10-25-2019 10:00 AM

Understanding Interdisciplinary Corroboration: Lessons from a Review Paper in the Mind-Brain Sciences

Jaclyn Lanthier, *The University of Western Ontario*

Supervisor: Dr. Jacqueline Sullivan, *The University of Western Ontario*
A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Philosophy

Abstract

A current and popular view of the relationship between areas of the mind-brain sciences is one where inter-communication and cross-disciplinary collaboration is required to advance epistemically strong claims about the mind-brain. My goal in this dissertation is to analyze what it means for information from different areas of science to fit together to produce strong epistemic claims by addressing how and to what extent claims about the mind-brain are corroborated in scientific practice. Philosophers of science have advanced various concepts of the notion of fitting together information from different areas of science and its relation to scientific progress (e.g., Bickle, 1998, 2003, 2006; Darden and Maull, 1977; Mitchell, 2002, 2003; Mitchell and Dietrich, 2006; Nagel 1949, 1970, 1979; Piccinini and Craver, 2011; and Roskies, 2010). However, each of these concepts are vague and subject to multiple clarificatory questions. To get a handle on the notion of fitting together and its ability to produce epistemically strong claims about the mind-brain, I introduce the term 'interdisciplinary corroboration' as a placeholder for the various accounts of fitting together to facilitate my investigation of how claims about the mind-brain are corroborated by interdisciplinary evidence in practice. To address this question, I argue that looking at review papers (specifically, *Trends in Cognitive Sciences* review papers) is a good place to begin analyzing interdisciplinary corroboration. Accordingly, I conduct a two-part analysis of a *Trends in Cognitive Sciences* review paper by Eichenbaum (2013) entitled 'Memory on Time' structured around what Eichenbaum does and does not do to support the claims he advances about spatio-temporal memory organization. My analysis revealed that critical information was missing to adequately assess the extent to which the claims about memory organization were corroborated in Eichenbaum's paper. I use the lessons from my analysis to develop and advance a methodology for philosophers of science interested in knowledge production for evaluating review papers for corroboration in the mind-brain sciences.

Key Words: Philosophy of Neuroscience, Philosophy of Science, Mind-Brain Sciences, Interdisciplinary Corroboration, Review Papers, Systematic Review Papers, Meta-Analysis, Spatio-Temporal Organization of Memory, Scientific Progress, Scientific Explanation.

Summary for Lay Audience

The mind-brain sciences comprises multiple areas of science (e.g., psychology, biology, computational modelling and artificial intelligence, neuroscience, etc.). It has been argued by philosophers of the mind-brain sciences that claims about the mind-brain stand on firm epistemic footing (i.e., that they are well supported) when they are the product of evidence from those respective areas of the mind-brain sciences fitting together. However, the philosophical accounts advocating the notion of fitting together lack a clear explanation as to what constitutes fitting together; what does it mean for pieces of evidence from different areas of science to fit together in such a way that produces epistemically strong claims about the mind-brain? How and to what extent are claims about the mind-brain corroborated in scientific practice? This is the main question I address in my dissertation. I introduce the term ‘interdisciplinary corroboration’ as a placeholder for the various philosophical accounts of fitting together to facilitate my investigation of what it means for claims to be corroborated by interdisciplinary evidence in the mind-brain sciences. My definition of interdisciplinary corroboration reads that claims about the mind-brain are corroborated when interdisciplinary pieces of evidence point towards a general claim and when those pieces of evidence are mutually supportive. I argue that review papers (specifically, *Trends in Cognitive Sciences* review papers) are a good place to start investigating interdisciplinary corroboration insofar as these papers often contain arguments demonstrating how interdisciplinary evidence combines to corroborate a given claim. To get clear on how exactly this occurs, I conducted a two-part analysis of a review paper by Eichenbaum (2013) entitled ‘Memory on Time’ on the topic of memory organization (i.e., how memories are organized in space and time). My analysis revealed that key information was missing in the review paper that is necessary to properly assess how and to what extent claims about memory organization are corroborated. I use the lessons from my analysis to develop and advance a methodology for evaluating review papers with an eye for corroboration in the mind-brain sciences.

Acknowledgements

Firstly, I'd like to thank my supervisor, Dr. Jacqueline Sullivan, for her continual guidance, feedback, and encouragement throughout this process. Her work on experimental protocols (2009) and construct stabilization (2016) motivated a great deal of my thinking about how pieces of evidence from different areas of science fit together to produce claims that stand on firm epistemic footing. Without her ongoing support from my time in the Master's in philosophy program and throughout the PhD in philosophy program, this dissertation would not be possible.

My prospectus committee, Dr. Angela Mendelovici and Dr. Robert DiSalle, provided me with invaluable feedback on earlier versions of the dissertation for which I am deeply grateful. I am grateful for the insightful questions and comments from my thesis examination committee, Dr. Robert DiSalle, Dr. Christopher Viger, Dr. Jody Culham, and Dr. Sarah Robins.

I am appreciative for the feedback and support of other members of the lab associates' program in the Rotman Institute of Philosophy, especially Robert Foley and Derek Oswick. I am thankful to the Rotman Institute and Brain and Mind Institute for making such a program possible, and to Dr. Jacqueline Sullivan, Dr. Chris Viger, and Dr. Michael Anderson (the current coordinator of the Rotman-BMI lab associates' program) for their hard work creating, maintaining, and facilitating the program. I am thankful for the number of thought-provoking talks held by the Rotman Institute of Philosophy and the Brain and Mind Institute, particularly, the talks held during the Rethinking the Taxonomy of Psychology Conference spearheaded by Dr. Jacqueline Sullivan, Dr. Michael Anderson, and Dr. Tim Bayne. In 2014, I had the honor to hear Dr. Howard Eichenbaum (1947-2017) speak on 'The Hippocampus in Space and Time', hosted by the Rotman Institute of Philosophy. His talk sparked my interest and curiosity about memory and the ways in which pieces of evidence from different areas of the mind-brain sciences fit together to corroborate claims about memory.

From the Brain and Mind Institute, I am thankful to have been able to attend numerous informative talks via their Psychology Colloquium Series/Brown Bag series as well as their Neuroscience Speaker Series. I am grateful to the Brain and Mind Institute and especially to Dr.

Rhodri Cusack and Dr. Stephan Köhler for welcoming me into their respective labs; I learnt a great deal from both concerning research methodologies and practice. From the Memory lab, I am especially thankful to Anna Blumenthal and Jane Kouptsova for insightful conversations about their research and methods, and for helping me work through some challenging scientific papers.

I am grateful to my fellow colleagues Derek Oswick (Western University), Robert Foley (Western University), Christine Barbieri (Memorial University), and David Fuzes (Concordia University) for their helpful feedback at various stages of the dissertation.

Most importantly, I'd like to thank my family. My late mother, Heather Lanthier, my father Pierre Lanthier and my stepmother Sharon Lanthier, my mother-in-law Susan Reitz and my father-in-law André Langevin, as well as my sister Nicole Lanthier, and stepbrothers Parker, Riley, and Cooper Tewksbury for their continual support and encouragement. Finally, I'd like to thank my best friend and husband, Andrew Baum, for believing in me and showing me complete support and encouragement throughout this process.

Funding for this research was provided by the Social Sciences and Humanities Research Council of Canada; and the Rotman Institute of Philosophy.

Table of Contents

List of Figures

Preface

Multiple areas of the mind-brain sciences (e.g., psychology, cognitive science, biology, neuroscience, etc.) aim to advance knowledge about how the brain realizes psychological functions (e.g., decision making, memory, language, etc.). Kaplan (2017) has noted both scientists and philosophers of science alike concede that understanding the relationship between the different areas comprising the mind-brain sciences is critical for scientific progress. The current and dominant view of the relationship between the different areas of the mind-brain science is one where inter-communication and cross-disciplinary collaboration is required to advance epistemically strong claims about the mind-brain[1].

In this dissertation, I am interested in the integrative and collaborative relationship between different areas of the mind-brain science and how this relationship results in epistemically strong claims about the mind-brain. My goal in this dissertation is to analyze what it means for information from different areas of science to fit together in such a way that results in epistemically strong claims by addressing the question of how and to what extent claims about the mind-brain are corroborated in scientific practice.

This question was prompted by numerous accounts of scientific progress emphasizing the notion of interdisciplinary pieces of evidence fitting together in such a way that the resulting claim about the mind-brain stands on firm epistemic footing. That is, I noticed many accounts of progress in the philosophy of the mind-brain sciences argue that strong epistemic claims are the result of interdisciplinary evidence fitting together in such a way that the resulting claim is corroborated.

---

[1] A note on how I use the term 'mind-brain' is necessary. I am concerned primarily with how claims from different areas of the mind-brain sciences fit together in such a way that a claim about the mind-brain is corroborated. Due to this, and the fact that much of my analysis deals with scientific studies, I use the word 'mind' in the way that scientists in the mind-brain sciences use the term, i.e., to refer to cognitive processes/functions and states occurring within an organism that are 'causally responsible for but not identical to' that organism's behavior (Sullivan 2014,48). It is this set of states and processes that scientists in the mind-brain sciences investigate. While the goal of the sciences is often stated in terms of 'understanding the mind', scientists typically break this down into 'understanding an aspect of the mind'. By aspect of the mind I mean either a cognitive state (e.g., feeling happy, feeling hungry, etc.) or a cognitive process (e.g., procedural memory, decision-making, etc.). Thus, when I discuss how the mind-brain (or an aspect of the mind-brain) works, I understand these questions as referring to how the brain realizes psychological functions.

To better understand this phenomenon, I begin in chapter one by laying down the relevant motivation and background for my dissertation by surveying the history of different accounts of scientific progress emphasizing the fitting together of results from different areas of science as an integral part of scientific progress and strong explanations. I review some reductionist (Nagel 1949, 1961, 1979; Bickle 1998, 2003, 2006) and non-reductionist (Darden and Maull, 1977; Mitchell 2002, 2003; Mitchell and Dietrich 2006; Piccinini and Craver 2011; Roskies 2010) accounts of the relationship between areas of science emphasizing the phenomena of fitting together. While the surveyed accounts advance different concepts and descriptions of fitting together and what it entails for scientific progress, ultimately, both reductionist and non-reductionist accounts alike recognize the importance of collaboration amongst different areas of science such that evidence from each area be fit together (in some respect) to yield strong epistemic claims about the mind-brain.

A note on the scope and purpose of this chapter. There are many more accounts of scientific progress and fitting together than I survey in chapter one; I flag these accounts in footnote seven of chapter one. Including every and all accounts emphasizing the notion of different areas of science fitting together for scientific progress would go beyond the scope of my first chapter. Accordingly, I selected some key reductionist and non-reductionist accounts of progress to demonstrate that despite different concepts of fitting together, a common assumption underlying these accounts is that epistemically strong claims about the mind-brain require evidence from different areas of science to fit together in some way.

Despite this consensus, the accounts of fitting together that I survey are vague and elicit numerous questions as to how exactly this fitting together occurs. For example, it is unclear what criteria must be met in order to confidently assert that pieces of evidence fit together in such a way that its resulting claim stands on strong epistemic footing.

To get a handle on the nature of fitting together in the mind-brain sciences, I introduce the term ‘interdisciplinary corroboration’ (hereafter IC) in chapter two as a placeholder for what I think philosophers of science have been gesturing towards but have yet to analyze in depth. I introduce this term as a working definition whereby IC occurs when the results of different areas of the mind-brain sciences (using different protocols, methods, etc.) corroborate some claim about the mind-brain insofar as they all generally point towards a similar claim and insofar as those pieces of evidence are mutually supportive.

Given that IC is a working definition, it too requires elaboration. To shed light on the nature of IC, I suggest looking for instances of it in scientific practice. Specifically, I recommend looking at review papers, which are units of analysis surveying literature from a variety of areas of science on a given topic. While some review papers consist solely of a literature summary, other review papers include explanations and arguments. These arguments, which may be implicit or explicit, can involve showing how a given claim is corroborated by the literature surveyed in the review paper; because authors of review papers frequently cite interdisciplinary findings in their summary, and because they sometimes demonstrate how those findings corroborate a claim about the mind-brain, review papers are strong candidates to begin investigating the nature of interdisciplinary corroboration. Put differently, they are a good place to start investigating how and to what extent claims about the mind-brain are corroborated in scientific practice.

I argue that while there are numerous types of review papers on offer in the mind-brain sciences, looking at a *Trends in Cognitive Sciences* review paper is an ideal starting point to begin my investigation insofar as it contains components indicative of IC (i.e., the author of a review paper in *Trends in Cognitive Sciences* surveys interdisciplinary findings and sometimes includes arguments about how those findings fit together to corroborate a claim about the mind-brain).

A note on the scope of this chapter. In chapter two, I consider two options for beginning my investigation of how claims about the mind-brain are corroborated in scientific practice – i.e., a *Trends in Cognitive Sciences* review paper and a *Frontier in Neuroscience* systematic review paper. I argue are both well-suited for the task of investigating IC. However, due to the scope of my dissertation, I decided to focus on a *Trends in Cognitive Sciences* review to begin the process of investigating how, and to what extent, claims about the mind-brain are corroborated in scientific practice. I acknowledge, in chapter four, different avenues of research one could undertake in exploring the epistemic role of units of analysis that function to corroborate claims about the mind-brain and their implications for scientific progress. This includes extending my project by comparing a *Trends in Cognitive Sciences* paper with a *Frontiers in Neuroscience* systematic review paper to discern whether the additional components of the latter (i.e., adhering to systematic review guidelines and/or including a meta-analysis) increase the integrity of the corroboration of the claims they advance about the mind-brain.

To better understand how claims about the mind-brain are corroborated in a *Trends in Cognitive Sciences* review paper, I conducted a two-part analysis of a review paper entitled ‘Memory on Time’ (2013) (hereafter referred to as the case study) by Eichenbaum in chapters three and four.

Critically, I began my analysis in chapter three with a general understanding of how a *Trends in Cognitive Sciences* review paper works; this came from the descriptive work I provided for the different kinds of review papers found in mind-brain science journals in chapter two. Furthermore, my analysis is guided by some cursory criteria such as the following questions: “what is the main aim of the review paper”, “what claims does Eichenbaum defend”, “how does Eichenbaum show support for his claims”, “what kind of arguments does Eichenbaum provide”, and “does the case study represent a strong example of interdisciplinary corroboration”?

Guided by these questions, I structured chapter three around the work Eichenbaum does in the case study to defend the claims he advances. This work primarily consists of Eichenbaum stating findings from various interdisciplinary studies which he states support a given claim about memory organization. His arguments for the corroboration of these claims are implied by the language and terms he used to introduce findings for a given claim and to advance specific claims about memory organization. For example, when introducing evidence for claim *x*, Eichenbaum states that there is ‘substantial evidence’ supporting *x*, thus conveying to the reader that there are good reasons to take *x* as corroborated (since it is supported by substantial interdisciplinary evidence) (Eichenbaum 2013, 83).[2] My analysis of these terms indicated that the kind of corroboration in the case study was one where the claims were not to be understood as true, but instead as highly probable, approximately true, etc. Further, I argue because the cited findings are interdisciplinary and because they are published in high-calibre mind-brain journals with high impact factors and citation rates, the cited findings constitute reliable pieces of information (thus increasing the integrity of the corroboration of the claims in the case study). I

[2] Where *x* represents one of the claims contained within ‘Memory on Time’ (Eichenbaum 2013).

note that, *prima facie*, the case study appears to represent a strong example of interdisciplinary corroboration.

As I conducted my analysis in chapter three, which indicated Eichenbaum's argument for the claims he advances in the case study depended on critical terms such as 'substantial evidence', I began to notice that certain pieces of information were missing from Eichenbaum's argument that could potentially undermine the case study representing a clear case of IC (ibid). Specifically, I began to wonder about the following questions: "what counts as substantial evidence", "what is meant by consistency in the case study", and "how are anomalies factored into the findings supporting the claims found in the case study". These questions led me to conduct part two of my analysis in chapter four, where I address these questions and their implications for interdisciplinary corroboration.

Thus, my analysis in chapter three structured around the work Eichenbaum did to support the claims he advances prompted questions concerning just how clearly the case study represented a clear instance of IC. To better analyze how and to what extent corroboration occurs in the case study, I structured my analysis in chapter four around what features were missing in the case study (i.e., what Eichenbaum does not do to support the claims he advances).[3] For example, I noted that definitions for terms like 'substantial evidence' and 'consistency' were lacking; this is critical because Eichenbaum's implicit arguments hinges on these terms (ibid). I argued that if these key terms are not defined, it is difficult to discern whether and to what extent the claims in the case study are corroborated. I elaborated on the questions "what counts as substantial evidence", "what is meant by consistency in the case study", and "how are anomalies factored into the findings supporting the claims in the case study" in this chapter, and, in elaborating on these, it struck me that another important factor missing in the case study was an evaluation of the methods used by the papers Eichenbaum cites. I then demonstrate how all of these missing pieces of information complicate the case study representing a clear case of interdisciplinary corroboration.

[3] These missing features were determined by the questions and concerns that arose in conducting my analysis in chapter three.

From my analysis, I extracted three key lessons. The first lesson is that the fundamental assumptions shared by both reductionist and non-reductionist accounts of progress that I surveyed in chapter one – i.e., that epistemically strong claims are the result of evidence from different areas of science fitting together – is complex, and not nearly as straightforward as the literature makes it out to be. The second lesson is that my analysis prompts a revised definition of interdisciplinary corroboration. Interdisciplinary corroboration should now be understood as a phenomenon whereby findings corroborate claims about the mind-brain insofar as they are interdisciplinary, mutually supportive, and rule out coincidence, errors, and alternative hypotheses, and all point towards a similar claim. Finally, the third lesson emerging from my analysis is that certain critical pieces of information are missing to adequately assess the extent to which the claims Eichenbaum advances in the case study are corroborated.

After conducting my analysis and extracting its key lessons, I recognized that the questions and concerns I raise in chapter four would be helpful in evaluating future/different review papers for corroboration. From this, I advanced a methodology comprising an itemized list of questions to direct towards the evaluation of review papers for corroboration.

A note on the nature of my methodology is warranted. The methodology I advance in chapter four is meant as *a working framework*. In chapter two, I highlighted multiple types of units of analysis that could be considered for investigating interdisciplinary corroboration. However, due to the scope of my dissertation, I began my investigation of interdisciplinary corroboration via a *Trends in Cognitive Sciences* review paper. Accordingly, the methodology I advance pertains primarily to this kind of unit of analysis. As more research is conducted on the different units of analysis in the mind-brain sciences and interdisciplinary corroboration, I expect my methodology to be revised, and my definition of IC to become sharper.

I conclude, in chapter five, by summing up the dissertation, acknowledging its limitations, and noting the contributions of my dissertation for the mind-brain sciences and, more generally, for the philosophy of science and empirically engaged philosophy of mind.

A final note on what I am not trying to accomplish in this dissertation is necessary. I am not trying to establish whether any of the main claims Eichenbaum defends are true; nor is it my goal to assess the scientific rigor of Eichenbaum's paper. The purpose of the dissertation is also not to ascertain which of the surveyed accounts in the first chapter most accurately captures the nature of interdisciplinary corroboration in the mind-brain sciences. Instead, my main goal is to analyze one aspect of the nature of the integrationist relationship of the different areas comprising the mind-brain sciences by asking what it means for different pieces of evidence to fit together in a way that yields strong epistemic claims about the mind-brain. I introduced a case study of a *Trends in Cognitive Sciences* review paper to see what lessons can be learned regarding the nature of interdisciplinary corroboration. My analysis revealed that critical information is missing to adequately assess the extent to which corroboration has occurred, and, accordingly, I advanced an itemized list as a methodology for philosophers of science interested in knowledge production to consult when analyzing review papers for corroboration.

Chapter 1

1 Introduction

Multiple fields of science comprising the mind-brain sciences (e.g., psychology, cognitive science, neuroscience, biology, etc.) aim to understand and advance knowledge about the mind-brain.[4] Kaplan (2017) notes there is consensus between both scientists and philosophers alike that "… understanding the complex relationship between neuroscience and psychological science …" is critical for achieving progress in the mind-brain sciences (1).[5]

During the pinnacle of computational cognitive psychology (i.e., the mid-20th century), the relationship between these fields was viewed as one of autonomy (ibid). In other words, the respective phenomena of investigation, methods, and theories of each area of science were held to be completely independent from each other. There was thus no need for interdisciplinary collaboration or cross-field projects (ibid).

This understanding of the relationship between the different fields of mind-brain sciences has recently undergone substantial change. Nowadays, the autonomy view of the mind-brain sciences is considerably less popular. Instead, the autonomy view has been replaced by a characterization which sees the relationship as one of 'complete disciplinary integration and interdependence' (ibid).

In discussing various positions on the debate over which ontology of the mind-brain is the correct one, Anderson (2015) notes that *all* positions in this debate reject the traditional view of the autonomy of psychology and neuroscience (70). Thus, there is currently a general rejection of the autonomy view of the relations between different areas of the mind-brain sciences. Churchland and Sejnowski (2000) state whereas it was previously held that discoveries and explanations at the neuronal and the cognitive level were considered completely separate and of

---

[4] As I note in my preface, my use of the term 'mind-brain' throughout this dissertation refers to understanding how the brain realizes psychological functions.

[5] Kaplan (2017) typically refers to the relationship in question as consisting of the relationship between neuroscience and psychology. Unless otherwise specified, when I refer to this relationship, I understand it in the broader sense of comprising the multiple areas housed under the mind-brain sciences. These include but are not limited to: computer science and artificial intelligence, cognitive psychology, biology, neurobiology, neurophysiology, linguistics, etc.

little 'academic significance' to the other, there is now consensus amongst scientists that the time is ripe for the convergence of the mind-brain sciences (14). Kaplan (2017) attributes this shift to the birth of cognitive neuroscience, where from the beginning, contributors of the cognitive neuroscience movement (such as Michael S. Gazzaniga and George A. Miller who coined the term 'cognitive neuroscience' in 1979) acknowledged the 'integrative and interdisciplinary' character of the emerging field which shared the goal of trying to understand the mind-brain (Kaplan 2017, 2).[6]

Similarly, Boone and Piccinini (2016) argue against the traditional cognitive science view of explanations as distinct and autonomous; instead, they maintain that explanations in cognitive neuroscience must integrate the "… computational and representational functions and structures across multiple levels of organization in order to explain cognition" (1509). The research strategy appealed to in cognitive neuroscience, according to Churchland and Sejnowski, is neither top-down nor bottom-up, but is instead best understood as an *interaction* between the various fields in the mind-brain sciences whereby research from one field provides "…constraints, corrections, and inspiration for research at other levels" (Churchland and Sejnowski 2000, 14).[7]

---

[6] In a letter correspondence to Michael S. Gazzaniga from George A. Miller, Miller defines cognitive neuroscience as the following: "Cognitive neuroscientists attempt to discover the molecular logic of organic knowledge systems, i.e., the principles that, in addition to the principles of physics, chemistry, biology, and psychology, govern the behavior of inanimate matter in living knowledge systems" (Gazzaniga 2000, 11).

[7] Churchland and Sejnowski note that there are at minimum three different definitions of the term 'levels' (Churchland and Sejnowski 2000, 14). They highlight: levels of analysis, levels of organization, and levels of processing. Briefly, levels of analysis refers to the conceptual distinctions in terms of the different types of questions one can ask about a given phenomenon. Marr and Poggio (1977) distinguish three respective levels of analysis: the computational level (which decomposes a task into its main parts), the algorithmic level (which stipulates a form procedure to perform the task) and the level of physical implementation (Churchland and Sejnowski 2000, 14). The level of organization refers to how those three levels of analysis 'map onto the nervous system'. The structure of the nervous system involves "… molecules, synapses, neurons, networks, layers, maps, and systems. The range of structural organization implies, therefore, that there are many levels of implementation and that each has its companion task description" (ibid, 15). The implications of the structural organization of the nervous systems are that because there are as many types of task descriptions as there are levels of structure, the algorithmic level as well as the level of implementation are 'over simplified' (ibid). Critically, the levels of structural organization, while conceptually separable, cannot be physically separated. For example, they note that attention can rely on mechanisms at various levels; some may be found at the level of a local neural network, while others may be housed within (a) broader neural system(s) found in multiple locations of the brain (ibid). Finally, the level of processing is described as there being different degrees of information processing depending on the distance between the responding cells and the sensory input. As an example of fruitful integration of levels between different fields of the mind-brain sciences, Churchland and Sejnowski advance a case study of color vision whereby different theories were proven correct, but at different levels of structural organization (ibid 17-18).

Thus, the current and popular view of the relationship between the different fields of the mind-brain sciences is integrationist, interdisciplinary, and non-autonomous. This characterization allots a fundamental role to the notion of 'fitting together' for explanation and progress in the mind-brain sciences. That is, to understand an aspect of the mind-brain, pieces of evidence from different areas of the mind-brain science must fit together in some way to produce epistemically strong claims about the mind-brain; this is part and parcel of the view of the relationship between the mind-brain sciences being integrative, collaborative, and non-autonomous.

My goal in this dissertation is to assess and analyze this integrative nature of the mind-brain sciences; to tackle the question of what it means for different areas of the mind-brain sciences to be fit together in such a way that yields epistemically strong claims, thereby advancing knowledge of the mind-brain.[8]

My aim in this chapter primarily motivational; here, I review the history of different accounts of scientific progress emphasizing the phenomena of fitting together and highlight a key problem shared by the surveyed accounts. I begin in section 1.1 and 1.1.2 by reviewing some reductionist (Nagel 1949, 1961, 1979; Bickle 1998, 2003, 2006) and non-reductionist (Darden and Maull, 1977; Mitchell 2002, 2003; Mitchell and Dietrich 2006; Piccinini and Craver 2011; Roskies 2010) accounts of the relationship between areas of the mind-brain sciences emphasizing the fitting together of evidence from different areas of science as critical for scientific progress and strong epistemic claims about the mind-brain.[9] In section 1.2, I evaluate the surveyed accounts

---

Thus, when Churchland and Sejnowski (2000) claim that different fields comprising the mind-brain sciences (which often work at different levels) interact with each other by providing "… constraints, corrections, and inspiration for research at other levels" they mean that this type of interdisciplinary interaction will lead to 'a chain of models linking adjacent levels', ultimately providing us with a more comprehensive picture of how some feature of the mind-brain works (ibid, 14, 21).

[8] My goal for my dissertation and the question of how pieces of interdisciplinary evidence fit together to yield strong claims about the mind-brain was motivated by Sullivan's (2009) work on experimental protocols and her (2016) work on construct stabilization. Specifically, Sullivan's (2009, 2016) skepticism towards accounts of scientific progress claiming that results from different areas of the mind-brain sciences fit neatly together to yield claims about the mind-brain that stand on firm epistemic footing, led me to think more critically about how the concept of fitting together actually occurs in scientific practice.

[9] Reviewing *each and every* reductionist and non-reductionist account of progress and interdisciplinary fitting together would be far beyond the scope of my dissertation. For more reductive accounts c.f., Churchland's notion of 'co-evolution' (1986), Hooker's 'general theory of reduction' (1981), Schaffner's 'general reduction paradigm' (1967), and his revised 'general reduction-replacement model' (1992); for non-reductive accounts c.f., the notion

and highlight the common assumption underlying all accounts – namely, that progress is made when results from different areas of science are fit together yielding complete and strong epistemic claims. I then identify a key problem with this assumption; specifically, the notions of fitting together advanced by the surveyed accounts are vague and under-analyzed. I sum up this chapter in section 1.3, by noting why it is important to analyze this assumption and to get a handle on the notion of fitting together.

1.1 Historical Reviews of Scientific Progress in the Mind-Brain Sciences (Some Reductionist Approaches)

Reductionism is a popular account of explanation that offers a model for the relationship between different areas of science. Historically, reductionism was brought to the forefront of philosophy of science with the alleged reduction of Newtonian mechanics to Einstein's special theory of relativity, chemistry to atomic physics, gas laws to statistical mechanics, etc. (Curd and Cover 1998, 903). Nagel spearheaded reductionism by proposing a model of inter-theoretic reduction (Nagel 1949, 1961, 1979). Nagel believed that reductionism enabled a special kind of understanding; one which constituted/led to an 'important scientific achievement' (Nagel 1949, 123).[10] Briefly, Nagel took reductionism as an explanatory relation holding between two theories whereby one theory is derivable from the other, with the help of bridge laws (when applicable). Such a relation holds only if two formal conditions plus an essential condition are met; if one or more conditions are not satisfied, the alleged reduction would fail (Nagel 1949, 1979). I will first consider each condition in detail, and then discuss Nagel's inter-theoretic model of reduction and its relation to the concepts of fitting together and scientific progress.

As stated above, Nagel construed reductionism as an explanatory relation holding between the theories of two sciences whereby one theory is derivable from the other. Nagel took the set of theories and/or experimental laws comprising the reduced theory as the 'secondary science', and the theory which enables the reduction as the 'primary science' (Nagel 1949, 107; Nagel 1979, 338). On Nagel's view, a reduction occurs when the claims of a secondary sciences'

---

of 'cross-disciplinary patterns' by Bechtel (1986), 'decomposition' by Bechtel and Richardson (1993), 'integration' by Boone and Piccinini (2016), 'coordinated contributions' by Hastings et al. (2014), and the notion of 'robustness' by Wimsatt (1980, 2007).

[10] Nagel contrasts 'important scientific achievement' with 'trivial scientific accomplishments', which are often *ad hoc* in nature (Nagel 1949, 123). I elaborate on the two types of achievements shortly.

theories/experimental laws are demonstrated to be *the logical consequence* of the fundamental claims of the primary science. This is the first formal condition of reduction – *the derivability condition*. That is, the claims from the secondary science are *logically derivable* from those of the primary science.

Nagel identifies two main types of reductions (i.e., homogenous and heterogenous), each of which have different implications for the derivability condition (Nagel 1979, 339). The former kind of reduction occurs when "… the laws of the secondary science employ no descriptive terms that are not also used with approximately the same meanings in the primary science. Reductions of this type can therefore be regarded as establishing deductive relations between two sets of statements that employ a homogenous vocabulary" (ibid). These reductions are relatively unproblematic since the two sciences share (more or less) the same descriptive terms. Accordingly, the derivability condition can be easily satisfied (Nagel 1949, 103).

As an example of a homogenous reduction, Nagel cites the reduction of Galilean science to Newtonian mechanics. Nagel notes while it is undoubtedly true that one can differentiate between the subject matters of Galilean science and Newtonian mechanics, there is an important sense in which these subject matters are 'homogenous and continuous' (ibid). That is, both sciences are inquiring into the motions of bodies and the causal underpinnings of these motions, and there are no descriptive terms that appear in Galilean science that do not also appear with approximately the same meaning in Newtonian mechanics (ibid). Thus, because there are no terms which are discontinuous (i.e., the terms in the former possess more or less the same meaning as those found in the latter), the derivation of the secondary science to the primary science is uncomplicated and uncontroversial (ibid 103-104; Nagel 1979, 338).

Heterogenous reductions, however, are more complicated and are a source of great 'intellectual discomfort' (Nagel 1979, 338). This discomfort is in part due to the fact that the subject matters between the primary and secondary science are 'qualitatively discontinuous' and that "… the secondary science employs in its formulations of laws and theories a number of distinctive descriptive predicates that are not included in the basic theoretical terms or in the associated rules of correspondence of a primary science" (ibid, 342).

As an example of a heterogenous case, Nagel cites the reduction of thermodynamics to statistical mechanics. Specifically, the terms 'temperature' 'heat' and 'entropy' as well as other general assumptions found in the former are not present in nor entailed by principles in mechanics. (Nagel 1949, 121, and Nagel 1979, 342-343). The discontinuity of subject matter and descriptive terms makes satisfying the derivability condition impossible. This is because, according to Klee (1999), the principles of logic stipulate that "… no term can appear in the conclusion of a formal demonstration unless the term also appears in the premises" (112). The only way the derivability condition can be satisfied for heterogenous reductions is by satisfying Nagel's second formal condition: *the connectability condition* (Nagel 1949, 119).

The connectability condition requires that one establish a 'proper kind of relation' between those expressions contained in the laws of the secondary science and the fundamental principles of the primary science (ibid, 120). This relation – also known as a bi-conditional bridge law – can be understood in one of two ways: either as a *logical relation* whereby a term in the secondary science is related to some expression in the primary science via synonymity or entailment, or, by adopting *a material or physical hypothesis* whereby "… the occurrence of the properties designated by some expression in the premises of the primary science is a sufficient, or necessary and sufficient, condition for the occurrence of the properties designated by the expression of the secondary discipline" (ibid). Thus, if there is some term/expression *x* in the experimental laws of the secondary science not found in the theoretical assumptions of the primary science, a suitable relation/bridge law must be established between *x* and the theoretical assumptions of the primary science. In the case of a material or physical hypothesis, however, because the meaning of the term/expression contained in the secondary science cannot be demonstrated as 'analytically related' to the expression in the primary science, the material or physical hypothesis is at best, 'contingently true' (ibid, 121).

Nagel examines the two possible relations via the derivation of the Boyle-Charles law from the kinetic theory of gases. Nagel notes that the term 'temperature' – which, for the sake of simplicity, Nagel considers as the sole term occurring in the laws of the secondary science that does not occur in the expressions of the primary science – as used in thermodynamics is neither synonymous with, nor entailed by the expression 'mean kinetic energy' (ibid). Accordingly, a successful deduction would require that an additional assumption be introduced, and in this case that assumption would come in the form of an empirical hypothesis stating "… a determinate factual connection between two properties of physical systems that are in principle independently identifiable …" (ibid).

Nagel considers the reduction of a secondary science to a primary science to be constitutive of an 'important scientific achievement' crucial for scientific progress (ibid, 123). However, the satisfaction of the two formal conditions does not guarantee an important scientific achievement. Nagel notes: "For if the premises of an alleged primary science could be selected quite arbitrarily, subject only to formal requirements that have been mentioned thus far, the logical deduction of the laws of a secondary science from such premises selected *ad hoc* would in most cases represent only a *trivial scientific accomplishment*" (ibid, second emphasis mine). Since a trivial scientific accomplishment would be largely inconsequential to scientific progress, Nagel introduced an essential condition, which requires the assumptions of the primary science be empirically supported by evidence that has 'some measure of adequacy' (ibid).[11] Nagel elaborates on the essential condition by returning to the example of the reduction of thermodynamics to mechanics:

> It is well known that the general assumption according to which physical bodies in different states of aggregation are systems of molecules is c*onfirmed by a large number of well-established facts of chemistry and of molar physics,* facts which are not primarily about thermal properties of bodies. Accordingly, the adoption of this hypothesis for the special task of accounting for the thermal behavior of gases is in line with the normal strategy of the natural sciences to *extend the use of ideas fruitful in one set of inquiries into related domains*. Similarly, the fundamental principles of mechanics, which serve as partial premises in the reduction of thermodynamics to mechanics, *are supported by evidence drawn from many fields of study distinct from the study of gases* (ibid, 124, emphases mine).

I will now unpack this quote in the context of Nagel's inter-theoretic reduction model to draw out some implications of his model for the notion of fitting and scientific progress. To see how fitting together contributes to scientific progress, it is helpful to consider Nagel's distinction between good and bad science.

[11] Nagel says no more here on what is meant by adequacy, stating that it is an issue beyond the scope of his work.

Bad science is the product of unsuccessful and *ad hoc* reductions; bad science results in trivial scientific accomplishments (ibid, 123). Good science, on the other hand, is the product of successful reductions and results in important scientific achievements (ibid). A successful reduction requires the two formal conditions *plus* an essential condition are satisfied. The essential condition requires assumptions in the primary science to be supported by empirical evidence that possesses a certain degree of adequacy (ibid). Nagel provides no explanation of what he means by 'adequacy', but at the very least, Nagel seems committed to the idea that assumptions are empirically adequate to the extent that they are well-confirmed, and assumptions are well-confirmed to the extent that there is *evidence from multiple fields of science* supporting them. In Nagel's example of the reduction of thermodynamics to mechanics (quoted above), he states the assumptions of mechanics (i.e., the primary science) are *supported by numerous different fields of study*; this is where the notion of fitting together enters into Nagel's model of inter-theoretic reduction and scientific progress. That is, if the assumptions of the primary science are not empirically supported by different areas of science, the essential condition is not satisfied. If this condition is not satisfied, the reduction is unsuccessful. By this, Nagel seems to imply that for the assumptions of the primary science to be empirically adequate, they must be supported by evidence from different areas of science. Thus, different areas of science must fit together in such a way that they support/confirm the assumptions of the primary science. To the extent that this occurs, the essential condition is satisfied. If the essential condition and formal conditions are satisfied, a successful reduction occurs. This is important because Nagel takes successful reductions as paradigmatic of important scientific achievements. This is the relation of the notion of fitting together and scientific progress for Nagel's inter-theoretic model of reduction.

Another more recent account of reduction has been proposed in the mind-brain sciences by Bickle (2003, 2006). Building on his 1998 book, *Psychoneural Reduction*, Bickle advances a 'reductive-in-practice' model in his (2003) *Philosophy and Neuroscience: A Ruthlessly Reductive Account*, and again in his (2006) 'Reducing Mind to Molecular Pathways: Explicating the Reductionism Implicit in Current Cellular and Molecular Neuroscience'. In these works, Bickle argues that a ruthless reduction is 'alive and thriving' in the field of cellular and molecular neuroscience and that in this field, low-level neuroscience can *directly explain* *behavior* via a methodology Bickle dubs 'intervening molecularly and tracking behaviorally' (hereafter the IMTB method) (Bickle 2003, 2-3; Bickle 2006, 411, 425).

One of the motivating factors behind Bickle's works is to dispel the anti-reductionist worry that low-level neuroscience is incapable of directly explaining behavior (Bickle 2006, 411). Drawing on Kandel, Schwartz, and Jessell's (2000) *Principles of Neural Science*, Bickle argues that most working neuroscientists agree with the main tenets of the textbook; namely, that it is, in fact, possible to link psychological concepts to 'molecular-biological mechanisms and pathways' (Bickle 2006, 412). For Bickle, these links are constitutive of a ruthless reduction whereby behavior is directly explained by low-level neuroscience.

A great deal of anti-reductionist push-back stems from a deep-seated doubt about the potential for neuroscience to bypass multiple levels in one shot. That is, reduction in the mind-brain sciences was traditionally understood as being composed of multiple levels such as the behavioral level, the information processing level, the neuroanatomical level, etc.[12] The thought was that in order to claim a mind-to-molecules reduction, intermediary bridges needed to be established between *each and every level*.[13] However, Bickle argues, this step-by-step reductionist picture is misguided and can be completely discarded once one shifts their focus to work in cellular and molecular neuroscience. In this field, Bickle states: "… molecular neuroscientists have developed experimental practices that bridge the behavioral to the molecular pathway levels *directly*; and these practices are common to all recent empirical successes …" (ibid, 414). Further, when one devotes attention to the field of cellular and molecular neuroscience, it is demonstrably clear that a step-by-step level-based reduction is unnecessary since there is a type of ruthless reduction that can and has bypassed multiple levels by explaining behavior directly at the cellular/molecular level (ibid).

While Bickle never explicitly states that this kind of ruthless reduction constitutes scientific progress, resulting in epistemically strong explanations about the mind-brain, it is implied by his account of reduction. For instance, Bickle quotes *the Principles of Neural Science* (1991),

[12] See Bickle's (2006) figure 1 for a visual representation of this reduction via levels model on page 413.

[13] This seems akin to the picture Churchland and Sejnowski paint when they describe interdisciplinary interaction as 'a chain of models linking adjacent levels', providing a more comprehensive explanation of the mind (Churchland and Sejnowski 2000, 14).

wherein Kandel, Schwartz, and Jessell assert that the ultimate explanandum of the mind-brain sciences is to discover and understand how the mind works (i.e., how it thinks, perceives, remembers, processes, etc.) (Bickle 2003, 3). As stated above, Bickle notes that neuroscientists already have experimental practices in place that *bridge* molecular pathways *directly* to behavior, and that further, *all successful empirical cases share these experimental practices* (Bickle 2006, 414). It is those 'bridges' or 'links' that do the work required by Kandel, Schwartz, and Jessell (1991, 2000) to answer the 'ultimate explanandum' (i.e., the mind) (Bickle 2003, 3; Bickle 2006, 414).

What experimental methods are employed by working neuroscientists enabling them to explain behavior to molecules directly? One of these methods is what Bickle dubs the 'intervene molecularly and track behaviorally' method (hereafter the IMTB method). This method is of great significance, since, for Bickle, this method is tantamount to, "… what "reductionism" amounts to within the current discipline" (Bickle 2006, 425). The IMTB method provides a new way to conceptualize reductionism contrasting with inter-theoretic options which dominated a fair part of modern philosophy of science (ibid). The IMTB method is comprised of two parts: first, one causally *intervenes* "… at the level of cellular activity or molecular pathways within specific neurons …", and second, one *tracks* "… the effects of these interventions under controlled experimental conditions using behavioral protocols well accepted within experimental psychology" (ibid). Bickle explains the implications of this method:

> … this methodology constitutes an implicit condition on *explanation* in this field. One only claims a successful *explanation*, a successful *search for a cellular or molecular mechanism*, or a successful *reduction*, of a psychological kind when one successfully intervenes at the lower level and then measures a statistically significant behavioral difference (ibid).

Bickle uses Kogan, Frankland, and Silva's (2000) work on social recognition memory in mammals to illustrate the IMTB method in practice. Kogan et al. (2000) intervened in the cAMP-PKA-CREB pathway in mice by inserting a *Neo gene* which does not allow the activator α and δ isoforms of the CREB molecule to develop (Bickle 2006, 420). Kogan et al. (2000) then applied a "… hippocampus-dependent fear-conditioning memory task (memory for a novel environmental context-foot shock association) …" to the CREBαδ− knock-outs which demonstrated impairment in long-term memory despite in-tact short-term memory (Bickle 2006, 421-422).[14] Kogan et al. (2000) applied a behavioral paradigm developed by Thor and Holloway (1982) to their work with CREBαδ− knock-out mice in addition to employing different experimental control groups (Bickle 2006, 422). Their results demonstrated:

> … CREBαδ—mutants displayed normal short-term social recognition memory for a juvenile conspecific presented earlier. Thirty minutes after the initial exposure to it, they spent the same amount of time (statistically) in stereotypic investigative interaction with the juvenile as did the wild-type control adults. … However, CREBαδ—mutants were seriously impaired in long-term recognition memory. In fact, the mutants spent the same amount of time (statistically) investigating the same juvenile 24 h after initial exposure as they did during the initial exposure. This result indicates complete long-term amnesia despite intact short term-recognition (ibid).

Bickle takes this result as *confirmation* of the claim that there is "… a molecular mechanism of the consolidation switch for a cognitively robust kind of memory" (ibid). Thus, because interventions at the molecular level (i.e., the cAMP-PKA-CREB pathway) allowed for the tracking of targeted behavior (i.e., social recognition memory) via the 'widely accepted' Thor and Holloway (1980) protocol for studying the nature of social recognition memory, Bickle concludes by stating this example constitutes the gold standard for how hypotheses and claims about the mind-brain can be confirmed reductively in a way that bypasses multi-level steps (by moving *directly* from the cellular/molecular level to behavior)[15] (ibid, 425-427).

However, if it is the case that Bickle's account of reduction does away with different levels involved in traditional accounts of reduction (e.g., work done at the information processing level, the neuroanatomical level, etc.), each of which is studied by different fields of science, and we can explain behavior directly by appeal to the cellular/molecular level, where does the concept of fitting together come into play? *Prima facie,* it appears to have little to no role at all. That is, if the 'explanatory relevance' of higher-level theorizing is abandoned in favor of the IMTB method

[14] See Bickle (2006) pages 420-425 for an in-depth explanation of Kogan et al.'s (2000) experiments and results.
[15] See Sullivan (2009) for a critique of Bickle's reduction-in-practice model.

and the levels held between behavior and molecules on the traditional view disappear, then there seems to be no place for higher-level theorizing from different areas of the mind-brain sciences.

Importantly, in arguing for a ruthless reduction that bypasses multi-level steps, Bickle does not advocate a wholescale elimination of high-level theorizing, but instead offers it a new job, so to speak. He maintains that the methods employed by higher-level theorizing in contemporary neuroscience such as fMRI, cognitive and computational modelling, neuropsychological assessment, etc., are crucial for successful reductions (and ultimately achieving progress with respect to understanding the mind-brain) (ibid, 427). Further, Bickle insists this kind of theorizing and their respective methods are *essential* for explaining the link from molecules to behavior, as the theorizing conducted at higher levels helps address critical questions such as the following: "What are good experimental protocols for tracking behavioral outcomes for the psychological phenomenon we seek the cellular and molecular mechanisms of … and "Where shall we begin inserting our cellular and molecular interventions? (The possibility space in both brains and intra-neuron molecular pathways is enormous!)" (ibid).

Bickle classifies these kinds of questions as heuristic, and accordingly, considers the role for higher-level theorizing as one of heuristics (ibid). Bickle argues it would be quite difficult to address these questions at the cellular/molecular level; that is, he takes these higher-level questions as fundamental and ineliminable in the search for cellular and molecular mechanisms (ibid). This interplay of higher-level theorizing about behavioral and/or cognitive phenomena, along with the work conducted at the lower-level neuroscience level via the IMTB method, demonstrates the importance of interdisciplinary fitting together for scientific progress and strong explanations about the mind-brain. The phenomenon of fitting together, however, does not sit at the level of the actual reduction itself (as it does with many inter-theoretic models of reduction) but instead precedes it. That is:

> … from a diachronic perspective on mind-brain science, investigation into a cognitive or behavioral phenomenon initially looks like the approach illustrated in Fig. 1 with investigations and explanations at multiple levels and the search for step-by-step linkages of features down the level's hierarchy. But when this methodology reaches candidate cellular or molecular processes in specific neurons linked in the ways revealed by neuroanatomy,

investigation shifts to the approach in Fig. 8. And when these investigations succeed, a reduction of mind to molecular pathways is taken as established (ibid, 428).[16]

The role of higher-level theorizing comes into play in the early stages of investigation by isolating the "… relevant neuroanatomy and the candidate cellular and molecular mechanisms …" so that the work at the cellular/molecular level can proceed via the IMTB method (ibid). While Bickle never explicitly argues this, his account seems to imply that *if* higher-level theorizing were eliminated, *then* work at the cellular/molecular level would not get off the ground, and thus the prospect of reducing mind to molecular pathways would be compromised.

In fact, Bickle (2003) argues that transdisciplinary projects (i.e., projects that incorporate questions, data, theories, etc., from different areas of the mind-brain sciences) are key for making headway in the search for cellular/molecular mechanisms. He cites his own project (Bickle et. al 2000) which combines methods, theories, data, etc., from single-cell physiology, functional neuroimaging and biological modelling to help answer what cellular/molecular mechanism is responsible for saccade sequences in the frontal cortex and how they are processed (Bickle 2003, 115-130). According to this transdisciplinary picture, aspects of both high and low-level neuroscience need to fit together for progress (i.e., for the reduction of mind to molecules) to occur.[17]

In sum, Bickle argues that low-level neuroscience can directly explain behavior by appealing to experimental practices such as the IMTB method commonly used by scientists in cellular/molecular neuroscience. However, low-level neuroscience still needs higher-level theorizing in order to adequately pick out and isolate the desired cellular/molecular mechanism in question; this is where the relationship between progress and fitting together lies.

Before moving onto non-reductionist accounts of progress, I will briefly recapitulate how the notion of fitting together plays into Nagel and Bickle's accounts of reductionism. For Nagel, the

---

[16] Figure 1 displayed the traditional model of reduction as a step-by-step hierarchy of levels (see Bickle 2006, page 413 for illustration) and figure 8 displayed Bickle's reductive-in-practice or ruthless reduction model which bypasses multiple levels by linking the behavioral level directly to the cellular molecular level (see Bickle 2006, page 426 for illustration).

[17] See Bickle (2003) for an example of a transdisciplinary project which combines questions, theories, and data from single-cell physiology, neurocomputational modelling, and functional neuroimaging (115-130).

role of fitting together comes into play with the essential condition which claims that the assumptions of the primary science must be well-supported by different areas of science. Despite Nagel failing to provide an account of what it means for evidence to be well-supported and how evidence from multiple fields confirms assumptions in the primary science, it is clear that the phenomenon of fitting together interdisciplinary evidence is required for successful reductions, which, on Nagel's view, are the key to scientific progress. That is, the results from these areas of science must fit together in such a way that lend empirical support to the assumptions of the primary science for the reduction to stand on firm epistemic footing. For Bickle, the notion of fitting together occurs between the interaction between high-level theorizing which occurs prior to actual investigations and discoveries at the cellular/molecular level. Similarly, Bickle does not provide an in-depth analysis of what fitting together would look like nor what it entails; nonetheless, the fitting together of the higher and lower-levels of the mind-brain sciences is critical for the reduction of mind-to-molecules (and thus for scientific progress and strong explanations of the mind-brain).

1.1.1 Historical Reviews of Scientific Progress in the Mind-Brain Sciences (Some Non-Reductionist Approaches)

Reductionist models are not the only accounts of scientific progress emphasizing the importance of different areas of science fitting together; non-reductive alternatives similarly emphasize this notion as central to scientific progress and strong explanations. Some non-reductive alternatives include the following concepts put forward to explain the nature of how different areas fit together in a way that is constitutive of progress and strong explanations: 'interfield theories' by Darden and Maull (1977), 'integrative pluralism' Mitchell (2002, 2003), 'convergence' by Roskies (2010), and 'seamless integration' by Piccinini and Craver (2011).

During Nagel's time, the focus was on relations between different *theories* within different areas of science (Darden and Maull 1977, 43). Nagel's inter-theoretic reduction model demonstrated the derivational reduction of a *theory* from a secondary science to a *theory* from the primary science (ibid). Instead of trying to understand the relations between different areas of science by focusing on their *theories*, Darden and Maull shift their focus to understanding the kinds of relations that exist between *fields* in their (1977) paper 'Interfield Theories'. Briefly, examples of fields include cytology, genetics, biochemistry, etc. (ibid). Fields may contain a theory or multiple theories; genetics, for example, contains the classical theory of the gene (ibid). Theories such as the classical theory of the gene are what they dub 'intrafield theories' – i.e., theories that exist *within* a particular field. Finally, interrelations between different fields can be generated via what they call 'interfield theories' (ibid). Interfield theories bridge different fields of science, resulting in scientific progress. To better understand how interfield theories constitute progress, I begin by reviewing what a field is, the kinds of relations that can exist between them, and the motivations for generating an interfield theory. I elaborate on their account of interfield theories by drawing on their example of the chromosome theory of Mendelian heredity and state its implications for scientific progress and the notion of fitting together.[18]

Darden and Maull define the term field as a field of science possessing the following characteristics: (a) a field must have a central problem, (b) it must possess a domain of items that are understood as facts relating to that problem, (c) it must possess general goals and explanatory factors guiding scientists on how to solve the central problem, (d) it must possess its respective techniques and methodologies, (e) they frequently employ their own special vocabulary, and sometimes (f) but not always, fields will have concepts, laws, and theories that are associated with the central problem (a) that help to realize its explanatory goals (c) (ibid, 44).

For example, in the field of physical chemistry, the central problem (a) was to determine how different parts of molecules interact, its domain (b) consisted of molecules, their respective parts and interactions, and the development of new techniques (d) such as mass spectrometry have allowed scientists to determine (c) the structure and interaction of molecules.[19] At any given

---

[18] While Darden and Maull (1977), Mitchell (2002, 2003), and Mitchell and Dietrich (2006) write and focus on examples in the biological sciences specifically, I have included them in my survey since the biological sciences are housed under the larger umbrella of cognitive neuroscience and/or (more generally) the mind-brain sciences. Recall that George A. Miller defined 'cognitive neuroscience' as the following: "Cognitive neuroscientists attempt to discover the molecular logic of organic knowledge systems, i.e., the principles that, in addition to the principles of physics, chemistry, *biology*, and psychology, govern the behavior of inanimate matter in living knowledge systems" (Gazzaniga 2000, 11, emphasis mine).
In a different characterization, Bechtel and Richardson (1993) use the term biology generally to characterize the 'life sciences' as comprising, "… physiology, cellular and molecular biology, biochemistry, genetics, evolutionary biology and neuroscience" (xvii). Therefore, I have included accounts focusing biology since biology is considered by many to be a part of the mind-brain sciences.

[19] For similar examples in the fields of cytology, genetics, and biochemistry, see pages 46-47 (Darden and Maull, 1977).

time, a field may contain either no theory, a working theory, multiple competing theories, or one main successful theory (ibid, 48).[20] Sometimes when theories are well confirmed they become part of the general domain, and one might consequently look for a 'more encompassing' theory to explain said domain (ibid). One final point worth noting is that fields cannot be reduced to another. Darden and Maull state: "… reduction in the sense of derivation would be impossible between such elements of a field as techniques and explanatory goals" (ibid). Presumably, this is because they consider fields to be distinct (insofar as they possess unique sets of vocabulary, a specific problem and set of tools to address the problem etc.) making it impossible for a classical reduction to occur.[21] Taken together, this is how Darden and Maull understand and use the term field.[22] [23]

Fields bear relations to other fields via what Darden and Maull dub 'interfield theories.[24] An interfield theory describes and explains relations that exist between different fields. Darden and Maull highlight four main types of relations that can hold between fields. The first type is the *part-whole relation* whereby one field demonstrates a specific physical location for some entity/process assumed in a different field. For example, the chromosome theory of Mendelian heredity posited that Mendelian genes were located either *in or on* the chromosomes; the physical location required specification. Cytology had the tools necessary to specify the physical location of genes (ibid, 49). With this knowledge from cytology (i.e., that genes are in the chromosomes), the relation between the two fields became complete, hence it being a part-whole relation (ibid). The second type of relation is the *physical nature relation*; here, a field may present the physical nature of some entity/process held by another field. For example, "… biochemistry provided the physical nature of the repressor, an entity postulated in the operon theory" (ibid). The third relation is a *structural relation* and occurs when one field investigates

[20] It is important to note here that although Darden and Maull (1977) use the term 'competing theories', they maintain that in general, fields do not compete and nor do theories from different areas of science.

[21] This remark clearly demonstrates the authors' wanting to distance and distinguish themselves from Nagel and inter-theoretic reduction type models.

[22] Cf. also Toulmin's (1972) term 'discipline', Lakatos's (1970) 'research programme', and Kuhn's (1962) 'paradigm' contrasting with their definition of a 'field'.

[23] See Darden and Maull's (1977) pages 45-46 for a detailed comparison between their concept of a field and Toulmin's (1972) 'discipline'.

[24] Interfield theories, Darden and Maull maintain, best describe the kinds of relations between fields in biology (contra reductive relations) (Darden and Maull 1977, 48).

the structural property of a given entity/process and another field investigates the functional properties of the same entity/process (ibid). For example, physical chemistry provided the structure of molecules and biochemistry provided the function of the molecules – this example constitutes a *structural relation* (ibid). Finally, a *causal relation* occurs when the entities/ processes assumed by one field are the causes of the effects found and studied by a different field. As an example, Darden and Maull cite the theory of allosteric regulation which provides "… a causal explanation of the interaction between the physicochemical structure of certain enzymes and a characteristic biochemical pattern of activity" (ibid). These relations are not mutually exclusive; a part-whole relation can also be a causal relation, a structural relation may also be a causal relation, etc. (ibid).

Darden and Maull note three motivating factors for generating an interfield theory. First, there is motivation to generate an interfield theory when there are relationships between two fields which are already known to exist. They note: "… prior to the proposal of operon theory, the fields of genetics and biochemistry were known to be related; to cite one of many instances, the physical nature of the gene was specified biochemically as DNA. Thus, further relations could be expected between the fields and might lead to the generation of an interfield theory" (ibid). Secondly, one might generate an interfield theory when two fields are investigating different aspects of the same phenomena (ibid).[25] Thirdly, one might generate an interfield theory if there are questions in each field that cannot be answered by their respective resources (i.e., by their methods, tools, techniques, etc.).[26]

To better understand the nature of interfield theories, Darden and Maull draw on examples from the chromosome theory of Mendelian heredity (which bridged the fields of cytology and genetics); the operon theory (which bridged the fields of genetics and biochemistry); and the theory of allosteric regulation (which bridged the fields of biochemistry and physical chemistry) (ibid, 48).

---

[25] This motivation will become clearer in the following paragraph wherein the example of the chromosome theory of Mendelian heredity is described.

[26] Similarly, I elaborate on this motivation in the following paragraph.

For the sake of brevity, I focus on their first example. When cytology emerged as a field in the early 1800's, one of its central problems was to answer *where* within the germ cell 'hereditary material' can be found (ibid, 51). By the 1900's, it was generally accepted that chromosomes answered that question. Theories of heredity emerged in the early 1900's in part due to the discovery of Mendel's law (ibid). Mendel noted the hereditary traits of garden peas and crossed them artificially and proposed a law of inheritance based on his findings (ibid). This one law is not solely responsible for the emergence of genetics, however. Instead, genetics emerged approximately between 1900 and 1905 with the "… independent discovery of Mendel's law by Hugo de Vries and Carl Correns and with the promulgation of Mendel's experimental approach by William Bateson" (ibid).

By 1903, cytology and genetics were investigating *the same phenomena* (i.e., heredity), but were asking importantly *different questions about it*.[27] For example, cytology was concerned with the location of hereditary material within the cell; genetics was concerned with explaining the patterns of inheritance within genes (ibid, 49). However, there were questions within each field that could not be answered from within their respective fields' resources (ibid, 51-52).[28] For example, genetics studied patterns of inheritance and their methods of investigation included artificial breeding; cytology studied the location of hereditary material within cells via methods such as various types of microscopy (ibid, 49-50). Thus, part of the motivation to generate an interfield theory in this instance was because genetics did not have the appropriate set of tools/methods to answer the question of *where* genes were physically located. On the other hand, cytology had the tools and resources required to answer this question and provide the specific physical location of genes inquired by geneticists (ibid, 52).[29]

[27] Recall that this is the second motivations listed by Darden and Maull as a reason to generate an interfield theory (Darden and Maull 1977, 49).

[28] This is the third motivation listed by Darden and Maull as a reason to generate an interfield theory; thus, there were numerous motivating factors around during the early 20th c for interfield theories.

[29] This nicely demonstrates one of the kinds of relations – i.e., the part-whole relation – that can exist between fields. As Darden and Maull (1977) explain: "… the chromosome theory of Mendelian heredity postulated that the Mendelian genes were in or on the chromosomes; cytology provided the physical location of the genes. With more specific knowledge, the theory explained the relation in more detail: the genes are part of (in) the chromosomes" (49).

The chromosome theory of Mendelian heredity is an instance of an interfield theory created to bridge the fields of cytology and genetics in a way that advanced knowledge of heredity. In bridging the two fields, the interfield theory unified knowledge in each field which allowed for deeper explanations such as the similarities between chromosomes and genes (ibid).

According to Darden and Maull (1977), when one shifts the attention away from the traditional view of science whereby philosophers focus on the *relations between theories* – specifically, theories viewed as being the same in kind (i.e., as axiomatic systems) and the relations between them as derivationally reductive – to focus on the *relations between fields*, a multitude of *virtues* emerge (ibid, 60). For one thing, a host of similarities shared between theories may have gone unnoticed with the Received View still in place.[30] For example, they note that their analysis revealed important relations between fields (ibid, 61). That is, there are "… important similarities between the generation and function of relationships which, on the older analysis, are in the different categories of "reductive" and "nonreductive" and would not have been seen as similar" (ibid). Thus, interfield theories can bring into awareness new items for a domain that previously went unmentioned (and would have continued this way without the introduction of an interfield theory). Secondly, interfield theories can solve some theoretical problems by introducing a new idea pertaining to the relations between two fields, such as the idea that chromosomes are the physical location of genes (ibid 52-53, 59). Thirdly, they aid in answering questions that are unanswerable from within one specific field (as was demonstrated in the example of the chromosome theory of Mendelian heredity) (ibid, 59). Fourthly, interfield theories can help predict new items for the domains of each respective fields. Fifthly, interfield theories can generate a *new, collaborative line of research* fields.[31] There are thus numerous virtues of interfield theories. Importantly, interfield theories involve the fitting together of different fields of science in a non-reductive way that enables progress and strong explanations.

---

[30] According to Suppe (1974), the Received View, which was almost universally accepted between the 1920's to the 1950's, held that the structure and nature of theories were the most fundamental topic in philosophy of science, since theories were considered to be a 'vehicle of scientific knowledge' and therefore must be involved in almost every part of science (3). The Received View held that theories were "… axiomatic calculi which are given a partial observational interpretation by means of correspondence rules" (ibid). Further, as Darden and Maull (1977) note, the relationship between theories was understood as derivationally reductive (60).

[31] By focusing on the reduction of theories, one potentially misses out on opportunities for collaborative research projects which can lead to new discoveries and insights.

On Darden and Maull's (1977) view, progress in modern biology occurs when interfield theories bridge together the biological and the physical sciences.[32] On their view, it seems that at least some explanations in the life sciences cannot be complete without interfield theories. Darden and Maull do not explicitly make this claim, however, it follows that if it is the case that there are questions and problems that are unanswerable from within one respective field F1, then F1's explanations and knowledge can only go so far. Eventually, due to conceptual, technical and methodological limitations, explanations and the understanding of some phenomena investigated by F1 will plateau. However, interfield theories have the potential to overcome this plateau by bridging F1 with F2 (which is equipped with the proper resources to help answer some of the questions F1 is unable to).[33]

Indeed, Darden and Maull strongly maintain that interfield theories are crucial for progress: "… the chromosome theory was the first of a series of interfield theories … that advanced our understanding of the relationship between the biological and physical sciences and resulted in *progress* in modern biology" (ibid, 62, emphasis mine). The aspect of fitting together comes into play in the different ways that two fields can be bridged together via an interfield theory. It could be the fitting together of a structure of one field to a function of another field, or an entity of one field with a physical location of another field, etc. Regardless of those details, the fitting together of different aspects from different fields via an interfield theory is important for progress (e.g., new discoveries, new predictions, a new and deepened understanding/explanation of a given phenomenon, etc.).

---

[32] This is what unification looks like on their view. They call this the "… progressive unification (bridging) of the biological and physical sciences (Darden and Maull 1977, 61).

[33] It seems that, up until this point, Bickle might agree. In fact, this general narrative fits with his description of and argument for the importance of transdisciplinary projects and higher-level theorizing (i.e., there are some questions that are simply unanswerable at the cellular/molecular level; one needs to look to different areas of science using different techniques such as single-cell recording in physiology, biological modelling, and FMRI in neuroscience to help get a deeper and complete view of some phenomena). According to Bickle's notion of a transdisciplinary project both high and low-level neuroscience need to fit together for progress (see my footnote 15 on where to find more details on Bickle et al.'s (2000) transdisciplinary project). At least at this point Darden and Maull's and Bickle's views seem compatible. Where they part ways is the details of fitting together where Bickle proposes a strongly reductive view and Darden and Maull a non-reductive view. Regardless of those details, however, a commonality linking the two is that there must be interaction(s) between different fields of science and aspects of those different fields must be fit together in order to have deeper and more complete explanations of some phenomenon.

In sum, the aspect of fitting together thus comes into play in the different ways that two fields can be bridged together via an interfield theory. Interfield theories enable the fitting together of different aspects of different fields. For example, an interfield theory might involve the fitting together of the structure of one field to a function in another field. This fitting together constitutes progress insofar as it results in new discoveries, predictions, and a more complete and epistemically strong explanation about a given phenomenon.

Like Darden and Maull (1977), Mitchell (2002, 2003) positions herself in opposition to reductionist accounts of the relationship between different areas of science and scientific progress. Mitchell rejects both classical inter-theoretic versions of reductionism (e.g., Nagel 1949, 1970) as well as more current, 'new wave' versions of reductionism (e.g., Bickle 1997, 2003; Churchland 1986), and presents an alternative, non-reductive account of the relationship between different areas of science and scientific progress (Michell 2003, 185-186).[34]

Alternatives to reduction tend to be pluralistic and Mitchell's account is no exception (ibid, 186). However, Mitchell maintains that most pluralist accounts are untenable. Mitchell rejects the 'anything goes' pluralism commonly associated with Feyerabend (1975, 1981), stating that condoning any and all 'possibly inconsistent' theories is just as unacceptable as reductionism (ibid). Non-Feyerabendian versions of pluralism tend to appeal to isolationist theories, such as Sherman's (1988) 'levels of analysis' account (hereafter LOA).

Sherman's LOA identified four levels of questions for the biological sciences: the level of evolutionary origin, of current reproductive function, of ontogeny, and of mechanisms (Mitchell 2002, 57).[35] According to LOA, scientists working on a phenomena – e.g., the division of labor in social insects – would investigate it from within their own respective level (which asks its own respective questions) (Mitchell 2003, 187). For example, scientists working in the evolutionary origin level would ask questions concerning why a certain behavior/trait arose at a certain period in time; those working at the mechanistic level would investigate the "… environmental triggers

[34] The biggest reason to reject reductionism stems from the fact that reductionism with its focus on *laws* and *theories* just does not seem to adequately capture the nature and relations of the life sciences (Mitchell 2003). For general criticisms of classical reduction see Darden and Maull (1977), Fodor (1974, 1975), Feyerabend (1962, 1966), Wimsatt (1972), Van Riel (2011), etc. For criticisms of current 'new wave' reductionism, see Schouten and De Jong (1999).

[35] See table 6.1 (page 201) for a visual representation of the LOA analysis (Mitchell 2003).

and hormonal or cognitive mechanisms that then issue in the behaviors", etc. (ibid). According to LOA, levels are *autonomous* and do not (and need not) interact with other levels. Accordingly, competition between explanations occurs solely *within* and not *between* levels; there can be a plurality of competing explanations *only* within a given level (Mitchell 2002, 57; 2003, 187). If there is no competition between levels, and each level is strictly autonomous, there is no need for any kind of collaborative interaction between levels; each scientist would work from within their respective level on their respective questions.

Mitchell draws on Darden and Maull (1977) to argue that this isolationist picture is fundamentally flawed (Mitchell 2003, 187). That is, it conflicts with the history of science which is ripe with examples of fruitful collaborative and interdisciplinary projects between scientists working on questions at different levels (ibid). While Sherman's (1988) account may be representative of how *some* scientists work, it does not represent how *all* scientists work. Mitchell argues that the LOA account of how science works is thus mistaken and that its key mistake lies in its understanding of the epistemological structure of answers (Mitchell 2002, 57). That is:

> … scientific explanations often are causal explanations, identifying the set of conditions that give rise to the phenomenon of interest. At the same time, *complex phenomenon harbor multiple interacting causal processes and multiple levels of organization which all may be involved in the generation of the feature to be explained*. By disambiguating the question to be answered by an explanation … one is still left with a plurality of potential causes acting at a number of levels of organization which may well constitute compatible answers to that single question (ibid, emphasis mine).[36]

Thus, due to the complex and dynamic nature of biological organisms, Mitchell argues pluralism is a better alternative to reductive options. However, that pluralism cannot be a Feyerabendian 'anything goes' pluralism, nor can it endorse isolationist views which conflict with the history of science. Instead, Mitchell (2002, 2003) advances the concept of integrative pluralism to account

[36] Bechtel and Richardson (1993) argue similarly against this kind of isolationist picture claiming that the causal dependence and interaction between different levels implies that narrow work in exclusively one level would be unable to provide an adequate explanation of the phenomenon of interest.

for the types of relations existing between levels which constitute strong explanations and scientific progress.[37]

To better understand Mitchell's concept of 'integrative pluralism', I begin by reviewing the earlier (2002, 2003) articulation of integrative pluralism, and then elaborate on Mitchell and Dietrich's (2006) later version of integrative pluralism. In articulating this concept, I introduce her distinction between competitive and compatible explanations, and elaborate on this via her example of the division of labor of social insect colonies. Afterwards, I explain how the concept of integrative pluralism relates to the concept of fitting together and scientific progress.

Mitchell's integrative pluralism hinges on the distinction between *competitive* and *compatible alternative explanations*. Mitchell notes that most philosophers of science have focused exclusively on the virtues of competitive alternatives (or hypotheses, explanations, etc.) such as Darwinian and Lamarckian theories of inheritance (Mitchell 2002, 56). The argument in favor of competitive pluralism is that competition amongst different alternatives concerning the same phenomena of interest creates an optimal space for testing alternative explanations. According to this view, science progresses when false alternatives are revealed (and accordingly ruled out and rejected) through this optimal testing space.[38] Therefore, a plurality of competing explanations is desirable in that they increase the likelihood of ruling out false explanations, thus narrowing the scope of true explanations (ibid).

Another argument in favor of competitive pluralism is advanced by Beaty (1987) and Kitcher (1991). They argue that "… maintaining multiple, competing theories and explanations is

---

37 Importantly, other than offering Sherman's (1988) description of levels in biology, Mitchell gives no description of what she means by levels. Thus, I operate on the assumption that she accepts Sherman's definition and uses it accordingly. While Sherman's 'levels' are described in the biological sciences, and Mitchell seems to accept this definition, it is unclear whether Mitchell would also accept the traditional picture of levels described by Bickle (see Bickle 2006, page 413 for illustration).

38 C.f. Bacon's ([1620] 1902) notion of crucial experiments whereby Bacon took crucial experiments (hereafter CE) as experiments that could be the neutral arbiter between two competing (rival) theories. This was done by finding a prediction which two (or more) theories disagreed on, then conducting an experiment to see which makes the correct prediction. Theory *A* predicts *x* and theory *B* predicts not-*x*; a CE whose result was *x* would demonstrate theory *A* to be the correct theory. As an example, consider the measurement of the bending of light by the sun. The general theory of relativity and classical mechanics differ by a factor of two. Eddington's experiment demonstrated Einstein's predictions to be correct, and this can thus be considered as a CE (on Bacon's account) which contributed to the early acceptance of relativity.

deemed the rational strategy to adopt for the scientific community as a whole in order to hedge its bets against empirical uncertainty" (Mitchell 2002, 56). While these accounts all advocate (a type of) pluralism, Mitchell argues their conceptualization of pluralism misconstrues how science is and should be done. Mitchell's main disagreement with these accounts is that they all seem to imply the goal of uncovering one true, overarching theory (ibid).[39] That is, these accounts of competitive pluralism all assume "… that pluralism is temporary and strategic, but ultimately eliminable. While this analysis correctly describes *some* of the diversity of models and explanations found in contemporary science, it fails to capture all of it. The remainder is constituted by *compatible, not mutually exclusive, alternatives*" (ibid, 56-57, emphases mine).

Mitchell argues many instances in the biological sciences are constituted by compatible explanations. For example, a social insect colony is compositionally, dynamically, and evolutionarily complex.[40] It is compositionally complex in that it is made up of many, non-random parts (ibid, 58). Dynamically, it hosts a multitude of interacting causes, some of those comprising linear functions, and some comprising non-linear functions (ibid). Evolutionarily, social insect colonies exhibit a diversity of "… historically contingent, adaptive responses to environmental challenges" (ibid). Mitchell considers three different self-organization explanations of the division of labor in social insect colonies: a 'genetic diversity for threshold response' by Page and Mitchell (1991, 1998), a 'foraging for work algorithm' by Tofts and Franks (1992), and a 'learning algorithm by Deneubourg, Pasteels, Fresneau, and Lachaud (1987) (Mitchell 2002, 60-61).

Each of these explanations focuses on *a single causal* factor in the explanation of division of labor in a social insect colony. Instead of viewing each explanation (all located at the ontogenetic level) as *competing*, Mitchell argues they are better understood as *compatible explanations*. This is because each explanation represents an idealized abstraction from the actual, concrete phenomenon (i.e., the division of labor of social insects) (ibid, 63). For

[39] As we shall see, Mitchell's account of integrative pluralism recognizes a plurality of 'true theories' not one all encompassing theory. To foreshadow, instead of one true theory of the division of labor in social insects, there will be one true theory for honeybees, one true theory for leafcutter ants, etc. (Mitchell 2002,67).

[40] She appeals to similar examples such as the origin of an infectious disease responsible for amphibian decline (Mitchell and Dietrich 2006), and Culver (1999) and Koonce and Locci's (1999) example of Lake Erie as instances of compatible, not mutually exclusive explanations for complex and dynamical biological phenomena.

example, Page and Mitchell's (1991, 1998) genetic model focuses *solely* on individual genetic diversity while ignoring any/all other factors (Mitchell 2002, 63).[41] Thus, despite residing at the same level, the three explanations do not actually compete because they each describe only what would occur in 'one non-overlapping world' (ibid, 64).[42]

What kinds of relations hold between compatible explanations at the same level? Mitchell (2002, 2003) believes the answer lies in what she calls integrative pluralism. Integrative pluralism rests on the distinction between theoretical modelling and the application of the model in the explanation of some concrete phenomenon (Mitchell 2002, 66). At the theoretical level of the social insect example, pluralism is justified since it presents three different explanations referring to three different idealized worlds. However, at the application level, integration is necessary, since no matter how complex and how many causes contribute to some biological phenomenon, this phenomenon has only one causal history (ibid).

Mitchell and Dietrich (2006) elaborate on the notion of integrative pluralism by defining it as: "… a view of the diversity of scientific explanations that endorses close study and modeling of different causes and different levels of organization but calls for integration of multiple accounts in the explanation of a concrete phenomenon" (S76). Therefore, if the three self-organization explanations are understood as partial explanations, then a correct theory of the division of labor of social insects would be one that properly *integrated all three partial accounts at the application level*.

However, Mitchell notes, while integration of partial explanations is required for the explanation of a concrete phenomenon, it is unlikely that integration will occur at the theoretical level (Mitchell 2002, 67). This is because the "… diversity of 'solutions' to adaptive problems and the historical contingencies influencing those variable paths that preclude global, theoretical unification" (ibid). For example, ants and bees display similar patterns of division of labor.

---

[41] The other two models similarly focus on one aspect of the division of labor in social insects while ignoring any/all other factors (Mitchell 2002, 64).

[42] C.f. Cartwright (1994, 2000) and Wimsatt (2007) for further discussion of idealized models and the falsity of models in science. While both views of the falsity of models/theories help make sense of compatible scientific models, Mitchell believes that her integrative pluralism offers a better solution for making sense of compatible model examples such as the one described above, and that more can be said on the nature of compatible models in the biological sciences and how they can be integrated (Mitchell 2002, 65).

Despite similarities, what may at first seem like a plausible explanation of the same phenomenon will be dependent on the "… particular features and pathways that occur in each case" (ibid). This, however, does not sanction a Feyerabendian 'anything goes' pluralism, because, as Mitchell notes, there will be only one 'true' integrated explanation for honeybees, and one 'true' integrated explanation for leaf-cutting ants, etc. (ibid). Mitchell concludes by noting that while explanations at the same level can be compatible due to their representing different abstract idealizations, the true nature of explanation requires integration. Importantly, given the multiple causal factors and historical contingencies of organisms in the biological sciences, integration at the application level will be piecemeal. Accordingly, pluralism with respect to compatible explanations "… can and should coexist with integration in the generation of explanations of complex and varied biological phenomena" (ibid).

How does integrative pluralism relate to the notion of fitting together and scientific progress? The answer to this involves the nature of organisms studied in the biological sciences. That is, scientists in the biological sciences must frequently make use of simplified and abstract models; these models usually focus on one feature while ignoring others. Accordingly, these models only account for partial causes. In such cases, Mitchell maintains that a more accurate explanation of a phenomenon will require the integration of "… diverse theories of all the relevant features that uniquely characterize a given phenomenon" (Mitchell 2003, 190). Thus, the integration of partial explanations constitutes the 'hallmark of better science' (ibid). To drive the point home, integration is the hallmark of better science. Good science involves the integration of partial explanations. To make progress into biological sciences – i.e., to advance accurate and complete explanations – involves integration, and the concept of integration is where the notion of fitting together comes into play. That is, in the social insect example, it is the partial causal explanations of self-organization that must be fit together at the concrete application level to produce an epistemically strong explanation of the division of labor in the social insect colony.

In sum, Mitchell (2002, 2003) and Mitchell and Dietrich (2006) advocate an alternative to reduction and non-reductive pluralist accounts of scientific explanations and progress, which focuses on the relations between compatible explanations. They call this integrative pluralism whereby different, yet compatible explanations can co-exist within a level, but must be integrated to provide a complete understanding of a concrete phenomenon.

A more recent alternative to reductionism is advanced by Roskies (2010) in her paper 'Saving Subtraction: A reply to Van Orden and Paap', where she argues that scientific progress occurs when different pieces of evidence in the mind-brain sciences converge on a conclusion via a method she dubs 'functional triangulation. I begin by laying out the motivation behind Roskies' proposal. I then introduce three distinct kinds of functional triangulation and show how Roskies' account relates to the notion of scientific progress and fitting together.

Roskies' main motivation is to demonstrate, contra Van Orden and Paap (1997), that contrastive methods (e.g., classical subtractive methods) commonly used in cognitive neuroscience can and do enhance our understanding of the mind.[43] [44] Roskies believes their dismissive attitude towards neuroimaging is due to their failure to properly contextualize neuroimaging evidence within the broader scope of the mind-brain sciences (Roskies 2010, 640). Far from conceding that individual neuroimaging experiments stand on their own by producing conclusive evidence, Roskies states they must instead be considered "… pieces in a larger puzzle of inference to the best explanation" (ibid). Roskies argues that neuroimaging advances knowledge of the mind-brain when they are involved in a method called functional triangulation. Functional triangulation works when "… information from other task comparisons and other studies are brought to bear on the interpretation of experimental data" (ibid, 640-641). Functional triangulation (hereafter FT) is a method by which claims about the mind-brain are substantiated via a mass task/study comparison when the results of which all converge on a given claim.[45]

A close reading of Roskies' (2010) paper reveals multiple kinds of FT. Despite this, no clear taxonomy is advanced; in what follows, I attempt to do this work by distinguishing two main

---

[43] Roskies explains the subtractive method in the following way: "In subtractive studies, task-related changes in blood flow are identified by subtraction of the image intensity of a blood flow profile due to the performance of one task from that of another. Thus, suppose task A involves processes X, Y, and Z and task B involves the processes X and Y. The blood flow profile generated by performing task A will indicate increased activity in the brain regions which support processes X, Y, and Z, while the profile generated by performing task B will show increased activation in those areas involved in supporting processes X and Y. … if the results of performing task B are subtracted from Task A, the resulting subtraction or 'difference image' will reflect the blood flow differences associated with differences in neural activity between the performance of task A and task B. In this case, the regions subserving process Z will be isolated" (Roskies 2010, 638).

[44] Roskies highlights four key criticisms that Van Orden and Paap (1997) wage against the use of neuroimaging to better understand the mind. I will not be addressing the four criticisms at large (as they do not directly relate to the main goal of my dissertation), but will refer to them only when necessary.

[45] C.f. Campbell and Fiske (1959) and Jick (1979) for earlier articulations of the concept of triangulation.

types of FT – micro and macro-FT. Under micro-FT is what Roskies calls 'within-subject' FT (hereafter WLFT).[46] WLFT employs different comparison tasks within the same lab experiment where they are (ideally) designed to "… enable comparisons between tasks within different combinations of psychological components of interest" (ibid, 641). As an example of how this works in practice, Roskies cites a semantic processing experiment from her own lab (Roskies et al. (2001)) where experimenters used several different comparison tasks – a hard and easy categorization task, a synonym task, and a rhyme task – to hone in more specifically on the brain region(s) involved in semantic processing. Each of these tasks were "… performed on pairs of words, matched for frequency and length, and both the hard and easy categorization tasks used the same words" (ibid). The lab researchers began by comparing the difference in activity in brain regions between the easy and the hard categorization task with the hypothesis that the latter would require more semantic analysis than the former (ibid). This comparison revealed "… a frontal region modulated by task …" (ibid). In order to functionally triangulate from within Roskies' lab/experiment and accrue greater confidence with respect to the nature of semantic processing and its associated brain region(s), additional task comparisons were performed via synonym and rhyme tasks. Data from these tasks were acquired and analyzed to identify whether the initial brain region was active. The experimenters found that the same frontal region identified earlier was active in the synonym task, but not for the rhyme task (ibid). Roskies et al. (2001) concluded that these findings together helped lend support to the initial hypothesis that the frontal region was involved in semantic processing (Roskies 2010, 641).

In this example, WLFT begins with a certain goal (e.g., to identify (a) brain region(s) involved in semantic processing) and a certain hypothesis (e.g., that the frontal region is implicated in

---

[46] It is not entirely clear here what Roskies means by the 'subject' in her 'within-subject' (Roskies 2010, 641). She could mean functional triangulation occurring within one individual test subject (i.e., within one human, one animal, one organism). For example, sometimes two or more different experimental conditions will be explored within the same person, animal, etc. However, Roskies could also mean functional triangulation as it occurs within one particular lab across subjects (since there are typically multiple 'subjects' in a given experiment). In my dissertation, I take her to be mean the latter, and will thus refer to her 'within-subject' FT as 'within-lab' FT (hereafter WLFT). It is important to note that adopting the latter interpretation does not rule out that the former interpretation is *also* required.

semantic processing). That hypothesis gains more support with each additional task comparison that points to the conclusion that the frontal region is implicated in semantic processing (ibid).[47]

While the conclusions reached via WLFT may be more accurate than conclusions from a lab employing only *one or even no task comparison*, WLFT is not sufficient for the stronger kind of claims Roskies wants to make (i.e., that functional triangulation and convergence imply that a claim about the mind-brain stands on firm epistemic footing). As Roskies notes, the results from one lab experiment in isolation represents nothing more than "… evidence for one way of parceling out cognitive function over another …" and alone cannot claim the status of 'conclusive evidence' (ibid). What would be required for 'conclusive evidence'? What demarcates weak from strong evidence? Roskies answers:

> Almost invariably, functional triangulation is also achieved with reference to other imaging experiments, usually those involving other tasks that would seem to involve a like component …. In such cases, the fact that imaging results are reported with reference to a canonical brain space helps the scientist to determine whether the same region as the one she has seen in her experiment is active in other experiments that are hypothesized to involve a similar psychological component. *Convergence across multiple experiments is key to epistemic warrant* when it comes to attributing function to anatomical regions (ibid, emphasis mine).[48]

The type of FT involved in this quotation is importantly different from WLFT (which focuses on comparisons within a single lab experiment); as such, I have chosen to classify it under macro-FT in that it involves the comparison of multiple tasks from *multiple different labs/experiments*.[49] I identify two types of FT under this category: ALFT-1 and ALFT-2. ALFT-1 involves cases where there are multiple task comparisons from multiple different labs/experiments *from within the same area of science* (e.g., the comparison of the results from 26 lab experiments in psychology). ALFT-2 involves cases where there are comparisons from *different*

---

[47] I classify this type of FT under micro-FT insofar as it involves/focuses on multiple task comparisons from within *one individual* lab/experiment.

[48] One important note concerning this quotation. Due to the aim of her paper, Roskies talks about the importance for attributing 'function to anatomical regions' (i.e., structure-function conclusions). However, it is likely that her views on convergence here would extend to cover other (non-structure-function) conclusions borne out of the mind-brain sciences.

[49] Due to this type of FT involving multiple labs/experiments, I will call it across-labs FT and abbreviate it as ALFT.

labs/experiments from *multiple different areas of science* (e.g., the comparison of results from different lab experiments focusing on semantic processing in psychology, molecular biology, physiology, neuroscience, etc.).[50] Roskies states: "… results from psychology, neurology, or neurophysiology are frequently referred to in interpreting imaging data and are also used to further test hypothesized conclusions from studies already performed" (ibid, 642).[51]

FT has different forms, but suffice it to say that for Roskies, FT plays a critical role making headway in understanding the mind-brain (ibid, 642). Roskies' picture of scientific progress is (at least, *prima facie*) compatible with the image of scientific progress endorsed by Darden and Maull (1977), and Mitchell (2002, 2003). That is, Roskies' picture of scientific progress is one of *bootstrapping*, whereby complete and strong epistemic claims about the mind-brain come together *piecemeal*. Roskies explains: "Knowledge derived from other techniques for investigating brain function is combined with insight from psychological theories and imaging studies in order to yield interpretations that cohere with our best theories" (Roskies 2010, 660). Both Darden and Maull (1977) and Mitchell (2002, 2003) – albeit in different ways – emphasize the fact that there are often questions that are unanswerable with the techniques, resources, and methods from within one field that another field has the resources to answer.[52] For Roskies', this is where ALFT-2 comes into play by building an understanding of the mind-brain drawing on the different techniques/methods from different areas of science all trying to understand some aspect of the mind-brain.

Thus, the notion of fitting together in Roskies' account comes into play with the results from either (a) different task comparisons from within a lab being fit together in a way that yields a convergent claim, (b) different task comparisons from multiple different labs/experiments are fit together in a way that yields a convergent claim, or finally (c) that different task comparisons from within different labs/experiments from different areas of science utilizing different methods

[50] It could be argued that ALFT-2 is best suited for establishing well-confirmed conclusions in the mind-brain sciences since it is (or could be) a combination of WLFT and ALFT-1. Roskies herself does not explicitly argue for this.

[51] Here Roskies directs the reader to see Cohen et al. (2008) for an example of this case.

[52] Even Bickle's ruthlessly reductive view concedes this point, which is particularly illuminated with his example of the importance of transdisciplinary projects.

and techniques are fit together in a way that yields a convergent claim.[53] [54] It is in this way that progress and strong explanations in the mind-brain sciences are achieved.

In sum, on Roskies (2010) account, when the results from different lab experiments from different areas of science (ALFT-2) converge on the same/similar result we gain new insights and knowledge about the mind-brain. The fitting together of those results is key to progress because when convergence is achieved, there are good reasons to believe that the claim in question stands on firm epistemic footing.

The last alternative to reduction I consider in this chapter is the concept of 'seamless integration' advanced by Piccinini and Craver (2011).[55] Piccinini and Craver (2011) argue that mechanistic explanations are superior to all other types on offer in the mind-brain sciences in that they offer "… deeper, better confirmed and more useful descriptions …" (ibid, 307). Mechanistic explanations produce reliable claims about the mind-brain by seamlessly integrating results from psychological and neuroscientific research. To flesh out their account, I begin by describing functional analysis and mechanistic explanation. Following this, I cite their example of the neurotransmitter and chemical neurotransmission to introduce their concept of seamless integration. I close by indicating how seamless integration relates to the notions of fitting together and progress.

Piccinini and Craver (2011) note there is consensus that the psychological capacities studied by psychology are explained functionally (ibid, 286). They call this explanation 'functional analysis' which they describe as "… the analysis of a capacity in terms of the functional properties of a system and their organization" (ibid). There are three main types of functional analysis, depending on the kind of functional properties being investigated: *task analysis* (Cummins 1975; Dennett 1978), *functional analysis via internal states* (Fodor 1965, 1968; Stich

---

[53] Importantly, no strong examples of (b) or (c) are provided in Roskies' (2010) work.

[54] One might also note some compatibility between (b) and (c) and what is required for adequate empirical support for the assumptions of the primary science in Nagel's inter-theoretic reduction via his essential condition. Recall that the essential condition requires that the assumptions of the primary field are well supported empirical evidence coming from different fields of study (which presumably use different methods and techniques) which fit together to provide strong empirical support.

[55] Again, there are, of course, other alternatives, but to review every and each reductive and non-reductive account of explanation and progress would be beyond the scope of this dissertation. See footnote 8 of this chapter for other non-reductive accounts not explored in detail here.

1983), and *boxology* (Fodor 1965, 1968) (Piccinini and Craver 2011, 286). Task analysis involves the decomposition of some capacity into its 'subcapacities and their organization'; function analysis via internal states involves analyzing a capacity in terms of its internal states and their interactions; and boxology involves the decomposition of some capacity into a 'set of functionally individuated components' (often represented in diagrams as black boxes), their processes and organization (ibid). Importantly, it is typically believed that to explain some phenomenon properly and fully, functional analysis need only appeal to the functional properties of a given phenomenon.[56]

A mechanistic explanation involves describing "… capacities (functions, behaviors, activities) of a system as a whole in terms of some of *its components*, their properties and capacities (including their functions, behaviors and activities), and the way they are organized together" (ibid, 291, emphasis mine). Components in a mechanism have *both* functional *and* structural properties; an explanation appealing to function with structure, or structure without function would be incomplete, according to Piccinini and Craver (ibid). The functional properties of a component may include "… their activities or manifestations of their causal powers, dispositions or capacities …" and the structural properties of a component may include things like shape, size, location, and the organization of their sub-components (ibid).[57]

Mechanistic explanations require identifying the crucial components of a given mechanism (ibid). These components are sometimes identified in terms of their structural properties; for example, one might use anatomical techniques to parse out the nervous system based on the

---

[56] Piccinini and Craver (2011) do not subscribe to the view that some capacity can be explained solely via analysis of functional properties and processes. This belief hinges on the argument that functional analyses of psychology and mechanistic explanations of neuroscience are not autonomous and distinct. For a full explication of this particular argument, see pages 286-291. For the sake of understanding the concept of integration as integral to producing good explanations, it will suffice to note that Piccinini and Craver hold that functional analysis and mechanistic explanation are not distinct and autonomous, and that their description of mechanistic explanation is "… rich enough to incorporate the kinds of functional properties postulated by functional analysis" (290).

[57] Piccinini and Craver (2011) note that the term 'structure' as used in mechanistic explanations does not entail that structural components are static, or possess one sole function, or are neatly organized and localizable, or stable. Instead, structural components can consist of dynamic and complex feedback relations with other components in the same or different mechanism. Further, on the potential complexity of structural components, they note that "… a structural component might be so distributed and diffuse as to defy tidy structural description, though it no doubt has one if we had the time, knowledge, and patience to formulate it" (291).

various types of neurons found in each structure and how they are connected (ibid).[58] Sometimes components are identified functionally; here, Piccinini and Craver appeal to how the use of fMRI has been used on cases of brain damaged patients to help identify areas of the brain that play a role in certain cognitive functions (ibid).[59] These functional components play a critical role for mechanistic explanation insofar as they help investigators identify the parts of the brain where different kinds of cognitive functions might be found and help allocate functions to different parts of the mechanism (ibid). Importantly, mechanistic explanation, on their account, involves a two-way, interdependent constraint:

> … the presence of certain functional properties within a mechanism constrains the possible structures and configurations that might exhibit those properties. … the presence of certain structures and configurations within a mechanism constrains the possible functions that might be exhibited by those structures and configurations (ibid).

Mechanistic explanations can be complete or incomplete. Piccinini and Craver call incomplete models of mechanisms 'mechanistic sketches' (ibid, 292). Mechanistic sketches contain gaps, questions marks, boxes and arrows, etc., and do not count as a complete mechanistic explanation insofar as critical details have been left out (ibid). Incomplete mechanistic models still have heuristic value, however. That is, mechanistic sketches can provide users with just enough information relevant to the goal/task at hand (e.g., for running an introductory tutorial, for presenting a defense in a courtroom, etc.) (ibid). While mechanistic sketches can be useful, they are nonetheless *elliptical/incomplete* insofar as they fail to provide a full-blown (i.e., consisting of all the necessary details) mechanistic explanation.

The goal of mechanistic explanation is "… to reveal the causal structure of a system. Explanatory models are *evaluated as good or bad* to the extent that they capture, even dimly at times, aspects of that causal structure" (ibid, emphasis mine). Piccinini and Craver argue that defenders of functional analysis are committed to the *very same norms* guiding mechanistic

[58] Piccinini and Craver (2011) cite Brodmann's decomposition of the cortex into various structural regions as an example of identifying components structurally (291).

[59] fMRI (or functional magnetic resonance imaging) picks up on changes in blood oxygenation levels (the BOLD signal) which occur as the result of changes in blood flow and other related factors (Roskies 2010, 637). Typically, changes in the BOLD signal are taken to correlate with neural activity.

explanations *precisely because* the former is an elliptical mechanistic explanation (ibid). Thus, the two types of explanation, on their account, are inextricably linked.

As an example of a mechanistic explanation, they choose the neurotransmitter and the mechanism of chemical neurotransmission (ibid, 304). They list six criteria that can be used to determine whether a molecule is a neurotransmitter (ibid).[60] They use this table to demonstrate how *attributing function to the neurotransmitter involves showing how it fits into the mechanism of chemical neurotransmission*. That is, each criterion is meant to demonstrate that the neurotransmitter is "… organized – spatially, temporally, actively and quantitatively – into the mechanisms of chemical transmission" (ibid, 305). For example:

> *Spatially*, the transmitter has to be located in the presynaptic neuron and contained within vesicles. *Temporally*, if the synapse has to convey signals passed by electrical signals that are hundreds of milliseconds apart, there needs to be some means for removing the transmitter from the cleft after each episode of release. The transmitter has to be *actively integrated* within the mechanisms of synthesis of the presynaptic cell … and with the receptor mechanisms of the post-synaptic cell … And *quantitatively* the molecule's release must be correlated with activation of the pre-synaptic cell, and it must be released in sufficient quantity to affect the post-synaptic cell (ibid, 305-306, emphases mine).

Importantly, many of these criteria (e.g., localization) appeal directly to the structural properties of the synapse; thus, understanding the neurotransmitter mechanistically involves *integrating its functional properties within the structural components of the mechanism of chemical transmission*. Piccinini and Craver argue that none of the properties alluded to in the example of the neurotransmitter are 'independent of the causal details' of the mechanism of which it is a part (ibid). Thus, they reason that this example is a prime illustration of how "… functional language at lower-levels of neuroscience involves implicit commitments to structural facts about the mechanism in question. The same, we claim, extends all the way up to the highest levels of the neuroscience hierarchy" (ibid, 306).

[60] See their table 1. For a visual representation and listing of the criteria for identifying a neurotransmitter (Piccinini and Craver 2011, 304).

It is at this point that they introduce their concept of seamless integration.[61] Piccinini and Craver claim that once the structural properties of some functional analysis are provided, the functional analysis becomes a full-blown mechanistic explanation. It is *this process* of fitting *the structural details* together with *the functional details* of some capacity that entails a *complete mechanistic explanation*: "… functional analyses can be *seamlessly integrated* with mechanistic explanations, and psychology can be seamlessly integrated with neuroscience" and this gives us a unified picture of cognition (308, emphases mine). Integration entails that "… explanatory unification will be achieved through the integration of findings from different areas of neuroscience and psychology into a description of a multilevel mechanism" (ibid, 285).[62]

While they do not explicitly state it, it follows from their arguments and the neurotransmitter/chemical neurotransmission example that it is the fitting together of different areas of the mind-brain sciences (in this example, psychology and neuroscience) that entails progress. That is, *if* it is the case that mechanistic explanation and functional analysis are 'inextricably linked', and *if* it is the case that when both explanations from each respective area of science are fit together they are seamlessly integrated in such a way that entails a complete explanation, *then* it would appear that progress requires integration of this sort.[63] In fact, Craver and Piccinini do claim that mechanistic explanations are the best type of explanation on offer in the mind-brain sciences, and that mechanistic explanations produce "… deeper, better confirmed and more useful descriptions …" (ibid, 307).

In sum, on their account, it is the functional properties of psychology that are being fit together with the structural components of a mechanism that yields epistemically strong claims about the mind-brain.

---

[61] Importantly, their concept of 'seamless integration' is brought up only peripherally at the end of their paper, with little description and elaboration on their concept.

[62] Thus, Craver and Piccinini (2011) too, like all other surveyed alternative non-reductive accounts position themselves against reduction both in its classic, inter-theoretic form, and its more current articulations, such as Nagel (1949, 1961, 1979) and Bickle's (2003, 2006) accounts of explanation and scientific progress.

[63] While they make no mention of Roskies 'convergence' (2010), Mitchell's (2002, 2003) and Mitchell and Dietrich's (2006) 'integrative pluralism', nor Darden and Maull's (1977) 'interfield theories', their (general account) of fitting together seems compatible with both non-reductive accounts of explanation and scientific progress.

1.2 Evaluating the Different Accounts of Fitting Together

There is consensus between both reductionist and non-reductionist accounts alike that claims about the mind-brain stand on firm epistemic footing when they are the product of information from different areas of science fitting together.[64] Whether this fitting together involves one theory being reduced to another, more primary theory (Nagel 1949, 1970), behavior being reduced to low-level neuroscience (Bickle 1998, 2003, 2006), one theory bridging two respective fields of science (Darden and Maull, 1977), or different areas of science being integrated (Mitchell 2002, 2003; Mitchell and Dietrich, 2006; Piccinini and Craver, 2011) or converging (Roskies 2010), the common thread joining all the surveyed accounts is that epistemically strong claims about the mind-brain are the product of interactive, collaborative, and integrative work between different areas of science whereby information from each area is fit together to produce a claim about the mind-brain. Despite the consensus on this assumption about the role of fitting together and scientific progress, what each account lacks is a clear explanation of what it means for information from different areas of science to fit together in such a way that demonstrates, to use Roskies' (2010) term, 'conclusive evidence' for a given claim (641). I will now elaborate on this problem and close by stating its implications for the nature of scientific progress (i.e., why is it important to get a handle on the nature of fitting together).[65]

On Roskies' (2010) account, scientific progress consists of a process whereby different areas of science triangulate (i.e., converge) on a claim; one might ask how convergence is meant to be understood? Similarly, if we consider Nagel's (1949) essential condition, which requires that the claims of the primary science be empirically supported by evidence from different areas of science, one might ask what it means for different areas of science to provide adequate empirical support for a given claim? Must they, to use Roskies terminology, *converge* on the same claim?

[64] Furthermore, it is worth noting all surveyed accounts implicitly embrace some form of scientific realism. While none of the authors explicitly state their stance on the matter, it follows that from the general belief that the fitting together of different areas of science leads to better, deeper, more complete, stronger, and epistemically warranted explanations that constitute scientific progress, that the authors probably all, to some degree or other, embrace some form of scientific realism.

[65] Although each account was surveyed historically (i.e., chronologically) in the first section, I will now consider these accounts in terms of which bear important similarities to the other (e.g., Roskies 2010 and Nagel 1949, 1970; Bickle 2003, 2006 and Darden and Maull 1977; and, Mitchell 2002, 2003, Mitchell and Dietrich 2006, and Piccinini and Craver 2011).

Neither Nagel nor Roskies elaborate on what it means for information from different areas of science to converge/support a claim. Accordingly, I flesh out a couple of possibilities for the notion of convergence/supporting to show how they are problematic, and state why their concepts of fitting together require philosophical analysis.

Let us assume that Nagel's essential condition requires something akin to Roskies' notion of convergence.[66] How is convergence meant to be understood? [67] I identify two possible ways to define convergence. The first involves *a strong identity claim* and the latter involves *the notion of consistency*. The strong version requires that evidence from different tasks (for WLFT), or evidence from different labs from (sometimes) different areas of science (ALFT-1/ ALFT-2) converge on *the exact same claim*. A weaker version of convergence can be described via the notion of consistency, whereby the evidence from different tasks and/or different labs (from different areas of science) all *generally point towards* a claim with room for *minor discrepancies*.[68] However, both characterizations are problematic.

With respect to the strong version, it is impossible for evidence from different lab experiments in different areas of science to converge on the exact same claim. This is because ALFT-2 involves experiments from different lab experiments in different areas of science.[69] Accordingly, if one lab is working on memory on the rat brain using lesioning techniques and another lab is working on memory in the human brain via fMRI, they would not converge on *the exact same claim* in virtue of their investigating memory on different species via different experimental protocols/techniques.

Instead, one might argue that pieces of evidence from different areas of the mind-brain sciences need only be *consistent* with each other – i.e., that they do not contradict each other and there is a general pointing towards a similar but not identical claim— and thus characterize convergence

[66] Since Nagel (1949) provides no elaboration on this point, it is at least a possibility that this could be what was meant.

[67] Again, Roskies (2010) herself provides little to no guidance here.

[68] It is already clear here, however, how terms like 'general pointing towards' and 'minor discrepancies' are vague and require elaboration.

[69] Nagel (1949) too required that the empirical support for the assumptions of the primary science come from evidence from different areas of science.

via the notion of consistency. However, the weak characterization of convergence runs the risk of reducing mind-brain claims to trivial conclusions. By requiring consistency *alone* all we seem to need is *a general compatibility* between the results from different areas of science. This kind of compatibility seems all too easy to achieve, and accordingly, it trivializes claims resulting from ALFT-2.[70] Consistency alone is insufficient for ensuring reliable and accurate claims; that evidence be consistent is one thing, but it is also necessary for evidence to be mutually reinforcing in a way that coincidence, errors, and alternative explanations can be ruled out. Thus, the concept of convergence is thus vague and prompts many clarificatory questions.

There is a lot of compatibility between Darden and Maull (1977) and Bickle's (2003, 2006) accounts. On Darden and Maull's view, progress occurs when interfield theories bridge two fields of science. One of their arguments for the importance of and generation of interfield theories is that there are often questions which are unanswerable from within one field (F1) that can be answered via the techniques, methods, and resources of another field (F2). Bickle (2003, 2006) concedes this point. In fact, this is the very basis for his encouragement of transdisciplinary projects and the importance of higher-level theorizing to help address questions and problems that cannot be addressed via the resources of low-level neuroscience. Thus, on both accounts, information from two (or more) fields need to be fit together to solve a problem/question that one field in isolation cannot answer. This fitting together of information from different areas of science, on both accounts, leads to strong explanations about the mind-brain[71]

While they concede these points, both offer different accounts of fitting together. Darden and Maull (1977) argue the key to complete and epistemically strong claims involves the generation of interfield theories. However, they fail to provide criteria for determining what defines a successful interfield theory; what demarcates a successful from an unsuccessful relation? What would an unsuccessful relation look like? Since there are multiple types of relations that can

[70] And in any case, who decides what entails a general compatibility? How is one to determine where general compatibility begins and ends. Is there a general standard/set of criteria for determining this or is it a 'to each (researcher) their own' scenario?

[71] Of course, there are critical differences between the two; most notably relating to the characterization of the relations as non-reductive (Darden and Maull) and reductive (Bickle). The point of emphasis here is to stress where their accounts are compatible with respect to the notions of fitting together, explanations, and progress.

hold between fields (Darden and Maull identify three distinct types in their (1977) paper 'Interfield Theories')), is there/should there be a hierarchical taxonomy of criteria addressing each type of relation?

Bickle (2003, 2006) argues that for a successful reduction of behavior-to-molecules, higher-level theorizing must first isolate a candidate cellular or molecular process. However, on this account, one might ask what kinds of norms guide this kind of relationship between high-level theorizing and low-level neuroscience? What criteria must high-level theorizing appeal to in order to adequately pick out a candidate cellular or molecular process; what criteria must a behavior-molecules reduction appeal to in order to ensure epistemically strong claims/explanations about the mind-brain? How many studies in low-level neuroscience are required for one to confidently state a ruthless reduction has occurred? How, when, and at what point can we know that a reduction of Bickle's sort has been achieved? Like the concept of convergence, the concept of interfield theories and ruthless reduction are also vague and subject to many clarificatory questions.

Finally, Mitchell (2002, 2003), Mitchell and Dietrich (2006), and Piccinini and Craver (2011) each propose the concept of 'integration' to account for the relations between the mind-brain sciences that constitute strong epistemic claims and scientific progress. Here too, however, similar questions and problems emerge. Specifically, their characterizations of 'integrative pluralism' and 'seamless integration' are vague. On Mitchell (2002, 2003) and Mitchell and Dietrich's (2006) account of integrative pluralism, one might ask what it would mean for different but compatible hypotheses to be integrated together? Mitchell (2002, 2003) appeals to three explanations of self-organization for the division of labor in social insects and argues that these different but compatible hypotheses must be integrated at the level of concrete application. However, Mitchell (2002, 2003) never actually explains *if* and *how* this occurred. What criteria must be satisfied for a proper integration? Must *all* compatible hypotheses be integrated, or only some? If the latter, how are the hypotheses to be integrated identified and chosen from the hypotheses that will be omitted from the integrated explanation; what norms guide this kind of integration of compatible hypotheses?

Piccinini and Craver (2011) argue that when neuroscience provides the structural details of some psychological capacity/process, the two fields are seamlessly integrated resulting in complete and epistemically strong explanations about the mind-brain. They appeal to the example of chemical neurotransmission as an instance of information from two fields being seamlessly integrated, but it remains unclear, from this example, just how exactly information from psychology and neuroscience was integrated. Seamless integration must require more than simply fitting (a) function(s) of some mechanism into its structure(s). How can one discern when a full-blown mechanism has been met? Piccinini and Craver (2011) provide no criteria for discerning *when* and *how* a seamless integration occurs. Admittedly, this is because Piccinini and Craver say nothing on what it means to be seamlessly integrated (only that it has occurred); no definition or elaboration is provided. Thus, the various accounts of integration are also vague and elicit numerous questions and concerns.

None of the surveyed accounts provide a strong definition or an in-depth analysis of what it would mean for information from different areas of science to be fit together, nor do they provide a clear set of criteria for determining when, how, and to what extent different areas of science have been successfully 'reduced', 'integrated', 'converged', etc.

1.3 Concluding Remarks on the Historical Survey of the Accounts of Scientific Progress and Fitting Together

In sum, all surveyed accounts involve the assumption that information from different areas of science must be fit together to yield complete and epistemically strong explanations. However, as I demonstrated in section 1.2, each notion of fitting together is vague and requires philosophical attention. If it is the case that the fitting together of information from different areas of the mind-brain sciences is required for claims about the mind-brain to stand on firm epistemic footing, it is crucial to get clear on the concept of fitting together by asking what fitting together entails, how it unfolds in scientific practice, and whether it is an achievable goal. Given that the notion of fitting together is emphasized across both reductive and non-reductionist accounts of scientific progress as a fundamental requirement for epistemically strong claims, and given that current accounts of fitting together are unclear, it is critical to isolate and analyze this concept.

To address this problem, I introduce the term ‘interdisciplinary corroboration’ in chapter two as a placeholder for the term ‘fitting together’ to capture what I think philosophers of science have been gesturing towards, but which has yet to be clearly articulated and analyzed. From this point, I identity a unit of analysis (i.e., the review paper) that exemplifies the fitting together of evidence from different areas of the mind-brain sciences in practice. I conduct a two-part analysis of a case study of a review paper in chapters three and four, respectively, to get a handle on how review papers work to corroborate claims about the mind-brain by drawing on evidence from different areas of the mind-brain sciences.

Chapter 2

2 Introduction

In chapter 1, I noted the current and popular view of the relationship between the different areas of the mind-brain sciences is integrationist, interdisciplinary and non-autonomous. A key assumption behind this view is the thought that evidence from different areas of the mind-brain sciences must be fit together in a certain way to yield complete and epistemically strong claims about the mind-brain. I surveyed some reductionist and non-reductionist accounts of scientific progress emphasizing the notion of fitting together, and closed by highlighting key problems held by each account. The overall problem I identified is that each concept of fitting together is vague and under-analyzed; since this concept is emphasized across the board by both reductionist and non-reductionist accounts of scientific progress for strong epistemic claims about the mind-brain, it is crucial to isolate and analyze this concept.

The aim of this chapter is to introduce various units of analysis that exemplify the fitting together of findings from different areas of the mind-brain sciences, and to argue that of those units of analysis, the *Trends in Cognitive Sciences* review paper is a good candidate to start investigating the nature of fitting together and its relation to strong epistemic claims about the mind-brain. I begin, in section 2.1, by introducing the term 'interdisciplinary corroboration' as a general placeholder for the notion of fitting together emphasized in the different accounts of scientific progress I surveyed in chapter one. Critically, 'interdisciplinary corroboration' is introduced as a working definition, and accordingly, requires elaboration and analysis. I argue a good starting point for investigating the nature of 'interdisciplinary corroboration' is to analyze review papers as these units of analysis survey literature on a given topic and often demonstrate how different findings fit together to support a claim about the mind-brain. To get a better understanding of the epistemic role of review papers in relation to 'interdisciplinary corroboration' and the mind-brain sciences, I begin by sketching the motivation, history, and development of review papers in section 2.2. In section 2.3, I advance a taxonomy for the kinds of review papers found in mind-brain science journals. In section 2.4, I note the epistemic advantages review papers boast; this sets the stage for section 2.5, wherein I argue that the *Trends in Cognitive Sciences* review paper is a good starting part for analyzing 'interdisciplinary corroboration'.

To better understand the notion of ‘interdisciplinary corroboration’ via a *Trends in Cognitive Sciences* review paper, I conduct a two-part analysis in chapters three and four, with the aim of extracting key lessons for the notion of ‘interdisciplinary corroboration’ and its relation to progress in the mind-brain sciences.

2.1 Interdisciplinary Corroboration: A Working Definition

In chapter 1, I identified a common assumption linking both reductive and non-reductive accounts of explanation and progress – i.e., that different areas of science must be fit together to generate strong epistemic claims about the mind-brain (or an aspect of the mind-brain, such as memory). The problem I identified with this assumption is that the notion of fitting together has been largely under-analyzed; as such, it may turn out that this fundamental assumption is flawed (or that it requires further philosophical attention). To address this problem, I introduce the term ‘interdisciplinary corroboration’ (hereafter IC) as a general placeholder for what I believe philosophers of science have been gesturing towards but have yet to articulate and analyze in detail. It is from this point that I will begin to analyze what IC means and what its implications are for scientific progress and strong explanations about the mind-brain.

My working definition of IC is that scientific findings corroborate claims about the mind-brain insofar as they are interdisciplinary, mutually supportive, and all point towards a general claim. The accounts of scientific progress I surveyed in chapter one all maintained that strong epistemic claims about the mind-brain required evidence from different areas of the mind-brain sciences to fit together. Implicit in their accounts is an argument stipulating claims about the mind-brain are *corroborated* by *the fitting together* of results from *different areas of the mind-brain sciences*. For this reason, I introduce IC as a general placeholder for the various concepts of fitting together presented in chapter one to facilitate my analysis of how interdisciplinary results fit together to corroborate claims about the mind-brain.

There are two key components comprising IC: an interdisciplinary component, and a corroboration component. The interdisciplinarity component relates to the different areas comprising the mind-brain sciences (e.g., computer science and AI, biology, neuroscience, psychology, genetics, etc.) and their respective findings on a given topic. The corroboration component relates to how evidence confirms or lends support to a given claim about the mind-brain.

How do interdisciplinarity and corroboration connect? Bechtel and Richardson (1993) discuss corroboration in the context of the necessity of independent measures:

> If we have available a number of independent but uncertain pieces of evidence, all of which point to the same conclusion, then that conclusion may in fact be *extremely probable*, even if each piece of evidence is uncertain. By contrast, even a highly reliable procedure will have some likelihood of failure. It is a simple exercise in probability to see that employing multiple independent modes of evidence is superior to relying on a chain of inferences from one piece of evidence. Afterall, the likelihood of getting an error in the latter case is the sum of the likelihood of errors at each step. Think of adding a long series of numbers. At each step there is a small chance of error, but the overall likelihood of error can be quite high. By contrast, with independent but unreliable pieces of evidence, each piece of evidence may have a relatively high likelihood of error, but the overall likelihood will tend to be low (xxvii, emphasis mine).[72]

Importantly, they note that for this to be true, it is critical for different pieces of evidence to be the product of independent measures. This is where the interdisciplinarity component comes into play. As noted in chapter one of this dissertation and above, the mind-brain sciences comprise multiple areas of science including but not limited to: psychology, cognitive science, cognitive neuroscience, biology, cellular and molecular neurobiology, computer science and artificial intelligence, genetics, etc. Reflected in this diversity of areas investigating the mind-brain are numerous different methods, tools, techniques, etc. These methods include but are no means limited to the following: questionnaires and surveys, imaging techniques such as fMRI or positron emission tomography (PET)), electrophysiological recordings, gene knockouts or genome editing, computational modelling, etc. Accordingly, it is likely that if there are multiple sources of evidence from different areas of the mind-brain sciences, which make use of many

[72] This is similar to the argument Roskies (2010) makes with her insistence on the necessity of functional triangulation for convergent (i.e., strong) explanations and scientific progress.

different methods, experimental protocols, and tools/techniques of analysis, then one can argue the pieces of evidence are the product of independent measures.

Bechtel and Richardson (1993) argue if multiple sources point to the same conclusion, and especially if those multiple sources employ independent measures, it is very likely that a conclusion has been corroborated (in the sense that it is well-confirmed/supported/highly likely).[73] [74] This is what I have in mind for IC: evidence from multiple areas of the mind-brain sciences (using different tasks, experimental protocols, methods, etc.) corroborating claim *x* insofar as that evidence all points towards *x* and that evidence is mutually supportive.

My working definition of IC leads to the question: how does IC function in the mind-brain sciences, and to what extent does it result in epistemically strong claims about the mind-brain? In answering this, it is helpful to look for instances of IC in practice; where are such instances found? [75]

Instances of IC are typically not found in the day-to-day science conducted in individual labs; in these cases, the goal is to answer very specific research questions from within a specific area of science. IC involves interdisciplinarity, and, as such, given that most day-to-day lab work is

---

[73] Wimsatt (2007) makes a similar argument for his concept of robustness. Building on Campbell's (1966) concept of triangulation in the social sciences, Wimsatt argues that things are robust to the extent that, "… *they are accessible (detectable, measurable, derivable, definable, producible, or the like) in a variety of independent ways*" (Wimsatt 2007, 196). Like Bechtel and Richardson (1993), the emphasis on independent measures is critical. Wimsatt notes that we can be more confident in entities, processes, relations, mechanisms, etc. that are the result of multiple independent measures because "… the chance that we could be simultaneously wrong in each of these ways declines with the number of independent checks we have" (Wimsatt 2007, 196).

[74] Hacking (1983) also makes a similar argument in the context of entity realism. Hacking argues that entity realism (specifically unobservable entities) is warranted if an entity is detected by multiple different and independent means of detection; the more independent means of detection; the more warranted said entity becomes. Hacking refers to the example of detecting the dense bodies of red blood platelets. In this case, "Two physical processes – electron transmission and fluorescent re-emission – are used to detect the bodies. These processes have virtually nothing in common between them. They are essentially unrelated chunks of physics. It would be a preposterous coincidence if, time and again, two completely different physical processes produced identical visual configurations which were, however, artifacts of the physical processes rather than real structures in the cell" (Hacking 1983, 201).

[75] In emphasizing the importance to practice I am aligning myself with the new mechanists, whom, towards the end of the 20th century, emphasized the importance of looking at the history of science and scientific practice to inform our theories. See Bechtel and Richardson (1993) for a brief historical summary of the shift from logical positivism's focus on scientific justification and the logical structure of explanations, to the focus on history with Kuhn's (1970) *Structure of Scientific Revolutions*, and then to the new mechanist focus on practice, history, and discovery in the late 1980's (xxi).

conducted within one specific area of science, it is not an ideal starting point for examining IC in practice.

The scientific research papers emerging out of those labs seem like a better option insofar as these papers often try to situate their findings within the broader context of the literature, and in doing so, may cite other studies to lend support to the claim(s) they advance. The problem with these papers is that the primary goal is to present original research, and, accordingly, the scope of the corroborative work is narrow. Ideally, one would want the scope of corroboration to be much wider and for the main goal of the paper to demonstrate how interdisciplinary findings on a given topic fit together to support a given claim about the mind-brain.

Instead, a unit of analysis with great promise for analyzing IC is various types of review papers. Interestingly, Roskies (2010) claims we can find instances of convergence in meta-analyses/review papers. Arguing against Van Orden and Paap's (1997) view that neuroimaging fails to find convergent results, Roskies states:

> Convergence across multiple experiments is the key to epistemic warrant when it comes to attributing function to anatomical regions. Results from any single study are viewed by scientists as providing evidence for one way of parceling out cognitive function over another, but not as conclusive evidence. Van Orden and Paap seem to disregard the importance of functional triangulation in part because they claim that experiments meant to triangulate function fail to find convergent regions of activity. *The majority of the literature belies this interpretation as recent meta-analyses have shown strong convergence in imaging results across tasks, across laboratories, and both across and within subjects* (Roskies 2010, 641, emphasis mine).

Roskies cites review papers by Krain et al. (2006), Miller et al. (2002), Steele and Lawrie (2004), and Van Overwalle (2009), as examples of literature demonstrating convergence (Roskies 2010, 624). However, Roskies does not provide an in-depth analysis of these units of analysis, nor their implications for convergence and scientific progress.

Briefly, as an introductory description, authors of review papers summarize literature on a given topic. Some review papers consist primarily of a literature summary and contain no arguments.

Critically for my dissertation, however, many review papers provide more than a summary; they also contain arguments – sometimes implicit, sometimes explicit. These arguments support claims about the mind-brain based on the author's summarized findings. These arguments often involve the author demonstrating how findings from different areas of science combine to corroborate a claim about the mind-brain. Sometimes these arguments are the product of meta-analysis, sometimes they are the result of a qualitative analysis. Thus, authors of review papers appear to do the work required by IC; that is, they demonstrate how findings from different areas of the mind-brain sciences (interdisciplinarity) fit together in a way that lends support to/confirms (corroboration) claims about the mind-brain. Therefore, review papers provide a strong starting point for beginning to investigate the nature of IC by analyzing how it unfolds in practice.[76]

Other philosophers of science have stated that progress in science is made when the findings from different areas of science fit together to support/confirm a claim about the mind-brain.[77] IC is a working definition I introduce to get a handle on the phenomena of fitting together. Review papers are a good starting point for this project insofar as they involve both the interdisciplinarity and corroboration aspects, and, accordingly, appear to produce epistemically strong claims about the mind-brain. In chapters three and four, I conduct an in-depth analysis of a case study review paper. Importantly, because I begin with a working definition of IC in this chapter, it is likely that a more complex and nuanced picture of IC will emerge post-analysis. Before conducting my analysis of a review paper, however, I begin by explicating what a review paper is, what elements it is composed of, what standards it must adhere to, and arguments for the epistemic superiority of claims emerging out of review papers.

2.2 Motivation, History, and Development of Review Papers

If one is interested in IC, and the ideal outcome of IC is knowledge production – specifically the production of epistemically strong claims about the mind-brain – in the mind-brain sciences, and if review papers are strong candidates for how IC works in practice, then one would imagine there are certain standards review papers must adhere to. Part of my goal in this chapter is to

[76] That is not to say that review papers represent the *only way* in which areas of science are demonstrated to corroborate a claim about the mind, only that it is a strong candidate to begin analyzing IC in practice.

[77] See my chapter one for a review of this literature.

explicate what a review paper is, when and how it developed in the sciences, and to elucidate what criteria are used to produce epistemically strong claims.

A few precursory notes before I begin. There are no specific organizations or lists of criteria for writing review papers in the mind-brain sciences.[78] For this reason, when I discuss the history, development, and criteria for writing review papers, I do so in the context of the social sciences and the health sciences. Accordingly, I reference different organizations (e.g., the Cochrane Collaboration) and standardized checklists for writing review papers (e.g., the Preferred Reporting Items for Systematic Reviews and Meta-Analysis (hereafter PRISMA)) within those areas of science. Critically, the standards for Cochrane Collaboration and PRISMA review papers are not representative of review papers found in mind-brain science journals. This is because review papers in the mind-brain sciences involve the synthesis of information from the different areas of science housed under the umbrella of the mind-brain sciences (e.g., psychology, biology, computer science and AI, genetics, neuroscience, etc.). These areas of science make use of different methods, experimental protocols, and tools and techniques of analysis, and, accordingly, the standards advanced by organizations such as the Cochrane Collaboration in the health sciences, for example, will not be applicable to review papers in the mind-brain sciences.[79] As such, I introduce and discuss the history, development, and criteria for writing review papers in the context of the health and social sciences while referring to organizations such as the Cochrane Collaboration for comparative reasons.

---

[78] I will return to this point later in this section.

[79] That is, review papers in the health sciences frequently analyze and synthesize information using similar methods, tools, and techniques of analysis. For example, a Cochrane review by Mehrholz, Thomas, and Elsner (2017) looked at treadmill training and body weight support interventions for humans post-stroke. Critically, this review looked at two main types of interventions – i.e., treadmill training and body weight support – and focused on human adults post-stroke. Of the 3105 human participants, the majority were approximately 60 years old and had previously had a stroke; thus, there was little variation in the type of species being considered in the review paper. Further, the methods used were all similar; all involved a combination of treadmill training with body weight support, treadmill training without body weight support, and treadmill training with some other physiotherapy support. Review papers in the mind-brain sciences, however, involve looking at information from different areas of the mind-brain sciences; consequently, the information being synthesized comes from research conducted on a wide variety of human and non-human animals, using very different and distinct methods, experimental protocols, techniques and tools of analysis. Consequently, the standards for review papers in the health sciences/social sciences are not applicable to the types of reviews found in mind-brain sciences journals.

Generally speaking, a review paper provides a condensed summary of the current state of research in a specific field (e.g., in literature, anthropology, cognitive neuroscience, etc.) in relation to a specific topic within that field (e.g., exile in Prussian literature, kinship in anthropology, episodic memory in cognitive neuroscience, etc.), and this summary will include multiple citations from relevant secondary sources. Review papers typically include some or all of the following features (depending on the kind of review paper they are classified as[80]): a *summary*, an *explanation*, a *synthesis*, and a *direction(s) for future research*. The first component involves summarizing recent work relative to some question within a field of research, and, as such, it provides an avenue for the reader (both the expert, but particularly the non-expert) to develop a general understanding of the topic and where things stand. In writing this summary, authors of review papers cite relevant published papers. Authors of the review papers might also explain how the various findings from their cited studies shape and impact the question or topic they are writing on, and they might also provide a synthesis of those findings such that a given claim is advanced as substantiated by the cited literature. Finally, authors might suggest, based on their summary and/or their analyses (via explanation and/or synthesis), directions for future research. This is a rough and ready sketch of what a review paper is. I will now contextualize review papers within the history and development of research synthesis, and then contextualize them further within the mind-brain sciences.

Gurevitch et al. (2018) trace the beginning of research synthesis with the birth of meta-analysis, which they date around the early 20th C. They define research synthesis as: "… synthesizing results across studies to reach an overall understanding of a problem and to identify sources of variation outcomes" (Gurevitch et al. 2018, 175). This began when Pearson attempted to combine knowledge from different studies to determine whether vaccination was an effective preventative measure against typhoid in soldiers in 1905 (ibid, 176). In 1925, Fisher developed a method for combining probabilities over multiple studies (ibid). The development of meta-analysis continued into the mid-20th century with the attempt to formalize Fisher's method via fixed and random-effects models by Yates and Cochran (1938) and Cochran (1954) (Gurevitch et al. 2018, 176). However, it was not until 1977 that meta-analysis became acknowledged as a *significant and invaluable tool for scientific research*. That is "… it was not until the insight of

80 I will elaborate on this classification below in section 2.3.

psychologists G. Glass and M. Smith in 1977 – that outcome measures from different experiments could be standardized and put on the same scale – that meta-analysis began to affect scientific research to a large extent" (ibid).

Approximately nine decades following Pearson's attempt to combine results across multiple studies, Chalmers and Altman introduced the term 'systematic review' in their 1994 foreword to their 1995 book, *Systematic Reviews*. In this book, they argue for the necessity of review papers in light of the "… ever-increasing tidal waves of new research evidence" (Chalmers and Altman 1994, foreword).[81] Specifically, they argue that what they call 'the systematic review' is best suited to help researchers navigate the 'tidal wave' of new evidence pouring in from multiple sources in the health sciences (ibid). They define the systematic review as a review which has "… been prepared using some kind of *systematic approach* to *minimizing biases and random errors*, and that the components of the approach will *be documented in a materials and methods section*" (ibid, emphases mine).[82] *Systematic Reviews* (1995) acts as an introductory handbook for understanding and writing systematic reviews by advising potential authors on how to pick relevant studies for a systematic review, how to identify problems within meta-analyses, provides checklists for scientifically sound systematic review papers, and guidelines on how to report, update, and correct systematic reviews.

Chalmers and Altman's seminal work *Systematic Reviews* (1995) changed the course of research synthesis (both with respect to how it was defined and conducted). Prior to 1995, research synthesis took the form of what Gurevitch et al. (2018) call the 'narrative review'. Gurevitch et al. (2018) provide no definition of a narrative review, claiming only that the main aim of a narrative review is to 'summarize results across studies" (175). The main problem they identify with narrative review papers for research synthesis is that it has become increasingly difficult to summarize hundreds of studies on a given problem in an *un-biased way* (ibid). This, they speculate, may account for why narrative reviews haven fallen out of favor and why they have been replaced by systematic reviews, which are allegedly more equipped to deal with large

---

[81] Importantly, this book is written in the context of the health sciences.

[82] They distinguish the systematic review from meta-analyses in terms of qualitative and quantitative analyses (Chalmers and Altman 1994, foreword). Systematic reviews do not necessarily need to incorporate a meta-analysis (i.e., a quantitative synthesis of data); instead, they may present a qualitative synthesis (this is especially popular in the social sciences).

volumes of research and problems of bias and random errors (ibid). Gurevitch et al. (2018) note that despite having fallen out of favor, narrative reviews are still useful for providing a summary of the development and progress of a given intervention, problem, or field of research; the problem is that they cannot adequately 'summarize results across studies' (ibid). Importantly, Gurevitch et al. do not explain what they mean by 'summarize results across studies' (ibid). On my reading, when they claim that narrative review papers cannot accurately summarize results across studies I take it they mean narrative review papers cannot *synthesize* results across studies. Accordingly, one of the key differentiating factors between narrative and systematic review papers is that the primary component of the former is *summary*, and the latter is *synthesis*. If this is the case, then narrative reviews simply list out/or summarize findings from different studies, but offer no explanations, arguments, or comments on their findings and how they shape the topic at hand, nor do they provide a synthesis of those findings. Instead, that kind of work is performed by authors of systematic reviews.

After 1995, systematic reviews were the main form of research synthesis in the medical/health and social sciences. Gurevitch et al. define systematic reviews as papers which:

> … aim to provide *a robust overview* of the efficacy of an intervention, or of a problem or field of research. They can be *combined with quantitative meta-analyses* to *assess the magnitude of the outcome across relevant primary studies* and to *analyze the causes of variation among study outcomes* (effect sizes)" (ibid, emphases mine).

Systematic reviews boast numerous advantages and are revered in the literature on research synthesis as being a critical tool for scientific research and scientific progress (ibid).

One particular feature – i.e., the reduction of bias – is held by many in the research synthesis literature (e.g., Chalmers and Altman 1995, Gurevitch et al. 2018, Mulrow 1995, Shamseer et al. 2015, etc.) as being one of the standout features that differentiates systematic reviews from traditional, narrative reviews, contributing to epistemically strong claims. None of these authors advance a full-fledged argument for the reduction of bias via the systematic review. However, on my reading, their claim hinges on what I identify as two key assumptions made in the literature on research synthesis. The first involves systematic reviews relying on an explicit set of methods and standardized guidelines for writing review papers, and the second involves the relationship between systematic reviews and open science; I will now review each assumption in turn.

The first assumption I identify is because systematic reviews rely on 'explicit methods' and 'formal methodological guidelines' via organizations such as the Cochrane Collaboration, the Campbell Collaboration, and PRISMA, the claims emerging out of systematic reviews stand on firmer epistemic footing than claims emerging out of papers which do not adhere to any standardized guidelines (Gurevitch et al. 2018, 175; Mulrow 1995, 1).[83] That is, these organizations produce checklists and guidelines for:

> … literature search, study screening (including critical appraisal of eligible studies according to *pre-defined criteria*), data extraction, coding and often statistical analysis (that is, meta-analysis), along with detailed, transparent documentation of each step. Software, protocols and reporting guidelines for systematic reviews and meta-analyses are well established in many fields; for example, PRISMA … (ibid, emphasis mine).

The key assumption I identity for the reason systematic reviews deliver unbiased claims hinges on this notion of 'pre-defined criteria' (ibid). That is, because, (ideally), all authors are being held to *the same standards* with respect to how to write, publish, and update systematic reviews, authors are less likely to let their own biases infiltrate their analyses. An author writing a systematic review without consulting well-documented and formalized guidelines would presumably be choosing their own methods for writing a review paper. Accordingly, the author in question would be choosing their own selection criteria, method(s) of data documentation and interpretation, etc., increasing the likelihood of the author's personal biases infiltrating each step of the review writing process. This would arguably make the claims emerging from said review less reliable.[84] [85] [86] [87]

[83] I will elaborate on these organizations and their guidelines shortly.

[84] C.f. also Oxman and Guyatt (1988) for the elimination/reduction of bias in systematic reviews.

[85] Again, this is my interpretation of what I think Gurevitch et al. (2018) are trying to get at in their claim that systematic reviews eliminate bias as they never explicitly back up and elaborate on this claim.

[86] These claims are, of course, not without their criticisms. See Eysenck (1995); Egger et al. (1997); Ioannidis (2016); Gurevitch et al. (2018), etc. for some challenges and limitations pertaining to systematic reviews, meta-analysis, and their ability to reduce bias in the health sciences.

[87] See Sullivan (2016) for a critique of meta-analysis in the mind-brain sciences.

How did organizations such as the Cochrane Collaboration emerge and develop in the history of research synthesis? The late 20th and early 21st century saw substantial procedural and methodological growth, with much of that growth being attributed to multi-disciplinary collaboration organizations geared towards the advancement of research synthesis for the sciences. Of notable mention for multi-disciplinary collaborations are the Cochrane Collaboration founded in 1993, the Campbell Collaboration created in 1999, the Society for Research Synthesis Methodology founded in 2005, and PRISMA (Gurevitch et al. 2018, 175). These organizations effectively "… oversee systematic reviews in the *medical and social sciences*, respectively, bringing practioners and methodologists together and setting standards for research synthesis publications and evidence-based guidelines for practice and policy" (ibid, 176, emphasis mine). To re-iterate a point made above, it seems for Gurevitch et al. (2018) and Mulrow (1995), systematic reviews reduce bias in scientific claims precisely because they follow the strict and well-documented methods/guidelines advanced by these organizations.

As an example of guidelines for writing sound systematic review papers, meta-analyses, or systematic reviews with a meta-analysis, consider PRISMA's 27-item checklist.[88] PRISMA provides potential authors with a minimum set of items required for writing systematic reviews (with and without meta-analysis) (http://www.prisma-statement.org/). PRISMA "… focuses on the reporting of reviews evaluating randomized trials but can also be used as a basis for reporting systematic reviews of other types of research, particularly evaluations of interventions" (ibid). The PRISMA checklist includes instructions for sections of review papers including the title, abstract, introduction, methods, results, discussion, and funding (ibid). Each of these sections contains sub-topics which are elaborated on; this elaboration provides authors with instructions on what they need to do to ensure a scientifically sound review paper. For example, under the section 'methods' in the sub-topic 'study selection', PRISMA states all authors must "… state the process for selecting studies (i.e., screening, eligibility, included in systematic review, and , if applicable, in meta-analysis" (http://prisma-statement.org/documents/PRISMA%202009%20checklist.pdf, 1). Under sub-section 'risk of

[88] Since PRISMA's checklist pertains to systematic reviews, meta-analysis, or systematic review with a meta-analysis, I will shorten this to 'review paper' in the remainder of this section. Unless specified otherwise, I use the term in this section of the chapter to encompass any of those types of reviews.

bias in individual studies', PRISMA states authors must "Describe methods used for assessing risk of bias in individual studies (including specification of whether this was done at the study or outcome level), and how this information is to be used in any data synthesis" (ibid). Under the sub-section 'synthesis of results', PRISMA requires that all authors "Describe the methods of handling data and combining results of studies, if done, including measures of consistency … for each meta-analysis (ibid). [89] Similar, but more elaborate guidelines are provided in the Cochrane handbook for the health/medical sciences and the Campbell Collaboration for the social sciences.[90]

PRISMA's checklist and the standards provided by organizations such as the Campbell collaboration deliver strict methods and guidelines which, on my interpretation, Mulrow (1995) and Gurevitch et al. (2018) hold responsible for the reduction of bias in systematic review papers.[91] That is, following these guidelines allows for the un-biased assessment of a question, problem, or debate, resulting in reliable claims because authors are *all held to the same standards for researching, reporting, writing, and updating review papers*. Importantly, I interpret both Gurevitch et al. (2018) and Mulrow's (1995) claims as *comparative*. That is, they are not claiming that following these guidelines entails the end-product of a systematic review will be *wholly unbiased*, but that *in comparison to* traditional reviews (i.e., narrative reviews), they are significantly *less biased*.

The second assumption I identify with respect to the arguments in the literature on research synthesis for the reduction of bias in systematic reviews involves the relationship between open-science and systematic reviews. Open science is the belief that the dissemination of scientific research ought to be transparent and openly accessible to the public (including both experts and non-experts). Gurevitch et al. (2018) discuss the nature and virtues of open-science practices:

---

[89] See http://prisma-statement.org/prismastatement/Checklist.aspx for the full PDF version of the checklist.

[90] Please see file:///C:/Users/Jackie/Downloads/Campbell_Policies_and_Guidelines_v4.pdf for guidelines by the Campbell Collaboration "Campbell Systematic Reviews: Policies and Guidelines" and http://handbook-5-1.cochrane.org/ for the Cochrane Handbook for Systematic Reviews of Interventions.

[91] Their claim thus importantly hinges on most/all journals/authors appealing to and abiding by these guidelines and standards. That is, if all are held to the same set of standards, then the author's personal biases are less likely to infiltrate the systematic review, thus resulting in less biased conclusions.

> Open-science practices have emphasized full and unbiased access to scientific data, which is of longstanding importance and central to future progress in meta-analysis. Pre-registration (called 'registration' in some fields) of planned studies can reduce selective reporting of outcomes; publication of 'registration reports' in which the methods and proposed analyses for a study are peer-reviewed and published before the research is conducted and can reduce publication bias (ibid, 180).

Gurevitch et al. (2018) say little more on how open science contributes to the reduction of bias. Accordingly, I draw on work by Nosek et al. (2015) to draw out the assumption I believe is behind Gurevitch et al.'s claims (ibid).

Nosek et al. (2015) elaborate on the nature of open-science practices by advancing eight standards they deem critical for openness and transparency in the scientific community.[92] For example, Nosek et al. (2015) note that design standards allow for an "… increase in transparency about the research process and reduce vague or incomplete reporting of the methodology" (1423). Complete transparency and proper documentation of one's design protocols allows others to better evaluate the claims emerging out of review papers. If one's design and methods protocols were vague, incomplete, and/or absent, this may lead to an inaccurate assessment of review paper's main claim. Accordingly, the epistemic status of that paper's main claim would stand on uncertain epistemic footing. That is, to properly assess whether the claims emerging out of a review paper stand on firm epistemic footing, one needs to know that the methods used to investigate the topic in question are sound/reliable. It is difficult to make that assessment if the protocols are not explicitly outlined in the review paper; therefore, increasing transparency and documentation standards for review papers may help alleviate this problem.

Another standard crucial for openness and transparency involves the *preregistration of studies* and *the preregistration of plans for the analysis*. Gurevitch et al. (2018) also stress the importance of this standard in relation to open-science practices. Here, they note preregistration of studies and analysis plans are peer-reviewed and published *prior to* the writing of the actual review. This reduces bias insofar as the author must adhere to those published plans, even if the

[92] I only address a few of their standards here for the sake of brevity. For a full list, see their diagram of standards and levels on page 1424 (Nosek et al. 2015).

resulting claim disconfirms their initial hypothesis (180). Nosek et al. (2015) add that the preregistration of studies makes the process of finding research (published or unpublished) much easier by making sure the study is recorded in a public registry, and that the preregistration of analysis plans "… certify the distinction between confirmatory and exploratory research …" (1423).

In addition to providing eight general standards, Nosek et al. propose levels from 0-3 where levels 0-3 represent increasingly strict standards where level 0 is offered for comparative reasons and level 0 does not meet any standard, and level 3 outlines the highest standards for open science (ibid, 1423-1424). Ideally, all review papers would fall under level 3 – i.e., they would adhere to strict standards to increase transparency and reduce bias in the review paper.[93] The assumption that systematic reviews reduce bias in virtue of their relationship to open-science practices is that because guidelines such as the PRISMA checklist are accessible to the public, and because the emphasis on transparency requires authors to document all steps of the review process in a public forum (e.g., to publish the plans for the analysis prior to conducting the review), the claims emerging out of systematic review papers are less likely to be biased compared to narrative review papers.

The relationship between open-science practices and systematic reviews that Gurevitch et al. (2018) point towards is one where instead of the traditional reviewer writing within the narrow confines of their own work space, with little to no interaction with other specialists, choosing their selection criteria and interpretation of the data as they see fit, systematic reviews (with or without meta-analysis) are written in a space that breaks the confines of the traditional workspace in virtue of the author consulting organizations that publish lists of standardized protocols and guidelines that their work must adhere to. Ideally, all authors of systematic review papers would adhere to the guidelines set by either the Cochrane Collaboration, Campbell Collaboration, and/or PRISMA to ensure that the claims emerging out of them stand on firm epistemic footing. It is thus the openness of science and the adherence to strict and well-

[93] See their diagram (page 1424) for a more elaborated explanation of levels and standards.

documented guidelines that reduces the likelihood of bias in review papers, thereby increasing the integrity of the epistemic status of the claims advanced by the authors of those papers. [94]

2.3 Taxonomy of Review Papers in the Mind-Brain Sciences

The history and development of research synthesis as provided by Chalmers and Altman (1995) and Gurevitch et al. (2018), Higgins and Green (2011), Moher et all (2009, 2015), Stewart et al. (2015), focus almost exclusively on *the medical and health sciences*. Handbooks and critical research on the nature of research synthesis, such as those by Cooper, Hedges and Valentine (2009) and Rosenthal (1991) focus on *the social sciences*. Further, some of the largest and most popular cross-disciplinary organizations, such as the Cochrane Collaboration, contextualize research synthesis and related methods/guidelines in *the health/medical sciences*. In their opening 'about us' statement, the Cochrane Collaboration states: "Cochrane is for anyone interested in using high-quality information to make *health decisions*. Whether you are a doctor or nurse, patient or caretaker, researcher or funder, Cochrane evidence provides a powerful tool to *enhance your healthcare knowledge* and decision making" (https://www.cochrane.org/about-us, emphases mine).

Shamseer et al. (2015) similarly highlight the importance of research synthesis for the health/medical sciences: "Systematic reviews hold a unique place in *healthcare*. They help form the basis for developing practice guidelines and they provide information on gaps in knowledge, thus informing future research efforts. This information is relevant to stakeholders across *the health system*" (https://www-bmj-com.proxy1.lib.uwo.ca/content/349/bmj.g7647, emphasis mine). They continue to note that systematic reviews are particularly trustworthy due to the 'a priori planning and documenting' of their methods and protocols (ibid). While this might help reduce bias and increase the reliability of claims emerging out of systematic reviews, it is important to note that most organizations for research synthesis are grounded in *the medical/health sciences* and/or the *social sciences*:

> Various international organizations such as the Cochrane and Campbell Collaborations and the Agency for Healthcare Research and Quality (AHRQ) regularly require and publish

[94] It is an open question whether and to what extent this occurs in research synthesis in the mind-brain sciences, especially since, a point I have emphasized and will re-emphasize later, there are no organization or checklists dedicated specifically to writing review papers in the mind-brain sciences.

> protocols. However, outside of such organizations, few protocols are published in traditional journals and most reports of completed reviews (89%) do not mention working from a protocol… *Outside of select systematic review organizations, little to no general guidance exists for preparing review protocols* (ibid, emphasis mine).

This is problematic because many disciplines outside the health/social sciences make use of research synthesis (in many forms) in the dissemination of knowledge. Of particular importance for my dissertation is the fact that while the mind-brain sciences make great use of research synthesis via review papers, there is no organization to date dedicated *solely* to outlining criteria/standards for writing review papers in the mind-brain sciences, nor has there been much analysis done on the nature of research synthesis within the context of the mind-brain sciences.[95] [96]

Now that the background and motivation for research synthesis is in place, I will survey the different types of review papers found in journals for the mind-brain sciences. Despite having no handbook or organization dedicated to laying out protocol for review papers in the mind-brain sciences, journals in the mind-brain sciences are replete with research synthesis in the form of review papers (both with and without meta-analysis). Given that the standards for review papers will be different than those in the health sciences, it is important to look at review papers in the mind-brain sciences to see how, if at all, claims are being corroborated by the synthesis of interdisciplinary evidence.

---

[95] Again, part of the reason for introducing the standards and guidelines advanced by organizations such as the Cochrane Collaboration was to set the historical background for the development of research synthesis and for comparative reasons, given that the standards set by organizations for writing review papers in the health and social sciences will not be applicable to the types of review papers found in mind-brain science journals.

[96] Of the many journals in the mind-brain sciences, few provide any guidance with respect to review paper writing standards, with the exception of *Frontiers in Neuroscience*, which states authors should adhere to the guidelines published in either the Cochrane Collaboration, the Campbell Collaboration, or the PRISMA checklist (https://www.frontiersin.org/journals/neuroscience#article-types). However, seeing as the mind-brain sciences are distinct from the social and health sciences, in that they bring together information from multiple areas of science, using different methods, experimental protocols, techniques and tools of analysis, one might argue for and advance a unique set of standards for writing review papers in the mind-brain sciences. I consider this prospect in chapter four of this dissertation.

There are numerous journals that researchers in the mind-brain sciences can submit their work to, including but not limited to: *Frontiers in Neuroscience*, *Trends in Cognitive Science* (*Elsevier*), *Neuron, the Journal of Neuroscience, Nature Neuroscience*, etc. I focus only on a subset of these top journals – specifically *Frontiers in Neuroscience* and *Trends in Cognitive Science* (*Elsevier*) – that contain review papers for two reasons. First, I must constrain my focus for the sake of space, and second, of the above-mentioned journals, these two provide the most detail about the different types of review papers that can be published within their journal and what each type entails.[97] Each journal classifies article types slightly differently, however, as I will demonstrate shortly, there are three main category types that review papers may fall under.

Frontiers in Neuroscience is a popular mind-brain science journal. In their 'scope and mission' section, they describe their journal as:

> … publishing *rigorously peer-reviewed research across a wide spectrum of specialties and disciplines*. …This *multidisciplinary open-access journal* is *at the forefront of disseminating and communicating scientific knowledge and impactful discoveries* to researchers, academics, clinicians and the public worldwide. … *we seek to integrate and cross-link studies and citations in related subfields*, providing an overview of the state-of-the-art in these fields and the ways in which they complement each other. After all, we may each of us look at the brain from particular standpoints - all the way from "genes to behavior" - yet the brain integrates all these levels seamlessly; this journal should reflect this in the best possible way. (https://www.frontiersin.org/journals/neuroscience#about, emphases mine).

Statistically, *Frontiers in Neuroscience* is at the forefront of research in the mind-brain sciences. That is, the journal is the number one most cited journal in the neurosciences (https://reports.frontiersin.org/reports/frontiers-in-neuroscience-report.html). *Frontiers in Neuroscience* boasts an impact factor (i.e., the average number of citations in a given year by papers that were published from the two preview years) of 3.566, a total of 13 520 508 article

[97] The other journals either provide little (in comparison to *Frontiers in Neuroscience* and *Trends in Cognitive Sciences*) to no breakdown of article types nor any advice on what makes for a strong review article.

views and downloads, 691 mentions in newspapers such as Forbes and the Washington Post, and 20,063 citations (ibid).[98]

Frontiers in Neuroscience provides a list of tier 1 article types published in their journal, and I elaborate on each kind of review paper below (https://www.frontiersin.org/journals/neuroscience#article-types). Non-review-based tier 1 articles found in the *Frontiers in Neuroscience* journal include but are not limited to: original research papers, methods papers, protocol papers, technology report papers, perspective papers, case report papers, etc.[99]

Tier 1 review-based articles found in *Frontiers in Neuroscience* include the following:

1. The systematic review
2. The review
3. The mini review

Authors of systematic review papers advance *syntheses* of published studies on a specific topic (e.g., episodic memory, retinotopic maps, semantic processing, spatial cognition, etc.) "… that uses *systematic and clearly defined methods* to identify, categorize, analyze and report aggregated evidence on a specific topic" (ibid, emphasis mine). Authors must address an important question or problem in their field of research, and their answer must be grounded on the "… sound empirical basis of studies which were conducted well scientifically" (ibid).[100] Authors of systematic review papers do more than provide a summary of the current state of some field of research or topic; they also *compare* and *explain* the findings they cite in their summary and *synthesize* those findings into a claim about the mind-brain. The findings involved in the systematic review may all belong to the same field (e.g., the author(s) may compare a number of studies in *the field of psychology* on the nature of mind-reading) and/or they may

[98] Cited statistics are based on the most recent Journal Citation Reports (2017) .

[99] I will comment on these and their relation to review papers shortly.

[100] No elaboration is provided with respect to what a 'sound empirical basis' entails nor how one would evaluate/discern between studies which are empirically sound and those which are not.

involve studies from a variety of fields (e.g., the author(s) may compare a number of studies from psychology, cellular and molecular biology, cognitive neuroscience, etc., on the nature of mind reading). Systematic reviews may or may not contain meta-analyses (i.e., the quantitative analysis of data from multiple independent studies on the same phenomenon conducted with the goal of discovering important trends). The author may use the results of their meta-analysis (or their qualitative analysis) to advance a synthesized claim about the mind-brain. Authors may also, in closing, suggest directions for future research. While systematic reviews include all four components listed in section 2.2 of this chapter – i.e., a summary, explanation, synthesis, and directions for future research – the primary component to these papers is the *synthesis* of various studies (which may include original research papers, books, case reports, perspective papers, other published review papers, etc.) via either a meta-analysis or qualitative analysis.

The second article type is the review paper (hereafter the *Frontiers* review paper to distinguish it from another type of review paper outlined in *Trends in Cognitive Science*). *Frontiers* review papers focus on topics that have undergone a significant change (or which have undergone substantial progress) and present a comprehensive overview of this topic (ibid). Instead of the main emphasis being synthesis (as it was for the systematic review), the main emphasis of a *Frontiers* review is the *summary* of the current state of a topic about the mind-brain. Authors of *Frontiers* review papers are to discuss critical aspects of their summary such as different sides of the debate, consensus and controversies, key concepts and problems (including gaps in the research), and they are also expected to close directions for future research (ibid). *Frontiers* review papers, like systematic review papers, are not to include any of their own unpublished research; instead, they draw on other published research when presenting their summary.

The third article type is the mini review paper. The mini review is similar to the *Frontiers* review paper but differs primarily with respect to its scope/length. That is, mini reviews tend to be more focused in scope, and home in on a particular facet of some topic (e.g., instead of *the general topic of memory*, a mini review might summarize findings on *the more specific topic* of *spatial memory in rats*) (ibid). Authors of mini review papers are to present a clear and concise *summary* of a topic in the mind-brain sciences. This brief but succinct summary allows readers to develop an understanding of the current state of the literature, its key concepts, controversies and problems, as well as future research possibilities (ibid). Much like authors of the *Frontiers* review paper, authors of the mini review paper cannot include any of their own unpublished research, citing other published research. As mentioned above, the key difference between review and mini review papers is scope and length of the review, with the former typically having a maximum word count of 12 000 words (enabling the author(s) to write on a more general topic like memory) and the latter 3000 (enabling the author(s) to home in on a very specific aspect of memory) (ibid).

Another popular academic journal for the mind-brain sciences is *Trends in Cognitive Sciences* (*Elsevier*). *Trends in Cognitive Sciences* (hereafter *Trends*) is another highly influential peer-reviewed journal focusing on review-based articles on topics in the cognitive sciences/neurosciences. *Trends* boasts an impact score (i.e., the average number of citations in a given year by papers that were published from the two preview years) of 15.557 and a total of 25391 citations according to the most recent 2017 Journal Citation Report (https://www.bioxbio.com/journal/TRENDS-COGN-SCI).

According to their 'Aims and Scopes' section, *Trends* is *a highly interdisciplinary journal*, bringing the fields of psychology, computer science, neuroscience, anthropology, evolution, economics, and political science, etc. together to provide both expert and non-expert readers with up-to-date knowledge and insight about important topics and debates in the mind-brain sciences (https://www.cell.com/trends/cognitive-sciences/aims). According to *Trends*:

> *Trends* journals seek to provide all scientists, from the tenured to the tenderfoot, with concise and curated updates on the latest research. It is our aim to highlight new scientific developments and their impact on the world outside the laboratory. Our high-caliber articles are cutting edge, provocative, yet accessible and are written by *the most authoritative voices in science today*. They are intended not only to bring readers up to speed on recent progress in the field, but also to serve as platforms for debate and to push the boundaries of conventional thinking. The *Trends* journals offer more than summaries, they *contribute insight* https://www.cell.com/trends/cognitive-sciences/authors).

Each issue of *Trends* includes what they call review articles (hereafter *Trend*s review papers, to distinguish them from *Frontiers* review papers) which function to bring readers up to date on recent debates and critical developments with respect to a mind-brain related topic.

Trends review papers are central to the monthly issues in *Trends* and are written by leading scholars – those whom they describe as having some of the 'most authoritative voices in science' in their respective fields (ibid). Authors of *Trends* review papers present the reader with a balanced overview of the current state of some topic in the mind-brain sciences. They do so by providing relevant background information and highlighting important developments and discoveries. While *Trends* review papers focus on the *summary* component, these papers often include *explanations* and *directions for future research*. That is, while authors of *Trends* review papers are not to include their own original research, *Trends* review papers do "… allow room for some speculation and debate …" (ibid). Authors can *explain* and *comment on* the findings presented in their summary and/or *take a stance with respect to a debate* (with the caveat that authors are clear when they are providing their own opinions) (ibid). It is within this space that authors can advance arguments regarding the findings contained in their summary. For example, an author might argue that the summarized findings support a given claim about the mind-brain. This claim can be read as a synthesized claim insofar as it is the amalgamation of the summarized findings; it is the product of a qualitative analysis of the cited literature which, taken together, suggests a given claim. Importantly, however, this synthesized claim is not the product of a quantitative meta-analysis or statistical modelling (like it is in the case of *Frontiers* systematic review), but is instead the result of a qualitative analysis. Authors of *Trends* review papers are also strongly encouraged to conclude with directions for future research by clearly indicating to the reader "… the most productive avenues for future research and by highlighting current and future limitations" (ibid). Critically, *Trends* review do not include original unpublished research, any new hypothesis, mathematic model(s) nor are they to include meta-analyses (ibid).

Before moving onto the advantages of review papers, something should be said about the nature of the relationship between review papers and the literature upon which they draw.[101] The types of literature review papers draw on include but are not limited to: original research papers, case reports, perspective papers, books, textbooks, clinical trial papers, conceptual analysis papers,

[101] When I use the term 'review papers' here I am referring to all four types surveyed thus far – i.e., the systematic review, the *Frontiers* review, the mini review and the *Trends* review.

opinion papers, etc.[102] There is a two-way influence when it comes to review papers and non-review papers. In writing a review paper, authors survey literature on a given topic. The findings from this survey may influence and shape the authors' understanding of the topic; this might manifest itself in the way in which the summary is presented, the nature of the explanations of the summarized findings, and/or the synthesis. Once a review paper is published, it has the potential to influence the course of future research. That is, it might influence a new wave of original research. If, to provide just one example, an author of a review paper closes with directions for future research indicating a gap within a theory, researchers may strive to fill in said gap, thus adding new research that may influence a future review paper (or might be used to update an existing review paper). Thus, there is a two-way influence between review and non-review papers, and those who herald the importance of research synthesis for science acknowledge this influential relationship as crucial for progress and better explanations in the sciences. For example, Gurevitch et al. (2018) state: "Evidence synthesis should become a regular companion to primary scientific research to maximize the effectiveness of scientific inquiry" (180).

In sum, three distinct types of review papers emerge from the types I survey: the systematic review, the *Trends* review, and what I call the general review paper (hereafter shortened to GRP, which is essentially an amalgamation of the *Frontiers* review and the mini review). The systematic review stands on its own insofar as its main component is *synthesis* (often the product of *a meta-analysis*). The GRP's are grouped together insofar as their primary component is *summary*. Finally, the *Trends* review stands on its own as it arguably has two main components – *a summary*, and *an explanation* of the summarized findings (which can sometimes take the form of *a synthesized claim*, although the synthesized claim will not be the product of a quantitative meta-analysis as is often the case with the systematic review). This classification is represented visually below in figure 1.

[102] When referring to these types of papers in general, I will simply call them non-review papers.

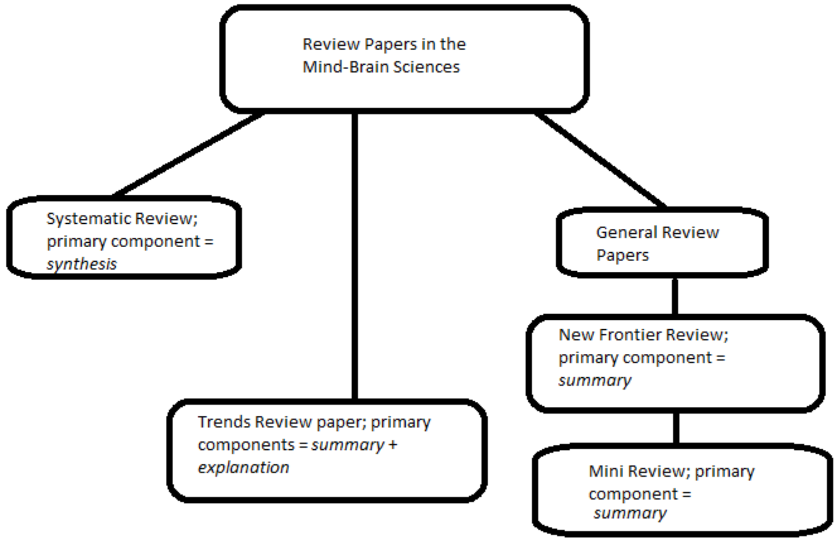

Figure 1 - Represents my taxonomy of review paper types in the mind-brain sciences based on their main components.

Now that I have laid out a taxonomy of review paper types in the mind-brain sciences, I turn to some of the advantages they possess, before closing this chapter with their implications for IC.

2.4 Advantages of Review Papers for Progress in the Mind-Brain Sciences

In general, review papers in the mind-brain sciences provide readers with current, up-to-date knowledge about the mind-brain (or an aspect of the mind-brain). [103] They help shed light on important concepts, distinctions, debates, and problems or gaps in the literature. Ideally, this allows readers to understand what is at stake and potentially influences their thinking and future research about the mind-brain.

These are the general features of review papers; what about the specific features of each individual type of review paper? Do the different types of review papers surveyed above possess distinct features over above the others? Assessing the unique features of each type of review paper will help set the stage for understanding why review papers are a good place to start investigating the nature of IC. In this section, I survey the unique advantages of GRP's, the

[103] I use the term reader here to mean readers and/or researchers (as often readers are themselves researchers).

systematic review, and the *Trends* review, respectively. Once this is complete, I draw out the implications review papers have for IC.

One major feature of the GRP I consider advantageous for scientific research is their ability to enhance knowledge and update readers on the current state of a topic in the mind-brain sciences. This feature is in part due to the primary component of the GRP being the *summary*. That is, a large portion of a GRP is dedicated to providing a general overview of a topic including things like historical background and development, critical debates and controversies, key concepts and theories, recent progress and problems, etc. All of this helps provide readers with a general understanding of the current state of a topic in the mind-brain sciences. This feature likely has different implications depending on the type of reader (e.g., non-expert layman, non-expert undergraduate, graduate student, expert, etc.). For the non-expert laymen and undergraduate students, this overview functions somewhat like an introductory handbook (which will include background information, key concepts and problems, etc.) on the topic, allowing them to achieve a general understanding of a topic they previously knew little to nothing about. From this point, the non-expert reader might be influenced to pick up on certain aspects of the summary; for example, an undergraduate student might find a certain problem interesting, and decide to pursue graduate studies dedicated towards investigating (potentially solving or deepening) said problem.

Expert/advanced readers may find more benefit in mini review papers, which tend to be less generalized, and more focused on a specific feature or problem pertaining to the mind-brain. The mini review therefore might provide the expert reader with new insight, potentially influencing them to fine-tune their own research (e.g., perhaps by attempting to fill in a gap in the research), or to conduct a very specific study in relation to a problem noted in the mini review, thereby contributing more original research on the topic.

Systematic reviews boast numerous advantages; however, I focus on the two most relevant to the goal of my dissertation (i.e., investigating the nature of IC in the mind-brain sciences and assessing whether and to what extent review papers corroborate claims about the mind-brain). The two advantages I identify are the ability to integrate large amounts of information and the potential to corroborate claims about the mind-brain. Starting with the former, as mentioned earlier in this chapter, the 21st century is inundated with an avalanche of information coming in from multiple areas of science. To provide a general example to set the stage, Mulrow (1995) notes in the medical sciences "… large quantities of information must be reduced into palatable pieces for digestion. Over two million articles are published annually in the biomedical literature in over 20 000 journals – literally a small mountain of information. … In a stack, two million such articles would rise 500 m" (1).[104] With such a large number of articles coming in annually, a method for sorting through this information in such a way that identifies and picks out the relevant and scientifically sound papers from the irrelevant and unsound articles is necessary. Mulrow believes this sorting method is accomplished by authors of systematic review papers who can critically analyze, evaluate, and sort through the mass amounts of information (ibid, 2). Mulrow does not argue how nor why these authors are able to do so, but I suspect once again the assumption behind her claim is the fact that these authors possess the proper sorting tools in virtue of their abiding by pre-established and standardized guidelines for writing review papers. That is, authors of review papers can consult the Cochrane handbook or the PRISMA checklist, for example, scroll to the criteria for study selection, and sort through papers, selecting only those which are scientifically sound and relevant according to pre-established guidelines. Having a pre-determined method for selecting studies and data, analyzing and interpreting data, etc., appears to make authors of systematic review papers mini experts of sorts on how to efficiently triage and integrate large amounts of information.[105]

Understanding how systematic review papers can potentially corroborate claims involves analyzing the notion of synthesis (which is the primary component of the systematic review). When preparing to write a systematic review, an author will examine papers on a topic in the mind-brain sciences (e.g., memory) to determine whether there are any trends or patterns pertaining to a specific question or problem they are investigating. Accordingly, if a meta-analysis is conducted (i.e., either on findings from the same area of science or findings from different areas of science) and a general trend/pattern *x* is found, the author might argue that *x* is well-supported (i.e., it has a high degree of corroboration because a number of independent

[104] These statistics are from the late 20th C, and surely have risen in the 21st C. Importantly, these statistics are from the biomedical literature, wherein a large majority of the work on research synthesis has taken place for a while, however, the mind-brain sciences also deal with a large amount of published works on the mind-brain annually, and their numbers are also likely to be high.

[105] Again, this is my interpretation of Mulrow's (1995) claim that systematic reviewers are best equipped to sort through and integrate large amounts of information into a claim that stands on firm epistemic footing.

studies all suggest/or are consistent with/or point towards the same conclusion).[106] This claim is thus the amalgamation/synthesis of all surveyed papers. For example, if a meta-analysis is conducted and finds that all the cited papers in the systematic review (and their respective findings) implicate the hippocampus for memory, the author might argue the claim that the hippocampus is a critical brain region for memory is corroborated by the literature. That is, this claim appears to be corroborated by findings from different areas of science all of which appear to point towards the same general conclusion about the hippocampus and memory.

If systematic reviews can corroborate claims, these papers have the potential to influence education and future research. For example, an allegedly corroborated claim may be taken as evidence against an alternative competing claim. That is, if some alternative claim *x* is supported by ten studies from one particular area of the mind-brain sciences, meanwhile claim *y* has been supported by over eighty studies from multiple areas in the mind-brain sciences, then claim *y* (hereafter shortened to '*y*') might be taken as evidence against claim *x* (hereafter shortened to '*x*'), since it has a substantially higher degree of corroboration than *x*. This may lead researchers to either abandon their work on *x*, or to try and save *x* by perhaps modifying it in some way or other.[107]

Further, corroborated claims emerging from systematic review papers may be taken as providing a foundation of sorts. If *y* has demonstrated a strong degree of corroboration (and this has lasted for a substantial amount of time), then one may take *y* as confirmed. This confirmation may potentially make it such that *y* becomes a starting point for future research to expand on. For example, if *y* represents 'the hippocampus is a critical brain region for memory' and *y* has received a high degree of corroboration, new research projects might emerge trying to establish the finer details of *y* (e.g., what other structures, regions, networks are involved, what other types of memory are realized in the hippocampus, how can one incorporate claim *y* within a dynamical brain systems theory, etc.).

---

[106] For simplicity, I use the term meta-analysis here to refer to both qualitative and quantitative analysis of multiple pieces of evidence.

[107] Of course, one could argue that in time claim *x* might accrue more support than claim *y*. Perhaps with newer and better technologies, and more research, time and funding being poured into claim *x*, it may turn out that claim *x* actually has a higher degree of corroboration than claim *y*.

In an education setting, a chapter on memory may base all of its teaching off the fundamental assumption that *y* is true, so that the undergraduate studying memory for the first time may come to accept this claim as true.[108] In sum, the two advantages of the systematic review, ideally, are its ability to integrate large amounts of information about the mind-brain and its ability to corroborate said claims via meta-analysis (either qualitative or quantitative, although more frequently the latter).

Trends reviews also have the potential to integrate information and corroborate claims. With respect to the former, because *Trends* reviews are written by leading scholars considered to be experts in their respective area of science, their expertise and knowledge on a given topic makes them particularly well-suited for triaging numerous studies in the literature and extracting only the most scientifically sound and relevant. Their ability to corroborate claims is grounded in one of the primary components of a *Trends* review: *explanation*. That is, authors of *Trends* reviews are provided room for 'speculation and debate' within their review (https://www.cell.com/trends/cognitive-sciences/authors). It is within this space that the author may explain and comment on the summarized findings and advance their own conjecture based on the current state of the literature (as per the author's summary). This conjecture can sometimes take the form of a *synthesized claim*; that *synthesis* would likely be the product of integrating the summarized findings into a cohesive explanation about an aspect of the mind-brain.[109] The caveat with *Trends* reviews advancing their own conjectures and speculations is simply that the author be clear that this is his/her/their personal opinion (https://www.cell.com/trends/cognitive-sciences/authors). Interestingly, because *Trends* claims to contribute deep insight into the mind-brain by producing reviews written by only the greatest scholars with the most 'authoritative voices in science today', the *synthesized claim* emerging from the *Trends* review appears to stand on strong epistemic footing insofar as it is the integrated product of a *qualitative analysis/synthesis of multiple studies* (recall that *Trends* reviews do not include quantitative meta-analysis or statistical models) from different areas of the mind-brain

[108] A critically minded student may challenge this assumption, but the point is simply that corroborated claims can sometimes take a foundational form whereby students are taught that this claim is confirmed.

[109] The author of the *Trends* review *may not necessarily* include their own opinion on the debate or problem at all, or if they do, it *may not necessarily* take the form of a synthesized claim (integrating the information from different areas of science). The point is merely that their opinion *can potentially* take the form of a synthesized claim, and when it does, this claim can appear to be corroborated.

sciences. Further, because the claim is advanced by an authoritative and well-respected voice within the mind-brain science community, that claim might gain more epistemic warrant (ibid). The allegedly corroborated claims emerging out of *Trends* reviews thus have the potential to influence future research in the same way that the systematic review does: that is, they can potentially act as evidence against another alternative claim or theory, and/or they can act as a foundation of sorts, both for research and in education.[110]

At this point, one might ask what it is about review papers and their specific features that make them uniquely suited to the project of exploring the nature of IC in the mind-brain sciences. The goal of the next section will be to make this connection succinct.

2.5 Implications of Review Papers for IC

To understand why review papers are a good place to start investigating the concept of IC – that is, what features they possess that makes them uniquely suited for the task – it is helpful to re-examine the main components of IC. In section 2.1 of this chapter, I provided a working definition of IC whereby IC occurs when findings from different areas in the mind-brain sciences corroborate a claim about the mind-brain in such a way that those findings point towards a claim and are mutually supportive and interdisciplinary. I noted there are two main components to IC – interdisciplinarity and corroboration. I will now briefly re-visit both aspects of IC.[111]

The interdisciplinarity aspect of IC relates to the different areas of science housed under the mind-brain sciences. Such disciplines include but are certainly not limited to computer science and AI, biology, psychology, genetics, linguistics, neuroscience, etc. Both systematic reviews and *Trends* reviews survey findings from multiple areas of the mind-brain sciences in their investigation of a given topic; therefore, both satisfy the interdisciplinarity component. Further, *Trends* prides itself on being a *highly interdisciplinary research journal*, and thus one would assume the articles published therein would be too (or at least contribute to that aspect).

Corroboration occurs when pieces of evidence confirm/support a claim. In section 2.1 of this chapter, I noted individual research papers provide a good first pass at demonstrating

[110] See above for elaboration on these two points.

[111] This is just a brief review; for a more complete description of 'interdisciplinarity' and 'corroboration', please see section 1 of this chapter.

corroboration in that they (sometimes) try to situate their original findings within the broader context of the literature by citing other findings that support their main claim/finding, but the degree of corroboration in these papers was too narrow in scope.

The role of individual research papers, however, is not insignificant to corroboration. As Gurevitch et al. (2018) note, “An individual primary study may now be seen as a contribution towards the accumulation of evidence rather than revealing the conclusive answer to a scientific problem” (178), and, as Mulrow (1995) puts the point, “The hundreds of hours spent conducting a scientific study ultimately contributes only one piece of an enormous puzzle” (7). Corroboration thus relies on different findings from original research papers; the more findings point towards *y*, the more one can argue *y* has been corroborated. Thus, for findings to corroborate *y*, multiple research sources must confirm/lend support to/point towards *y*.

Further, a point emphasized by philosophers of science (e.g., Bechtel and Richardson (1993); Campbell (1966); Hacking (1983); Wimsatt (2007), etc.), those findings must be the result of ‘independent measures’ (Bechtel and Richardson 1993, xxvii). This point connects both the interdisciplinarity component and the corroboration component in that when the different areas of the mind-brain sciences (interdisciplinarity) all employ different measures (i.e., they appeal to specialized tools, methods, technology, software, etc.), and they all point toward *y*, then *y*, it might be argued, has been corroborated. Thus, the more original research from different areas of the mind-brain sciences points towards *y*, the more corroboration *y* accrues.[112]

---

[112] This process of multiple independent research papers contributing one piece of the puzzle which is completed when multiple pieces of evidence from different areas of science point towards some claim is strikingly similar to Roskies’ (2010) notion of convergence via functional triangulation. In this case, the different studies from different areas of science can be said to triangulate/converge on the same claim, and insofar as this is done, the claim stands on firm epistemic footing; this, arguably, could be what Nagel’s (1949, 1979) ‘essential condition’ requires.

Further, one might interpret the synthesis component of the systematic review and the *Trends* review to be akin to the kinds of ‘integration’ required by Mitchell (2002, 2003), Mitchell and Dietrich (2006) and Piccinini and Craver (2011). That is, the product of the review paper – *the synthesis* – is the act of integrating all the results of different areas of science into one claim; this resulting claim is considered to be well-supported/confirmed and accordingly stands on firm epistemic footing.

On this point, Bickle (2003, 2006) might claim that it is the analysis of different studies from different labs in cellular and molecular neuroscience that have been synthesized. In fact, in Bickle’s (2006) example of the cAMP-PKA-CREB pathway as the molecular mechanism for social recognition memory in mammals, Bickle points to the fact that this conclusion has been corroborated by a number of other independent original studies on the topic. In other words, a large number of studies revealed a general trend – i.e., that this pathway is the molecular

The concepts of interdisciplinarity and corroboration are echoed in *both* the systematic review and the *Trends* review. Both papers survey research papers from different areas of the mind-brain sciences on a topic (interdisciplinarity).[113] Authors of systematic review papers will analyze these papers and their respective findings to determine any overall patterns/trends; this is often performed via meta-analysis and statistical modelling. If a trend is identified, authors advance a claim which can be considered *synthesis* of the surveyed findings involved in the meta-analysis.

Authors of *Trends* reviews summarize the current state of some problem or topic about the mind-brain, citing multiple papers from different areas of the mind-brain sciences. They may also comment on the nature of their findings and this comment (the authors' opinion) may take the form of a *synthesized claim*.[114] This *synthesis* provides the reader with a more complete understanding of the mind-brain and is considered more reliable insofar as it is supported by numerous interdisciplinary findings and is advanced by an author considered to be an influential and authoritative voice on the matter.

The synthesized claims falling out of both the systematic and the *Trends* review (as the result of either a quantitative meta-analysis or a qualitative analysis) relates to the corroboration aspect. That is, if the interdisciplinary papers (and their respective findings) surveyed by both types of papers reveal a general trend – which is identified either by meta-analysis in the case of the systematic review, or qualitative analysis in the case of the *Trends* review – whereby that trend suggests a claim about the mind-brain, one might argue that that claim has been corroborated (in that it is supported by considerable interdisciplinary evidence).

---

mechanism for social recognition memory in mammals – and this trend is taken to be corroborated by Bickle (Bickle 2006, 423-424).

Interpreting this through Darden and Maull's (1977) notion of interfield theories, one might say that an interfield theory such as the chromosome theory of Mendelian genetics has been corroborated via numerous studies which confirm the interfield theory.

[113] Authors of course are not required to survey studies from multiple areas of science; some authors of review papers may choose to focus solely on evidence from one particular area of science. However, systematic review papers and *Trends* review papers frequently cite evidence from different areas of science, and accordingly, my claims extend only to papers which do so.

[114] Not all authors of *Trends* reviews will take their own stance on a debate, nor comment on those findings, nor will those comments (if they provide them) necessarily all take the form of a synthesized claim. My argument pertains only to those papers that do.

In section 2.3 of this chapter, I surveyed three main types of review papers (i.e., the systematic review paper, the *Trends* review paper, and the GRP). Of those mentioned, the systematic review and the *Trends* review paper do what is required for IC: that is, they review findings from different areas of the mind-brain sciences (i.e., interdisciplinarity), and if a trend/pattern is identified, authors of the review paper advance a claim taken to be a *synthesis* of those findings (i.e., a corroborated claim in that multiple pieces of independent evidence from different areas of science support this claim). Accordingly, both would provide a good starting point for investigating the nature of IC and analyzing whether and to what extent claims about the mind-brain are corroborated in practice.

However, due to the length of this project, it is not possible to analyze *both* a *Trends* review paper and a systematic review. Instead, I begin what I hope to be an ongoing project of investigating the epistemic role of review papers in the mind-brain sciences by analyzing a *Trends* review paper; because the *Trends* review contains less features than a systematic review (which often additionally includes meta-analyses), my analysis will create a baseline against which to compare future units of analysis for corroboration. Furthermore, is important to consider not just the systematic review, since journals in the mind-brain sciences contain and make use of different types of review papers (not *just* systematic reviews). Ioannidis (2016) notes that papers tagged as 'reviews' are published at a much larger scale than papers tagged as 'systematic reviews' in PubMed (489). Thus, it is important to not rush this investigation by immediately jumping to a systematic review; again, starting with a *Trends* review allows a baseline against which to compare different units of analysis with different/additional features (such as a meta-analysis) to determine how and whether the additional features increase the integrity of the claims about the mind-brain emerging from those papers. Since *Trends* claims to provide deep insight into the mind-brain via review papers from some of the most authoritative scholars, the *Trends* review paper, which contains both the interdisciplinarity and the corroboration component of IC, provides a strong starting point for investigating the nature of IC.

In chapters three and four, I analyze a case study of a *Trends* review paper by Eichenbaum (2013) entitled 'Memory on Time' as a potential instance of IC with the goal of extracting critical lessons for the notion of IC and its relation to strong epistemic claims in the mind-brain sciences.

Chapter 3

3 Introduction

In chapter one, I surveyed popular accounts of the relationship between different areas of science emphasizing the notion of fitting together as essential for scientific progress. I noted a problem with these accounts is their notions of fitting together were vague and subject to multiple clarificatory questions. In chapter two, I introduced the term ‘interdisciplinary corroboration’ (hereafter IC) as a placeholder for the notion of fitting together to get a better understanding of the assumption underlying the surveyed accounts of scientific progress. I noted, however, IC was a working definition which needed elaboration and analysis. I argued a good place to begin investigating IC and assessing the extent to which claims are corroborated in the mind-brain sciences was to look for instances of it in practice via a *Trends* review paper. As a case study for this investigation I have selected a paper by Eichenbaum (2013) entitled ‘Memory on Time’.

My investigation of this review paper comprises a two-part analysis structured around what Eichenbaum does (chapter three) and does not do (chapter four) to show support for the main claims he defends in this paper. The aim of this chapter is to showcase what Eichenbaum does to support the main claims he advances in ‘Memory on Time’ and how it is representative of IC. My analysis focuses on the language and terms used to describe the findings which he cites in support of his claims. I begin section 3.1 by explicating why I selected Eichenbaum’s review paper as my case study, what cursory criteria I use to guide my analysis, and what methods I use to analyze Eichenbaum’s paper. I provide a brief summary of ‘Memory on Time’ and the main claims advanced by Eichenbaum in section 3.2. In section 3.2.1, I analyze the ways in which Eichenbaum shows support for the main claims he defends by focusing on the language he used to describe the cited findings, such as the findings ‘show that’ or ‘suggest that’ *x* (Eichenbaum 2013, 83, 86)[115]. I assess the extent to which the case study conforms to IC in sections 3.2.2 (focusing on the interdisciplinary component), and 3.2.3 (focusing on the corroboration component), and argue that, *prima facie*, the case study appears to represent a strong instance of IC where interdisciplinary findings fit together in such a way that lends support to Eichenbaum’s main claims.

[115] Where *x* represents one of the claims Eichenbaum (2013) advances in ‘Memory on Time’.

However, in conducting this analysis, certain questions began to emerge such as "what counts as substantial evidence", "what does it mean for evidence to be consistent with each other", and "how are anomalies factored into the 'substantial evidence' that appears to be consistent with each other". These questions began to reveal cracks in what appeared to be a clear instance of IC in Eichenbaum's review paper. This led me to conduct part-two of my analysis in chapter four, wherein I deepen my analytic work in chapter three by explicitly addressing these questions and their implications for IC. As my aim in this chapter is to constrain my focus on the work Eichenbaum does to support the claims he defends, I bracket these questions and concerns for chapter four.

3.1 Methods

The case study I selected for this analysis is a paper by Eichenbaum (2013) entitled 'Memory on Time'. As noted in chapter one, I am interested in claims made in the mind-brain sciences about how the brain realizes psychological functions. Memory organization, specifically, spatio-temporal memory (i.e., memory for how events are organized in time and space) is a psychological function and is the main topic of my case study.

The paper I selected is a *Trends* review paper with a substantial number of citations.[116] However, more than a high citation count, I selected this paper because of the tremendous influence Eichenbaum has had on memory research. Eichenbaum (1947-2017) is considered by many to have been one of the leading pioneers in memory research. Hasselmo and Stern (2017) described Eichenbaum as a '*vital leader*' in illuminating the role of the hippocampus for memory and learning (875, emphasis mine). Eichenbaum held an impressive amount of positions in the scientific community including being a William Fairfield Warren Distinguished Professor, the director of BU's Center for Memory and Brain and the Laboratory of Cognitive Neurobiology, and the founder of "… what is now called the Graduate Program for Neuroscience …" (Brown 2017, https://www.bu.edu/today/2017/howard-eichenbaum-obituary/). Additionally, in 2015, he

[116] There were different papers on memory organization by Eichenbaum I could have selected as my case study. However, I chose 'Memory on Time' (2013) which has 179 citations to date as a middle ground between his (2017) 'On the Integration of Space, Time, and Memory' paper, which has 54 citations, and his (2014) 'Time Cells in the Hippocampus: A New Dimension for Mapping Memories' paper, which has 247 citations. Eichenbaum's (2013) review 'Memory on Time' has enough citations that a case can be made that the paper is influential, but not too many that tracking the citations and cited studies (the 2014 paper has approximately 70 more cited studies than the 2013 paper) would be too lengthy.

was elected into the American Academy of Arts and Science and was a Fellow of the American Association for the Advancement of Science, where he acted as chair in the neuroscience division and the Association for Psychological Science (ibid). He also served on the Council of the Society for Neuroscience as well as the National Institute of Mental Health (ibid). Finally, Eichenbaum was the editor-in-chief of the journal *Hippocampus* until his passing in 2017 (Hasselmo and Stern 2017, 875).

Eichenbaum was an '*internationally recognized figure*' known for making headway and advancing knowledge of the hippocampus and the neural mechanisms for memory organization (Brown 2017, https://www.bu.edu/today/2017/howard-eichenbaum-obituary/). Commenting on Eichenbaum's work, Dr. David Somers, professor and chair of psychological and brain sciences, states:

> His extensive empirical findings … his *integrative approach*, which *synthesized results across species*, *across methods*, and *across levels of analysis*; and his important theoretical advances concerning multiple memory systems in the brain, all helped advance our understanding of how memory works and how it is organized in the brain" (ibid, emphases mine).

Within memory research, Eichenbaum is considered a 'pioneer', a 'titan', and an 'internationally recognized figure', making him a top scholar on memory research. Given his authoritative status within memory research, the impressive and deep impact his research has had and continues to have on the mind-brain science community, as well as the fact that much of his research, as Somers points out, includes *an integrative approach* which *synthesizes results* across different areas of science (an aspect emphasized by both reductive and non-reductive accounts of progress in chapter one of this dissertation), Eichenbaum's paper is a great candidate for my case study.

My goal for my analysis of Eichenbaum's (2013) 'Memory on Time' is to evaluate how and to what extent claims about the mind-brain are corroborated within this review paper. From my analysis, I hope to draw out important lessons concerning the nature of IC and how claims about the mind-brain are corroborated in units of analysis such as the *Trends* review paper. These lessons might serve as a springboard for future research; I discuss such prospects in section 4.4. of chapter four.

Importantly, I begin this analysis with a general understanding of how review papers function from my descriptive work in chapter two. That is, *Trends* review papers are packaged in such a way that they include a *summary* of a topic in the mind-brain sciences, an *explanation* of findings cited in the summary section, and *directions for future research*. I noted, in chapter two, that the explanation component can manifest in a number of ways. For example, authors of review papers can explain the relevance of the scientific findings they cite, they can explain those findings by contextualizing them within a debate or within the broader picture of a mind-brain topic, they can comment on the nature of those findings, they can argue how those findings fit together to support a given claim, etc. In providing explanations, authors sometimes argue for claims that are a *synthesis* of the findings they survey in their *summary*. This, however, is a superficial understanding of how review papers work; the purpose of my analysis is to do the analytic work required to discern how and to what extent claims about the mind-brain are being corroborated in units of analysis such as the *Trends* review paper.

Accordingly, I come to the table with some cursory criteria to help address these questions and guide my analysis. These criteria includes the following questions: "what is the aim of Eichenbaum's review paper", "what claims does Eichenbaum argue for", "how does Eichenbaum show support for those claims", "are Eichenbaum's arguments implicit or explicit", and "does the case study represent a strong case of IC, and if so, how"?

To evaluate how and to what extent claims are corroborated in Eichenbaum's (2013) 'Memory on Time', I begin by laying down some descriptive work by providing a brief synopsis of Eichenbaum's paper, including the paper's purpose and its main claims. I consider the ways in which Eichenbaum shows support for those claims. In analyzing this, I notice that a crucial aspect of Eichenbaum's arguing for the claims he defends involves citing interdisciplinary findings on memory organization and using certain descriptive terms/phrases indicative of corroboration. Accordingly, I focus my analytic work around the language and terms Eichenbaum uses to describe the findings he cites in support of his claims. When relevant, I appeal to the qualitative and quantitative analysis of important terms via a software program called NVIVO12, as well as other important statistical information such as citation counts and impact factors.[117] [118] [119]

3.2 Summary of 'Memory on Time'

Trends reviews involve a *summary* of the current literature on a given topic as well as explanations/comments that *sometimes* take the form of a *synthesized claim* (provided the author clearly states when the claim in question is their own opinion). The main aim of 'Memory on Time' (hereafter referred to as 'the case study') is to establish the role of the hippocampus as creating a spatio-temporal framework for the organization of memories. In the case study, Eichenbaum summarizes findings from recent literature on the spatio and temporal aspects of memory organization. He explains the significance of the findings from this literature by stating they all fit together to support claims about memory organization. Specifically, he states that in addition to summarizing literature on memory organization, he will relate the findings he cites on temporal representation to literature on spatial organization, and argue that, 'taken together', these findings suggest a role for the hippocampus which functions to create a spatio-temporal framework for the organization of memories (Eichenbaum 2013, 81).

My analysis yielded what I consider to be five synthesized claims advanced by Eichenbaum in the case study. Importantly, Eichenbaum only flags what I identify as the fifth claim as a synthesized claim; this is the main claim that Eichenbaum explicitly argues for in the case study. However, because Eichenbaum states what I identify as claims 1-4 are supported by findings from different areas of science which, when combined, results in a specific claim about the nature of memory organization, I read each claim as a unique synthesis of the findings from Eichenbaum's cited research papers. Thus, on my reading, the case study contains four implicit claims (claims 1-4) and one explicit claim (claim 5). The five synthesized claims I take Eichenbaum to advance and defend in the case study are:

---

[117] NVIVO 12 is one of the leading qualitative research software programs on the market for determining insights in unstructured data and is used internationally by academia, government, non-profit, and commercial organizations (Kolodny, online tutorial 2019-03-18).

[118] All citation counts were acquired through Google Scholar, which enables users to "… gauge the visibility and influence of recent articles in scholarly publications" (https://scholar.google.ca/intl/en/scholar/metrics.html).

[119] All impact factors were acquired through the most recent statistics released by the Journal Citation Report (JCR) on their website and/or the website of the journal in question.

Claim 1: the hippocampus is responsible for the memory of the order of events (ibid, 83).

Claim 2: the hippocampus disambiguates overlapping sequences of events (ibid).

Claim 3: hippocampal neurons replay sequences of events by encoding and retrieving the order of events (ibid).

Claim 4: hippocampal time cells possess the properties of context cells, allowing them to bridge the order of events[120] [121] (ibid, 86).

Claim 5: the combination of claims 1-4, in addition to previously unmentioned literature on memory organization, entail that the role of the hippocampus is to provide a spatio-temporal framework for the organization of memories (ibid, 86-87).

Of the synthesized claims listed above, I make a distinction between two types of syntheses.[122] I identify claims 1-4 as mini syntheses, and claim 5 as the macro synthesis (i.e., the main claim advanced by Eichenbaum in the case study). My reasoning for the distinction is that claim 5 not only synthesizes findings from previously uncited studies, but also syntheses claims 1-4 (and the papers/findings cited for each of those). In this sense, claim 5 is the culmination of the work done to support claims 1-4, as well as additional references added to provide further support to claim 5.

In the case study, Eichenbaum argues that claims 1-5 are well supported by interdisciplinary research on memory organization. *Prima facie*, it appears that the synthesized claims are corroborated by results from research spanning multiple areas of the mind-brain sciences investigating memory organization. In the following section, I demonstrate what kind of work Eichenbaum does to provide support for claims 1-5 (i.e., how does he argue for their corroboration), and explore this analysis within the context of IC in sections 3.2.2. and 3.2.3.

---

[120] Eichenbaum defines time cells as: "neurons that fire at a specific moment within a temporally structured episode" (Eichenbaum 2013, 81).

[121] Eichenbaum defines context cells as: "theoretical neurons whose activity bridge events and therefore could link their representations in order. These may pre-exist or arise independent of specific events, or could arise as an integration of event information over time" (Eichenbaum 2013, 81).

[122] For length purposes, I shorten 'the synthesized claims 1-5' to 'claims 1-5' in the remainder of this dissertation.

3.2.1 What Eichenbaum Does to Show Support for his Main Claims

In the case study, Eichenbaum advances what I identify as claims 1-5; what does Eichenbaum do to show these claims are well supported? Put differently, what kind of work does Eichenbaum do to demonstrate that claims 1-5 stand on firm epistemic footing? How does Eichenbaum argue for each claim; what kind of arguments does he advance? In analyzing the case study, I noticed the majority of this work is done by stating findings from different research papers on memory organization and using certain terms to describe or introduce those findings and the claims that they suggest. The terms used to describe the findings Eichenbaum states in support of claims 1-5 are indicative of corroboration. To better understand how Eichenbaum provides support for his claims by using these key terms, I focus on claim 3 and claim 5 as examples.[123]

Claim 3 states hippocampal neuronal ensembles replay sequences of events in memory by encoding and retrieving the order of events (ibid, 83). Eichenbaum cites findings from six research papers on memory organization to demonstrate the substantiation of claim 3. Eichenbaum provides no critical analysis of these findings. Instead, he briefly states findings from each paper and states what he takes to be their implication(s) for memory organization. He begins by stating there is '*substantial evidence*' for the claim that hippocampal neuronal ensembles are responsible for representing the order of past events (ibid, emphasis mine). Importantly, however, Eichenbaum provides no explanation as to what constitutes 'substantial evidence', how it is quantified, etc. (ibid). He goes on to claim a great deal of evidence backing this claim comes from studies working on rats in a task where the animals routinely pass through the same route while investigators record place cell activation in the animals (ibid).[124]

Within this paradigm, Eichenbaum notes '*numerous studies*' have reported both forward and reverse replays of place cell sequences (ibid, emphasis mine). The 'numerous studies' Eichenbaum refers to consists of four papers investigating memory for order in rats (ibid). The first study Eichenbaum cites is by Karlsson and Frank (2009), who reported consistent replays of

---

[123] The analysis of claim 3 has similar results for claims 1, 2, and 4, since they all share the same structure (i.e., Eichenbaum reviews a certain amount of studies, which, he argues provide strong support for the mini syntheses). For length purposes I chose to constrain my focus on claim 3 as an illustration of what Eichenbaum does to show support for his mini synthesized claims 1-4.

[124] Eichenbaum defines a place cell as: "neurons in the hippocampus that fire when an animal (or human) is at a particular location in a spatially defined environment" (Eichenbaum 2013, 81).

past experience in animals both during periods of wakefulness and rest (Eichenbaum 2013, 83). After stating this, Eichenbaum cites a paper by Ji and Wilson (2007), who found hippocampal replays are synchronized with cortical replays, and a paper by Jadhav et al. (2012), which demonstrated that interfering with replays lead to impairments on memory for maze learning tasks (Eichenbaum 2013, 83). Eichenbaum claims the two findings together *suggest* that replays are an essential aspect of memory consolidation (ibid). Eichenbaum immediately cites a study by Johnson and Redish (2007), who found when rats undergo maze learning, hippocampal neuronal ensembles replayed possible route options (ibid). Eichenbaum states this finding *suggests* that replays represent memory for past events (ibid). Eichenbaum again does not explain how this finding fits together with the previous findings from Karlsson and Frank (2009), Ji and Wilson (2007), and Jadhav et al. (2012), but seems to take these studies as representing the 'substantial evidence' implicating hippocampal neuronal replays as essential for memory for sequences of events.

Eichenbaum continues to show support for claim 3 by claiming there is not only 'substantial evidence' from rat studies, but also evidence from human studies implicating hippocampal replays in memory consolidation. Specifically, a study from Gelbard-Sagiv et al. (2008) found neuronal ensembles within the human hippocampus replayed firing sequences when human participants were asked to recall scenes in a movie, and a study from Paz et al. (2010) found hippocampal neuronal ensembles demonstrated reliable firing patterns while human participants learned sequences from a movie (Eichenbaum 2013, 83). As with the rat studies, Eichenbaum does not explain the relation between the two human studies nor how they fit together with the rat studies other than immediately concluding the section by claiming: "These findings provide *compelling evidence* that hippocampal neuronal ensembles can encode and retrieve the temporal structure of sequential events" (ibid, emphasis mine).

This is what Eichenbaum does in the case study to substantiate claim three; I will now unpack this work by establishing what kind of an argument Eichenbaum provides for claim 3. Eichenbaum wants to demonstrate that claim 3 is supported by evidence from human and non-human animals. This evidence consists of four research papers investigating memory organization in rats, and two papers on memory organization in humans, for a total of six research papers. Generally speaking, all six findings implicate hippocampal replays in memory consolidation (in some shape or form). Eichenbaum's argument for claim 3 standing on firm epistemic footing is *implicit*; that is, he does not explicitly advance an argument explaining how and in what ways the cited findings support claim 3, but it seems implied that because all six findings are consistent with each other, and further, because the cited findings are from experiments on different species that all implicate hippocampal replays in memory consolidation (i.e., because they all point towards the same general claim), claim 3 stands on firm epistemic footing.

This seems especially true considering the terms Eichenbaum used to describe/introduce the findings he cites. That is, the terms used to describe the overall evidence in favor of claim 3 include: 'there is substantial evidence for', 'provides compelling evidence that', and 'suggests that' (ibid). These statements convey a great deal of confidence with respect to the epistemic status of claim 3. That is, if there is substantial and compelling evidence suggesting claim 3, it would seem that one has good reasons to believe claim 3 stands on firm epistemic footing. Again, however, the argument Eichenbaum advances to support claim 3 is implicit; the reader may be able to piece together Eichenbaum's argument, perhaps by noting as I do that all six findings seem to point towards the same conclusion and are consistent with each other, however, Eichenbaum himself never explicitly makes this point in the case study. Instead, Eichenbaum simply states there is 'substantial' and 'compelling' evidence for claim 3, cites findings from six interdisciplinary studies on memory organization, and states that they all suggest hippocampal replays are involved in memory consolidation (ibid, 83).

In sum, Eichenbaum shows support for claim 3 by stating there is compelling and substantial evidence for it by citing six findings from different scientific research papers on human and non-human animals, and stating what they, together, suggest for memory organization (which is claim 3, according to Eichenbaum). I will now analyze the work done by Eichenbaum to support claim 5.

The main claim Eichenbaum defends in the case study (i.e., claim 5) is that the fundamental role of the hippocampus is to create a spatio-temporal framework for the organization of memories (ibid, 81). Eichenbaum begins by stating there is 'considerable evidence' implicating the hippocampus in *spatial representation*; here he cites McNaughton et al. (1996) and Shapiro et al.

(2006) (Eichenbaum 2013, 86). However, he notes claims 1-4 implicate the hippocampus in *temporal representation*. Eichenbaum asks how the two streams of research can be reconciled? Does the hippocampus function to represent spatial *or* temporal memory information? Eichenbaum answers this by stating time and place cells are not, in fact, distinct, and that whether their activity is associated with temporal or spatial representations depends on the context of learning within which the human/non-human animal is engaged in (this is claim 5) (ibid). His answer thus *synthesizes* both streams of research on the role of the hippocampus for memory organization. How does Eichenbaum provide support for claim 5? How does he argue claim 5 stands on firm epistemic footing?

Eichenbaum argues claim 5 is supported by two key features of place and time cells. The first feature is that the two cell types share a host of similarities. He cites findings from MacDonald et al. (2011) – a study Eichenbaum cites heavily in support of claim 4 – stating changes to the primary dimension of a task creates a 'partial re-timing' of hippocampal time cells in much the same way that changes to 'spatial cues' causes 'partial re-mapping' of place cells (Eichenbaum 2013, 86). Drawing again on MacDonald et al. (2011, 2012) and newly on Komorowski et al. (2009), Eichenbaum states both time and place cells encode specific events from within their respective frameworks (Eichenbaum 2013, 86). Despite there being multiple sources of temporal information in the brain (here Eichenbaum cites two new papers from and Buhusi and Meck (2005) and Mauk and Buonomano (2004)), Eichenbaum draws on Kraus et al. (2012) to show the main source of spatial and temporal information reaches the hippocampus via the medial entorhinal cortex (Eichenbaum 2013, 86). Eichenbaum continues to highlight a final similarity between the two types of cells. Specifically, he cites McNaughton et al. (1996) who found that place cells seem to be 'pre-configured' and applied to representations of spatial dimensions, just as Dragoi and Tonegawa (2011) (previously cited in support of claim 4) found that time cells "… reflect a sequential configuration of the passage of time, onto which specific events are serially attached" (Eichenbaum 2013, 86).

The second feature of time and place cells Eichenbaum identifies as key to substantiating claim 5 is demonstrated by MacDonald et al. (2011) (previously cited in support of claim 4) who found that time cells tend to incorporate spatial information, and place cells tend to incorporate temporal information (Eichenbaum 2013, 86). Further, a finding by Kraus et al. 2011 (also cited to support claim 4) in experiments where rats traverse through a T-maze task demonstrated that many of the same neuronal ensembles 'fire selectively' at periods of time also display 'spatially specific activity' while the rat traverses the maze (Eichenbaum 2013, 86). Thus, while Eichenbaum does not comment further on these findings, it would seem as though time and place cells are the same type of cell (or at least that they share many of the same properties).

After listing traits shared between time and place cells and demonstrating hippocampal cells represent both spatio-temporal aspects of memory organization, Eichenbaum comments critically on the distinction between the two cells in the literature as being "… a reflection of different experimental designs that emphasize time or place representation, respectively, rather than a fundamental distinction between time cells and place cells" (ibid). Finally, Eichenbaum sums up in his 'concluding remarks' section by claiming

> … *a synthesis* of emerging knowledge *suggests that* the same population of hippocampal neurons encodes both the spatial and temporal regularities of experience, creating a framework that organizes the spatio-temporal context of experience and puts each event into its time and place (ibid, 87, emphases mine).

This is a brief description of what Eichenbaum does in the last two sections of the case study where he introduces claim 5. Like claim 3, Eichenbaum wants to demonstrate that claim 5 is supported by evidence from human and non-human animals. He begins by stating there is 'considerable evidence' implicating the hippocampus in spatial memory; for this, he cites one research paper on rats, one review paper including data on human and non-human animals, and two studies on human animals (ibid, 86). Thus, the 'considerable evidence' consists of four papers on human and non-human animals implicating the hippocampus in spatial memory.

Eichenbaum asks how one stream of research (i.e., *the spatial aspects* of memory in the hippocampus) can be reconciled with literature suggesting the hippocampus is responsible for *temporal aspects* of memory (i.e., claims 1-4 and all the research papers he cites to support those claims)? His answer is to *synthesize* the two streams of research and argue that the role of the hippocampus is to establish a spatio-temporal framework for memory organization (ibid, 86).

According to Eichenbaum, there are two main features of time and place cells suggesting claim 5 (ibid). The first feature is that time cells and place cells share similar properties; here Eichenbaum cites three rat studies, two review papers on human and non-human animals, and one paper on mice, for a total of five papers. While each cited finding refers to a different similarity, generally speaking, all six of these studies show that time and place cells share properties such as 're-timing' and 're-mapping', for example (ibid).

The second feature is that time cells incorporate spatial information and place cells incorporate temporal information; here, Eichenbaum cites a paper on rats and a review paper on human and non-human animals, for a total of two papers. Eichenbaum's argument for claim 5, like claim 3, is *implicit*. That is, Eichenbaum does not explicitly state what the two features entail, but it seems implied that if time and place cells share a host of similarities, and if each type can incorporate spatial and/or temporal information, then the two cell types are not distinct. Since they are not distinct (the implicit argument seems to go), whether their activity is associated with time and/or space is dependent on the context within which the human/non-human animal finds itself (ibid, 86).

To re-iterate, the implicit argument rests on the two features of time and place cells described by Eichenbaum. These two features are supported by a total of seven papers (three of which are review papers) on human and non-human animals. When Eichenbaum reports these features, he does so by briefly stating what each paper found, generally speaking, without introducing them with any terms like 'compelling evidence', 'shows that', 'suggests that', or 'indicates that', etc. (ibid 83-86-87). Eichenbaum only uses terms indicative of corroboration at the beginning of his section entitled 'time and place cells', when claiming there is 'considerable evidence' implicating the hippocampus in spatial representations in non-human animals and 'strong evidence' for it in human animals, and when he introduces claim 5 as an answer to how the spatial and temporal aspects of the hippocampus and memory can be reconciled. With respect to the latter, he argues there are two main features of time and place cells which 'suggest that' the role of the hippocampus is to establish a spatio-temporal framework for the organization of memories (ibid). In his section entitled 'concluding remarks', Eichenbaum states that a 'synthesis of emerging knowledge' – presumably that which he reviewed in the case study – 'suggests that' the same population of hippocampal neurons encode both spatial and temporal aspects of experience, therefore time and place cells are not distinct, and the role of the hippocampus is to create a spatio-temporal framework for memory organization (ibid, 87). Thus, the main terms used to introduce and advance claim 5 are 'suggest that' in Eichenbaum's section entitled 'time cells and place cells' and 'suggests that' in his section entitled 'concluding remarks. What suggests claim 5? Again, Eichenbaum does not make this argument explicit, but it seems that what *suggests* claim 5 is 'compelling', 'substantial', 'considerable' evidence from claims 1-4, along with evidence from a few papers which were not cited in support of claims 1-4 in his section 'time cells and place cells' showing time and place cells are not distinct (ibid, 83, 86-87). Importantly, the term 'suggests' implies that while claim 5 may not be completely true, it is very likely, given that it is well supported by interdisciplinary research on multiple species in the mind-brain sciences investigating memory organization.

This descriptive analysis leads to the question: "does the case study represent a strong case of IC, and, if so, how"? I address these questions in the following sections.

3.2.2 In What Ways 'Memory on Time' is a Prime Example of IC in Practice: The Interdisciplinarity Component

Thus far I have showcased what Eichenbaum does in the case study to argue for claims 1-5 by focusing on the language and terms he used to describe the findings he reported to substantiate claims 1-5. That is, he used terms like 'compelling' and 'substantial evidence' to describe the cited findings, and (sometimes) commented on some of the findings by stating what he thinks they suggest (ibid, 83). I now want to address how and in what ways the case study adheres to the requirements of IC; does the case study represent a strong case of interdisciplinary corroboration?

In answering this, I begin with the interdisciplinarity component of IC insofar as it is relatively straight-forward to define and assess. The interdisciplinarity component held that the papers surveyed in one's review paper must come from different areas of science, which often make use of different methods of experimentation, different protocols, different types of analysis, on different animals, etc., (instead of reviewing studies *only* in psychology, *only* on humans, using *only* questionnaires). The 59 articles mentioned in Eichenbaum's review paper span the fields of neuroscience, cognitive neuroscience, neuropsychology, biology, computational biology, neurobiology, behavioral biology and cognitive science, cognitive neurobiology, psychology, etc.[125] Eichenbaum thus does not constrain his focus to *one sole area of science*, but instead reviews studies from *multiple areas of science* housed under the mind-brain sciences researching memory organization.

Critically, the cited papers not only span different *areas of science*, but different *species*. For instance, the reviewed papers involve experimental research on *humans* (e.g., Ekstrom and Bookheimer, 2007; Lehn et al., 2009; Shimamura, 1990; Spiers et al., 2001; Vargha-Kadem et al., 2007, etc.), *rats* (e.g., Ergorul and Eichenbaum, 2006; Fortin et al., 2002; Karlsson and Frank, 2009; Kesner et al., 2002, etc.), *mice* (e.g., Devito and Eichenbaum, 2011; Dragoi and Tonegawa, 2011Suh et al., 2011), and *monkeys* (e.g., Naya and Suzuki, 2011; Templer and Hampton, 2012) (see figure 4 for a breakdown of the type of species studied by each of the cited papers).

Further, while the cited papers involve work on four main species (i.e., humans, rats, mice, and monkey), each species type is varied. For example, in humans, control groups were compared against humans with global retrograde amnesia (e.g., Vargha-Khadem et al., 2007), patients with Korsakoff's syndrome (e.g., Shimamura, 1990), or with selective bilateral hippocampal lesions (e.g., Mayes et al., 2001). Of the two papers on monkeys (i.e., Naya and Suzuki, 2011; Templer and Hampton, 2012), the former used macaque monkeys and the later rhesus monkeys.

Different *tasks* were used to investigate memory organization such as *serial order for list tasks* (e.g., Shimamura, 1990; Vargha-Khadem et al., 2007), *memory of order for action and stimuli sequences* (e.g., Ekstrom and Bookheimer, 2007), *a taxi cab virtual reality game task*, (Lehn et al., 2009), *retrieving a series of action sequences from movie clips task*, (e.g., Fortin et al., 2002), *overlapping sequences of odors task* (e.g., Agster et al., 2002) *retrieval for overlapping images of faces task* (e.g., Ross et al., 2009), and various *T-maze tasks* (e.g., Ainge et al., 2007; Ferbinteanu and Shapiro, 2003; Frank et al. 2000; Wood et al., 2000, etc.).

---

[125] While Eichenbaum cites 61 studies in his 'References' section, I have chosen to exclude the first two – Aristotle's (1930) 'On Memory and Reminiscence' and Tulving's (1984) 'Precis of Elements of Episodic Memory' – insofar as *their sole citational purpose* is to set up the appropriate background/history of temporal memory research. Accordingly, because I'm interested in the findings cited for the purpose of providing evidence for claims 1-5, they do not fit this description, and I have thus excluded them from any further analysis in this chapter and chapter four.

These tasks involved *varied methodologies*. One common method of analysis in human and non-human animals is the comparison of a control group versus a group with some kind of brain damage. The type of and the extent of brain damage varies from paper to paper. For example, hippocampal lesions/damage in human animals are often naturally occurring (i.e., they are the product of some disease/pathology such as Korsakoff's syndrome and not the product of a researcher directly intervening by creating the lesions themselves). In non-human animals, however, brain damage is a product of direct intervention and is produced in a number of ways. For example, many of the research papers on rats (whether the task is the disambiguation of overlapping odors or the navigation of a T-maze) compare the behavior of a group of control rats to rats with various types of brain damage. In some cases, this damage involved brain lesions which were produced in a number of ways; for instance, Agster et al. (2002) produced human-made lesions by either *injecting infusions of the neurotoxin ibotenic acid* or via the application of a *radiofrequency current*, and Fortin et al. (2002) *surgically created lesions in the hippocampus* (Agster et al. 2002, 5760-5761; Fortin et al. 2002, 458).

Finally, the behavior/results of the tasks are analyzed via a *variety of techniques/tools* including but not limited to *MRI* (e.g., Vargha-Khadem et al., 1997), *fMRI* (e.g., Schendan et al. 2003; Staresina and Davachi, 201), *positron emission tomography* (PET)[126] (e.g., Gill et al., 2011), *electroencephalogram* and *magnetoencephalographic imaging* (EEG and MEG)[127] (e.g., Jensen and Lisman, 2005), and *electrophysiological recording/ single-cell recording* (e.g., Ji and Wilson, 2007; Johnson and Redish, 2007; Karlsson and Frank, 2009; Komorowski et al. 2009; Wood et al., 2000; etc.).

In sum, the case study clearly adheres to/includes the interdisciplinary component of IC. That is, it includes findings from different areas of science, performed on different species, using different experimental tasks and protocols, the results of which were analyzed via different techniques/tools. While the case study has a strong interdisciplinary component, summarizing findings from different areas of science is not sufficient for IC. The cited interdisciplinary

---

[126] PET uses a radioactive substance to detect disease in the body.

[127] EEG is a tool used to detect brain waves and MEG is an imaging technique that allows researchers to detect and record "… the magnetic field associated with electrical activity in the brain" (https://www.merriam-webster.com/dictionary/electroencephalogram, and https://www.merriam-webster.com/dictionary/magnetoencephalography).

research papers and their respective findings must also *corroborate* claims 1-5. In the next section, I analyze the case study in light of the corroboration requirement.

3.2.3 In What Ways 'Memory on Time' is a Prime Example of IC in Practice: The Corroboration Component

In chapter two, I noted a standard definition of corroboration is one where evidence either *confirms* or *supports x*.[128] In the context of the case study, we might say claims 1-5 are corroborated if multiple pieces of evidence confirm or lend support to one or more of them. To confirm *x* is akin to saying the evidence indicates the truth of *x*. To claim evidence lends support to *x*, however, is epistemically weaker. One can claim evidence supports *x* without making the additional claim that *x* is true. That is, in saying *a*, *b*, and *c*, lend support to *x*, we can accept that *x* is not necessarily true, but instead approximately true, likely, unlikely to be false, highly probable, etc.[129] Accordingly, analyzing the case study within the context of IC requires asking what kind of epistemic claims (i.e., either that it confirms or lends support to) claims 1-5 are making.

To answer this, I further investigated the language/terms Eichenbaum used before advancing claims 1-5. The terms used before stating a claim that appear to imply corroboration are: the studies *show* that *x* (a remark made with respect to claim 1 and 2), there is *compelling* evidence for *x* (a remark made with respect to claim 3), these findings *strongly support x* (a remark made with respect to claim 4), and that a synthesis of knowledge on time and place cells *suggests* that *x* (a claim made with respect claim 5 in two separate sections of the case study) (ibid 83,86-87). Since claim 5 is mentioned twice, both in the section of the case study entitled 'time and place cells' and again in the section entitled 'concluding remarks', there are a total of 6 terms indicative of corroboration used to advance a given claim. A text frequency search on NVIVO12 revealed the two most common terms of the six are 'show that' (shortened in fig 2. below as

---

[128] Here I take *x* to represent a claim about the temporal organization of memory, where *x* could thus be any of claims 1-5.

[129] Bechtel and Richardson (1993) define corroboration in probabilistic terms. They claim when different pieces of evidence (using different and independent measures) point to some claim *x*, *the likelihood* of *x* being an error is low. Thus, their epistemic understanding is not one of confirmation, but something weaker along the lines of probability (xxvii). Similarly, Wimsatt's (2007) concept of robustness also presents a case for corroboration as something weaker than confirmation. For Wimsatt, something is robust if it is assessible, detectible, etc., by a variety of different and independent measures; however, to say x is robust is not to say that *x* is true, instead, the likelihood of *x* being robust increases with each piece of evidence supporting it (Wimsatt 2007, 196).

'show') and 'suggest that ' (shortened in fig 2. below as 'suggest'), each of which make up a respective weighted percentage of 33.3%.[130] To show *x* implies *x* is true or has been *confirmed*, but to suggest *x* merely entails it being a *likely possibility*, *approximately true*, etc. While it seems that the kind of corroboration occurring in the case study is undetermined, insofar as the two most frequent terms are tied, the two remaining terms break this tie. That is, there being 'compelling evidence' for *x* (shortened in fig 2. below as 'compelling') and there being evidence that 'strongly supports' *x* (shortened in fig. 2 below as 'supports') each make up a weighted percentage of 16.67%. Neither of these terms entail *x* has been confirmed; there can be compelling evidence for *x* without *x* being true, just as evidence can strongly support *x* without *x* being true. These terms simply imply there are good reasons to believe *x* stands on firm epistemic footing insofar as it is well supported by multiple distinct pieces of evidence; but it need not and does not entail *x* is true.

The terms 'compelling' and 'supports', along with the term 'suggest' support a reading of corroboration understood in terms of evidence *lending support* to *x*. This seems especially true given the term used with respect to the claim 5 (i.e., the main claim Eichenbaum argues for in the case study) is 'suggest that', which implies something weaker than confirmation. While there are aspects both of a confirmation and a lending support reading throughout the case study, ultimately, the terms falling under a 'lending support to' reading of corroboration (i.e., suggest, compelling, and supports) make up a weighted percentage of 66.64%, making it predominantly on the side of 'lending support to' reading of corroboration.[131] Corroboration in the case study should therefore be understood along the lines of 'lending support', whereby the cited findings lend support to claims 1-5 in such a way that one has good reasons to believe the claims are probable, likely, etc.

[130] All weighted percentages were automatically generated via the NVIVO12 text frequency analysis.

[131] If one considers the terms used to with respect to claim 5 (i.e., the term 'suggests' once in Eichenbaum's section entitled 'time cells and place cells' and once in Eichenbaum's 'concluding remarks'), the percentage jumps to 100% in support of a corroboration defined as pieces of evidence lending support to claim *x* (Eichenbaum 2013, 86).

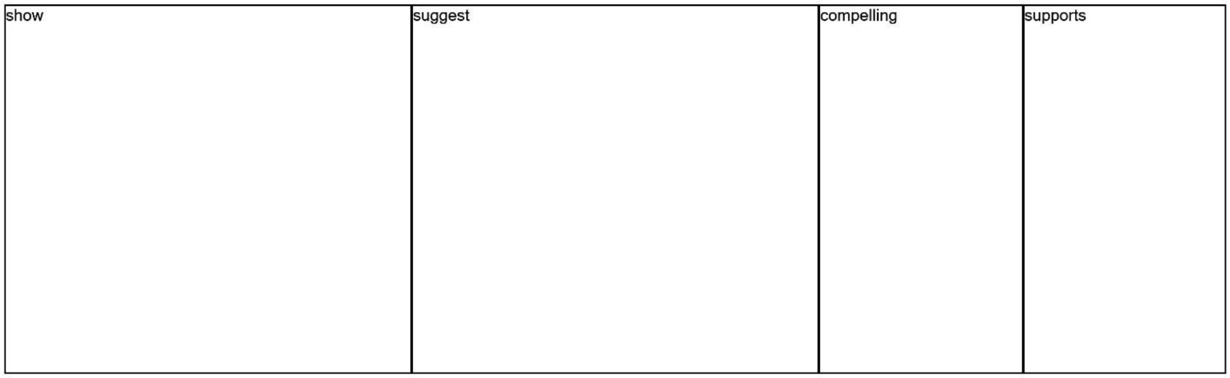

Figure 2 – Tree map produced via text frequency query on NVIVO12 demonstrating the most common terms to the least common terms suggesting corroboration. I've shortened all the phrases down to their key terms (e.g., 'studies show that' to 'show').

I noted in section 3.2.1 the phrase 'suggests that' used by Eichenbaum before advancing claim 5 left open the possibility that claim 5 may not be completely true, even if there was 'substantial' or 'compelling evidence' in favor of it (ibid, 83). Applying the new definition of corroboration to the case study, one might argue claims 1-5 have been corroborated insofar as the findings cited by Eichenbaum all lend support to claims 1-5 and do so in such a way the claims are to be understood as approximately true, probable, likely etc. That is, the cited findings, when combined, appear to lend support to claims 1-5.

Do the cited findings possess any other features that may increase the likelihood of the corroboration of claims 1-5? What characteristics might the cited findings possess such that one can confidently say a reliable claim about memory has been advanced? I answer this by considering two key features of the cited findings in the case study: the journals in which the findings were published, and the number of citations each of the cited papers received. As Seglen (1997) notes, evaluating the quality of scientific research papers is notoriously challenging. Impact factors and citation statistics are frequently called upon in the evaluation of scientific research papers insofar as they are considered to be a "… quantitative and objective indicator directly related to published science" (Seglen 1997, 498).

One important feature of the papers Eichenbaum cites in support of claims 1-5 is that they are published in high-calibre and well-respected journals within the mind-brain science community. To see a breakdown of the cited journals, I've created a chart below indicating the name of the journal, how many studies Eichenbaum cites from each journal, and the journal's impact factor.[132] [133]

Name of Journal	Number of Cited Studies in Case Study	Impact Factor
Academic Press	1	Book
Annual Review of Neuroscience	2	14.675
Science	6	37.205
Neuropsychologia	1	2.888
Cognitive Neuropsychology	1	1.84
Hippocampus	4	3.945
Learning & Memory	3	2.906
Journal of Neuroscience	11	***5.971***
PLOS Biology	1	9.163
Cerebral Cortex	1	6.308
Nature Neuroscience	3	19.912
Frontiers in Behavioral Neuroscience	3	3.104
Neuron	9	***14.318***
Behavioral Brain Research	2	3.002
Proceedings of the National Academy of Sciences of the United States of America	1	9.504
Trends in Neurosciences	2	17.755
Psychological Review	1	7.230
Journal of Experimental Biology	1	3.32
Nature	3	41.577
Current Opinion in Neurology	1	5.307
Frontiers Integrative Neuroscience	1	3.23
Nature Reviews Neuroscience	1	28.880

Figure 3 – Chart including one column for the name of the journal, a second column for the number of studies in the case study that are published in said journal, and a final column stating the impact factor of the journal.

[132] Again, I've only included the 59 studies as they are the ones doing the epistemic work in the case study.

[133] An impact factor is the outcome of a quantitative analysis measuring how frequently articles in a journal have been cited for a given year (https://researchguides.uic.edu/if/impact). An impact factor is represented numerically and is used to determine the rank or status of a given journal; the higher the number, the higher ranked the journal is (ibid). Specifically: "… The calculation is based on a two-year period and involves dividing the number of times articles were cited by the number of articles that are citable" (ibid). Journal Citation Reports (hereafter JCR) are responsible for providing impact factors for journals in science, technology and the social sciences (ibid).

The journals highlighted in yellow in figure 3 represent the three journals with the highest number of cited papers in the case study. Specifically, Eichenbaum cites six studies from *Science*, nine studies from *Neuron,* and eleven studies from *the Journal of Neuroscience*, for a total amount of 26 studies out of 59 (making up almost half of the cited studies at 44%).[134] All three of these journals are well-established within the mind-brain science community. *Science,* founded in 1880, claims to publish only 'the very best' research articles in science and those articles "… consistently rank among the most cited in the world" (https://www.sciencemag.org/about/about-science-aaas?_ga=2.10191605.1947000741.1556017401-385344628.1554129536). Indeed, out of all the journals listed above, *Science* has the second highest impact factor rating (see my fig. 3 above).

Neuron is another well-respected, multidisciplinary journal in the mind-brain sciences. Over the last fifteen years, *Neuron*, founded in 1988, has established its place in the scientific community as "… as one of *the most influential* and *relied upon* journals in the field of neuroscience. … Neuron serves as one of *the premier intellectual forums of the entire neuroscience community*" (https://www.cell.com/neuron/aims, emphases mine).

Finally, *the Journal of Neuroscience* claims to advance neuroscientific research by "… publishing and widely disseminating *the highly rigorous research* representative of the breadth of neuroscience …" (http://www.jneurosci.org/content/about-jneurosci, emphasis mine). *The Journal of Neuroscience* prides itself on five main qualities: first, delivering articles of excellent quality and scientific rigor, by ensuring that studies reflect 'unbiased and repeatable experiments' and reporting procedures to ensure progress in neuroscience (ibid). Second, the journal represents the wide-scope and ever-evolving field of neuroscience (ibid). Third, the journal ensures diversity amongst its editorial board (ibid). Fourth, the journal prides itself on its quick and efficient dissemination of peer-reviewed articles by publishing fifty weeks out of the year and is committed to publishing articles as quickly as possible while maintaining scientific rigor and excellence (ibid). Finally, the journal prides itself of its constructive and fair peer-review process whereby: "Decisions are made by our journal editors who are active research scientists in the field" (ibid). In sum, all three of the journals with the highest number of papers

[134] Percentages in this section are rounded up for simplicity.

referenced by Eichenbaum are from internationally recognized, multidisciplinary, cutting-edge, and highly influential journals producing some of the best and most scientifically rigorous and sound articles on the mind-brain.

The impact factors highlighted in turquoise in figure 3 represent the three most highly ranked journals of Eichenbaum's 59 references. These journals are *Science*, *Nature*, and *Nature Reviews Neuroscience*. *Science*, which is also ranked in the top three journals with the highest number of scientific papers referenced by Eichenbaum, is the second highest ranked journal of those included in the case study (in terms of its impact factor). As stated above, *Science* claims to publish and disseminate only the very best research articles in science, making it a high-ranking, peer-reviewed, international journal in the mind-brain sciences (https://www.sciencemag.org/about/about-science-aaas?_ga=2.10191605.1947000741.1556017401-385344628.1554129536).

Nature, which has the highest impact factor of those included in the case study, is an international journal that publishes weekly, producing, "… *the finest peer-reviewed research*" (https://www.nature.com/nature/about, emphasis mine). In their mission statement, *Nature* states that they contribute to advancing scientific knowledge by the "… prompt publication of significant advances in any branch of science … to ensure that the results of science are rapidly disseminated to the public throughout the world …" (ibid). *Nature* is therefore a multidisciplinary, peer-reviewed, and highly ranked journal in the mind-brain sciences producing some of the best research concerning the mind-brain.

The third highest ranking journal according to its impact factor is *Nature Reviews Neuroscience*. In their 'aims and scope' section, *Nature Reviews Neuroscience* claims similarly to provide some of the best peer-reviewed research from a variety of branches in neuroscience (https://www.nature.com/nrn/about).

The journals with the highest number of cited research papers (i.e., *Science, the Journal of Neuroscience*, and *Neuron*) are considered prestigious and high-calibre mind-brain science journals. The papers cited from those journals make up almost 44% of Eichenbaum's references.[135] Of the top-three highest ranking impact factors, 6 papers were cited from *Science*, 3 papers from *Nature*, and 1 paper from *Nature Reviews Neuroscience*. If these papers are added to those from *The Journal of Neuroscience* and *Neuron* (omitting *Science,* as its numbers are already factored into the 44%) we get a total of 30/59 papers, making a total percentage of 50% of the cited papers. Thus, at least half of the studies Eichenbaum cites in support of claims 1-5 are from internationally top-ranked, multidisciplinary journals that purport to disseminate only the best and most scientifically sound articles.

Further, there are other highly ranked journals just shy of the top three. For example, Eichenbaum cites two studies from *Annual Review Neuroscience* (which has an impact factor of 14.75), three studies from *Nature Neuroscience* (which has an impact factor of 19.912), and two studies in *Trends in Neurosciences* (which has an impact factor of 17.755). Adding these papers to those above results in 37/59 papers, making a total percentage of 63%. Accordingly, more than half of Eichenbaum's references are from trustworthy and reputable sources. A final point worth noting is all journals listed in fig. 3 (even those with lower impact factors) are peer-reviewed journals in the mind-brain science.[136]

I inquired above whether the findings cited by Eichenbaum possess features which might increase the likelihood of claims 1-5 being corroborated. The analysis of journal impact factors suggests that if the cited papers and their respective findings are published in reputable and high-calibre journals that allegedly produce the most scientifically sound articles, one might argue this lends more credibility to Eichenbaum's findings, which would not be the case if he were citing findings from random websites or non-peer reviewed journals. This credibility increases the likelihood of the corroboration of claims 1-5 in the case study.

Another feature to consider is the number of citations each individual cited paper has received. To visualize the breakdown of citations per study, I have included a chart (see figure 4) containing the author(s) of each article, the journal in which it was published, the number of citations it received, and the primary species studied.

---

[135] Percentage was rounded up for simplicity.

[136] Peer-reviewed means any submitted article for potential publication must pass a board of reviewers who are knowledgeable in the topic to ensure top quality research.

Author(s)	**Journal**	**# of Citations**	**Primary Species Studied**
1. N/A			
2. N/A			
3. Eichenbaum et al. (2007)	Annual Review Neuroscience	2004	Human and non-human animals
4. Vargha-Khadem (1997)	Science	1849	Humans
5. Shimamura (1990)	Neuropsychologia	649	Humans
6. Mayes et al. (2001)	Cognitive Neuropsychology	146	Humans
7. Spiers et al. (2001)	Hippocampus	198	Humans
8. Ekstrom and Bookheimer (2007)	Learning & Memory	134	Humans
9. Lehn et al. (2009)	Journal of Neuroscience	155	Humans
10. Kumaran and Maguire (2006)	PLOS Biology	236	Humans

11. Tubridy and Davachi (2011)	Cerebral Cortex	110	Humans
12. Staresina and Davachi (2006)	Journal of Neuroscience	280	Humans
13. Fortin et al. (2002)	Nature Neuroscience	740	Rats
14. Kesner et al. (2002)	Behavioral Neuroscience	281	Rats
15. Templer and Hampton (2012)	Hippocampus	21	Mice
16. Devito and Eichenbaum (2012)	Journal of Neuroscience	116	Rhesus Monkeys
17. Ergorul and Eichenbaum (2004)	Learning & Memory	234	Rats
18. Schendan et al. (2003)	Neuron	522	Humans
19. Ergorul and Eichenbaum (2006)	Journal of Neuroscience	66	Rats

20. DeCoteau and Kesner (2000)	Behavioral Neuroscience	126	Rats
21. Fouquet et al. (2010)	Behavioral Brain Research	36	Rats
22. Agster et al. (2002)	Journal of Neuroscience	251	Rats
23. Kumaran and Maguire (2006)	Neuron	109	Humans
24. Ross et al. (2009)	Hippocampus	67	Humans
25. Brown et al. (2010)	Journal of Neuroscience	94	Humans
26. Karlsson and Frank (2009)	Nature Neuroscience	436	Rats
27. Ji and Wilson (2007)	Nature Neuroscience	1115	Rats
28. Jadhav et al. (2012)	Science	410	Rats
29. Johnson and Redish (2007)	Journal of Neuroscience	680	Rats
30. Gelbard-Sagiv et al. (2008)	Science	382	Humans

31. Paz et al. (2010)	Proceedings of the National Academy of Science of the United States of America	74	Humans
32. Jensen and Lisman (2005)	Trends in Neurosciences	369	Humans
33. Mehta et al. (2000)	Neuron	382	Rats
34. Levy (1989)	Academic Press	217	Computational Modelling
35. Wallenstein et al. (1998)	Trends in Neurosciences	518	Computational Modelling – Human and non-human animals
36. Rolls (2010)	Behavioral Brain Research	211	Monkeys
37. Dargoi and Buzsaki (2006)	Neuron	450	Rats
38. Howard et al. (2005)	Psychological Review	250	Computational Modelling
39. McNaughton et al. (1996)	Journal of Experimental Biology	852	Rats

40. Dragoi and Tonegawa (2011)	Nature	380	Mice
41. Wood et al. (2000)	Neuron	709	Rats
42. Frank et al. (2000)	Neuron	573	Rats
43. Ferbinteanu and Shapiro (2003)	Neuron	429	Rats
44. Ainge et al. (2007)	Journal of Neuroscience	114	Rats
45. Shapiro et al. (2006)	Current Opinion in Neurobiology	68	Human and non-human animals
46. Ginther et al. (2011)	Journal of Neuroscience	45	Rats
47. Barnes et al. (1997)	Nature	345	Rats
48. Mann et al. (2007)	Neuron	288	Rats
49. Pastalkova et al. (2008)	Science	733	Rats
50. Kesner et al. (2005)	Behavioral Neuroscience	118	Rats

51. Farovik et al. (2010)	Learning & Memory	111	Rats
52. Suh et al. (2011)	Science	200	Mice
53. MacDonald et al. (2011)	Neuron	593	Rats
54. Yin and Troger (2011)	Frontiers of Integrative Neuroscience	41	Rats/Mice
55. Gill et al. (2011)	Hippocampus	71	Rats
56. Naya and Suzuki (2011)	Science	160	Monkeys
57. Erkstrom et al. (2013)	Nature	1114	Humans
58. Maguire et al. (1997)	Journal of Neuroscience	761	Humans
59. Komorowski et al. (2009)	Journal of Neuroscience	240	Rats
60. Mauk and Buonomano (2004) 61. Buhusi and Meck (2005)	Annual Review Neuroscience Nature Review Neuroscience	793 1618	Human and non-human animals Human and non-human animals

Figure 4 – Chart with one column for the name of the author(s) and year of publication, a second column for the journal in which it was published, a third column for the number of citations each received, and a final column for what kind of species the study (primarily) investigated.

I noted above that Eichenbaum selected research papers from high-calibre journals in the mind-brain sciences, but have those papers themselves had much impact in terms of their citations? An analysis of the 59 studies reveals a total of *24 274* citations for the combined 59 studies, giving the references *a mean of 411 citations*.[137] The three papers with the most citations (highlighted in yellow in fig. 4) are Eichenbaum et al.'s (2007) paper with *2004* citations, Vargha-Khadem's (1997) paper with *1849* citations, and Buhusi and Meck (2005) paper with *1618* citations. On the lower end of the spectrum, the three least cited papers (highlighted in turquoise in fig. 4) are Templer and Hampton's (2012) paper with *21* citations, Fouquet et al.'s (2010) paper with *36* citations, and Yin and Troger's (2011) paper with *41* citations. Despite these low numbers, the majority of the cited papers hit a minimum of one hundred citations, with many of those almost reaching (and some breaking) one thousand citations. As an example of the former, Fortin et al.'s (2002) has *740* citations, and as an example of the latter, Ji and Wilson's (2007) article surpassed 1000 citations with a total of *1115* citations.

Further, the three papers with the most citations were also published in high-calibre journals with high impact factors. For example, Vargha-Khadem's (1997) paper was published in *Science*, which has an impact factor of 37.205, Buhusi and Meck's (2005) paper was published in *Nature Review Neuroscience*, which has an impact factor of 28.880, and finally Eichenbaum et al.'s (2007) paper was published in *Annual Review Neuroscience*, which has an impact factor of 14.675, all of which are high impact factors.

Thus, one might argue a mean of 411 citations for Eichenbaum's cited papers, in addition to their being published in high-calibre and internationally recognized mind-brain journals adds more credibility to claims1-5 (since they are grounded in the reliability of cited findings).

[137] Mean average was rounded down for simplicity.

Furthermore, the cited papers having high citation counts and being published in journals with high impact factors, in addition to the fact that the cited findings are highly interdisciplinary adds another level of credibility to the findings corroborating claims 1-5. That is, a point made by Bechtel and Richardson (1993), Campbell (1966), Hacking (1983), Roskies (2010) Wimsatt (2007), is that when multiple pieces of evidence from different areas of science, using different tools, techniques, and methods of analysis all point towards a general conclusion, then that conclusion is more probable than one reached by one individual piece of evidence or pieces of evidence from the same area of science using the same general protocols and methods of analysis.[138] In section 3.1.1, I noted the ways in which the case study adheres to the interdisciplinary component of IC by including papers on memory organization from different areas of the mind-brain sciences, on different species, using different tasks, and methods of analysis. In section 3.1.2, I noted that claims 1-5 appear to be corroborated by the cited findings, but in a way that claims 1-5 are understood as likely, probable, etc., but not necessarily full-stop true. Further, because each the cited findings are from papers published in high-calibre mind-brain science journals with high impact factors and citations, the findings can be considered reliable pieces of information. Therefore, the case study appears to represent a prime example of IC in practice insofar as the interdisciplinary papers/findings cited in the case study are considered reliable pieces of information, which, when taken together, corroborate claims 1-5 in such a way that makes them very likely/probable.

In sum, I began my analysis in this chapter with a general understanding of how *Trends* review papers work based off of my research presented in chapter 2 (i.e., they contain a *summary*, *explanation(s)*, which can take the form of a *synthesis*, and directions for future research). My analysis in this chapter was based off of this information and was guided by cursory criteria such as the following questions: "what is the aim of Eichenbaum's review paper", "what claims does Eichenbaum argue for", "how does Eichenbaum show support for those claims", "are Eichenbaum's arguments implicit or explicit", "and "does the case study represent a strong case of IC, and if so, how"?

[138] Each author articulates this point somewhat differently. For those differences, refer to chapter two of this dissertation.

I stated that the main aim of the case study was to establish the role of the hippocampus as creating a spatio-temporal framework for the organization of memories. I identified a total of five synthesized claims advanced in the case study; claims 1-4 were implicit, and claim 5 was explicitly flagged by Eichenbaum. Eichenbaum supported claims 1-5 by citing findings from papers suggesting the claim in question. For example, claim 3 stated hippocampal neuronal ensembles replay sequences of events in memory by encoding and retrieving the order of events (ibid, 83). To support this claim, Eichenbaum cited findings from human and non-human animals studies which all generally implicated hippocampal replays in memory consolidation. That is, while there were minor discrepancies in the findings – for example, intervening with hippocampal replays impairs the *rats memory for maze learning* (Jadhav et al. 2012) and that hippocampal replays are observed in *humans during the recall of scenes from a movie* (Gelbard-Sagiv et al. 2008) – all six findings observed hippocampal replays during memory and learning tasks (Eichenbaum 2013, 83). Thus, the combined interdisciplinary findings suggest claim 3. Further, by describing and introducing findings using terms like 'substantial' and 'compelling evidence', Eichenbaum gives the reader the impression that there is good reason to take claim 3 (and by proxy, claims 1, 2, 4, and 5) as corroborated by the interdisciplinary findings he cites.

Critically, the arguments for the corroboration of claims 1-5 are implicit. That is, Eichenbaum never explicitly spells out the arguments for his claims; instead, he cites a number of studies which he takes as representing 'substantial' or 'compelling' evidence for a claim, and states that taken together, they suggest a given claim 1-5 (ibid). Thus, his arguments for the corroboration of claims 1-5 are implied by the terms he uses to describe the findings he cites in support of his claims.

My goal in this chapter was to get a deeper understanding of how review papers work to corroborate claims about the mind-brain and to assess the extent to which the claims in the case study were corroborated. In section 3.1.2, I argue that *prima facie*, the case study appears to represent a strong case of IC, insofar as claims 1-5 are supported by reliable and interdisciplinary findings.

However, analyzing the case study – especially the implicit arguments for claims 1-5 – prompted a number of questions and concerns about the case study representing a strong case of IC. That is, part of Eichenbaum's support for claims 1-5 is that there is 'substantial' and 'compelling' evidence or that there are 'numerous studies' corroborating a given claim (ibid, 83). This introduced worries about what Eichenbaum meant by 'substantial' or 'numerous studies' (ibid). Eichenbaum states there are 'numerous studies' that show forward and reverse hippocampal replays in animals (ibid). However, he only cites four papers in support of this claim; are the findings from four paper sufficient for representing 'numerous studies' in this instance (ibid)?

In Eichenbaum's introduction, entitled 'the temporal organization of memory', Eichenbaum states he will relate the cited findings in such a way that shows how they all combine to corroborate claim 5 (ibid, 81). This 'relating' thus involves showing *how* the findings fit together in a way that corroborates claim 5 (ibid). The problem is Eichenbaum only shows how they fit together implicitly; he never explicitly explains how all of the findings fit together to support a given claim. Implicitly, it seems that Eichenbaum is arguing that because all the findings point towards a given claim, and because those findings are consistent with each other, they corroborate claims 1-5. This, however, introduced worries about the term 'consistency'; how is 'consistency' defined, how is it measured, and further, is this even what Eichenbaum meant to imply?

Worries about consistency and 'substantial evidence' and its relation to IC prompted questions about anomalous findings (ibid). That is, I began to wonder about Eichenbaum's selection criteria; why were there next to no anomalies or contradictory findings presented in the case study? Of the few anomalous findings that are included in the case study, how do they fit together with the rest of findings which are seemingly all consistent with each other; it was not immediately clear to me that they did fit together with the rest of the findings in this way.

In the following chapter, I elaborate on these problems and their implications for IC in the mind-brain sciences; I then use these questions and concerns to advance a methodology for evaluating review papers with an eye towards assessing their claims for corroboration in the second half of chapter four.

Chapter 4

4 Introduction

In chapter three, I conducted part-one of my analysis of a *Trends* review paper 'Memory on Time' by Eichenbaum (2013). My aim in chapter three was to focus on the work Eichenbaum did to support what I identified as claims 1-5.[139] I noted this work consisted of citing interdisciplinary findings, which, taken together, corroborated one of the five claims I identified. This corroboration was implied by the terms Eichenbaum used to introduce findings and to advance a given claim. For example, with respect to claim 3, when he introduced relevant findings, he prefaced their description by stating 'there is substantial evidence' for hippocampal replays representing the order of past events, and that 'numerous studies' have shown both forward and reverse hippocampal replays in memory consolidating (Eichenbaum 2013, 83). After all the findings for a given claim have been cited, Eichenbaum stated 'these findings provide compelling evidence' for claim 3; this language suggests there are good reasons to think claim 3 has been corroborated (ibid). An analysis of the language used in the case study revealed that the kind of corroboration at play in the case study is one where a given claim is not full-stop true, but instead approximately true, probable, likely, etc. Contextualizing this analysis within the context of IC revealed that *prima facie*, the case study appeared to represent a strong example of IC; that is, interdisciplinary findings corroborate claims 1-5 in such a way that those claims are understood to be approximately true, probable, likely, etc.

However, the process of conducting this part of my analysis elicited a number of questions and concerns such as "what counts as substantial evidence", "what does it mean for evidence to be consistent with each other", and "how are anomalies factored into the 'substantial evidence' that appears to be consistent with each other" (ibid). These preliminary questions suggested the support provided for claims 1-5 in the case study is much more questionable than it initially appeared in chapter 3. Thus, my analysis in chapter 3 using my understanding of how a *Trends* review works from chapter two and my cursory criteria prompted a number of questions and concerns I deem important to have answered in order to better assess corroboration in the case study.

[139] I noted that the claims 1-4 were implicit in the case study, and claim 5 was explicit in that it is the main claim Eichenbaum argues for and because he explicitly flags it as a synthesis.

The aim of this chapter is to continue my analysis by focusing on what is absent in the case study and to advance a methodology for assessing review papers for corroboration. I begin in section 4.1 by addressing the question of "what counts as substantial evidence", what it means for findings to be consistent with each other in section 4.1.2, and how anomalous findings feature in the case study in 4.1.3. At the end of each, I note how these questions complicate evaluating the extent to which claims 1-5 have been corroborated in the case study. As I conducted this part of the analysis, it struck me that another factor missing in the case study was an evaluation of the methods used by the papers Eichenbaum cites. I discuss this problem in section 4.1.4, and show how this point, together with the other problems I mention, demonstrate the corroboration of claims 1-5 is uncertain, at best. From this, I call attention to the lessons learned from my two-part analysis of the case study in section 4.2. In section 4.3, I use the problems and lessons from my analysis to advance a methodology for evaluating review papers with an eye for corroboration. I close, in section 4.4., with some questions and directions for future research.

4.1 Substantial Evidence, Numerous Studies, Suggesting, and IC

As I noted in chapter 3, the arguments for claims 1-5 are implicit. That is, Eichenbaum never explicitly advanced an argument to support them; he never says something akin to "here are the findings that support claim *x*, this is how they fit together", and/or "the findings corroborate claim *x* because they have features *y* and *z*", etc. Instead, Eichenbaum shows support for claims 1-5 by stating evidence that appears to point towards or suggest a given claim, and notes that there is 'substantial' and 'compelling' evidence to supporting it (ibid). Using such language indicates to the reader that there are good reasons to believe claims 1-5; that is, if there is *substantial* and *compelling* evidence suggesting claims 1-5, it would seem that the claims stand on firm epistemic footing.

This logic might hold if it were the case that there is substantial and compelling evidence supporting *x*. However, as I conducted my analysis in chapter three focusing on how Eichenbaum supports claims 1-5, I began to wonder about what counted as substantial evidence. Throughout the case study, Eichenbaum uses terms like 'substantial evidence', 'compelling evidence', 'considerable evidence' to introduce findings that allegedly support a given claim (ibid 83, 86). These terms convey a powerful epistemic message; that is, a given claim is not simply supported by a couple scientific findings, but it is backed by *substantial empirical* *evidence*. This conveys to the reader that the evidence behind a given claim is strong, reliable, and significant.

However, one might ask: "what does substantial evidence mean"? "Is 'substantial evidence' defined and meant to be understood qualitatively or quantitatively" (ibid, 83). If the former, then some definition would be required for the reader to understand what and how Eichenbaum understands and uses the term in the case study. If the latter, how much evidence is required for 'substantial evidence'; how many scientific findings are required to be considered 'substantial'. One might also wonder how other key terms like 'compelling evidence' and 'numerous studies' are defined (ibid). One might ask: "to what extent and how exactly is the evidence compelling"? "Can evidence be compelling without it also being substantial?" "How many studies must a review paper cite to constitute or merit the term 'numerous studies"?

Consider claim 3 again. In showing support for claim 3 – i.e., that hippocampal neurons replay sequences of events by encoding and retrieving the order of events – Eichenbaum states there is 'substantial evidence' that hippocampal neuronal ensembles represent the organization of past experiences (ibid). The 'substantial evidence' for claim 3 consists of *six findings* from research papers on memory organization. One might wonder whether six studies is sufficient to count as *substantial evidence*. Relatedly, after advancing this statement, Eichenbaum goes on to state there are 'numerous studies' within this 'substantial evidence' that demonstrate forward and reverse hippocampal replays in rats (ibid). The 'numerous studies' consist of four findings from Jadhav et al. (2012), Ji and Wilson (2007), Johnson and Redish (2007), and Karlsson and Frank (2009) on rats. Again, one might wonder whether four papers sufficiently represents the term 'numerous studies' (Eichenbaum 2013, 83). If four papers do sufficiently represent the term 'numerous studies', some justification seems to be required.

Further, a big part of how Eichenbaum implicitly argues for claims 1-5 is by stating there is 'substantial' and 'compelling' evidence *suggesting* a given claim (ibid). However, Eichenbaum never explicitly explains *how* the evidences suggests what it does. Instead the argument is implicit, and the reader is required to try and piece together Eichenbaum's thought process. I noted in chapter three that Eichenbaum seems to imply that because there is substantial evidence that is consistent with each other, and because that evidence seems to point towards claim 3, that claim 3 stands on firm epistemic footing. Claims 1-5 are suggested by the substantial cited evidence. However, if it turns out that the evidence is not substantial, perhaps the support for claim 3 is weaker than it appears when prefaced with the term 'substantial evidence'; perhaps the evidence does not suggest what Eichenbaum claims it does (ibid).

It is important to flesh out these key terms as the corroboration of claims 1-5 hinges on them. If there is no 'substantial evidence', if the evidence is not, in, fact, compelling, then the claims may not be as corroborated as they appear on a first reading/analysis. Thus, because Eichenbaum provides no definition for terms like 'substantial' and 'compelling evidence', and because he does not provide an explanation as to how the evidence suggests what he claims it does, it is unclear how/whether claim 3 – and claims 1, 2, 4, and 5 by proxy – are corroborated (ibid 81-87).

4.1.1 Consistency and IC

Eichenbaum's arguments for claims 1-5 are implicit. In chapter three, I noted my reading of the case study implied that another factor (in addition to there being substantial evidence suggesting a given claim) implying corroboration is the fact that the cited findings are *consistent with each other*. In the example I use in chapter three, I note that all six findings cited in support of claim 3 are consistent with each other insofar as they do not contradict each other and generally seem to suggest that hippocampal replays are an essential part of memory for past events (i.e., they all seem to suggest a similar role of hippocampal replays and memory organization). Importantly, with respect to claim 3, Eichenbaum does not explicitly make this point nor does he state that the findings are consistent with each other; this is my interpretation of how I think his argument for claim 3 would go.

Eichenbaum does, however, in other parts of the case study use the terms 'consistent with' and 'complemented by' with respect to claim 1 and claim 4; this is part of what led me to believe that the consistency of the findings was crucial to establishing the corroboration of all the claims I identify in the case study (ibid 82-84, 86). I review claim 1 and claim 4 in turn to make this connection concise.

Claim 1 stated the hippocampus is responsible for the memory of the order of events (ibid, 81, 83). Eichenbaum begins by reviewing findings implicating the hippocampus for the memory of the order of events in human animals; he cites eight studies on human animals suggesting claim 1. Eichenbaum then stated these findings from experiments on human animals are '*complemented by*' studies on non-human animals (e.g., rats, mice, and monkeys) (ibid, 82, emphasis mine). Eichenbaum does not explain what he means by evidence from human studies is 'complemented by' evidence from animals studies (ibid). One might surmise that Eichenbaum means that the studies on animals enhance the argument for claim 1 in the following way: *not only* is there evidence for the hippocampus being involved for memory of the order of events in *human animals*, but there is *also* evidence from a number of studies on different *non-human animals* suggesting the same general claim. The evidence from animals studies thus enhances the evidence by providing further support for the claim that it allegedly suggests. This is my reading of what Eichenbaum means by 'complementary' and how it factors into his argument for claim 1; importantly, Eichenbaum does not explain what he means by complementary here nor how it functions to corroborate claim 1.

Eichenbaum states that findings are 'consistent with' each other multiple times while providing support for claim 4 (ibid, 83-84, 86). Claim 4 states hippocampal time cells possess the properties of context cells, allowing them to bridge the order of events (ibid, 86). Eichenbaum provides support for claim 4 in the same way he does for the other claims: by citing findings from different areas of the mind-brain sciences, on different species, using different tasks and methods of analysis, and relating them in such a way that he *synthesizes* them to show that they, taken together, suggest claim 4. This work consists primarily of stating findings by introducing them via terms like there is 'substantial evidence' showing that *x* (ibid). However, Eichenbaum does more than this with claim 4; here, he states multiple times that various pieces of evidence not only suggest a given claim, but that they are *consistent with each other* (ibid, 82-84, 86). On my reading, this implies that claims 1-5 hinge on the cited evidence being consistent with each other. It is thus important to understand what is meant by consistency. If no explanation of consistency is provided, it is difficult to assess whether the evidence is, in fact, consistent, and consequently, it is challenging to evaluate whether a given claim as been corroborated.

For example, with respect to claim 4, Eichenbaum cites a finding by McNaughton et al. (1996) demonstrating: "…the existence of pre-configured contextual representations to which specific behavioral events were linked to compose spatial memory" (Eichenbaum 2013, 84). He then relates this finding to another finding by Dragoi and Tonegawa (2011) stating: "*Consistent with this proposal* … hippocampal ensembles tend to 'pre-play' sequences that will subsequently reflect serial place cell firings in a new environment, which suggests that pre-existing ensemble structure contributes to new sequence coding" (Eichenbaum 2013, 84, emphasis mine). What is meant by consistency in this example? Unfortunately, nowhere in the case study does Eichenbaum provide an explanation of what he means by evidence being consistent with each other, nor does he provide a measure of consistency (i.e., how is consistency measured in the case study – qualitatively or quantitatively?).

One might speculate that what Eichenbaum has in mind is that evidence is consistent with each other to the extent that it does not contradict each other, and that all findings consistently demonstrate a similar pattern/trend with respect to memory organization. Consider a couple of the findings cited for claim 4. McNaughton et al. (1996) found 'pre-configured contextual representations' during memory tasks, and Dragoi and Tonegawa (2011) found a 'pre-existing ensemble structure' which replayed certain sequences during memory tasks (Eichenbaum 2013, 84). Both findings demonstrate a 'pre-existing' neuronal ensemble critical for memory organization. Thus, one could argue the findings are consistent in that their findings both revealed a similar trend (i.e., a pre-existing neuronal structure that replays during memory tasks), and the findings do not contradict each other. Perhaps this is what Eichenbaum meant by consistency and its relation to the corroboration of claims 1-5.

The point I want to emphasize is that because Eichenbaum provides no explanation nor measure of consistency in the case study, it is difficult to ascertain whether the evidence is consistent, and accordingly, it is difficult to evaluate the whether a given claim has been corroborated. Since Eichenbaum's arguments for claims 1-5 are implicit, and because they appear to hinge on terms like 'consistency' and 'substantial evidence', the fact that neither term is explained is problematic for evaluating the extent to which claims 1-5 are corroborated in the case study.

Furthermore, in chapter one of my dissertation, I flagged a potential problem with the consistency of evidence entailing the corroboration of a claim. That is, in discussing a potential way of fleshing out the concept of convergent evidence for Nagel (1949, 1979) and Roskies' (2010) accounts of fitting together and strong explanations, I advanced a conceptualization of the convergence of evidence via the notion of consistency wherein findings do not contradict each other and point towards a similar but not identical conclusion. I argued, however, that understanding convergence via consistency runs the risk of trivializing the claims which it allegedly corroborates. That is, by requiring the consistency of evidence in corroborating a claim, all that is required is a general compatibility between different findings; this seems all too easy to achieve. An author could easily scan the literature and choose to cite only those findings consistent with each other and with the main claim(s) they want to argue for and ignore contradictory or anomalous findings. Thus, *consistency alone* is insufficient for ensuring strong epistemic claims; that evidence be consistent is one thing, but it is also important that the cited evidence to be mutually supportive and consistent in a way that coincidence, errors, and alternative explanations can be ruled out. Thus, without an explanation and measure of consistency provided in the case study, evaluating whether a given claim has been corroborated is difficult.

This leads to the next problem I want to address, which is the general lack of anomalous findings and the lack of substantial explanations as to how the few anomalous findings that do feature in the case study fit together with the rest of the cited findings.

4.1.2 Anomalies and IC

During the analysis I conducted in chapter three, it struck me that the majority of the cited findings all suggested a (general) role of the hippocampus in memory organization. I began to wonder why the cited findings were primarily in favor of claims 1-5. That is, of the 59 findings Eichenbaum cites, only a couple of them are anomalous; why was no contradictory evidence included in the case study? Why were there so few anomalies cited in the case study? Are adequate explanations provided for why the cited anomalies can be dismissed? If claims 1-5 depend on the cited evidence being consistent with each other, then anomalous findings have the potential to complicate a clear case of consistent evidence, which would consequently undermine the corroboration of claims 1-5.

One clear anomaly is cited in Eichenbaum's discussion of claim 4. Recall claim 4 states hippocampal time cells possess the properties of context cells, allowing them to bridge the order of events. After summarizing allegedly consistent findings from human and non-human animals which all suggest "…a temporal context signal containing common elements that encode the common temporal structure and distinct elements that encode the differences in memory", Eichenbaum flags a potential complication via a finding by Naya and Suzuki (2011) (Eichenbaum 2013, 86). Naya and Suzuki's (2011) research on monkeys found:

> …hippocampal neurons in monkeys performing a visual paired association task, where the stimuli are separated in time, provide a temporal signal during the interstimulus interval. Neurons in the higher order visual areas (inferotemporal and perirhinal cortex) encode the visual cues, and entorhinal neurons encode both the visual cues and time, whereas the hippocampal neuronal ensemble signals only the temporal organization, without specificity of items held in memory (Eichenbaum 2013, 86).

This complicates the narrative Eichenbaum delivers in support of claim 4 insofar as it is inconsistent with the literature he cites on rats (ibid). Specifically, Eichenbaum notes that it is unclear why studies on rats have consistently demonstrated the existence of *memory-specific time cells*, but Naya and Suzuki's (2011) work on monkeys demonstrated "…the same time cell sequences for *all* memories" (Eichenbaum 2013, emphasis mine).

How does Eichenbaum explain this potentially anomalous finding? Eichenbaum briefly explains the discrepancy between the findings from studies on rats and the findings on monkeys in this example as the product of different experimental protocols (ibid). That is, in the experiments on rats, association tasks involved a small amount of repetitive sequence exposures, whereas the monkeys performed a task which, conversely, involved a large number of sequence exposures where those sequences were rarely repeated (ibid). Eichenbaum therefore reasons: "… *possibly*, after many exposures, stimulus sequences are embodied in the regularities of the temporal structure, whereas when diverse stimulus sequences seldom repeat, only the reliable temporal structure of the task is represented by the hippocampus" (ibid, emphasis mine). This explanation is unsatisfactory for two reasons. First, his reasoning that perhaps the discrepancy is the multiple exposures that the monkeys undergo is only a possibility. His use of the term 'possibly' implies only that this *might* explain the differences between the studies; it is also possible that if the tasks were the same as in the rat paradigm, the findings from the monkey studies might still be incompatible with the findings from rat studies, thus undermining or complicating the corroboration of claim 4 (ibid). Secondly, offloading the inconsistency between findings in rats and on monkeys as the product of differences in experimental protocols is unsatisfying given that the majority of the case study comprises findings from different areas of science, studying different animal species, via different tasks, methods, and experimental protocols. The cited findings for claims 1-5 all involve different tasks and protocols; nonetheless, they are deemed consistent with each other, and further, the fact that they employ different experimental protocols is seen as bolstering the corroboration of claims 1-5. Why is it the case that different experimental protocols are considered a virtue when they are consistent with and point towards claims 1-5, but seen as problematic when they yield a result inconsistent with the rest of the cited findings?

The example above represents one anomaly Eichenbaum introduces in the case study; this appears to be the only anomalous finding in the case study. One might argue, however, that the research implicating the hippocampus in spatial representations can be read as anomalous research insofar as it suggests a different role for the hippocampus than that suggested by claims 1-4 (which all suggest a temporal role). How does Eichenbaum handle this anomalous data?

Eichenbaum handles the discrepancy between research suggesting a spatial role for the hippocampus and memory and research suggesting a temporal role of the hippocampus and memory by *synthesizing* the two streams of research by advancing claim 5. His synthesis is grounded in two features of time and place cells. He cites evidence from findings cited in support of claim four, as well as previously uncited findings to illustrate that time and place cells are not distinct cell types. Importantly, there is no critical analysis of the cited findings supporting claim 5. In demonstrating the time and place cells are similar, Eichenbaum simply lists findings indicating similarities between time and place cells.

One interesting and potentially anomalous findings is introduced during this process. That is, two different review papers by Mauk and Buonomano (2004) and Buhusi and Meck (2005) found that there are *numerous origins of temporal signals in the brain* (Eichenbaum 2013, 86). Immediately after introducing this finding, Eichenbaum dismisses this potential complication by citing evidence from a poster presentation by Kraus et al. (2012) who found time and place cells are similar in that both receive information from the medial entorhinal cortex (hereafter MEC) (Eichenbaum 2013, 86). Eichenbaum does not, however, explain *how* this finding dismisses the potential anomaly introduced by the finding from Mauk and Buonomano (2004) and Buhusi and Meck (2005). In fact, Eichenbaum does not even flag this finding as an anomaly or source of worry at all. That is, Eichenbaum does not introduce it as an anomalous finding in the way that he does with the Naya and Suzuki (2011) finding (Eichenbaum 2013, 86).

On my reading, however, the finding from the two review papers does constitute an anomalous finding; at least, it constitutes an anomalous finding in the way the Eichenbaum presents the findings from their respective papers. That is, according to Eichenbaum, Mauk and Buonomano (2004) and Buhusi and Meck (2005) review papers state: "…there are multiple origins of *temporal signals* in the brain…" (Eichenbaum 2013, 86, emphasis mine). Since this comment pertains solely to temporal signals, this finding would seem to pertain to time cells (but not place cells) and the various sources from which they receive information. The Kraus et al. (2012) finding stated immediately after introducing the Mauk and Buonomano (2004) and Buhusi and Meck (2005) reviews found that a *key source of information* for time and place cells is the MEC (Eichenbaum 2013, 86).

However, if it is the case that time cells receive information from the MEC in *addition to multiple other sources* – as the Mauk and Buonomano (2004) and Buhusi and Meck (2005) finding seems to suggest – but place cells *only* receive information from the MEC, this would demonstrate that time and place cells are not distinct (or at least put pressure on the statement they that are distinct). Another possibility is that while both receive information from the MEC, time and place cells may also receive information from other *different sources in the brain*; this could also undermine the claim that the two place cells are distinct if the other sources of information for time and place cells were different for each cell type.[140] Due to the fact that Eichenbaum does not explicitly address this complication, the possibility remains open that the two cells types are not as similar as the case study makes them out to be.

---

[140] It is unclear why Eichenbaum introduced the Mauk and Buonomano (2004) and Buhusi and Meck (2005) finding in the way that he does (i.e., with the strict focus on the temporal signal), especially since both papers are in agreement that time and space are inextricably linked. To quote Buhusi and Meck (2005): COINCIDENCE DETECTION is also used by bats, owls and frogs to form an accurate, topographic representation of space from INTERAURAL TIME DIFFERENCES. For these species, *telling space is telling time*" (755, italicized emphasis mine).

Interestingly, neither cited paper specifically stated: "…there are multiple origins of temporal signals in the brain…"; this is Eichenbaum's interpretation of their review papers (Eichenbaum 2013, 86). Put differently, this is the aspect Eichenbaum focused on in their papers and decided to include in his review. However, the Mauk and Buonomano (2004) and Buhusi and Meck (2005) review papers both comprise a review of literature in human and non-human animals on temporal organization that, on my reading, complicates claim 5. Both papers distinguish between multiple types of memory organization. For example, Buhusi and Meck (2005) differentiate between circadian timing, interval timing, and millisecond timing (756). Buhusi and Meck proceed to cite different results from imaging, lesions studies and electrophysiological studies investigating different aspects of timing. Unlike Eichenbaum's review paper, which highlights and focuses almost exclusively on hippocampal activity in human and non-human animals during various timing tasks, Buhusi and Meck (2005) – as well as Mauk and Buonomano (2004) – found and highlighted activation of a *whole host of brain areas/networks*. For example, electrophysiological studies (e.g., Lewis and Miall, 2003) demonstrated activation of the thalamo-cortical striatal circuits, which include the basal ganglia, the prefrontal cortex (hereafter the PFC) and the posterior parietal cortex (hereafter PPC), during interval timing tasks (Buhusi and Meck 2005, 760). Imaging studies focusing on millisecond and interval timing (e.g., Macar et al., 2002 and Lewis and Miall, 2003) found activation of the basal ganglia and the cerebellum, as well as a majority of the supplementary motor area (hereafter SMA), the dorsolateral prefrontal cortex (hereafter DLPFC), anterior cingulate cortex and the right parietal cortex (Buhusi and Meck 2005, 761). In their concluding remarks, Buhusi and Meck (2005) state:

> In mammals, the circadian clock that drives metabolic and behavioral rhythms is located in the SCN. Another timer, which is responsible for automatic motor control in the millisecond range, relies on the cerebellum. In contrast to these relatively localized timing mechanisms, a general-purpose, flexible, cognitively controlled timer that operates in the seconds-to-minutes range involves the activation of *a network of brain areas* that form part of the thalamo-cortico-striatal circuits, notably the SMA, the PFC, and the PPC. … As these areas are involved in several cognitive phenomena, it is likely that *this circuit is not limited to temporal*

processing, but is also involved in other processes, such as the estimation of quantity and numerosity (763, emphases mine).[141]

It is interesting to me that while Eichenbaum cited both papers in his review, he chose to selectively focus on one vague finding, i.e., that temporal processing has multiple origins in the brain, and then immediately dismisses it, without justification, by the fact that Kraus et al. (2012) found that both time and place cells receive information from the MEC (Eichenbaum 2013, 86). The fact that both the Mauk and Buonomano (2004) and Buhusi and Meck (2005) highlight a multitude of brain areas/networks involved in various types of timing tasks was not included in the review is problematic. That is, while these findings do not necessarily falsify claim 5, they complicate the clear picture painted by Eichenbaum in the case study whereby the hippocampus is the main brain structure responsible for the realization of spatio-temporal aspects of memory.

Eichenbaum introduces few anomalous/contradictory findings throughout the case study. This prompts a number of questions about his selection criteria. For example, why were certain findings included while others were omitted? Were many anomalous or contradictory findings encountered during his review process; if so, how were those handled? It could be the case that the reason why so few anomalies are included in the case study is simply because there were not many in the literature to begin with. However, because Eichenbaum's selection criteria is not know to the reader, one cannot know this for certain.

Further, many of the cited papers include more than one finding; what criteria did Eichenbaum use to decide which findings to focus on and which to omit when reading through individual research papers. For example, why did Eichenbaum choose to exclude findings which involve the activation of structures/networks other than the hippocampus (such as the many findings from the Buhusi and Meck (2005) paper implicating a multitude of brain regions/networks))? Again, because Eichenbaum's selection criteria and review writing strategy is unknown to the reader, one cannot know with certainty why certain findings were included and others omitted.

[141] Mauk and Buonomano (2004) make a similar point. They note that while there is still debate about whether temporal processing is localizable or distributed, it seems that if any neural circuit can "… intrinsically process temporal information, then virtually any circuit could be involved, and the location of temporal processing would depend on the nature and modality for the task at hand" (334).

I stated above that the lack of definitions for terms like consistency was problematic for evaluating the corroboration of claims 1-5. I argued that we want evidence to be consistent in a way that coincidence, errors, and alternative explanations can be ruled out; because few anomalous findings are introduced in the case study, and because the explanations of those few anomalies are unsatisfactory, it is a) not clear that the evidence is consistent, and b) it is not clear that coincidence, errors, and alternative explanations have been ruled out. Adding this to the complications arising from the lack of definitions for terms like 'substantial', 'compelling evidence', and 'suggests', it is unclear how the findings in the case study fit together to corroborate claims 1-5.

4.1.3 Evaluation of Methods, Citation-Based Bibliometric Statistics, and IC

As I conducted this part of the analysis (i.e., what important aspects are missing in the case study in section 4.1., 4.1.2., and 4.1.3), it struck me that another missing factor was an evaluation of the methods used by the papers Eichenbaum cites. That is, as I analyzed claim 1 and claim 4 in greater detail in section 4.1.2 of this chapter, I noticed that Eichenbaum treated claims 2, 3 and 5 slightly differently than claims 1 and 4. Eichenbaum spends more time elaborating on the findings cited in support of claims 1 and 4. Eichenbaum's showing support for claims 2, 3, and 5 consisted primarily of Eichenbaum citing findings and immediately stating what they, taken together, suggest. Eichenbaum does the same thing for claim 1 and 4, but also includes brief descriptions of the tasks used in some of the papers that he cites.

Importantly, while Eichenbaum sometimes describes the tasks involved in the papers he cites, he does not critically evaluate the methods and protocols used to yield the findings he cites in the case study. For example, consider his description of findings from Kumaran and Maguire (2006) and Lehn et al.'s (2009) research on hippocampal activation for memory for the order of events in humans:

> Another study showed striking selective activation of the hippocampus when humans reconstructed the correct ordering of scenes previously viewed in a movie clip, as compared to when they inferred the logical order of related scenes from the same movie … Using a different approach that exploited the well-known observation of enhanced neural response to novel stimuli, Kumaran and Maguires [10] demonstrated hippocampal activation when subjects were presented with items out of order late in a familiar series; activation was not

observed in the hippocampus … when an entirely novel ordering of familiar pictures was presented. Thus, hippocampal activation signaled a mismatch of expectation in temporal order per se, and not merely a novel sequences of pictures (Eichenbaum 2013, 81).

In this example, Eichenbaum describes what kind of task is involved in each paper (i.e., a virtual reality driving game (Lehn et al. (2009) and the presentation of pictures (Kumaran and Maguire (2006)) and the results from each task (i.e., human hippocampal activation in the virtual reality task, and human hippocampal activation when participants were shown items out of order later on in a 'familiar series') (Eichenbaum 2013, 81). What is absent in this description is *an evaluation of the methods/protocols used to conduct and analyze said tasks*. This leaves open the possibility that the methods used by the cited studies are unsound. If the methods of the cited findings in the case study are unsound, this would discredit or at least put pressure on the corroboration of claims 1-5.

However, one might argue that because the papers cited in support of claims 1-5 are from papers published in mind-brain sciences journals with high impact factors and citations that the methods/protocols used in the cited papers must be sound (I make a similar point in section 3.2.3. of this dissertation wherein I argue that because the cited findings are published in mind-brain journals with high impact factors/citations, the cited findings can be considered reliable pieces of information).

Citation-based bibliometric measures, such as a journal's impact factor, are frequently consulted in the evaluation of scientists and their research (Seglen 1994, 1; Seglen 1997, 498).[142] One of the attractions of this measure is that it is considered to be quantitative, simple to calculate, and objective (Seglen 1997, 498). Assuming this measure is indeed objective, one might have good reason to believe that the methods and protocols used by papers published in journals with high impact factors and citations are scientifically sound.

Should one take it on good faith that the methods/protocols used in the cited papers are scientifically sound due to their being published in journals with high-impact factors and

[142] Funding allocation based on impact factors has been reported in countries such as Canada and Hungary (Seglen 1997, 498). In many Nordic countries, impact factors have been appealed to in the evaluation of individual scientists as well as scientific institutions (ibid).

citations? Drawing on Seglen (1994, 1997), I argue we should not, because impact factors are not always reliable indicators of sound scientific research. I raise three points to support this argument.[143] Firstly, a journal's impact factor is not 'statistically representative' of individual research papers published within it (Seglen 1997, 498-499). An impact factor is defined as "… the mean annual citedness of all articles published in a given journal in a given year" (Seglen 1994, 1). For a journal's impact factor to be fairly representative of its papers, the citation rate for each individual article must demonstrate a "… narrow distribution, preferably a Gaussian distribution around the mean value (the journal's impact factor)" (Seglen 1997, 498-499). However, Seglen's case study of three different biomedical journals showed skewed distributions of citation rates; only a select few papers came close to the mean (ibid, 499). Further analysis of Seglen's case study demonstrated: "… the most cited half of the articles are cited, on average, 10 times as often as the least cited half" (ibid). Assigning the 'same score' (i.e., the impact factor) to all papers within a given journal disguises this important difference (ibid). That is, an article with low to no citations will receive "… full credit for the impact of the few highly cited articles that predominantly determine the value of the journal impact factor" (ibid). Therefore, impact factors do not reliably represent the quality of individual papers. This is one reason to not rely (solely) on a journal's impact factor/citations when determining or assessing the quality and scientific rigour of its papers.

Secondly, authors consider many factors other than a journal's impact factor when submitting a paper for publication. Authors do not necessarily submit their best work to the most prestigious journals, nor are their papers necessarily representative of the impact factor of the journal in which they are published (Gordon, 1984; Seglen 1997, 499). While some authors may decide which journal to submit their paper to based on a journal's impact factor, there are many other factors that can contribute to an author's decision on where to submit and/or resubmit a paper for publication. Other factors include but are not limited to: the journal's area of science and how well it aligns with the author's own specialization, how quickly the editing team works, the fairness and helpfulness of feedback from peer-reviewers, the likelihood of acceptance, publication turn around times, etc. (Seglen 1997, 499). Gordon (1984) performed a sociological

[143] For more arguments on why impact factors are not reliable measures of the quality of scientific research, see also Allison (1980); Folly et al. (1981); Hargens and Felmlee (1984); Moed et al. (1985); and Seglen (1992,1994, 1997).

experiment in which 98 biomedical authors were asked to fill out a questionnaire including questions such as which journal they had recently submitted to, which other journals were considered, any journals that previously rejected their submission, and if rejected, would they resubmit to that journal and why (29).

Gordon's analysis demonstrated a great diversity of responses among the 98 authors involved in the study with respect to reasons for journal selection criteria, submission, and resubmission. The data acquired from Gordon's study provided "… no support to the assumption that scientists necessarily chose the most prestigious journal from amongst those which they feel would accept their papers" (ibid, 40). Accordingly, a paper of high scientific quality may end up in a journal with a lower impact factor because that journal's area of science aligns more specifically with the author's area of expertise, or because the journal is known to have quick publication turnarounds, etc. Therefore, it is not the case that scientifically sound papers are published *exclusively* in journals with high impact factors/high citation rates. This is yet another reason one should not rely solely on bibliometric measures like a journal's impact factor/citations in assessing the soundness/quality of a research paper.

Thirdly, self-citations are frequently not corrected for in the calculation of a journal's impact factor (Seglen 1997, 500). That is, databases that record and calculate impact factors such as the Science Citation Index and the Journal Citation Report include citation types of all sorts such as citations in editorials, meeting abstracts, etc. (ibid). Further, citations for translated journals are often listed twice (ibid). Accordingly, journals including items like editorials and meeting abstracts will have an impact factor that is 'greatly inflated' compared to journals excluding these items (ibid). Editors who want to increase the impact factor of their journal will often make 'frequent references' to these items since the databases do not correct for self-citations (ibid).

This is particularly relevant to my analysis of the case study insofar as of the 59 papers cited to support claims 1-5, Eichenbaum co-authored 12/59, for a total of 20% of his references.[144] Recall in chapter two of this dissertation I noted one of the alleged advantages of research

[144] Percentage is rounded down for simplicity.

synthesis is its ability to minimize bias (Chalmers and Altman 1995, Gurevitch et al. 2018, Mulrow 1995, Shamseer et al. 2015, etc.). While none of these authors advanced a specific argument for the minimization of bias in research synthesis, I noted their claims about bias reduction hinge on systematic reviews adhering to some kind of official and standardized protocol for writing review papers (whatever kind of review paper it may be).

However, because the case study does not include a methods section stating which standardized guidelines were used, if any, and because Eichenbaum co-authored 20% of his references, the corroboration of claims 1-5 is jeopardized by potential author bias. Some of the cited findings are not as objective as they initially appear. The fact that citation databases do not correct for self-citations opens the threat of objectivity and worries about bias. This is another reason why impact factors/citations should not be considered purely objective nor a reliable indicator of the soundness of one's scientific research. Thus, because one cannot be certain that the methods cited by the paper's Eichenbaum cites are reliable, the integrity of the substantial evidence for claims 1-5 is questionable. This point makes it hard to assess the extent to which claims 1-5 have been corroborated in the case study due to a lack of evaluation of the methods used in the papers Eichenbaum cites.

In sum, I have highlighted four key components absent in the case study. First is a definition of terms like 'substantial evidence' 'numerous studies' 'compelling evidence' and 'suggests' (83. 86-87). Eichenbaum's arguments in the case study are implicit; these arguments, on my reading, hinge on terms like 'substantial evidence' and 'suggests' (ibid). If claim *x* is backed by substantial evidence suggesting *x*, one might be confident in the corroboration of that claim. However, to properly evaluate the extent to which claims 1-5 have been corroborated, a statement describing what Eichenbaum means by each term and how they factor into the corroboration of the claims is required. If these terms are undefined, it is difficult to assess whether claims 1-5 has been corroborated.

Secondly, on my reading, Eichenbaum's (implicit) arguments for claims 1-5 also hinge on the terms 'complemented by' and 'consistent with' (ibid, 82-84, 86). That is, part of Eichenbaum's implicit argument is that there are numerous and substantial evidence backing claims 1-5. Additionally, the cited evidence comprising the substantial and compelling evidence, on my reading, requires consistency. The term consistency is used throughout the case study to support claims 1 and 4 and suggests that for these claims to be corroborated by the interdisciplinary cited findings, those findings must fit together in such a way that they are consistent with each other and with the resulting claim. Although Eichenbaum does not use the term consistency explicitly for claims 2, 3 and 5, it seems necessary for those claims as well that the cited findings be consistent with each other. However, as I note above, terms like consistency are also not defined. It is therefore unclear what Eichenbaum means by consistency and how it is measured in the case study; accordingly, it is open to question that the cited findings fit together in a way indicative of corroboration.

This second point is further complicated by the third feature lacking in the case study. Namely, there is a general lack of anomalous findings, and the few times anomalous findings are mentioned, they are given little explanation with respect to how they fit in with the rest of the cited findings. This is especially important because if it is true that the cited findings need to be consistent with each other for the corroboration of claims 1-5, ideally there would be more than just a general compatibility between findings. Ideally, the cited findings would be consistent with each other in a way that also rules out random error, coincidence, alterative hypotheses etc. Therefore, the introduction of anomalous findings and how they fit together with the rest of the findings is critical. If Eichenbaum introduced anomalous findings for each claim, and if he also then explained how they fit together with the other findings (e.g., how the other findings reliably rule out the anomaly), the epistemic integrity of claims 1-5 would increase.

However, in the few instances that Eichenbaum does mention anomalous findings, they are either immediately dismissed (another finding is introduced and the anomalous finding is left with no explanation), or the explanation is unsatisfactory (such as Eichenbaum claiming the discrepancy between the rat and monkey studies cited in claim 4 can be explained away by the use of different experimental protocols) (ibid, 86).

The final feature I flagged as lacking in the case study was an evaluation of the methods used in the papers he cites. This is important because claims 1-5 depend on the reliability of the findings cited in support of them. That is, it is not sufficient for the corroboration of claims 1-5 being supported by substantial evidence; that substantial evidence must consist of reliable and scientifically sound findings. If the substantial evidence cited in support of claims 1-5 comprised of findings from unsound scientific research, the epistemic integrity of claims 1-5 would be compromised. One way to demonstrate the reliability of the cited findings in the case study would be to include an evaluation of the methods used by the cited papers. This might involve including statements like "these are the methods used in this cited paper and this is why they can be considered reliable and sound methods" to help increase the reliability of the findings supporting claims 1-5.

One might argue, however, that this is unnecessary in virtue of the fact that the majority of the cited findings were published in papers from journals with high impact factors and citation rates; because such journals allegedly produce research of high scientific rigour, one can confidently assume the respective findings from those papers are the product of reliable methods. However, as I argue above, drawing on Seglen (1994, 1997), citation-based bibliometric measures such as impact factors are not reliable indicators of sound scientific research. It is therefore possible that some of the methods used in the cited papers were unsound, discrediting the reliability of the cited findings. Furthermore, the fact that Eichenbaum co-authored 20% of his cited references generates worry about author bias – again, potentially reducing the reliability of the cited findings.

Together, these four missing components put pressure on the case study representing a clear case of IC. Towards the end of my analysis in chapter three, I began to notice some questions and problems complicating the corroboration of claims 1-5. My analysis in this chapter explored those problems and showed that given the lack of definitions such as 'substantial evidence' and 'consistency', along with sufficient explanations of anomalies and the methods used in the cited papers, it is unclear how the cited findings fit together to support claims 1-5; Eichenbaum does not tell us this. Accordingly, it is unclear that claims 1-5 have been corroborated. I will now highlight some key lessons from this analysis.

4.2 Lessons from a Review Paper

The first lesson I draw from this analysis is the fundamental assumption underlying the reductionist and non-reductionist accounts of progress I surveyed in chapter one – i.e., that pieces of evidence from different areas of science must fit together in some way to yield epistemically strong claims – is not as simple or as clear as it is assumed to be in the literature. That is, in the case study I analyzed, Eichenbaum argued the claim that the role of the hippocampus is to create a spatio-temporal framework for the organization of memories is well supported by pieces of evidence from different areas of science. This unit of analysis – the *Trends* review paper – provided an example of an instance of a claim gaining corroboration in virtue of different pieces of interdisciplinary evidence supporting it.

However, as my analysis in this chapter revealed, it is unclear exactly how pieces of evidence fit together to corroborate the claims advanced in the case study, and accordingly, it is difficult to assess whether or not its claims stand on strong epistemic footing. Put differently, because it is unclear how the pieces of evidence fit together in the case study, it is uncertain how and to what extent claims 1-5 have been corroborated. My analysis thus reveals that the assumption I highlight in chapter one of this dissertation is complex and messy. When one analyzes units of analysis that appear to do the work required of IC, one sees that pieces of evidence do not necessarily fit together as neatly as one might expect. This assumption therefore requires more philosophical attention and work.

The second lesson I draw from my analysis is my working definition of IC (which I introduce in chapter two) requires revision. Initially my working definition of IC read that findings corroborate claims about the mind-brain insofar as those findings are interdisciplinary, mutually supportive, and together all points towards a given claim. However, as I note above, consistency and mutual support alone is not sufficient for corroboration. Pieces of evidence need to be more than consistent and mutually supportive; they also need to rule out coincidence, error, and alternative explanations. Thus, my revised definition of IC reads: interdisciplinary corroboration is a phenomenon whereby findings corroborate claims about the mind-brain when said findings are interdisciplinary, mutually supportive and consistent in a way that rules out coincidence, errors, and alternative hypotheses, and where the findings all point towards a similar (but not identical) claim.

This leads to the third lesson. My analysis revealed not only that the assumption held by reductionist and non-reductionist accounts of progress is complex and messy, but that certain fundamental pieces of information are missing to adequately assess the extent to which corroboration occurs in the case study. For example, definitions for key terms like 'substantial evidence', 'suggests', and 'consistency' were absent. This absence makes it difficult to assess whether claims 1-5 have been corroborated since it is unclear what Eichenbaum meant by evidence being consistent with each other and if he even considered it to be a requirement for the corroboration of claims 1-5. Further, because the reader does not know Eichenbaum's selection criteria nor his review process, it is unclear why certain findings were included and others were omitted. Was the general lack of anomalous findings simply because Eichenbaum did not encounter many in his search process? Or were some encountered but omitted; if this was the case, what was the justification for this omission? Of the few anomalies that were included in the case study, Eichenbaum does not provide a satisfactory explanation of how those anomalies fit together with the other cited findings and how those cited findings rule out the anomaly. If one cannot rule out the possibility of potential anomalies, then it is difficult to evaluate whether claims 1-5 have been corroborated. If one knew, on the other hand, that Eichenbaum encountered no other anomalies than those cited in the case study, and if Eichenbaum explained why the cited anomalies could be reliably dismissed, one would be better positioned to assess the corroboration of claims 1-5. However, since that information is not provided, it is unclear whether claims 1-5 have been corroborated. Finally, the lack of an evaluation of the methods used by the cited papers in the case study makes it challenging to assess the corroboration of claims 1-5. If claims 1-5 have been corroborated by the cited findings, those findings need to be scientifically sound/reliable. However, because Eichenbaum does not provide an evaluation or justification for the soundness of the methods used in the cited papers, it is unclear whether the cited findings constitute reliable pieces of information. Since this cannot be ascertained, it is difficult to evaluate the extent to which corroboration has occurred in the case study.

The point I want to emphasize is not that Eichenbaum's main claim is false, nor that the case study represents poor quality scientific research. The goal of the dissertation was to analyze and assess how pieces of evidence fit together in such a way that is indicative of corroboration. The lessons emerging from the analysis of the case study is not that claims 1-5 are false. Instead, the lessons are that the fundamental assumptions shared by reductionist and non-reductionist accounts of progress and fitting together is more complicated than assumed by the literature (and that this assumption requires further philosophical attention), that my working definition of IC required refinement (i.e., the addition of the fact that the findings must be consistent in such a way that they rule out coincidence, error, alternative hypotheses, etc.), and that certain critical pieces of information are lacking to adequately address the extent to which the corroboration of claims 1-5 has occurred.

4.3 A Methodology for Evaluating Review Papers for Corroboration

If philosophers are interested in knowledge production and want to assess how pieces of evidence from different areas of science amalgamate to produce better and epistemically stronger claims, there are certain questions and concerns that should be raised towards the units of analysis (e.g. different types of review papers) which allegedly do this work. In virtue of this, I use the questions and concerns prompted by my analysis in this chapter to advance the following itemized list as a methodology for analyzing review papers with an eye towards corroboration in the mind-brain sciences. Importantly, I consider this methodology to be *a working framework*; as more research is conducted on different units of analysis and IC, I suspect the list (as well as my current definition of IC) will change.

Factors to consider in evaluating review papers for corroboration:

1.1. What is the main aim of the review paper?

1.2. What are the main claims advanced in the review? Are those claims implicit, explicit, or both?

1.3. How does the author argue for those claims? Are their arguments implicit or explicit? Do the arguments depend on any specific terms? What other important factors do their arguments hinge on?

1.4. What kind of language does the author use to describe the findings they cite in support of their claims? What are the key terms used in the paper to relate the different cited findings? Does the author provide a definition of those terms? Are those terms qualitatively or quantitatively defined?

1.5. Are the cited findings interdisciplinary? I.e., does the author include evidence from multiple areas of the mind-brain sciences? Are the species, tasks, and methods of analysis varied in the cited papers?

1.6. How current are the cited findings? Does the author draw primarily on current research on a topic, or do they include older papers as well?

1.7. Does the author primarily report findings from research papers, or do they elaborate on those findings, explaining the tasks and methods used in each and how they fit together with other cited findings?

1.8. Does the author report the conclusion of a given research paper as the cited finding, or is it one of many findings from various tasks in the paper? Are there other findings in the cited paper that might complicate the picture the author of the review paper is painting?

1.9. Does the author evaluate the methods involved in the cited findings and state how they constitute reliable pieces of information? Are there other aspects of the review that make up for the lack of an assessment of the methods of the cited papers (if such a lack is present in the review paper) such that the reliability of the cited findings remains intact?

1.10. How is coincidence ruled out in the review paper? How are random errors and alternative hypotheses ruled out?

1.11. Does the author include any contradictory or anomalous pieces of evidence? If yes, how many findings fall under these categories? How are they dealt with, and does the author provide sufficient reasons for why those anomalous findings can justifiably be dismissed?

1.12. Does there review explicitly adhere to a standardized guideline (e.g., the Cochrane Handbook)?

1.13. If no such standardized guideline is used, does the author include a methods section wherein they describe and justify the methods they use in their review.

1.14. Does the author describe their selection criteria and search process wherein they state which journal databases they consulted, their criteria for inclusion and omission, etc.?

1.15. What measures does the author use to reduce risk of bias in the individual papers they cite in the review?

1.16. Does the author co-author any of the cited papers? If so, how many? How does this potential risk of author bias factor into corroboration?

1.17. What methods does the author use to combine (i.e., synthesize) the findings in the review paper? Is a meta-analysis involved or is the combining qualitative? If qualitative, how is the combining justified? How are anomalous and/or contradictory findings factored into this synthesis?

1.18. What kind of corroboration is represented in the review paper? Is it one of confirmation, lending support, or something else?

1.19. Does the author include/flag any limitations their review may be subject to? Do they flag any gaps in their arguments? Do they acknowledge any potentially complicating factors that decrease the likelihood of their main claims being corroborated?

This list is meant to be read as a working framework. As I note, in section 4.4. of this chapter, future research on different types of units of analysis will likely yield important implications for this methodology (i.e., it may be the case that the list is expanded, certain questions are refined, or that different types of units of analysis require their own methodology, thus resulting in a plurality or a hierarchy of methodologies for evaluating corroboration in the mind-brain sciences, etc.).

4.4 Questions and Directions for Future Research

This dissertation only begins to scratch the surface of the epistemic role of review papers for IC in the mind-brain sciences. In this dissertation, I analyzed one specific case study; however, one case study is not sufficient for drawing conclusive claims about the nature of IC and its relation to scientific progress. Future projects might involve comparing multiple *Trends* reviews to discern whether my findings were simply coincidental, then comparing a *Trends* review paper with a systematic review paper to see whether and to what extent the additional component of a methods and materials section increases the integrity of corroboration at play in the case study. I see the next step of this process as looking at systematic reviews with meta-analysis.[145] Another

[145] Review papers are one way to explore whether interdisciplinary evidence fits together in a way that corroborates claims about the mind-brain. I conducted an analysis of a *Trends* review paper in my dissertation and

project might be analyzing these units of analysis in the of context big data projects (e.g., BrainMap) and open science.

As qualitative (and quantitative) analysis software programs such as NVIVO12 develop over time, a more fine-tuned analysis of the case study I analyze in this dissertation may also be possible; this may result in a more nuanced definition of IC and a modified methodology for evaluating review papers for corroboration.

As more research is conducted on the nature of IC and different units of analysis that do this work, I suspect my current definition of IC will undergo refinement, and that my methodology will similarly undergo revision. Moreover, I did not elaborate on the methodology I advance in section 4.3 or provide different methodologies for assessing different types of units of analysis for research synthesis. Perhaps, as more comparative analyses are conducted (e.g., a *Trends* review paper with a systematic review, or a systematic review with a systematic review including a meta-analysis, etc.), it might be revealed that a plurality of methodologies is best suited for the task of evaluating different kinds of units of analysis that synthesize research across different areas of science.

Another important question with respect to evaluating IC in the mind-brain sciences via units of analysis such as review papers is whether the mind-brain sciences require their own distinct guideline/handbook/checklist such as the Cochrane handbook for the health sciences. I noted in my second chapter that the standards for organizations such as the Cochrane Collaboration in the health sciences are not applicable to the mind-brain sciences, in virtue of the latter synthesizing information from a variety of areas of the mind-brain sciences, using different methods, protocols, and techniques. As such, future research might involve the advocacy of such an organization or list of criteria and identify what standards must change and why. In doing so, one might look to the standards of organizations such as the Cochrane Collaboration and identify

---

found that the pieces of interdisciplinary evidence in the case study did not fit together as neatly nor as clearly as would be required for corroboration. I see the next stage of my research as systematically moving through various units of analysis that appear to do the work required for IC; this includes units of analysis such as systematic reviews with meta-analysis/various types of meta-analyses.
For a different view of where to look for the fitting together of interdisciplinary evidence, see Sullivan (2016), who is skeptical of meta-analyses and thinks the best strategy is to look at actual ground-work practices in labs.

whether any of their standards should be included in the criteria for writing reviews in the mind-brain sciences and which new ones need to be introduced in virtue of the mind-brain sciences being distinct from disciplines like the health sciences.

Future research might also address the question of what aspects of scientific papers are being cited in the review paper. That is, scientific research papers typically include detailed descriptions of different tasks and the results of those tasks; there are frequently more than just one very specific finding in a given scientific paper. For example, Buhusi and Meck (2005) highlighted multiple brain structure/networks involved in numerous timing tasks; despite this, Eichenbaum selectively focused on one vague aspect – i.e., the fact that there are multiple sources of temporal information in the brain – while excluding information about the many areas and networks in the brain that were activated during various timing tasks. One might therefore investigate what parts of the cited scientific research papers are being included in the review in question and what parts are omitted. This is an important question because it has the potential to impact the corroboration of a given claim. If a claim states something along the lines of "the hippocampus is the key structure involved in spatio-temporal memory", and the author selectively focuses on aspects of papers highlighting hippocampal activation while excluding findings demonstrating activation in other areas, one might question just how epistemically warranted that claim is. This leads to my final point.

Finally, future research might involve advocating for a methods section in future *Trends* review papers. That is, especially if the review does not adhere to standardized guidelines such as those included in the Cochrane handbook, the inclusion of a methods section wherein the author states their selection criteria, defines their key terms, identifies their main claims, states how the studies they cite were screened for bias and error in individual papers, etc., may help the reader better assess the extent of corroboration in review papers, thereby increasing the integrity of the review in question.

In sum, I demonstrated that a lack of definitions for critical terms like 'substantial evidence', 'numerous studies', 'consistency', etc., makes it difficult to assess the extent of corroboration at play in the case study. This, in addition to a lack of adequate explanations for anomalous findings and an evaluation of the methods used by the papers Eichenbaum cites in support of claims 1-5 complicate the case study representing a clear case of IC. I showcased three implications of my analysis. First, the assumption held by the accounts of progress I surveyed in chapter one is complex and requires further philosophical analysis. Secondly, my working definition as introduced in chapter two required revision; my new definition added that IC required that interdisciplinary evidence be consistent in such a way that rules out coincidence, errors, and alternative explanations. Thirdly, I argued that critical information was missing in the case study to adequately assess the extent to which claims 1-5 were corroborated. I used the lessons and the questions/concerns I addressed in this chapter to serve as criteria for my methodology for evaluating review papers with an eye towards corroboration.

In the following chapter, I conclude my dissertation by re-capitulating my main aim, how I achieved it, the results of my analysis, as well as how I think my results generalize out into other areas of philosophy and science.

Chapter 5

5 Conclusion

My goal in this dissertation was to analyze what it means for information from different areas of science to fit together in such a way that results in strong epistemic claims by addressing the question of how and to what extent claims about the mind-brain are corroborated in scientific practice.

To address how and to what extent claims about the mind-brain are corroborated in scientific practice, I began by establishing the relevant historical and motivational background for my dissertation in chapter one. This involved surveying a number of reductionist (Nagel 1949, 1961, 1979; Bickle 1998, 2003, 2006) and non-reductionist accounts (Darden and Maull, 1977; Mitchell 2002, 2003; Mitchell and Dietrich 2006; Darden 2006; Roskies 2010; Piccinini and Craver 2011) of scientific progress emphasizing the notion of fitting for claims about the mind-brain to stand on firm epistemic footing. I noted that while they all concede that interdisciplinary evidence must fit together to produce strong epistemic claims, each notion of fitting together (e.g., ruthless reduction for Bickle (2003, 2006), 'convergence' for Roskies (2010), etc.)) was vague in that it was unclear from their accounts and their examples what it meant for pieces to fit together in such a way that 'conclusive evidence' is provided for a claim about the mind-brain (Roskies 2010, 641).

To better understand how evidence fits together to yield strong epistemic claims, I introduced the term 'interdisciplinary corroboration' (hereafter IC) as a general placeholder for what I think philosophers of science have been gesturing towards, but which has yet to be clearly articulated and analyzed in chapter two. My working definition of IC stated findings corroborate claims about the mind-brain when they all point towards a general claim and when said findings are interdisciplinary, consistent, and mutually supportive. The accounts of progress I survey in chapter one all claim that the fitting together of interdisciplinary evidence results in epistemically strong claims; implicit in their accounts is an argument that claims about the mind-brain (e.g., the frontal region is involved in semantic processing (Roskies 2010, 641)) are corroborated by the fitting together of interdisciplinary evidence. For this reason, I introduced IC as a general placeholder for these accounts to facilitate my analysis of how interdisciplinary evidence fits together to corroborate claims about the mind-brain.

Since my definition of IC was introduced as a working definition, it required further analysis to better understand the nature of IC in the mind-brain sciences. To better understand how IC functions to corroborate claims about the mind-brain, I suggested looking for instances of it in practice. I argued that a good place to start assessing IC in practice was by looking at review papers, which are units of analysis surveying literature from a variety of disciplines on a given topic. Some review papers simply contain a literature review; however, many review papers also contain arguments (sometimes implicit, sometimes explicit). These arguments can involve showing how a given claim about the mind-brain is supported by the evidence cited in the review paper (i.e., showing how the interdisciplinary evidence fits together to corroborate a claim about the mind-brain). Review papers are therefore ideal units of analysis to begin investigating the nature of IC and addressing the extent to which corroboration occurs in these papers.

To address the questions of how and to what extent corroboration of claims about the mind-brain occur in review papers, I conducted a two-part analysis in chapters three and four of a case study of a *Trends* review paper entitled 'Memory on Time' by Eichenbaum (2013).

I began my analysis with some cursory criteria. That is, I came to the table with a general understanding of how *Trends* reviews function through my descriptive work in chapter two. Further, my analysis was guided by cursory criteria such as the following questions: "what is the main aim of the review", "what claims does Eichenbaum defend", "how does Eichenbaum show support for his claims", "what kind of arguments does Eichenbaum provide", and "does the case study represent IC"? My analysis in chapter three, guided by this criteria and knowledge of *Trends* review papers, was structured around the work Eichenbaum did in the case study to support claims 1-5. I stated Eichenbaum's arguments for claims 1-5 were implicit and that his implicit arguments hinged on terms like 'substantial evidence' and 'consistency' (Eichenbaum 2013, 82-84, 86). That is, claims 1-5 are corroborated insofar as they are supported and suggested by 'substantial evidence' consisting of findings from different areas of the mind-brain sciences, on different species, using different tasks and methods of analysis, all of which are consistent with each other and its associated claim. An analysis of the language used in the case study revealed that the kind of corroboration at play was one where claims 1-5 are best understood as being highly probable, but not full-stop true. I further argued that because the cited findings were published in mind-brain science journals with high impact factors and citation numbers, the cited findings appear to constitute reliable sources of information. This increases the strength of corroboration in the case study. On a first pass, the case study appeared to represent a strong case of IC.

However, my analysis structured around what Eichenbaum does to support claims 1-5 gave rise to questions and concerns such as: "what counts as substantial evidence", "what is meant by consistency in the case study", and "how are anomalies factored into the findings supporting claims 1-5". These questions led me to ask how well the case study represented IC after all.

I used these questions to guide the second part of my analysis, which I conducted in chapter four. While my analysis in chapter three focused on the work Eichenbaum did to support claims 1-5, my analysis in chapter four focused on what was absent in the case study. That is, definitions for key terms like 'substantial evidence' and 'consistent with' were lacking; because Eichenbaum's arguments for claims 1-5 implicitly depend on these terms, it is crucial to understand how Eichenbaum uses these terms. However, because Eichenbaum provides no definitions or explanations of these terms, it is challenging to evaluate the extent to which claims 1-5 have been corroborated. Further, because the anomalous findings introduced in the case study are inadequately addressed/explained by Eichenbaum, it is unclear how all the cited findings fit together; Eichenbaum does not tell us. In elaborating on these three questions, I noticed another aspect missing in the case study was an evaluation of the methods used by the cited findings; because of this, it is hard to discern whether the methods used in the cited papers are reliable. If the methods used by the cited papers were unsound, the cited findings would comprise unreliable pieces of information, thus discrediting or putting pressure on the corroboration of claims 1-5. Taken together, these questions and the concerns they raise complicated the case study representing a clear case of IC.

I then stepped back from the analysis in the second half of chapter four to showcase three key lessons emerging from the analysis. The first lesson is that the fundamental assumptions shared by both reductionist and non-reductionist accounts of progress I surveyed in chapter one – i.e., that epistemically strong claims are the result of evidence from different areas of science fitting together – is complex and not nearly as straightforward as the literature makes it out to be. The second lesson is that my working definition of IC needed to be revised. IC should now be understood as a phenomenon whereby findings corroborate claims about the mind-brain insofar as they are interdisciplinary, and mutually supportive/consistent in such a way that rules out coincidence, errors, and alternative hypotheses, and all point towards a similar claim. The third lesson from my analysis is that certain critical pieces of information were missing to adequately assess the extent to which claims 1-5 were corroborated. I used the missing information and the concerns that arose during my analysis of the case study to advance a methodology for evaluating review papers in the mind-brain sciences for corroboration.

The contributions of my dissertation to philosophy of the mind-brain sciences are twofold. First, I demonstrated that a fundamental assumption shared by numerous accounts of progress in the philosophy of the mind-brain sciences – i.e., that epistemically strong claims are the result of evidence from different areas of science fitting together – was not as clear as the literature presumed it to be. Accordingly, I showcased that more philosophical attention is required to better understand how interdisciplinary pieces of evidence fit together to corroborate claims about the mind-brain. I began this work by advancing the concept of IC and analyzing a case study of a review paper to start to get a handle on how IC works in practice. My analysis revealed critical information was missing in the case study to properly determine whether and to what extent its claims were corroborated. The implications of this for philosophers of the mind-brain sciences interested in knowledge production is that more work needs to be conducted on different units of analysis that function to corroborate claims about the mind-brain (such as comparing a *Trends* review with a *Frontiers* review, or a Frontiers review with a review with meta-analyses, etc.). The more research conducted on the different units of analysis that function to corroborate claims about the mind-brain, the more one can elaborate and fine-tune both my working definition of IC as well as the methodology I advance to evaluate review papers for corroboration. The point is not that the assumption held by reductionist and non-reductionist accounts of progress is false; instead, my contribution is to demonstrate that the assumption (which to date has been under-analyzed and unquestioned) is complex, and that future research should continue to examine this assumption and its implications for knowledge production and scientific progress.

My second contribution to the philosophy of the mind-brain sciences is the methodology I advance in my fourth chapter for evaluating review papers for corroboration. That is, if philosophers of the mind-brain sciences are interested in how claims about the mind-brain are corroborated by evidence from different areas of science, there are certain questions and concerns that should be addressed in assessing the corroboration of claims in units of analysis such as review papers. My methodology outlines the questions and concerns I deem important to be addressed in evaluating the corroboration of claims in review papers.

It should be acknowledged my methodology is meant to be a working framework. My hope is that as more research is conducted, my methodology will become more fine-tuned. One limitation of my methodology is that since the dissertation is the beginning of what will hopefully be an ongoing investigation of the role of IC in the mind-brain sciences, my methodology pertains primarily to *Trends* review papers. Future research might address how different kinds of units of analysis corroborate claims about the mind-brain. For example, such projects might involve examining a systematic review, or comparing a *Trends* review paper with a systematic review, or comparing reviews with different types of meta-analyses, etc. As these units of analysis involve different components (i.e., some include meta-analysis, some do not; some include a methods section, some do not, etc.), it may turn out that a plurality of methodologies for approaching IC in the mind-brain sciences is necessary.

More generally, my dissertation has implications for philosophy of science and empirically engaged philosophy of mind. With respect to the former, other areas of science similarly draw on interdisciplinary evidence to corroborate claims. For example, multiple pieces of evidence from cosmology, chemistry, and particle physics come together to corroborate the existence of dark matter.[146] Here, one might ask whether and to what extent these pieces of evidence corroborate the existence of dark matter. To address this, one might consider various units of analysis demonstrating the corroboration of dark matter and appeal to some of the questions I flag in my methodology to evaluate whether the hypothesis about the existence of dark matter has been corroborated. Conversely, one might appeal to this list to evaluate alternative theories concerning dark matter such as the Modified Newtonian Dynamics theory and the Tensor-

[146] I am grateful to Dr. Marie Gueguen for this point.

Vector-Scalar gravity theory.[147] Appealing to my methodology in other areas of science outside the mind-brain sciences may therefore help assess units of analysis for the corroboration (or lack thereof) of dark matter hypotheses in physics.[148]

My dissertation also has implications for empirically engaged philosophy of mind. That is, in the philosophy of mind, philosophers sometimes appeal to empirical research on a given topic in order to argue for (or against) a philosophical position. For example, in their paper 'Neural Plasticity and Consciousness', Hurley and Noë (2003) argue for a dynamic sensorimotor account of consciousness that allegedly bridges some versions of the explanatory gap problem. In arguing for their account of consciousness and how it solves some explanatory gap problems, Hurley and Noë appeal to scientific evidence to support the main claims behind their account. I will not get into the specifics of their arguments, nor how it deals with explanatory gap problems, but instead want to highlight the fact that their account of consciousness hinges on empirical work on neural plasticity demonstrating that the cortex both dominates (here they cite Ramachandran and Blakeslee 1998; Ramachandran and Hirstein 1998) and defers (here they cite Buchel et al. 1998; Cohen et al. 1997; Sadato et al. 1996, 1998) (Hurley and Noë 2003, 133-134).

Since their account depends on this cited empirical evidence, if a philosopher of mind wanted to argue against Hurley and Noë's account, a crucial step would be demonstrating that the empirical evidence upon which their account is based is weak. Using some of the questions I outline in my methodology might help the reader better assess the empirical literature for strong epistemic claims, therefore making a stronger case against Hurley and Noë (2003) if it turns out that the papers they cite and their respective findings are not well-established. Conversely, a philosopher of mind might appeal to my methodology to identify strong empirical claims from review papers and cite those to support some of their own arguments about the mind.

In sum, via my analysis of Eichenbaum's (2003) 'Memory on Time', I demonstrated that pieces of evidence may not fit together as neatly as one would require for the corroboration of a claim. This does not entail that the assumption that strong conclusions require evidence from multiple

[147] Again, I am grateful to Dr. Marie Gueguen for this point.

[148] This example generalizes out in similar ways to other hypotheses in physics and other areas of science.

fields of science is false, but that the assumption is complex and requires more philosophical attention to better understand how claims are corroborated in the mind-brain sciences. The methodology I advance at the end of chapter four is meant as a guide for philosophers of science to better assess the epistemic strength of claims emerging out of review papers. Future research will help expand on this methodology as more work on the nature of IC is conducted.

References

"About JNeurosci". *The Journal of Neuroscience*. Accessed April 27th, 2019. http://www.jneurosci.org/content/about-jneurosci

"About Science & AAAS". *Science*. Accessed April 27th, 2019. https://www.sciencemag.org/about/about-science-aaas?_ga=2.10191605.1947000741.1556017401-385344628.1554129536

"About the Journal". *Nature International Journal of Science*. Accessed April 27th, 2019. https://www.nature.com/nature/about

"About the Journal". *Nature Reviews Neuroscience*. Accessed April 27th, 2019. https://www.nature.com/nrn/about

"About Us". *Cochrane Organization*. Accessed January 15th, 2019. https://www.cochrane.org/about-us.

Agster, Karla. L., Nobert J. Fortin, and Howard Eichenbaum. 2002. "The Hippocampus and Disambiguation of Overlapping Sequences". *Journal of Neuroscience* 22 (13): 5760-5678.

"Aims and Scope". *Neuron*. Accessed April 24th, 2019. https://www.cell.com/neuron/aims.

"Aims and Scope". *Trends in Cognitive Sciences*. Accessed February 5th, 2019. https://www.cell.com/trends/cognitive-sciences/aims.

Ainge, James. A., Minija Tamosiunaite, Florentin Woegoetter, and Paul A. Dudchenko. 2007. "Hippocampal CA1 Place Cells Encode Intended Destination on a Maze with Multiple Choice Points". *Journal of Neuroscience* 27 (36): 9769-9779.

Albus, James. S., George A. Bekey, John H. Holland, Nancy. G. Kanwisher, Jeffery L. Krichmar, Mortimer Miskin, Darmendra S. Modha, Marcus E. Raichle, Gordon M. Shepherd, and Guilio Tononi. 2007. "A Proposal for a Decade of the Mind Initiative". *Science* 317 (5843): 1321.

Allison, Paul D. 1980. "Inequality and Scientific Productivity". *Social Studies of Science* 10 (2): 163-179.

Anderson, Michael L. 2015. "Mining the Brain for a New Taxonomy of the Mind". *Philosophy Compass* 10 (1): 68-77.

Aristotle. 1930. "On Memory and Reminiscence". In *The Works of Aristotle*, Volume 3, translated by J. I. Beare, edited by W. D. Ross. Oxford: Clarendon Press.

"Article Types". *Frontiers in Neuroscience.* Accessed January 15th, 2019. https://www.frontiersin.org/journals/neuroscience#article-types

Bacon, Francis. 1902. *Novum Organum*. Edited by Joseph Devey. New York: Collier.

Barnes, Carol. A., Matthew S. Suster, Jiemin Shen, and Bruce L. McNaughton. 1997. "Multistability of Cognitive Maps in the Hippocampus of Old Rats". *Nature* 388: 272-275.

Bechtel, William. 2002a. "Decomposing the Mind-Brain: A Long-Term Pursuit". *Kluwer Academic Publishers, Brain and Mind* 3 (2): 229-242.

——— 2002b. "Aligning Multiple Research Techniques in Cognitive Neuroscience: Why It Is Important?". *Philosophy of Science* 69: S48-S58.

——— 2007. "Reducing Psychology While Maintaining Its Autonomy via Mechanistic Explanation". In *The Matter of the Mind: Philosophical Essays on Psychology, Neuroscience and Reduction*, edited by Maurice Kenneth Davy Schouten and Huibert Looren de Jong, 172- 198. Oxford: Blackwell Publishing.

———2008. *Mental Mechanisms: Philosophical Perspectives on Cognitive Neuroscience*. New York: Psychological Press.

Bechtel, William, and Adele Abrahamsen. 2005. "Explanation: A Mechanistic Alternative". *Studies in History and Philosophy of the Biological and Biomedical Sciences* 36 (2): 421-441.

Bechtel, William, and Robert C. Richardson. 1993. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton: Princeton University Press.

——— 2002. "Neuroimaging as a Tool for Functionally Decomposing Cognitive Processes". In *Foundational Issues in Human Brain Mapping*, edited by Stephen José Hanson and Martin Bunzl, 241-262. Cambridge, MA: MIT Press.

Bickle, John. 1992. "Mental Anomaly and the New Mind-Brain Reductionism". *Philosophy of Science* 59 (2): 271-230.

———1998. *Psychoneural Reduction: The New Wave*. Cambridge: MIT Press.

———2003. *Philosophy and Neuroscience: A Ruthlessly Reductive Account*. Dordrecht: Kluwer Academic Publishing.

——— 2006. "Reducing Mind to Molecular Pathways: Explicating the Reductionism Implicit in Current Cellular and Molecular Neuroscience". *Synthese* 151 (3): 411-434.

——— 2007. "Ruthless Reductionism and Social Cognition". *Journal of Physiology-Paris* 101 (4-6): 230-235.

——— 2008. "Real Reduction in Real Neuroscience: Metascience, Not Philosophy of Science (and Certainly Not Metaphysics!)". In *Being Reduced*, edited by Jakob Hohwy and Jesper Kallestrup, 34-51. Oxford: Oxford University Press.

Bickle, John, Cindy Worley, and Marica Berstein. 2000. "Vector Subtraction Implemented Neurally: A Neurocomputational Model of some Sequential Cognitive and Conscious Processes". *Consciousness and Cognition* 9 (1):117-144.

Boone, Worth, and Gualtiero Piccinini. 2016. "The Cognitive Neuroscience Revolution". *Synthese* 193 (5):1509-1534.

Bower, Gordon H. "A Brief History of Memory Research". In *The Oxford Handbook of Memory*, edited by Endel Tulving and Fergus I. M. Craik, 3-32. Oxford: Oxford University Press, 2000.

Brigandt, Ingo. 2010. "Beyond Reduction and Pluralism: Toward an Epistemology of Explanatory Integration in Biology." *Erkenntnis* 73 (3): 295-311.

Brown, Joel. 2017. "BU's Howard Eichenbaum Dies at 69: CAS Professor Conducted Groundbreaking Research in the Science of Memory". *BU Today*. Last updated July 24th, 2017. https://www.bu.edu/today/2017/howard-eichenbaum-obituary/.

Brown, Thackery, Robert S. Ross., Joseph B. Keller, and Chantal E. Stern. 2010. "Which Way Was I Going? Contextual Retrieval Supports the Disambiguation of Well-Learned Overlapping Navigational Routes". *Journal of Neuroscience* 30 (21): 7414-7422.

Buchel, Christian, Cathy Price, R. S. J. Frackowiak, and Karl Friston. 1998. "Different Activation Patterns in the Visual Cortex of Late and Congenitally Blind Subjects". *Brain* 121 (3): 409-419.

Buhusi, Catalin. V., and Warren H. Meck. 2005. "What Makes us Tick? Functional and Neural Mechanisms of Interval Timing". *Nature Reviews Neuroscience* 6 (10): 755-765.

Button, Katherine S., John P. A. Ioannidis, Claire Mokrysz, Brain A. Nosek, Jonathan Flint, Emma S. J. Robinson, and Marcus R. Munafo. 2013. "Power Failure: Why Small Sample Size Undermines Reliability of Neuroscience". *Nature Reviews Neuroscience* 14: 365–376.

Campbell Donald. T., and Donald W. Friske. 1959. "Convergent and Discriminant Validation by the Multitrait-Multimethod Matrix". *Psychological Bulletin* 56 (2): 81-105.

Cartwright, Nancy. 2000. *The Dappled World: A Study of the Boundaries of Science*. Cambridge: Cambridge University Press.

Cat, Jordi. "The Physicists' Debates on Unification in Physics at the End of the 20th Century". 1998. *Historical Studies in the Physical and Biological Sciences* 28 (2): 253-299.

Causey, Robert L. 1977. *Unity of Science*. Dordrecht: Reidel Publishing Company.

Chalmers, Ian, and Douglas G. Altman, eds. 1995. *Systematic Reviews*. London: BMJ Publishing Group.

"Chapter One: Overview and History of Memory Research". Accessed March 20th, 2019. https://catalogue.pearsoned.ca/assets/hip/us/hip_us_pearsonhighered/samplechapter/0205734820.pdf.

Churchland, Patricia. S. 1986. *Neurophilosophy*. Cambridge, MA: MIT Press.

Churchland, Patricia. S. and Terrence Sejnowski. 1992. *The Computational Brain*, Cambridge, MA: MIT Press.

——— 2000. "Perspectives on Cognitive Neuroscience" In *Cognitive Neuroscience: A Reader*, edited by Michael S. Gazzaniga, 14-24. Malden, Mass: Blackwell.

Churchland, Paul. M. 1979. *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press.

——— 1981."Eliminative Materialism and the Propositional Attitudes". *Journal of Philosophy* 78 (2): 67-90.

——— 1985. "Reduction, Qualia and the Direct Introspection of Brain States". *Journal of Philosophy* 82 (1): 8-28.

———1987. *Matter and Consciousness* (revised edition). Cambridge, MA: MIT Press.

———1989. *A Neurocomputational Perspective*. Cambridge, MA: MIT Press.

Clark, Austen. 1980. *Psychological Models and Neural Mechanisms: An Examination of Reductionism in Psychology*. Oxford: Clarendon Press.

Cochran, William G. 1954. "The Combination of Estimates from Different Experiments". *Biometrics* 10: 101-129.

Cohen, Leonardo G., Pablo Celnik, Alvaro Pascual-Leone, Brian Corwell, Lala Faiz, James Dambrosia, Manabu Honda, Norihiro Sadatok, Christian Gerloff, M. Dolores Catala and Mark Hallett. 1997. "Functional Revelance of Cross-Modal Plasticity in Blind Humans". *Nature* 389 (6647): 180-183.

Cohen, Michael X., K. Richard Ridderinkhof, Sven Haupt., Christian E. Elger, and Juergen Fell. 2008. "Medial Frontal Cortex and Response Conflict: Evidence from Human Intracranial EEG and Medial Frontal Cortex Lesion". *Brain Research* 1238 (31): 127-142.

Coltheart, Max. 2004. "Brain Imaging, Connectionism, and Cognitive Neuropsychology". *Cognitive Neuropsychology* 21 (1): 21-25.

——— 2006a "Perhaps Functional Neuroimaging Has Not Told Us Anything about the Mind (So Far)". *Cortex* 42 (3): 422-427.

——— 2006b. "What Has Functional Neuroimaging Told Us about the Mind (So Far)", *Cortex* 42 (3): 323-331.

Cooper, Harris, Larry V. Hedges, and Jeffrey C. Valentine, eds. 2009. *The Handbook of Research Synthesis and Meta-Analysis* (2nd edition). New York: Russell Sage Foundation.

Craver, Carl. F., and Lindley Darden. 2001. "Discovering Mechanisms in Neurobiology". In *Theory and Method in the Neurosciences*, edited by Peter K. Machamer, Rick Grush and Peter McLaughlin, 112-137. Pittsburgh, PA: University of Pittsburgh Press.

Craver, Carl F. 2006. "When Mechanistic Models Explain. *Synthese* 153 (3): 355-376.

———2007. *Explaining the Brain*. Oxford: Oxford University Press.

Cummins, Robert. "Functional Analysis". 1975. *Journal of Philosophy* 72 (20): 741-765.

———1983. *The Nature of Psychological Explanation*. Cambridge, MA: Bradford/MIT Press.

——— 2000. "'How Does It Work?' versus 'What Are the Laws?': Two Conceptions of Psychological Explanation". In *Explanation and Cognition*, edited by Frank Keil & Robert A. Wilson, 117-144. Cambridge, MA: MIT Press.

Cummings, Johnathon N., and Sara Kiesler. 2005. "Collaborative Research Across Disciplinary and Organizational Boundaries". *Social Studies of Science* 35 (5): 703-722.

Culver, David. A. "Ecological Modeling of Lake Erie Trophic Dynamics". 1999. *Ohio Sea Grant College Program as OHSU-TS-054*. Results of Ohio Sea Grant Project R/EM-0201999. Accessed June 15th, 2019. http://ohioseagrant.osu.edu/archive/_documents/publications/TS/TS-054%20Ecological%20Modeling%20of%20Lake%20Erie%20Trophic%20Dynamics.pdf

Curd, Martin, and J. A. Cover, eds. 1998. "Intertheoretic Reduction: Introduction". In *Philosophy of Science: The Central Issues*, 1st Edition, edited by Martin Curd, and J.A. Cover, 903-904. New York: W.W. Norton.

Darden, Lindley. 2006. *Reasoning in Biological Discoveries: Essays on Mechanisms, Interfield Relations, and Anomaly Resolution*. Cambridge: Cambridge University Press.

Darden, Lindley and Nancy Maull. 1977. "Interfield Theories". *Philosophy of Science* 44 (1): 43-64.

DeCoteau, William E., and Raymond P. Kesner. 2000. "A Double Dissociation Between the Rat Hippocampus and Medial Caudoputamen in Processing Two Forms of Knowledge". *Behavioral Neuroscience*. 114 (6): 1096-1108.

Deneubourg, Jean-Louis., Simon Goss, Jacques Pastells, Dominique Fresneau and Jean Paul Lachaud. 1987. "Self-Organization Mechanisms in Ant Societies (II): Learning in Foraging and Division of Labor". *Experientia Supplementum* 54: 177-196.

Dennett, Daniel C. 1978. *Brainstorms*. Cambridge, MA: MIT Press.

Devito, Loren. M. and Howard Eichenbaum. 2011. “Memory for the Order of Events in Specific Sequences: Contributions of the Hippocampus and Medial Prefrontal Cortex”. *Journal of Neuroscience* 31 (9): 3196-3175.

Dragoi, George, and György Buzsáki. 2006. “Temporal Encoding of Place Sequences by Hippocampal Cell Assemblies”. *Neuron* 50 (1): 145-157.

Dragoi, George, and Susumu Tonegawa. 2011.“Preplay of Future Place Cell Sequences by Hippocampal Cellular Assemblies”. *Nature* 469: 397-401.

Dupré, John. 1993. *The Disorder of Things. Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.

Duval, Shelley, and Richard L. Tweedie. 2000. “Trim and Fill: A Simple Funnel-Plot-Based Method of Testing and Adjusting for Publication Bias in Meta-Analysis”. *Biometrics* 56 (2): 455-463.

Egger, Matthias, George Davey Smith, Martin Schneider, and Christoph Minder. 1997. “Bias in Meta-Analysis Detected by a Simple, Graphical Test”. *BMJ* 315 (629): 629-634.

Eichenbaum, Howard. “Memory on Time”. 2013. *Trends in Cognitive Sciences* 17 (2): 81-88.

——— 2014. “Time Cells in the Hippocampus: A New Dimension for Mapping Memories”. *Nature Reviews Neuroscience* 15 (11): 732-744.

——— 2017. “On the Integration of Space, Time and Memory”. *Neuron* 95 (5): 1007-1018.

Eichenbaum, Howard, Andrew R. Yonelinas, and Charan Ranganath. 2007. The Medial Lobe and Recognition Memory”. *Annual Reviews Neuroscience* 30: 123-152.

Ekstrom, Arne D., and Susan Y. Bookheimer. 2007. “Spatial and Temporal Episodic Memory Retrieval Recruit Dissociable Functional Networks in the Human Brain”. *Learning & Memory* 14 (10): 645-654.

Ergorul, Ceren, and Howard Eichenbaum. 2004. “The Hippocampus and Memory for ‘What’, ‘Where’, and ‘When’”. *Learning & Memory* 11 (4): 397-405.

——— 2006. “Essential Role of the Hippocampus in Formation in Rapid Learning of Higher-Order Sequential Associations”. *Journal of Neuroscience* 26 (15): 4111-4117.

Erkstrom, Arne D., Michael J. Kahana, Jeremy B. Caplan, Tony A. Fields, Eve A. Isham, Ehren L. Newman, and Itzhak Fried. 2003. “Cellular Networks Underlying Human Spatial Navigation”. *Nature* 425: 184-188.

Eysenck, H. J. “Problems with Meta-Analysis”. 1995. In *Systematic Reviews*, edited by Ian Chalmers and Douglas G. Altman, 64-74. London: BMJ Publishing Group.

Farovik, Anja, Laura M. Dupont, and Howard Eichenbaum. 2010. "Distinct Roles for Dorsal CA3 and CA1 in Memory for Sequential Nonspatial Events". *Learning & Memory* 17 (1): 12-17.

Fazekas, Peter. 2009. "Reconsidering the Role of Bridge-Laws in Intertheoretical Reductions". *Erkenntnis* 71 (3): 303-322.

Ferbinteanu, Janina, and Matthew L. Shapiro. 2003. "Prospective and Retrospective Memory Coding in the Hippocampus". *Neuron* 40 (6): 1227-1239.

Feyerabend, Paul K. 1962. "Explanation, Reduction, and Empiricism". In *Scientific Explanation, Space and Time, Minnesota Studies in the Philosophy of Science*, volume III, edited by Grover Maxwell and Herbert Feigl, 28-97. Minneapolis: University of Minnesota Press.

——— 1966. "The Structure of Science". *The British Journal for the Philosophy of Science* 17 (3): 237-249.

———1975. *Against Method*, London: New Left Books.

———1981. *Philosophical Papers*, Cambridge: Cambridge University Press, 1981.

Fodor, Jerry A., 1965. "Explanations in Psychology". In *Philosophy in America*, edited by Max Black, 161-179. Ithaca: Cornell University Press.

———1968a. *Psychological Explanation*. New York: Random House.

———1968b. "The Appeal to Tacit Knowledge in Psychological Explanation". *Journal of Philosophy* 65 (20): 627-640.

——— 1972. "Psychological Explanation". *Philosophy of Science* 39 (1): 95-97.

——— 1974. "Special Sciences: (Or the Disunity of Science as a Working Hypothesis)". *Synthese* 28 (2): 97-115.

———1975. *The Language of Thought*. New York: Thomas Crowell.

Folly, G., B. Hajtman, J. I. Nagy, and I. Ruff. 1981. "Some Methodological Problems in Ranking Scientists by Citation Analysis". *Scientometrics* 3 (2): 135-147.

Fortin Norbert J., Kara L. Agster, and Howard Eichenbaum. 2002. "Critical Role of the Hippocampus in Memory for Sequences of Events". *Nature Neuroscience* 5: 458-462.

Foster, Jonathan K. 2009. *Memory: A Very Short Introduction*. Oxford: Oxford University Press.

Fouquet, Céline, Christine Toubin, and Laure Rondi-Reig. 2010. "A New Approach for Modeling Episodic Memory from Rodents to Humans: The Temporal order Memory". *Behavioral Brain Research* 215 (2): 172-179.

Frank, Loren M., Emery N. Brown, and Matthew A. Wilson. 2000. "Trajectory Encoding in the Hippocampus and Entorhinal Cortex". *Neuron* 27 (1): 169-178.

Friedman, Kenneth. “Is Intertheoretic Reduction Feasible?”. 1982. *The British Journal for the Philosophy of Science* 33 (1):17-40.

Gazzaniga, Michael. S., ed. 2000. *Cognitive Neuroscience: A Reader*. Malden, Mass: Blackwell.

Gelbard-Savig, Hagar, Roy Mukamel, Michal Harel, Rafael Malach and Itzhak Fried. 2008. “Internally Generated Reactivation of Single Neurons in Human Hippocampus During Free Recall”. *Science* 322: 96-101.

Giere, Ronald N. 1999. *Science Without Laws: Science and its Conceptual Foundations*. Chicago: Chicago University Press.

Gill, Patrick R., Sheri J. Y. Mizumori, and David M. Smith. 2011. “Hippocampal Episode Fields Develop with Learning”. *Hippocampus* 21 (11): 1240-1249.

Ginter, Matthew R., Devin F. Walsh, and Seth J. Ramus. 2011. “Hippocampal Neurons Encode Different Episodes in an Overlapping Sequences of Odors Task”. *Journal of Neuroscience* 31 (7): 2706-2711.

Glymour, Clark. 1969. “On Some Patterns of Reduction”. *Philosophy of Science* 37 (3): 340-353.

“Google Scholar Metrics”. *Google Scholar*. Last updated July 2018. https://scholar.google.com/intl/en-US/scholar/metrics.html.

Gordan, Michael D. 1984. “How Authors Select Journals: A Test of the Reward Maximization Model of Submission Behavior”. *Social Studies of Science* 14 (1): 27-43.

Gorman, Michael E. 2002. “Levels of Expertise and Trading Zones: A Framework for Multidisciplinary Collaboration”. *Social Studies of Science* 32 (5/6): 933-938.

Grene, Marjorie, ed. 1969. *Toward a Unity of Knowledge*. New York: International Universities Press.

———1971. *Interpretations of Life and Mind: Essays around the Problem of Reduction*. London: Routledge & Kegan Paul, New York: Humanities Press.

“Guide to Nature Reviews”. *Nature Reviews*. Accessed April 27th, 2019. https://media.nature.com/full/nature-assets/reviews/pdf/naturereviews.pdf

Gurevitch, Jessica, Julia Koricheva, Shinichi Nakagawa, and Gavin Steward. 2018. “Meta-Analysis and the Science of Research Synthesis”. *Nature* 555: 175-182.

Hacking, Ian. 1983. *Representing and Intervening*. Cambridge: Cambridge University Press.

Hargens, Lowell L., and Diane H. Felmlee. 1984. “Structural Determinants of Stratification in Science”. *American Sociological Review* 49 (5): 685-697.

Hasselmo, Michael, and Chantal Stern. 2017. “Howard Eichenbaum (1947–2017)”. *Science* 357 (6354): 875.

Hastings, Janna, Gwen A. Frishkoff, Barry Smith, Mark Jensen, Russell A. Poldrack, Jane Lomax, Anita Bandrowski, Fahim Imam, Jessica A. Turner, and Maryann E. Martone. 2014. "Interdisciplinary Perspectives on the Development, Integration, and Application of Cognitive Ontologies." *Frontiers in Neuroinformatics* 8 (62): 1-6.

Hempel, Carl.G., 1969. "Reduction: Ontological and Linguistic Facets". In *Philosophy, Science, and Method: Essays in Honor of Ernest Nagel*, edited by Sidney Morgenbesser, Patrick Suppes, and Morton Gabriel White, 179-199. New York, NY: St. Martin's Press.

Higgins, Julian P. T., and Sally Green, eds. 2011. *Cochrane Handbook for Systematic Reviews of Interventions*. Version 5.1.0. Last updated March 2011. http://handbook-5-1.cochrane.org/

Howhy, Jakob. 2003. "Capacities, Explanation and the Possibility of Disunity". *International Studies in the Philosophy of Science* 17 (2):179-190.

Hooker, Cliff. 1981a. "Towards a General Theory of Reduction. Part I: Historical and Scientific Setting". *Dialogue* 20 (1): 31-59.

———1981b. "Towards a General Theory of Reduction. Part II: Identity in Reduction". *Dialogue* 20 (2): 201-236.

———1981c. "Towards a General Theory of Reduction. Part III: Cross-Categorical Reduction". *Dialogue* (3): 496-529.

Howard, Marc W., Mrigankka S. Fotedar, Aditya V. Datey, and Michael E. Hasselmo. 2005. "The Temporal Context Model in Spatial Navigation and Relational Learning: Toward a Common Explanation of Medial Temporal Lob Function Across Domains". *Psychological Review* 112 (1): 75-116.

"Instructions for Authors". *Trends in Cognitive Sciences*. Accessed January 17th, 2019. https://www.cell.com/trends/cognitive-sciences/authors.

Ioannidis John P. A. 2016. "The Mass Production of Redundant and Misleading, and Conflicted Systematic Reviews and Meta-Analyses". *The Milkbank Quarterly* 94 (3): 485-514.

Jadhav, Shantanu P., Caleb Kemere, P. Walter German, and Loren M. Frank. 2012. "Awake Hippocampal Sharp-Wave Ripples Support Spatial Memory". *Science* 336 (6087):1454-1458.

Jensen, Ole, and John E. Lisman. 2005. "Hippocampal Sequence-Encoding Driven by a Cortical Multi-Item Working Memory Buffer". *Trends in Neurosciences* 28 (2): 67-72.

Ji, Daovun, and Matthew A. Wilson. 2007. "Coordinated Memory Replay in the Visual Cortex and Hippocampus during Sleep". *Nature Neuroscience* 10 (1): 100-107.

Jick Todd D. 1979. "Mixing Qualitative and Quantitative Methods: Triangulation in Action". *Administrative Science Quarterly* 24 (4): 602-611.

Johnson, Adam. and A. David Reddish. 2007. “Neural Ensembles in CA3 Transiently Encode Paths Forward of the Animal at a Decision Point”. *Journal of Neuroscience* 27 (45): 12716-12189.

“Journal Citation Reports”. *Clarivate Analytics*. Accessed April 21st, 2019. https://clarivate.com/products/journal-citation-reports/.

“Journal Impact”. *Trends in Cognitive Sciences*. Accessed July 30th, 2019. https://www.bioxbio.com/journal/TRENDS-COGN-SCI

“Journal Report”. *Frontiers in Neuroscience*. Accessed July 30th, 2019. https://reports.frontiersin.org/reports/frontiers-in-neuroscience-report.html

Kandal, Eric R., James H. Schwartz, and Thomas M. Jessell, eds. 1991. *Principles of Neural Science*, 3rd edition. New York: McGraw-Hill.

Kandal, Eric R., James H. Schwartz, and Thomas M. Jessell, eds. 2000. *Principles of Neural Science*, 4th edition. New York: McGraw-Hill.

Kaplan, David. M. 2017. *Explanation and Integration in Mind and Brain Science*. Oxford: Oxford University Press.

Karlsson, Mattias P., and Loren M. Frank. 2009. “Awake Replay of Remote Experiences in the Hippocampus”. *Nature Neuroscience* 12 (7): 913-918.

Kemeny, John G, and Paul Oppenheim. 1956. “On Reduction”. *Philosophical Studies* 7: 6-19.

Kesner, Raymond P., Paul E. Gilbert, and Lindsay A. Barua. 2002 “The Role of the Hippocampus in Memory for the Temporal Order of a Sequence of Odors”. *Behavioral Neuroscience* 116 (2): 286-290.

Kesner, Raymond P., Michael R. Hunsaker, and Paul E. Gilbert. “The Role of CA1 in the Acquisition of an Object-Trace-Odor Paired Associate Task”. *Behavioral Neuroscience* 119, no. 3 (2005): 781-786.

Kim, Jaegwon. “Multiple Realization and the Metaphysics of Reduction”. 1992. *Philosophy and Phenomenological Research* 52 (1): 1-26.

——— 2008. “Reduction and Reductive Explanation. Is One Possible Without the Other?” In *Being Reduced. New Essays on Reduction, Explanation and Causation*, edited by in Jesper Kallestrup and Jakob Hohwy, 93-114. Oxford, UK: Oxford University Press.

Kincaid, Harold. 1996. *Individualism and The Unity of Science*. Lanham: Rowman and Littlefield.

——— 1997. *Philosophical Foundations of the Social Sciences*. Cambridge: Cambridge University Press.

Kitcher, Philip. 1981. “Explanatory Unification”. *Philosophy of Science* 48 (4): 507-531.

——— 1984. “1953 and All That: A Tale of Two Sciences”. *Philosophical Review* 93 (3): 335-373.

——— 1986. “Projecting the Order of Nature”. In *Kant's Philosophy of physical Science*, edited by Robert E. Butts, 201-235. Dordrecht: Reidel.

——— 1989. “Explanatory Unification and the Causal Structure of the World”. In *Scientific Explanation*, edited by Philip Kitcher and Wesley C. Salmon, 410-505. Minneapolis: University of Minnesota Press.

——— 1991. “The Division of Social Labor”. *Journal of Philosophy* 87 (1): 5-22.

Klee, Robert, ed. 1999. *Scientific Inquiry: Readings in the Philosophy of Science*. New York: Oxford University Press.

Klein, Colin. 2009. “Reduction without Reductionism: A Defence of Nagel on Connectability”. *Philosophical Quarterly* 59 (234): 39-53.

Kockelmans, Joseph J. 1979. *Interdisciplinarity and Higher Education*. University Park: The Pennsylvania State University Press.

Kogan, Jeffrey H., Paul W. Frankland, and Alcino Silva. 2000. “Long-Term Memory Underlying Hippocampus-Dependent Social Recognition in Mice”. *Hippocampus* 10 (1): 47-56.

Kolodny, Dave. 2019. “Welcome to NVIVO: The Leading Solution for Unlocking Insights in Unstructured Data”. *QSR International*. Online webinar accessed March 18th, 2019.

Komorowski, Robert W., Joseph R. Manns, and Howard Eichenbaum. 2009. “Robust Conjunctive Item-Place Coding by Hippocampal Neurons Parallels Learning What Happens”. *Journal of Neuroscience* 29 (31): 9918-9929.

Koonce, Joseph R., and Ana B. Locci. 1999. “Resolution of Issues of Scope and Detail in the Development of the Lake Erie Ecological Model”. *Proceedings of the Lake Erie Modelling Summit, International Joint Commission*.

Krain, Amy L., Amanda M. Wilson, Robert Arbuckle, F. Xavier Castellanos, and Michael P. Milham. 2006. “Distinct Neural Mechanisms of Risk and Ambiguity: A Meta-Analysis of Decision-Making”, *NeuroImage* 32 (1): 477-484.

Kraus, B.J., M. P. Brandon, M. A. Connerney, R. J. Robinson, S. Eriksson, C.P. Libby. J. A. White, M. E. Hasselmo, and H. Eichenbaum. 2012. “Medial Entorhinal Cortical Neurons Exhibit Temporally-Modulated Firing Patterns during Stationary Treadmill Running”. *Society for Neuroscience Annual Meeting*. Abstract 230.14/CCC23, poster presentation.

Kuhn, Thomas S. 1962. *The Structure of Scientific Revolutions*, Chicago: University of Chicago Press.

Kumaran, Dharshan, and Eleanor A. Maguire. 2006a. “An Unexpected Sequence of Events” Mistmatch Detection in the Human Hippocampus”. *PLOS Biology* 4. https://doi.org/10.1371/journal.pbio.0040424.

———2006b. "The Dynamics of Hippocampal Activation During Encoding of Overlapping Sequences". *Neuron* 49: 617-629.

Lakatos, Imre. 1970. "Falsification and the Methodology of Scientific Research Programmes". In *Criticism and the Growth of Knowledge*, edited by Imre Lakatos and Alan Musgrave, 91-195. Cambridge, England: University Press.

Lehn, Hanne, Hill-Aina Steffenach, Niels M. van Strien, Dick J. Veltman, Menno P. Witter and Asta K. Håberg. 2009. "A Specific Role of the Human Hippocampus in Recall of Temporal Sequences". *Journal of Neuroscience* 29 (11): 3475-3484.

Leimu-Brown, Roosa, and Julia Koricheva. 2004. "Cumulative Meta-Analysis: A New Tool for Detection of Temporal Trends and Publication Bias in Ecology". *Proceedings of the Royal Society B: Biological Sciences* 271, no. 1551 (2004): 1961-1966.

Levy, William B. "A Computational Approach to Hippocampal Function". 1989. In *Computational Models of Learning in Simple Neural Systems*, edited by Robert D. Hawkins and Gordon H. Bowers, 243-305. Academic Press.

Lewis, Penelope A, and Christopher R. Miall. 2003. "Distinct Systems for Automatic and Cognitively Controlled Time Measurement: Evidence from Neuroimaging". *Current Opinion in Neurobiology* 13 (2): 250-255.

'Literature Review Paper'. 2011. *The University of Wisconsin-La Crosse and the Board of Regents of the University of Wisconsin*. Last updated 2011. http://websites.uwlax.edu/biology/ReviewPapers.html

Longino, Helen. 1998a. *Science as Social Knowledge*, Princeton, NJ: Princeton University Press.

——— 1998b "Values and Objectivity". In *Philosophy of Science: Central Issues*, edited by Martin Curd and J. A. Cover, 170-191. New York: W. W. Norton & Company.

———2001. *The Fate of Knowledge*. Princeton, NJ: Princeton University Press.

Macar, Françoise, Hervé Lejeune, Michel S. Bonnet, Antonella Ferrara, Viviane Pouthas, Franck Vidal and Pierre A. Maquet. 2002. "Activation of the Supplementary Motor Area and of Attentional Networks during Temporal Processing". *Experimental Brain Research* 142 (4):475-485.

MacDonald, Christopher J., Kyle Q. Lepage, Uri T. Eden, and Howard Eichenbaum. 2011. "Hippocampal 'Time Cells' Bridge the Gap in Memory for Discontiguous Events". *Neuron* 71 (4): 737-749.

Maguire, Eleanor A., Richard S. Frackowiak, and Christopher D. Frith. 1997. "Recalling Routes around London: Activation of the Right Hippocampus in Taxi Drivers". *Journal of Neuroscience* 17 (18): 7103-7110.

Manns, Joseph R., Marc W. Howard, and Howard Eichenbaum. 2007. “Gradual Changes in Hippocampal Activity Support Remembering the Order of Events”. *Neuron* 56 (3): 530-540.

Marr, David, and Tomaso Poggio. 1977. *A Theory of Human Stereo Vision*. Cambridge, Mass: Massachusetts Institute of Technology Artificial Intelligence.

Mauk, Michael D., and Dean V. Buonomano. 2004. “The Neural Basis of Temporal Processing”. *Annual Review Neuroscience* 27 (1): 307-340.

Mayes, A. R., Isaac, C. L, Holdstock, N. M., Montaldi, D., Downes, J. J., MacDonald, C., Cezayirli, E. and J. N. Roberts. 2001. “Memory for Single Items, Word Pairs, and Temporal Order of Different Kinds in a Patient with Selective Hippocampal Lesions”. *Cognitive Neuropsychology* 18 (2): 97-123. doi:10.1080/02643290042000008.

McGivern, Patrick. 2008. “Reductive Levels and Multi-Scale Structure”. *Synthese* 165 (1): 53–75.

McNaughton, B. L., Barnes, C. A., Gerrard, J. L., Gothard, K., Jung, M. W., Knierim, J. J., Kudrimoti, H., Qin, Y., Skaggs, W. E., Suster, M., and K. L. Weaver. 1996. “Deciphering the Hippocampal Polyglot: The Hippocampus as a Path Integration System”. *Journal of Experimental Biology* 199 (1): 173-185.

“Measuring Your Impact: Impact Factor, Citation Analysis, and other Metrics: Journal Impact Factor (IF)”. 2018. *UIC University Library*. Last updated November 23rd, 2018. https://researchguides.uic.edu/if/impact

Mehta, Mayank R., Michael C. Quirk, and Matthew A. Wilson. 2000. “Experience-Dependent Asymmetric Shape of Hippocampal Receptive Fields”. *Neuron* 25 (3): 707-715.

Mehrholz, Jan, Simone Thomas, Bernhard Elser. 2017. “Treadmill Training and Body Weight Support for Walking Afterstroke”. *Cochrane Database of Systematic Reviews* 8 (CD002840): 1-228.

Miller, Michael B., John Darrell Van Horn, George L. Wolford, Todd C. Handy, Monica Valsangkar-Smyth, Souheil Inati, Scott Grafton, and Michael S. Gazzaniga. 2002. “Extensive Individual Differences in Brain Activations Associated with Episodic Retrieval are Reliable Over Time”. *Journal of Cognitive Neuroscience* 14 (8): 1200-14.

Mitchell, Sandra D. 2002. “Integrative Pluralism”. *Biology and Philosophy* 17 (1): 55-70.

———2003. *Biological Complexity and Integrative Pluralism*. Cambridge: Cambridge University Press.

———2009. *Unsimple Truths: Science, Complexity and Policy*. Chicago: University of Chicago Press.

Mitchell, Sandra D., and Michael R. Dietrich. 2006. "Integration without Unification: An Argument for Pluralism in the Biological Sciences". *The American Naturalist* 168 (S6): S73-S79.

Milton, Jami L. "The Loss of a Neuroscience Giant: Howard Eichenbaum". *Inscopix*. Last July 24th, 2017. http://media.inscopix.com/the-loss-of-a-neuroscience-giant-howard-eichenbaum.

Moed, Henk F., W. J. M. Burger, J. G. Frankfort, and Ton Van Raan. 1984. "The Application of Bibliometric Indicators: Important Field and Time-Dependent Factors to be Considered". *Scientometrics* 8 (3-4): 177-203.

Moher, David, Alessandro Liberati, Jennifer Tetzlaff, Douglas G. Altman. 2009. "Preferred Reporting Items for Systematic Reviews and Meta-Analyses: The PRISMA Statement". *PLoS Medicine* 6 (7). doi:10.1371/journal.pmed1000097

Morrison, Margaret. 2008. "Reduction, Unity and the Nature of Science: Kant's Legacy?". *Royal Institute of Philosophy Supplement* 83 (63): 37-62.

Moulines, Carlos-Ulises. 1984. "Ontological Reduction in the Natural Sciences". In *Reduction in Science: Structure, Examples, Philosophical Problems*, edited by Wolfgang Balzer, David A. Pearce and H-J. Schmidt, 51-70. Dordrecht, NL: Reidel.

Mulrow, Cynthia D. 1987. "The Medical Review Article: The State of the Science". *Annals of Internal Medicine* 106 (3): 485-488.

——— 1995. "Rationale for Systematic Reviews". In *Systematic Reviews*", edited by Iain Chalmers and Douglas G. Altman, 1-8. London: BMJ Publishing Group.

Nagel, Ernest. 1935. "The Logic of Reduction in the Sciences". *Erkenntnis* 5 (1): 46–52.

——— 1949. "The Meaning of Reduction in the Natural Sciences" In *Science and Civilization*, edited by Robert C. Stauffer, 99-135.Madison: University of Wisconsin Press.

——— 1970. "Issues in the Logic of Reductive Explanation". In *Mind, Science and History*, edited by Howard Evans Kiefer and Karl Milton Munitz, 117-137. New York: SUNY Press.

———1979. *The Structure of Science: Problems in the Logic of Explanation*. Cambridge: Hackett Publishing Company.

Naya, Yuji, and Wendy A. Suzuki. 2011. "Integrating What and When Across the Primate Medial Temporal Lobe". *Science* 333 (6043): 773-776.

Nickles, Thomas. 1973. "Two Concepts of Intertheoretic Reduction". *Journal of Philosophy* 70 (7): 181–201.

Hurley, Susan, and Alva Noë. 2003. "Neural Plasticity and Consciousness". *Biology and Philosophy* 18 (1): 131-168.

O'Malley, Maureen A. 2013. "When Integration Fails: Prokaryote Phylogeny and the Tree of Life". *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 44 (4, part A): 551-562.

Oxman Andrew D., and Gordon H. Guyatt. 1988. "Guidelines for Readings Literature Reviews". *Canadian Medicine Association Journal* 138 (8): 697-703.

Page Jr., Robert E., and Sandra D. Mitchell. 1990. "Self Organization and Adaptation in Insect Societies". In *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 2: 289-98. http://www.jstor.org/stable/193075.

——— 1998."Self Organization and the Evolution of Division of Labor". *Apidologie* 29(1-2): 101-120.

Pastalkova, Eva, Vladimir Itskov, Asohan Amarasingham, and Gyorgy Buzaki. 2008. "Internally Generated Cell Assembly Sequences in the Rat Hippocampus". *Science* 321 (5894): 1322-1327.

Paz, Rony, Hagar Gelbard-Sagiv, Roy Mukamel, Michal Harel, Rafael Malach, and Itzhak Fried. 2010. "A Neural Substrate in the Human Hippocampus for Linking Successive Events". *Proceedings of the National Academy of Science US* 107(13): 6046-6051.

Pearson, Christopher H., and Matthew P. Schunke. 2015. "Reduction, Explanation, and the New Science of Religion". *Sophia* 54 (1): 47-60.

Piccinini, Gualtiero, and Carl Craver. 2011. "Integrating Psychology and Neuroscience: Functional Analyses as Mechanism Sketches". *Synthese* 183 (3): 283-311.

Poldrack, Russell A., and Anthony D. Wagner. 2004. "What Can Neuroimaging Tell Us about the Mind? Insights from Prefrontal Cortex". *Current Directions in Psychological Science* 13: 177-181.

Poldrack, Russell A. 2006. "Can Cognitive Processes be Inferred from Neuroimaging Data?". *Trends in Cognitive Sciences* 10 (2): 59-63.

Popper, Karl R. 1959. *The Logic of Scientific Discovery*. New York: Harper Torchbooks, Harper & Row.

Repko, Allen F., and Rick Szostak. 2012. *Interdisciplinarity: Process and Theory*, third edition. Thousand Oaks: Sage.

Ramachandran, Vilayanur Subramanian, and Sandra Blakeslee. 1998. *Phantoms in the Brain: Probing the Mysteries of the Human Mind*. New York: William Morrow.

Ramachandran, Vilayanur Subramanian, and William Hirstein. 1998 . "The Perception of Phantom Limbs". *Brain* 121: 1603-1630.

Richardson, Robert C. 1979. "Functionalism and Reductionism". *Philosophy of Science* 46 (4): 533–558.

Rolls, Edumund T. 2010.“A Computational Theory of Episodic Memory Formation in the Hippocampus”. *Behavioral Brain Research* 215 (2): 180-196.

Rosenberg, Alexander. 2006. *Darwinian Reductionism: or, How to Stop Worrying and Love Molecular Biology*. Chicago: University of Chicago Press.

Rosenthal Robert. 1991. *Meta-Analytic Procedures for Social Research*, Beverly Hills: Sage.

Roskies, Adina L., Julie Fiez, D. A. Balota, Marcus Raichle, and Steve E. Peterson. 2001. “Task-Dependent Modulation in Left Inferior Frontal Cortex during Semantic Processing”. *Journal of Cognitive Neuroscience* 13 (6): 1-16.

Roskies, Adina. 2010. “Saving Subtraction: A Reply to Van Orden and Paap”. *British Journal for the Philosophy of Science* 61 (3): 635-665.

Ross, Robert S., Thackery I. Brown, and Chantal E. Stern. 2009. “The Retrieval of Learned Sequences Engages the Hippocampus: Evidence From fMRI”. *Hippocampus* 19 (9): 790-799.

Sadato, Norihiro, Alvaro Pascual-Leone, Jordan Grafman, Vicente Ibanez, Marie-Pierre Delber, George Dold, and Mark Hallett. 1996. “Activation of the Primary Visual Cortex by Braille Reading in Blind Subjects”. *Nature* 380 (6574): 526-528.

Sadato, Norihiro, Alvaro Pascual-Leone, Jordan Grafman, Marie-Pierre, Vincent Iban, and Mark Hallett. 1998. “Neural Networks for Braille Reading by the Blind”. *Brain* 121 (7): 1213-1229.

Salmon, Wesley C. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

———1989. *Four Decades of Explanation*. Minneapolis: University of Minnesota Press.

Schaffner, Kenneth F. 1967. “Approaches to Reduction”. *Philosophy of Science* 34 (2): 137–147.

——— 1974. “The Peripherality of Reductionism in the Development of Molecular Biology”. *Journal of the History of Biology* 7 (1): 111-139.

——— 1992. “Philosophy of Medicine”. In *Introduction to the Philosophy of Science,* edited by Merrilee H. Salmon, John Earman, Clark Glymour, James G. Lennox, Peter Machamer, J.E. McGuire, John D. Norton, Wesley C. Salmon, and Kenneth F. Schaffner, 310-344. Englewood Cliffs, NJ: Prentice-Hall.

———1993. *Discovery and Explanation in Biology and Medicine*, University of Chicago Press.

———2006. “Reduction: the Cheshire Cat Problem and a Return to Roots”. *Synthese* 151 (3): 377–402.

——— 2012. “Ernest Nagel and Reduction”. *Journal of Philosophy* 109 (8-9): 534–565.

Shamseer Larissa, David Moher, Mike Clarke, Davina Ghersi, Alessandro Liberati, Mark Petticrew, Paul Shekelle, Lesley A. Stewart. 2015. "Preferred Reporting Items for Systematic Review and Meta-Analysis Protocols (PRISMA-P) 2015: Elaboration and Explanation". *BMJ* 349. doi:10.1136/bmj.g7647.

Shapiro, Matthew L., Pamela J. Kennedy, and Janina Ferbinteanu. 2006. "Representing Episodes in the Mammalian Brain". *Current Opinion in Neurobiology* 16 (6): 701-709.

Schendan, Haline E., Meghan M. Searl, Rebecca J. Melrose, and Chantal E. Stern. 2003. "An FMRI Study of the Role of the Medial Temporal Lobe in Implicit and Explicit Sequence Learning". *Neuron* 37 (6): 1013-1025.

Schouten, Maurice Kenneth Davy, Huibert Looren De Jong, eds. 2007. *The Matter of the Mind: Philosophical Essays on Psychology, Neuroscience and Reduction*. USA: Blackwell Publishing.

"Scope and Mission". *Frontiers in Neuroscience*. Accessed July 30th, 2019. https://www.frontiersin.org/journals/neuroscience#about

Seglen, Per O. 1992. "The Skewness of Science". *The Journal of the American Society of Information Science* 43: 628-638.

——— 1994. "Causal Relationship between Article Citedness and Journal Impact". *The American Society for Information Science* 45 (1):1-11.

——— 1997. "Why the Impact Factor of Journals Should Not be Used for Evaluation Research". *British Medical Journal* 314 (7079): 498-502.

Sherman, Paul W. 1988. 'The Levels of Analysis'. *Animal Behavior* 36 (2): 616–619.

Shimamura, Arthur P., Jeri S. Janowsky, and Larry R. Squire. 1990. "Memory for the Temporal Order of Events in Patients with Frontal Lobe Lesions and Amnesiac Patients". *Neuropsychologia* 28 (8): 803-813.

Simpson, R. J. S., and Karl Pearson. Report on Certain Enteric Fever Inoculation Statistics". 1904. *The British Medicine Journal* 2 (2288): 1243-1246.

Sklar, Lawrence. 1967. "Types of Inter-Theoretic Reduction". *British Journal for Philosophy of Science* 18 (2): 109-124.

Sleutels, Jan. 2009. "Beyond Reduction: Philosophy of mind and Post-Reductionist Philosophy of Science. *International Studies in the Philosophy of Science* 23 (2): 233 – 236.

Spector, Marshall. 1978. *Concepts of Reduction in Physical Science*. Temple University Press.

Spiers, Hugo J., Neil Burgess, Tom Hartley, Faraneh Vargha-Khadem, and J. O'Keefe. 2001. "Bilateral Hippocampal Pathology Impairs Topographical and Episodic Memory but not Visual Pattern Matching". *Hippocampus* 11 (6): 715-725.

Spinner, Helmut F. 1973. “Science without Reduction”. *Inquiry*: *An Interdisciplinary Journal of Philosophy* 16 (1-4): 16 – 94.

Staresina, Bernhard P., and Lila Davachi. 2006. “Differential Encoding Mechanisms for Subsequent Associative Recognition and Free Recall”. *Journal of Neuroscience* 26 (36): 9162-9172.

Steele, J. Douglas, and Stephen M. Lawrie. 2004. “Segregation of Cognitive and Emotional Function in the Prefrontal Cortex: A Stereotactic Meta-analysis”. *NeuroImag*e 21 (3): 868-875.

Stewart, Lesley A., Mike Clarke, Maroeska Rovers, Richard D. Riley, Mark Simmonds, Gavin B. Stewart, Jayne F. Tierney, and Catrin Tudur Smith. 2015. “Preferred Reporting Items for a Systematic Review and Meta-Analysis of Individual Participant Data: The PRISMA-IPD statement”. *The Journal of the American Medical Association* 313 (16): 1657-1665. doi:10.1001/jama.2015.3656.

Suh, Junghyup, Alex Rivest, Toshiaki Takashiba, and Susumu Tonewaga. 2011. “Entorhinal Cortex Layter III Input to the Hippocampus is Crucial for Temporal Association Memory”. *Science* 334 (6061): 1415-1420. doi: 10.1126/science.1210125.

Sullivan, Jacqueline A. 2009. “The Multiplicity of Experimental Protocols: A challenge to Reductionist and Non-Reductionist Models of the Unity of Neuroscience”. *Synthese* 167 (3): 511-539.

——— 2010. “Realization, Explanation and the Mind-Body Relation: Editor’s Introduction”. *Synthese* 177 (2): 151-164.

——— 2014. “Is the Next Frontier in Neuroscience a ‘Decade of the Mind’?”. In *Brain Theory: Essays in Critical Neurophilosophy*, edited by Charles T. Wolfe, 45-67. Basingstoke, Hampshire: Palgrave Macmillan.

——— 2016. “Stabilizing Constructs through Collaboration across Different Research Fields as a Way to Foster the Integrative Approach of the Research Domain Criteria (RDoC) Project”. *Frontiers in Human Neuroscience* 10, no. 309. doi:10.3389/fnhum.2016.00309

——— 2017. “Coordinated Pluralism as a Means to Facilitate Integrative Taxonomies of Cognition. *Philosophical Explorations* 20 (2): 129-145.

Suppe, Frederick. 1974. *The Structure of Scientific Theories*. University of Illinois Press, Chicago.

Templer, Victoria L., and Robert R. Hampton. 2013. “Cognitive Mechanisms of Memory for Order in Rhesus Monkeys”. *Hippocampus* 23 (3): 139-201.

Thor, Donald H., and William R. Holloway. 1982. “Social Memory of the Male Laboratory Rat”. *Journal of Comparative and Physiological Psychology* 96 (6): 1000-1006.

Tofts, Chris, and Nigel R. Franks. 1992. “Doing the Right Thing: Ants, Honeybees and Naked MoleRats”. *Trends in Ecology and Evolution* 7 (10): 346–349.

Toulmin, Stephen. 1972. *Human Understanding*, vol. I. Princeton: Princeton University Press.

Tulving, Endel. 1984. “Précis of *Elements of Episodic Memory*”. The Behavioral and Brain Sciences, 7: 223-268.

Tulving, Endel, and Fergus I. M Craik, eds. 2000. *The Oxford Handbook of Memory*. Oxford: Oxford University Press.

Turbidy, Shannon, and Lila Davachi. 2011. “Medial Temporal Lobe Contributions to Episodic Sequence Encoding” *Cerebral Cortex* 21 (2): 272-280.

Uttal, William R. 2001. *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain*. Cambridge: MIT Press.

Van Orden, Guy C., and Kenneth R. Paap. 1997 “Functional Neuroimages Fail to Discover Pieces of Mind in Parts of the Brain”. *Philosophy of Science* 64 (4): S85-S94.

Van Overwalle, Frank. 2009. “Social Cognition and the Brain: A Meta-analysis”. *Human Brain Mapping* 30 (3): 829-858.

Van Riel, Raphael. 2011. “Nagelian Reduction Beyond the Nagel-Model”. *Philosophy of Science* 78 (3): 353–375.

Vargha-Khadem, Faraneh, David G. Gadian, Kate E. Watkins, Alan Connelly, Win Van Paesschen, and Mortimer Mishkin. 1997. “Differential Effects of Early Hippocampal Pathology on Episodic and Semantic Memory”. Science 277 (5324): 376-380.

Wallenstein, Gene V., Michael E. Hasselmo, and Howard Eichenbaum. 1998. “The Hippocampus as an Associator of Discontiguous Events”. *Trends in Neurosciences* 21 (8): 317-323.

Wimsatt, William C. 1972. “Complexity and Organization”. *PSA 1972: Proceedings of the Biennial Meeting of the Philosophy of Science Association*: 67–86.

———1980. “Reductionist Research Strategies and their Biases in the Units of Selection Controversy”. In *Scientific discovery: Case Studies*, edited by Thomas in T. Nickles, 213-260. Dordrecht: Reidel Publishing Company.

——— 2006. “Reductionism and its Heuristics: Making Methodological Reductionism Honest”. *Synthese* 151 (3): 445–475.

———2007. *Re-Engineering Philosophy for Limited Beings: Piecemeal Approximations to Reality*. Cambridge: Harvard University Press, 2007.

Wood, Emma R., Paul A. Dudchenko, R. Jonathan Robitsek, and Howard Eichenbaum. 2000. “Hippocampal Neurons Encode Information about Different Types of Memory Episodes Occurring in the Same Location”. *Neuron* 27 (3): 623-633.

Woodward, James F. 2003. *Making Things Happen: A Causal Theory of Explanation*. Oxford: Oxford University Press.

Yates, Frank and William. G. Cochran. 1938. "The Analysis of Groups of Experiments". *Journal of Agricultural Science* 28: 556-580.

Cochran, William G. 1954. "The Combination of Estimates from Different Experiments". *Biometrics* 10: 101-129.

Yin, Bin Zhu, and Andrew B. Troger. 2011. "Exploring the 4th dimension: Hippocampus, Time, and Memory Revisited". *Frontiers of Integrative Neuroscience* 5 (36): 1-5.

Curriculum Vitae

Name:	Jaclyn Lanthier
Post-secondary Education and Degrees:	Concordia University Montreal, Quebec, Canada 2007-2013 B.A.
	The University of Western Ontario London, Ontario, Canada 2013-2014 M.A.
	The University of Western Ontario London, Ontario, Canada 2015-2019 Ph.D.
Honors and Awards:	Dean's Entrance Scholarship, Western University 2013-2014
	Rotman Doctoral Entrance Scholarship, Western University 2016-2019
	Social Science and Humanities Research Council (SSHRC) Doctoral Fellowship 2016-2019
	Doctoral Excellence Research Award 2016-2019
Related Work Experience	Research Assistant Rethinking the Taxonomy of Psychology Conference 2016

Publications:

1. Lanthier, Jaclyn (2016). Outside Color: Perceptual Science and the Puzzle of Color in Philosophy MAZVIITA CHIRIMUUTA Cambridge: The MIT Press, 2015; 239 pp.; $42.00 (hardback). *Dialogue,56*(1), 190-192. doi:10.1017/S0012217316000433

2. Lanthier, Jaclyn (2017). Surfing Uncertainty: Prediction, Action, and the Embodied Mind CLARK ANDY Oxford: Oxford University Press, 2016; 401 pp.; $29.95 (hardback). Dialogue, 1-2, doi:10.1017/S0012217317000270