed Tau AV1451-PET data derivatives were downloaded from ADNI following processing by standard protocols [125, 126] and is available from (adni.loni.usc.edu). Each PET image was co-registered to a corresponding high-resolution T1 weighted magnetic resonance image and segmented into 68 cortical and 15 subcortical regions according to the DKT atlas [127]. The resulting 83 regions align with the

Demographics				
Amyloid status	CN	SMC	MCI	total
A^+	$n = 16$	$n = 9$	$n = 13$	$n = 38$
A^-	$n = 8$	$n = 5$	$n = 10$	$n = 23$
total	$n = 24$	$n = 14$	$n = 23$	$n = 61$

Table 2.1: **Subject demographics.** Demographics by cognitive diagnosis for whole cohort and each amyloid group. Key: CN, cognitively normal; SMC, significant memory concern; MCI, mild cognitive impairment.

nodes of the brain network described in Section 2.5.1. ADNI provides regional SUVR for all considered subjects and time points, normalised to the inferior cerebellum. We map the SUVR into a zero-to-one interval by applying a tau SUVR cutoff of 1.1. For all values beneath this cutoff, tau SUVR is set to 0. The filtered tau SUVR is then normalised by the maximum SUVR value [76], allowing for direct comparison with our model output. For each subject, the initial conditions for the tau field are given by the tau concentration measured in the baseline PET scan, \mathbf{C}_0^n for $n = 1, \dots, N$.

Atrophy data preparation

We use Freesurfer [128] in combination with the Clinica [129] t1-freesurfer-longitudinal pipeline to extract regional volume information from the sMRI images. For every included subject and visit, we compute volume measures for all 83 brain regions contained in our network model. Many of the subjects included in this study underwent a number of study visits at which only structural image data was obtained, before ADNI started to routinely include tau PET. Therefore, we include additional information about regional brain volumes for up to twelve years before the first tau PET baseline scan. For each subject, we use the earliest available structural scan to determine regional reference volumes $\mathcal{V}^n \in \mathbb{R}^R$ to which we normalize the regional volumes of all follow-up visits within each subject $V_{ij}^n = V_{ij}^n / \mathcal{V}_i^n$ for $i = 1, \dots, R$, $j = 1 \dots T_n$ and $n = 1 \dots N$. We define a measure of nodal atrophy as the relative reduction in volume, $A_{ij}^n = 1 - V_{ij}^n$ for $i = 1 \dots R$, $j = 1 \dots T_n$ and $n = 1 \dots N$. For each subject, we set the initial conditions for the atrophy field of our model to the relative atrophy values measured at time of the first tau PET, \mathbf{A}_0^n .

Inference Set-up

We run inference separately for the A^- and A^+ cohorts. For each cohort with N subjects, we have data $\mathbf{Y} = \{(\mathbf{C}^n, \mathbf{A}^n)\}_{n=1}^N$. Each subject has initial conditions $\mathbf{Y}_0 = \{(\mathbf{C}_0^n, \mathbf{A}_0^n)\}_{n=1}^N$ and T_n total scans at time points $\mathbf{t}^n = \{t_1^n, \dots, t_{T_n}^n\}$ for $n = 1, \dots, N$. For inference, we define subject-specific model parameters $\Theta = \{(\rho_n, \alpha_n, \eta_n)\}_{n=1}^N$, for tau transport, tau production and tau induced atrophy, respectively. We also use hierarchical priors to group information across subjects, with population-level parameters $\Omega = \{\rho_\mu, \rho_\sigma, \alpha_\rho, \alpha_\sigma, \eta_\mu, \eta_\sigma\}$, on which each of the individual subject parameters depends. We assume independent and identically distributed Gaussian noise shared across subjects for each modality, $\sigma = \{\sigma_P, \sigma_A\}$. Using this model construction, we compute the posterior distributions for parameters, Θ and Ω , given data, \mathbf{Y} .

$$\mathbf{C}^n \sim \mathcal{N}(\mathbf{f}_P(\mathbf{Y}_0^n, \Theta^n, \mathbf{t}^n), \sigma_P^2 \mathbf{I}), \quad (2.20)$$

$$\mathbf{A}^n \sim \mathcal{N}(\mathbf{f}_A(\mathbf{Y}_0^n, \Theta^n, \mathbf{t}^n), \sigma_A^2 \mathbf{I}). \quad (2.21)$$

$n = 1, \dots, N$ subjects, and where $\mathbf{f}_P(\mathbf{Y}_0^n, \Theta^n, \mathbf{t}^n)$ and $\mathbf{f}_A(\mathbf{Y}_0^n, \Theta^n, \mathbf{t}^n)$ are the solutions to the coupled ODEs, Equation (2.10) and Equation (2.19), respectively. The unnormalised posterior is:

$$p(\Theta, \Omega, \sigma \mid \mathbf{Y}) \propto p(\mathbf{Y} \mid \Theta) p(\Theta \mid \Omega) p(\sigma, \Omega). \quad (2.22)$$

Here, $p(\mathbf{Y} \mid \Theta) p(\Theta \mid \Omega)$ denotes the likelihood, $p(\sigma, \Omega)$ denotes the prior distribution for noise and population-level parameters. The hierarchical structure allows us to gain personalised posterior distributions while simultaneously accounting for commonalities between subjects [130]. We assume a different set of hyperparameter depending on amyloid status, allowing us to account for potential differences in tau and atrophy dynamics between amyloid groups, as well as likely similarities within amyloid groups. We select informative priors for our hyperparameters Ω based on previous results [124, 76]. The full list of priors is summarized in Table 2.2. Priors relating to population-level and individual-level transport and production parameters are informative about the parameter ranges and scales. Specifically, we expect ρ to be positive, whereas α can be negative or positive, representing clearance or production. We use half-Normal priors for population-level scale parameters to allow for partial pooling across subjects and to constrain individual variability to the expected parameter scales. Additionally, we expect both parameters to be on the unit scale, providing dynamics on the scale of years.

Parameter	Prior distribution	Support
ρ_μ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
ρ_σ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
α_μ	$\mathcal{N}(0, 1)$	$[-\infty, \infty]$
α_σ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
η_μ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
η_σ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
ρ_i	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
α_i	$\mathcal{N}(0, 1)$	$[-\infty, \infty]$
η_i	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
σ_t	$\Gamma^-(2, 3)$	$[0, \infty]$
σ_a	$\Gamma^-(2, 3)$	$[0, \infty]$

Table 2.2: **Prior distributions.** Prior distributions for the personalised model parameters and corresponding hyperparameters, and the noise associated with tau and atrophy data.

To perform inference we solve Equation (2.10) and Equation (2.19) using the DifferentialEquations.jl library [131] and perform inference using the Turing.jl probabilistic programming library [132]. We use the a No-U-Turn-Sampler [96] to sample four chains each with 1000 tuning samples and 2000 posterior samples per chain and a target acceptance rate of 0.8 for NUTS.

After inference, we simulate posterior predictive curves for a subset of subjects to examine the model fit to longitudinal biomarker data. Specifically, we select those subjects for which a linear regression on the globally averaged tau and atrophy data indicates positive slopes for both tau and atrophy dynamics. For these subjects, the data aligns with our model assumptions of increasing tau and atrophy over time, thus we expect optimal conditions for our model performance. We propagate the uncertainty from the personalised posterior parameter distributions through the model to create posterior predictions of global tau and atrophy dynamics, including credible intervals, that can be compared to the observed data.

2.5.2 Results


Atrophy Follows Tau Progression

Figure 2.11a illustrates regional misfolded tau concentrations and amounts of atrophy after averaging across all subjects and longitudinal scans in each amyloid group. The processed image data confirm several expected trends. We observe higher tau and atrophy values in the $n = 38$ A^+ subjects than in the $n = 23$ A^- subjects. In the A^+ group, there is a strong topographic relationship between elevated concentrations of misfolded tau and elevated atrophy. Figure 2.11b summarizes the regionally averaged atrophy and highlights differences in atrophy dynamics between regions and amyloid status. The results show notably higher atrophy in subjects with positive amyloid status than in those with negative amyloid status, with a significant difference in the hippocampus ($p = 0.015$). The plot also illustrates that the atrophy in our data roughly follows the spatiotemporal evolution described in the literature for tau and atrophy [26, 24]. Atrophy rates are highest in the temporal lobe, with especially pronounced rates in the hippocampus and entorhinal cortex. Consistent with the pattern of spatiotemporal disease progression, regions of the frontal lobe exhibit lower atrophy rates than temporal regions, closely followed by parietal and occipital regions, mirroring the progression of tau.

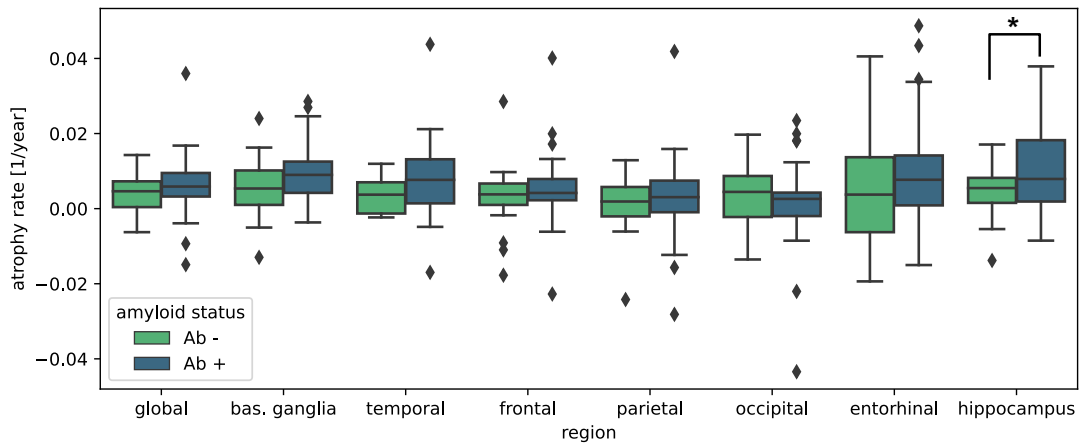
Posterior Distributions Reveal Differences in tau and Atrophy Dynamics

The inference results in converged posterior distributions for all three model parameters at the group and subject levels. All population-level and individual-level posteriors have a high effective sample size, show strong mixing, measured by $0.99 \leq \hat{r} \leq 1.01$, and significant movement away from the weakly informative priors.

Figures 2.12a to 2.12c and Table 2.3 summarize the resulting posterior distributions. There are no significant differences in the posterior distributions for the group mean transport coefficient ρ_μ between amyloid groups, with both distributions being concentrated close to zero. This trend is similarly reflected in the similarity between the profiles of the individual posterior distributions in Figure 2.12a. In contrast, the posterior distributions for the group level growth coefficient α_μ exhibit clear differences between amyloid groups. Overall, A^+ subjects exhibit a significantly ($p=0.0036$) higher tau protein growth rate than A^- subjects, with a population-level mean of 0.016/year and individual distribution means ranging from -0.637/year to 0.446/year for the A^+ group, and a population-level mean of -0.204/year and individual distribution means ranging from -0.846/year to 0.313/year for the A^- group.



(a)



(b)

Figure 2.11: **Regional tau and atrophy for A^+ and A^- groups.** **2.11a** Regional tau concentrations and atrophy values averaged across all subjects and visits for A^+ and A^- groups illustrated on a template brain. **2.11b** Median atrophy rates across individuals between visits separated by regions for A^+ and A^- groups. Asterisk indicates significant difference between amyloid groups.

Interestingly, the individual production rate posteriors exhibit two clusters for the A^- group, while the majority of individuals have a negative production rate, there are some subjects with a positive production rate, these likely reflect individuals with primary age related tauopathy [37]. The population-level parameter for the tau-induced atrophy coefficient η shows much overlap for both amyloid groups, there are noticeably more subjects with higher atrophy coefficients in the A^+ group. In fact, a comparison between individual distribution means in an independent t-test shows that atrophy coefficients are significantly higher ($p=0.0048$) for A^+ subjects than for A^- subjects.

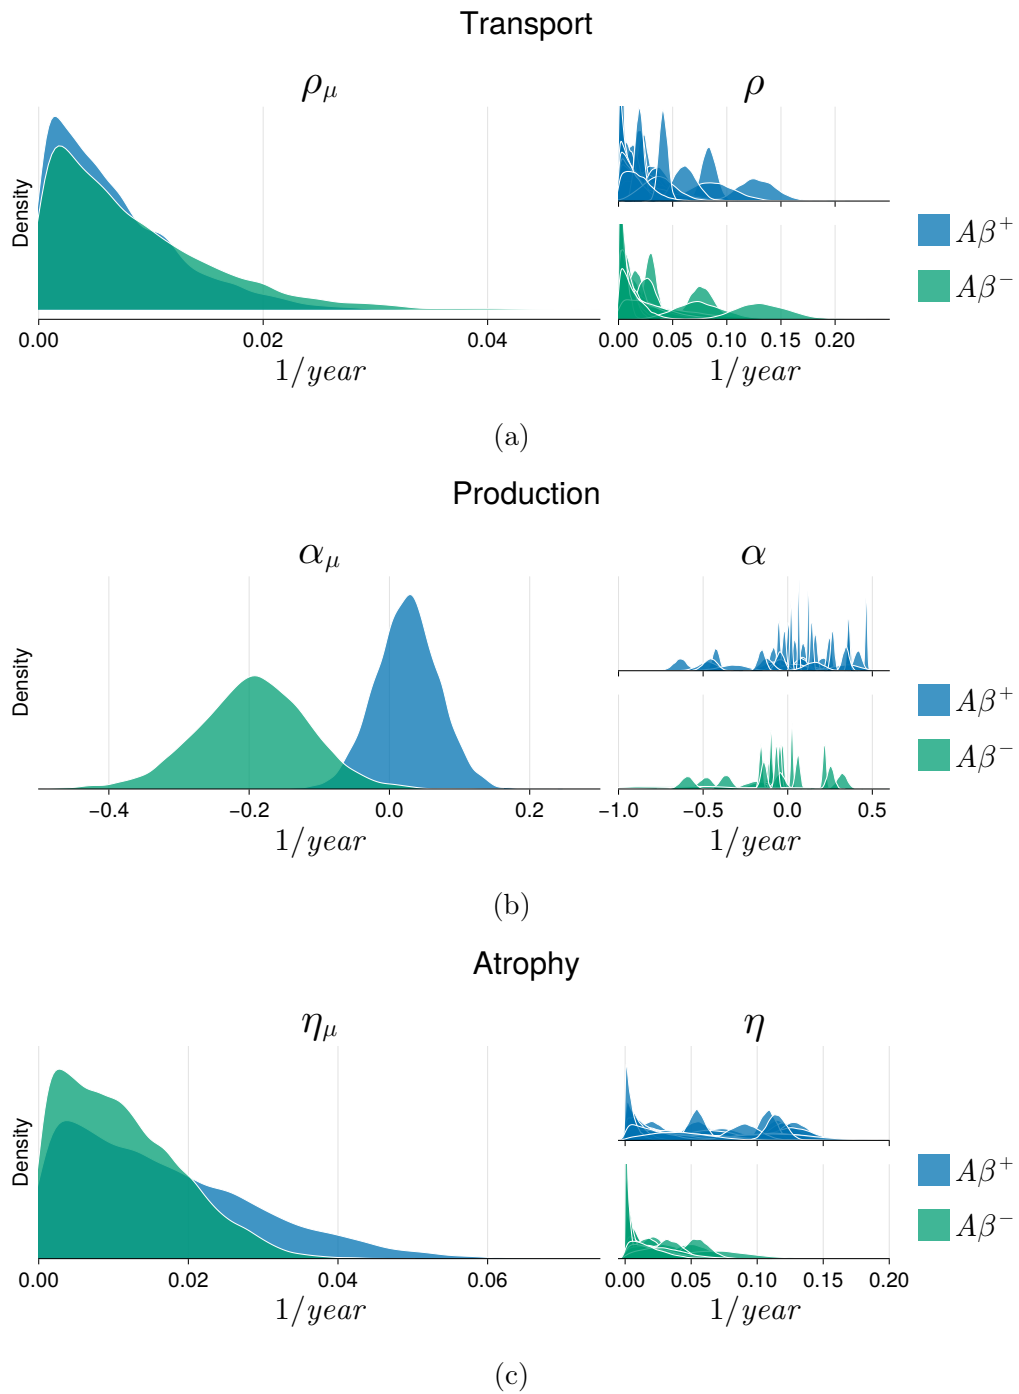


Figure 2.12: **Posterior distributions for model parameters.** Group and subject level posterior distributions for transport coefficient (2.12a), production coefficient (2.12b) and atrophy coefficient (2.12c)

Posterior summary				
	$A\beta^+$		$A\beta^-$	
Parameter	mean	std	mean	std
ρ_μ	0.0184	0.0152	0.0265	0.0232
α_μ	0.0161	0.0440	-0.2042	0.0791
η_μ	0.0165	0.0122	0.0111	0.0080
σ_t	0.0193	0.0002	0.0114	0.0001
σ_a	0.0412	0.0003	0.0300	0.0003

Table 2.3: **Posterior summary.** Mean values and standard deviations for model parameter population-level distributions and noise estimates for tau and atrophy data for A^+ and A^- groups.

Posterior predictive simulations

Our cohort contains $N = 24$ subjects, $N = 21$ A^+ and $N = 3$ A^- , for which both global tau and atrophy data trajectories exhibit an increasing slope. These subjects align with our expectations that tau pathology and tissue atrophy should be increasing in Alzheimer’s disease patients.

Figures 2.13 and 2.14 show our model predictions for tau and atrophy dynamics in the entorhinal cortex during the first four years after baseline tau PET in direct comparison to the observed data. The entorhinal cortex is one of the first regions affected by Alzheimer’s pathology. Overall, the model performs well in capturing the tau and atrophy dynamics in the selected subjects. The model cannot capture non-monotonic trajectories in tau or atrophy, so it performs less well in those cases. Notably, the credible intervals are relatively narrow and do not always capture all data points. There are several potential explanations for this. First, shrinkage due to the hierarchical model structure; the hierarchy allows us to incorporate information across subjects to determine disease dynamics and prevents overfitting to particular subjects. Second, we assume the same noise distribution for each subject. This choice was made to aid parameter identifiability but comes at the cost of decreasing the quality of individual subject fitting. Third, and perhaps most importantly, the model parameters are grouped over nodes and therefore are less sensitive to regional variations production dynamics or noise. To better understand these limitations, we perform a residual analysis of mean posterior fit to the last in-sample scan, shown

in Figure 2.15. The analysis shows that the FKPP model is prone to over or underestimate changes in production depending on region, and atrophy is typically underestimated. This deficiency suggests that the model is incapable of correctly capturing the regional heterogeneity present in the model. Furthermore, we examine how well the posterior noise estimate accounts for observation error. For both tau and atrophy the residuals are symmetric around zero, in-keeping with the Gaussian assumption. The estimate for atrophy measurement error captures the atrophy residuals well, however, the residuals for the tau data are super-gaussian, with a large peak centred at zero and broad tails and are less well captured by the Gaussian error model. The deficiency of the tau error model possibly stems from the model mismatch, with the FKPP model systematically either overestimate or underestimate regional tau burden. Issues relating to regional heterogeneity will be further examined in the next chapter.

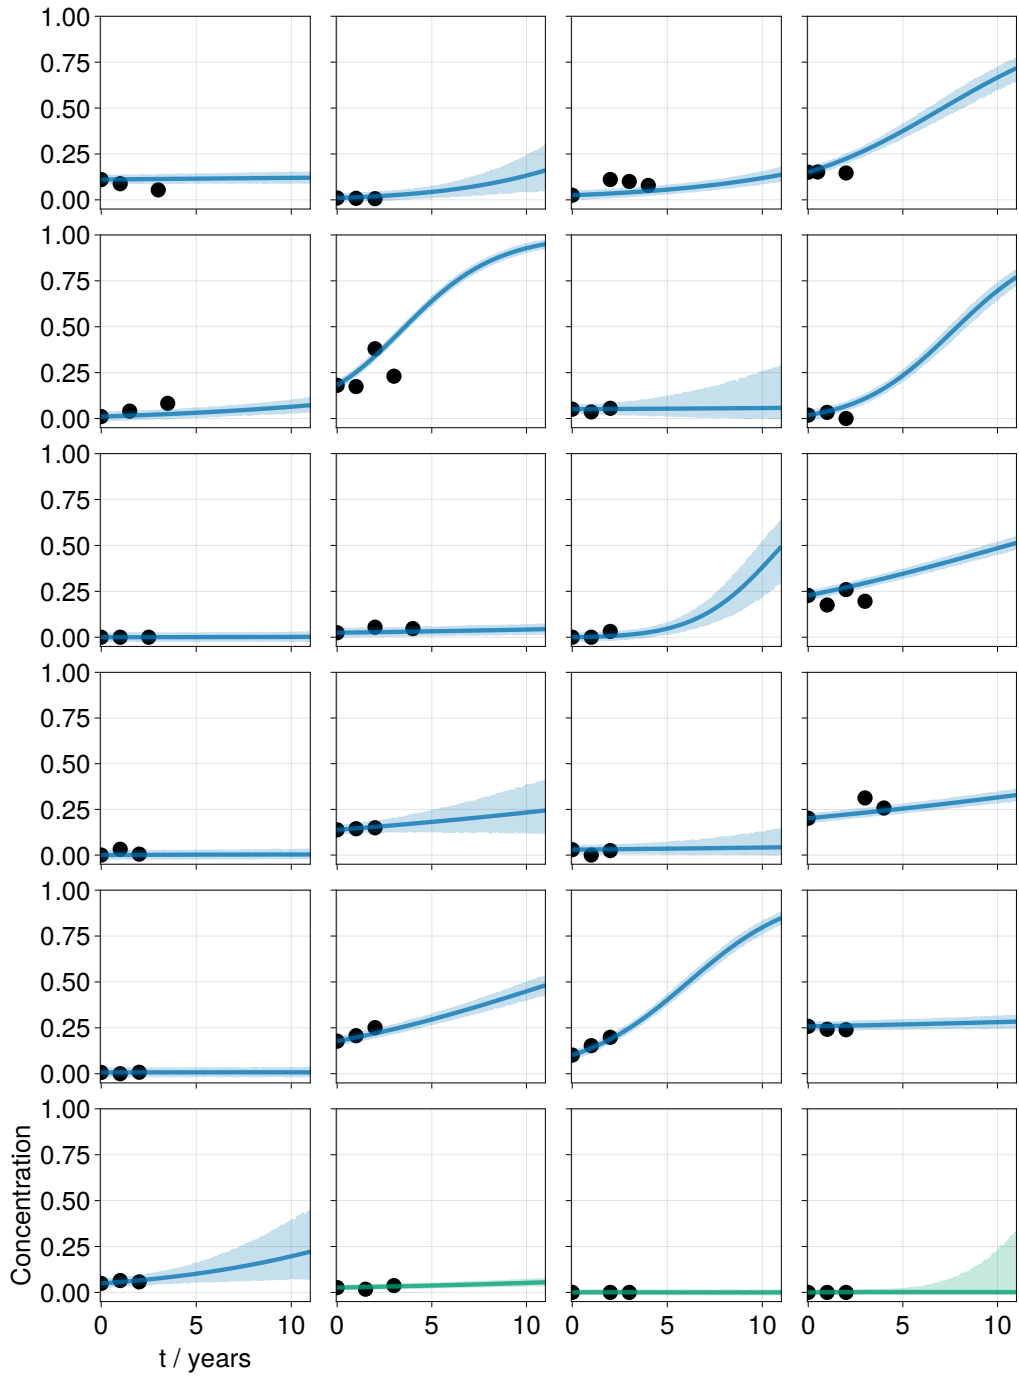


Figure 2.13: **Posterior predictive simulations for tau.** Model predictions for misfolded tau concentration in the entorhinal cortex compared to observed data from longitudinal tau PET. Each subplot represents one subject with circles indicating data points, solid red lines showing the median model prediction, and shaded areas representing the 95% credible intervals. Individuals with blue trajectories correspond to A^+ individuals and green trajectories correspond to A^- individuals.

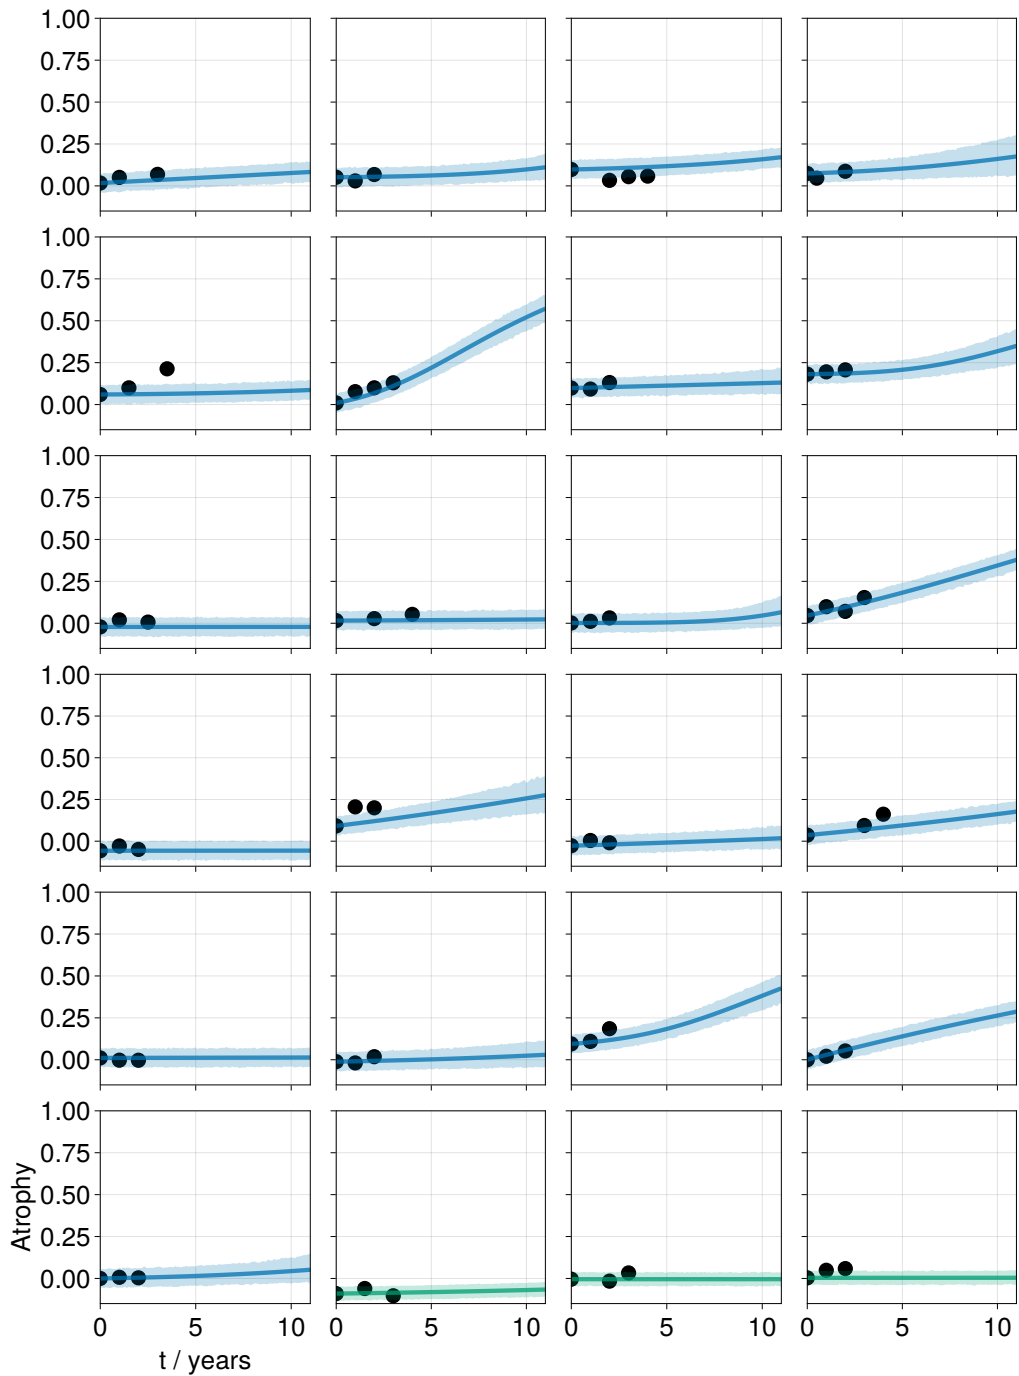


Figure 2.14: **Posterior predictive simulations for atrophy.** Model predictions for atrophy in the entorhinal cortex compared to observed data from longitudinal sMRI imaging. Each subplot represents one subject with circles indicating data points, solid blue lines showing the median model prediction, and shaded areas representing the 95% credible intervals. Individuals with blue trajectories correspond to A^+ individuals and green trajectories correspond to A^- individuals.

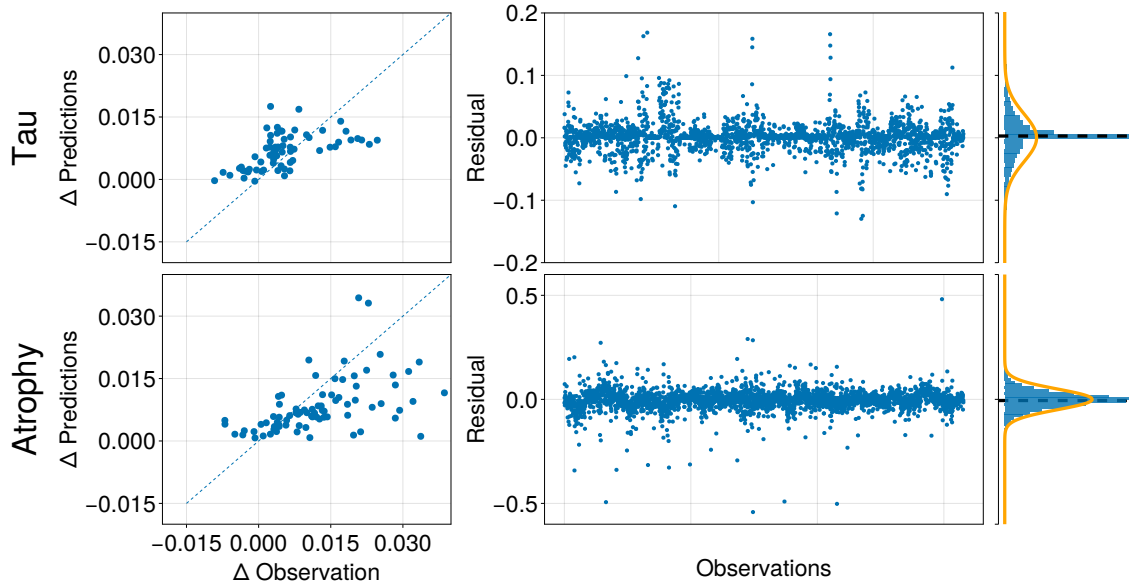


Figure 2.15: **Posterior predictive residual analysis.** Left panels show change in predictions vs change in observation for the last scan, using the first scan as the initial conditions the mean posterior estimate, for tau (top) and atrophy (bottom). Panels on the right show residuals for the tau (top) and atrophy (bottom), using predictions of the final scan based on mean posterior estimates, with histograms of residuals and the estimated mean Gaussian error model.

2.6 Discussion

In this chapter we have developed a framework for modelling and inference in Alzheimer’s disease and demonstrated an application using real tau PET and sMRI data from ADNI.

First, we reviewed some common models used in the literature and analysed the structural and practical identifiability of models. We show that the reduced order network FKPP model provides an accurate qualitative description of the prion-like reaction-diffusion process underlying AD and that it is robustly identifiable from partially observed data. Demonstrating this allows us to use the model for calibration with patient neuroimaging data. Second, we provide an analysis of tau PET and sMRI data from ADNI A^+ and A^- subjects and introduce a coupled model of tau dynamics and atrophy. We show that tau deposition is correlated with atrophy and that atrophy is higher in A^+ groups compared to A^- groups in regions of the brain particularly affected by AD. This brief analysis shows that the data supports the dynamics proposed by the coupled tau and atrophy model, in which atrophy follows tau accumulation. We also show that the inferred personalised parameters can be

used to generate longitudinal predictions at the regional level, including uncertainty estimates.

In this chapter, we developed a coupled model of tau and atrophy informed by clinical observations and personalised its model parameters to multi-modal neuroimaging data of $N = 61$ subjects. We employed a hierarchical Bayesian model for inference, which allowed us to identify significant differences between A^+ and A^- groups for two model parameters, the misfolded tau growth coefficient, α , and the tau-induced atrophy coefficient, η . For the growth coefficient, we identified group-level mean values of 0.0161/year and -0.2042/year for A^+ and A^- groups, respectively. In the context of the FKPP model we chose for describing tau misfolding and propagation, a negative growth rate implies that protein clearance dominates over production and that the model is stable in the healthy state. For the atrophy coefficient, we identified group-level mean values of 0.0165/year and 0.0111/year for A^+ and A^- groups, respectively. These group level differences between A^+ and A^- subjects support our hypothesis that the presence of amyloid plaques in the brain has a magnifying influence on tau and atrophy dynamics. These findings are consistent with other studies observing that the presence of amyloid significantly increases misfolded tau aggregation and tau-induced neuronal loss [38, 39]. We identified a very low transport coefficient independent of amyloid status, consistent with previous findings [76, 133].

Our analysis of the volume data indicates that any current or prospective Alzheimer’s patients in the cohort are in very early stages of the disease, when atrophy is mostly pronounced in the temporal lobe. The atrophy data confirm two main hypotheses. First, atrophy and atrophy rates are more pronounced in A^+ than in A^- subjects. Second, the spatiotemporal progression of atrophy mirrors the known topographic pattern of tau. The lack of representation of more advanced disease stages in our cohort may prevent us from testing our model performance and predictive capacity on more advanced pathology. However, as more longitudinal scans become available for our subjects, we can naturally address this potential limitation.

The posterior predictive simulations for a subset of $N = 24$ subjects show overall good performance of the model in capturing the observed data. By design, the model fails to describe tau or atrophy trajectories that are non-monotonic. Our model also performs weakly when there is a small increase in tau pathology but steep increase in atrophy or vice versa. The linear coupling between tau and atrophy through the coefficient η in our model fails to reflect scenarios in which tau pathology is increasing slowly and linearly, but atrophy is increasing fast and exponentially. This limitation may be resolved by developing more complex atrophy models that allow

for non tau-related avenues of atrophy. We chose a simplistic model here because there is currently not enough longitudinal multi-modal imaging data available to ensure parameter identifiability of more complex models with more parameters during inference.

There are several limitations of the current study. The first limitation of the study presented here are the relatively low numbers of subjects, in particular those with cognitive impairment and AD, This shortage limits the capacity of this work to make inferential claims about mechanisms in AD and how heterogeneity in the A^- and A^+ populations manifest. Second, although we are able to personalise parameters given individual patient data, we are unable to make further personalisations, such as including subject specific connectomes. The use of such connectomes may enhance predictive capabilities stemming from individual differences in connectome topology, however, would significantly increase the computational cost of inference. Limited data availability lead us to make simplifications in our inference methods. We did not attempt to infer subject-specific noise parameters. We also did not include any potential noise in the initial conditions extracted from the baseline images. Both of these simplifications reduce the number of parameters to infer and make our approach a viable compromise between model expressiveness and identifiability. Naturally, the growing amount of subject data in the future will allow us to improve our procedure and increase statistical certainty in our conclusions. To this end, Bayesian methods are the optimal tool to work with continuously updated data, and explore more complex models and assumptions in the future.

This study extends our previous work by personalizing a familiar coupled tau-atrophy model to a larger data set [124]. The model intrinsically captures known features of atrophy including the early acceleration, late deceleration of atrophy rates and the regional heterogeneity of atrophy that closely follows the spatiotemporal pattern of tau neurofibrillary tangle invasion. Extending our data set for model validation to more subjects allows us to confirm our hypothesis that amyloid status affects tau and atrophy dynamics. This is manifested in the distinct group-level posterior distributions for two out of three model parameters. Our Bayesian approach provides personalised model parameters, uncertainties, and model predictions and allows us to characterize the tau- and atrophy-related pathology in single individuals and in groups of A^+ and A^- subjects. A better understanding of the interplay of amyloid-beta, tau, and atrophy, fuelled by the ability to measure these biomarkers in vivo and non-invasively in the living brain, could open doors to advance diagnosis and early treatment in Alzheimer’s disease.

Chapter 3

Regional Modelling of Tau Transport and Production

As discussed in Section 1.2.1, there is growing evidence indicating that the progression of tau in AD depends on two key factors: 1) the transport throughout the structural brain network, and (2) local production of tau. In Chapter 2 we developed a simple model of these dynamics and demonstrated how it can be applied to relevant data. However, the extent to which transport and production each contribute to AD progression and whether their contributions change over time has yet to be determined. In this chapter, we build on the work presented in Chapter 2 to derive a model of tau PET that more accurately predicts longitudinal data and can be used to examine longitudinal changes in production and transport dynamics.

3.1 Introduction

There is now substantial evidence that tau propagation follows a prion-like mechanism, progressively forming toxic oligomeric seeds and neurofibrillary tangles through an autocatalytic production process [2, 36]. The prion-like nature of tau has been demonstrated with transgenic animal models in which cortical injections of tau seeds induces the formation of tau aggregates that grow in concentration over time at the site of injection and surrounding areas [33, 134]. In 2012, studies from Liu et al. [42] and de Calignon et al. [43] showed that transgenic mice overexpressing pathological human tau in the entorhinal cortex exhibit accumulation of tau aggregates and that tau invades axonally connected regions through trans-synaptic transport to form seeds in otherwise healthy regions. The prion-like aggregation and axon-based transport of tau has also been suggested in human post-mortem studies [90] and by in-vivo studies using structural connectome-based models of tau PET capable of reproducing

observed tau aggregation and spreading [72, 77, 78, 76, 79, 81]. In a recent investigation, Meisl et al. [133] analysed multimodal tau data from Braak stage 3 onward and showed that tau production, not transport, is the main contributor of tau progression. However the study does not account for the spatial progression across individual brain regions or estimate dynamics across the full AD progression timeline. It remains unanswered whether there are changes in the rates of tau production and transport over time and whether the balance of these two processes changes along the disease timeline. To address these outstanding questions, we develop a whole-brain model capable of accurately describing longitudinal tau PET data and conduct a multi-cohort study to analyse tau dynamics across the full disease progression timeline.

To answer questions about temporal changes in AD tau dynamics in the human brain requires an accurate and reliable model of longitudinal tau observations. There have been numerous efforts over the past decade to use mathematical models to better understand the spatiotemporal properties of AD pathology, ranging from linear diffusion models of tau [77, 78] to infinite dimensional spatiotemporal models of toxic protein aggregation [83]. Each of these models make different assumptions about the physical mechanisms of tau spread, however, there has not been a unifying effort to rigorously compare commonly used models to identify which are best able to accurately describe longitudinal tau PET observations. In addition, the models currently described in the literature fail to account for regional variations in tau dynamics, which has been shown to influence the progression of tau [135, 136, 137] and are incapable of predicting longitudinal changes at a regional level. Here, we present a novel model that provides a qualitative account of regional vulnerability and its effect on tau progression. Using the pipeline developed in Chapter 2 for performing inference on longitudinal tau modelling, we perform hypothesis-driven model selection on a family of common models from the AD modelling literature, including a new model accounting for regional dynamics. We show that models relying only on network diffusion or homogeneous tau production dynamics are not sufficient to model regional longitudinal data, whereas models accounting for regional variations in tau dynamics are able to accurately model longitudinal tau observations at a regional level.

To determine whether there are changes in the rates of tau transport and production during AD progression, we apply our model to three groups amyloid negative, A^- , amyloid positive tau negative (non-zero tau concentration but below tau positivity threshold, see ??), A^+T^- , and amyloid positive tau positive, A^+T^+ subjects, representing different stages of the AD timeline [138]. We show that the transport of tau is faster in early stage disease (A^+T^-), and that there are primary and secondary

increases in production dynamics of tau along the disease timeline. We validate these results on an independent dataset using a different tau tracer, BioFINDER-2, on which the same results are obtained, further showing that the model and results are robust and generalisable across datasets and choice of tau tracer. Finally, we validate our model by showing that it can *predict* regional rates of tau accumulation over time for individual patients. The combination of these methods provide a novel pipeline for analysing and understanding longitudinal tau data, allowing for us to compare changes in disease dynamics across the AD timeline and predict subject specific, region specific changes in tau over time.

3.2 Results

3.2.1 Deriving A Generative Model Of Tau Dynamics

We seek to derive a model that is tailored to longitudinal tau progression as measured through tau PET. There are two sources of regional variation that need to be accounted for: (1) regional differences in tracer binding and uptake [68, 65], and (2) differential vulnerability to tau invasion and proliferation [135, 136, 137]. We derive this model from the heterodimer model through a similar process to that presented in Chapter 2. Recall that the heterodimer model of healthy tau concentration, $p_i = p_i(t)$, and toxic tau concentration, $\tilde{p}_i = \tilde{p}_i(t)$, for $i = 1, \dots, R$ nodes, and graph Laplacian, \mathbf{L} , is

$$\frac{dp_i}{dt} = -\rho \sum_{j=1}^R L_{ij} p_j + k_0 - k_1 p_i - k_{12} p_i \tilde{p}_i, \quad i = 1, \dots, R, \quad (3.1a)$$

$$\frac{d\tilde{p}_i}{dt} = -\rho \sum_{j=1}^R L_{ij} \tilde{p}_j - \tilde{k}_1 \tilde{p}_i + k_{12} p_i \tilde{p}_i \quad i = 1, \dots, R. \quad (3.1b)$$

and can be linearised around a steady, healthy and homogeneous state to:

$$\frac{d\tilde{p}_i}{dt} = -\rho \sum_{j=1}^R L_{ij} \tilde{p}_j + \beta \tilde{p}_i - \alpha \tilde{p}_i^2, \quad i = 1, \dots, R, \quad (3.2)$$

with

$$\beta = \frac{k_0}{k_1} k_{12} - \tilde{k}_1 \quad \text{and} \quad \alpha = \frac{k_0 k_{12}^2}{k_1^2}.$$

To accommodate the observed non-zero healthy baseline tau PET signal, we make a change of variables $s_i = \tilde{p}_i + s_{0,i}$ for $i = 1, \dots, R$. This change of variable simply states that the observed tau PET signal, $s_i = s_i(t)$, is a sum of a regional toxic protein,

\tilde{p}_i and a baseline healthy baseline of tau PET tracer binding, $s_{0,i}$, for $i = 1, \dots, R$. Incorporating this change of variables, we have:

$$\frac{ds_i}{dt} = -\rho \sum_{j=1}^R L_{ij} (s_j - s_{0,j}) + \beta (s_i - s_{0,i}) - \alpha (s_i - s_{0,i})^2, \quad i = 1, \dots, R. \quad (3.3)$$

$$\beta = \frac{k_0}{k_1} k_{12} - \tilde{k}_1 \text{ and } \alpha = \frac{k_0 k_{12}^2}{k_1^2}.$$

To determine regional carrying capacities, we assume a healthy, homogenous state, therefore:

$$0 = \beta (s_i - s_{0,i}) - \alpha (s_i - s_{0,i})^2 \quad i = 1, \dots, R, \quad (3.4)$$

$$= \alpha (s_i - s_{0,i}) (\beta \alpha^{-1} - (s_i - s_{0,i})) \quad i = 1, \dots, R. \quad (3.5)$$

Then, either $s_i = s_{0,i}$, at the minimum, or s_i reaches its maximum when $s_i = \beta \alpha^{-1} + s_{0,i}$, which we define as a new variable $s_{\infty,i}$, the carrying capacity. However, note that $s_{\infty,i} = \beta \alpha^{-1} + s_{0,i}$ implies that the carrying capacity varies regionally only as a function of baseline values, $s_{0,i}$, and not as any physiological vulnerability to tau invasion. The simplest assumption we can make to reflect regionally varying carrying capacities, given the heterodimer model Equation (3.1), is to introduce a regionally varying clearance rate of toxic protein, $\tilde{k}_1 \rightarrow \tilde{k}_{1,i}$, since it appears only in our in carrying capacity through β but not the uniform growth rate α . Therefore, the new definition for our carrying capacity is $s_{\infty,i} = \beta_i \alpha^{-1} + s_{0,i}$ and the *local FKPP* model is then for $i = 1, \dots, R$

$$\frac{ds_i}{dt} = -\rho \sum_{j=1}^R L_{ij} (s_j - s_{0,j}) + \alpha (s_i - s_{0,i}) [(s_{\infty,i} - s_{0,i}) - (s_i - s_{0,i})], \quad (3.6)$$

$$s_i(0) = s_{i,0}, \quad (3.7)$$

which defines a generative model of tau SUVR, grounded upon a dynamical model of tau that accounts for regional variations in tau PET baseline values, $s_{0,i}$, and carrying capacities, $s_{\infty,i}$. We note, however, that the choice of regionally varying clearance rates is the only the simplest assumption that provides a homogeneous production rate, α . It is likely that all parameters vary regionally, however, we choose to enforce a regionally homogeneous production rate, α , to aid identifiability.

We can also see that Equation (3.6) implies regionally heterogeneous production rates of tau. Introducing a change of variables to $q_i = (s_i - s_{0,i}) / (s_{\infty,i} - s_{0,i})$ for

$i = 1, \dots, R$, then Equation (3.6) becomes a standard (network) FKPP equation:

$$\frac{dq_i}{dt} = -\rho \sum_{j=1}^R L_{ij} q_j + \hat{\alpha}_i q_i (1 - q_i) \quad i = 1, \dots, R, \quad (3.8a)$$

$$q_i(0) = q_{i,0} \quad i = 1, \dots, R. \quad (3.8b)$$

with $\hat{\alpha}_i = \alpha (s_{\infty,i} - s_{0,i})$, Therefore, regions with greater difference between carrying capacity and baseline values will tend toward their carrying capacities at a faster rate than those with a lower difference, making the notion of regional vulnerability more explicit.

To ensure a transport process respects mass conservation across regions of varying volumes, we weight the graph Laplacian by regional volumes,

$$\mathcal{L} = \mathbf{V}^{-1} \mathbf{L} \quad (3.9)$$

where $\mathbf{V} = \text{diag}(\mathbf{v}/v_r)$, and $\mathbf{v} = (v_1, v_2, \dots, v_R)$ is a vector of regional volumes and v_r is a reference region. For a given cohort of N subjects, we define the normalised volume of a region v_i as

$$v_i = \frac{1}{N} \sum_n \frac{v_i^n}{v_r^n}, \quad (3.10)$$

taking v_i^n as the initial volume of the i th region and n th subject and v_r^n is as the maximum initial regional volume for the n th subject. Then \mathbf{v} is the average normalised volume per subject in a cohort

In this new model, we introduce two novel parameter vectors, regional baseline values, \mathbf{s}_0 , and carrying capacities \mathbf{s}_∞ , that represent a healthy state and a late-stage AD state, respectively, which add information about regional variations in production dynamics. We estimate the fixed parameters for \mathbf{s}_0 and \mathbf{s}_∞ using a two component Gaussian mixture model calibrated to cross-sectional regional tau PET SUVR. For regions in which a reliable measure of tau SUVR can be obtained, we expect to see two separate distributions, a T^- distribution capturing the expected normal tau load in a given region, and a T^+ distribution describing the pathological tau load [72]. For subcortical nuclei it is not possible to obtain reliable tau PET signal due to off-target binding [68, 65] and we therefore exclude these regions from our model, leaving a total of $R = 72$ regions, comprising the cortical regions of the DKT atlas, in addition to the hippocampus and amygdala. Since these parameters are estimated from tau PET, they also encode specific features of the tracer, such as regional differences in tracer uptake, on-target binding and off-target binding, and allow us to model tau SUVR directly.

For ADNI, we use the Gaussian mixtures calibrated from the multi-tau cohort of AV1451 PET data in [72], since the same tracer is used and there is much more available data. In Section 3.2.4, we also use tau PET data from BF2 using the RO948 tracer. The Gaussian mixture modelling is reapplied to BF2 data to derive tracer specific parameters. For all regions used and for both ADNI and BF2 cohorts a Gaussian mixture with two components provided a better fit to data compared to a single component, compared using the AIC score. Using the fitted two component Gaussian mixture models, we approximate $s_{0,i}$ as the mean of the T^- distribution for the i -th region and $s_{\infty,i}$ as the 99-th percentile of the T^+ distribution for the i -th region, for $i = 1, \dots, R$. An example of regional Gaussian mixture modelling is shown in Figure 3.1b for the left inferior temporal lobe, showing two separate distributions for the T^- and T^+ components. The carrying capacities, \mathbf{s}_{∞} , for the right hemisphere are shown in Figure 3.1c. The simulated transport and production dynamics are shown in Figure 3.1. We can see in Figure 3.1a that regions progress at different rates, highlighting the heterogeneous vulnerability of each region, and capture the expected Braak staging. The ability to account for these regional variations extends previous models with homogeneous dynamics across regions [82, 81] and provides a picture of tau progression that is more consistent with observed tau PET staging.

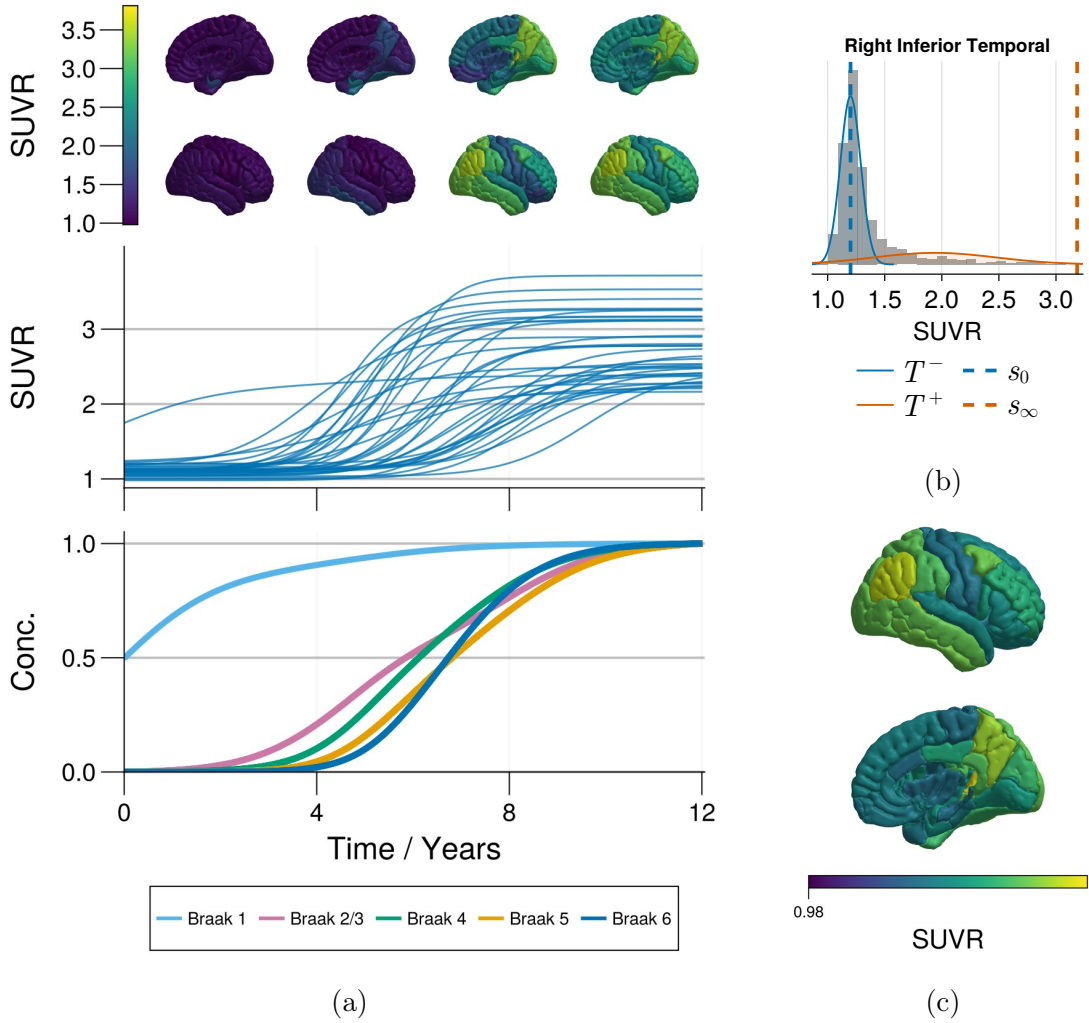


Figure 3.1: Simulated transport and production dynamics in the local FKPP model. **3.1a** Simulation from the local FKPP model using carrying capacities derived from Gaussian mixture models (shown in 3.1b). Simulations are initialised with a seed value of $s_{i,0} = (s_{0,i} + s_{\infty,i})/2$ in the bilateral entorhinal cortex, $i = \{27, 63\}$ and $s_{i,0} = 0$ elsewhere, with $\rho = 0.1/yr$ and $\alpha = 0.8/yr$. Each line in the middle panel represents the SUVR trajectory of one DKT atlas brain region. Values at time points $t \in \{0, 4, 8, 12\}$ years are projected onto a cortical rendering in the top panel. Each line in the bottom panel represents concentration averaged over Braak regions, after rescaling simulated SUVR as $q_i = (s_i - s_{0,i}) / (s_{\infty,i} - s_{0,i})$. **3.1b** Two component Gaussian mixture model fit to a multi-cohort tau PET dataset [72] and data from ADNI for the right inferior temporal lobe. For each node i , the baseline values are taken as the mean of the T^- distribution and the carrying capacity as the 99-th percentile of the T^+ distribution, and these are used to simulate the model in 3.1a and throughout this paper. **3.1c** Right hemisphere cortical rendering of the SUVR carrying capacities as determined through Gaussian mixture modelling.

3.2.2 Regional Heterogeneity Is Necessary for Longitudinal Forecasting

To determine whether the local FKPP is capable of fitting observed AD trajectories, we compare it to tau PET data. For comparison, we also consider three simpler models that can be obtained from the local FKPP model Equation (3.6), namely the *diffusion* model, *global FKPP* model and *logistic* model. The diffusion model is given by assuming there is no toxic protein production, $\alpha = 0$ in the local FKPP model, so that

$$\frac{ds_i}{dt} = -\rho \sum_{j=1}^R L_{ij} (s_j - s_{0,j}), \quad i = 1, \dots, R, \quad (3.11)$$

$$s_i(0) = s_{i,0} \quad i = 1, \dots, R. \quad (3.12)$$

The global FKPP model is given by assuming $s_{0,i}$ and $s_{\infty,i}$ do not vary regionally, then, for $i = 1, \dots, R$

$$\frac{ds_i}{dt} = -\rho \sum_{j=1}^R L_{ij} (s_i - s_0) + \alpha (s_i - s_0) [(s_{\infty} - s_0) - (s_i - s_0)], \quad (3.13)$$

$$s_i(0) = s_{i,0} \quad (3.14)$$

where s_0 and s_{∞} are taken as $\min(\mathbf{s}_0)$ and $\max(\mathbf{s}_{\infty})$, respectively. Finally, the logistic model is obtained by assuming there is no transport between regions, taking $\rho = 0$ in Equation (3.6):

$$\frac{ds_i}{dt} = \alpha (s_i - s_{0,i}) [(s_{\infty,i} - s_{0,i}) - (s_i - s_{0,i})], \quad i = 1, \dots, R, \quad (3.15)$$

$$s_i(0) = s_{i,0} \quad i = 1, \dots, R. \quad (3.16)$$

We use tau PET from ADNI to compare the different models. We download the fully processed tau PET and sMRI data, summarised as SUVR and volumes for each of the cortical regions in the DKT atlas, with the addition of the bilateral hippocampus and amygdala $R = 72$. We re-normalise individual SUVR using an inferior cerebellum SUVR reference region. For model comparison, we select A^+T^+ subjects who have at least three scans. Amyloid status is provided from ADNI and based on $A\beta$ PET SUVR. Tau PET SUVR. Tau status is determined using tau PET SUVR in two composite regions, the medial temporal lobe (MTL, defined as the mean of the bilateral entorhinal and amygdala), and another for neocortical positivity (defined as the middle temporal and inferior temporal gyri). For each composite region, we average the SUVR values from the constituent regions and fit a two component Gaussian

mixture model. The threshold for the composite region is then set to the SUVR at which there is a 50% chance of being T^+ . For ADNI, the thresholds are 1.375 and 1.395 for the MTL and neocortical regions. We define a subject as being T^+ if their last scan is suprathreshold in either the MTL or cortical tau PET SUVR and T^- if the SUVR value is below both SUVR thresholds.

To calibrate the models to longitudinal ADNI data, we use hierarchical Bayesian model. For each of $n \in \{1, \dots, N\}$ A^+T^+ subjects, there are T_n scans, summarised over $R = 72$ regions. The observations times, i.e. scan dates, are denoted by $\mathbf{t}^n = \{t_1^n, \dots, t_{T_n}^n\}$ for $n = 1, \dots, N$. We denote the full data set for a group as \mathbf{Y} and individual subject data as Y_{ij}^n , corresponding to the n th subject, at scan j and region i . For a single subject, we have the following data generating function:

$$\mathbf{Y}^n = \mathbf{f}(\mathbf{y}_0^n, \theta^n, \mathbf{t}^n) + \boldsymbol{\epsilon}. \quad (3.17)$$

where \mathbf{Y}^n is the individual data for R regions and T_n time points, with initial conditions \mathbf{y}_0 , model parameters θ , and observations times \mathbf{t} . The data are generated by the vector function, \mathbf{f} , which are solutions to the dynamical systems for tau propagation, with observation error $\boldsymbol{\epsilon}$. To derive a likelihood function from Equation (3.17), we assume the observations errors are independently and identically distributed and sampled from a Gaussian distribution with standard deviation σ . The data generating distribution for a single observation from a subject is then:

$$\boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2) \quad (3.18)$$

$$\mathbf{Y}^n \sim \mathcal{N}(\mathbf{f}(\mathbf{y}_0^n, \theta, \mathbf{t}^n), \sigma^2 \mathbf{I}) \quad (3.19)$$

To extend this to a hierarchical population model, we define random variables, $\Theta = \{(\rho_i, \alpha_i)\}_{i=1}^N$, encoding subject specific model parameters and hierarchical population parameters, $\Omega = \{\rho_\mu, \rho_\sigma, \alpha_\mu, \alpha_\sigma\}$, upon which each Θ_i depends. In Sections 3.2.2 and 3.2.4, we assume fixed initial conditions, \mathbf{y}_0^n , and observations times, \mathbf{t}^n , taken as the first tau PET scan and scan dates, respectively. Then the likelihood function for a single subject under the hierarchical model is

$$p(\mathbf{Y}^n, \Theta^n \mid \Omega, \sigma, \mathbf{y}_0^n, \mathbf{t}^n) = \prod_j^{T_n} \prod_i^R p(Y_{ij}^n \mid \Theta^n, \sigma, \mathbf{y}_0^n, t_j^n) p(\Theta^n \mid \Omega), \quad (3.20)$$

where the first term inside the product on the right hand side is the contribution of the subject level model and the second term is the hierarchical model. Then the

posterior for all subjects, hierarchical parameters, subject specific parameters and observation noise is

$$p(\Theta, \Omega, \sigma \mid \mathbf{Y}, \mathbf{t}, \mathbf{y}_0) \propto \prod_n^N p(\mathbf{Y}^n, \Theta^n \mid \Omega, \sigma, \mathbf{y}_0^n, \mathbf{t}^n) p(\Omega, \sigma). \quad (3.21)$$

To sample from the posterior distributions, we use a Hamiltonian Monte Carlo No-U-Turn Sampler (NUTS). We use the prior distributions provided in table 3.1. The priors are informative about the support and scale we expect each parameter to take. More specifically, we expect transport to be strictly positive and we choose scales that provide trajectories on the scale of years to match longitudinal tau PET data. For the transport parameter, we use *Lognormal* distributions since they are continuous across their support, facilitating efficient inference, and avoid significant prior mass at zero, emphasising that we expect transport to occur. For population level scale parameters, we test two different parameterisation: Lognormal priors for population-level scale parameters and the observation noise parameter, shown in Table 3.1; and using half-Normal priors for population-level scale parameters and inverse-Gamma priors for the observation noise parameter, provided in Table 3.2. For population-level scale parameters, the former has a wider tail and due to it’s low mass at zero, emphasises variation among individuals, while the latter allows for partial-pooling among subjects and limits variance across parameter scales [139, 140]. In all cases, The NUTS sampler is initialised with a unit diagonal Euclidean metric and a target acceptance ratio of 0.8. For each patient group, we collected four chains each with 2000 samples. All chains showed good convergence (measured by $0.99 < \hat{r} < 1.01$) with no post warm-up numerical errors associated with the NUTS sampler.

Parameter	Prior	Support
ρ_μ	<i>Lognormal</i> (0, 1)	$[0, \infty]$
ρ_σ	<i>Lognormal</i> (0, 1)	$[0, \infty]$
α_μ	$\mathcal{N}(0, 1)$	$[\infty, \infty]$
α_σ	<i>Lognormal</i> (0, 1)	$[0, \infty]$
ρ_i	$\mathcal{N}(\rho_\mu, \rho_\sigma)$	$[0, \infty]$
α_i	$\mathcal{N}(\alpha_\mu, \alpha_\sigma)$	$[-\infty, \infty]$
σ	<i>Lognormal</i> (0, 1)	$[0, \infty]$

Table 3.1: Prior distributions for hierarchical model parameters.

Parameter	Prior	Support
ρ_μ	$Lognormal(0, 1)$	$[0, \infty]$
ρ_σ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
α_μ	$\mathcal{N}(0, 1)$	$[\infty, \infty]$
α_σ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
ρ_i	$\mathcal{N}^+(\rho_\mu, \rho_\sigma)$	$[0, \infty]$
α_i	$\mathcal{N}(\alpha_\mu, \alpha_\sigma)$	$[-\infty, \infty]$
σ	$\Gamma^{-1}(2, 3)$	$[0, \infty]$

Table 3.2: Prior distributions for hierarchical model parameters, using half-Normal priors for population-level scale parameters.


We use two metrics to compare the four models, the maximum log-likelihood and the expected log predictive density (ELPD). The maximum log-likelihood is used to compare the in-sample accuracy of each of the model’s fit to the data using a NUTS sampler. The maximum log-likelihood estimate is taken as the maximum log-likelihood from the posterior samples.

The ELPD is used for estimating the out-of-sample predictive accuracy and is adapted from [85]. To do this, we use A^+T^+ subjects who have more than three tau PET scans ($N = 10$), using only the first three scans for training and remaining scans to measure predictive accuracy. For our model, the ELPD is then calculated as

$$ELPD = \sum_{n=1}^N \left[\log \left(\frac{1}{S} \sum_{s=1}^S p(\mathbf{Y}^n | \Theta_s^n, \mathbf{y}_0^n, \mathbf{t}^n) \right) \right], \quad (3.22)$$


where \mathbf{Y}^n are the unobserved data, $\Theta_s^n = \{(\rho, \alpha)\}_{s=1}^S$ are posterior samples of model parameters, \mathbf{y}_0^n are subjects initial condition and \mathbf{t}^n are scan dates, each for $n = 1 \dots N$ subjects.

Both metrics for each of the four models are provided in Table 3.3. Figure 3.2a shows the in-sample longitudinal fit for each of the four models. The out-of-sample fit for the local FKPP model is shown in Figure 3.2b.



(a) In-sample model fits.

H



(b) Out-of-sample model fits.

Figure 3.2: **Model fit for in-sample and out-of-sample data.** **3.2a** Goodness of fit for four models, local FKPP, global FKPP, diffusion and logistic. For all panels, each point represents a region in the connectome model, averaged over subjects per scan. Top row shows average estimated vs observed regional SUVR values. Bottom row shows average estimated change vs observed change in regional SUVR from the first scan to last scan. Only the local FKPP and logistic model are able to accurately capture longitudinal changes, while the global FKPP and diffusion models are each structurally incapable of describing heterogeneous production. **3.2b** Out-of-sample fits for four models of proteopathy Top row: predicted vs observed out-of-sample SUVR. Bottom row: predicted change vs observed change from first in-sample scan to last out-of-sample scan. Each point represents a region in the connectome model, averaged over subjects.

	Local FKPP	Global FKPP	Diffusion	Logistic
L_{max}	17510.628	16977.588	10554.256	16855.024
ELPD	638.216	622.676	-175.579	607.139

Table 3.3: Assessment of model fit using the maximum log-likelihood L_{max} and the expected log predictive density; higher values correspond to better models. The local FKPP model performs best in both metrics.

	Local FKPP	Global FKPP	Diffusion	Logistic
L_{max}	17510.875	16977.946	10554.256	16855.024
ELPD	647.131	630.237	-172.215	610.033

Table 3.4: Assessment of model fit using the maximum log-likelihood L_{max} and the expected log predictive density; higher values correspond to better models. The local FKPP model performs best in both metrics.

The local FKPP performs best for both in-sample fit and out-of-sample predictive accuracy, followed by the global FKPP and logistic models. The results show that the the diffusion model is ill-equipped for longitudinal modelling of tau PET data, shown clearly in Figure 3.2. The diffusion model underestimates SUVR evolution in time since there is no mechanism for tau production or clearance and therefore total concentration is conserved. The global FKPP model is able to capture the changes in tau load, but it cannot describe the regional heterogeneity in tau production, resulting in a larger residuals, calculated as the difference between regional model predictions and final scans for each subject, mean residual = -0.004 , s.d. of residuals = 0.11 , shown in Figure 3.3. In contrast, the logistic model is able to capture heterogenous tau production dynamics and shows more concentrated residuals compared to the global FKPP, mean residual = -0.02 , s.d. of residual = 0.02 , shown in Figure 3.3. However, the logistic model systematically underpredicts change in regions with low SUVR compared to the local FKPP and global FKPP models, shown in Figure 3.4, suggesting transport may play an important role for initial local tau deposition. These deficiencies are addressed by the local FKPP, which incorporates both transport of tau and heterogenous local tau production, producing small and concentrated residuals (mean residual = -0.01 , s.d. of residuals = 0.02). The residual analysis also shows that assumption of independent and identically distributed Gaussian noise on data is well justified, since the residuals for the local FKPP, global FKPP and logistic models

are all approximately centred at zero and are well characterised by the inferred noise distribution. The model selection results support the role of both transport and regional tau production playing an important role in driving AD pathology. We additionally test perform model selection using the priors provided in Table 3.2, with the results provided in Table 3.4. The results are consistent with the priors used in table Table 3.1, indicating the priors do not significantly alter the model fit to data. Overall, the data support the use of local FKPP model, evidenced by it being the most capable at describing in-sample and out-of-sample data, while also capturing the role of both tau transport and local tau production.




Figure 3.3: **Residual analysis.** Residual analysis for the local FKPP, global FKPP and logistic models. Residuals are shown between the last scan and the corresponding model prediction for all regions and subjects. A histogram of the residuals is shown on the right panel and the inferred Gaussian error model is shown in orange. The dashed black line shows average residual.




Figure 3.4: **Regional Residual analysis.** Regional residuals averaged over subjects for the local FKPP, global FKPP and logistic model. On the y-axis are the average regional residual for the final scan and the corresponding model prediction. On the x-axis is the average regional SUVR over subjects. The dashed black line shows the average residual across regions.

3.2.3 Local Model Accurately Forecasts Regional Tau Progression

A major possible benefit of mathematical modelling of AD lies in its application to pharmacological and clinical research, particularly through predicting patient disease progression. Here, we show that the local FKPP model can be used to accurately *forecast* the trajectory of tau PET. We divide the ADNI A^+T^+ cohort into train and test cohorts. We define a fixed training set comprising $n = 41$ subjects who have three longitudinal scans. For the remaining $N = 16$ subjects with greater than three longitudinal scans, we run inference three times using the Equation (3.21) and the priors in Table 3.1, each time adding an additional scan to the training set, starting with a single scan. The resulting posterior predictive trajectories are displayed in Figure 3.5 for the left inferior temporal lobe and Figure 3.6 for the left entorhinal cortex. Here, we neglect the influence of noise to highlight the uncertainty in model parameters. Posterior trajectories with noise are provided in Figure 3.7 and Figure 3.8, for the inferior temporal lobe and entorhinal cortex, respectively.

The results show that a single scan is often insufficient to provide meaningful forecasting accuracy, despite benefiting from information pooled across subjects in the hierarchical Bayesian model. This greatly improves with the addition of a second data point, however in some cases this produces inaccurate forecasts of future data if there is a decrease in tau SUVR, perhaps due to atrophy. With the addition of a third data point, the results generally converge with low uncertainty and accurately forecast future observations. The results suggest that a more constrained approach which more heavily weights the effect of population priors on forecasts may prove fruitful in circumstances in which there is insufficient data to capture an individual's trajectory. Nonetheless, these results demonstrate the power of a simple model with two free parameters in accurately forecasting the regional progression of a individual tau PET progression that may provide benefit for clinical and pharmaceutical researchers.

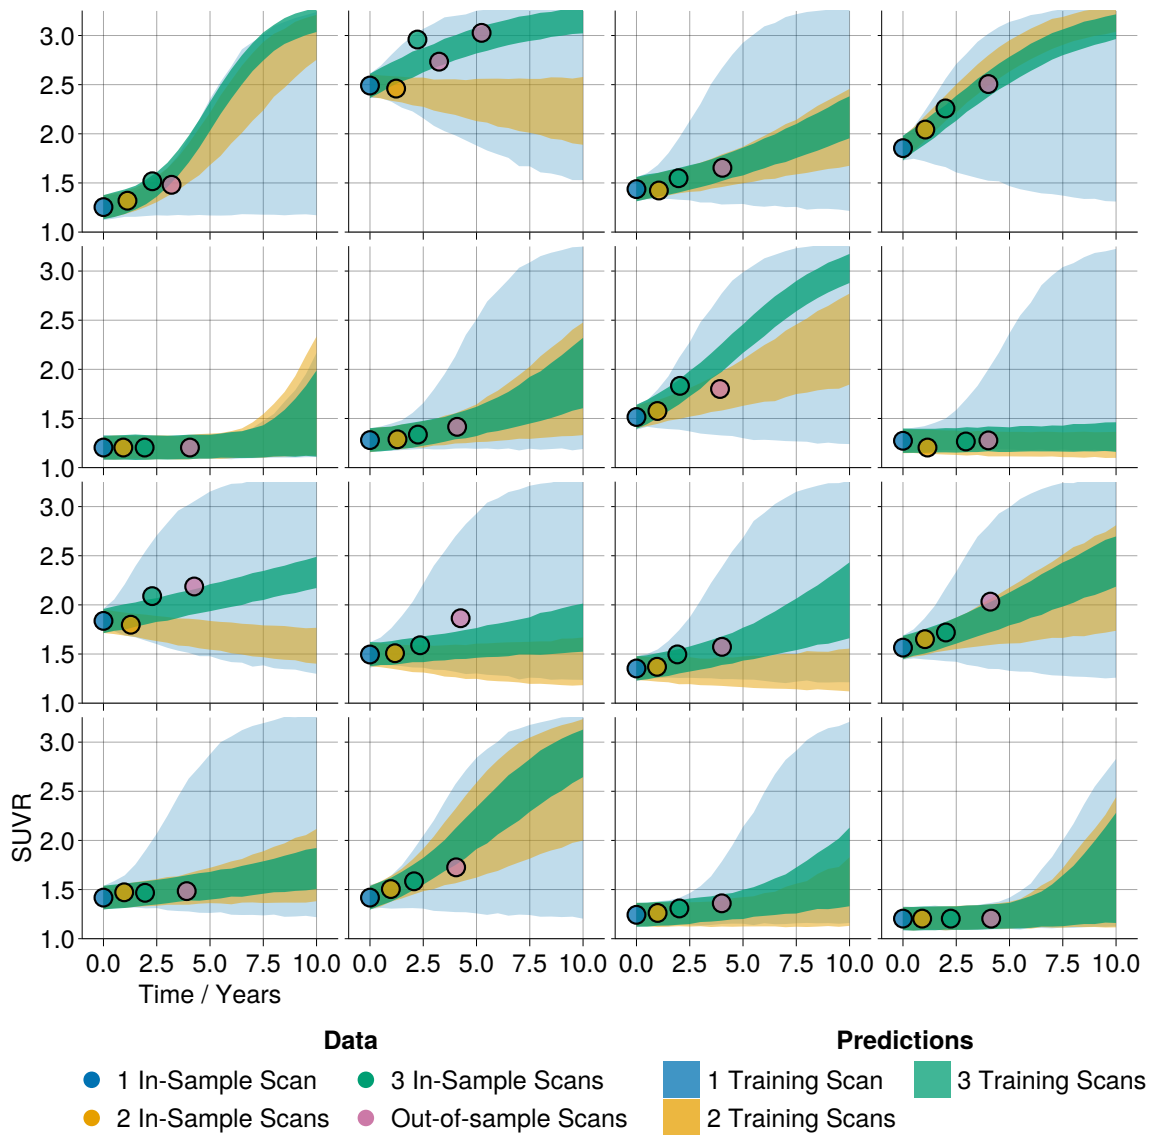


Figure 3.7: **Posterior predictive plots with noise for the left inferior temporal lobe.** The local FKPP model was iteratively calibrated to a A^+T^+ ADNI cohort with 41 in-sample subjects and 16 test subjects. Three iterations were run where for each iteration an additional scan from the test subjects were included, starting with a single scan. Posterior predictive trajectories for left inferior temporal lobe are shown for each iteration, including inferred observation noise. In the above figure, each panel represents each of the 16 test subjects. Each point represents a data point added for training iteration; trajectories are color matched to correspond to the number of longitudinal data points included for training.

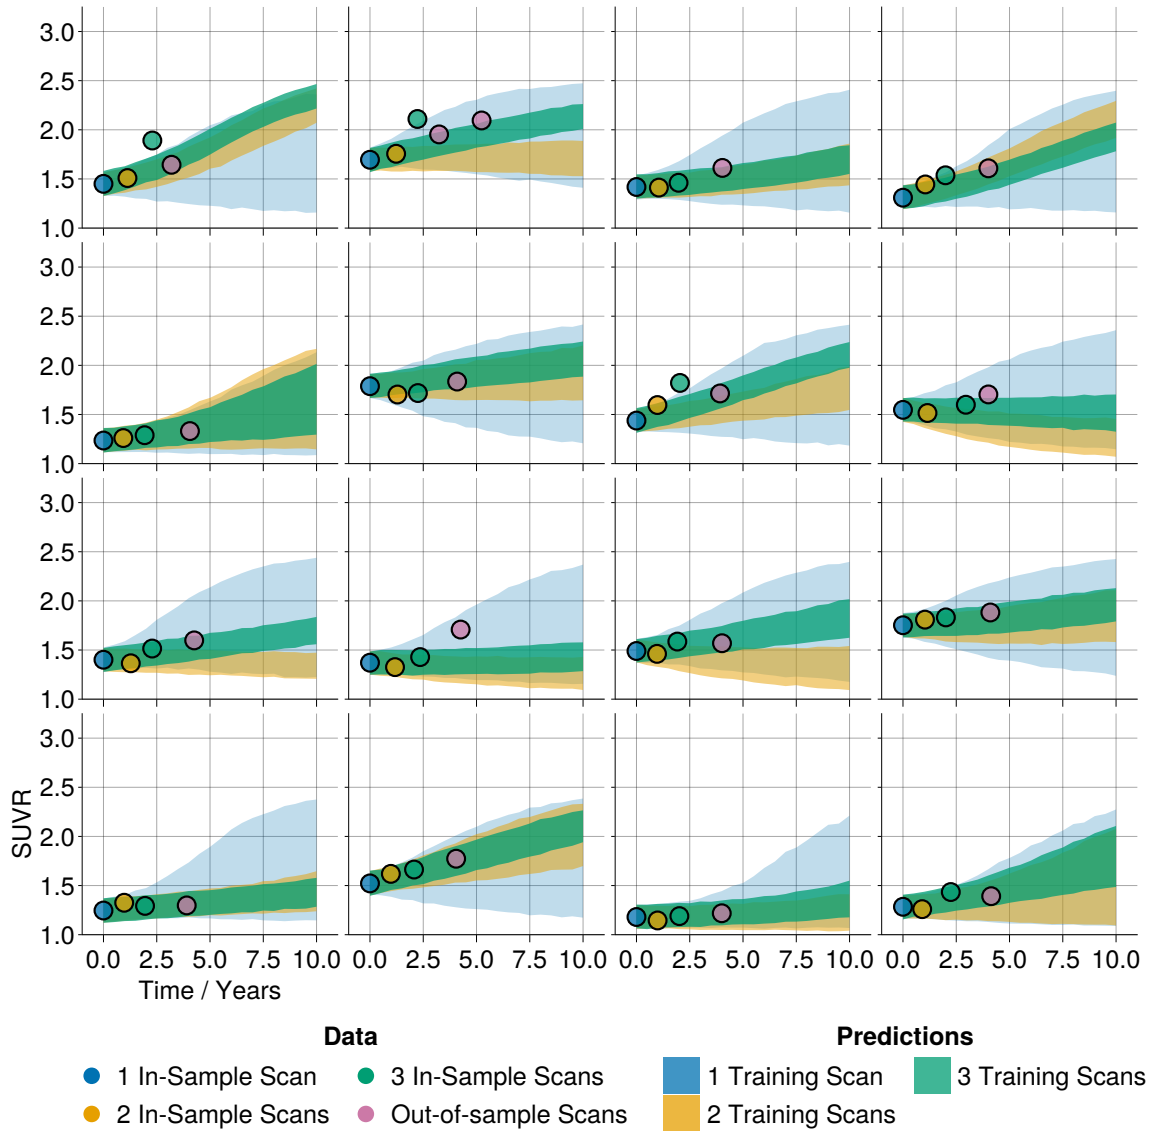


Figure 3.8: **Posterior predictive plots with noise for the left entorhinal cortex.** The local FKPP model was iteratively calibrated to a A^+T^+ ADNI cohort with 41 in-sample subjects and 16 test subjects. Three iterations were run where for each iteration an additional scan from the test subjects were included, starting with a single scan. Posterior predictive trajectories for left inferior temporal lobe are shown for each iteration, including inferred observation noise. In the above figure, each panel represents each of the 16 test subjects. Each point represents a data point added for training iteration; trajectories are color matched to correspond to the number of longitudinal data points included for training.

3.2.4 Early AD progression is driven by tau transport

We next sought to determine whether there are any changes in tau production and transport dynamics across the AD progression timeline. To do so, we use two cohorts of tau PET data, ADNI and BioFINDER-2 (BF2), each divided into three groups, A^- , A^+T^- , A^+T^+ , representing different stages of AD.

Data from the Swedish BioFINDER-2 study uses the RO948 tau PET radiotracer. All participants were recruited at Skåne University Hospital and the Hospital of Ängelholm, Sweden and the cohort covers the full spectrum of AD, ranging from cognitively normal individuals, patients with mild cognitive impairment (MCI) and with dementia. All details about the cohort have been described previously [141]. Amyloid status was determined by amyloid-PET (flutemetamol) based on a previously established cutoff from Gaussian mixture modelling as detailed in [142]. Patients with dementia do not undergo amyloid-PET and thus amyloid status was based on the CSF $A\beta$ 42/40 ratio [143]. The tau data are processed using the analysis pipeline detailed in [141]. Briefly, SUVR images were generated using the inferior cerebellum as reference region and average SUVR were extracted from the DKT atlas. Four subjects from the BF2 A^+T^- group are removed due to high off-target binding in the skull/meninges or MRI registration issues. Since BF2 uses a different tau PET radiotracer, we rerun the Gaussian mixture modelling analysis detailed in Section 3.2.1 to recover tracer specific baseline values, carrying capacities and tau positivity thresholds. For BF2, the tau positivity thresholds in the MTL and neocortex are 1.248 and 1.451, respectively. In both cohorts we select only subjects who have at least three tau PET scans to allow for inference on the time-series model. ADNI and BF2 tau PET data are summarised in Table 3.5.

	Group	N. Subjects	Age	Female	Education	CN	MCI	AD
ADNI	A^+T^+	57	72.8	0.56	16.3	0.28	0.50	0.21
	A^+T^-	39	75.7	0.52	15.9	0.65	0.27	0.07
	A^-	65	71.6	0.49	16.4	0.56	0.40	0.03
BF2	A^+T^+	54	73.9	0.65	11.6	0.2	0.43	0.37
	A^+T^-	18	72.1	0.5	13.2	0.61	0.31	0.08
	A^-	53	66.9	0.5	12.8	0.81	0.19	0.0

Table 3.5: Demographics for ADNI and BF2 cohorts.

We fit the hierarchical Bayesian model presented in Section 3.2.2 to the A^+T^+ , A^+T^- and A^- groups using the priors in Table 3.1 across all groups, and examine the population level production and transport parameters. The population parameter distributions for the A^- , A^+T^- and A^+T^+ groups are shown in Figure 3.9a for ADNI and Figure 3.10 for BF2 and summarised in Table 3.6. The posterior distributions across cohorts are qualitatively the same, with changes likely reflecting differences in cohort and tracers. The inferred parameters show an increase in the transport rate for the A^+T^- group relative to the A^+T^+ and A^- groups, suggesting that tau more readily spreads between regions during early stages of disease (A^+T^-), and is minimal in later stages of AD (A^+T^+). The inferred posterior distributions for the production parameter show a progressive increase in production rate along the disease timeline, with a primary increase from A^- to A^+T^- group and a secondary increase from the A^+T^- group to the A^+T^+ group. The negative production rate for the A^- group indicates that the signal, on average, decreases. This could reflect changes in noise due to off-target, nonspecific binding or atrophy from non-AD related neurodegeneration. Therefore, any observed clearance dynamics likely reflect fluctuations in noise or age-related atrophy. Additionally, we check the bivariate posterior distributions between the transport and production parameters to examine whether there are any correlations between parameters. These are shown in Figure 3.11. For all biomarker groups in both ADNI and BF2, the population production and transport parameters are uncorrelated, showing that changes in transport and production across groups do not result from correlations between these parameters. We additionally test the effect on posterior distributions of changes to prior distributions, in particular, changes to population scale parameters, dictating variation among individuals, using the priors provided in Table 3.2. The results of this analysis are shown in Figure 3.12 for ADNI data and Figure 3.13 for BF2 data. The posteriors show very little difference to the inferred parameters in Figure 3.9a and Figure 3.10 (for ADNI and BF2, respectively), suggesting the change in prior has little influence over inferred parameters. These results suggest that in early AD, tau begins in a transport dominated phase ($\rho > \alpha$) and later switches to a production dominated phase ($\alpha > \rho$).

To confirm that the parameter distributions reflect meaningful dynamics present in the data and are not as a result of statistical patterns, we re-run the analysis on the A^+T^- and A^+T^+ groups using spatially shuffled data [87]. We perform this test on the A^+T^+ and A^+T^- groups by performing a random spatial shuffling of the data, using the same shuffle across subjects. The same random permutation are applied to the regional baseline values and carrying capacities to ensure numerical stability.

The same inference algorithm used for unshuffled data was applied to the shuffled dataset and 1000 posterior samples were collected. This process was repeated 10 times for each group. In the A^+T^+ positive group, one chain failed to converge and was discarded. The results of these for the production and transport parameters are shown in Figures 3.9b and 3.9c. Note that we only spatially shuffle the data and therefore expect minimal changes to the estimated production parameters. We note a marked difference in estimated transport dynamics in both the A^+T^- and A^+T^+ groups between the true and shuffled data, confirming that dynamics present in the data are not a consequence of statistical patterns in the data but represent tau dynamics as measured through PET. The transport dominated phase of the early AD subjects supports evidence showing tau seeds are present throughout the cortex before symptom onset [90, 144] and, together with the small role of transport in the A^+T^+ groups, helps explain the strong performance of the logistic model in Section 3.2.2. Overall the results reveal temporal changes in the dynamics of tau progression, with an initial transport dominated phase, perhaps in which seeds are deposited around the cortex, followed by a production dominated phase indicative of secondary tauopathy, likely due to spatial colocalisation with $A\beta$ catalysing tau production.

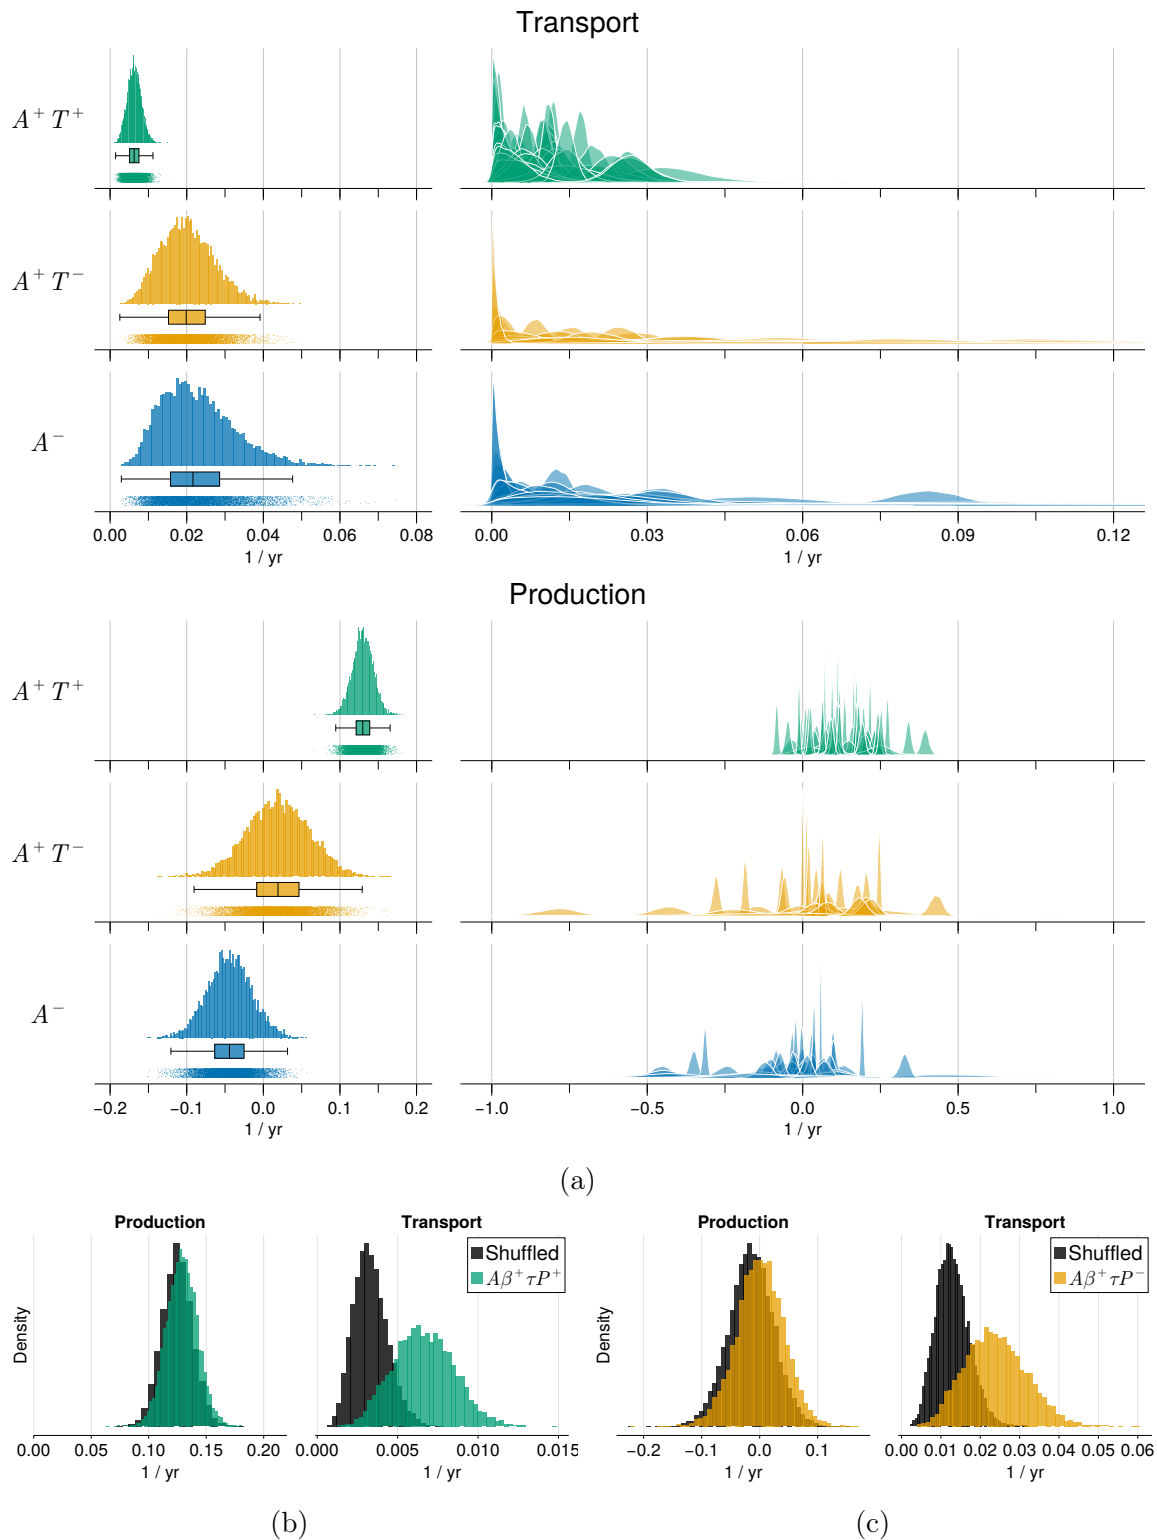


Figure 3.9: **Inferred population level parameters using ADNI data.** **3.9a** Population (left) and individual (right) level transport and production parameters for A^+T^+ , A^+T^- and A^- groups in the ADNI cohorts. **3.9b & 3.9c** Inferred population production and transport parameters from spatially shuffled data (shown in grey) for A^+T^+ (3.9b), A^+T^- (3.9c) ADNI groups.

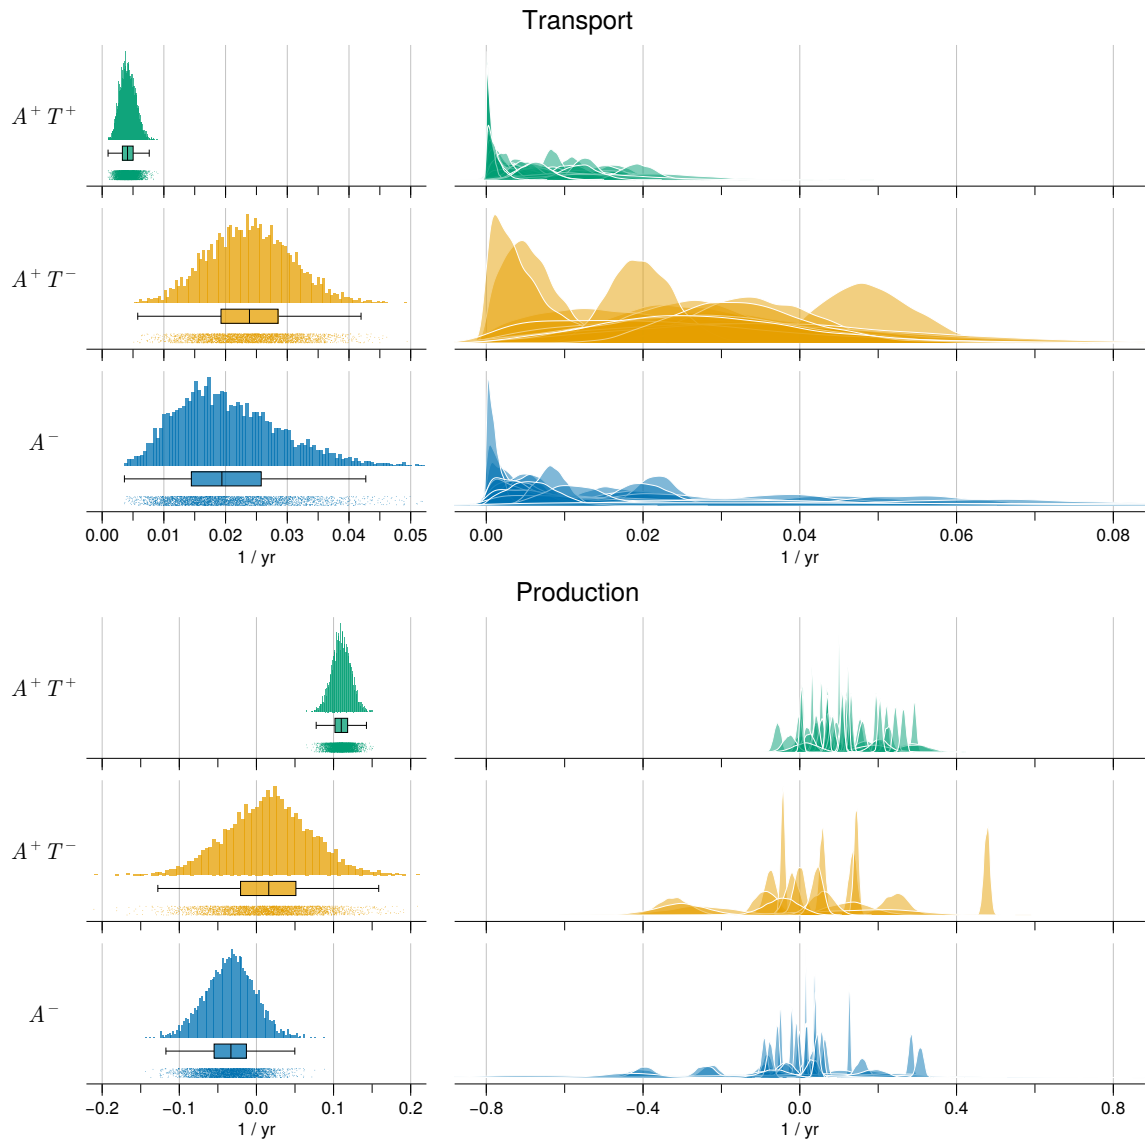
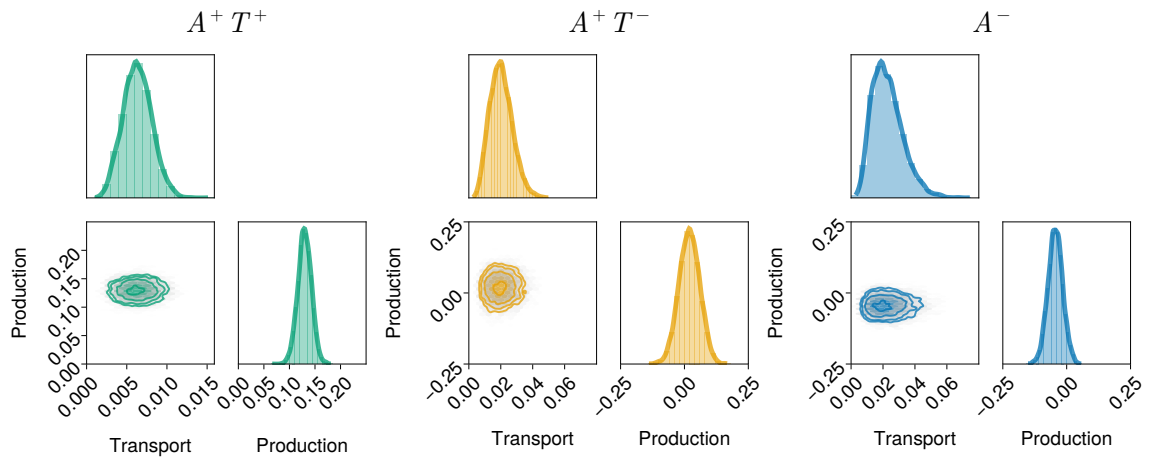
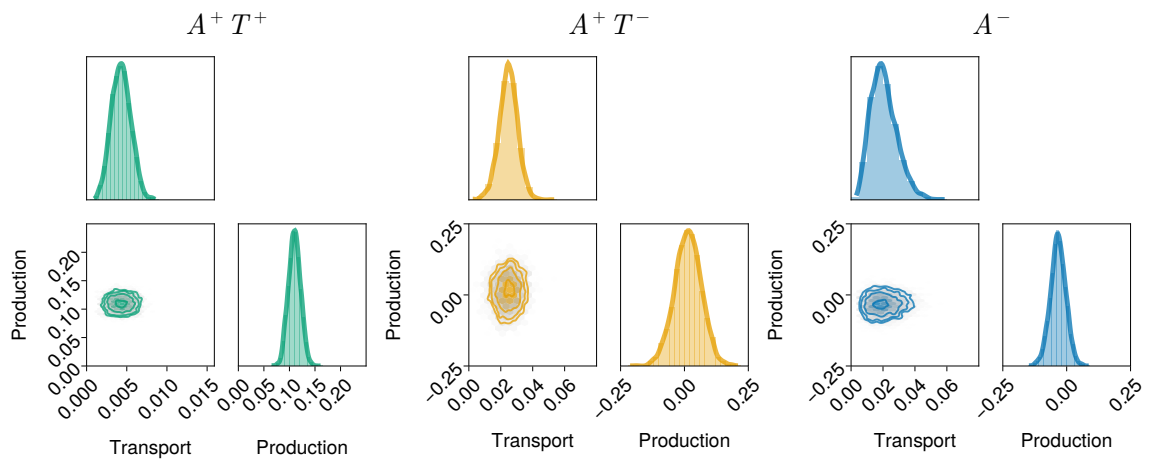


Figure 3.10: **Inferred population level parameters using BF2 data. 3.9a** Population (left) and individual (right) level transport and production parameters for A^+T^+ , A^+T^- and A^- groups in the BF2 cohorts.



(a)



(b)

Figure 3.11: **Bivariate posterior distributions for production and transport.** Population level bivariate posterior distributions for transport and production parameters in the A^+T^+ , A^+T^- , and A^- groups for ADNI (3.11a) and BF2 (3.11b).

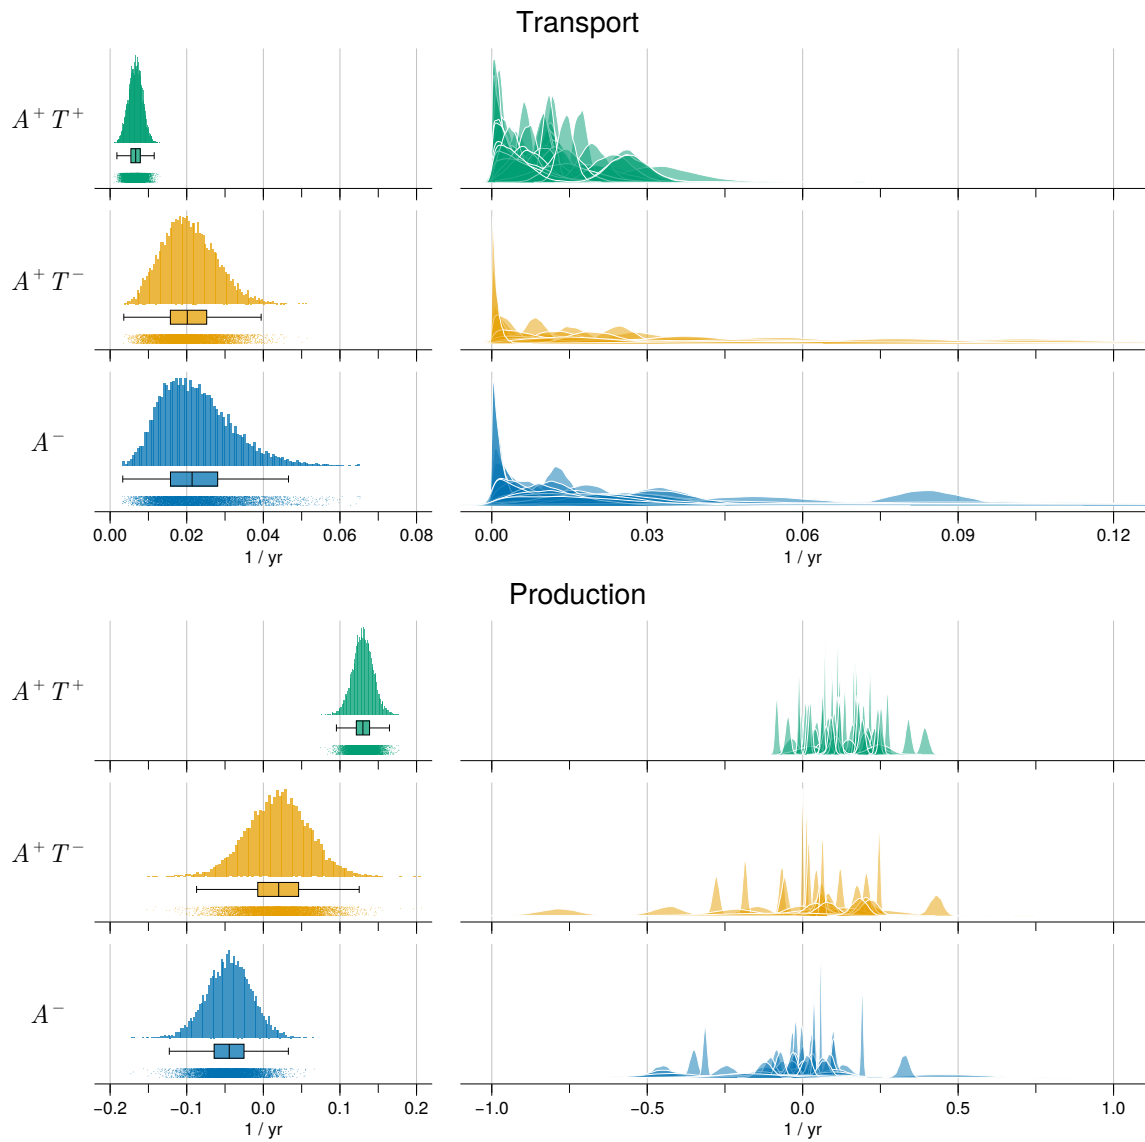


Figure 3.12: Inferred population level parameters using ADNI data with changes to population-level prior distributions. 3.9a Population (left) and individual (right) level transport and production parameters for A^+T^+ , A^+T^- and A^- groups in the ADNI cohorts using priors in Table 3.2.

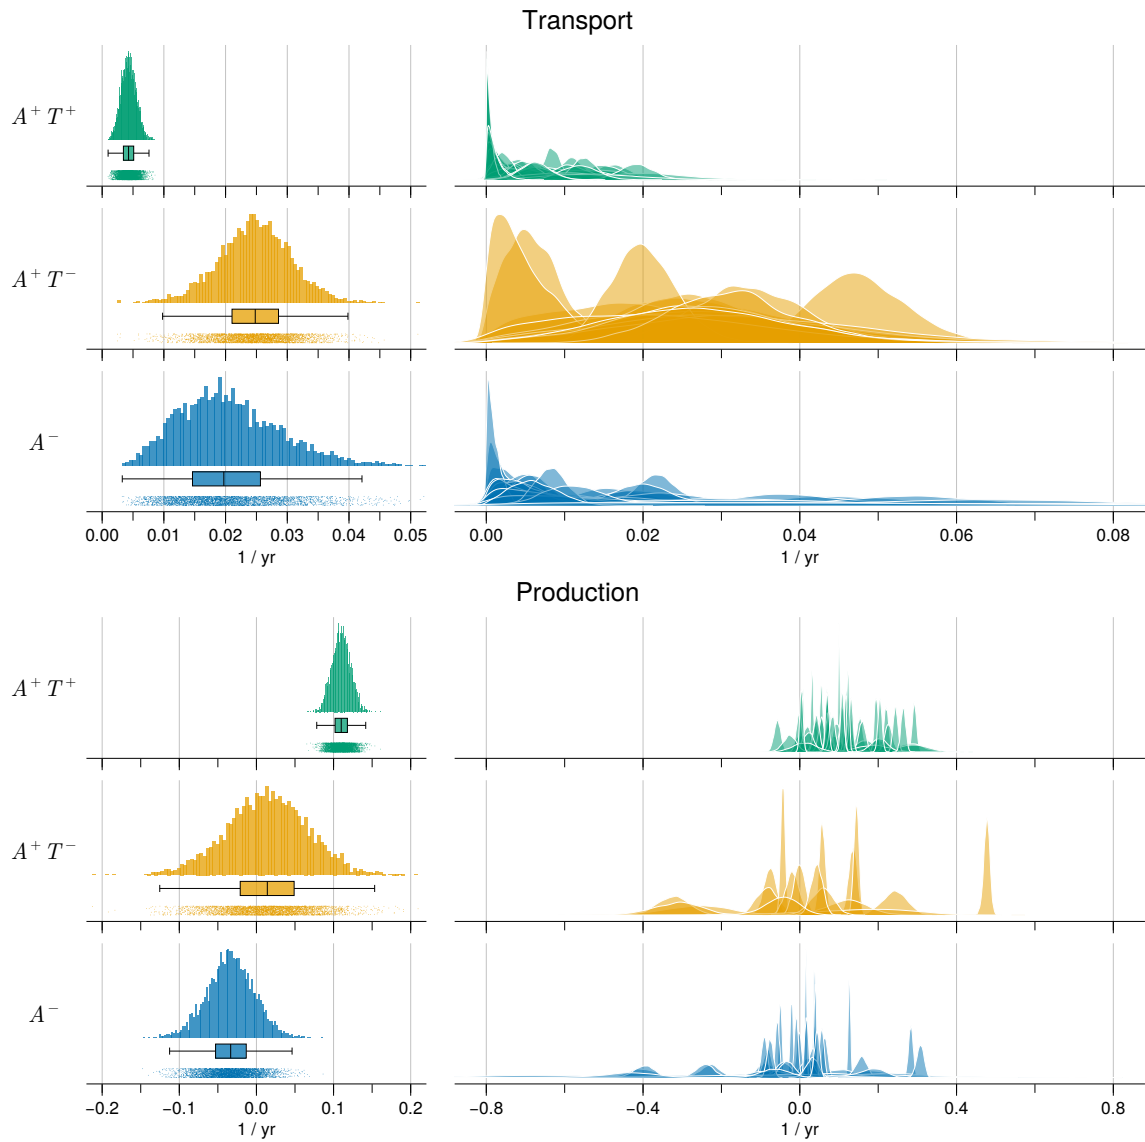


Figure 3.13: Inferred population level parameters using BF2 data with changes to population-level prior distributions. 3.9a Population (left) and individual (right) level transport and production parameters for A^+T^+ , A^+T^- and A^- groups in the BF2 cohorts using priors in Table 3.2.

	Group	ρ_μ		ρ_σ		α_μ		α_σ	
		mean	s.d.	mean	s.d.	mean	s.d.	mean	s.d.
ADNI	A^+T^+	0.01	0.003	0.02	0.003	0.12	0.02	0.11	0.02
	A^+T^-	0.04	0.01	0.15	0.03	-0.005	0.05	0.26	0.04
	A^-	0.02	0.01	0.05	0.01	-0.05	0.03	0.19	0.02
BF2	A^+T^+	0.004	0.001	0.01	0.001	0.11	0.01	0.09	0.01
	A^+T^-	0.02	0.01	0.02	0.01	0.02	0.05	0.22	0.04
	A^-	0.02	0.01	0.08	0.01	-0.03	0.03	0.19	0.03

Table 3.6: Posterior summary providing the means of inferred population parameters for ADNI and BF2.

3.3 Discussion and Conclusions

In this chapter, we have derived a mechanistic model to describe tau PET data in terms of underlying tau dynamics and applied it data from ADNI and BF2 to understand how it compares to other models present in the literature and how modelling can inform us about tau transport and production dynamics in the human brain. We have shown that this model can accurately forecast longitudinal regional tau PET progression in AD subjects. Furthermore, by performing inference across different patient groups across the AD disease timeline, we uncover temporal changes in transport and production dynamics, showing an initial transport dominated phase associated with primary tauopathy and seeding, followed by an accelerated production dominated phase indicative of secondary tauopathy.

There have been several studies proposing different models of proteopathy in AD, a key difference among them being descriptions of the tau production process, which vary widely in complexity [72, 75, 78, 76, 79, 83, 133]. Here we present a parsimonious model of tau transport that relies on regionally specific carrying capacities and we showed through model selection, Section 3.2.2, that it is able to outperform other models proposed in the literature. A possible cause of regional heterogeneity in carrying capacities is heterogeneity in regional risk factors that promote tau proliferation, the most likely of which is $A\beta$. $A\beta$ has its own spatial topography within AD patients, being particularly present throughout the fronto-parietal-temporal default mode network and stimulating neuronal hyperactivation [6, 145, 146]. The presence of $A\beta$ will have a two-fold effect on tau dynamics, first through a catalysing effect on

the production of tau [38, 39, 40] and second by promoting the activity dependent spread and production through the functional networks [147, 148]. In Thompson et al 2020., [89] we formulated a model describing the dynamic interaction between $A\beta$ and tau, that predicts an increase in carrying capacities based on $A\beta$ concentration, however, further work toward simplifying the model will be necessary before it can be used for inference with patient data.

Another key set of factors contributing to regional vulnerability are genetic markers. It has already been shown in mice models of AD that gene expression patterns can inform spreading of tau [137] and human modelling of Parkinson’s disease has shown how gene expression patterns can inform regional vulnerability to create a model of toxic protein spread in Parkinson’s disease [149]. There are several candidate genes for modelling regional vulnerability in AD, most notably microtubule association protein tau (MAPT), as a proxy for relative baseline tau vulnerability [150] and apolipoprotein-E (APOE) for those patients with the APOE ϵ 4 mutation [151, 135, 152]. While there are many other candidate genes that may influence regional vulnerability, care should be taken to avoid creating overdetermined models. In sum, while the work here provides compelling evidence for the necessity of regional vulnerability, further work should seek to explain the mechanisms through which regional carrying capacities emerge from a culmination of regional risk factors, such as $A\beta$ deposition and gene expression patterns.

There is extensive evidence of tau transport and production throughout the brain, however, it has not yet been determined whether one of these processes dominate the other and whether their relative contributions to disease progression changes over time. To this end, we sought to determine whether inferred parameters of our model change in groups across the disease timeline. We find that during early stages of disease, when there is a low tau concentration in the medial temporal lobe, tau dynamics are transport dominated but become production dominated later in disease. This supports previous work by Meisl et al. [133] who show through an analysis of multiple datasets and methods of tau quantification that tau dynamics are production dominated from Braak stage 3 onwards. This is consistent with our work, considering individuals who are positive on tau-PET in early Braak stage regions may already show fairly advanced Braak stages at autopsy [153] analogous to individuals at middle Braak stages used by Meisl et al. 2021. [133] These results suggest that, in early AD, tau seeds invade connected regions from the medial temporal lobe, but the overall concentration does not grow substantially. Only in later stages of AD is

the spread production-dominated and driven by fast increases in concentration gradients, leading to progressive Braak-like staging. These results also support the utility of the logistic production model in being able to describe longitudinal A^+T^+ data (Figure 3.2a & Figure 3.2b), since seeds would already be densely present around the cortex and progression is driven by tau production. The results are also consistent with experimental evidence showing that tau seeds are present before tau pathology [90, 144].

Together the results are indicative of an intrinsically spatiotemporal process, with a primary, transport dominated tauopathy resulting in tau seeds spreading from the medial temporal lobe to axonally connected regions, followed by a secondary, production dominated tauopathy, likely resulting from $A\beta$ interaction, resulting in fast regional accumulation and sequential, Braak-like staging. This is in slight contrast to the largely temporal process of $A\beta$, described by [154], where $A\beta$ begins as being diffusely present throughout the brain, but increases in concentration at different rates due to regional vulnerabilities. These results suggest that the early period of AD during which tau is more readily transported between brain regions may be a critical time for intervention. Many immunotherapies currently being developed act on extracellular tau [155] and should therefore interrupt tau transmission through the extracellular space of the synaptic junction. If AD is a consequence of first tau spread and then production, it will be crucial for these immunotherapies to be administered early in the AD process to halt the widespread transmission of tau before accelerated local production can occur. In contrast, therapies that act to reduce intracellular tau concentration should be effective in slowing AD progression across the AD continuum, regardless of whether widespread tau transmission has occurred [156].

This work presents a step forward in whole brain tau modelling, however, there are still many obstacles that are not addressed here. First, there are limitations that pertain to the sparsity of longitudinal data. In this work, we fix a number of parameters to ensure the practical identifiability of the models given the available data. In particular, we fix baseline values and carrying capacities in the dynamical system, and subject initial conditions in the probabilistic model. By fixing baseline values and carrying capacities, we are unable to determine whether these also undergo dynamical changes. This also limits the direct application of the model to other tauopathies that exhibit different tau PET profiles. This will, in part, be addressed by including dynamical regional risk factors, however, the time over which changes in these factors can be observed will remain limited until more longitudinal data is available. Second, there are limitations related to the scale at which we are modelling.

While the model we present here is derived from a mechanistic model, the model reduction comes at the cost of a loss in mechanistic insight into precise transport and production mechanisms. This will remain a hard limitation while we work with macroscale brain data. In addition, tau PET data is intrinsically limited by resolution, inability to detect early changes, and non-specific and off-target binding sources, collectively providing a source of uncertainty that affects parameter identifiability of intricate processes such as transport. Therefore, while the modelling results suggest changes to transport and production across the AD continuum, our conclusions are limited by the nature of PET measurements and require experimental validation. A potential avenue to address this limitation will be the development of multiscale models that rely on in-vitro or animal studies for calibration and permit macro-scale reduced order models.

The primary contribution of this chapter has been to supply a parsimonious account of regional tau dynamics in AD. Furthermore, the study sheds light on potential avenues of clinical investigation into anti-tau therapies by showing how targeting different tau processes (transport or local production) at different times during the AD continuum may be essential for effective intervention. In the following chapter, we will expand on this work to develop multimodal model of amyloid, tau and neurodegeneration with the aim of describing unexplained features of AD pathology, such as how regional variations in tau dynamics arise, the role of primary age related tauopathy and tau related atrophy.

Chapter 4

Dynamical Modelling of the ATN Pathway

In the previous chapters, we have developed the foundation for dynamical modelling of neuroimaging data for AD, so far dealing specifically with tau pathology. In Chapter 2, we focussed on developing a full pipeline for connecting longitudinal individual subject neuroimaging data with mathematical models of disease progression. In Chapter 3, we specifically addressed progression modelling of tau PET and highlighted the importance of regionally specific tau dynamics in the longitudinal modelling of tau dynamics. In this chapter, we will combine these toward an overarching goal of modelling interactions between three key elements of the AD pathway, $A\beta$, tau and neurodegeneration. Imaging and fluid biomarkers of $A\beta$, tau and neurodegeneration comprise the amyloid-tau-neurodegeneration (ATN) biomarker framework, currently used by the National Institute of Ageing for quantitatively classifying the stages of Alzheimer's disease according to biomarker abnormality (see Section 1.2.3 for a summary of the ATN framework or [10] for a review). In so doing, we will show that macro-scale brain modelling combined with probabilistic modelling, can provide a greater understanding of dynamical interactions between key elements in AD pathology and facilitate longitudinal prediction of key ATN biomarkers.

4.1 Introduction

Alzheimer's disease has two key driving factors, $A\beta$ and tau. Thus far, we have focussed primarily on the role of tau, however, in this chapter we will provide a more comprehensive picture of AD by modelling the copathology of $A\beta$ and tau and their effect on neurodegeneration.

The role of $A\beta$ in AD pathology is well documented and the primary theory of AD remains the $A\beta$ cascade hypothesis (see Section 1.2.2 for more details or [44, 45] for review). The $A\beta$ cascade hypothesis proposes that $A\beta$ is the primary causative agent in AD and that $A\beta$ leads to the development of tau pathology and neurodegeneration [1, 8, 44]. This hypothesis has been repeatedly updated to account for criticisms, particularly around its interactions with tau and the role of primary age related tauopathy (PART) in the AD continuum [45, 32]. In its current iteration, the $A\beta$ cascade hypothesis still maintains a critical role for $A\beta$ in the development of AD and posits that $A\beta$ is necessary for tau expansion into the neocortex and that PART is a risk factor, or perhaps even a prerequisite for $A\beta$ induced secondary tauopathy [45]. The $A\beta$ cascade hypothesis also predicts the biomarker progression that is described by the ATN framework, in which substantial $A\beta$ load increase precedes tau proliferation and atrophy [46, 10]. The ATN framework, which is now the standard diagnostic criteria issued by the National Institute for Ageing, place only those with A^+ status on the AD continuum, with characterisations associated with PART ($A^-T^+N^-$ and $A^-T^+N^+$) categorised as non-AD pathological change (see Table 1.1 for ATN categories), at odds with those who place PART as an early event on the AD continuum [37, 157]. In this chapter, we aim to illuminate the role of $A\beta$ in facilitating downstream events in AD, namely tau pathology and subsequent atrophy.

It is worth commenting explicitly on those interactions which we will investigate. Using animal models of tau and $A\beta$ pathology two important interactions between the toxic agents have been identified. First, $A\beta$ has a catalytic effect on tau tangle formation [38, 39, 40]. Second, $A\beta$ accelerates tau spreading between axonally connected regions [38, 158]. Therefore, $A\beta$ has both local and non-local effects on tau pathology, affecting both transport and local production of tau. The third interaction we wish to examine more closely is the relationship between tau and neurodegeneration, since tau, but not $A\beta$, is most well correlated with the atrophy patterns observed in AD subjects [25, 26, 117]. This was investigated in Chapter 2 and in the following chapter we will determine how well $A\beta$ -tau co-pathology explains observed atrophy.

The overarching research aim in this chapter is to provide a mathematical model of the ATN framework that predicts biomarker progression and interactions. More specifically, we aim for the model to account for $A\beta$ -induced acceleration in tau expansion and of the sequential staging of $A\beta$, tau and neurodegeneration. In doing this, we find that $A\beta$ accounts for regional heterogeneity in tau dynamics, tau staging, and atrophy onset. Furthermore, we show that the coupled model is able to predict the region and time $A\beta$ and tau first colocalise, causing an acceleration of tau pathology

in expansion into the neocortex. Finally, we show that gradients in A β deposition predicts the seeding of tau in the entorhinal cortex.

4.2 Modelling the ATN Pathway

We start with a modified version of the network heterodimer model that incorporates A β /tau interactions [89]. For variables $u_i = u_i(t)$ and $\tilde{u}_i = \tilde{u}_i(t)$ representing toxic and healthy A β concentration, respectively, and $p_i = p_i(t)$ and $\tilde{p}_i = \tilde{p}_i(t)$ healthy and toxic tau concentration, respectively, at each node $i = 1, \dots, R$, the coupled model is:

$$\frac{du_i}{dt} = -\rho_1 \sum_{j=1}^R L_{ij} u_j + k_0 - k_1 u_i - k_2 u_i \tilde{u}_i, \quad (4.1a)$$

$$\frac{d\tilde{u}_i}{dt} = -\rho_1 \sum_{j=1}^R L_{ij} \tilde{u}_j - \hat{k}_1 \tilde{u}_i + k_2 u_i \tilde{u}_i, \quad (4.1b)$$

$$\frac{dp_i}{dt} = -\rho_2 \sum_{j=1}^R L_{ij} p_j + b_0 - b_1 p_i - b_2 p_i \tilde{p}_i - b_3 \tilde{u}_i p_i \tilde{p}_i, \quad (4.1c)$$

$$\frac{d\tilde{p}_i}{dt} = -\rho_2 \sum_{j=1}^R L_{ij} \tilde{p}_j - \hat{b}_1 \tilde{p}_i + b_2 p_i \tilde{p}_i + b_3 \tilde{u}_i p_i \tilde{p}_i, \quad (4.1d)$$

with $u_i(0) = u_{i,0}$, $\tilde{u}_i(0) = \tilde{u}_{i,0}$, $p_i(0) = p_{i,0}$, and $\tilde{p}_i(0) = \tilde{p}_{i,0}$ for $i = 1, \dots, R$. The model comprises two heterodimer models for (u_i, \tilde{u}_i) and (p_i, \tilde{p}_i) , for $i = 1, \dots, R$, with an additional one-way coupling between them, $\pm b_3 \tilde{u}_i p_i \tilde{p}_i$, describing the catalytic effect of A β on the tau prion-like process of production. A full analysis of the model is given in [89], highlighting how the model predicts secondary tauopathy based on the spatial colocalisation of A β and tau and how a linearisation of the model produces travelling wave dynamics. The model has the same pitfalls as the heterodimer model, that is, while it is informative about the mechanisms of AD proteopathy, there is insufficient information in tau PET data from which to calibrate the parameters. Therefore, we seek to make an approximation to this model with fewer parameters that can be calibrated with existing A β and tau PET data.

We follow a similar procedure to previous chapters and [159]. Equations Equation (4.1a)–Equation (4.1b) correspond to a heterodimer model identical to Equation (2.4) and since they are uncoupled from Equation (4.1c)–Equation (4.1d), they can be linearised in the same way (see Section 3.2.1). We further assume that the

transport of $A\beta$ is negligible across the brain network, following evidence from cross-sectional analysis showing that $A\beta$ is likely to be densely deposited around the cortex during early stages of AD [160]. Therefore, we have:

$$\frac{d\tilde{u}_i}{dt} = \alpha\tilde{u}_i(\tilde{u}_{\infty,i} - \tilde{u}_i), \quad (4.2)$$

for $i = 1, \dots, R$, where $u_{\infty,i}$ is the regional carrying capacity for $A\beta$, to describe the evolution of toxic $A\beta$.

To linearise Equation (4.1c)–Equation (4.1d), we assume a healthy and homogeneous state, then rewrite Equation (4.1c) to find v as a function of (\tilde{u}, \tilde{v}) ,

$$p_i(\tilde{u}_i, \tilde{p}_i) = \frac{b_0}{b_1 - b_2\tilde{p}_i - b_3\tilde{u}\tilde{v}}, \quad (4.3)$$

to which we can make a first order approximation,

$$p_i(\tilde{u}_i, \tilde{p}_i) = \frac{b_0}{b_1} + \tilde{p}_i \left(-\frac{b_0b_2}{b_1^2} - \frac{b_0b_3\tilde{u}_i}{b_1^2} + \mathcal{O}(\tilde{u}_i^2) \right) + \mathcal{O}(\tilde{p}_i^2). \quad (4.4)$$

Substituting this expression into Equation (4.1d), we have the following expression

$$\frac{d\tilde{p}_i}{dt} = -\rho_2 \sum_{j=1}^R L_{ij}\hat{v}_j + \left(\frac{b_0b_2}{b_1} - \hat{b}_1 \right) \tilde{p}_i + \frac{b_0b_3}{b_1}\tilde{u}_i\tilde{p}_i - \frac{b_0b_2^2}{b_1^2}\tilde{p}_i^2 + \mathcal{O}(\tilde{p}_i^2\tilde{u}_i). \quad (4.5)$$

Following the assumption of a healthy state, we drop higher order terms of toxic species, since they will be small relative to quadratic terms. The, grouping like powers of \tilde{p}_i ,

$$\frac{d\tilde{p}_i}{dt} = -\rho_2 \sum_{j=1}^n L_{ij}\hat{v}_j + \gamma\tilde{p}_i (v_{\infty,i} - \tilde{p}_i), \quad (4.6)$$

$$v_{\infty,i} = \kappa + \beta\tilde{u}_i, \quad (4.7)$$

where

$$\gamma = \frac{b_0b_2^2}{b_1^2}, \quad \kappa = \frac{b_0b_2}{b_1\gamma} - \frac{\hat{b}_1}{\gamma}, \quad \beta = \frac{b_0b_3}{b_1\gamma}, \quad (4.8)$$

and $v_{\infty,i}$ is the regional carrying capacity in the absence of transport. This model provides a linear approximation to Equations (4.1c) to (4.1d). Therefore, our reduced order approximation for Equations (4.1a) to (4.1d) is given by a system of equations in variables $(\tilde{u}_i, \tilde{p}_i)$ for $i = 1, \dots, R$,

$$\frac{d\tilde{u}_i}{dt} = \alpha\tilde{u}_i(\tilde{u}_{\infty,i} - \tilde{u}_i), \quad (4.9a)$$

$$\frac{d\tilde{p}_i}{dt} = -\rho \sum_{j=1}^R L_{ij}\hat{v}_j + \gamma\tilde{p}_i (v_{\infty,i} - \tilde{p}_i), \quad (4.9b)$$

$$\hat{v}_{\infty,i} = \kappa + \beta\tilde{u}_i. \quad (4.9c)$$

Variables	Fixed Point	ATN Status
(u, p)	$(0, 0)$	$A^-T^-N^-$
(u, p)	$(u_\infty, 0)$	A^+T^-
(u, p, a)	$(0, p_\infty, 1)$	$A^-T^+N^+$
(u, p, a)	$(u_\infty, p_\infty, 1)$	$A^+T^+N^+$

Table 4.1: Fixed points of Equation (4.11) and ATN signatures.

Note that we drop subscript notation for ρ , since we neglect the transport the effect on $A\beta$. To derive a full model of the $A\beta$ -tau-neurodegeneration (ATN) framework, we can introduce a simple model of atrophy. Since atrophy is primarily correlated with tau [117, 25], we describe the evolution of atrophy, $a_i = a_i(t)$ that depends on tau, \tilde{p}_i for $i = 1, \dots, R$:

$$\frac{da_i}{dt} = \eta\tilde{p}_i(1 - a_i). \quad (4.10)$$

Together with Equation (4.9a)–Equation (4.9c), this defines a full dynamical model $A\beta$, tau and neurodegeneration and can be used to model the ATN framework.

4.2.1 Linear Stability and Bifurcation Analysis

Considering only dynamics in the absence of transport, we can analyse the system and see how the model relates to the ATN framework. We first examine the dynamics a single node, with variables (u, p, a) for $A\beta$, tau and atrophy, respectively, and we drop the $\tilde{\cdot}$ notation for $A\beta$ and tau, since we are now only modelling toxic proteins. The system is given by the coupled equations:

$$\frac{du}{dt} = \alpha u(u_\infty - u), \quad (4.11a)$$

$$\frac{dp}{dt} = \gamma p([\kappa + \beta u] - p), \quad (4.11b)$$

$$\frac{da}{dt} = \eta p(1 - a). \quad (4.11c)$$

The system has four fixed points, given by $(u_1, p_1, a_1) = (0, 0, 0)$, $(u_2, p_2, a_2) = (u_\infty, 0, 0)$, $(u_3, p_3, a_3) = (0, \kappa, 1)$, $(u_4, p_4, a_4) = (u_\infty, p_\infty, 1)$, each of which has a meaningful interpretation as stations along the ATN spectrum, given in Table 1.1.

We can examine the stability of the system by looking at the Jacobian matrix of

Equation (4.11), given by:

$$\begin{pmatrix} \alpha(u_\infty - 2u) & 0 & 0 \\ \beta\gamma p & \gamma(\kappa + \beta u - 2p) & 0 \\ 0 & \eta - \eta z & -\eta p \end{pmatrix}. \quad (4.12)$$

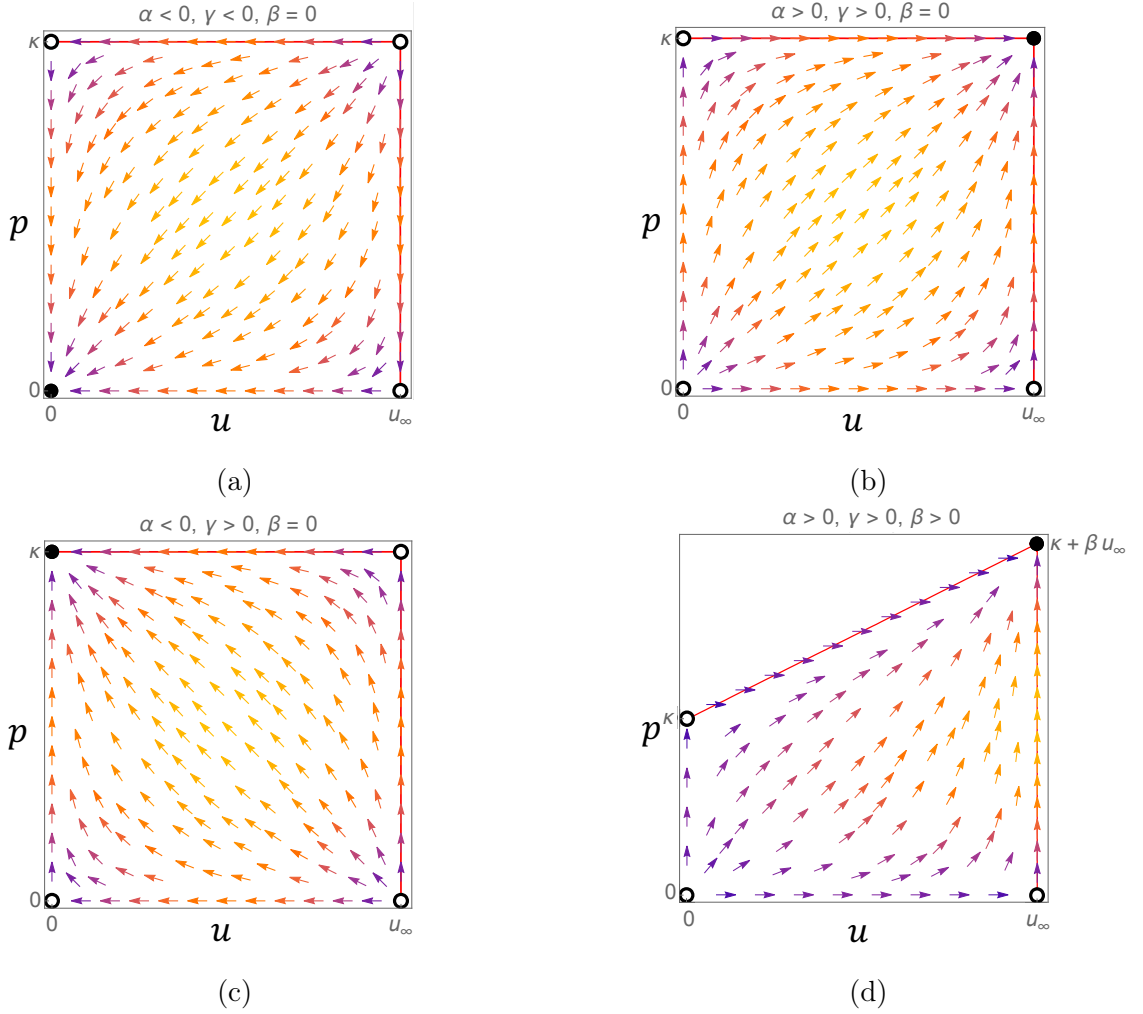


Figure 4.1: **Phase diagrams for the coupled $A\beta$ -tau model.** Phase diagrams for the $A\beta$ -tau components of the ATN model Equation (4.11) using different parameter configurations.

We start by examining the $A\beta$ -tau subsystem given by $\eta = 0$ to gain a better understanding of the implied protein dynamics. Then, for the first solution, $(u, p) = (0, 0)$, the eigenvalues of the Jacobian matrix are $(\alpha u_\infty, \gamma \kappa)$, therefore the fixed point is unstable for both u and p , if $\alpha > 0$ and $\gamma > 0$, or stable if $\alpha \leq 0$ and $\gamma \leq 0$. Similarly, for the solution $(u, p) = (u_\infty, p_\infty)$, the eigenvalues are $(-\alpha u_\infty, -\gamma p_\infty)$, and

Fixed Point	Jacobian Eigenvalues
$(0, 0)$	$(\alpha u_\infty, \gamma \kappa, 0)$
$(u_\infty, 0)$	$(-\alpha u_\infty, \gamma p_\infty, 0)$
$(0, \kappa, 1)$	$(\alpha u_\infty, -\gamma \kappa, -\eta \kappa)$
$(u_\infty, p_\infty, 1)$	$(-\alpha u_\infty, -\gamma p_\infty, -\eta p_\infty)$

Table 4.2: Linear stability analysis of single node ATN model.

the stability depends on the sign of α and γ . Therefore, α and γ act as transcritical bifurcation parameters, exchanging stability between (u_1, p_1) and (u_4, p_4) . The phase portrait for each of these parameter configurations are shown in Figures 4.1a and 4.1b. The fixed points (u_2, p_2) and (u_3, p_3) are both saddle points for α and γ both positive or negative.

Of particular interest to us are the fixed points $(u_3, p_3) = (0, \kappa)$ and $(u_4, p_4) = (u_\infty, p_\infty)$. For $\alpha < 0$ and $\gamma > 0$, the former corresponds to a state where there is tau pathology in the absence of $A\beta$ pathology, as is the case in primary age related tauopathy (PART), where there is a natural accumulation of tau in the medial temporal lobe without the presence of $A\beta$. The dynamics for this case are shown in Figure 4.1c. From Equation (4.1) and Equation (4.9), we see that PART arises naturally and can act as a precursor to secondary tauopathy associated with $A\beta$ interaction. Focussing on Equation (4.11), we see that secondary tauopathy occurs when $\beta > 0$ and the new stable fixed point is given by (u_4, p_4) , instead of (u_3, p_3) , as shown graphically in Figure 4.1b and Figure 4.1d. There are two important features of the system representing secondary tauopathy. First, the carrying capacity of tau is higher as $A\beta$ grows and, second, as a result the velocity at which tau grows is larger in the presence of higher $A\beta$ concentration. This can be shown by computing the eigenvalue around (u_3, u_3) in tau p -axis, which are $\gamma(\kappa + \beta u_\infty)$ for $\beta > 0$ and $\gamma \kappa$ for $\beta = 0$. Furthermore, when we consider the full system, inclusive of atrophy, we see that amyloid has a similar effect on atrophy as on tau. For the fixed points $(u_3, p_3, a_3) = (0, \kappa, 1)$ the eigenvalue corresponding to atrophy is $-\eta \kappa$ and for $(u_4, p_4, a_4) = (u_\infty, p_\infty, 1)$ the eigenvalues for atrophy is $-\eta(\kappa + \beta u_\infty)$. Therefore, for $\beta > 0$ there will be a higher rate of tau-induced atrophy in the presence of $A\beta$.

We know experimentally that $A\beta$ acts to catalyse tau production and therefore the single node condition of $(\alpha > 0, \gamma > 0, \beta = 0)$ is contrived since it is necessarily the case that $\beta > 0$. However, in the spatially extended case, Equation (4.9), under

both $A\beta$ and tau pathology, both (u_2, p_2) and (u_4, p_4) will arise naturally depending on whether $A\beta$ and tau are spatially colocalised. If $A\beta$ and tau are not present in the same region then this is representative of an uncoupled model where $\beta u = 0$ and therefore $p_\infty = \kappa$. Similarly, when $A\beta$ and tau are colocalised, the coupling becomes non-zero, $\beta u > 0$, therefore elevating the carrying capacity of tau.

Overall, from the linear analysis of the ATN model we can make several predictions about the dynamics between $A\beta$, tau and neurodegeneration. First, $A\beta$ increases carrying capacity of tau. Second, $A\beta$ increases the rate of tau accumulation and atrophy. Third, this $A\beta$ affect on atrophy occurs when $A\beta$ and tau are spatially colocalised. In the following sections we will examine the consequences of the these dynamics on parameter identifiability and fit to longitudinal neuroimaging data.

4.3 A Measurement Model for ATN Biomarkers

Similarly to our work in Chapter 3, we aim to account for the baseline properties of PET tracers to allow us to directly model longitudinal neuroimaging data. From a modelling perspective, this requires specifying a measurement model to shift u_i and p_i for $i = 1, \dots, R$ to account for regional baseline values of $A\beta$ and tau PET SUVR. To extend Equation (4.9)–Equation (4.10) to one that accounts for baseline $A\beta$ and tau PET values, we introduce the variables

$$x_i = x_{0,i} + u_i, \quad i = 1, \dots, R, \quad (4.13)$$

$$y_i = y_{0,i} + p_i, \quad i = 1, \dots, R. \quad (4.14)$$

where, x_i and y_i represent $A\beta$ and tau SUVR at the i -th region, and $x_{0,i}$ and $y_{0,i}$ are baseline SUVR values at the i -th region. In addition to this change of variables, we investigate the relationship of regional atrophy with tau SUVR. In Figure 4.2, we show mean atrophy in a A^+T^+ ADNI cohort compared to baseline corrected tau SUVR, $(y_i - y_{0,i})$, Figure 4.2a and tau concentration, SUVR normalised by baseline values and carrying capacities, $(y_i - y_{0,i})/(y_{\infty,i} - y_{0,i})$. This comparison is shown in Figure 4.2. We see that atrophy is more closely correlated tau concentration than baseline corrected tau SUVR and we therefore parametrise tau induced atrophy using regional tau concentration. Incorporating these change into Equation (4.9)–

Equation (4.10), we have for $i = 1, \dots, R$:

$$\frac{dx_i}{dt} = \alpha (x_i - x_{0,i}) [(x_{\infty,i} - x_{0,i}) - (x_i - x_{0,i})], \quad (4.15a)$$

$$\frac{dy_i}{dt} = -\rho_2 \sum_{j=1}^R L_{ij} (y_j - y_{0,j}) + \gamma (y_i - y_{0,i}) [(y_{\infty,i} - y_{0,i}) - (y_i - y_{0,i})], \quad (4.15b)$$

$$y_{\infty,i} = \kappa_i + \beta (x_i - x_{0,i}) \quad (4.15c)$$

$$\frac{da_i}{dt} = \eta \frac{y_i - y_{0,i}}{y_{\infty,i} - y_{0,i}} (1 - a_i), \quad (4.15d)$$

with $x_i(0) = x_{i,0}$, $y_i(0) = y_{i,0}$, $a_i(0) = a_{i,0}$, and where κ_i is the regional tau carrying capacity in the absence of amyloid, $y_{\infty,i}$ is the regional tau carrying capacity in the presence of amyloid, $x_{0,i}$ and $x_{\infty,i}$ are the baseline values and carrying capacities for A β SUVR. Note that we further assume κ , representing the natural vulnerability for tau accumulation, of varies regionally, consistent with non-pathological tau accumulation seen in PART.

As in Section 3.2.1, we wish to fix baseline values and carrying capacities from cross-sectional data to facilitate inference. Note that $y_{\infty,i}$ depends on the unknown parameter β and therefore cannot be estimated ahead of inference. There are four parameters we are able to fix from cross-sectional data, baseline A β and tau fields $\mathbf{x}_0, \mathbf{y}_0$, carrying capacities for A β PET, \mathbf{x}_{∞} , and the PART field κ .




Figure 4.2: **Comparison of atrophy with tau.** **4.2a** Atrophy compared to baseline corrected tau SUVR. **4.2b** Atrophy compared to tau concentration, normalised to baseline values and carrying capacities from Section 3.2.1. Highlighted in red are regions in the atrophy composite regions, the bilateral entorhinal, amygdala, middle temporal, inferior temporal, inferior parietal and precuneus regions [161].

4.3.1 Estimating Fixed Parameters

To estimate \mathbf{x}_0 and \mathbf{x}_∞ we reproduce the analysis of Jack et al. [162] and Whittington et al. [154] using the most recent A β PET data available from ADNI from subjects with at least two scans ($N = 913$). Full demographics for this group are provided in Table 4.3. A β PET were normalised by the ADNI composite reference region and the ADNI summary composite region was used to measure individual A β load [163], both available from the ADNI online portal. Using this measure, we approximate the derivative at each point by calculating the forward difference between an individuals baseline and final scan.

$$\frac{dA\beta}{dt} \approx \frac{A\beta_{\Delta t} - A\beta_{t_0}}{\Delta t}, \quad (4.16)$$

where $A\beta_{t_0}$ is the initial A β load for a subject, Δt is the time between initial and final scan, and $A\beta_{\Delta t}$ is the A β load at the final scan. A β load is taken as the mean A β SUVR in the summary composite region provided by ADNI. The result is shown in Figure 4.3a. The phase-space data are subsequently binned to smooth the noisy derivative estimates and second-degree polynomial is fit by minimising the least squares error (see Figure 4.3a). By integrating the resulting polynomial we retrieve the logistic curve shown in Figure 4.3b and individuals can be assigned a *time* on this curve corresponding to their baseline A β level. Using this temporal ordering, we can reconstruct data at a regional level, as shown in Figure 4.3c for the right inferior temporal lobe, showing sigmoidal dynamics over the A β timeline. From here, we can estimate regional dynamics by fitting a sigmoid curve at the regional level that is parametrised in terms of a baseline value and carrying capacity,

$$f_i(t) = \frac{x_{\infty,i} - x_{0,i}}{1 + \exp[-\varphi_i(t - \tau_i)]} + x_{0,i}, \quad (4.17)$$

where φ_i is a regional production rate, τ_i is the time at which the half-maximal concentration is reached, $x_{0,i}$ is regional baseline value, and $x_{\infty,i}$ is the regional carrying capacity. This model was fit to each region by minimising the least squares error. The carrying capacity field is shown over the right cortex in Figure 4.3e and is consistent with A β staging observed in A β PET data [6].

To determine the PART field, κ , we apply the Gaussian mixture modelling approach used in Section 3.2.1 and [72] to A $^-$ subjects from ADNI ($N = 568$), redsee Table 4.3 for group demographics. A β status was determined using the ADNI composite region for A β PET. Tau positivity threshold are identical to those used in Chapter 3. We restrict our analysis to temporal regions only, since these are regions affected by PART [37]. To examine whether pathological tau SUVR can be reliably

observed, we fit both a one component and a two component Gaussian mixture model and compare them using the Akaike information criteria (AIC) [72]. Regions are considered T^+ if the AIC for the two component model is lower than the one component model. For regions that are T^+ , the PART carrying capacity is then given by the 99th percentile of the T^+ distribution, capturing pathological tau accumulation. An example of this is shown in Figure 4.3d for the right entorhinal cortex. We do not expect significant interhemispheric differences in PART carrying capacities, however, we are limited by small sample size since most PART pathology occurs with $A\beta$ pathology and cannot be easily disentangled from $A\beta$ pathology. Therefore, we enforce symmetry of the PART related carrying capacities by taking the average regional value between hemispheres. The estimated PART field is shown in Figure 4.3f and as expected shows small amounts of tau invasion in the medial temporal lobe and little to no invasion in the lateral temporal lobes.

Group	N. Subject	Age	Gender	Education	CN	MCI	AD
$A\beta$	913	72.78	0.50	16.37	0.50	0.45	0.05
A^-	567	72.63	0.51	16.56	0.66	0.31	0.02
A^+T^-	71	73.88	0.49	16.78	0.66	0.29	0.04
A^+T^+	37	71.13	0.59	16.13	0.29	0.43	0.27

Table 4.3: **Demographics for ADNI cohorts.**

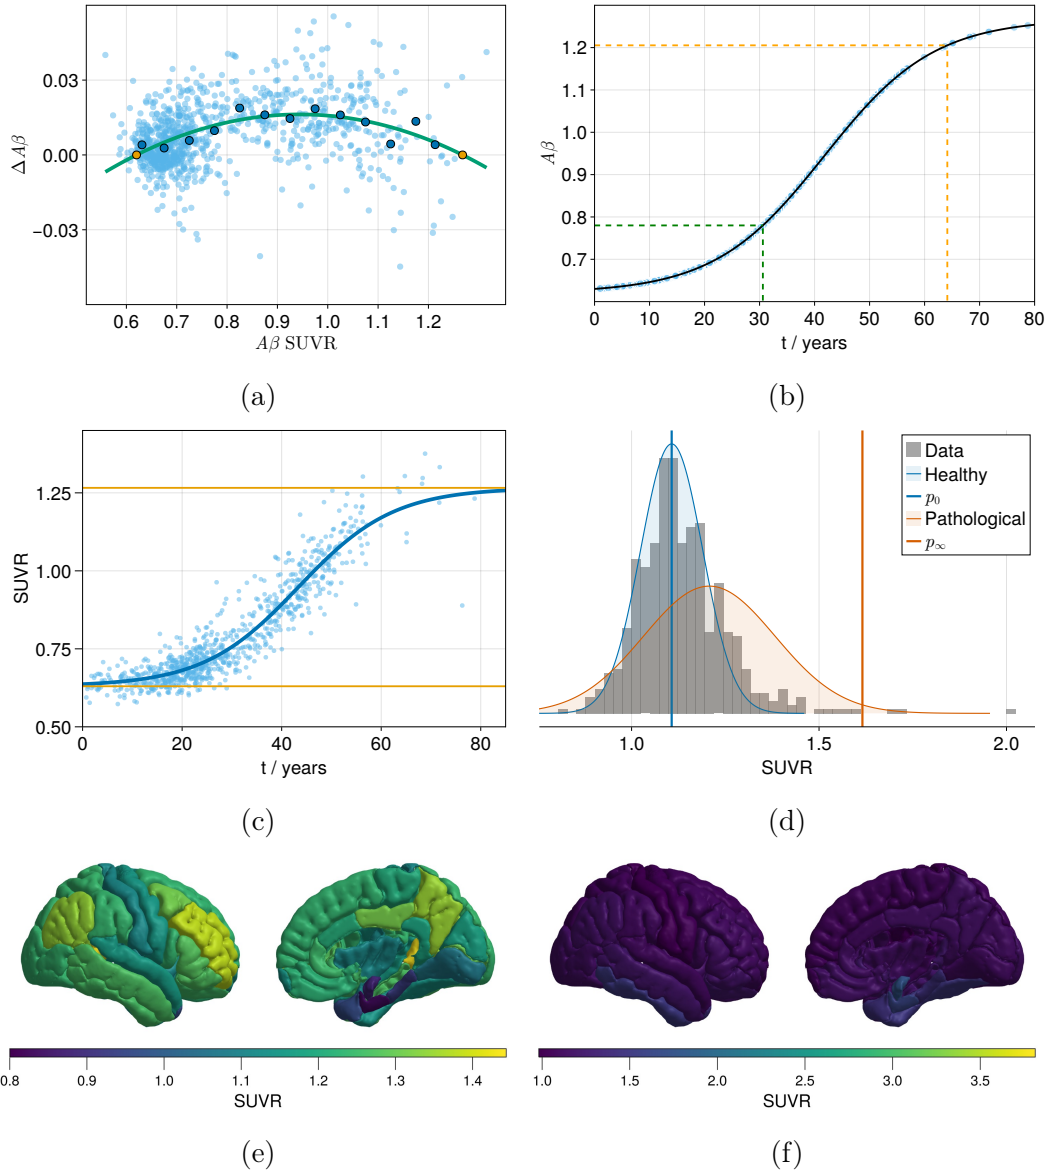


Figure 4.3: **Population analysis for estimated fixed parameters of the ATN model.** **4.3a** $A\beta$ data vs change in $A\beta$ from baseline for $N = 913$ subjects from ADNI with at least two scans. Light blue points represent individual data. Dark blue points are binned data (see methods). Green line shows a fitted second degree polynomial, with roots highlights in yellow, to estimate sigmoidal function of $A\beta$ progression. **4.3b.** Integrated polynomial with subjects ordered according to their time on the $A\beta$ trajectory. Green dashed line shows $A\beta$ positivity threshold (defined by ADNI). Yellow dashed line shows 90% threshold. **4.3c** SUVR in the right inferior temporal region using the subject ordered by $A\beta$ load. Blue line shows a logistic model fit, yellow lines highlight the regional baseline and carrying capacities. **4.3d** Two component Gaussian mixture model applied to right entorhinal cortex from A^{-T+} in ADNI. Blue shows the T^{-} negative distribution and orange shows the T^{+} distribution, with vertical lines highlighting the mean and 99th percentile of the distributions, respectively. **4.3e-4.3f,** right cortical rendering of the $A\beta$ carrying capacities and tau PART fields, respectively.

4.3.2 Simulating ATN Dynamics

Using these fixed parameters, we simulate the dynamics from the ATN model using synthetic parameters, shown in Figure 4.4. For this simulation we assume a uniform seeding concentration for $A\beta$ throughout the cortex of $x_{i,0} = 0.2(x_{\infty,i} - x_{0,i})$ for $i = 1, \dots, R$, and tau seeding of $y_{i,0} = 0.2(y_{\infty,i} - y_{0,i})$ in the bilateral entorhinal cortex $i = 27, 62$ and $y_{i,0} = 0$ elsewhere, and a uniform zero atrophy level, to be representative of early stage AD. The full list of initial conditions and parameters for the simulation is provided in Table 4.4. The model predicts and is consistent with several key features of AD pathology.

Variable	Value
$x_{i,0}$	$0.2(x_{\infty,i} - x_{0,i})$
$y_{i,0}$	$0.2(y_{\infty,i} - y_{0,i})$
$a_{i,0}$	0
Parameter	Value
α	0.75/yr
ρ	0.015/yr
γ	0.5/yr
β	3.75
η	0.1/yr

Table 4.4: **Simulation Setup.** Initial conditions and parameters used for simulating the ATN model.

First, the model is consistent with the staging sequence of $A\beta$, tau and atrophy detailed by the $A\beta$ cascade hypothesis and the ATN biomarker framework [44, 10]. Amyloid quickly accelerates first reaching saturation before cortical tau expansion and widespread atrophy.

Second, and relatedly, amyloid triggers the expansion of tau pathology into the cortex. This can be seen first by comparing Figure 4.4, with $A\beta$ -tau coupling, $\beta > 0$, and Figure 4.5, without $A\beta$ -tau coupling, $\beta = 0$. In Figure 4.5, tau and atrophy are restricted to the MTL, with negligible expansion into the cortex, as seen in PART, whereas in Figure 4.4, tau quickly expands into the cortex in the presence of $A\beta$ pathology, with atrophy following at an accelerated rate.

Third, the model predicts that colocalisation at a threshold of 90% $A\beta$ and 10% tau occurs first in the inferior temporal lobe (closely followed by the precuneus and fusiform) and that the time of this colocalisation (illustrated by the dashed black line in Figure 4.4) coincides with tau acceleration and cortical atrophy.

Fourth, the model predicts that regional $A\beta$ explains regional tau heterogeneity. The end stage distribution of tau shown in Figure 4.4 closely resembles that which has been observed in tau PET studies [6, 4] and in the analysis in Section 3.2.1 and that the staging resembles the Braak staging of tau as observed in tau PET [6]. Overall, the model is consistent with the main predictions of the $A\beta$ cascade hypothesis and closely follows staging dictated by the ATN biomarker staging framework.

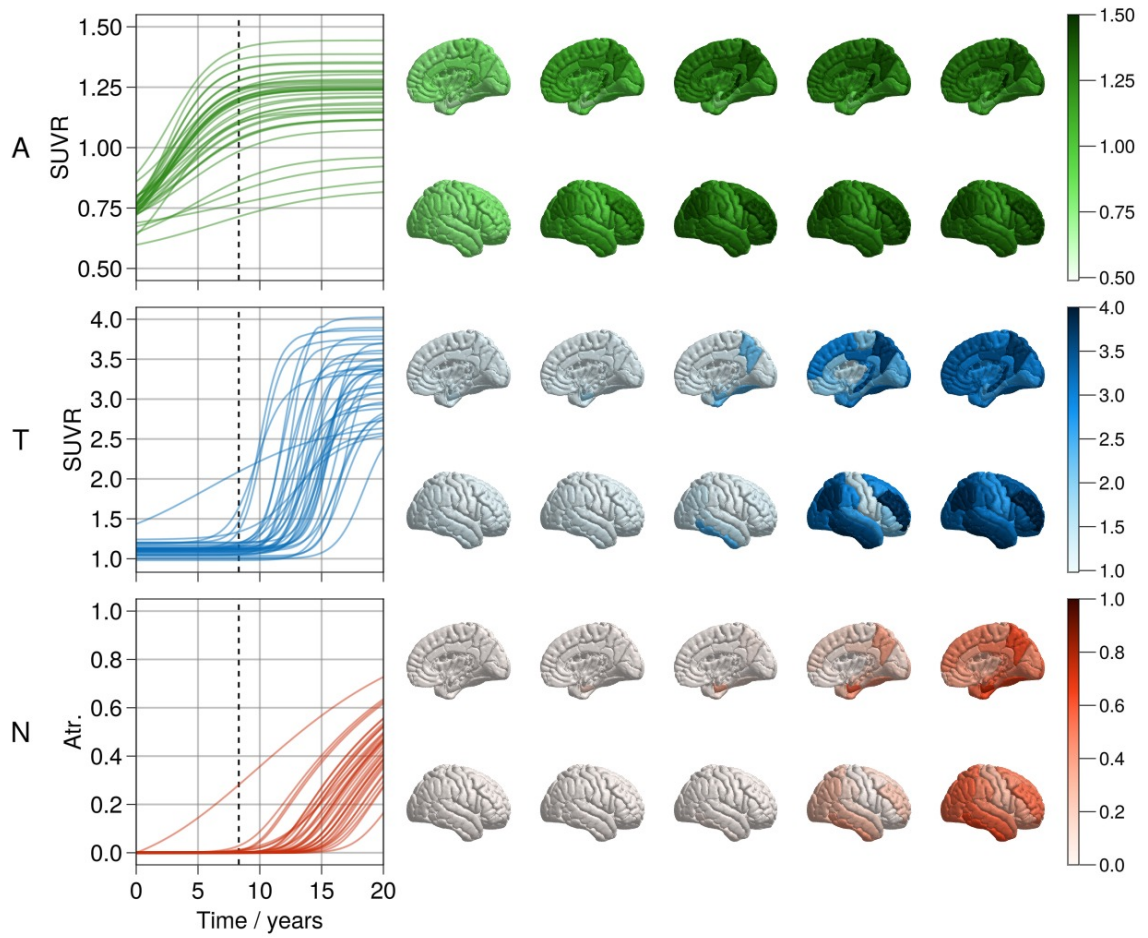


Figure 4.4: **Simulation from the ATN model.** Simulation from the ATN model using seeding of $A\beta$ throughout the cortex of $x_{i,0} = 0.2(x_{\infty,i} - x_{0,i})$ for $i = 1, \dots, R$, and a tau seeding of $y_{i,0} = 0.2(y_{\infty,i} - y_{0,i})$ in the bilateral entorhinal cortex $i = 27, 62$ and $y_{i,0} = 0$ elsewhere. We biophysical synthetic parameters, with tau in a production-dominated regime, $\alpha = 0.75/yr$, $\rho = 0.015/yr$, $\gamma = 0.5/yr$, $\beta = 3.75$, $\eta = 0.1/yr$. The trajectories for the right cortex are shown over a 20 year period and cortical renderings on the right cortex at $t = \{0, 5, 10, 15, 20\}$.

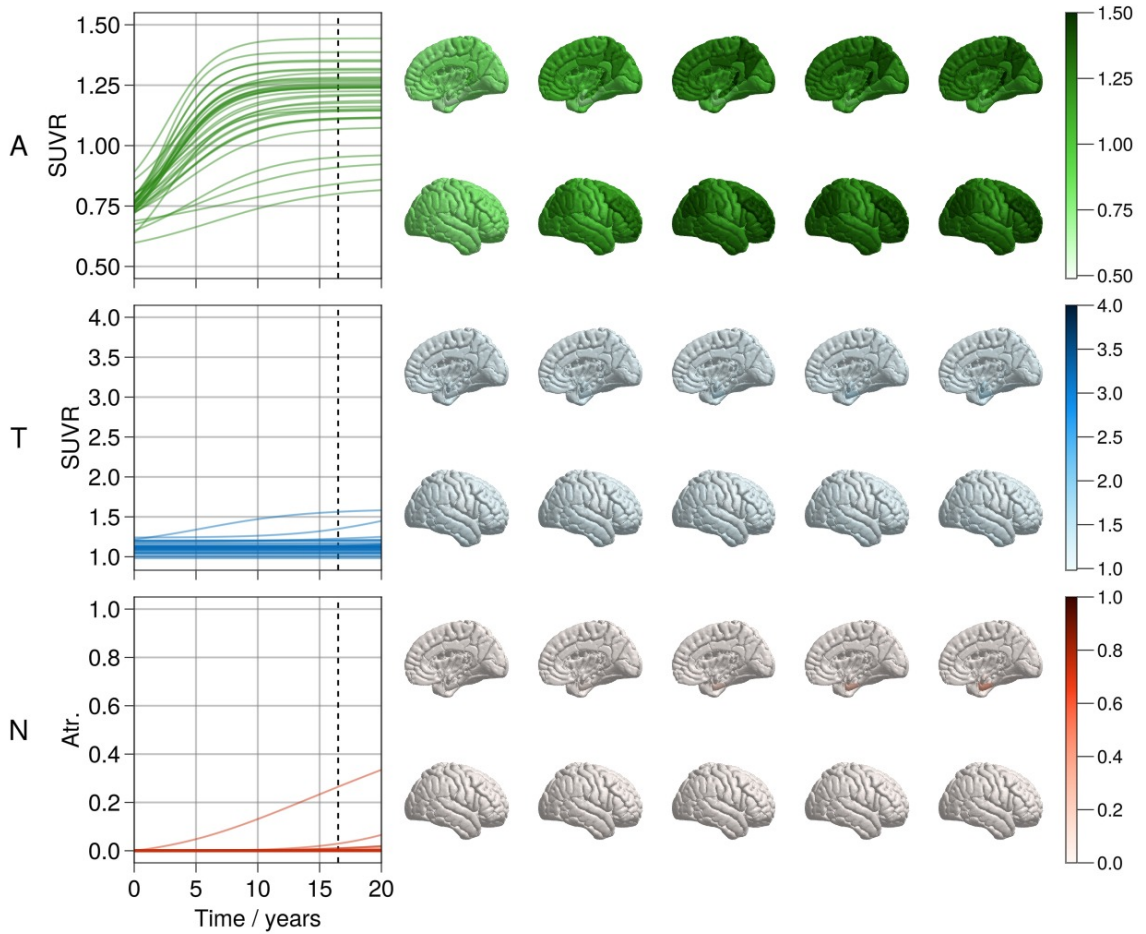


Figure 4.5: **Simulation from the ATN model with no $A\beta$ -tau coupling.** Simulation set up is identical to Figure 4.4 (full details in Table 4.4), with the exception that there is no $A\beta$ -tau coupling, $\beta = 0$. The trajectories for the right cortex are shown over a 20 year period and cortical renderings on the right cortex at $t = \{0, 5, 10, 15, 20\}$.

4.4 Identifiability Analysis

Before performing inference against patient groups, we first examine the structural and practical identifiability of the model. We assess the structural identifiability using the `StructuralIdentifiability.jl` package in the Julia programming language [99]. Global structural identifiability is extremely computationally expensive and we therefore seek to restrict our analysis to a smaller subnetwork of the full brain network. In Chapter 2, we show that it is necessary for local identifiability of tau transport to observe nodes in a connected network. Additionally, we expect the most meaningful signal from real-world data to come from the temporal lobe, since this is where tau first seeds

[41, 6], where amyloid and tau first colocalise [164], and where AD-related atrophy first occurs [25, 26]. Therefore, for this analysis, we choose a subnetwork network bilateral temporal lobes ($R = 22$), since this is well-connected subnetwork across which the bulk of amyloid, tau and neurodegeneration changes are observed across in real-world data. The analysis shows all model parameters, $\alpha, \rho, \gamma, \beta, \eta$, are globally identifiable with a correctness probability of 0.99, and therefore should be practically identifiable provided sufficient data.

To assess the practical identifiability of the model parameters, we look at the sensitivity profiles and synthetic inference using different time windows across observed dynamics. We calculate local sensitivities $D_\theta(\mathbf{p})$ for $\theta \in \{\alpha, \rho, \gamma, \beta, \eta\}$ and $\mathbf{p} \in \{\mathbf{x}, \mathbf{y}, \mathbf{a}\}$, using forward-mode automatic differentiation, as implemented in the DifferentialEquations.jl software package. The simulation procedure is identical to that in Figure 4.4 (see Table 4.4), but over a longer time period, $t = [0, 30]$ to allow for full saturation of $A\beta$, tau and atrophy. The initial conditions and parameters are representative of what we expect to see from observed data, since $A\beta$ is typically seeded throughout the cortex by people before early tau accumulation begins.

The results from the local sensitivity analysis are provided in Figure 4.6. The production parameter for $A\beta$, α , shows sensitivity across the output range and is highest during the exponential growth phase. For the remaining parameters, there is a time delay before any substantial effect are observed, since there is not widespread tau invasion or atrophy until $A\beta$ and tau colocalise in the neocortex. After this delay, all of the tau specific parameters and the amyloid production rate, $\alpha, \rho, \gamma, \beta$ show high sensitivity with respect to tau, and a to a lesser degree, atrophy, decaying with the saturation of tau and atrophy. The sensitivity of β with respect to tau, remains constant after tau has saturated since a change in β will affect the asymptotic solutions for tau. The atrophy coefficient, η only affects the atrophy output, and therefore exhibits no sensitivity with respect to $A\beta$ or tau. Sensitivities for η are positive at an early stage due to initial tau seeding in the entorhinal cortex and after the amyloid related delay, the remaining cortical regions show sensitivity to the parameter. Overall, this analysis indicates that parameters are likely to be identifiable from data after there is sufficient mixing of amyloid and tau.

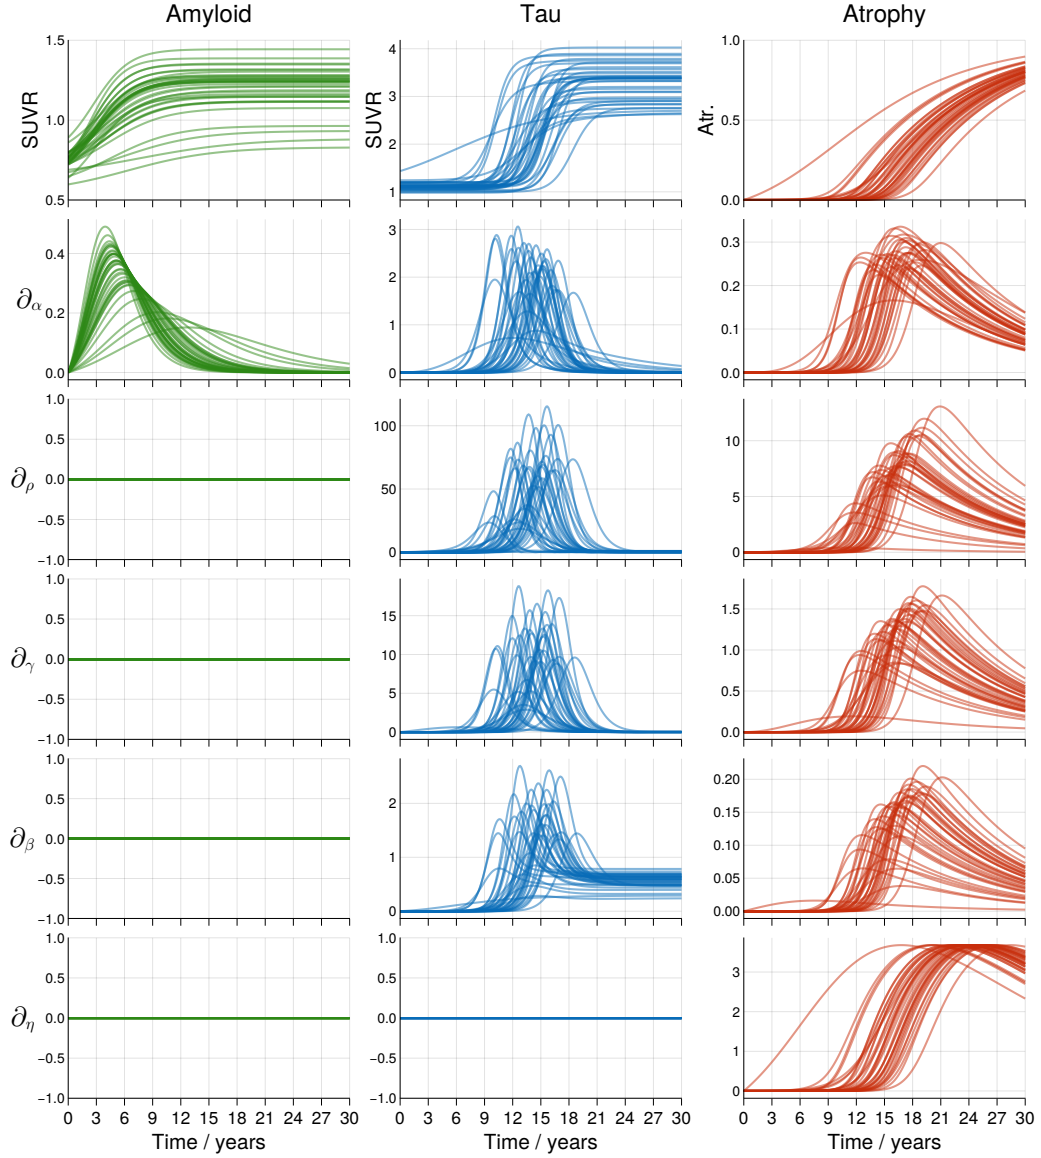


Figure 4.6: **Sensitivity analysis for ATN model.** Top row shows simulated trajectories using the simulation set-up in Table 4.4. Proceeding rows show sensitivities of each variable relating to $A\beta$, tau and atrophy, with respect to model parameters, $\alpha, \rho, \gamma, \beta, \eta$.

As a more informative test of practical identifiability of the model, we simulate data across the dynamical space of the ATN model and see how well parameters are inferred given noisy and sparsely sampled data. Using the same parameters and initial conditions as in Figure 4.4 and Figure 4.6, we generate 10 synthetic data sets by sampling the trajectory at times $\mathbf{t} \in \{(3n, 3n + 1.5, 3n + 3)\}_{n=0}^9$ and adding independent and identically distributed noise sampled from $\sigma \sim \mathcal{N}(0, 0.025)$. We perform inference on these synthetic datasets using NUTS with a target acceptance

rate of 0.8. For each dataset, we collect 4 chains of 2000 samples. We use the priors listed in Table 4.5.

Parameter	Prior	Support
α	$\mathcal{N}(0, 1)$	$[0, \infty]$
ρ	$\mathcal{N}(0, 1)$	$[0, \infty]$
γ	$\mathcal{N}(0, 1)$	$[\infty, \infty]$
β	$\mathcal{N}(0, 3)$	$[0, \infty]$
η	$\mathcal{N}(0, 1)$	$[0, \infty]$
σ	$\Gamma^-(2, 3)$	$[0, \infty]$

Table 4.5: Prior distributions for model parameters for testing identifiability.

The results from this analysis are shown in Figure 4.7. We see that the results mirror the insights gained from the sensitivity analysis. The production parameter for $A\beta$ is identifiable across the entire domain of the simulation, however, as more regions are fully saturated, the posterior distributions become broader. The remaining parameters are well identified after a delay period and converge on the true parameters after $t_0 = 9$ years. Similarly to $A\beta$ parameters, the posterior distributions for the tau and atrophy parameters ρ , γ and η become more broad as tau and atrophy saturate, from $t_0 = 21$ years onwards. The increase in uncertainty is less pronounced compared to the α parameter, and this is likely because of the time delay in tau and atrophy accumulation following $A\beta$ accumulation. Additionally, atrophy rates are informative about tau dynamics rates, therefore aiding identifiability after tau saturates and while atrophy is still progressing. The $A\beta$ -tau coupling parameter β is identifiable from $t_0 = 9$ years onwards does not suffer from increased uncertainty as progression saturates, consistent with our sensitivity analysis, since it uniquely determine the tau carrying capacities. From this analysis we can identify the disease states from which we can reliably infer model parameters, namely those where $A\beta$ and tau are colocalised and neither are saturated. The most important factor determining identifiability whether there is sufficient spatial colocalisation between $A\beta$ and tau. Therefore, this restricts the groups on which we can perform subject-specific inference, precluding A^+T^- groups who do not have tau expansion outside of the medial temporal lobe.

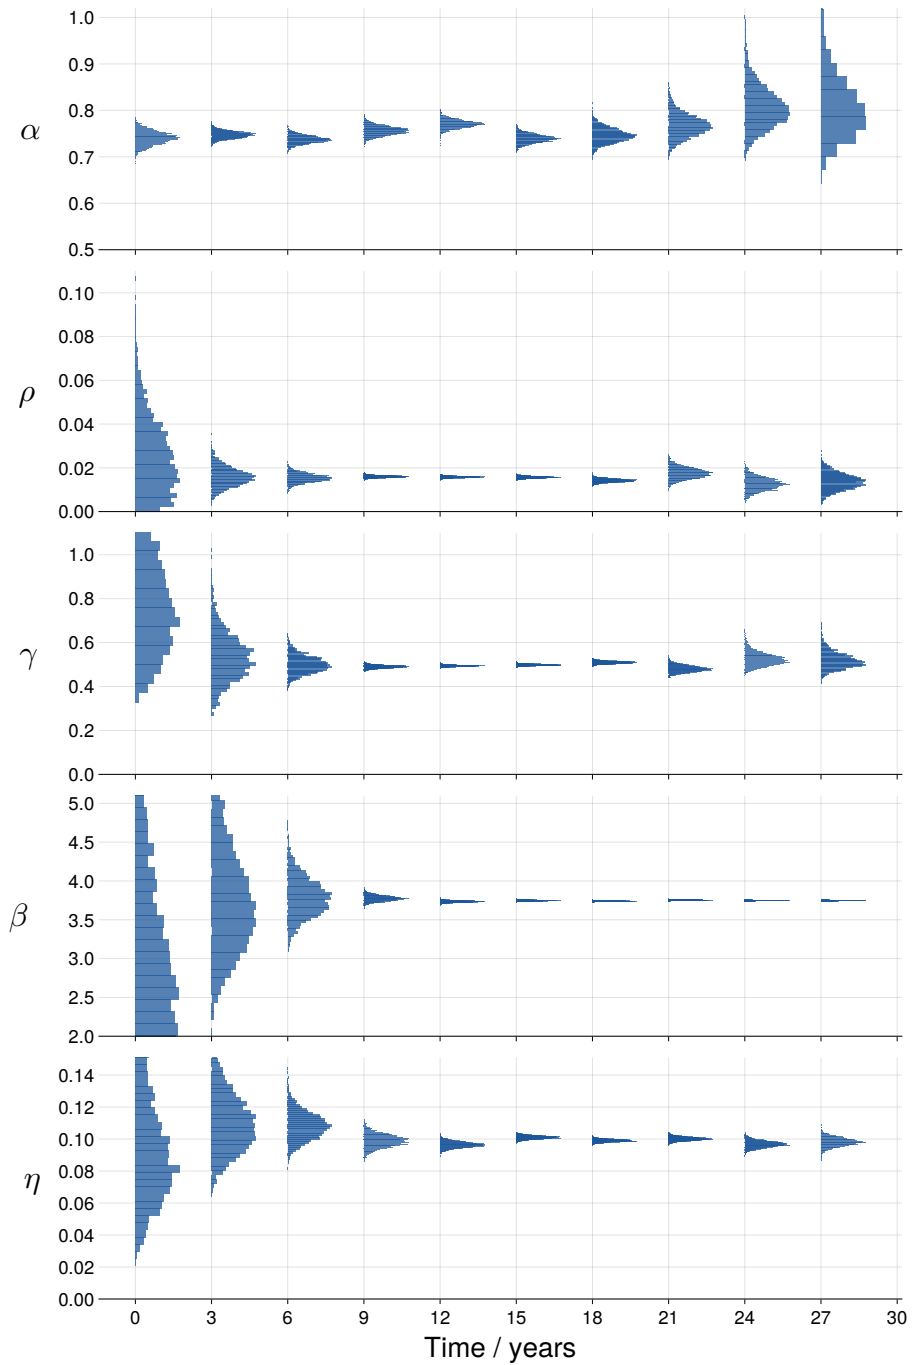


Figure 4.7: **Identifiability the ATN model.** Simulations from the ATN using the initial conditions and parameters in Table 4.4, with samples saved at intervals of three time points starting from such that $\mathbf{t} \in \{(3n, 3n + 1.5, 3n + 3)\}_{n=0}^9$. Samples trajectories had i.i.d Gaussian noise sampled from $\mathcal{N}(0, 0.025)$. The posterior distributions for each of the parameters are aligned vertically at the initial time for each synthetic data set.

4.5 Inference with the Coupled Model

We now compare our ATN model with longitudinal biomarker data from ADNI. We use longitudinal A β PET, tau PET and sMRI to calibrate each of the elements in the ATN model.

We select A^+T^+ subjects with at least two A β scans and at least three tau scans. We retrieve data corresponding to the regions in the DKT atlas. Since A β PET provides particularly poor signal in the bilateral hippocampus, these regions are excluded from the analysis, leaving $R = 70$ regions. Regional A β PET data, denoted \mathbf{Y}_A , are normalised to lie between the baseline and carrying capacities, \mathbf{x}_0 and \mathbf{x}_∞ respectively, determined from cross-sectional population-level analysis in Section 4.3. Similarly, tau PET data, denoted \mathbf{Y}_T , are normalised relative to their baseline values, \mathbf{Y}_0 only, as calculated in Section 3.2.1, since their carrying capacities are unknown before inference. Regional volumetric data, $\mathbf{V}^n \in \mathbb{R}^{R \times T_T^n}$ for $n = 1, \dots, N$, are provided with each PET scan and these are used to measure regional neurodegeneration, \mathbf{Y}_N . Regional volumes are normalised to their baseline total intracranial volume. Then, all volume measures for an individual divided by their values at the baseline scan and neurodegeneration is measured as $\mathbf{Y}_N = \mathbf{1} - \mathbf{V}_t^n$ for $t = 1 \dots T_T^n$ and $n = 1, \dots, N$, where $\mathbf{V}_t^n \in \mathbb{R}^R$ is the t th scan for the n th subject. Therefore, at the initial scan for each subject, the measured atrophy per region is 0 and increases as brain volumes decrease.

In total, we have $N = 37$ A^+T^+ subjects, denoted $\mathbf{Y} = \{(\mathbf{Y}_A^n, \mathbf{Y}_T^n, \mathbf{Y}_N^n)\}_{n=1}^N$, where subscripts A, T and N are used to denote biomarker data for each of the ATN biomarkers. Each subject has T_A^n A β PET scans, and T_T^n tau PET scans and sMRI scans, for $n = 1, \dots, N$. Note that we use the corresponding sMRI scans obtained with tau PET scans. Subjects may have different measurement times for A β PET scans and tau PET scans. Therefore, we have $\mathbf{t}_A^n \in \mathbb{R}^{T_A^n}$ and $\mathbf{t}_T^n \in \mathbb{R}^{T_T^n}$ for $n = 1, \dots, N$ subjects A β PET and tau PET scan times, respectively, and $\mathcal{T} = \{(\mathbf{t}_A^n, \mathbf{t}_T^n)\}_{n=1}^N$, total observation times for each of N subjects. Additionally, we denote the initial conditions as the vector $\mathbf{y}_0^n \in \mathbb{R}^{3R}$, comprising the initial A β PET, tau PET and sMRI values for the n th subject.

For the n th subject the data generating function is:

$$\mathbf{Y}_A^n = \mathbf{f}_A(\mathbf{Y}_0^n, \Theta^n, \mathbf{t}_A^n) + \epsilon_A, \quad (4.18)$$

$$\mathbf{Y}_T^n = \mathbf{f}_T(\mathbf{Y}_0^n, \Theta^n, \mathbf{t}_T^n) + \epsilon_T, \quad (4.19)$$

$$\mathbf{Y}_N^n = \mathbf{f}_N(\mathbf{Y}_0^n, \Theta^n, \mathbf{t}_T^n) + \epsilon_N \quad (4.20)$$

where $\mathbf{f}_A, \mathbf{f}_T, \mathbf{f}_N$ are solutions of ATN model, Equations (4.15a) to (4.15d), respectively, $\Theta = \{(\alpha^n, \rho^n, \gamma^n, \beta^n, \eta^n)\}_{n=1}^N$ are the model parameters for the n th subject, and $\epsilon_A, \epsilon_T, \epsilon_N$ represent measurement noise process specific to each biomarker.

We assume that each biomarker has independent and identically distributed Gaussian noise, parametrised by $\sigma = \{\sigma_A, \sigma_T, \sigma_N\}$. Our generative model for the n th and each biomarker is then:

$$\mathbf{Y}_A^n \sim \mathcal{N}(\mathbf{f}_A(\mathbf{Y}_0^n, \theta^n, \mathbf{t}_A^n), \sigma_A^2 \mathbf{I}), \quad \epsilon_A \sim \mathcal{N}(0, \sigma_A), \quad (4.21)$$

$$\mathbf{Y}_T^n \sim \mathcal{N}(\mathbf{f}_T(\mathbf{Y}_0^n, \theta^n, \mathbf{t}_T^n), \sigma_T^2 \mathbf{I}), \quad \epsilon_T \sim \mathcal{N}(0, \sigma_T), \quad (4.22)$$

$$\mathbf{Y}_N^n \sim \mathcal{N}(\mathbf{f}_N(\mathbf{Y}_0^n, \theta^n, \mathbf{t}_N^n), \sigma_N^2 \mathbf{I}), \quad \epsilon_N \sim \mathcal{N}(0, \sigma_N). \quad (4.23)$$

As in ??, we extend this model to a hierarchical Bayesian model to explicitly model inter-individual differences by introducing population-level parameters $\Omega = \{\alpha_\mu, \alpha_\sigma, \rho_\mu, \rho_\sigma, \gamma_\mu, \gamma_\sigma, \beta_\mu, \beta_\sigma, \eta_\mu, \eta_\sigma\}$, upon which each of the respective model parameters depend. The log likelihood for a single subject under the hierarchical model is given by:

$$\begin{aligned} \log p(\mathbf{Y}, \Theta^n \mid \Omega, \sigma, \mathbf{Y}_0^n, \mathcal{T}^n) &= \log p(\mathbf{Y}_A^n \mid \Theta^n, \sigma_A, \mathbf{Y}_0^n, \mathbf{t}_A^n) + \\ &\quad \log p(\mathbf{Y}_T^n \mid \Theta^n, \sigma_T, \mathbf{Y}_0^n, \mathbf{t}_T^n) + \\ &\quad \log p(\mathbf{Y}_N^n \mid \Theta^n, \sigma_N, \mathbf{Y}_0^n, \mathbf{t}_N^n) + \\ &\quad \log p(\Theta^n \mid \Omega), \end{aligned} \quad (4.24)$$

where $\mathbf{Y} = \{(\mathbf{Y}_A^n, \mathbf{Y}_T^n, \mathbf{Y}_N^n)\}_{n=1}^N$ is the collection of all ATN biomarker data in the cohort. The first three terms represent the log-likelihood contribution for each of the ATN components in the model given their respective longitudinal data, and the last term represents the contribution from the population-level model. This can then be extended to the full cohort of A^+T^+ subjects by accumulating the log likelihood. The full posterior is then given by:

$$\log p(\Theta, \Omega, \sigma \mid \mathbf{Y}, \mathcal{T}, \mathbf{Y}_0) \propto \sum_{n=1}^N \log p(\mathbf{Y}, \Theta^n \mid \Omega, \sigma, \mathbf{Y}_0^n, \mathcal{T}^n) + \log p(\Omega, \sigma). \quad (4.25)$$

In addition to this model, where each individual has a shared coupling parameter, β , we test a model in which individuals share a coupling parameter. For this model, the log likelihood for a single subject is given by:

$$\begin{aligned} \log p(\mathbf{Y}, \Theta^n \mid \Omega, \sigma, \beta, \mathbf{Y}_0^n, \mathcal{T}^n) &= \log p(\mathbf{Y}_A^n \mid \Theta^n, \sigma_A, \beta, \mathbf{Y}_0^n, \mathbf{t}_A^n) + \\ &\quad \log p(\mathbf{Y}_T^n \mid \Theta^n, \sigma_T, \beta, \mathbf{Y}_0^n, \mathbf{t}_T^n) + \\ &\quad \log p(\mathbf{Y}_N^n \mid \Theta^n, \sigma_N, \beta, \mathbf{Y}_0^n, \mathbf{t}_N^n) + \\ &\quad \log p(\Theta^n \mid \Omega), \end{aligned} \quad (4.26)$$

and the full posterior is

$$\log p(\Theta, \Omega, \sigma, \beta, | \mathbf{Y}, \mathcal{T}, \mathbf{Y}_0) \propto \sum_{n=1}^N \log p(\mathbf{Y}, \Theta^n | \Omega, \sigma, \beta, \mathbf{Y}_0^n, \mathcal{T}^n) + \log p(\Omega, \sigma, \beta). \quad (4.27)$$

We will refer to model Equation (4.25) as the *individualised* model and Equation (4.27) as the *pooled* model. We place weakly informative priors on the model parameters, provided in Table 4.6. For model parameters that are strictly positive we use standard Lognormal priors as continuous positive distributions skewed toward smaller parameter values that produce dynamics on the scale of years. For the coupling parameter β , we use a half-Normal prior of $N^+(3, 3)$. This represents our belief that carrying capacities for tau are dictated by regional $A\beta$. In Chapter 3, we estimate carrying capacities for regional tau with a mean regional SUVR of 2.7 and standard deviation of 0.45, however, these values are inclusive of atrophy effects and therefore are lower than we might expect the asymptotic values of tau SUVR to be in the absence of atrophy. The prior on β therefore represents our belief that the carrying capacity given by $A\beta$ -tau interactions should be slightly higher than the carrying capacities estimated in Chapter 3.

Parameter	Prior	Support
α_μ	Lognormal(0, 1)	$[0, \infty]$
α_σ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
ρ_μ	Lognormal(0, 1)	$[0, \infty]$
ρ_σ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
γ_μ	Lognormal(0, 1)	$[0, \infty]$
γ_σ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
β_μ	$\mathcal{N}^+(3, 3)$	$[0, \infty]$
β_σ	$\mathcal{N}^+(0, 3)$	$[0, \infty]$
η_μ	Lognormal(0, 1)	$[0, \infty]$
η_σ	$\mathcal{N}^+(0, 1)$	$[0, \infty]$
α_i	$\mathcal{N}^+(\alpha_\mu, \alpha_\sigma)$	$[0, \infty]$
ρ_i	$\mathcal{N}^+(\rho_\mu, \rho_\sigma)$	$[0, \infty]$
γ_i	$\mathcal{N}^+(\gamma_\mu, \gamma_\sigma)$	$[0, \infty]$
β_i	$\mathcal{N}^+(\beta_\mu, \beta_\sigma)$	$[0, \infty]$
β	$\mathcal{N}^+(3, 3)$	$[0, \infty]$
η_i	$\mathcal{N}^+(\eta_\mu, \eta_\sigma)$	$[-\infty, \infty]$
σ_A	$\Gamma^-(2, 3)$	$[0, \infty]$
σ_T	$\Gamma^-(2, 3)$	$[0, \infty]$
σ_N	$\Gamma^-(2, 3)$	$[0, \infty]$

Table 4.6: Prior distributions for hierarchical model parameters.

To sample from the posterior distributions we use the NUTS algorithm with an acceptance rate of 0.8. For the individualised model we collect a single chain with 1000 samples. For the pooled model we collect four chains starting from different initial parameters with 1000 samples each. For all chains we observe no numerical errors associated with the Hamiltonian and $0.99 \leq \hat{r} \leq 1.01$, showing the chains converged.

The posterior distributions for both models are shown in Figures 4.8 and 4.9. The parameters are broadly similar between the two models, with the only exception coming from β and γ . In the individualised model, β is lower and γ is higher compared to the model with pooled β parameters. The individualised model shows little movement away from the prior for the β parameter. This indicates that there

is insufficient information in the data estimate of $A\beta$ -tau coupling coefficients at the individual-level. In contrast, the pooled model shows considerably less uncertainty in the pooled coupling coefficient, β . Interestingly, this may indicate that the coupling dynamics between $A\beta$ and tau do not vary widely and are largely conserved across the population. To further assess this, we next examine the posterior predictive model fit.

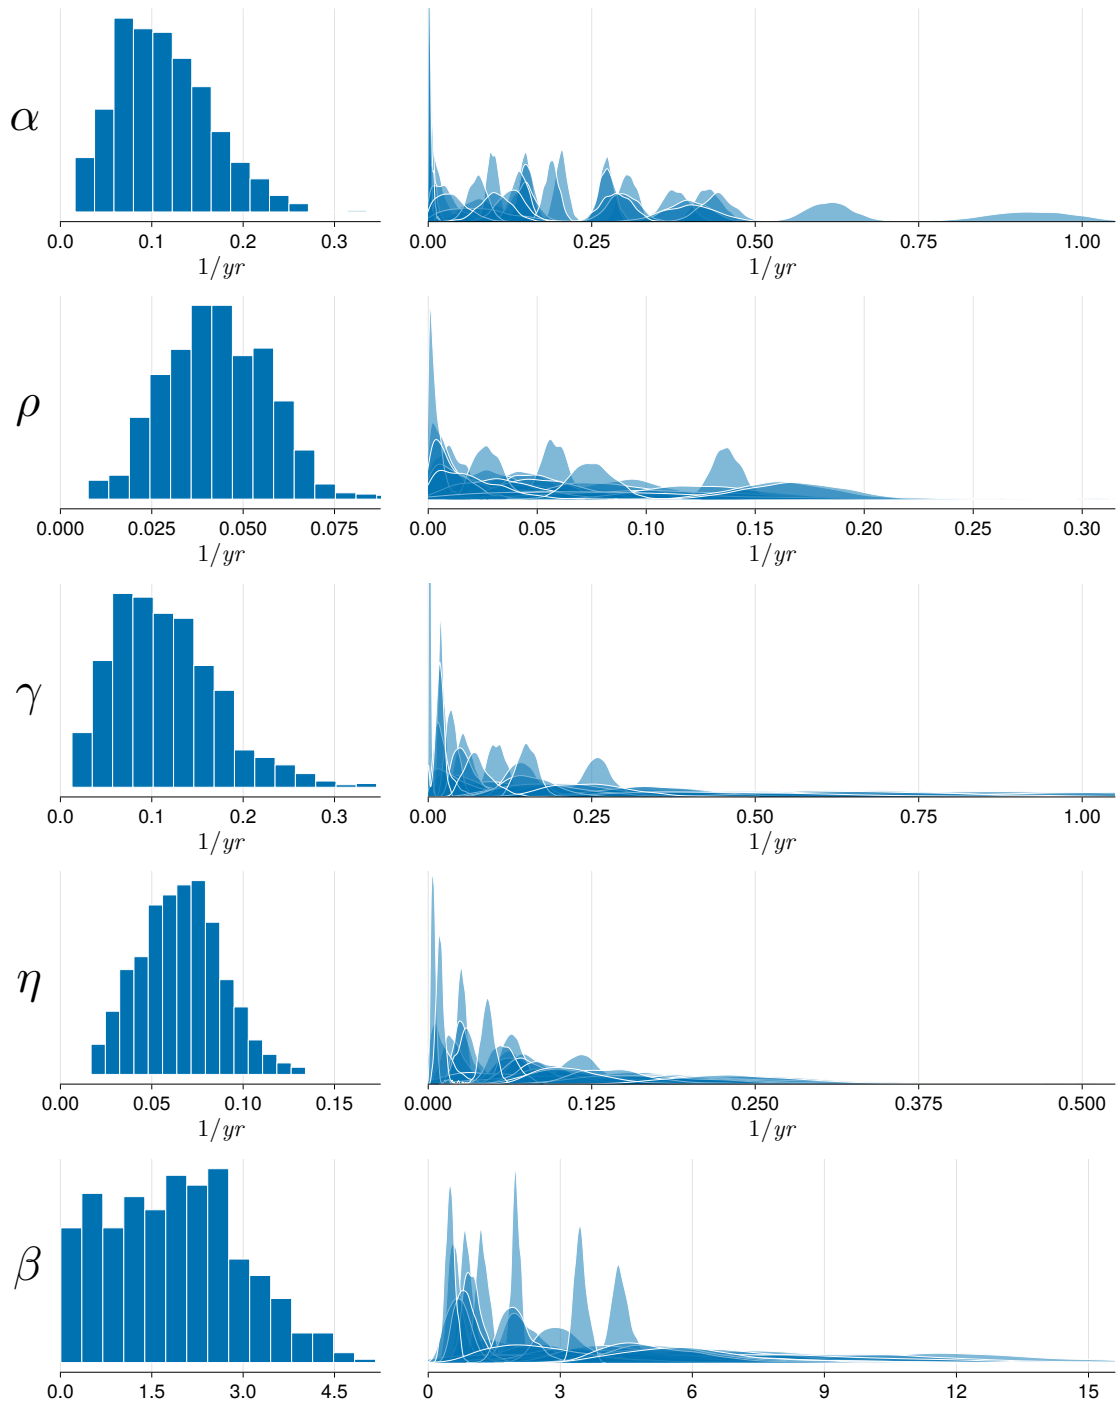


Figure 4.8: **Posterior distributions from individualised model.** Posterior distributions for ATN model parameters using a probabilistic model containing individual $A\beta$ -tau coupling parameters. On the left are population level mean parameters, $\{\alpha_\mu, \rho_\mu, \gamma_\mu, \eta_\mu, \beta_\mu\}$, with corresponding individual-level model parameters on the right.

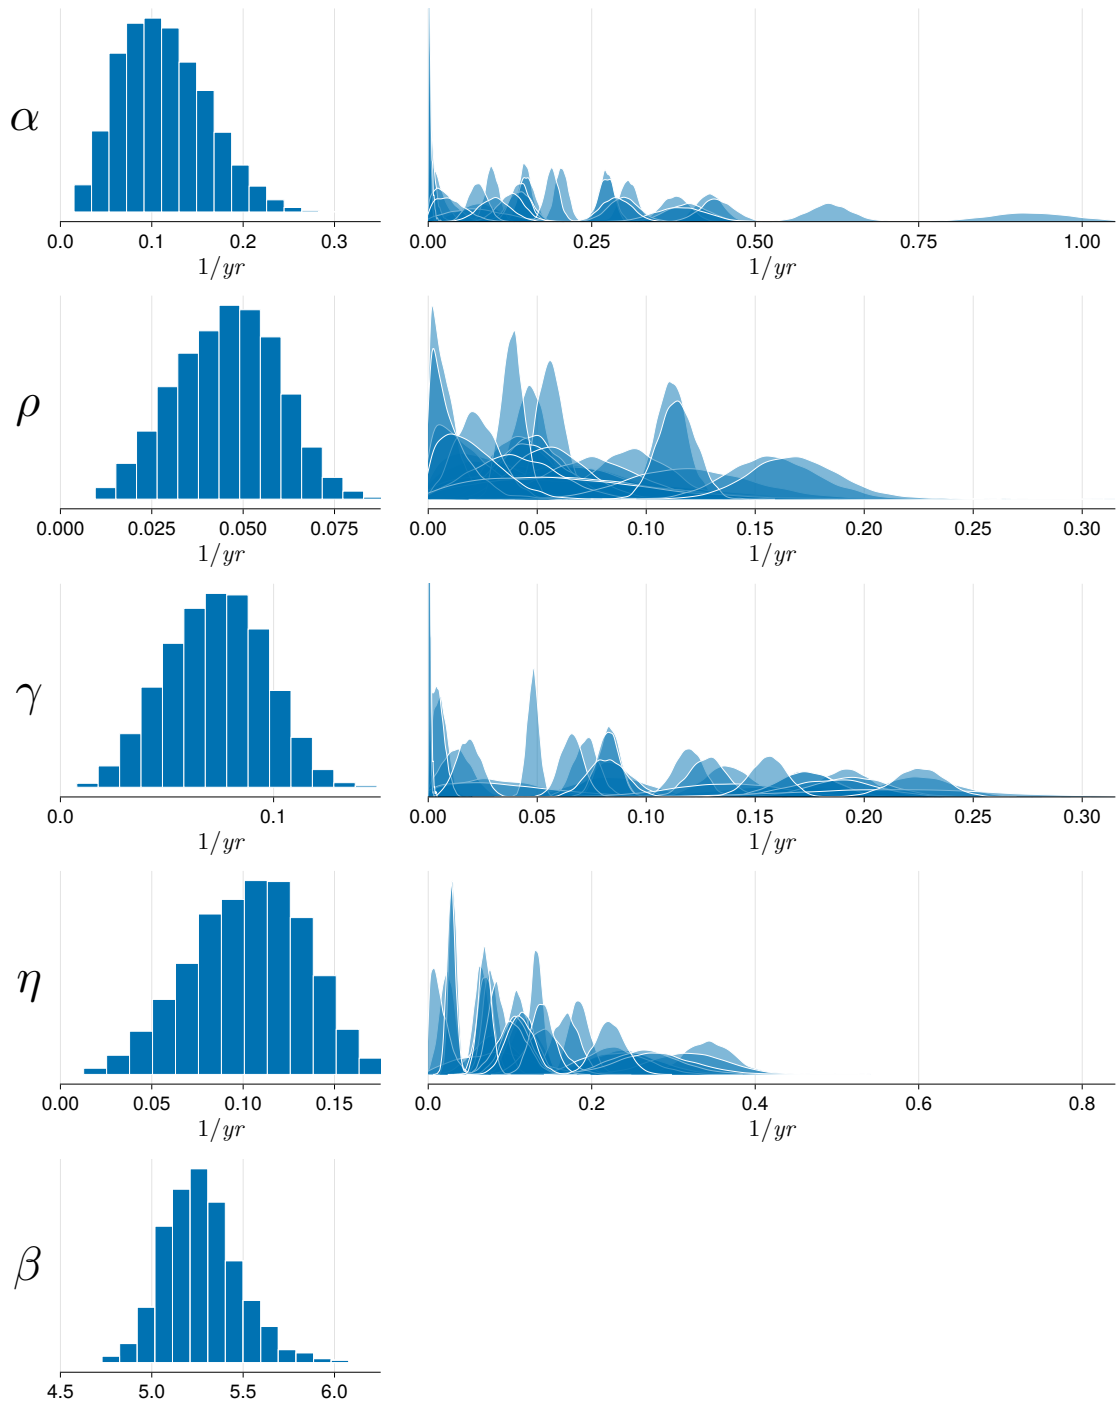


Figure 4.9: **Posterior distributions from pooled model.** Posterior distributions for ATN model parameters using a probabilistic model with shared $A\beta$ -tau couplign parameter. On the left are population-level mean parameters and the shared coupling parameter, $\{\alpha_\mu, \rho_\mu, \gamma_\mu, \eta_\mu, \beta_\mu\}$, with corresponding individual-level model parameters on the right.

Given the substantial uncertainty present in the model with individual coupling we do not examine the posterior predictive fits from this model. Instead, we focus our attention on the pooled coupling model, Equation (4.27), which demonstrates better identifiability. The model fit for the pooled models is shown in Figure 4.10. $A\beta$ PET data is captured well by the model, with increasing error for higher changes in SUVR. The mean predictions for tau PET show that regions tend to be under-predicted as SUVR change increases, but are well captured for small changes. The model is particularly good at predicting structural changes in key regions that are affected by tau-induced atrophy, (highlighted in red).

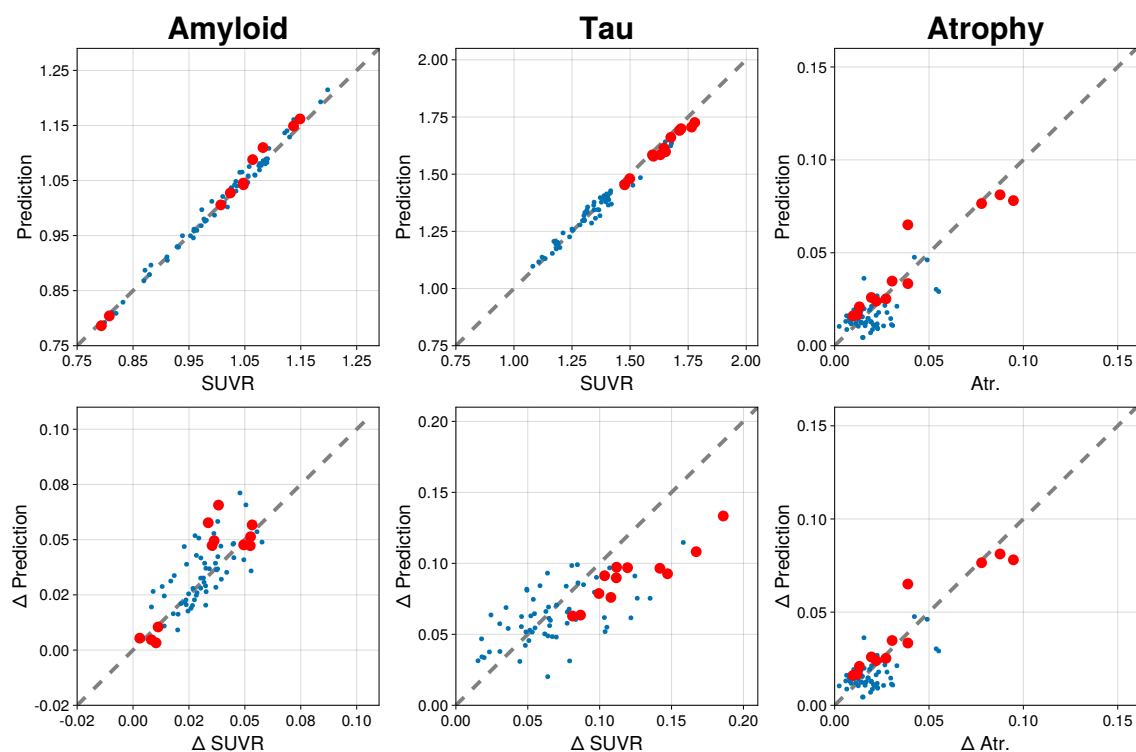


Figure 4.10: **Posterior predictions from the ATN model calibrated to A^+T^+ subjects from ADNI.** Top row shows the predicted vs observed SUVR for $A\beta$, tau and atrophy. Bottom row are predicted change vs observed change for $A\beta$, tau and atrophy. Each point represents a region in the parcellation averaged over subjects. Highlighted in red are regions in the atrophy composite regions, the bilateral entorhinal, amygdala, middle temporal, inferior temporal, inferior parietal and precuneus regions [161].

4.6 Epicentre of Amyloid-Tau Colocalisation

Using the ATN model, we seek to further understand the spatio-temporal relationship of $A\beta$ and tau interaction. It is widely understood that the entorhinal cortex is the initial seeding site for tau pathology [6, 41]. However, there is substantial evidence suggesting that tau pathology does not spread into the neocortex unless there is substantial amyloid burden already present [37, 39, 158, 165, 166, 167]. However, there is a spatial complementarity in the deposition of $A\beta$ and tau during early AD, with $A\beta$ in the neocortex and tau in the medial temporal lobe [6, 41, 168], suggesting that AD accelerates only when $A\beta$ and tau become spatially colocalised. Furthermore, the ATN model predicts $A\beta$ -tau colocalisation facilitates tau expansion into the neocortex (Section 4.3) and for the synthetic parameters predicts the inferior temporal as the site of initial colocalisation (with the precuneus and fusiform closely following). In this section, we seek to investigate $A\beta$ -tau colocalisation using real subject parameters and initial conditions from subjects on the AD spectrum.

We apply the ATN model and posterior samples obtained in Section 4.5 to A^+T^- subjects, i.e. subjects who are early in the disease progression and for whom $A\beta$ and tau deposition are not yet colocalised and therefore provide realistic biophysical initial conditions for the $A\beta$ and tau fields. To correct for spurious signal quantification in the frontal pole, temporal pole and banks of the superior sulcus [55], we average the signal in those regions with those of neighbouring regions. The regions used for correction are provided in Table 4.7.

Region	Neighbouring Regions
Temporal Pole	inferior temporal, middle temporal, superior temporal, entorhinal
Frontal Pole	lateral orbitofrontal, medial orbitofrontal, rostral middle frontal, superior frontal
Banks. Temp Sulcus.	Inferior parietal, middle temporal, superior temporal

Table 4.7: Table of regions that are affected by poorly defined boundaries and signal leakage [55], and adjacent regions that are used to adjust their signal.

Using these initial conditions, we simulate trajectories using samples of the posterior distributions from the pooled model obtained in Section 4.5. From these sim-

ulations, we calculate the arrival time of $A\beta$ and tau in each region for a given concentration threshold. The *concentration* of $A\beta$ or tau in a region is given by the normalisation to $[0, 1]$ interval,

$$\bar{x}_i = \frac{x_i - x_{0,i}}{x_{\infty,i} - x_{0,i}}, \quad \bar{y}_i = \frac{y_i - y_{0,i}}{y_{\infty,i} - y_{0,i}}, \quad (4.28)$$

where $x_{0,i}$, $x_{\infty,i}$, $y_{0,i}$ and $y_{\infty,i}$ are the baseline values and carrying capacities of $A\beta$ and tau, respectively, for $i = 1, \dots, R$. In this case, $y_{\infty,i}$ represents the asymptotic state of y_i inclusive of the transport effect, that is, the asymptotic solution to Equation (4.15b) for particular values posterior parameters. We calculate the regional tau carrying capacity, \mathbf{y}_{∞} , which depends on the amyloid-tau coupling, for each posterior sample of model parameters using a non-linear root-finding algorithm.

Using $N = 37$ A^+T^- subjects from ADNI, we initialise the $A\beta$ field using the average $A\beta$ SUVR, normalised to the ADNI composite REF. We initialise the tau field using the average tau SUVR across subjects, and set regions with less than 1.5% signal set to zero to avoid spurious signal due to noise. For both $A\beta$ and tau, we symmetrise the initial conditions across hemispheres by setting the regional signal to the average of each region across hemispheres and conduct further analysis on the right hemisphere. The initial conditions for $A\beta$ and tau on the right hemisphere are shown in Figure 4.11a and Figure 4.11b, respectively, normalised to their relative regional concentration, (Equation (4.28)). We see that $A\beta$ shows high activation across all regions, with particularly strong elevation in the parietal, frontal, and cingulate cortices. The tau field shows low binding restricted to the medial temporal lobe (entorhinal cortex, amygdala, parahippocampus), with the highest concentration in the entorhinal cortex.

Using samples from the posterior distribution, we sample posterior predictive trajectories using the initial conditions in Figures 4.11a and 4.11b to determine regional *colocalisation times*. Simulated trajectories using the mean posterior parameters are shown in Figure 4.11c and Figure 4.11d, respectively, with trajectories from the entorhinal cortex highlighted in blue and the inferior temporal lobe in red. For both $A\beta$ and tau, the entorhinal cortex houses a higher initial concentration, however, in both cases the trajectory for the inferior temporal lobe surpasses it due to greater acceleration resulting from having higher $A\beta$ carrying capacities. To calculate the time at which $A\beta$ and tau colocalise, we calculate the arrival times in a region given a certain threshold. For $A\beta$ to have a significant catalytic effect on tau, there should be a high concentration of $A\beta$, coupled with a low concentration of tau. Therefore, we

use an initial detection threshold of 90% concentration for $A\beta$ and 10% concentration for tau.

Using the mean parameters from posterior samples, we first calculate the *colocalise order*, that is, the temporal sequence in which each region reaches the $A\beta$ and tau thresholds, shown in Figure 4.11e. The colocalisation order for all regions in the right hemisphere is listed in Table 4.8. For the mean posterior parameters and the average A^+T^- initial conditions, the first region to reach the colocalisation threshold is the inferior temporal lobe. Invasion thereafter largely follows surrounding areas in the lateral occipital and parietal regions, then followed by lateral frontal cortex. The last regions to colocalise are those in the medial temporal lobe. This is likely due to the absence of significant amyloid involvement until late stages, therefore decreasing the production of tau.

To characterise the uncertainty in determining colocalisation sites, we simulate trajectories for each set of parameter samples in the posterior distribution and calculate the region at which $A\beta$ and tau colocalise at a threshold of 90% $A\beta$ and 10% tau. Figure 4.11f shows probability of a given site being the first region of $A\beta$ -tau colocalisation. The full list of regions and their colocalisation probability are provided in Table 4.9. There are four hubs in which colocalisation occurs, the medial temporal lobe, lateral temporal lobe, lateral parietal cortex and frontal cortex, with the fusiform and inferior temporal lobe dominating the probability of first colocalisation.

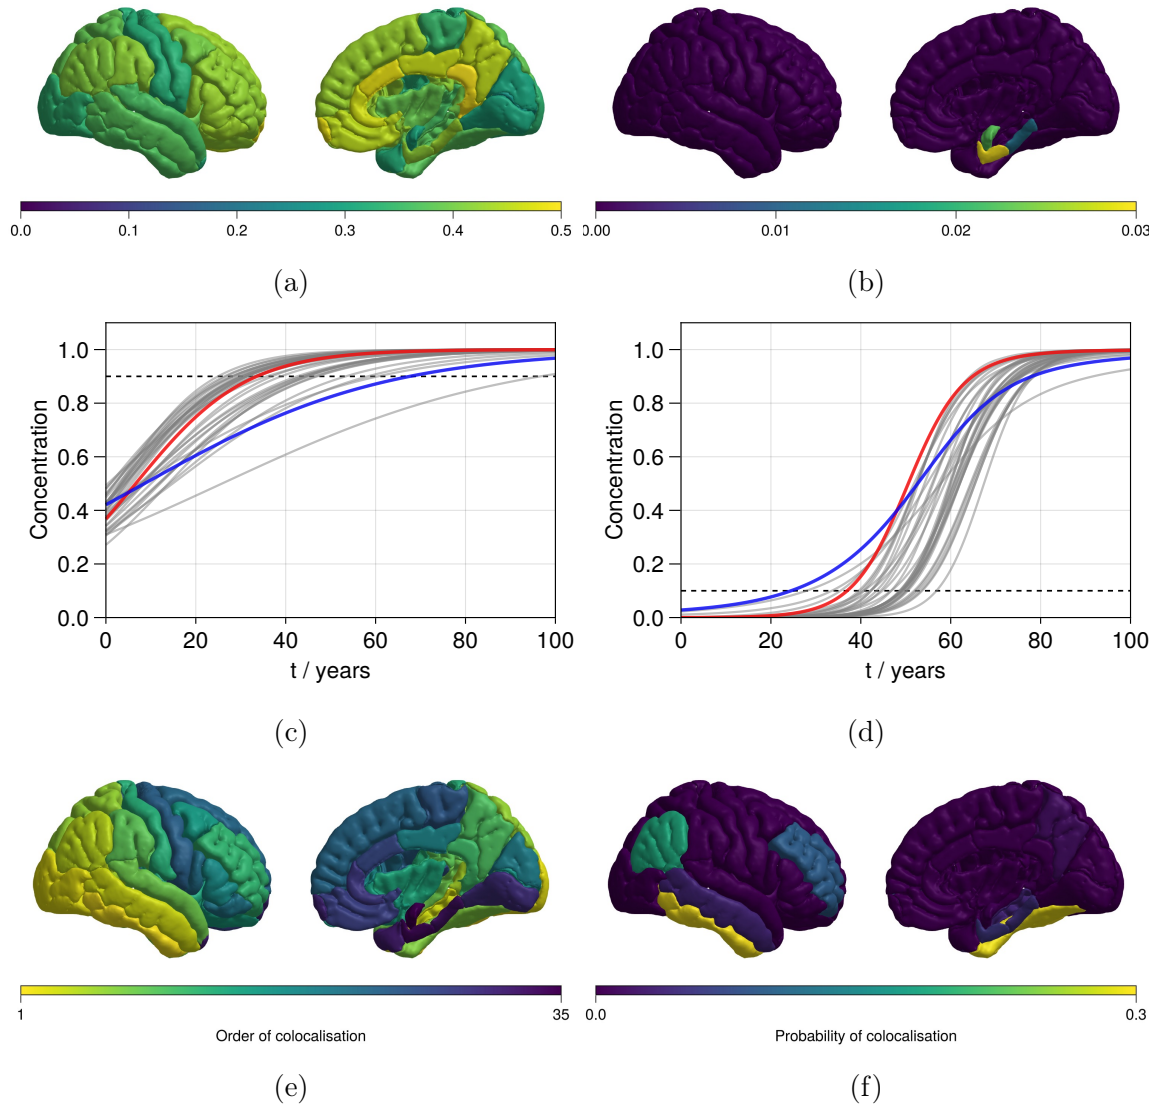


Figure 4.11: **Quantifying A β -tau colocalisation sites.** **4.11a** shows the average concentration of A β in the A^+T^- ADNI cohort. **4.11b** shows the average regional concentration of tau in the A^+T^- ADNI cohort. **4.11c**—**4.11d** show simulated trajectories for A β and tau respectively, using the mean A β and tau fields as initial conditions and the mean posterior model parameters from the pooled probabilistic model with shared A β -tau coupling. **4.11e** Shows the order of A β -tau colocalisation using a threshold of 90% A β and 10% tau, with 1 being the initial region of colocalisation and 35 being the final region, using mean posterior model parameters from the pooled probabilistic model. **4.11f** Probability of a region being the initial site of A β -tau colocalisation determined by iterating over samples of posterior model parameters.

Order	Region	Time / years
1	inferior temporal	37.1
2	lateral occipital	39.6
3	middle temporal	40.2
4	inferior parietal	40.9
5	bankssts	42.3
6	superior parietal	42.5
7	fusiform	42.6
8	superior temporal	43.5
9	supramarginal	44.5
10	precuneus	45.3
11	isthmus cingulate	46.9
12	rostral middle frontal	47.0
13	post central	48.6
14	insula	48.8
15	caudal middle frontal	48.9
16	lateral orbitofrontal	49.0
17	pericalcarine	49.2
18	cuneus	49.3
19	parstriangularis	49.4
20	posterior cingulate	49.5
21	parsorbitalis	50.0
22	parsopercularis	50.3
23	superior frontal	50.3
24	precentral	50.5
25	paracentral	52.0
26	transverse temporal	52.2
27	caudal anterior cingulate	52.9
28	lingual	53.4
29	medial orbitofrontal	53.6
30	rostral anterior cingulate	53.9
31	frontal pole	56.9
32	temporal pole	58.3
33	parahippocampal	60.5
34	entorhinal	67.6
35	Right-Amygdala	96.7

Table 4.8: **Colocalisation order.** Order and time of $A\beta$ -tau colocalisation and given initial conditions from the A^+T^- ADNI cohort and mean posterior model parameters.

We further characterise the seeding probabilities using different thresholds, shown in Figure 4.12, for $A\beta$ -tau thresholds, $(A\beta^*, \tau^*) \in \{(1 - p, p) \mid p \in \{0.05, 0.1, 0.15, 0.2\}\}$. The results show that for high tau thresholds and low amyloid thresholds, $A\beta$ -tau colocalisation occurs more frequently in the medial temporal

lobe, and conversely for low tau thresholds and high amyloid thresholds, colocalisation occurs most frequently in cortical (parietal and frontal regions) and never in the medial temporal lobe. This is to be expected due to the spatial pattern of initial amyloid and tau deposition. Since tau is initially present in the entorhinal cortex, the concentration in the entorhinal cortex will initially rise quickly, conversely, since $A\beta$ is present in the neocortex first these regions will reach the high threshold burden first while satisfying the low tau burden to be suprathreshold due to transport. Consistent across these are regions in the lateral temporal lobe, namely the the fusiform and inferior temporal lobe and support results from Figures 4.11e and 4.11f that these regions are the first regions in which $A\beta$ and tau colocalise to accelerate tau expansion into the cortex.




Figure 4.12: **Variable arrival time thresholds for $A\beta$ -tau colocalisation.** Probability of regions being the initial site of colocalisation with variations in detection thresholds, given on the left ($A\beta^*/\tau^*$).

Seed	Probability	Lobe
fusiform	0.295	lateral temporal
inferior temporal	0.287	lateral temporal
inferior parietal	0.196	parietal
rostral middle frontal	0.085	frontal
entorhinal	0.051	medial temporal
parahippocampal	0.035	medial temporal
middle temporal	0.026	lateral temporal
precuneus	0.016	parietal
bankssts	0.003	lateral temporal
amygdala	0.003	medial temporal

Table 4.9: **Colocalisation probability.** Probability of seed region given 90% $A\beta$ threshold and 10% tau threshold, given samples from the posterior distribution.

4.7 Amyloid Induced Symmetry Breaking Promotes EC Vulnerability

So far we have discussed how $A\beta$ and tau interact after there is seeding in the medial temporal lobe. In Section 4.6 we establish a framework for determining the epicentre of secondary tauopathy in AD as the point of $A\beta$ and tau colocalisation. In this section, we propose a novel theory for how $A\beta$ might influence the initial seeding of tau in the EC, causing primary tauopathy.

Although the staging of tau throughout AD is well characterised, it remains unknown how tau seeds in the entorhinal cortex. Amyloid is known to primarily invade the neocortex before reaching deep brain structures such as the entorhinal cortex and medial temporal lobe and is particularly abundant in the parietal and frontal cortex, with substantially less involvement in the medial temporal lobe. This is consistent with our findings in Figure 4.3e, where we show through cross-sectional modelling that the asymptotic distribution of $A\beta$ is highest in the medial and lateral parietal lobes and lowest in the medial temporal lobe, the epicentre of tau seeding. In this section, we investigate how the early spatial distribution of $A\beta$ might facilitate tau seeding in the medial temporal lobe.

We investigate a cohort of A^+T^- subjects from ADNI ($N = 37$; full demographics in Table 4.3). Figure 4.13a shows the average distribution of $A\beta$ in the cortex with regional baseline $A\beta$ PET values subtracted. As expected, there is a high distribution of $A\beta$ in the medial and lateral parietal lobe and lateral frontal cortex. Similarly to the asymptotic state of $A\beta$ in Figure 4.3e, subjects at this stage show a relatively

low binding in the medial temporal lobe (entorhinal cortex and amygdala). There are numerous ways that $A\beta$ deposition might affect tau spreading, the most likely of which is through $A\beta$ induced effects on neuronal hyperactivation [169, 170, 171], in turn facilitating tau spreading [38, 39], or through the direct effects of $A\beta$ on tau spreading, which have been demonstrating in animal models [172, 40, 158, 165]. We model these effects by weighting the undirected graph Laplacian by the cortical $A\beta$ field.

$$\hat{\mathbf{L}} = \mathbf{L}\mathbf{U}^\varphi, \quad (4.29)$$

where $\mathbf{U} = \text{diag}(\mathbf{u})$ is a diagonal matrix containing the mean centred average regional $A\beta$ load present in A^+T^- subjects, and $\varphi \in \mathbb{R}$ is a weighting coefficient controlling the effect of $A\beta$ on the structural network. Weighting the graph Laplacian in this way breaks the symmetry present in the undirected graph. Therefore, while mass conservation is preserved, Fick's law does not apply and transport can occur in the absence of a concentration gradient. To examine how symmetry breaking in this way affects tau spreading we first examine the spectrum of the weighted graph Laplacian, $\hat{\mathbf{L}}$. A cortical rendering of the first eigenvector (corresponding to the zero eigenvalue) is shown in Figure 4.13b. The first eigenvector is proportional to the stationary distribution of tau under linear transport dynamics, $\hat{v} = -\hat{\mathbf{L}}v$. It can be seen in Figure 4.13b that the flow induced by the directed graph Laplacian is driven to the medial temporal lobe, and particularly to the entorhinal cortex. This shows that, under linear transport dynamics on a network, the distribution of $A\beta$ creates a natural vulnerability for the EC to accumulate tau.

Next, we investigate how the linear transport dynamics may affect regional non-linear production processes. We revisit the heterodimer model using the weighted Laplacian:

$$\frac{dp_i}{dt} = -\rho \sum_j^R \hat{L}_{ij} p_j + k_0 - k_1 p_i - k_{12} p_i \tilde{p}_i, \quad p_i(0) = p_{i,0}, \quad i = 1, \dots, R \quad (4.30a)$$

$$\frac{d\tilde{p}_i}{dt} = -\rho \sum_j^R \hat{L}_{ij} \tilde{p}_j - \tilde{k}_1 \tilde{p}_i + k_{12} p_i \tilde{p}_i, \quad \tilde{p}_i(0) = \tilde{p}_{i,0}, \quad i = 1, \dots, R. \quad (4.30b)$$

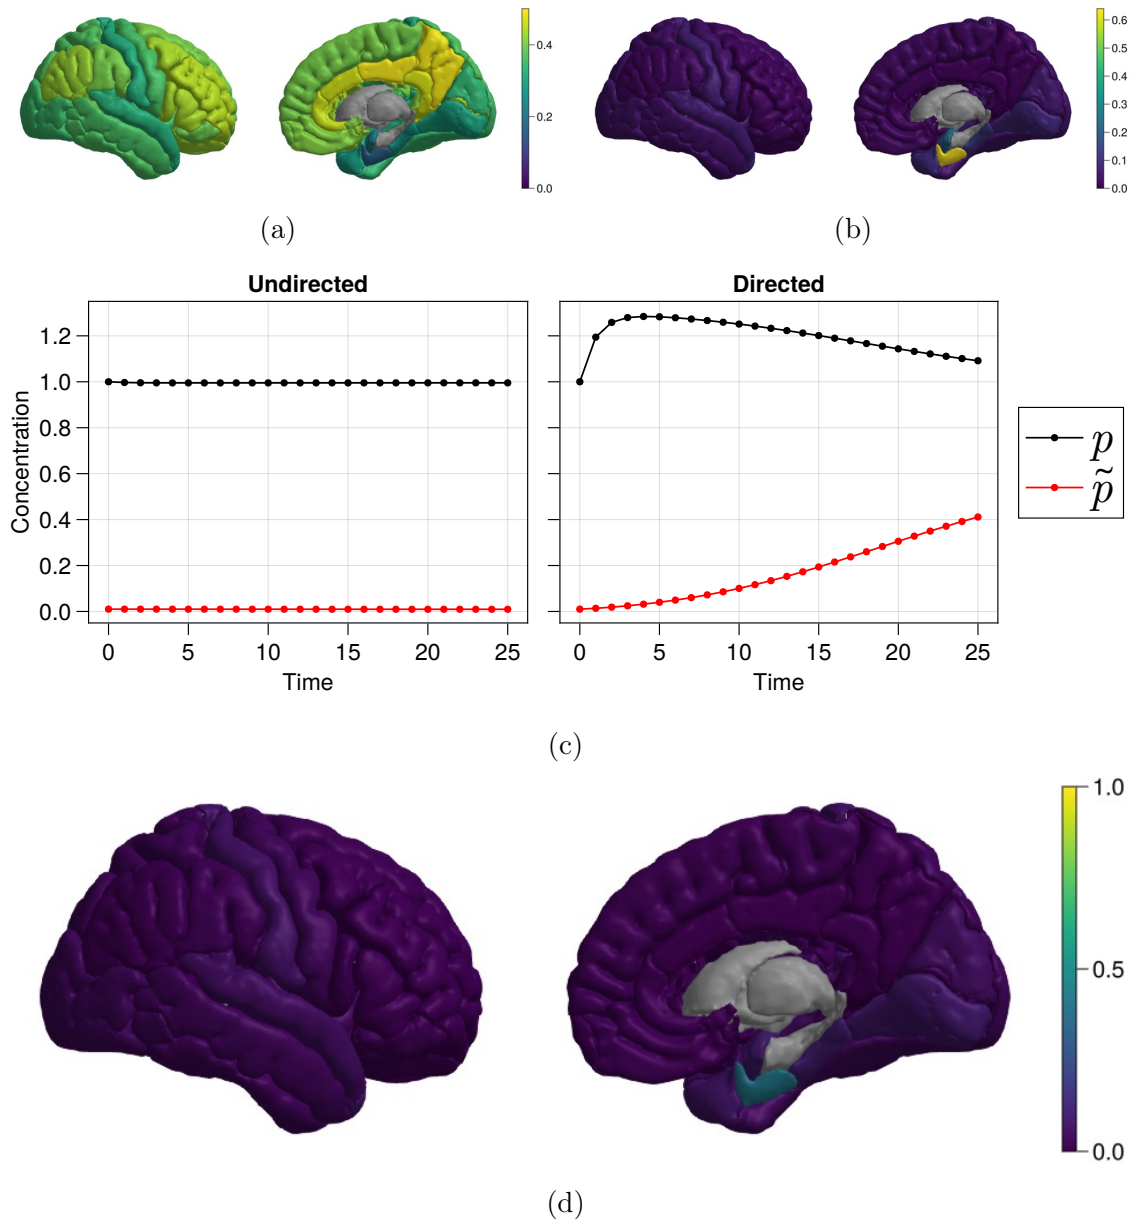


Figure 4.13: **Posterior distributions from the ATN model calibrated to A^+T^+ subjects from ADNI.** **4.13a.** Regional $A\beta$ SUVR for the right hemisphere, corrected to baseline $A\beta$ PET values. **4.13b.** First eigenvector of the graph Laplacian weighted by mean centred average baseline corrected $A\beta$ from A^+T^- ADNI subjects with $\varphi = 3.5$. **4.13c.** Simulations from the heterodimer model with using the $A\beta$ weighted network with $\mathbf{p}(0) = \mathbf{1}$, $\tilde{\mathbf{p}}(0) = \mathbf{0.01}$, $\rho = 1.0/yr$, $k_0 = 1.0/yr$, $k_1 = 1.0/yr$, $k_{12} = 0.5/yr$, $\tilde{k} = 0.5/yr$. **4.13d.** Cortical representation of simulation in 4.13c at $t = 25$ years.

Recall from Section 2.2.2 that k_0 , the natural source term for healthy protein, p_i , acts a bifurcation parameters, shifting the dynamics from healthy steady state to disease steady state. Under directed flow, the transport term is no longer uniform,

but instead acts as a source term for regions that acts as sinks for directed flow, in this case, the entorhinal cortex. To test whether this directed flow is sufficient to induce a regional bifurcation we simulate the heterodimer model using an undirected and directed graph (weighted by $A\beta$) using a uniform distribution of \mathbf{p} and $\tilde{\mathbf{p}}$, where $p_i \gg \tilde{p}_i$ and parameters in a stable regime $\rho = 1.0/yr$, $k_0 = 1.0/yr$, $k_1 = 1.0/yr$, $k_2 = 0.5/yr$, $\tilde{k} = 0.5/yr$. The results of these simulations in the entorhinal cortex are shown in Figure 4.13c for the undirected graph (left) and directed graph (right). As expected, in the undirected case, the concentration of both p_i and \tilde{p}_i remain stable. However, in the directed case, the system becomes unstable and \tilde{p}_i begins to grow. We show the end concentration field in Figure 4.13d, which shows that this instability is confined to the entorhinal cortex, since this is the region which experiences the greatest in-flow of p_i . This analysis shows that the asymmetry induced by regional $A\beta$ distribution confers specific vulnerability to the entorhinal cortex that, when coupled with regional production and clearance dynamics, allows tau seeds to amplify in the entorhinal cortex.

4.8 Discussion

Alzheimer’s disease is a complex disease with a multifaceted pathology centred around $A\beta$ and tau. In this chapter we extend the results of Chapter 2 and Chapter 3 to develop a novel, mechanistic model of the interactions between $A\beta$, tau and neurodegeneration, describing the ATN framework [10]. We show that the model predicts several aspects of the coupled ATN pathology, including tau heterogeneity, Braak staging of tau, $A\beta$ induced acceleration in tau and atrophy, and ATN biomarker staging. Furthermore, we apply the model to longitudinal multimodal ATN data from a A^+T^+ cohort in ADNI and show that the model can accurately explain biomarker development using a pooled $A\beta$ -tau coupling parameter. Using this calibrated model, we show that by modelling both $A\beta$ and tau we are able to identify colocalisation sites of $A\beta$ and tau. Finally, we show that gradients in $A\beta$ deposition create a natural vulnerability for the entorhinal cortex to accumulate tau, possibly providing an explanation for tau seeding in the entorhinal cortex. Overall, we have provided a complete account of the ATN biomarker framework that facilitates prediction and understanding of key factors in AD.

In Chapter 3 we demonstrated the effectiveness of a regionally heterogeneous model of tau production and spread and highlight the importance of regionally specific dynamics in accurately simulating the progression of tau pathology. A key shortfall of

the model presented in Chapter 3 was that carrying capacities of tau were fixed from population data and were therefore qualitative in nature. $A\beta$ is the most likely candidate for providing a mechanism for heterogeneous tau dynamics since $A\beta$ itself has distinct spatial topography, and is correlated with neocortical tau pathology [6], and is known to enhance tau pathology [39, 166, 158, 40]. In this chapter, we test whether $A\beta$ explains tau heterogeneity by developing a model that explains $A\beta$ induced tau acceleration as an elevation in regional tau carrying capacities. In Figure 4.4, we show that the staging and topology of tau correlates with Braak staging and heterogeneous tau deposition. Furthermore, in Figure 4.10 we show that the coupled ATN model provides a good explanation for tau progression. However, on average, the progression of tau is underpredicted, meaning $A\beta$ does not entirely explain regional tau dynamics and other factors might influence regional tau. There are several other candidates that might act together with $A\beta$ to influence tau, including genetic risk factors [137], such as regional apolipoprotein-E [135, 151, 152] and microtubule association protein tau [150] expression, and gradients in functional activation that has been shown to influence tau progression [148]. Future work should aim to include multiple risk factors influencing tau pathology, in addition to regional $A\beta$ load to explain.

There have been several studies showing that $A\beta$ is necessary for the acceleration of tau and the expansion of pathology into the neocortex [166, 167, 37, 164]. Due to the spatial segregation of $A\beta$ and tau during AD [168], $A\beta$ does not have an immediate effect on tau pathology. Here we propose that tau acceleration occurs when $A\beta$ and tau colocalise. By modelling $A\beta$ -tau simultaneously in the ATN model, we are able to infer the time and location at which $A\beta$ and tau colocalise. First, we show in Section 4.3.2 that the ATN model predicts colocalisation first in the inferior temporal lobe which temporally coincides with tau acceleration into the neocortex. In addition to the inferior temporal lobe, the fusiform and precuneus regions show early acceleration following $A\beta$ -tau colocalisation. These predictions were validated in Section 4.6 in which we use subjects who are early in the AD spectrum (A^+T^-) and samples from the posterior distribution obtained in in Section 4.5 to simulate trajectories using biophysical initial conditions and model parameters and determine $A\beta$ -tau colocalisation sites. In so doing, we show that the inferior temporal lobe, fusiform, middle temporal, precuneus, inferior parietal and rostral middle frontal and medial temporal lobe are all candidate regions for $A\beta$ seeding, with the most likely locations being the fusiform and inferior temporal. Interestingly, these regions coincide with those that are particularly affected by atrophy [161], suggesting that early $A\beta$ -tau colocalisation correlates with accelerated atrophy. By identifying these

regions and being able to simulate trajectories for when $A\beta$ and tau interact, we are able to identify a critical window before which AD pathology accelerates. Identifying the location and time of this interaction could be crucial for optimal application of disease modifying interventions in AD. Current therapies, such as Lecanemab and Donanemab, focus on the reduction of $A\beta$ [48, 50]. A possible explanation for the limited effects of $A\beta$ therapies on AD progression is that drugs are not administered soon enough in the AD timeline. This is consistent with the ATN model, which predicts that $A\beta$ reduction is necessary before there is tau acceleration, since after this time, tau can accumulate and cause atrophy. Stopping $A\beta$ after tau acceleration will result in slowing of disease progression (as observed in clinical trials), whereas reducing $A\beta$ before tau acceleration will halt the progression of AD into secondary tauopathy. At present the ATN model is not able to capture colocalisation at the individual level, with the addition of further longitudinal ATN biomarker data, it could be possible to predict individual $A\beta$ and tau trajectories to determine when individuals are likely to experience $A\beta$ induced tau acceleration, and therefore identify a critical window for therapeutic intervention.

In addition to showing how $A\beta$ and tau interact after there is seeding of $A\beta$ and tau, in Section 4.7 we propose a novel theory for how $A\beta$ facilitates tau seeding in the medial temporal lobe by inducing specific vulnerability in the entorhinal cortex, driven by spatial patterns of $A\beta$ deposition and the structural brain network. We propose that $A\beta$ induces the accumulation tau seeds in the entorhinal cortex given linear network dynamics weighted by $A\beta$ topology, triggering an instability between the non-linear processes of clearance and production that facilitates specific amplification of toxic tau in the entorhinal cortex. We use the heterodimer model to simulate non-linear dynamics, and although this model was initially proposed to simulate healthy and toxic tau fibrils [80], it can equally well be applied to endogenous tau seeds, such as oligomeric tau, and toxic tau fibrils [144]. The mechanism through which $A\beta$ drives tau spread could either be through a direct effect of $A\beta$ on tau [172, 158, 39], through the effect of $A\beta$ hyperactivation on activity dependent tau transport [42, 148, 169, 170], or more indirectly by affecting local clearance and production parameters. In support of the joint role of $A\beta$ and activity effects, Giorgio et al. (2023) have recently shown an association between $A\beta$ induced hyperexcitability and early tau accumulation [173], however, the authors were unable to identify particular mechanisms through which this accumulation occurs. We propose that future work address this through a combination of multiscale modelling of $A\beta$ dependent effects on activity and tau spreading.

Although the work here advances on the work in Chapter 3, we face similar limitations. First, we still rely on fixing important parameters, such as baseline and carrying capacities of $A\beta$ and the PART field for tau. The former are likely to be robustly identified from cross-sectional analysis, given the abundance of $A\beta$ PET data, however, by fixing their value we ignore potential inter-subject variability. In estimating the PART field, we are limited by the number subjects in whom PART can be detected in the absence of $A\beta$, namely the $A^{-}T^{+}$ ADNI cohort. This limitation might be addressed by increasing the amount of data available or by validation against different datasets with more sensitive tau tracers [143]. Additionally, in fixing these values we are unable to explain how heterogeneities in $A\beta$ deposition and early tau burden occur, and if or how they change over time. The importance of $A\beta$ heterogeneity is marked by recent studies demonstrating the existence of $A\beta$ spatiotemporal subtypes [174, 175] that overlap spatial subtypes of tau and cortical atrophy [176, 177, 74]. Future work should examine whether $A\beta$ deposition is predictive of tau deposition across spatial types and correlations between parameters exist that can be used to further characterise subtypes. Such valuable information could be included to provide more constrained prior information on parameters based on individual $A\beta$ deposition. As in Chapter 3, the model reduction we perform in Section 4.2 comes at the cost of losing mechanistic insight. To address more specific questions such as those posed in Section 4.7, we require more detailed modelling of $A\beta$ -tau interactions, cellular transport and interactions between oligomers of different sizes. Such modelling could be applied to animal models and future work might benefit from addressing not only macroscale modelling applied to multimodal human neuroimaging data but also microscale modelling of animal and in-vivo models. Another major limitation of the current study is the limited sample size. Despite this, we note that posterior densities for individual parameters are generally identifiable, showing little variation. Therefore, while individual parameters are unlikely to change dramatically provided more data, population-level parameters may change. Therefore, in the absence of more data, future work should validate the results presented here on other datasets, such as BF2, which may contain greater longitudinal $A\beta$ and tau PET for the $A^{+}T^{+}$ group. In addition to these limitations, we note the heavy data burden required to calibrate and apply not only the ATN model but models in Chapter 2 and Chapter 3 that will greatly limit the utility these models have in clinical practice, where longitudinal PET data is rarely available. However, such data is routinely collected for clinical trials and pharmaceutical research and the ATN model may provide useful applications for patient selection and pharmacodynamic modelling.

In this chapter we concluded our quest for advancing macroscale disease progression modelling of AD by combining methods in Chapter 2 and Chapter 3 to develop and apply a mechanistic model of the ATN biomarker framework. In so doing, we have built a tool that can be used for the prediction of critical AD biomarkers used in clinical research and practice and the foundation on which to build multimodal models of neuroimaging data in AD. In the next chapter, we will conclude this thesis and offer remarks on how we believe modelling can progress to tackle the remaining obstacles.

Chapter 5

Concluding Thoughts

The overarching aim of this thesis has been to develop a mathematical and inferential modelling framework for neurodegeneration that enables the application of disease progression modelling in clinical practice. With that ultimate goal in mind, we have made progress on several issues pertaining to disease progression modelling of AD. In this chapter, we reflect on the progress we have made toward our initial aims and discuss avenues in which this research can be extended to facilitate the application of modelling in clinical settings.

We have had several research aims throughout this thesis. First, we sought to connect existing network based models of pathological protein spread with neuroimaging data, using a Bayesian framework that characterises uncertainty in the modelling process. Second, we develop ways to assess different models against neuroimaging data to ensure the correct mechanisms are being described. Third, we built new mechanistic models of multi-modal neuroimaging data to describe important AD biomarkers.

In Chapter 2, we introduced simple network models of toxic proteins and developed a probabilistic modelling framework to assess differences in disease mechanisms in cohorts across the AD timeline. We showed that a simple model is able to predict individual patient trajectories and discern differences in disease dynamics between healthy and unhealthy subject cohorts. In Chapter 3 we focussed our attention on modelling tau, building a generative model of tau PET neuroimaging data that accounts for variations in regional PET tracer uptake and vulnerability to tau invasion. We compared the model to a family of related models and in so doing highlighted the importance of regional heterogeneity. Furthermore, we demonstrated the ability for mechanistic models to predict accurately regional out-of-sample data, potentially providing a valuable tool for patient specific, regional specific prediction. In Chapter 4 we extended the generative framework from Chapter 3 to include interactions between $A\beta$, tau, and neurodegeneration, thereby allowing predictions of biomarkers in

the ATN framework, which is the biomarker staging framework used by the National Institute on Ageing for qualitatively characterising stages of AD [10]. These results are the most clinically relevant, demonstrating the importance of $A\beta$ -tau interaction, the temporal importance of $A\beta$ -tau colocalisation and the prediction of multiple AD biomarkers, including neurodegeneration. Despite this progress toward modelling the progression of Alzheimer’s disease there remain challenges yet unmet which prevent the widespread utility of models in clinical research.

The main obstacle preventing the widespread utilisation of mechanistic models of AD progression is the heavy data burden required for parameter calibration. The models presented in Chapters 1 and 2 require at least three longitudinal PET scans, and the model in Chapter 4 requires at longitudinal scans for $A\beta$ PET, tau PET and structural MRI. However, PET scans are difficult to acquire due to their prohibitive cost, availability PET scanners and relevant radiotracers. Furthermore, PET scans are not routinely collected in clinical practice when patients present with typical AD progression, typically measured through cognitive tests [178, 179, 180]. As such, longitudinal neuroimaging data are rarely available for individual patients in a clinical setting. For these reasons, the mechanistic models implemented in this work are important research tools, but are likely to be of limited use to clinicians in their current form.

As a quantitative measure of AD pathology, fluid biomarkers for $A\beta$ and tau pathology, such as those derived from cerebrospinal fluid (CSF), are increasingly being used to in the diagnosis of Alzheimer’s disease [181, 182]. Furthermore, fluid biomarkers correlate with rates of $A\beta$ and tau accumulation, as well as the spatial deposition of tau PET [183, 138, 142]. These results suggest that modelling of fluid biomarkers might be a fruitful next step for predicting individual disease progression. However, fluid biomarkers provide only temporal point values for $A\beta$ or tau pathology, where as PET biomarkers provide additional information about the spatial deposition of pathology. Throughout this work we have emphasised the importance of modelling the spatial progression of tau and in Chapter 4, we highlighted the critical importance of the spatial deposition of $A\beta$ and tau, showing that the spatial colocalisation of $A\beta$ -tau drives the progression of AD. In addition, spatial information is necessary to determine individual subtype designations [74], which are not accessible through fluid biomarkers, and which may prove important as an avenue for personalised treatment strategies. Thus, if possible, we wish to preserve a spatial component to AD progression modelling. This might be achievable given certain assumptions. In Chapter 3 we showed that the addition of network spread

provides limited predictive information over regionally heterogeneous production for mid-late stage AD subjects (A^+T^+ subjects), indicating that a model relying on the heterogeneous tau production may suffice for A^+T^+ subjects. Additionally, recent modelling work by Putra et al. (2022) has shown that, given the small transport parameter, reliable asymptotic approximations can be made to derive network spread models of tau that do not rely on a transport parameter, only a production parameter [184]. These asymptotic models may be suitable for modelling subjects earlier in the disease process (A^+T^-) for whom tau transport is more important, for $A\beta$ induced tau seeding, cortical expansion of tau, and $A\beta$ -tau colocalisation. We propose that future modelling research should seek to test the application of combined modelling of longitudinal fluid and PET biomarkers. The former may provide a less data intensive way for calibrating production rates of tau as synthesised in transport-free or asymptotic models of tau progression, while the latter can provide data from which to initialise predictive spatial models. Together, these may provide the opportunity for the spatial prediction of ATN biomarkers using a limited number of PET scans with longitudinal fluid biomarkers.

There is greater potential for immediate utilisation of mechanistic modelling in pharmaceutical research, where longitudinal PET data are more readily available. In recent years, several anti- $A\beta$ therapies have been shown to be effective in reducing $A\beta$ burden, typically measured through longitudinal PET scans [48, 49, 50]. Despite reduction in $A\beta$ load, there have been only modest slowing in cognitive decline. There are several reasons for the limited efficacy of anti- $A\beta$ therapies, such as poor blood brain penetration, ineffective dosing regimes, or suboptimal administration times. Particular attention has been paid to the need for early intervention, even in prodromal or preclinical stages of AD to halt the progression of AD before there is widespread proteopathy across the cortex [185, 186]. In Chapter 4, we provide evidence to support early intervention, with our results suggesting that preventing substantial $A\beta$ -tau colocalisation may prevent widespread cortical tau progression. Therefore, modelling of the ATN pathway may help facilitate trials in prodromal or preclinical subjects by helping to identify suitable subjects, namely those with confined but growing tau progression or limited $A\beta$ -tau colocalisation, in so doing isolating those individuals who are likely to most benefit from anti- $A\beta$ therapies. Furthermore, when paired with pharmacokinetic-pharmacodynamic models of anti- $A\beta$ therapies, modelling of ATN biomarkers enables the use of synthetic clinical trials for optimising dosing strategies, patient selection criteria, administration periods, and analysis of effect size in responders. Through this combination of optimising trial

design and identifying most efficacious windows of administration, the modelling opportunities afforded by the ATN framework may help significantly reduce the price of clinical trials, which largely stems from the recruitment and scanning of participants [185]. At present, longitudinal clinical trial data is not publicly available and therefore the pharmaceutical applications of the models presented in this thesis to real-world pharmaceutical trial data are limited. However, the ability for spatiotemporal models of AD to accurately fit longitudinal biomarker data, as in Chapter 4 provides the opportunity for in-silico exploration of optimal pharmaceutical intervention strategies. Pertinent questions that may be addressed with in-silico studies include optimal intervention times as a function of individual amyloid and tau burden, the importance of intervention before $A\beta$ -tau colocalisation, dosing strategies for $A\beta$ reduction, including the importance of maintenance dosing, the existence of subtype specific dosing strategies, and many others. We anticipate this will be a fruitful avenue for future research.

In addition to novel applications of network based disease progression models, future work should also address limitations present in the modelling and inference pipeline described throughout this thesis. In Section 1.4, we highlighted some of the pressing issues in network-based disease modelling, namely uncertainty in model correctness, parameters, and network topology. Throughout this thesis, we have addressed issues of model correctness and parameter identifiability, however, we have not fully addressed uncertainty about network topology. The suitability of the network should be not underestimated since it defines the spatial domain of the model, including the discrete regions used, which may affect regionally heterogeneous dynamics, and how transport between them occurs. Problems around network topology are partially addressed by the use of null models [87], which can be used to test the effect of particular network features, such as topology, degree, centrality. Furthermore, in Putra et al. (2021), we use multi-scale network parcellations to test how discretisation size and the imbued network topology affects tau staging, increasing the chance of observing Braak staging under FKPP dynamics [86]. We have not tested how models with larger networks compare to longitudinal PET data. There are two challenges with using larger networks. First, reaction-diffusion models require greater computational effort to solve and might prohibitively slow down inference over large patient cohorts. Second, while tractography has a resolution that can facilitate fine grained connectomes, PET has a much lower resolution and may not benefit from the use high resolution connectomes. Future work should test how different

connectomes and parcellations affect the use of modelling disease progression for predicting longitudinal neuroimaging data. Furthermore, throughout this work we have used connectomes generated from a young healthy cohort in the Human Connectome Project. However, several studies now indicate that AD caused network failures causing changes in activity and structural decline [187, 188, 189, 190, 191]. This issue can be readily addressed by testing the application of reaction-diffusion network models using connectomes generated from older subjects, subjects with AD, or personalised connectomes from patients with AD.

Overall, we are enthusiastic about the progress made by the disease progression modelling community and are hopeful that further work will enable the application of modelling to reduce the cost of pharmaceutical development and provide precision clinical care to patients suffering with Alzheimer's disease.

Bibliography

- [1] John A Hardy and Gerald A Higgins. “Alzheimer’s disease: the amyloid cascade hypothesis”. In: *Science* 256.5054 (1992), pp. 184–186.
- [2] Mathias Jucker and Lary C Walker. “Propagation and spread of pathogenic protein assemblies in neurodegenerative diseases”. In: *Nature neuroscience* 21.10 (2018), p. 1341.
- [3] Rik Ossenkoppele et al. “Tau PET patterns mirror clinical and neuroanatomical variability in Alzheimer’s disease”. In: *Brain* 139.5 (Mar. 2016), pp. 1551–1567. ISSN: 0006-8950. DOI: [10.1093/brain/aww027](https://doi.org/10.1093/brain/aww027). eprint: <https://academic.oup.com/brain/article-pdf/139/5/1551/17348814/aww027.pdf>. URL: <https://doi.org/10.1093/brain/aww027>.
- [4] Michael Schöll et al. “PET Imaging of Tau Deposition in the Aging Human Brain”. In: *Neuron* 89.5 (2016), pp. 971–982. ISSN: 0896-6273. DOI: <https://doi.org/10.1016/j.neuron.2016.01.028>. URL: <http://www.sciencedirect.com/science/article/pii/S0896627316000532>.
- [5] Val J Lowe et al. “An autoradiographic evaluation of AV-1451 Tau PET in dementia”. In: *Acta Neuropathologica Communications* 4.1 (2016), p. 58.
- [6] Hanna Cho et al. “In vivo cortical spreading pattern of tau and amyloid in the Alzheimer disease spectrum”. In: *Annals of neurology* 80.2 (2016), pp. 247–258.
- [7] Val J Lowe et al. “Widespread brain tau and its association with ageing, Braak stage and Alzheimer’s dementia”. In: *Brain* 141.1 (Dec. 2017), pp. 271–287. ISSN: 0006-8950. DOI: [10.1093/brain/awx320](https://doi.org/10.1093/brain/awx320). eprint: <https://academic.oup.com/brain/article-pdf/141/1/271/24175722/awx320.pdf>. URL: <https://doi.org/10.1093/brain/awx320>.
- [8] Dennis J Selkoe. “The molecular pathology of Alzheimer’s disease”. In: *Neuron* 6.4 (1991), pp. 487–498.
- [9] Lary C Walker and Mathias Jucker. “Neurodegenerative diseases: expanding the prion concept”. In: *Annual review of neuroscience* 38 (2015), pp. 87–103.
- [10] Clifford R Jack et al. “NIA-AA Research Framework: Toward a biological definition of Alzheimer’s disease”. In: *Alzheimer’s & Dementia* 14.4 (2018), pp. 535–562.

- [11] George G. Glenner and C W Wong. “Alzheimer’s disease: initial report of the purification and characterization of a novel cerebrovascular amyloid protein.” In: *Biochemical and biophysical research communications* 120 3 (1984), pp. 885–90. URL: <https://api.semanticscholar.org/CorpusID:39050164>.
- [12] Dennis J Selkoe. “Amyloid β -protein and the genetics of Alzheimer’s disease”. In: *Journal of Biological Chemistry* 271.31 (1996), pp. 18295–18298.
- [13] Warren J. Strittmatter et al. “Apolipoprotein E: high-avidity binding to beta-amyloid and increased frequency of type 4 allele in late-onset familial Alzheimer disease.” In: *Proceedings of the National Academy of Sciences of the United States of America* 90 (1993), pp. 1977–1981. URL: <https://api.semanticscholar.org/CorpusID:20937018>.
- [14] Elizabeth H Corder et al. “Gene dose of apolipoprotein E type 4 allele and the risk of Alzheimer’s disease in late onset families”. In: *Science* 261.5123 (1993), pp. 921–923.
- [15] Holly M Brothers, Maya L Gosztyla, and Stephen R Robinson. “The physiological roles of amyloid- β peptide hint at new ways to treat Alzheimer’s disease”. In: *Frontiers in aging neuroscience* 10 (2018), p. 118.
- [16] Guo-fang Chen et al. “Amyloid beta: structure, biology and structure-based therapeutic development”. In: *Acta Pharmacologica Sinica* 38.9 (2017), pp. 1205–1235.
- [17] Yun-wu Zhang et al. “APP processing in Alzheimer’s disease”. In: *Molecular Brain* 4 (2011), pp. 3–3. URL: <https://api.semanticscholar.org/CorpusID:8469387>.
- [18] David R. Borchelt et al. “Familial Alzheimer’s Disease–Linked Presenilin 1 Variants Elevate A β 1–42/1–40 Ratio In Vitro and In Vivo”. In: *Neuron* 17 (1996), pp. 1005–1013. URL: <https://api.semanticscholar.org/CorpusID:18315650>.
- [19] Donalyn L Scheuner et al. “Secreted amyloid beta-protein similar to that in the senile plaques of Alzheimer’s disease is increased in vivo by the presenilin 1 and 2 and APP mutations linked to familial Alzheimer’s disease.” In: *Nature medicine* 2 8 (1996), pp. 864–70. URL: <https://api.semanticscholar.org/CorpusID:39017461>.
- [20] Sneham Tiwari et al. “Alzheimer’s disease: pathogenesis, diagnostics, and therapeutics”. In: *International Journal of Nanomedicine* 14 (2019), pp. 5541–5554. URL: <https://api.semanticscholar.org/CorpusID:199545358>.
- [21] Lary C. Walker, Juliane Schelle, and Mathias Jucker. “The Prion-Like Properties of Amyloid- β Assemblies: Implications for Alzheimer’s Disease.” In: *Cold Spring Harbor perspectives in medicine* 6 7 (2016). URL: <https://api.semanticscholar.org/CorpusID:21687699>.

- [22] Tomas T Olsson, Oxana Klementieva, and Gunnar K Gouras. “Prion-like seeding and nucleation of intracellular amyloid- β ”. In: *Neurobiology of disease* 113 (2018), pp. 1–10.
- [23] John Hardy and Dennis J Selkoe. “The amyloid hypothesis of Alzheimer’s disease: progress and problems on the road to therapeutics”. In: *science* 297.5580 (2002), pp. 353–356.
- [24] Theresa M Harrison et al. “Longitudinal tau accumulation and atrophy in aging and alzheimer disease”. In: *Annals of Neurology* 85.2 (2019), pp. 229–240.
- [25] Chenjie Xia et al. “Association of in vivo [18F] AV-1451 tau PET imaging results with cortical atrophy and symptoms in typical and atypical Alzheimer disease”. In: *JAMA neurology* 74.4 (2017), pp. 427–436.
- [26] L. Pini, M. Pievani, and G. et. al Frisoni. “Brain atrophy in Alzheimer’s Disease and aging”. In: *Ageing Res. Rev* 30 (2016), pp. 25–48.
- [27] Fong Ping Chong et al. “Tau proteins and tauopathies in Alzheimer’s disease”. In: *Cellular and molecular neurobiology* 38 (2018), pp. 965–980.
- [28] Eva-Maria Mandelkow and Eckhard Mandelkow. “Tau in Alzheimer’s disease”. In: *Trends in cell biology* 8.11 (1998), pp. 425–427.
- [29] Carlo Ballatore, Virginia M-Y Lee, and John Q Trojanowski. “Tau-mediated neurodegeneration in Alzheimer’s disease and related disorders”. In: *Nature reviews neuroscience* 8.9 (2007), pp. 663–672.
- [30] Inge Grundke-Iqbal et al. “Microtubule-associated protein tau. A component of Alzheimer paired helical filaments.” In: *The Journal of biological chemistry* 261 13 (1986), pp. 6084–9. URL: <https://api.semanticscholar.org/CorpusID:961777>.
- [31] Christian Bancher et al. “Accumulation of abnormally phosphorylated τ precedes the formation of neurofibrillary tangles in Alzheimer’s disease”. In: *Brain Research* 477 (1989), pp. 90–99. URL: <https://api.semanticscholar.org/CorpusID:23857381>.
- [32] Tara. L. Spires-Jones, Johannes Attems, and Dietmar Rudolf Thal. “Interactions of pathological proteins in neurodegenerative diseases”. In: *Acta Neuropathologica* 134 (2017), pp. 187–205. URL: <https://api.semanticscholar.org/CorpusID:21054901>.
- [33] Florence Clavaguera et al. “Brain homogenates from human tauopathies induce tau inclusions in mouse brain”. In: *Proceedings of the National Academy of Sciences* 110.23 (2013), pp. 9535–9540.
- [34] Cristian A. Lasagna-Reeves et al. “Alzheimer brain-derived tau oligomers propagate pathology from endogenous tau”. In: *Scientific Reports* 2 (2012). URL: <https://api.semanticscholar.org/CorpusID:18045415>.

- [35] Florence Clavaguera et al. “Transmission and spreading of tauopathy in transgenic mouse brain”. In: *Nature cell biology* 11 (2009), pp. 909–913. URL: <https://api.semanticscholar.org/CorpusID:7144501>.
- [36] Mathias Jucker and Lary C Walker. “Self-propagation of pathogenic protein aggregates in neurodegenerative diseases”. In: *Nature* 501.7465 (2013), pp. 45–51.
- [37] John F Crary et al. “Primary age-related tauopathy (PART): a common pathology associated with human aging”. In: *Acta neuropathologica* 128 (2014), pp. 755–766.
- [38] Amy M. Pooler and Bradley T. et. al. Hyman. “Amyloid accelerates tau propagation and toxicity in a model of early Alzheimer’s disease”. In: *Acta Neuropathologica Communications* 3.1 (2015), p. 14. ISSN: 2051-5960. DOI: [10.1186/s40478-015-0199-x](https://doi.org/10.1186/s40478-015-0199-x). URL: <https://doi.org/10.1186/s40478-015-0199-x>.
- [39] Rachel E Bennett et al. “Enhanced tau aggregation in the presence of amyloid β ”. In: *The American journal of pathology* 187.7 (2017), pp. 1601–1612.
- [40] Zhuohao He et al. “Amyloid- β plaques enhance Alzheimer’s brain tau-seeded pathologies by facilitating neuritic plaque tau aggregation”. In: *Nature medicine* 24.1 (2018), p. 29.
- [41] Heiko Braak and Eva Braak. “Neuropathological staging of Alzheimer-related changes”. In: *Acta neuropathologica* 82.4 (1991), pp. 239–259.
- [42] Li Liu et al. “Trans-synaptic spread of tau pathology in vivo”. In: *PloS one* 7.2 (2012), e31302.
- [43] Alix De Calignon et al. “Propagation of tau pathology in a model of early Alzheimer’s disease”. In: *Neuron* 73.4 (2012), pp. 685–697.
- [44] Dennis J Selkoe and John Hardy. “The amyloid hypothesis of Alzheimer’s disease at 25 years”. In: *EMBO molecular medicine* 8.6 (2016), pp. 595–608.
- [45] Giovanni Battista Frisoni et al. “The probabilistic model of Alzheimer disease: the amyloid hypothesis revised”. In: *Nature Reviews Neuroscience* 23 (2021), pp. 53–66. URL: <https://api.semanticscholar.org/CorpusID:244529357>.
- [46] Clifford R Jack Jr et al. “Tracking pathophysiological processes in Alzheimer’s disease: an updated hypothetical model of dynamic biomarkers”. In: *The Lancet Neurology* 12.2 (2013), pp. 207–216.
- [47] Sergey Shcherbinin et al. “Association of Amyloid Reduction After Donanemab Treatment With Tau Pathology and Clinical Outcomes”. In: *JAMA Neurology* 79 (2022), pp. 1015–1024. URL: <https://api.semanticscholar.org/CorpusID:252198780>.
- [48] Christopher H. van Dyck et al. “Lecanemab in Early Alzheimer’s Disease.” In: *The New England journal of medicine* (2022). URL: <https://api.semanticscholar.org/CorpusID:254094094>.

- [49] Samantha L. Budd Haeberlein et al. “Two Randomized Phase 3 Studies of Aducanumab in Early Alzheimer’s Disease”. In: *The Journal of Prevention of Alzheimer’s Disease* 9 (2022), pp. 197–210. URL: <https://api.semanticscholar.org/CorpusID:247488889>.
- [50] John R. Sims et al. “Donanemab in Early Symptomatic Alzheimer Disease: The TRAILBLAZER-ALZ 2 Randomized Clinical Trial.” In: *JAMA* (2023). URL: <https://api.semanticscholar.org/CorpusID:259946737>.
- [51] Prashanthi Vemuri and Clifford R. Jack. “Role of structural MRI in Alzheimer’s disease”. In: *Alzheimer’s Research & Therapy* 2 (2010), pp. 23–23. URL: <https://api.semanticscholar.org/CorpusID:16432354>.
- [52] Ross M Lawrence et al. “Standardizing human brain parcellations”. In: *Scientific data* 8.1 (2021), pp. 1–9.
- [53] Pantea Moghimi et al. “A Review on MR Based Human Brain Parcellation Methods”. In: *arXiv preprint arXiv:2107.03475* (2021).
- [54] Bruce Fischl et al. “Automatically parcellating the human cerebral cortex”. In: *Cerebral cortex* 14.1 (2004), pp. 11–22.
- [55] Arno Klein and Jason Tourville. “101 labeled brain images and a consistent human cortical labeling protocol”. In: *Frontiers in neuroscience* 6 (2012), p. 171.
- [56] Susan Landau and W Jagust. “Flortaucipir (AV-1451) processing methods”. In: *Alzheimer’s Disease Neuroimaging Initiative* (). URL: https://adni.bitbucket.io/reference/docs/UCBERKELEYAV1451/UCBERKELEY_AV1451_Methods_Dec2019.pdf (visited on 11/21/2022).
- [57] Carmen Echávarri et al. “Atrophy in the parahippocampal gyrus as an early biomarker of Alzheimer’s disease”. In: *Brain Structure and Function* 215.3-4 (2011), pp. 265–271.
- [58] Timothy EJ Behrens et al. “Characterization and propagation of uncertainty in diffusion-weighted MR imaging”. In: *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* 50.5 (2003), pp. 1077–1088.
- [59] Timothy EJ Behrens et al. “Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?” In: *Neuroimage* 34.1 (2007), pp. 144–155.
- [60] Adam J Schwarz et al. “Regional profiles of the candidate tau PET ligand 18 F-AV-1451 recapitulate key features of Braak histopathological stages”. In: *Brain* 139.5 (2016), pp. 1539–1550.
- [61] Chun-Fang Xia et al. “[18F] T807, a novel tau positron emission tomography imaging agent for Alzheimer’s disease”. In: *Alzheimer’s & Dementia* 9.6 (2013), pp. 666–676.
- [62] Marta Marquié et al. “Validating novel tau positron emission tomography tracer [F-18]-AV-1451 (T807) on postmortem brain tissue”. In: *Annals of neurology* 78.5 (2015), pp. 787–800.

- [63] Michael Honer et al. “Preclinical evaluation of 18F-RO6958948, 11C-RO6931643, and 11C-RO6924963 as novel PET radiotracers for imaging tau aggregates in Alzheimer disease”. In: *Journal of Nuclear Medicine* 59.4 (2018), pp. 675–681.
- [64] Rik Ossenkoppele et al. “Accuracy of tau positron emission tomography as a prognostic marker in preclinical and prodromal Alzheimer disease: a head-to-head comparison against amyloid positron emission tomography and magnetic resonance imaging”. In: *JAMA neurology* 78.8 (2021), pp. 961–971.
- [65] Ruben Smith et al. “Head-to-head comparison of tau positron emission tomography tracers [18F] flortaucipir and [18F] RO948”. In: *European journal of nuclear medicine and molecular imaging* 47.2 (2020), pp. 342–354.
- [66] Val J Lowe et al. “An autoradiographic evaluation of AV-1451 Tau PET in dementia”. In: *Acta neuropathologica communications* 4.1 (2016), pp. 1–19.
- [67] Laetitia Lemoine et al. “Tau positron emission tomography imaging in tauopathies: the added hurdle of off-target binding”. In: *Alzheimer’s & Dementia: Diagnosis, Assessment & Disease Monitoring* 10 (2018), pp. 232–236.
- [68] Jae Yong Choi et al. “Off-target 18F-AV-1451 binding in the basal ganglia correlates with age-related iron accumulation”. In: *Journal of Nuclear Medicine* 59.1 (2018), pp. 117–120.
- [69] Keith A Johnson et al. “Tau positron emission tomographic imaging in aging and early Alzheimer disease”. In: *Annals of neurology* 79.1 (2016), pp. 110–119.
- [70] Hiroto Kuwabara et al. “Evaluation of 18F-RO-948 PET for quantitative assessment of tau accumulation in the human brain”. In: *Journal of Nuclear Medicine* 59.12 (2018), pp. 1877–1884.
- [71] Dean F Wong et al. “Characterization of 3 novel tau radiopharmaceuticals, 11C-RO-963, 11C-RO-643, and 18F-RO-948, in healthy controls and in Alzheimer subjects”. In: *Journal of Nuclear Medicine* 59.12 (2018), pp. 1869–1876.
- [72] Jacob W Vogel et al. “Spread of pathological tau proteins through communicating neurons in human Alzheimer’s disease”. In: *Nature Communications* 11.1 (2020), pp. 1–15.
- [73] Davina Biel et al. “Tau-PET and in vivo Braak-staging as prognostic markers of future cognitive decline in cognitively normal to demented individuals”. In: *Alzheimer’s research & therapy* 13.1 (2021), pp. 1–13.
- [74] Jacob W Vogel et al. “Four distinct trajectories of tau deposition identified in Alzheimer’s disease”. In: *Nature medicine* 27.5 (2021), pp. 871–881.
- [75] Amelie Schäfer, Elizabeth C Mormino, and Ellen Kuhl. “Network Diffusion Modeling Explains Longitudinal Tau PET Data”. In: *Frontiers in Neuroscience* 14 (2020), p. 1370.

- [76] Amelie Schäfer et al. “Bayesian Physics-Based Modeling of Tau Propagation in Alzheimer’s Disease”. In: *Frontiers in Physiology* (2021), p. 1081.
- [77] Ashish Raj, Amy Kuceyeski, and Michael Weiner. “A network diffusion model of disease progression in dementia”. In: *Neuron* 73.6 (2012), pp. 1204–1215.
- [78] Ashish Raj et al. “Network diffusion model of progression predicts longitudinal patterns of atrophy and metabolism in Alzheimer’s disease”. In: *Cell reports* 10.3 (2015), pp. 359–369.
- [79] Amelie Schäfer et al. “Correlating tau pathology to brain atrophy using a physics-based Bayesian model”. In: *Engineering with Computers* (2022), pp. 1–11.
- [80] Johannes Weickenmeier, Ellen Kuhl, and Alain Goriely. “The multiphysics of prion-like diseases: progression and atrophy”. In: *Phys Rev Lett* 121.158101 (2018).
- [81] Yasser Iturria-Medina et al. “Epidemic spreading model to characterize misfolded proteins propagation in aging and associated neurodegenerative disorders”. In: *PLoS Comput Biol* 10.11 (2014), e1003956.
- [82] Sveva Fornari et al. “Prion-like spreading of Alzheimer’s disease within the brain’s connectome”. In: *Journal of the Royal Society Interface* 16.159 (2019), p. 20190356.
- [83] S. Fornari et al. “Spatially-extended nucleation-aggregation-fragmentation models for the dynamics of prion-like neurodegenerative protein-spreading in the brain and its connectome”. In: *Journal of Theoretical Biology* 486 (2020), p. 110102.
- [84] Sara Garbarino, Marco Lorenzi, and Alzheimer’s Disease Neuroimaging Initiative. “Investigating hypotheses of neurodegeneration by learning dynamical systems of protein propagation in the brain”. In: *NeuroImage* 235 (2021). URL: <https://api.semanticscholar.org/CorpusID:232766910>.
- [85] Andrew Gelman, Jessica Hwang, and Aki Vehtari. “Understanding predictive information criteria for Bayesian models”. In: *Statistics and computing* 24.6 (2014), pp. 997–1016.
- [86] Prama Putra et al. “Braiding Braak and Braak: Staging patterns and model selection in network neurodegeneration”. In: *Network Neuroscience* (2021), pp. 1–41.
- [87] Frantisek Vasa and Bratislav Misic. “Null models in network neuroscience”. In: *Nature Reviews Neuroscience* 23.8 (2022), pp. 493–504.
- [88] Jacob W. Vogel et al. “Connectome-based modelling of neurodegenerative diseases: towards precision medicine and mechanistic insight”. In: *Nature Reviews Neuroscience* 24 (2023), pp. 620–639. URL: <https://api.semanticscholar.org/CorpusID:261121199>.

- [89] Travis B. Thompson et al. “Protein-protein interactions in neurodegenerative diseases: A conspiracy theory”. In: *PLoS Computational Biology* 16.10 (Oct. 2020), pp. 1–41. DOI: [10.1371/journal.pcbi.1008267](https://doi.org/10.1371/journal.pcbi.1008267). URL: <https://doi.org/10.1371/journal.pcbi.1008267>.
- [90] Sarah L DeVos et al. “Synaptic tau seeding precedes tau pathology in human Alzheimer’s disease brain”. In: *Frontiers in neuroscience* 12 (2018), p. 267.
- [91] David C. Van Essen et al. “The Human Connectome Project: A data acquisition perspective”. In: *NeuroImage* 62 (2012), pp. 2222–2231. URL: <https://api.semanticscholar.org/CorpusID:33357464>.
- [92] Stamatios N. Sotiropoulos et al. “Advances in diffusion MRI acquisition and processing in the Human Connectome Project”. In: *NeuroImage* 80 (2013), pp. 125–143. URL: <https://api.semanticscholar.org/CorpusID:207187214>.
- [93] Edward T. Bullmore and Danielle S. Bassett. “Brain graphs: graphical models of the human brain connectome.” In: *Annual review of clinical psychology* 7 (2011), pp. 113–40. URL: <https://api.semanticscholar.org/CorpusID:31287876>.
- [94] Stanley B Prusiner. “Molecular Biology of Prion Diseases”. In: *Science* 252.5012 (1991), pp. 1515–1522. DOI: [10.1126/science.1675487](https://doi.org/10.1126/science.1675487).
- [95] Stanley B Prusiner. “Prions”. In: *Proceedings of the National Academy of Sciences* 95.23 (1998), pp. 13363–13383.
- [96] Matthew D Hoffman, Andrew Gelman, et al. “The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo.” In: *J. Mach. Learn. Res.* 15.1 (2014), pp. 1593–1623.
- [97] Radford M Neal et al. “MCMC using Hamiltonian dynamics”. In: *Handbook of markov chain monte carlo* 2.11 (2011), p. 2.
- [98] Michael Betancourt. “A conceptual introduction to Hamiltonian Monte Carlo”. In: *arXiv preprint arXiv:1701.02434* (2017).
- [99] R. Dong et al. “Differential Elimination for Dynamical Models via Projections with Applications to Structural Identifiability”. In: *SIAM Journal on Applied Algebra and Geometry* 7.1 (2023), pp. 194–235. URL: <https://doi.org/10.1137/22M1469067>.
- [100] Alexandre Sedoglavic. “A probabilistic algorithm to test local algebraic observability in polynomial time”. In: *International Symposium on Symbolic and Algebraic Computation*. 2000. URL: <https://api.semanticscholar.org/CorpusID:10790056>.
- [101] Alejandro Fernández Villaverde, Antonio Barreiro, and Antonis Pappachristodoulou. “Structural Identifiability of Dynamic Systems Biology Models”. In: *PLoS Computational Biology* 12 (2016). URL: <https://api.semanticscholar.org/CorpusID:3614656>.

- [102] Andreas Raue et al. “Joining forces of Bayesian and frequentist methodology: a study for inference in the presence of non-identifiability”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 371.1984 (2013), p. 20110544.
- [103] Franz-Georg Wieland et al. “On structural and practical identifiability”. In: *Current Opinion in Systems Biology* 25 (2021), pp. 60–69.
- [104] Ivo Siekmann, James Sneyd, and Edmund J Crampin. “MCMC can detect nonidentifiable models”. In: *Biophysical journal* 103.11 (2012), pp. 2275–2286.
- [105] Sanjay Pant and Damiano Lombardi. “An information-theoretic approach to assess practical identifiability of parametric dynamical systems”. In: *Mathematical biosciences* 268 (2015), pp. 66–79.
- [106] Matthew Bobinski et al. “The histological validation of post mortem magnetic resonance imaging-determined hippocampal volume in Alzheimer’s disease”. In: *Neuroscience* 95.3 (1999), pp. 721–725.
- [107] Bradford C Dickerson et al. “The cortical signature of Alzheimer’s disease: regionally specific cortical thinning relates to symptom severity in very mild to mild AD dementia and is detectable in asymptomatic amyloid-positive individuals”. In: *Cerebral cortex* 19.3 (2009), pp. 497–510.
- [108] Sean M Nestor et al. “Ventricular enlargement as a possible measure of Alzheimer’s disease progression validated using the Alzheimer’s disease neuroimaging initiative database”. In: *Brain* 131.9 (2008), pp. 2443–2454.
- [109] Mert R Sabuncu et al. “The dynamics of cortical and hippocampal atrophy in Alzheimer disease”. In: *Archives of neurology* 68.8 (2011), pp. 1040–1048.
- [110] Nick C Fox and Peter A Freeborough. “Brain atrophy progression measured from registered serial MRI: validation and application to Alzheimer’s disease”. In: *Journal of Magnetic Resonance Imaging* 7.6 (1997), pp. 1069–1075.
- [111] Basil H Ridha et al. “Tracking atrophy progression in familial Alzheimer’s disease: a serial MRI study”. In: *The Lancet Neurology* 5.10 (2006), pp. 828–834.
- [112] CR Jack et al. “Atrophy rates accelerate in amnesic mild cognitive impairment”. In: *Neurology* 70.19 Part 2 (2008), pp. 1740–1752.
- [113] WJP Henneman et al. “Hippocampal atrophy rates in Alzheimer disease: added value over whole brain volume measures”. In: *Neurology* 72.11 (2009), pp. 999–1007.
- [114] Jasper D Sluimer et al. “Accelerating regional atrophy rates in the progression from normal aging to Alzheimer’s disease”. In: *European radiology* 19.12 (2009), pp. 2826–2833.
- [115] CR McDonald et al. “Regional rates of neocortical atrophy from normal aging to early Alzheimer disease”. In: *Neurology* 73.6 (2009), pp. 457–465.

- [116] Keith A Josephs et al. “Protein contributions to brain atrophy acceleration in Alzheimer’s disease and primary age-related tauopathy”. In: *Brain* 143.11 (2020), pp. 3463–3476.
- [117] Renaud La Joie et al. “Prospective longitudinal atrophy in Alzheimer’s disease correlates with the intensity and topography of baseline tau-PET”. In: *Science translational medicine* 12.524 (2020), eaau5732.
- [118] Sandhitsu R Das et al. “Longitudinal and cross-sectional structural magnetic resonance imaging correlates of AV-1451 uptake”. In: *Neurobiology of aging* 66 (2018), pp. 49–58.
- [119] Paul M Thompson et al. “Dynamics of gray matter loss in Alzheimer’s disease”. In: *Journal of neuroscience* 23.3 (2003), pp. 994–1005.
- [120] Giovanni B Frisoni et al. “In vivo mapping of incremental cortical atrophy from incipient to overt Alzheimer’s disease”. In: *Journal of neurology* 256.6 (2009), pp. 916–924.
- [121] Vivek Singh et al. “Spatial patterns of cortical thinning in mild cognitive impairment and Alzheimer’s disease”. In: *Brain* 129.11 (2006), pp. 2885–2893.
- [122] Clifford R Jack Jr et al. “Steps to standardization and validation of hippocampal volumetry as a biomarker in clinical trials and diagnostic criterion for Alzheimer’s disease”. In: *Alzheimer’s & Dementia* 7.4 (2011), pp. 474–485.
- [123] Susan M Landau et al. “Comparing positron emission tomography imaging and cerebrospinal fluid measurements of β -amyloid”. In: *Annals of Neurology* 74.6 (2013), pp. 826–836.
- [124] Amelie Schäfer et al. “Predicting brain atrophy from tau pathology: A summary of clinical findings and their translation into personalized models”. In: *Brain Multiphysics* 2 (2021), p. 100039.
- [125] ADNI. *Alzheimer’s Disease Neuroimaging Initiative*. <http://adni.loni.usc.edu>. accessed July 31, 2020.
- [126] S.L. Baker et al. “Reference Tissue-Based Kinetic Evaluation of 18F-AV1451 for Tau Imaging”. In: *J. Nucl. Med.* 58.2 (2017), pp. 332–338.
- [127] Rahul S Desikan et al. “An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest”. In: *Neuroimage* 31.3 (2006), pp. 968–980.
- [128] FreeSurfer. *FreeSurfer Software Suite*. <http://surfer.nmr.mgh.harvard.edu>. accessed July 31, 2020.
- [129] Alexandre Routier et al. “Clinica: An Open-Source Software Platform for Reproducible Clinical Neuroscience Studies”. In: *Frontiers in Neuroinformatics* 15 (2021).
- [130] Andrew Gelman et al. *Bayesian data analysis*. Chapman and Hall/CRC, 1995.

- [131] Christopher Rackauckas and Qing Nie. “Differential equations. jl—a performant and feature-rich ecosystem for solving differential equations in julia”. In: *Journal of Open Research Software* 5.1 (2017).
- [132] Mohamed Tarek et al. “DynamicPPL: Stan-like Speed for Dynamic Probabilistic Models”. In: *arXiv preprint arXiv:2002.02702* (2020).
- [133] Georg Meisl et al. “In vivo rate-determining steps of tau seed accumulation in Alzheimer’s disease”. In: *Science advances* 7.44 (2021), eabh1448.
- [134] Florence Clavaguera et al. “Transmission and spreading of tauopathy in transgenic mouse brain”. In: *Nature cell biology* 11.7 (2009), pp. 909–913.
- [135] Dunja Mrdjen et al. “The basis of cellular and regional vulnerability in Alzheimer’s disease”. In: *Acta Neuropathologica* 138 (2019), pp. 729–749.
- [136] Hongjun Fu, John Hardy, and Karen E Duff. “Selective vulnerability in neurodegenerative diseases”. In: *Nature neuroscience* 21.10 (2018), pp. 1350–1358.
- [137] Eli J Cornblath et al. “Computational modeling of tau pathology spread reveals patterns of regional vulnerability and the impact of a genetic risk factor”. In: *Science Advances* 7.24 (2021), eabg6677.
- [138] Rik Ossenkoppele et al. “Amyloid and tau PET-positive cognitively unimpaired individuals are at high risk for future cognitive decline”. In: *Nature Medicine* (2022), pp. 1–7.
- [139] Andrew Gelman. “Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper)”. In: (2006).
- [140] Andrew Gelman, Daniel Simpson, and Michael Betancourt. “The prior can often only be understood in the context of the likelihood”. In: *Entropy* 19.10 (2017), p. 555.
- [141] Antoine Leuzy et al. “Diagnostic performance of RO948 F 18 tau positron emission tomography in the differentiation of Alzheimer disease from other neurodegenerative disorders”. In: *JAMA neurology* 77.8 (2020), pp. 955–965.
- [142] Alexa Pichet Binette et al. “Amyloid-associated increases in soluble tau relate to tau aggregation rates and cognitive decline in early Alzheimer’s disease”. In: *Nature Communications* 13.1 (2022), p. 6635.
- [143] Sebastian Palmqvist et al. “Discriminative accuracy of plasma phospho-tau217 for Alzheimer disease vs other neurodegenerative disorders”. In: *Jama* 324.8 (2020), pp. 772–781.
- [144] Martí Colom-Cadena et al. “Synaptic oligomeric tau in Alzheimer’s disease — A potential culprit in the spread of tau pathology through the brain”. In: *Neuron* 111 (2023), 2170–2183.e6. URL: <https://api.semanticscholar.org/CorpusID:257623562>.
- [145] Sebastian Palmqvist et al. “Earliest accumulation of β -amyloid occurs within the default-mode network and concurrently affects brain connectivity”. In: *Nature communications* 8.1 (2017), p. 1214.

- [146] Randy L Buckner et al. “Molecular, structural, and functional characterization of Alzheimer’s disease: evidence for a relationship between default activity, amyloid, and memory”. In: *Journal of neuroscience* 25.34 (2005), pp. 7709–7717.
- [147] Amy M Pooler et al. “Physiological release of endogenous tau is stimulated by neuronal activity”. In: *EMBO reports* 14.4 (2013), pp. 389–394.
- [148] Jessica W Wu et al. “Neuronal activity enhances tau propagation and tau pathology in vivo”. In: *Nature neuroscience* 19.8 (2016), pp. 1085–1092.
- [149] Ying-Qiu Zheng et al. “Connectome architecture, gene expression and functional co-activation shape the propagation of misfolded proteins in neurodegenerative disease”. In: *bioRxiv* (2018), p. 449199.
- [150] MGSMS Goedert et al. “Multiple isoforms of human microtubule-associated protein tau: sequences and localization in neurofibrillary tangles of Alzheimer’s disease”. In: *Neuron* 3.4 (1989), pp. 519–526.
- [151] Aylin Dincer et al. “APOE ϵ 4 genotype, amyloid- β , and sex interact to predict tau in regions of high APOE mRNA expression”. In: *Science translational medicine* 14.671 (2022), eabl7646.
- [152] Jorge Sepulcre et al. “Neurogenetic contributions to amyloid beta and tau spreading in the human cortex”. In: *Nature medicine* 24.12 (2018), pp. 1910–1918.
- [153] Vikas Kotari et al. “Early tau detection in flortaucipir images: validation in autopsy-confirmed data and implications for disease progression”. In: *Alzheimer’s Research & Therapy* 15.1 (2023), p. 41.
- [154] Alex Whittington, David J Sharp, and Roger N Gunn. “Spatiotemporal distribution of β -amyloid in Alzheimer disease is the result of heterogeneous regional carrying capacities”. In: *Journal of Nuclear Medicine* 59.5 (2018), pp. 822–827.
- [155] LA Sandusky-Beltran and EM Sigurdsson. “Tau immunotherapies: Lessons learned, current status and future considerations”. In: *Neuropharmacology* 175 (2020), p. 108104.
- [156] Catherine J Mummery et al. “Tau-targeting antisense oligonucleotide MAP-TRx in mild Alzheimer’s disease: a phase 1b, randomized, placebo-controlled trial”. In: *Nature Medicine* (2023), pp. 1–11.
- [157] Charles Duyckaerts et al. “PART is part of Alzheimer disease”. In: *Acta Neuropathologica* 129 (2015), pp. 749–756. URL: <https://api.semanticscholar.org/CorpusID:2596840>.
- [158] Luis Aragão Gomes et al. “A β -induced acceleration of Alzheimer-related τ -pathology spreading and its association with prion protein”. In: *Acta Neuropathologica* 138 (2019), pp. 913–941. URL: <https://api.semanticscholar.org/CorpusID:199576991>.

- [159] PG Kevrekidis, Travis B Thompson, and Alain Goriely. “Anisotropic diffusion and traveling waves of toxic proteins in neurodegenerative diseases”. In: *Physics Letters A* 384.36 (2020), p. 126935.
- [160] Alex Whittington, David James Sharp, and Roger N. Gunn. “Spatiotemporal Distribution of β -Amyloid in Alzheimer Disease Is the Result of Heterogeneous Regional Carrying Capacities”. In: *The Journal of Nuclear Medicine* 59 (2017), pp. 822–827. URL: <https://api.semanticscholar.org/CorpusID:23458115>.
- [161] Christopher G. Schwarz et al. “A large-scale comparison of cortical thickness and volume methods for measuring Alzheimer’s disease severity”. In: *NeuroImage : Clinical* 11 (2016), pp. 802–812. URL: <https://api.semanticscholar.org/CorpusID:143596>.
- [162] Clifford R Jack Jr et al. “Brain β -amyloid load approaches a plateau”. In: *Neurology* 80.10 (2013), pp. 890–896.
- [163] Susan Landau et al. “Florbetapir (AV-45) processing methods”. In: *Alzheimer’s Disease Neuroimaging Initiative* (). URL: https://adni.bitbucket.io/reference/docs/UCBERKELEYAV45/UCBERKELEY_AV45_Methods_01.14.21.pdf (visited on 11/21/2022).
- [164] Wha Jin Lee et al. “Regional A β -tau interactions promote onset and acceleration of Alzheimer’s disease tau spreading”. In: *Neuron* 110 (2022), 1932–1943.e5. URL: <https://api.semanticscholar.org/CorpusID:248232554>.
- [165] Cristina Vergara et al. “Amyloid- β pathology enhances pathological fibrillary tau seeding induced by Alzheimer PHF in vivo”. In: *Acta Neuropathologica* 137 (2019), pp. 397–412. URL: <https://api.semanticscholar.org/CorpusID:57189668>.
- [166] Amy M. Pooler et al. “Amyloid accelerates tau propagation and toxicity in a model of early Alzheimer’s disease”. In: *Acta Neuropathologica Communications* 3 (2015). URL: <https://api.semanticscholar.org/CorpusID:14285186>.
- [167] Vincent Doré et al. “Relationship between amyloid and tau levels and its impact on tau spreading”. In: *European Journal of Nuclear Medicine and Molecular Imaging* 48 (2021), pp. 2225–2232. URL: <https://api.semanticscholar.org/CorpusID:231706561>.
- [168] Dietmar R Thal et al. “Phases of A β -deposition in the human brain and its relevance for the development of AD”. In: *Neurology* 58.12 (2002), pp. 1791–1800.
- [169] Marc Aurel Busche et al. “Clusters of Hyperactive Neurons Near Amyloid Plaques in a Mouse Model of Alzheimer’s Disease”. In: *Science* 321 (2008), pp. 1686–1689. URL: <https://api.semanticscholar.org/CorpusID:206514927>.

- [170] Marc Aurel Busche et al. “Tau impairs neural circuits, dominating amyloid- β effects, in Alzheimer models in vivo”. In: *Nature Neuroscience* 22 (2018), pp. 57–64. URL: <https://api.semanticscholar.org/CorpusID:56171563>.
- [171] Samuel S. Harris et al. “Tipping the Scales: Peptide-Dependent Dysregulation of Neural Circuit Dynamics in Alzheimer’s Disease”. In: *Neuron* 107 (2020), pp. 417–435. URL: <https://api.semanticscholar.org/CorpusID:219984953>.
- [172] Marc Aurel Busche and Bradley T. Hyman. “Synergy between amyloid- β and tau in Alzheimer’s disease”. In: *Nature Neuroscience* 23 (2020), pp. 1183–1193. URL: <https://api.semanticscholar.org/CorpusID:221101316>.
- [173] Joseph Giorgio et al. “Amyloid induced hyperexcitability in default mode network drives medial temporal hyperactivity and early tau accumulation”. In: *Neuron* 112 (2023), 676–686.e4. URL: <https://api.semanticscholar.org/CorpusID:266200525>.
- [174] Lyduine E. Collij et al. “Spatial-Temporal Patterns of β -Amyloid Accumulation”. In: *Neurology* 98 (2022), e1692–e1703. URL: <https://api.semanticscholar.org/CorpusID:247474796>.
- [175] Seun Jeon et al. “Topographical Heterogeneity of Alzheimer’s Disease Based on MR Imaging, Tau PET, and Amyloid PET”. In: *Frontiers in Aging Neuroscience* 11 (2019). URL: <https://api.semanticscholar.org/CorpusID:201116608>.
- [176] Daniel Ferreira and et al Westman Eric. “Distinct subtypes of Alzheimer’s disease based on patterns of brain atrophy: longitudinal trajectories and clinical applications”. In: *Scientific Reports* 7.1 (2017), pp. 2045–2322.
- [177] Melissa E Murray et al. “Neuropathologically defined subtypes of Alzheimer’s disease with distinct clinical characteristics: a retrospective study”. In: *The Lancet Neurology* 10.9 (2011), pp. 785–796.
- [178] Keith A. Johnson et al. “Appropriate use criteria for amyloid PET: A report of the Amyloid Imaging Task Force, the Society of Nuclear Medicine and Molecular Imaging, and the Alzheimer’s Association”. In: *Alzheimer’s & Dementia* 9 (2013), E1–E16. URL: <https://api.semanticscholar.org/CorpusID:38647266>.
- [179] “Evidence-based indications for the use of PET-CT in the United Kingdom 2016.” In: *Clinical radiology* 71 7 (2016), e171–88. URL: <https://api.semanticscholar.org/CorpusID:6320934>.
- [180] Magdalena A Kolanko et al. “Amyloid PET imaging in clinical practice”. In: *Practical Neurology* 20 (2020), pp. 451–462. URL: <https://api.semanticscholar.org/CorpusID:221912291>.
- [181] Oskar Hansson. “Biomarkers for neurodegenerative diseases”. In: *Nature Medicine* 27 (2021), pp. 954–963. URL: <https://api.semanticscholar.org/CorpusID:235336120>.

- [182] Oskar Hansson et al. “CSF biomarkers of Alzheimer’s disease concord with amyloid- β PET and predict clinical progression: A study of fully automated immunoassays in BioFINDER and ADNI cohorts”. In: *Alzheimer’s & dementia : the journal of the Alzheimer’s Association* 14 (2018), pp. 1470–1481. URL: <https://api.semanticscholar.org/CorpusID:3639199>.
- [183] Gemma Salvadó et al. “Novel CSF tau biomarkers can be used for disease staging of sporadic Alzheimer’s disease”. In: *medRxiv* (2023). URL: <https://api.semanticscholar.org/CorpusID:259923474>.
- [184] Prama Putra et al. “Front propagation and arrival times in networks with application to neurodegenerative diseases”. In: *bioRxiv* (2022). URL: <https://api.semanticscholar.org/CorpusID:245810229>.
- [185] Paul S. Aisen et al. “Early-stage Alzheimer disease: getting trial-ready”. In: *Nature Reviews. Neurology* 18 (2022), pp. 389–399. URL: <https://api.semanticscholar.org/CorpusID:247924562>.
- [186] Harald Hampel et al. “Designing the next-generation clinical care pathway for Alzheimer’s disease”. In: *Nature Aging* 2 (2022), pp. 692–703. URL: <https://api.semanticscholar.org/CorpusID:251708176>.
- [187] Meichen Yu, Olaf Sporns, and Andrew J. Saykin. “The human connectome in Alzheimer disease — relationship to biomarkers and genetics”. In: *Nature Reviews Neurology* 17 (2021), pp. 545–563. URL: <https://api.semanticscholar.org/CorpusID:236095831>.
- [188] David Scott et al. “Default Mode Network failure is associated with increased tau PET uptake”. In: *Alzheimer’s & Dementia* 18 (2022). URL: <https://api.semanticscholar.org/CorpusID:254879482>.
- [189] David T. Jones et al. “Cascading network failure across the Alzheimer’s disease spectrum”. In: *Brain* 139 (2015), pp. 547–562. URL: <https://api.semanticscholar.org/CorpusID:5144299>.
- [190] Nick Corriveau-Lecavalier et al. “Default mode network failure and neurodegeneration across aging and amnesic and dysexecutive Alzheimer’s disease”. In: *Brain Communications* 5 (2023). URL: <https://api.semanticscholar.org/CorpusID:257415803>.
- [191] Marc Aurel Busche et al. “Tau impairs neural circuits, dominating amyloid- β effects, in Alzheimer models in vivo”. In: *Threshold* 30.40 (2019), p. 50.