Toward a Naturalistically Explicable Folk Psychology

by

Deryck Simon Hougaard

Thesis presented in fulfilment of the requirements for the degree of Master of Philosophy in the Faculty of Humanities at Stellenbosch University.

Supervisor: Dr Tanya de Villiers-Botha

December 2019

DECLARATION

By submitting this thesis electronically, I declare that the entirety of the work contained therein is my own, original work, that I am the sole author thereof (save to the extent explicitly otherwise stated), that reproduction and publication thereof by Stellenbosch University will not infringe any third party rights and that I have not previously in its entirety or in part submitted it for obtaining any qualification.

Deryck Simon Hougaard

December 2019

ABSTRACT

This thesis investigates the gap between what our science says and how many theorists and everyday people have characterised how we conceive of mental states. I argue that looking at our folk psychology (FP) in light of an understanding of real-world, current science yields beneficial philosophical results. FP, in the iteration that I shall be concerned with here, refers to every person's ability to apply reason explanations to conspecifics' behaviour. I focus predominantly on the underlying processes that we allegedly pick out in our folk psychologising, those of beliefs and desires (the propositional attitudes). The reason for this focus lies in the gap between our intuitive beliefs and understanding of our mental processes and the picture painted by the empirical sciences. I first explore some issues concerning traditional theorising on the topic, before discussing current scientific research into pur cognitive processes in the form of predictive processing (PP) as advocated by Friston (2003, 2008, 2010), Hohwy (2013), and Clark (2016). PP depicts our brains not as passive, stimulus-driven organs, but as active constructors of our environment. An implication of this approach is that the way in which we represent the environment within our mind is different to how it is typically conceived within traditional FP. I also explore Hutto's (2008a; Hutto & Myin 2013a, 2017) claim that our minds are wired to be attuned to the environment in terms of minimal content, which allows me to develop a minimal conception of representation in terms of content, one in which direct correspondence between mental states and the supposed representation of the environment need not obtain for the mental to do causally efficacious work. I conclude that the beliefs and desires utilised in FP are socio-cultural impositions upon the neural substrate with no counterpart in reality. This has clear implications for our understanding of how we think about mental states within the cognitive sciences and philosophy of mind, if what we are aiming toward is a clarification of just what the mental is. Additionally, these new insights may ensure that the cognitive sciences are better informed about what it is that is being explained and where to focus further research concerning the mental.

OPSOMMING

In hierdie tesis word ondersoek ingestel na die gaping tussen huidige wetenskaplike modelle van kognisie en die alledaagse verstaan van hierdie prosesse wat deur beide leke en sommige teoretici aangehang word. Ek voer aan dat dit ons veral sal baat om ons volksielkundigevermoëns in die lig van toepaslike, onlangse wetenskaplike teorieë in heroënskou te neem. In hierdie skrywe maak ek gebruik ŉ verstaan van "volksielkunde" (VS) wat verwys na ons almal se vermoë om rede-verklarings vir ons eweknieë se optrede te bied. Ek fokus hoofsaaklik op die onderliggende kognitiewe prosesse waarna ons volksielkundige uitlatings oënskynlik verwys, te wete, onderliggende oortuigings en begeertes (proposisionele houdings). Die rede vir hierdie benadering is die gaping wat tans bestaan tussen ons intuïtiewe verstaan van hoe ons kognisie werk en die beeld wat die empiriese wetenskappe daarstel. Die doel is dus om om vas te stel hoe ons ons volkverstaan en –diskoers rondom ons verstandelike prosesse moet bejeën in die lig van hierdie nuwe wetenskaplike modelle. Eerstens word tradisionele volksielkundige teorieë bespreek en van die problematiek wat hiermee gepaardgaan uitgewys. Dan word die kontemporêre kognitiewe-wetenskaplike teorie van voorspellende verwerkings bespreek, soos voorgestel deur Friston (2003, 2008, 2010), Hohwy (2013) en Clark (2016). Die voorspellendeverwerkingsteorie beeld ons breine uit as aktiewe samestellers van ons omgewingsverstaan, eerder as die passiewe, stimulus-gedrewe organe van die tradisionele verstaan. Gevolglik blyk dit dat die wyse waarop ons kognitiewe prosesse in der werklikheid ons omgewings uitbeeld drasties verskil van die alledaagse beeld wat ons van hierdie fenomeen het. Hierdie insig ondermyn daardie kognisieteorieë wat op ons volksielkunde gebaseer is. Verder word die werk van Hutto (2008a; Hutto & Myin 2013a, 2017) gebruik om te toon hoe ons kognitiewe verwerkingsprosesse waarskynlik ingestel is op ons omgewing by wyse van minimale kognitiewe "inhoud".. Gevolglik voer ek aan dat daar nie ŉ direkte ooreenkoms tussen ons verstandelike afbeeldings en die omgewing hoef te wees vir ons om effektief met die wêreld om te gaan nie. Ek kom tot die gevolgtrekking dat die "oortuigings" en "begeertes" van tradisionele volksielkundige teorieë niks anders is as ŉ sosiaal-kulturele konstruk nie wat nie met die neurale werklikheid ooreenstem nie. Dus, indien ons ŉ akkurate beeld van ons kognitiewe funksionering wil vestig behoort ons die gebruik van hierdie konstruksies in ons kognitiewe en filosofiese teorieë te heroorweeg. Verder kan hierdie insig ook leiding bied ten opsigte van wat presies die fokus van die kognitiewe wetenskappe behoort te wees indien ons ons werklike, onderliggende verstandelike werking wil verstaan.

ACKNOWLEDGEMENTS

This was an extensive and trying undertaking, and I thank the following for laying the passage to the other side:

My supervisor, Dr Tanya de Villiers-Botha, who was exposed to more mental issues than I feel she was expecting when signing onto this. I'm eternally grateful for shepherding me across the line. I'll remember your efforts more than whatever it was that I was writing about.

My internal and external examiners: you'll never get your time back. Apologies.

Desré Barnard, for being the rock I attempted to be during your thesis writing. It would have been far darker without you there to lie about how okay everything was. I also learnt what to avoid by observing the trauma you went through. My thesis, and the experience thereof, wouldn't be if it weren't for you. You are the reason for my suffering, and the source of my strength.

Fabio Tollon, for being the type of intellectual friend who makes you a better and more honest academic. Intellectually, you are the agent, I am the patient.

The familial unit: Frikkie, my best friend and person who created what I am. This is what happens when you don't use the instruction manual. Sue, for not asking what my thesis is about, and seeming genuinely impressed when I mentioned anything about it. She also helped create what I am, but unwittingly used the incorrect instruction manual. Andrea, for keeping the parents occupied so that they didn't realise how long it took. My parents definitely used the correct instruction manual for her. I also thank the parentals for funding this batshit experience. Love you guys.

Baloo and Luci. Yes, the cats. Having one of those bastards walk over everything to get your attention is the best excuse to stop working. It is also happiness.

The local coffee shop, for making me feel as if I was socialising when I was doing everything but. And lastly: that Russian website.

TABLE OF CONTENTS

Page

LIST OF ACRONYMS

Action Coordination Routines	ARCs
Dispositional Folk Psychology	DFP
Extended Mind Hypothesis	EMH
Folk Psychology	FP
Free Energy Principle	FEP
Language of Thought	LoT
Local Indexical Guides	LIGs
Narrative Practice Hypothesis	NPH
Prediction Error Minimisation	PEM
Predictive Acquisition of Cognitive Capacities	PACC
Predictive Processing	PP
Radically Enactive, Embodied Account of Cognition	REC
Theory of Mind	ToM

Chapter One: Introduction

1. Beginnings

The science and understanding concerning the mind, consciousness, and their relation to the body, has progressed in leaps and bounds over the past century. This, however, does not take away from the fact that there is a difference between what our science says and what many have taken it to mean when applying it to phenomena in the world. In other words, individuals do not always have an accurate understanding of what the science is even saying, including philosophers. By way of example, one could ask a scientist whether Cartesian dualism[1] is dead, and she would more than likely answer in the affirmative. But, as Dennett (1991a) has so forcefully argued, we are fooling ourselves if we believe this to be the case. Even advocates of materialism struggle to emancipate themselves from speaking in terms that belie dualistic thinking[2]. The philosophy of mind literature has not in all instances managed to rid itself of its dualistic inclinations, a problem compounded if the relevant science is not taken into account when discussing the mental. This, in turn, has also potentially affected the layperson's ability to apply the science to ourselves and incorporate it into our general discourse. These considerations have led me to delve further into the literature surrounding a multitude of discourses on the philosophy of mind, which has arguably exerted an influence on our conceptualisations of the mind. How is our language and truncated understanding preventing a comprehensive image of the sciences relating to the mind/brain? How has this fostered outmoded ways of thinking? How has this thinking influenced us, even in our day-to-day interactions with others? Taking a cue from Dennett's logical behaviourism[3] (1991a; 2005), I argue that looking at our folk psychology (FP) in light of an understanding of real-world, current science would yield beneficial results for philosophers and laypeople alike.

[1] Cartesian dualism finds its origins in the writings of Descartes (2008), and maintains a view of the universe which has a dual substrate, that of the mental and the physical. It has been the main source of argumentation with regards to the mind-body problem through the years.

[2] Hence Dennett's formulation of what he calls Cartesian materialism (1991a: 107), a seemingly paradoxical concept which is nevertheless extant in the world. It is what we are left with once one throws away Cartesian dualism, but keeps the notion of a centrally-located, but still importantly material, area within the brain where consciousness is localised.

[3] More specifically, Dennett endorses the primacy of verbal reports in the introspection of mental states. It is all we currently have to be able to tease out the enigmas of the mind from a personal level.

FP, in the iterations that I shall be concerned with here, refers to every person's ability to apply reason explanations[4] to others' behaviour, and *also* to how we tend to conceive of the intrinsic processes in the brain which constitute the mental states that we tend to pick out in others. These two levels must not be conflated, and I shall focus predominantly on the underlying processes that we allegedly pick out in our folk psychologising, which can confuse and blur conceptual lines. The reasons for this focus lies in the gap between our intuitive beliefs and understanding of our mental processes and the picture painted by the empirical sciences. Much like dualism still often has a stranglehold on our everyday intuitions, I argue that the way in which we view and utilise our FP is lacking in theoretical accuracy. Therefore, my project will be to incorporate the sciences and scientifically-inclined philosophy into further elucidations on how our minds potentially work. But my efforts will not be as general as that, as my project will be aimed at exploring how the folk conception of our minds and the social discourse surrounding it should be understood in light of these highly important fields of research. The purpose of this will be to tie the empirical, the philosophical, and the everyday conceptions of ourselves closer together. The benefits of this would be a clearer scientific understanding of the mental states we ascribe to each other, what they pick out in the world in juxtaposition to what we think we are doing.

I contend that philosophy is in the process of clarifying, refining, and elucidating the conceptual hurdles which we come across within our exploration of our world[5], and nothing more. To clarify, the philosophical discipline should reach toward descriptive claims in assisting the empirical sciences, and normative claims within the political, legal, and moral realms. To my mind, the good work that philosophy does can be subsumed into variations of the above categories. Incorporating this study into the preceding, reconceptualising the descriptive claims of what we are claiming mental states to be is a worthy extension of this broad approach to philosophy. Furthermore, and central to the aims of this thesis, I shall endorse a naturalistic worldview, which, in its broadest and most general sense, advocates theories which fit our best natural sciences.

More specifically, I will endorse a form of pragmatic physicalism, which limits what we take the universe to be made of, but provides a far more practical platform for conceptualising our

[4] Reason explanations refer to the narrative reasons we all utilise for the explanation of the behaviour of conspecifics.
[5] By "world" I mean more than the terrestrial, but instead refer to the broader implications of our lived experience in relation to the universe.

reality[6]. Physicalism, in its most broad iteration, refers to the idea that nature consists of only one kind of thing, i.e. the physical. Pragmatism, on the other hand, points to the philosophical position of treating the majority of fields and issues within philosophy in such a way as to focus on their practical benefits. There is an emphasis on applicability toward human experience, and thought directed toward this end need not necessarily represent reality as such.

This means that there is space for indeterminate positions in our epistemology and ontology, which need not affect the task of the pragmatic philosopher unnecessarily. An indeterminate position is that which holds a fundamentally undecidable element within its makeup, where the borders are not easily drawn. To be pragmatic is to bracket issues of contention if they hamper effective conceptualisations. Pragmatic physicalism has become more prevalent within the philosophical literature. The "pragmatic turn" (Engel *et al*. 2013: 202) in philosophy of mind has shown a profusion of philosophers taking up the mantle of pragmatism and applying it to their theories (Ullman 1991; Varela *et al*. 1991; Churchland & Ramachandran 1994). This change in theoretical perspective has been due to the advantages provided by the practicality of shedding the straight-jacketing effect of sticking to conceptualisations overly reliant on mental "representation". While not every philosopher has shifted toward this perspective, it is certainly true that many have thrown out or amended more traditional approaches to the mental with regards to concepts such as content, representation, and the like. Such traditional approaches often tend to focus on the content-rich nature of how the environment is represented within the mind. Content, in this technical sense, refers to the accuracy conditions of our representations. It is therefore taken to be a marker of how accurate our sensory impingements are in depicting the world. This thesis will advocate a minimal conception of representation, one in which direct correspondence between mental states and the representation of the environment need not obtain for the mental to do causally efficacious work[7]. I will claim that content-rich representations of mental states is a conceptualisation that has grown vastly out of hand over the decades, and at the very least needs to be reined in. Pragmatic philosophers have less need for content-rich representations, and may relax the pursuit thereof to explicate mental states (or the content inherent within them). This thesis shall go some way to elucidate the *why*

[6] My position is reductive or deflationary with regards to specific phenomena such as the special properties of the mental and propositional thinking (as explained below). I therefore endorse a view which does not elevate the importance of these phenomena.

[7] As will be explained further on and throughout this thesis, this deflates the notion of a content-rich representational theory of mind to something more akin to selective bodily engagement with the environment. This is all without a corresponding internal "code" standing in for what the mental state is meant to be representing.

and the *how*. As we shall see, mental representation does not seem to do as much work within our cognitive processing as has been traditionally conceived. Hence, pragmatic approaches to the mental can potentially assist us in elucidating a more nuanced understanding of the continual interaction between mind, body, and world than most approaches that have come before.

Not all theories within the philosophy of mind are the antithesis to what I propose, with some softer approaches ideal for integration, but some ideas must be done away with. For the purposes of my thesis, chief among these is the ascription of propositional attitudes (see below), as well as the conceptual content that such attitudes are purported to contain, *at least when it comes to most modes of cognition*. In other words, there is a possibility that there is not as much content in our apparent propositional thinking as has been proposed, and even the existence of the majority of propositional attitudes themselves are up for debate. Propositional attitudes, in their full-blooded guise, are closely aligned with content-rich representational theories of mind. If content-rich representational theories of mind are in themselves problematic, a deflationary approach to propositional attitudes is warranted. To tackle propositional attitudes is to tackle any representational theories of mind which are aligned to an internalised "language of thought", trading in content that closely corresponds with that which it represents. I shall propose, after engaging with the literature on fine-grained mental processing and bodily engagement with the environment, that "propositional attitudes" are a gloss on what is actually instantiated at the neuronal level. It will be shown that propositional attitudes are more post-hoc rationalisations of fine-grained mental processes that deal with consistent uncertainty. This will not result in an eliminativism with regards to propositional attitudes in terms of how we use them in day-to-day discourse, but it will reconceptualise what we are actually speaking about when we ascribe propositional attitudes to others. This has clear implications for our understanding of how we think about mental states within the cognitive sciences and philosophy of mind, especially if what we are aiming toward is a clarification of just what the mental is. In addition to this, these new insights may ensure that the cognitive sciences are better informed about what it is that is being explained and where to focus further research concerning the mental.

2. FP as Propositional Attitude Ascription

The importance of FP for the philosophy of mind has arisen as an expression of the issue of how it is that we can reliably know other peoples' mental states (in other words, how we can read their minds) (Stich & Nichols 2003: 1). This issue relates to the complications surrounding other peoples' minds and, more specifically, being able to extrapolate their mental states through their behaviour. In the philosophical field, FP is said to consist in our everyday ability to ascribe mental states to others in terms of their "beliefs", "desires", and so on. More specifically, a particularly influential view of FP has been that of conceiving of this ability in terms of propositional attitude ascription (Hutto 2008a: ix). A propositional attitude, first and foremost, is understood as an intentional state. Intentionality, within philosophy, refers to the directed nature of mental states, or the attitudes thereof: what they are "about" or "directed toward" (Hutto 2008a: 1). *Propositional* attitudes are typically in the form of a platitude such as "S believes that p", where the agent in question is "S", and "p" is a particular proposition, such as "this is long-winded" (Hutto 2008a: 1). Propositions are believed to be content-rich in their relation to the world, in that they closely correspond that which they pick out in the world. This is in contrast to the issues alluded to above by the more pragmatic school of thought, which advocates relaxing the need for mental states that correspond accurately to that which they refer. The problem of propositional attitude ascription will be analysed in detail in Chapter Two, but for now it is sufficient to know that this is the paradigmatic understanding that has informed much of the debate over FP. In other words, traditional FP presupposes accurate, inherent content that corresponds to propositional attitude ascription. In light of this, I shall proceed by setting up my argument for a view of FP which may rely on propositional attitudes as central to its operation, but does not claim to pick out what actually occurs in reality. As described earlier, content-rich propositional attitudes are looking increasingly hollow in their scientific accuracy.

To develop my argument, a suitably comprehensive overview of FP and of the related cognitive processes and their interactions is needed. I shall be drawing from three broad philosophical and scientific fields in order to develop such an overview, namely:

(1) The philosophical literature that has built up surrounding FP or, more accurately, *what has been surmised it is we do when we folk psychologise*. My discussion will serve as the grounds from which these views shall be problematised. Furthermore, I shall incorporate a potentially more accurate depiction of just what enables our FP abilities in the first place.

(2) The latest paradigm in neuroscientific theory as put forward by proponents of predictive processing (Clark 2013a, 2016; Hohwy 2013), which will be shown to amend the theories mentioned in (1) above by undermining the view of FP as propositional attitude ascriptions as accurately reflecting reality. Predictive processing appears to provide a scientifically more valid depiction of our fine-grained mental states, in opposition to propositional-style depictions of the very same.

(3) An enactivist account of cognition, primarily form the writings of Hutto and Myin (Hutto 2008a; Hutto & Myin 2013a, 2017; Hutto & Satne 2015). Enactivism, to be explained in Chapter Four, is an umbrella term which groups all theories that endorse a view that depicts cognition as a dynamic interaction between body and environment. This specific enactivist account argues against the plausibility of content-rich propositional thinking by advocating a minimalist view of the content of mental states. This theory is a compelling and plausible alternative to traditional thinking on the nature of the intentionality of the mental. This section shall be further subdivided into:

    a. A brief sketch of the problem of content ascription in the majority of our folk psychological reason explanations.
    b. An outline of the development of enactive theories of content. This shall then segue into a thorough unpacking of teleosemiotics (to be explained further below) as the most suitable theory of content for most forms of cognition, including supposed propositional thinking as depicted in FP.

I shall end by introducing the potential determinants of how we pick out our apparent folk psychological propositional attitudes. But, for now, in terms of the problems surrounding FP as a theory of our mental states, there is a great worry that if the problems highlighted here hold weight, there is a high possibility that FP could be *false*. If the platitudinous ascriptions it envisions are false, as eliminativists about propositional thinking would have us believe (such as Churchland (1981)), we may have reason to do away with FP in terms of propositional attitudinal thinking. Perhaps the fixed content of the propositions which are said to pick out our mental states are not well designated by our folk psychologising, as propositional attitudes are potentially of an etirely different nature[8]. If this is true, if we cannot achieve empirical validity for our claims, how can it be said that we "correctly" folk psychologise? In other words, if the

[8] I remain neutral for now as to whether this is a knock-down argument for our use of propositional attitudes within our day-to-day discourse.

propositional attitudes that we ascribe do not pick out reality, then we are clearly misunderstanding what is going on in people's "minds". Does this mean that our everyday folk psychologising is wrong? With this understanding of the resultant debates surrounding FP, philosophers are in need of formulating a more precise account of the nature of FP and what it picks out.

With this in mind, Chapter One will first situate FP historically, highlighting the specific ways in which the concept "folk psychology" has been utilised in the literature. The focus will primarily be on FP as propositional attitude ascription—what it is and what a couple of theorists on FP have specifically proposed in this regard. A critique of FP understood as propositional attitude ascription will follow in Chapter Two. Thereafter, in the same chapter, I shall describe how narratives during our ontogenetic development, those laden with folk psychological terminology and concepts, appear to be the basis from which we develop our habit of folk psychologising. This implies that what occurs during our folk psychologising has more to do with what is forced upon us in our youth rather than with accurate hypothesising as to someone's actual mental states. Social norms are communicated and internalised through these narratives, and inform our propositional thinking. This theory is articulated by Hutto (2007a; 2007b; 2008a; 2008b) as the Narrative Practice Hypothesis (NPH), and it opens the door for my own views on the social realm and how our FP is propped up primarily through a form of social scaffolding or the norms through which social communication holds. Further implications of this will be explored in Chapter Four when I tie the chapters together with regards to the interrelations between the themes of this thesis.

3. Predictive Processing

In Chapter Three, I shall investigate the ever-growing empirical work on the predictive capacities of the mind. I shall focus on the latest research in the neuro- and cognitive sciences relating to predictive processing. As the central focus of predictive processing is about the sufficiently accurate prediction of the organism's environment, it overlaps with what we do from the macro-perspective of social mindreading. I introduce this body of work to show that our base mental states are not as semantically-loaded as previously thought. In addition to this, it is a prevalent discipline in the sciences which has been picking up steam over the past number of years, with an empirically strong foundation (Clark 2013a; Hohwy 2013; Seth 2015). Fundamentally, it is an empirical view of the underlying processes of the mind, which depicts

the brain as a predictive machine first and foremost. Granted, the explanandum here is more fine-grained than our ability to predict the behaviour of conspecifics, but there is a case to be made that these underlying processes influence the nature of our macro-predictions. It has fundamental implications for our understanding of the role of propositional attitudes within our reason explanations, depicting propositional attitudes as coarse-grained impositions upon the more probabilistic nature of our cognitive structures. "Probabilistic", here, denotes that our mental states and processes are more concerned with finding likely environmental causes with a given range of inputs.

As I work through what predictive processing entails, it shall become clear that I endorse a particular formulation of it, as proposed by Clark (2013a, 2016). The central idea of this model of the mind is the unification of cognition, perception, and action, as well as their interrelations. This is achieved through the application of a broad, computational paradigmatic principle, providing a conceptually coherent narrative in terms of which cognition can be elucidated. The important repercussions of this for FP relates to the effect that this theory would have on the extant view of our FP as accurately designating propositional attitudes on the cognitive level. If we are seeking to elucidate and explain reason explanations, and the existence thereof, in our folk psychologising, the predictive processing framework (if accurate) should inform any suppositions that we make in this regard. As we shall see, general propositional attitudes such as "beliefs" and "desires" do not easily fit into the predictive processing picture *on the fine-grained cognitive level*. This leads some proponents of predictive processing (Clark 2016; Dewhurst 2017) to endorse a view of cognitive processing that depicts the effect of probabilistic models of the environment instantiated in the brain as what is *actually* being picked out by our folk psychologising. As I shall argue, while this need not significantly alter our everyday habit of folk psychologising as such, it significantly undermines the philosophical plausibility of much formal theorising on the topic. General FP may be a coherent heuristic in our social lives, but it should not be taken as empirically well-founded. In addition, the predictive processing framework could eventually seep into public discourse, altering the way that we describe our reason explanations. This, however, is not something that can be predicted with any certainty as of yet.

Generally, the predictive processing view of cognition does not describe neurophysiological phenomena *as such*, but rather describes the computational capacities of the brain, the understanding of which can *then* be used to describe scaled-up behaviours of both brain and body (Clark 2013a: 187). This means that, on the predictive processing view, the possibility

remains that the sub-personal could potentially alter phenomena on the agentive level in ways that we have not yet grasped. As already mentioned, the predictive processing picture undermines the status of our folk psychological concepts as concepts that accurately pick out something in the world. In addition, predictive processing could potentially span the conceptual gap between science and folk understanding, and also do positive work on the (perpetually rickety) bridge between the hard and soft sciences. This conceptual gap concerns the mismatch between understanding how things *actually* work in the world (determined through the empirical sciences), as opposed to the everyday conception of how things tend to work (determined through heuristic social learning). I contend that aligning folk conceptions of the world more closely with the empirical would help at insulating ourselves from the harmful effects of misattributed causal reasoning[9]. This type of empirical conceptual clarification should be cause for celebration, and much hope has been placed on it by its more philosophical adherents (Hohwy 2013; Clark 2016).

As we shall see, predictive processing is at 180 degrees to the more conventional idea of the brain as a passive sensory machine, driven purely by stimulus. Predictive processing endorses a view of the brain as constructively perceiving its environment in an active and continually selective way. This involves *top-down* processing when assimilating environmental stimulations, as opposed to the classical view of stimulus-driven, *bottom-up* processing. This process gets cashed out in the form of continually formulated *predictions* of imminent environmental sensory experiences (Clark 2013a: 181-182). Any errors in prediction provide updates of existing representations for incorporation into future predictions, which is known as prediction error minimisation (ibid.: 186). This occurs, for the most part, below the level of phenomenal experience. Simply put, these predictions stand in the place of representations of the environment and are utilised by the brain for predicting future real-world sensory impingements, as well as their potential sources. Any mismatch between the prediction and the actual sensory input lead to prediction errors, which then need to be rectified. These predictive models replace the content-rich representations of the environment of classical theories. In Chapter Three, I will argue for minimally contentful basic predictive processing. This will serve to call into question the existence of the majority of propositional mental states.

[9] One need look no further than how the elucidation of the complex interactions of neuronal wirings has affected our understanding of personality type and the effects thereof. For a typical example, elucidating processes at the neuronal level has helped explain the reasoning ability of autistic and similar individuals on the spectrum.

Moreover, predictive processing is also understood to emphasise action over perception, leading many to describe predictive processing as *action-oriented*, with perception mainly serving as the means by which knowledge is acquired from the environment in order to enable contextually-efficient action for organismic survival[10] (ibid.: 185-186). I shall be going into further detail on this highly important point in Chapter Three. Briefly, an important idea that underlies all of the above, and which helps explain the evolution of the brain toward predictive processing, is that of the free energy thesis as developed by Friston and colleagues (Friston 2003, 2009, 2010; Friston & Stephan 2007). The *minimisation* of free energy lies at the heart of the minimisation of prediction error within the neuronal dynamics of the predictive processing systems of our brain. This minimisation of free energy is thought to be the "path of least resistance" in the gradual evolutionary development of the mind/brain. It is free energy minimisation which is thought to drive much of the processing of the mind, whether at a fine-grained level or at a more macro-perspective. A take-home point from Friston's thesis is the *frugal* nature of the brain that we are left with (Friston 2010: 131-132; Clark 2015c: 14-18). If it is true that free energy is kept to a minimum, a view of the brain which utilises content-rich representation of the environment, in the form of propositional thinking, is further undermined. My discussion of predictive processing in Chapter Three will serve as necessary background for my critique of FP as traditionally conceived. I will provide a brief overview of some of the literature on predictive processing (PP) as a foil to my view of FP, which is somewhat at odds with tradition in that I will be downplaying the existence of propositional attitudes as actual cognitive mental states. I shall go into detail concerning the characteristics of this view of FP to provide a suitable platform for an understanding of FP as conventionally understood. The background concerning predictive processing shall serve to elucidate a view of brain processing that I *do* endorse due to its empirically-grounded approach. Once these fields of research have been adequately unpacked, the possible effect that predictive processing could have on our understanding of FP shall be expanded upon. Specifically, I will focus on the ways in which predictive processing would broadly affect our view of the scientific plausibility of propositional attitude psychology. To take our folk understandings of belief and desire as prime examples, their clearly delineated and determinate nature when utilised within folk psychological discourse would not readily fit into the more fine-grained, probabilistic framework endorsed by predictive processing, but this shall become clearer in Chapter Three.

[10] This also fits neatly into how many enactivist theorists view the role of cognition to be.

4. Teleosemiotics

In Chapter Four, a further view of the mind as endorsed by Hutto (primarily in Hutto 2008a; Hutto & Myin 2013a, 2017, Hutto & Satne 2015) will be explicated. Primarily, this is an enactivist view of basic[11] cognition and its relation to the apparent content of such cognition. The need for introducing this field of research lies in how it explains what is represented within our mental states, and in particular how it affects our understanding of propositional attitudes. Predictive processing goes a long way in subverting our understanding of how our mental world is constructed, but an enactivist-style teleosemiotic approach adds to the debate by depicting our base mental states as only minimally contentful (Hutto 2008a; Hutto & Myin 2013a, 2017). Minimum content refers to the arguable fact that we need only have few accuracy conditions within our mental states which delineate how a mental state can misrepresent something in the world. If this view is correct, this has negative ramifications for propositional attitude psychology. If there is minimal content on not only the sub-personal level, but also the personal, agentive level, this inevitably affects traditional FP theories. As already mentioned, propositional attitude psychology predominantly entails a view of cognition that is constituted by suitably rich content of the environment and, as such, cannot operate with minimal content. Recall that "content" refers to the accuracy conditions of propositions, and the accuracy conditions of propositions require rich content to be able to say anything about the world, in the sense of whether they accurately represent what is in the world or not. Minimal content theories have clear ramifications for our understanding of the cognitive paradigm, and – along with the problematisation that predictive processing also places on the propositional attitudes – alters our understanding of our cognitive selves.

To explicate this view of minimal content in our cognition, Chapter Four will proceed in a series of steps. Firstly, the problem of content-ascription for the majority of our cognitive processes will be introduced and expanded upon. Thereafter, a step-by-step guide through the literature concerning the emergence of alternative enactivist theories of cognition will be unpacked in order to provide a suitable grounding for a more minimal construal of cognitive content in the brain. Then the problem concerning one of the most prominent theories of mental content, namely teleosemantic theory (primarily from Dretske 1988; Millikan 1984; Papineau 1987), will be unpacked, and I will try to address this problem by means of my own amendment to the theory. The reason for introducing teleosemantic theories of mental content lies in the

[11] As will be reiterated in Chapter Four, "basic" here refers to the foundational processing of our minds, and not to a "simplistic" mind.

fact that they are the most accepted naturalised theories of content, while still not going quite far enough in their purported project of naturalising content. The "-semantic" suffix refers to the inherent *content* of our mental states. Teleosemantic theories focus on the "normative" aspects of this content within our intentional states, therefore relating them to *function* first and foremost, in that our intentional states are directed toward distal stimuli for a particular reason (Hutto & Myin 2013a: 74). In teleosemantic theories a particular intentional attitude serves a "normative" function through appeal to the "correct" teleological reason for its existence. The evolutionary history of the intentional state, for example, embeds its function into the cognitive makeup of the organism (ibid.: 76). Teleosemantic theories can be understood as general approaches as opposed to specific theoretical explanations of how minds obtain content in the first place. This content serves the function of representing the world in an accurate fashion, which makes representation a key element within theorising of this ilk, with propositional thinking often being central (ibid.: 75). The semantics of our internal mental states are seen as derivative of these originary functions.

Taking this understanding, I shall introduce and argue for an amendment to these teleosemantic theories. This shall take the form of a teleo*semiotic* account of content, as endorsed by Hutto (2008a; Hutto & Myin 2013a; 2017). Teleosemiotics, on Hutto's construal, performs a similar role to teleosemantic theories in terms of intentionality, but he emphasises a lack of semantic content within the intentional attitude of the organism. According to the teleosemiotic approach to mental content, when an organism tracks features of the environment, semantic content need not exist. This view of the brain as lacking in semantic content sits well with the predictive processing view as minimal content speaks to the need for minimal processing power.

Therefore, once the suitable groundwork has been laid, I will expand on theories of brain processing which may appear to be counterintuitive at first glance, but which will be shown to potentially have more scientific validity than theories that make use of tacit, semantic, representational, propositional attitude ascriptions. Here, I will specifically explore the plausibility of the claim that extensive content is actually instantiated within our mental states. This is why I will first present an argument for minimal-to-no- semantic content for the majority of our cognitive functions by backing this with an amendment to teleosemantic theories of content which, while admirable in their naturalism[12], have a supposition at their core that renders them problematic in their veracity. This will serve as a further basis for my

[12] By admirable I mean their predilection for empirically-valid scientific theories of the human mind.

thesis, which endorses a view of content in day-to-day cognition that is minimal and fundamentally of a more covariant nature. Covariance speaks to a relational function between stimuli and mental states, rather than informationally-accurate content. How this relates to FP is in the way that minimal content acquisition affects the alleged content of the propositional attitudes. Here, I will draw primarily from the work of Hutto and Myin (Hutto 2008a; Hutto & Myin 2013a, 2017; Hutto & Satne 2015), as their minimal theory of content appears to cohere well with the predictive processing account of the brain with its emphasis on the minimisation of free energy. Hutto, specifically, views content as minimal, as he is, by his own admission, a radical enactivist, endorsing a view of cognition as constituted through highly dynamic interaction between the body and environment with minimal content acquisition in the majority of our intentional states. This pushes him to endorse a highly pragmatic, naturalised theory of cognition, whereby the mind/brain utilises resources in a frugal manner to incorporate features of the environment, be it the physical or social environment. On top of this, action and behaviour come about in service of the survival of the organism, wherein accuracy conditions do not allow for the necessary frugal processing, linking it with the literature on PP.

Onn Hutto's view, talk of content representation is highly problematic as it appears to be adding phenomena to the brain which need not exist (Hutto & Myin 2013a, 2017). Throughout the literature there is an understanding (indeed a fundamental assumption) that mental content is sourced from the environment through the senses, taken into the brain, and then represented in some fashion for further discrimination, either immediately or in the form of episodic memory and the like. It is a view, whether tacitly or overtly accepted, that envisages the brain as typically operating in more-or-less input-output terms. This is the cognitivists' line of thinking, which treats the brain as a closed circuit of mental processing that has as its primary function the processing of a continual stream of sensory information (Branquinho 2001: xv). Hutto rejects this wholeheartedly, with very few reservations. In Chapter Four it will become clearer why he does so, and I shall be using this core divergence from the conventional cognitivist view of cognition as the primary hook from which to hang the rest of this thesis. The conventional cognitivist who thinks that we continually take in sensory information (which informs our propositional attitudes) is diametrically opposed to Hutto's view, as well as to the predictive processing account of cognition. Hutto (Hutto & Myin 2017; Chapter Five) argues that the core approach of the teleosemantic view of mental content certainly does do good work, but that we need to go a little further and ditch talk of "representation" altogether. Teleosemantic theories are useful for explaining the base tracking of environmental stimuli,

but there is no need to posit additional semantic properties over and above this to elucidate what we are trying to explain. Instead, according to him, we should appeal to a form of ur-intentionality (Hutto & Myin 2017, Chapter Five; originally introduced in Hutto & Satne 2015), which strips intentionality down to its core instrumental functionality of being focused on distal and internal stimuli, but without richly contentful representation thereof. He proposes that our mental states do not represent content in as detailed a fashion as is generally assumed. If our mental states are not saturated with detailed content and operate on a more streamlined and visceral level of interaction with the environment, then it is difficult to support a view that insists on their accurate representation of the environment. For our purposes, this is of interest, as this would apply to the content of supposed propositional attitudes as well. On the predictive processing construal, with its streamlined approach to sensory intake, it seems plausible that at most an attenuated form of representation is taking place.

5. The Social Mind

The view on representation discussed above will inform my view on FP, predictive processing, and my views on how our cognition is shaped in general, which is what I will address towards the end of this thesis. As mentioned above, utilising Hutto's views, in Chapter Four, I shall use a predictive processing lens along with teleosemiotics to propose a scientifically-plausible view of our mental states that has no need of full-blown representation and hence of propositional attitudes. This is all in order to overturn the philosophical biases that have crept into our common discourse and to offer an alternative view of our minds which is more scientifically grounded than has traditionally been the case in philosophy of mind. If we accept a model of our cognitive processes where these primarily involve predictive processing and minimal content, we shall need to revise our understanding of propositional attitude ascriptions and what role they play in our folk psychologising. However, I shall argue that this does mean that we need to do away with our practice of *folk psychologising* as such, partly because propositional attitude ascription is but one stone in the greater edifice that is our folk psychologising. Furthermore, everyday folk psychologising works, in the sense that it allows us to get by in our social environment. However, we need to be very careful not to conclude that this success is due to us accurately picking out beliefs and desires in others.

To end this thesis, in Chapter Five, the importance of assimilated cognitive norms and concepts such as social scaffolding and narratives, as briefly mentioned above, will be shown to

harmonise with the predictive processing framework as well, *even with* a view of cognition that has minimal content-ascription. Our social and ontogenetic development will be shown to be the basis for much of what is assumed by the propositional attitude view of our folk psychologising. It shall be proposed that enculturation and the narratives within social discourse shape how we come to understand and develop our folk discourse. It will be argued that reason explanation in the folk psychological paradigm need not assume the actual existence of content-laden propositional attitudes in the majority of cases of social cognition, and can instead be viewed as a cultural affectation that is only afforded to us as a gloss on our fundamental mental processes. It will be shown that reason explanation is more than likely retrospective instrumental[13] thinking that affords us a grip on an otherwise cognitively fuzzy environment, providing platitudes that guide us along pathways that work in the macro-perspective of general social interaction. I shall argue for a deeper, embodied view of our minds that does not work on much more than assimilated social norms incorporated into the predictive processing framework, delineating a sub-version of FP which I shall call Dispositional Folk Psychology (DFP).

The conclusions to be derived from the preceding bodies of work will describe a view of the mind as a frugal, or productively lazy (Clark 2015c: 9-12, 14-18), machine in a loose sense. Evolutionary affordances (Gibson 1979) over centuries of human development, so it shall be argued, have provided us with a brain/body system which seeks to maximise available resources toward the survival of the organism, by minimising free energy. The canonical definition of an affordance treats it as something within the environment of the organism which *matters* to it, for better or worse (ibid.: 127). The minimising of free energy lies at the base of the predictive processing account of the brain, as stated above, which means that an organism's affordances are constrained by the underlying principle of frugality. Accessing such affordances would require as direct a line to the body's environment as possible, something an a priori prediction machine utilising non-semantic covariant information could do far better than a constructivist machine relying on incoming sensory content. What appears to count here is the minimisation of free energy for maximum gain, with organismic survival being paramount. The evolutionary tricks to minimise this free energy have accumulated over millennia to arrive at complicated, complex processes[14], which have aggregated to provide a

[13] Instrumental thinking is that which incorporates appropriate cognitive means to a particular mental task. It is therefore crucial for intentionality, and by extension desire, traditionally conceived.

[14] As most in the biological fields stress, this in *no way* amounts to a guaranteed teleological outcome of perfection, and it *certainly* does not mean that humanity sits on top of the pile, superior in all the ways that count. We don't,

motley arrangement of functionalities. Not all of these are positive, as the body can only work with what has been given, which means that often the body's adaptions are jury-rigged over time for alternative purposes than originally intended. These shifts in functionalities are visible in the way that the brain has adapted over time, utilising less-than-optimal material to weave a tapestry that works *well enough*. There is a continually-expanding body of research which highlights how previous adaptations have been co-opted for new purposes, and occasionally into *multiple* functionalities. Exaptation is the technical term used to describe phenomena such as these. They can range from the innocuousness of a particular inert gene suddenly binding to a novel hormone, to the complexity of our neuronal wiring being utilised for functions other than those which it evolved for (Sapolsky 2017: 381).

It is from this work that a growing understanding of the brain as a flawed yet impressive machine has emerged. From here, it is not so far a leap to my view of the mind as a flawed prediction machine, co-opting previously evolved processes of environmental navigation toward this end. Everything that follows, be it language, norms, culture in general, are secondary to base survival, utilising environmental tracking mechanisms which have been *further* co-opted for these more nuanced features of our environment. Staying with the notion of survival, no longer is it dependent solely upon prudent reaction to physical environmental phenomena, but also upon prudent reaction to *social* environmental phenomena, in the form of cultural affordances (Ramstead *et al*. 2016). This is where my formulation of DFP plays a role as the assimilated social variant of the evolutionary game. Our social environment, as intangible as it may appear to be, has potentially facilitated a late-developing capacity in our brains to predict social affordances in just the way that other, more tangible, affordances have facilitated prediction capacities earlier in our developmental history. Again, it is plausible that these social predictive capacities have co-opted the already extant general predictive capacities of the brain toward this end, more than likely in a less-than-optimal way, and this has served us well enough in our evolutionary history. The end of Chapter Four will primarily be concerned with this view of our mind/brain.

By the end of this thesis, a more nuanced understanding of these brain processes, as well as their possible effects on social cognition, will have been developed. The counter-intuitive nature of much of what will be proposed might be taken to count against my argument, but our intuitions should not affect our acceptance of the science, and this includes current theories on

and we aren't. Having said that, advanced brain processing just so happens to be among the most complex phenomena known to man.

the ways in which predictive processing and teleosemiotics inform our understanding of our everyday cognition. The minimisation of free energy thesis within the predictive processing framework can complement a frugal teleosemiotic account of cognition, whereby minimal content exists in the brain for the majority of cognitive processes that utilise such content. This translates into a view of cognition that does not trade in content-rich propositional attitudes of many traditional folk psychological theories, which is the major point of interest and contribution of this thesis. While some traditional FP deals in content-rich propositional attitudes, I will conclude that it is an abstraction of our neural reality, regardless of its efficacy within social discourse. The implications of this are that it gives us a greater understanding of what we are actually doing when we ascribe beliefs and desires to others. Furthermore, it follows where the science is taking us, as opposed to hanging on to philosophical baggage. Regardless of the eventual validity of these claims, scientific work only progresses with a concerted effort toward adopting, absorbing, or eliminating theories. It is this drive for soundness in hypothesising that shall encapsulate much of what my project will entail.

Chapter Two: Folk Psychology as Propositional Attitude Ascription

"Jerry Fodor is my favorite philosopher.
I think that Jerry Fodor is wrong about nearly everything."
- Anthony Chemero (2009: ix)

1. Introduction

To begin a reworked assessment of what people actually pick out in the world when predicting other peoples' behaviour, we need to look at what has been said regarding folk psychology (FP). An attempt to ground our capacities to predict others' behaviour through empirical research and scientific methodologies was made, marking the pertinent distinction between our common-sense understanding utilised in our everyday explanations, and the *how* of the very same[15]. In other words, they began to separate the way in which people understood and explained each other's behaviour and the neurological processes which underlie these abilities. As shall become clearer later on in this thesis, not only are the methods by means of which we arrive at reason explanations[16] difficult to determine, but the very common-sense "understanding" of behaviour that people have is conceptually muddled at the very least. Compare the following: a highly contextually-bound conversation incorporating complex instrumental thinking[17] with someone's interaction within a crowd of people (be it with brief smiles or finer expressions of bodily intent), instantiating a minimal construal of instrumental thinking. What will become clearer in later chapters is that most of our cognition follows thinking, in its attenuated sense, of the second sort. This type of thinking does not involve representational, content-rich propositions of the sort often associated with instrumental cognition. This type of thinking will be shown to fold into a processing system within the brain, as discussed in Chapter Three, which generally enables successful navigation of the most pertinent environments for everyday human functioning (which includes the social

[15] The debate on *how* we folk psychologise is a diverse topic, with many adherents falling predominantly in either of the Theory-Theory, Simulation Theory, or the Hybrid Theory construals of FP (Stich & Nichols 2003: 238-251). The details of these need not concern us for this thesis, as these point more toward the methodologies we allegedly employ in our folk ascriptions, rather than the issues surrounding what these ascriptions entail.
[16] "Reason explanation" here refers to the elucidation of the behaviour of conspecifics within particular contexts.
[17] In this paper, instrumental thinking is construed as the explicit "rational" thinking that people are familiar with as they overtly and purposefully think about specific things.

environment). But for now, I shall simply preface what is to follow with the caveat that when thinking of our brains as generally interpreting information in this second sense, it is very difficult to accept the sketches of how the mind predicts behaviour as often put forward by cognitive theorists. Nevertheless, reviewing some prominent traditional views of Folk Psychology will inform our subsequent theorising on the topic and will aid us in reformulating a Theory of Mind (ToM) that draws on contemporary science. A ToM, here, refers to the ability in all of us to ascribe mental states (such as our beliefs and desires) to ourselves as well as others (Scholl & Leslie 1999). Furthermore, it is the understanding that others have differing mental states to our own.

There is a very influential view of ToMs within the fields of the cognitive sciences, philosophy of mind, and others. Whether explicit or not, it endorses a particular viewpoint on what occurs during interpersonal folk understanding[18] when we enact our ToMs. As Mölder (2016: 6-7) describes it, this is borne out by what he calls the Folk-Mentalist Thesis, an implicit view in these fields which describes the mind as exclusively comprised "of mental states as individuated through folk psychology" (ibid.: 7)[19]. Mölder (ibid.: 7) stresses that this is an *implicit* assumption that is widespread within the literature. The term "folk understanding" can also be substituted for "common-sense understanding", as these denote the same phenomenon (Ratcliffe 2007b: 56). They speak to a view of human interpersonal understanding as an everyday ability whereby people can derive an understanding of the behaviour of others through the attribution of beliefs and desires (amongst others) to their conspecifics. An example would be the attribution of a belief to someone of the effect "DSbelieves that p", but this will be expanded upon below. These, assuming a baseline cognitive ability in those involved, we tend to attribute to anyone. Churchland (1998a: 3) elucidates this quite clearly when he states that FP is

> …the prescientific, commonsense conceptual framework that all normally socialized human beings deploy in order to comprehend, predict, explain, and manipulate the behaviour of humans and the higher animals.

[18] It will become clear that the "interpersonal" is the baseline from which many studies of FP work, but concepts within FP are often utilised in a broader sense than simply interpersonal belief/desire attribution, which will be discussed.

[19] This Folk-Mentalist Thesis is derived from two implicit theses, the one being what Mölder calls The Folk Assumption, that "mental states are individuated through folk psychology" (2016: 6), the other being The Mentalist Assumption, that "the mind is solely composed of mental states" (ibid: 7).

This ability appears to manifest itself, whether through genetically-endowed innate processes or in ontogenetic development[20], independently of what we may eventually learn from scientific studies concerning the mind and its functioning. In other words, people *acquire or learn* the folk psychological framework, despite not *actually understanding* the base neurological processes that it is assumed to describe. Despite folk psychological *practice* centring on this acquired understanding of the interpersonal domain, prevalent folk psychological *theories* are aimed at describing the mental states on which the practice is predicated.

The point here is that there is an accepted, highly influential view of folk psychology which trades in belief-desire attributions (to be further explained below), and which rests on the assumption that these have something interesting to say about what actually occurs at a mental/neural level, and is referred to when we folk psychologise. It is this that will predominantly concern us for the rest of this chapter.

I shall begin this chapter with an historical overview of such theories of folk psychology, namely, those that construe it as entailing propositional attitudes. I will describe what propositional attitudes are, and how they are thought to function and work within our folk psychologising. I will then look at a couple of the most influential theorists who have argued for and against propositional attitudinal thinking in their writing. This will serve to highlight some of the central assumptions of such folk psychological theories, which I will go on to critique. My main criticism will be aimed at conceiving of folk psychology primarily as accurately picking out propositional attitudes that somehow exist at the neuronal level. A prominent example here is Fodor who, despite claiming that he does not hold this view, seemingly ends up arguing for it nevertheless (2003: 155). This will lay the groundwork for a more scientifically-valid theory of what occurs in our common-sense understanding of others, which is a primary concern of this thesis. Hutto's notion of the Narrative Practice Hypothesis (Gallagher & Hutto 2008; Hutto 2004; 2007a; 2007b; 2008a; 2008b), and the work that narratives perform in our cognitive development and subsequent reasoning, will be shown to be a cogent explanation of the temporally-extended effects that culture has on the development of our FP abilities. This further undermines the validity of describing the propositional attitude ascriptions in our folk psychologising as accurately picking something out in the world. The preceding will open up space for the following chapter on the predictive abilities of the brain,

[20] Ontogeny is the development of any organism from its origin to its mature form, whatever that may consist in for the particular organism.

which needs to be understood before introducing teleosemiotics and an alternative interpretation of how the mind tracks environmental cues later on**.** For now, folk psychological thinking as propositional attitude ascription that accurately picks out something in the world will be called into question.

2. Propositional Attitude Psychology

2.1 Propositional Attitudes and Folk Psychology

When people folk psychologise, they attribute mental states to others primarily on the basis of the observation of behavioural patterns. These mental states are taken to be the intentional states[21] of the others under observation. The concept of an intentional state has a long history in the philosophical field[22], but can simply be described as the "being aboutness" of a mental state (Hutto 2008a: 1). In other words, it is what that particular mental state is "directed toward", or what it is "about". For example, when a mental state involves a belief that there is an empty coffee cup on a table, we can assign the simple fact that that mental state is *directed toward* the existence of that empty coffee cup. Therefore, the mental *intentional state* is somehow directed toward the empty coffee cup. Traditionally in the folk psychological literature, talk about intentional states predominantly focuses on two types, those of "belief" and "desire", but there are many others (such as fear, hope and so on). Fodor (1987: x, and Chapter One) is indicative of this tendency when he describes our common-sense psychology in terms of the application of a belief/desire dichotomous pairing[23]. To understand behaviour, according to Fodor, one *must* ascribe a belief and/or desire to the agent in question in order to make sense of the agent's behaviour. He goes so far as to say that "there *are* no alternative

[21] In the broad field of philosophy, *intention* should not be confused with intension (with an s), which refers to the non-extensionality of specific words or phrases. To use an example, the *extension* (what the concepts denote in the world) of the phrases "organism with a liver" and "organism with a pair of lungs" are the same in that, in reality, organisms that have lungs also have kidneys. In contrast, however, the *intension* (the semantic content) of the concepts "liver" and "pair of lungs" are different. "Liver" and "lungs" have different meanings. Moreover, the shared extension here is contingent; these phrases could have had differing extensions, had animals evolved differently. This is not terribly important in this treatment, but some philosophers of language would perhaps argue otherwise.

[22] It is a term that was revived by Franz Brentano in the nineteenth century from medieval philosophy, and taken up again by late twentieth century philosophers of mind (Dennett 1987: 19-20).

[23] A mental state consisting of a belief is understood to guide behaviour in some form, while a mental state consisting of a desire is a motivator for behaviour through the integration of end-directed aims. This pair of mental states, belief and desire, are what are usually used within the folk psychological literature. After my description of predictive processing is introduced in Chapter Three, it shall be shown how the clear distinction between these two is problematic.

theories [of behavioural explanation] available" (Fodor 1987: x). Searle (1983: 1) states this very same idea even earlier in his project on the intentionality of the mind. To Searle, our intentional states can *only* be those that contain beliefs and desires with inherent content. They must be *about something*, and that *something* informs what amounts to a propositionally-laden attitude. While intentional states are more than just these paradigmatic examples[24], talk of FP trades almost exclusively in them. These paradigmatic mental states are expressed in the form of "propositional attitudes", where a person "D" is being mind-read by a conspecific, and this mindreading takes the form (within the conspecific) "D desires c", "D believes that p", and so forth, where "c" and "p" denote propositions such as "I see a cup of strong coffee" or "strong coffee will fix this" (Hutto 2008: 1-2). For the concepts of "belief" and "desire" to be efficacious, their meaning must be analysable in terms of sentences specifying the mental states under which people would behave in ways appropriate to the desires or beliefs in question (Stich & Nichols 2003: 2).

The ways in which these concepts are used within our folk psychologising fix their meaning within the broad theoretical framework of our FP (in other words: their causal role is important). On this view, FP (as a theoretical framework analogous to any science) explains our mental states by attributing beliefs and desires that are taken to accurately represent the mental states in question (ibid.: 4). What is certainly indisputable is the centrality of the *use* of propositional attitude ascription within FP *discourse* and its utilisation by people for behavioural understanding in interpersonal communication (Hutto 2008a: 2). What is also taken for granted is that those utilising FP genuinely interpret such propositional attitudes as causally determinant internal mental states that those under observation possess (ibid.: 5-6). These mental states are seen as causally determinant in that it is these propositional states that are taken to be the true instigators of behaviour. Dispositional states have seemingly been ignored in much writing on the topic for this very reason (Ratcliffe 2007b: 205-211). Dispositional states refer to a form of mental state which does not deal specifically with propositional content that is actually entertained, but that nevertheless has the potential to give rise to a particular belief or desire. In other words, they somehow have the capacity to become full-blown, intentionally-directed states if the requisite environmental prompting is present (ibid.: 204). It is not easy to attribute a propositional attitude to that which is not very clearly attributable on the dispositional construal, and hence the problem has been largely ignored. Our language, in the form of describing propositional attitudes, seemingly cannot capture this state

[24] For a good discussion on this, see Grzankowski (2012).

of affairs except through fairly broad holistic descriptions of the mental states under question (most prominently in the form of the beliefs and desires that we all ascribe to people). In other words, our descriptions of these particular mental states cannot be delineated further than the usual propositional attitudinal ascriptions, despite these not necessarily accurately describing these mental states. As I will argue, FP is accurate only insofar as it is utilised for coarse-grained behavioural understanding. Despite this, there is a sense of realism in talk attributing content to mental states for the purposes of prediction. Folk practicing FP generally manage to separate states involving beliefs and/or desires when appropriate and, while true in a holistically pragmatic sense, these are taken as accurately depicting the internal mental states of conspecifics (ibid.: 210-211).

Therefore, while theorists of FP believe that the processes involved in attributing propositional attitudes are mostly tacit, they also assume that explicitly elucidating our propositional attitudes can be indicative of what occurs beneath the surface (ibid.: 8). Proponents argue that it is possible to gain insight into the internal mental states of individuals through more rigorous philosophical analyses of what is expressed. Propositional attitude ascriptions (the utterances referred to above), such as "S desires p", are deemed to be a window into this murky world, and they are termed "platitudes", and anyone with a suitably operating cognitive ability would not deny their claims[25]. What is important to take from this though, is that these platitudes, vast in number, all appear to interrelate in suitably structured ways between the mental states possessed by individuals and their resultant actions (Malle 2004, see Chapter Four). And due to the seemingly obvious relationships that the platitudes seem to indicate between mental states and their impact on the world, it encourages the view that the study of our FP offers a conceptually unified (up to a point) understanding of our inner cognitive processes. This understanding is made manifest, it is thought, through these aforementioned platitudes that operate in a law-like fashion (Ratcliffe 2007b: 6) and is achieved with minimal effort on our parts. We all know these generalised platitudes as objects of common knowledge[26].

In light of the above, FP consists of our ability to predict and attribute these mental states to others, all in the service of explaining their behaviour. This is known as the mindreading aspect of our FP (first alluded to, but not explicitly named, in Stich & Ravenscroft 1994). Mindreading refers to the ability of individuals to attribute propositional mental states to conspecifics, which

[25] For extensive and seminal treatments of this conception of our FP, see Lewis' (1966, 1970, 1972, 1994).
[26] These platitudes, as explained in Chapter One of this thesis, come in the form: "S believes that p".

are then used to predict and explain their behaviour[27]. Generally, these mental states are thought to be *information-rich* (i.e content-rich in the form of propositional attitudes) and to be represented in some fashion within the minds of others, which explains their behaviour. Hence, propositional attitudes, as bearers of content-rich information, function as the reason for particular behaviour, which can lead to prediction by others through the mindreading aspect of FP (Stich & Nichols 2003: 6)[28]. Crucially, the platitudes (and by extension the propositional attitudes) mentioned above are seen as identical to the content-rich mental states that enable mindreading (ibid.).

2.2. Fodorian Folk Psychology

2.2.1. The Language of Thought and Fodorian Folk Psychology

A prominent thinker who endorses a view which is quintessentially of the nature set out above is Fodor (1975, 1978, 1984, 1987, 1990, 1994, 2008b). I make particular mention of him as his views on theory of mind, and in particular his Language of Thought (LoT) hypothesis (Fodor 1975, 2008b), presuppose a scientific as well as a propositional view of the mind. His thought on the topic trades in talk of generalisations of a law-like nature, modules[29], and the predictive power that these all afford (as with the rest of this thesis, the importance of the power of prediction always bubbles to the surface). The importance of this, for our purposes, rests on the scientific view that Fodor holds, but which nevertheless stops short of incorporating functional indeterminacy within his theory of mental processes. Here, functional indeterminacy refers to the indeterminate nature of mental processing mechanisms or, more specifically, how these even came to be within environments with a plethora of distal causes. Indeterminacy of mental processes denotes the inability to always accurately fix the content of mental states to the environment. Arguably, it is difficult to entertain the view that mental states accurately represent particular environmental stimuli if there are potentially multiple causes of these representations within the environment. If there are multiple causes which set the accuracy conditions of a mental state, we run into problems when needing to elucidate their definitive causal histories. This does not bode well for rich propositional content that trades in robust accuracy conditions, and this shall be a prominent point within the rest of this thesis.

[27] A clear example would be the mental states attributed to a pedestrian at a road crossing with a lot of traffic passing by. We attribute the mental state that she believes that she should stop and, disaster notwithstanding, we are often correct in our assumptions.
[28] It is for this reason that Stich & Nichols (2003: 6) point toward the "information-rich" aspect of FP.
[29] Talk of delineated modules and law-like generalisations can easily be transposed to the world of folk ascriptions.

Thinking, on Fodor's view, operates in a form of mentalistic language (his LoT), instantiated in a complex but systematic web of representations, which "is at the heart of [the] systematicity of thought" (Fodor 2008b: 20). The question of neuronal correlations aside, these representations are combined in a syntax[30], along with attendant semantic[31] relations. These allow for the processing of these representations in a causal manner. In other words, representations are capable of *causing* other representations within the syntactic chain (ibid.: 199). Much emphasis is put on the semantic relations that are afforded by this combinatorial system, and thinking is effectively deemed to be a "tokening" within the brain of the requisite structured representations (which include both syntactic and semantic characteristics) (ibid.: 5-6). On top of this, Fodor argues that individuals are realist when it comes to propositional folk ascriptions, and uses his fully representational theory of mind to defend the validity of this view (Ratcliffe 2007b: 6). It is worth noting that the incredible predictive power of the brain is made vastly more understandable if this view of how thinking occurs within the mind is at least partially correct. The type of mental states that inform our FP abilities would naturally follow from this fully representational depiction of mental states. This is one of the reasons why it has been so popular since its inception.

Fodor has always advocated a computational, strictly cognitivist, model of the mind (one can scan any of his works to arrive at the conclusion that he is a strict functionalist). Cognitivism (briefly alluded to in Chapter One) has been the dominant view detailing the mind and its relation to the brain within the literature since midway through the 20th century. While the details do not matter as much to this thesis, these days it has detractors (e.g. Clark 2013a, 2016; Hutto & Myin 2017) that take issue with its overly simplistic notion of passive sensory uptake and representation of this sensory information within the mind (a view to be problematised in the following chapters). Drawing from this, Fodor's thesis maintains that *every* mental state is identical to what could be called a monadic property within the mind (Fodor 1994: 12-16). This property can be characterised as the connective *relation* to an explicit *mental representation*. As these mental states are seen to be type-identical to computational processes, the mental states can be interpreted as *symbolic*, with their representations involving the compositional qualities/properties of syntax and semantics (Fodor & Pylyshyn 1988: 12). Within his LoT, we can begin to see the value this would have with regards to how the syntactic and semantic structure of the mind could form a system of innate "Mentalese" (Fodor 1994: 167) which

[30] Syntax refers to the structural relations which exist between words and sentences, which provide the principles upon which sentences are regulated within a language.
[31] Semantic refers to the inherent meaning of words.

supervenes[32] in some form upon the structural underpinnings of the neuronal processes (Fodor 2008b: 15). Therefore, the propositional attitudes, in turn, are instantiated within this complex web of Mentalese, providing just the causal and explanatory power needed to make thinking *intelligible* to us, while also bearing a relation to the monadic state mentioned above (Fodor 1994: 13). In other words, the propositional attitudes bear special relations to explicit mental representations, which have content *identical* to the propositional attitudes themselves. This provides the systematicity of thought that all (non-compromised) human beings seem to have (Fodor & Pylyshyn 1988: 40). The type-identity conflation of propositional attitudes and mental representations as somehow supervenient upon the neurological substrate of the brain will be critiqued (or problematised) throughout this thesis.

2.2.2 The Disjunction Problem and Adaptational Denialism

The problem, on my construal, is that while Fodor captures a sense of continuity between the properties of the environment (the "content") and mind with his LoT, he makes use of a view of mental content which is too "narrow" to account for its causal and explanatory powers. Narrow mental content is defined as mental content which is not determined by the individual's environment, but purely by the intrinsic properties of the individual. In other words, if there were to be a duplicate of said individual in a significantly different environment, this duplicate would *necessarily* still have the same content within her cognitive apparatus. This is in contradistinction to "broad" mental content, which incorporates the environment of the individual (including elements of the individual herself) to determine mental content. Accordingly, a duplicate of this individual would have differing mental content to the original. This is partly why Fodor struggles to conceptualise the derivation of content from the environment.

Fodor's ideas on propositional thinking are highly dependent on the "intact" cognitive perceiver, whose perception and cognition are functioning properly (1987: 126-127). Errors of misrepresentation are allegedly ruled out when the cognitive apparatus functions in a "proper" fashion (ibid.: 101-102). The emphasis, on my construal, on the "intact" cognitive perceiver seems to dismiss the accepted notion of ourselves as fundamentally flawed in our perceptual

[32] Supervenience is a term with a lot of weight within specific philosophical spheres. For the simplest explanation, if something supervenes on something else, a change in the latter necessitates a change in the former. Therefore, if Fodorian Mentalese supervenes upon the neuronal substrate, a change in the neuronal substrate would necessarily lead to a change in the Mentalese instantiated by these structures.

and reasoning abilities at even the conscious level[33]. While he has dropped an overtly narrow view of content since his (1994), as noted within his recent work, Fodor's views on evolutionary processes provide insight into his misgivings concerning how content is fixed within mental states.

Fodor (Fodor 2008a; Fodor & Piatelli-Palmarini 2010b) endorses a distinctive view of evolutionary adaptationism which seeks to undermine standard, in his terminology, "Darwinism"[34] (ibid.). On his view, current theories on natural selection have fundamental flaws which are in need of rectification. Fodor argues that these flaws apply to all phenotypic traits (Fodor & Piatelli-Palmarini 2010b: 154), but its import for content determination within our mental states are what fundamentally concern us here. His argument follows the contour of his renowned disjunction problem[35], namely, that it is difficult to distinguish between causes which are content-determinant and those that are not (Fodor 1994: 40). These causes have a direct bearing on the content (accuracy conditions) of mental representations. On Fodor's construal, a theory on representational content *must* make it clear how it is that a representation is capable of *mis*representing. In other words, what would make a particular representation a mistake, in terms of representing something in the world as having properties it does not have or being something other than it is. The disjunction problem underscores the difficulties in the acquisition of specific mental states through adaptation, as there could not be an easy mind-to-world fit (as explained below).

If one stimulus, let us call it S1, accounts for a particular mental representation M, delineating an adaptational attunement of the mental to that stimulus, we can say that there is no disjunction evident. But, if there is a second stimulus S2 that *also* triggers mental representation M, how are we to separate the adaptational attunement between mental representation M and either stimulus S1 or S2? By way of (an admittedly timeworn) example, if both chicken (S1) and duck-in-the-dark (S2) cause the representation M, this representation is instead constituted by the disjunctive chicken-or-duck-in-the-dark (S1 v S2). The issue, here, is that if both disjuncts reliably trigger the "chicken" representation, what is it about the "duck-in-the-dark" triggering that makes it a misrepresentation? This is the argument underlying Fodor's issue with adaptational and holistic depictions of content acquisition (ibid.). But in fact, the entirety of

[33] Kahneman (2011), Greene (2015), and Sapolsky (2017) are excellent contemporary sources detailing this very topic.
[34] The term itself should ring alarm bells as there are a multitude of evolutionary theories on offer which would fall under the bracket of Neo-Darwinism. The name alone appears to be introduced to construct a straw man argument for easy dispatchment, despite the ever-adaptable nature of evolutionary theory.
[35] This idea was first formulated in his (1984), and thereafter developed in subsequent works (1987, 1994).

Fodor's thesis on the topic is a masterstroke in wonderfully-argued obfuscation. According to him, what M actually represents is not either a case of S1 or S2, but rather the actual disjunction (S1 v S2) (ibid.). To him, there is a gap in our understanding if we are forced to attribute the result of adaptational processes to a disjunction as opposed to either of the disjuncts. If both the disjunction and either of the disjuncts are on the table, invoking correctness conditions does not get us out of the problem as there is no principled way, according to Fodor, to distinguish between the two using an adaptational theory of content. An adaptational theory of content can account for a representation, but cannot account for one (eg. chicken) being the correct representation and the other (eg. duck-in-the-dark) being a *mis*representation. No matter how fine-grained our analysis, we are still left with an indeterminacy at the core of our theory of content—we cannot account for mistaken representations. As a result, we need to accept that there is a fundamental problem at the heart of naturalising our content by means of adaptationist theory if we cannot undo this kind of result (ibid.: 48). Therefore, the functional indeterminacy engendered by the disjunctive relation is untenable on his view, which is the reason for his dropping of the "Darwinist" program, as selection-for (using the evolutionary term) is overdetermined[36].

To explain away the disjunction problem, Fodor has introduced his theory of "asymmetric dependency" (Fodor 1987, 1994). According to this theory, while duck-in-the-dark may cause a "chicken" tokening within the brain while not seeing an actual chicken, duck-in-the-dark only instantiates "chicken" tokens because chickens *actually* do. To be more precise, if chickens don't instantiate "chicken" tokens, then duck-in-the-dark (or any other objects which also cause a misrepresentation) would not do so either. But, even if duck-in-the-dark did not cause "chicken" tokens, chickens would still instantiate "chicken" tokens. If one is able to sever the connection between the chicken and "chicken" tokening, one would sever the duck-in-the-dark and "chicken" tokening, but not the other way around. This is why the duck-in-the-dark to "chicken" tokening is *asymmetrically* dependent upon the chicken to "chicken" tokening relation. It is not clear to me how this truly gets to the heart of the issue. This seems more semantic handwaving than a fundamental game-changer in explicating the accuracy conditions of our mental states. It stipulates in purely conceptual terms the way in which we can reliably misrepresent an object, but it does not introduce a robust explanation of the causal history of the misrepresentation which the disjunction problem embodies. An alternate path we can take

[36] In other words, there are potentially more than one possible cause of a particular representation, and we cannot distinguish between an accurate and an inaccurate representation of a particular phenomenon in the world.

is to explain the disjunction problem away with an appeal to the way in which our minds have evolved to pick out information in the environment.

As Rosenberg (2013: 5-8) correctly points out, it is selection-*against* which evolutionary pressures actually dictate, without purpose or end-directed aims. In other words, while evolutionary processes are blind to the end result or consequences of their pressures, Darwin's main theoretical accomplishment was that of

> …identifying the mechanism—random, i.e. blind, unforesighted variation and passive, environmental filtration that sculpts the appearance of purpose in nature, even though there is no reality of purpose operating in it (ibid.: 6).

If evolutionary theory is understood as passive and as a "mechanism of filtration" (ibid.), we can accept the view, *contra* what Fodor has written on the topic, which is that seemingly disjunctive properties can filter through despite no clear reason for the selection *for* these traits. The passive mechanisms of filtration in nature are only in the business of preventing traits from persisting which are not conducive to the continued survival of the organism. If we take natural selection to be an imprecise and gradual change in an organism, the selection for particular mental states at all times seems to be an illogical extension of evolutionary theory. Evolutionary processes need not select for a trait at all times, it need simply select *against* that which is not pertinent to the survival of the organism, and this also means that a lot of residual evolutionary baggage would not simply disappear. Extending this to dispositional mental states, the organism does not need to evolve mechanisms that instantiate determinate mental states consistently as there is seemingly no sufficient evolutionary pressure to evolve many of these mechanisms in the first place. This does not rule out determinate mental states, but it does allow for less determinate mental states in our cognitive processing, contrary to what Fodor would want with his representational theory of mind and LoT. Beliefs and desires can exist on this view, but not in the rigid fashion that Fodor would insist upon. Fodor's theories simply do not align with evolutionary theory, and worries about disjunctive properties are misplaced. For example, if we take a probabilistic and minimally contentful view of mental content[37] seriously, we need not worry about the accuracy conditions of each mental state if dealing with uncertainty is inherently built into our cognitive architecture. Our mental states can misrepresent as long as it does not impair our more pertinent interactions with the environment.

[37] To be properly expanded upon in Chapters Three and Four.

To make this clearer, Fodor argues that the concepts of propositions cannot be broken down further than the constituents of language, meaning that these concepts cannot be individuated into even simpler constituent concepts (Fodor & Pylyshyn 1988: 48-50). Mental representations, to him, are the individuated, linguistic components of propositional thinking, which affords the seeming systematicity of our thought and general cognition (ibid.: 48). For him, the unsystematicity that would proliferate in our cognition if we were to take any other theory would be "*preposterous*" (ibid.: 49, emphasis in original). If we were to ignore the semantic and syntactic relations afforded by a LoT, all we are left with is a "list" of representations[38] (ibid.). While holding the view that the conceptual meanings of propositional attitudes are atomistic (or, cannot be broken down into further constituent parts) (Fodor & Lepore 1992: 206), he has expanded this view into the aforementioned generalised viewpoint of how adaptational accounts do not suffice to allow for the selection of content through the adaptational history of the organism (Fodor & Piatelli-Palmarini 2010b). To grasp Fodor's views a little more, it is worth understanding his use of the constrast between "selection-of" and "selection-for" (ibid.: 110; Sober 2010: 595). His example is that of a population of theoretical entities in which some have hearts and a significant proportion do not. The hearts not only pump blood, but also make sounds. Needless to say, those without hearts do not have either of these traits. Both of these traits are called "coextensive" as they are contained within the same organisms within the given population. If, after time, we find that all of the organisms now have hearts, the question arises as to what caused this change in circumstance. Assuming natural selection occurred, which Fodor certainly does despite his misgivings, we can assume that organisms which had hearts to pump their blood had increased capacity for survival and reproduction (i.e. fitness) in relation to those that did not. As the aforementioned traits were coextensive, it can be stated without controversy that those organisms with the heart noises were fitter than those without.

This is all well and good, until Fodor (and Piatelli-Palmarini) attempt to problematise the distinction between selection-of and selection-for. The problem goes as follows: natural selection occurred, selecting *for* the pumping of blood, but there was no selection *for* the creation of the heart noises. Again, this is not problematic. But, while there was selection *of*

[38] This paper was specifically in response to connectionist theories of mental cognition, where they instead advocated for a symbol-based LoT. Connectionism refers to the application of approaches in artificial intelligence to the understanding of the brain, whereby multiple connections between brain cells enable simultaneous processes to function. Furthermore, many processes are nested in parallel hierarchies that eventually form mental states and assorted phenomena.

both of these traits, there is only selection *for* one of the traits (the pumping of blood). This is where Fodor (and Piatelli-Palmarini) take issue: they state that any claims that natural selection selected *for* one trait but not the other of two coextensive traits cannot be true. To the authors, all that natural selection can truly achieve is selection-*of* two coextensive traits. If one is selected, then so is the other, and there is nothing else that can be said about it (ibid.: 154). This appears to necessitate a rejection of how the sciences manage to separate causal relations (the selection for the pumping of blood increased fitness) and mere correlations (the "thump-thump" noises that merely came along for the ride) (Block & Kitcher 2010; Papineau 2010; Sober 2010: 596). The authors do not see it this way, and maintain that there is something about selection-for that is problematic (Sober 2010: 596). They would maintain that there is a "fact of the matter" concerning which trait caused increased fitness, but they deny that natural selection as currently understood could possibly distinguish between the actual cause of the increased fitness and the accidental, co-extensive trait (Fodor & Piatelli-Palmarini 2010a). Therefore, they assert that it is not simply that we are barred from truly knowing that one trait (the pumping of blood) was selected for, but that any such claims are simply false. There simply is no such fact of the matter. Taking it further, there are simply no laws about selection-for due to the high context-dependence of each case of adaptation. The complex nature of the environment makes it impossible for natural selection to provide sufficient laws which govern selection-for, as all instances of selection-for would need to have laws which dictate what they all have in common.

The above is a form of adaptational denialism. The general framework provided by Fodor is too simplistic a case to make against adaptational processes, and is ultimately unhelpful. I argue that Fodor's approach toward natural selection is a hangover from his view of the indispensability of the atomistic constituents of the propositional attitudes. Arguably, his argument about the disjunction problem is the proto-argument for his subsequent discussion on the inability of selective processes to fix adaptational traits. The above selection-for issues map directly onto Fodor's worries about the disjunction problem in that they are clear cases of undecidability that need to be explained, at least on Fodor's construal. Instead of accepting that misrepresentation need not be a fundamental issue when discussing the disjunction problem, or that selection for either of two traits need not be an issue when discussing adaptions, Fodor bites the bullet and advocates a clear grounding for each. In light of the above, I argue that Fodorian thinking on propositional attitudes is a deeply problematic mindset if we want to think

adequately about the constituents of our mental lives[39], especially when it informs our evolutionary history. Fodor's arguments fly in the face of the work done in naturalising philosophical theories within the cognitive sciences, where intentional attitudes and the content thereof can plausibly be shown to be either suitably grounded in plausible scientific theories, or explained away. Without further incorporation of scientific analyses into his theorising, it is also difficult to falsify neuroscientifically. For instance, as discussed in Chapter Four, it will be shown how we can allow for sensitivity to the environment without needing to adhere to strict content-determination. Arguably, suitably naturalised theories will be grounded in an evolutionary understanding of our brains, which can account for the adaptation of our mental states. Fodor would not support the proposals in this thesis due to his seeming adaptational denialism, but to naturalise propositional attitudes, and to solve (or more accurately explain away) the disjunction problem, the issue of the ostensible indeterminacy of content needs to be entertained. I propose that the disjunction problem falls away if we incorporate the teleosemiotic account of intentionality expanded upon in Chapter Four of this thesis. Furthermore, the atomistic view of the propositional attitudes, as supervenient upon our cognitive architecture, will be undermined through a theory of the indeterministic nature of our fine-grained mental processes, as expanded upon in both Chapters Three and Four of this thesis.

2.3. Churchland

2.3.1. Churchland's eliminativism

Another prominent thinker on the attribution of propositional attitudes within our folk psychologising is Churchland, who set out to dismantle FP as we know it (primarily in (1981), but also see (1985), (1998a), (1998b), (1998c)). I choose to discuss his views in order to show the opposing poles of the FP debate, and to therefore show the scope thereof. To Churchland (1981: 67), if FP is to be taken as an empirical theory that reliably describes the regular correlation between stimuli and behaviour, it is not obvious that propositional attitudes are realistic depictions of individuals' mental states. His central concern is the possibility of the scientific *falsehood* of much of our theoretical folk psychologising. To him, there remains the looming possibility that the disparate phenomena within our FP could fail to reduce to our actual mental states in a future science, which would necessitate dropping these concepts

[39] Or, as Dennett (2013b) would say, we need to pump our intuitions in another direction.

altogether[40]. The broad terminology we use to describe the mental states of others could then be subsumed and replaced by this mature science's construal of the inner workings of the mind. This, according to Churchland (ibid.: 70), would not bode well for the efficacy of FP as a whole, but *especially* for thinking about our beliefs and desires as taking the form of the much-vaunted propositional attitudes. Therefore, invoking the propositional attitude "S believes that p" could fail to describe any concrete mental state at all, beyond a false, pre-determined holistic ascription of an individual's behaviour[41]. What is meant by this is that an adequate fine-grained depiction of what is being picked out by our folk psychologising will not necessarily equate to the depiction that we use in our everyday behavioural ascriptions. Therefore, the "theory" used to explain behaviour (the propositional attitudes) will be false, strictly speaking, and only depicts the holistic dispositions of the agent (ibid.). FP as a theory must therefore be eligible as a "candidate for elimination" (Churchland 1981: 212) as more accurate theories are developed. This view has led Churchland to proclaim that there are *no such things as beliefs or desires*[42] (ibid.: 89). This provocative claim has often been used as a club with which to beat the back of Churchland's reasoning[43], as taking it at face value has led to emotive retorts ridiculing the very idea. But this response has generally been unfair, as the actual focus of his sustained attack has been the view that our beliefs and desires are *propositional attitudes in the conventional sense*. On his view, beliefs and desires (and whichever other relevant mental states under investigation) cannot exist *in this form*. He does not rule out the possibility that what we take to be our beliefs and desires can be reinstated in some future science, but in a different theoretical terminology that accurately depicts what is happening at a physical level in the brain (ibid.: 86). This approach could be taken to mean the death of beliefs and desires (in the traditional sense), and therefore the eliminativist's creed appears to be doing away with something integral to our self-conception (ibid. 1998b). However, to give a fair assessment of his view, a more comprehensive critique should be attempted.

[40] I am not personally enamoured with this all-or-nothing reductionism. I do not view it as a useful methodology within our epistemological concerns, including the eventual basis of what occurs in the mind with regards to our folk psychologising. Simply explicating our underlying architecture does not *necessitate* a sea-change in the way in which we operate at a social level. I advocate tempering our argumentation and letting the chips fall where they may.

[41] By "pre-determined" I refer to the socio-culturally acquired meaning of what a propositional attitude entails, and by "holistic" I refer to the notion that these meanings are adequate for our social interaction, but do not accurately pick out our mental states.

[42] This point is argued even more forcefully by Stich's early writings (Stich 1983: 231).

[43] Any response from Fodor drips with sarcasm, see for example Fodor & Lepore's Chapter Seven in (1992).

2.3.2. Churchland's eliminativist thesis

On Churchland's (1998a: 11-12) construal, it is not evident how the vast set of general "laws" involved in our FP can be applied in a rolling and continual fashion through the manipulation of sentential propositions in the brain. The task appears to be unnecessarily complex for such nuanced interaction and subsequent comprehension of others' behaviour[44]. Much like the notion of minimising free energy, which shall be expanded upon in the next chapter, these complex interactions seem unnecessary to provide the benefits of our folk psychologising. In his words, it is wholly "mysterious" how these "prodigious feats of retrieval" could function optimally in our "ongoing social commerce"[45] (Churchland 1998a: 12). Churchland, in a move that ties in with the spirit of this thesis, asks us to look at the emerging account of the way in which brains are understood to "encode" information (ibid. 1995: 21; 1998a: 10; 1998b: 31; 1998c). In contradistinction to Fodor's LoT, he does not endorse a view of the brain embodying disparate information in sentence-like structures. Instead, he attempts to bring this view into the modern scientific age with his endorsement of vector coding[46] within the brain, which is potentially a more plausible theory of how the brain "represents"[47] information (ibid. 1993; 1995: Chapters Two and Three; 1998c).

One of his favoured analogies to illustrate the power of vector coding is that of the television screen in relation to our own retinas (ibid. 1995: 7-8; 1998a: 13-14). Television screens, in his words, depict "a sequence of representations which are non-sentential" in both their "syntax as well as their semantics" (ibid.: 1998a: 13). The "semantics" within this analogy refers to the intensity of the activation patterns of the screen pixels, designating the content, while the "syntax" is rules governing the arrangement of the many pixels in just such a pattern. There is no logical structure beyond the activation of the many brightness intensities along the screen pixels. In much the same way, our photo-receptors have similar activation patterns. In this example, the "semantics" can be understood as pictorial, but different sensory modalities yield differing semantic relations. For example, taste codes activation patterns across the tongue, just

[44] Churchland (1998a: 12) also pointed out that individuals can seldom express how it is that they attribute their folk abilities. Ratcliffe (2007b: 46-52) performed an informal experiment to detect just what the folk think understanding (on the interpersonal level) means, leading to a similar conclusion.
[45] Of course, at face value, simply because it doesn't seem possible should never discourage a research paradigm. It is conceivable, for example, that various underlying processes take a lot of the burden enabling a sentential propositional language in the brain at some higher level.
[46] Vector coding refers to the mapping of activation patterns across the brain.
[47] The reason for the scare quotes shall be explained once I go into the issues surrounding representation further in this thesis, but I shall refrain from using them excessively from now on as it is Churchland's terminology.

as sound codes activation patterns across the auditory canal. The semantics of the above are clearly very different, but the underlying processes encode a vastly powerful tool of representation, regardless of the form in which this representation takes. The power of this picture of representation as vector coding is made evident once one realises that one singular vector representation is only one arrangement of the myriad values available, and the potential values available (or things able to be represented) increase on a logarithmic scale as more values are incorporated (ibid. 1995: 7; 1998a: 13). Compare the hundreds of thousands of pixels that a television screen can display, and the millions that our photo-receptors embody. Not only do these millions of values activate a vector across the target population of retinal neurons, but these can then be transformed across multiple neuronal populations, ending in the effector motor neurons responsible for action in response to the original sensory input (ibid. 1998a: 14). Here the television analogy breaks down, as the sheer number of potential neurons, ranging from the millions to the billions depending on where you draw the functional line, are smaller and incorporated into more than simply the two dimensions of a television screen.

On Churchland's view, the connections along the uncountable synaptic pathways embody the stored knowledge and learned ability of the individual, enabling a form of learning that does not trade in sentential structures, but multiply combinatorial vector "prototypes" (ibid. 1995: 27-34; 1998a: 14). These prototypes comprise the "learned perceptual and explanatory" capacities of the individual. Furthermore, these prototypes inform our perception over and above what is taken in through our sensory apparatus at any time[48] (ibid.). It is these prototypes that enable, indeed *embody*, our explanatory capabilities, and inform our perception. The massively parallel processing[49] capabilities of the brain, utilising the background perception that these activation vectors afford, cause behavioural responses to occur within milliseconds of response time. Thus,

> [t]o learn a theoretical framework is to configure one's synaptic connections in such a fashion as to partition the space of possible neuronal patterns into a system or hierarchy of prototypes. And to achieve explanatory understanding of an event is to have activated an appropriate prototype vector from the waiting hierarchy (ibid. 1998a: 15).

[48] Interestingly, by articulating this, Churchland had hit on something akin to what shall be introduced in the subsequent chapter in terms of the percepts that predictive models are thought to represent in the brain. He even states that these prototypes "constitute ampliative interpretations of that input, interpretations that place the input into an antecedently prepared context and fund expectations of features so far unperceived" (Churchland 1998a: 14). In other words: these prototypes can be interpreted as predictive models upon which we subsequently act upon the world.

[49] Parallel processing is a technological term which refers to simultaneous processing within a singular computer. This has been extended to the operations of the brain.

In this way, our perception, prediction, understanding, and action, are all interlinked in a mutual operation centred on vector processing[50]. But what does this lend to Churchland's view of the untenable nature of our current understanding of FP? He states that it is this very vector processing paradigm which informs FP, or more precisely the "family of learned vectorial prototypes" (ibid.) is simply what FP is constituted by. On my construal of his work, what he is alluding to here is that, when attributing a "belief" while folk psychologising, we are instead picking out the end result of the vastly complex internal dynamical relations of the brain instead of a canonical "belief". The prototypes are the higher order neuronal arrangements which broadly correlate to these attributed beliefs, but do not represent them in their delineated form as understood by the folk. It cannot be assumed that "the intentional idioms of FP" (ibid. 1981: 82) actually represent what is important within our cognitive processes, not to mention that we may have inherited a partly incoherent framework through our folk learning.

Churchland also makes an interesting point regarding the normative aspect of FP: the regularities (or "logical relations") that are said to hold between propositions does not say anything about their essential nature, it simply says what we *value* about the patterns we ascribe (ibid.: 82-83). What we *value* is simply what can be predicted from a macro-perspective in order to understand the behaviour of our conspecifics. Furthermore, looking at the evolutionary development of humanity and the relatively late emergence of language, it seems implausible to model our internal mental states as *fundamentally and essentially* constituted along the lines of language (ibid.: 83). We appear to have mistaken an intuitive learnt grasp of others as an insight into mental necessity[51]. As a result, the old view of FP as sententially-mediated in the form of propositional attitudes does not have a space in a future ontology. It is this which underlies his eliminative materialist project, arguing for the possible reducibility of our current understanding of our FP. Folk psychologising can instead be reconstrued as describing "a kinematics of activation patterns and a dynamics of vector-to-vector transformations driven by learned configurations of synaptic connections" (ibid.). Vectors could arguably be used to account for the compositionality and systematicity of thought, as opposed to the propositional thinking that Fodor endorses. To me, however, Churchland does not make it clear enough just how he demonstrates that propositional attitudes are effectively eliminated on this view. Vector coding is an admirable theory and explanation of how our brains may represent the

[50] This unitary conception of cognition shall be a theme which recurs in Chapter Three's account of the predictive nature of the brain.
[51] To paraphrase an excellent line from Dennett (1991a: 401).

environment, but it serves more of a descriptive function than an eliminativist one. To describe mental states in terms of vector coding, we are not doing sufficient work toward eliminating propositional attitudes. I am not sure that he would deny the existence of FP as a practice, and talk of "elimination" seems to jump the gun somewhat. Elucidating the internal dynamics of the brain would not necessarily necessitate a complete overhaul of our folk practices, but he appears to endorse this rather vehemently[52].

I am, however, partial to Churchland's problematisation of the propositional attitudes due to my interest in the internal workings of our cognitive systems, and the rest of this thesis will be in line with Churchland's original project. I shall show in Chapter Three how a different view of the underlying mental processes could perhaps be the language that Churchland was alluding to when he proposed that an entirely different one replace our talk of beliefs and desires in the literature. I am, however, unsure as to whether this language will suffice as a suitable replacement for propositional attitude ascription. Whatever theoretical terms Churchland argues we will/should inherit from an accurate scientific theory will only be useful insofar as they filter into public discourse. I argue we need not necessarily overhaul our folk terminology, but that these issues will iron themselves out in time.

Ultimately, Fodorian thinking appears to take folk psychological propositional attitudes as the primary mode of cognitive explanation and functioning, whereby they are the atomistic constituents of our cognition, which need to be explained and elucidated. As we have seen, if Churchland is even half right, Fodorian FP is far too basic and literalist for an accurate portrayal of our mental states (Dewhurst 2017: 6-8). On the other hand, Churchland and others who follow the eliminative materialist paradigm seem committed to the idea that if FP fails to sufficiently reduce to an accurate structural portrayal of cognition, then there is no recourse but to eliminate it from our ontology. I view neither extreme on this position wholly useful, as amendment of our understanding of FP can take over from Churchlandian eliminativism just as easily to account for Fodor's failings in the realm of FP. As will be seen later in this thesis, we can "save" propositional attitudes, but not in a way that is commonly thought. The eliminativist is close to a more coherent picture of the inner workings of the brain, but more nuance needs to be incorporated into the framework.

[52] There is also still a sense that his "vector-to-vector transformations" are broadly similar to a LoT, even if the mode of explanation shifts from sentential structures to the interrelations of vectors.

3. Decentralising the Propositional Attitudes

3.1 Problematising the Assumption

There have been other proposals focusing on the propositional attitudes of our FP throughout the past few decades, but it is not worth going through them all here. The above two views served as the prime examples of the debate on our folk psychologising. Despite Churchland's misgivings, the axiomatic nature and predictive capabilities of Fodorian propositions make his an attractive theory, which goes some way to reaching what Botterill (1996: 106) deems is necessary for something to fall into the sphere of valid theoryhood. Botterill makes a cogent point in stressing the need for a theory that is more general in its framework than that which it purports to describe, all while incorporating its disparate elements into a cohesive unit. Arguing for a theory more akin to the laws of Newtonian physics, he rather humorously pointed out a need for more than simply "gardening lore" (ibid. 1996: 109-110) when attempting as comprehensive a theory as possible. As he goes on to explain, Botterill (1996: 106) lists a number of criteria that would be essential for this undertaking, among others, the need to reference "unobservable" phenomena, being able to incorporate counterfactuals, conceptual integration, and the ability to predict phenomena. Ratcliffe (2007b: 9) points out that this all sounds rather similar to folk psychology in its present guise[53]. I argue that the best way of achieving a sound theory would be to start from the bottom up, describing the disparate phenomena in fine-grained terms before moving upward to explicate the multitude of these phenomena successfully under a common banner. This is precisely what I begin to do in Chapters Three and Four. We should perhaps let the chips fall where they may, and reconstruct a coherent theory to replace canonical FP by suspending as many preconceptions as we can, rather than shoehorning them into existing theories, weighted by their bias.

Despite the obvious differences in opinion concerning our folk psychological ascriptions, what appears to be ever-present is the view of propositional attitude ascriptions as central and primary to FP. Looking at Fodorian theories, we see a pattern of suppositions whereby propositional attitudes are elevated in importance and are central to our folk psychological abilities. Bartsch and Wellman (1995), by way of example, construct a behavioural view of how adults mediate an understanding of each other through unobserved and internalised propositional attitudes, but without an argument for why this should be so. Therefore, the

[53] By way of example, the behaviours of conspecifics are incorporated into the folk psychological system in order to predict and/or explain their behaviour, and folk psychological talk also trades in conceptual placeholders such as the typical "belief" and "desire" we are always faced with. On top of this, mental states are also unobservable and often elicit talk of a counterfactual nature (Ratcliffe 2007b: 9).

received FP view is taken as as a given for their subsequent analysis of *how* people use the prototypical FP concepts of belief and desire, with no exploration of whether they *do* (ibid.: 11). As Ratcliffe (2007b: 52) rightly states, there is no actual scientific discovery of the constructs of FP within experiments of this ilk, as FP is already inscribed as a constraining factor within the experiment which *confirms the bias* toward an elucidation of the data in terms of these constructs. Data selection is always already skewed from the start.

To reiterate the central issue at play, there is insufficient questioning of the assumption of the existence of propositional attitudes themselves within our mental lives. Instead of looking closely at how warranted we are at placing propositional attitudes in this primary role, many theorists have begun with the (tacit) assumption that propositional attitudes are central to our cognition, and then have focused instead on whether these attitudes are an innate ability or acquired. Ratcliffe (2007b: 45) maintains that it is at least "arguable" that assent to the existence of internalised propositional attitudes are instead products of philosophical theorising, as opposed to their alleged "commonsense" nature. Arguably, propositional attitudes are *not* very good at picking out the types of thinking that individuals "actually" do. For one thing, on Goldman's (1989: 167) view, there is the worry that the philosophical investigation into propositional thinking does not pick out actual representations within folk thinking. Rather, the literal construction of new propositional platitudes to "accommodate philosophical preconceptions" are further utilised to theorise about these alleged theoretical concepts, despite these being theoretical impositions in the first place (ibid.). At base, philosophers who assume propositional attitudes are central and expand their theorising from there, as opposed to the incorporation or amendment thereof from the outset, are playing language games. Worse, this does not fall within the remit of philosophical circles alone, but expands toward real-world application within cognitive and psychological experimentation and theorising, as with the Bartsch and Wellman (1995) example above.

3.2 The False Belief Task as Emblematic of the Problem

More than any others, much alleged empirical verification of propositional attitudes has been drawn from the false belief task experiment, first proposed and conducted by Wimmer and Perner (1983). They first conceived the experiment to ascertain the age at which children are alleged to ascribe a mistaken belief to other people. The subsequent experiments were careful to pick out when children were ascribing *false beliefs*, rather than simply acting on a lack of

requisite information to describe the state of affairs. Wimmer and Perner (ibid.: 103) state that a suitable conception of someone's false belief arises when there is an "explicit representation of the wrongness of this person's belief in relation to one's own knowledge", and "the ability to represent the relationship between two or more persons' epistemic states emerges" (ibid.: 126) at this point.

The first set of experiments involved children, along with a puppet called Maxi, who were faced with a pile of cookies. Both the children and Maxi observe the cookies being placed in a cupboard, before Maxi exits the stage. Those controlling the experiment then move the cookies to a second cupboard, before asking the children their opinion on Maxi's perception of where the cookies will be once he enters the stage again. Concepts such as "belief" were not introduced to the children during the task, and instead of verbally responding the children were asked to point. This was in order to distinguish between an actual ability to detect a false belief and merely an ability to describe themselves as detecting one. The results have become academically well known, as only children of a specific age[54], with sufficient cognitive capacities, are capable of shifting their perspective to that of Maxi's (ibid.: 122-123). The children who failed the test were deemed unable to grasp the notion that the puppet has his own "perspective", with requisite beliefs, and therefore *could not ascribe a false belief* to Maxi (ibid.: 123-126). The children who failed the experiment were unable to accomplish the requisite perspectival shift, as they did not have the cognitive resources to be able to adopt the multi-perspectival stance that this required. From this, it was surmised that the children did not possess an adequate conception of "belief"[55]. As Frith and Happé (1999: 3) put it, the false belief task is defined as the "litmus test" for any folk psychological ability. Fodor (1995) himself has used it to defend his representational theory of mind by arguing that it shows a form of innateness with regards to propositional modules in the brain. In other words, on this construal children appear to already have the cognitive tools to be able to "understand" belief in a holistic sense, and simply need further brain development to be able to ascribe context-sensitive beliefs[56].

As alluded to above in the applied approach by Bartsch and Wellman (1995), the false belief task becomes a paramount case whereby interpersonal understanding is imposed upon the

[54] Originally from between the ages of four and six, but this has been shown in subsequent experiments to be even lower on occasion (Dunn 1991; Surian & Leslie 1999).
[55] There have been a proliferation of the methods and types of experiments which purport to show similar conclusions to the original false belief task, one prominent example being Baron-Cohen *et al*.'s (1985) application thereof to autistic children.
[56] Scholl & Leslie (1999: 147) as well as Nichols & Stich (2003: 91) have similarly endorsed this view.

structure of the experimentation itself by the experimenters. The experimenters, in other words, assume that full-fledged propositional thinking occurs, and then seek to explain the experiment *in light of what is assumed*. Presupposing internalised mental state attributions, such as beliefs and desires, as central to every person's intersubjective abilities can therefore be seen as problematic (Ratcliffe 2007b: 54). The highly structured environment of these experiments lend themselves as important constitutive elements of what is extrapolated from the data (ibid.: 53-54). By way of example, the original experiments involved a series of debatable suppositions upon the part of the experimenters, including the fact that a puppet could be interpreted as a "person" by the children, which presupposes an already considerable cognitive ability within them. Furthermore, the inclusion of an extensive narrative explaining the situation to the children again presupposes a considerable interpretative ability on their part. Also, the simple fact that the children must be able to understand the experimenter to begin with points to a more socially nuanced interaction between experimenter and child (Gallagher 2001: 99; Ratcliffe 2007b: 54). Introducing propositional attitudes as central to these abilities seems at the very least questionable.

Taken together, these do not point toward grasping a propositional belief in the particular sense that the experimenters are attempting to coax from the children. One wonders instead whether the children respond due to the prompting of the experimenter(s) or whether there is an actual grasping of the epistemic relations at play. The children display a (relatively) considerable ability to even grasp their role within the task, drawing into question the centrality of false belief attribution within the social environment (Ratcliffe 2007b: 53). As has been noted (ibid.: 54), the experiments only required of the children that they point, and not extrapolate their epistemic condition in the form of a propositional attitude attributed to Maxi. This makes the prominent incorporation of the experiment into the literature hardly indicative of vindicating FP as propositional attitude ascription. Therefore, it cannot be unequivocally claimed that the children are in fact successfully attributing a "belief" or not. To successfully infer a particular phenomenon, in this case a false belief, does not exclude the possibility that one is instead focusing on other factors that are typically associated with the phenomenon. It can therefore be argued that the children learn their ability to attribute beliefs within contextually-bound environments of interaction (Hutto 2008a). If these contexts are removed, the ability would be attenuated if what partly constitutes the ability is taken away. Similarly, we could say that the attribution of internalised mental states, as advocated by FP, do not capture the intricacies of interpersonal understanding (Ratcliffe 2007b: 53). If a child is asked to experience the

unfolding of a narrative rather than the more familiar day-to-day interactions with adults (or whomever cares for the children), it should not be assumed that they can infer internal mental states. The children could instead only be able to infer the mental states of others through the highly structured environments of interaction and reinforcement elicited by the children's caregivers. If these structured environments are eliminated from the experiment, we will be doing away with important compositional elements which enable the ability of children to ascribe the alleged propositional attitudes (ibid.: 54). The understanding of the interpersonal would not necessarily be elucidated through experiments which presuppose elements of our FP.

At the end of the day, the false belief task seems to suggest that to grasp a mental state, such as a belief, it is sufficient enough to simply re-enact what amount to "rather crude behavioural patterns" (ibid.). If a child is able to detect beliefs fairly reliably, it does not necessarily extend to the fact that the child is understanding what is being detected. The experiments fail in supporting the assumption that it is merely the inferential reading of internal mental states that is of central value, while ruling out dispositional attunement toward others, and certain sensitivities such as to gaze and expression playing roles in eventual intersubjective understanding (ibid.: 54). These additional factors provide visceral data for bodily reaction *to* the information to occur, circumventing the propositional thinking we are alleged to engage in. For instance, it has been shown how children react to particular language use *without actual comprehension of the words or concepts used*. For example, it has been shown how even at the age of six months, children are more predisposed to paying attention and attending to speakers of the cultural group from which they belong[57] (Kinzler *et al*. 2007; Mahajan & Wynn 2012; Greene 2015: 50-51). Perhaps as the basis for this, it has been shown how neurons are able to anticipate particular environmental attunements dependant upon experience (Shuler & Bear 2006). Therefore the affective responsitivity of these neurons occur at the same time sensory stimulation occurs, and not necessarily after[58] (Barrett & Bar 2009: 1325). This can be extended to facial recognition of others, even at a young age (Debruille *et al*. 2012), wherein facial expressions and other subtle patterns can provide environmental information for subsequent behaviour. This does not necessarily support the explanation given by Wimmer and Perner for what the false belief task actually achieves. The main take-home point from this is that while

[57] There are linguistic cues which dictate social preference and foster group identity. It appears that children tend to be predisposed to favour those who lack a foreign accent (fostering the "Us" vs "Them" dichotomy).
[58] This is more than likely a result of the predictive processing architecture of the brain, introduced properly in Chapter Three.

the false belief task seems to support the existence of propositional attitudes in these interactions, it fails to make the strong case that adherents claim. As Greene puts it, these "arbitrary differences can serve a nonarbitrary function" (Greene 2015: 51), as seemingly innocuous factors which we do not take into account can affect our understanding of what we are trying to model.

3.3 The Emergence of Alternative Viewpoints

Ratcliffe has attempted more than most to problematise the importance of propositional-style theorising (2007a, 2007b).The above has led him to consider whether FP is a "theoretical imposition" (2007b: 45) upon our day-to-day lives in that the sciences and philosophy can result in "a projection of the framework of the latest scientific speculations into the common sense picture of the world" (Sellars 1963: 171, from Ratcliffe 2007b). Said differently, the attribution within interpersonal understanding of "beliefs" and "desires" is an "abstraction from social life" (Ratcliffe 2008: 445). Science and philosophy are understood to elucidate the world to us, but there can be slippage if we are not careful. In other words, even if we do value the intuitive understanding that the folk have concerning their mental lives in the form of their beliefs and desires, we must still be wary of what many of the experts have claimed occurs beneath the surface. Dennett (1991c: 137) has also made a similar point in his distinction between "craft" and "ideology" within FP. There is the "craft" element of FP, defining it as simply what people actually do, and the "ideology" aspect, which is what people *think* that they do. Using Dennett's distinction within FP, Ratcliffe (2007b: 55-57) points out that we should separate the responses from, for example, children and the academic experimenters' interpretation of their responses, which is already informed by a particular ideological basis. Ratcliffe argues that we should distinguish this further into four categories of analysis:

> "(i) What people do
>
> (ii) What people think that they do
>
> (iii) What philosophers and cognitive scientists think people do
>
> (iv) What philosophers and cognitive scientists think that people think that they do" (ibid.: 55).

The first category, (i) what people do, is how people interact in typical interpersonal interactions at the broad social level, while the second category, (ii) what people think that they do, entails the average person's intuitive understanding of what it is that is occurring during their interpersonal interactions. The second does not imply clear facts of the matter, but rather assumes them through folk learning. The third, (iii) what philosophers and cognitive scientists think people do, refers to the scientific/philosophical understanding that philosophers and cognitive scientists have concerning interpersonal understanding (importantly *not* how people typically understand what they are doing during interpersonal interactions). The fourth, (iv) what philosophers and cognitive scientists think that people think that they do, is the empirical or armchair extrapolation of the available data concerning the reports of individuals from the viewpoint of the scientists/philosophers. What Ratcliffe (ibid.: 56) proposes is that *even if* we do not take people at their word, theoretical FP is still an equivocal intermingling of (iii) and (iv) above, an imposition of theorising upon daily life, and not an accurate description of reality. Whereas what theories like the false belief task are actually picking out is simply (i), and occasionally aspects of (ii), in other words the base understanding of the average person of their interpersonal interactions. One could take this further and claim that (i) is interpreted through the lens of (iii) and (iv), which effectively feeds into Dennett's point above. None of these options *necessarily* entail actually picking out any actual mental belief or desire states.

But Fodor (1987: 9-10) has gone so far as to claim that our common-sense belief/desire psychology is so ingrained in our everyday practice that it *should* be seen as transcendentally derived in a Kantian sense, and empirical enquiry had better vindicate the workings of FP as currently understood. In other words, what the false belief task should be doing, on this construal, is *proving* the validity of our FP as we understand it. Despite this claim, and the extensive defence of it (ibid.: 10-26), just which elements of our FP should be preserved on these grounds is still not clear. The veracity of FP in its typical guise is, at the very least, debatable, but Fodorian FP, problematically, refuses to let go of the importance of the propositional attitudes within its framework, as has been noted. However, arguably, describing the relevant intentional states in the form of propositional attitudes overly simplifies intentionality, since general dispositional states, such as emotional reactions to bodily irritations or a general undirected anxiety, are notoriously difficult to express verbally in the form of propositions. Perhaps this is more the "failing" of our language range than anything else, but propositional attitudes do not necessarily capture the requisite information to

accurately describe many mental states. This only further problematises the accuracy of propositional attitudes within our FP.

Due to the issues highlighted above, this drives me to the need to de-essentialise propositional attitudes in our talk of FP, to bring forth the physical and social nuances at play. This would entail a look under the hood of our neuronal brain processes, not to mention the interactions between body and environment, in order to tease out more of an understanding of the mental states at issue when we folk psychologise. It should be stressed that this is not explicitly an eliminativist project[59], but rather a delving below the surface of a common philosophical and social practice. In so doing, we may discover a plethora of factors that inform our social understanding and which influence our use of the concepts of "beliefs" and "desires" in folk psychological discourse. A perfect analogue to the general project of de-essentialising the core elements of folk psychology is that of Hayes's (1978, 1980) original deconstruction of our folk physics[60] as axioms of knowledge that we all utilise in our daily lives. These axioms could be as banal as knowing (expecting) a loud bang to occur as a door is about to slam shut, the absorptive effect of cloths, that the throwing of a ball toward a wall will result in its bouncing back toward you (depending on the properties of the ball and the wall), et cetera. We are, in a very literal sense, dependent upon these axiomatic, often predictively correct, and simplified understandings of the physical environment around us. The immediacy and productive elements of these understandings provide prolific predictive abilities within our general environment. As Dennett (2005: 31-32) points out, in a piece on the efficacy of his heterophenomenological[61] approach to understanding behaviour, one cannot simply turn off one's "expectations" of what will occur[62]; it is ingrained upon our mind's understanding of the world. Hayes (1978) attempted an interestingly scientific and anthropological study of the counter-intuitive reality of physics in juxtaposition to the, what Dennett calls, "aprioristic anthropology of naïve physics" (Dennett 2005: 32). The aprioristic element of this concept implies the stark resistance of the folk to any counter-intuitive implications to their learnt understanding of the physical world[63]. Dennett (ibid.: 33) proposes just such an approach to

[59] Although, in effect, beliefs and desires as many take them to be will be problematised.
[60] Hayes' term for this was naïve physics.
[61] Heterophenomenology is the practice of applying scientific principles to the reports of individuals and other relevant information in order to ascertain the individual's mental state. The "hetero-" prefix indicates that it is not a first-person perspective, as with more traditional phenomenological approaches. The point is to ascertain the individual's perspective without attributing a high sense of accuracy to their reports.
[62] I shall be returning to this point in Chapter Three of this thesis with reference to the predictive nature of our minds.
[63] By way of example, we all tend to experience the remarkable effect that capillary forces exert within a straw. The liquid within the straw gets drawn up, against gravity, by the surface tension and adhesion of the liquid to the

our folk psychological assumptions on these grounds. We can predict the behaviour of the folk (and objects), but this does not mean we understand the more fine-grained reality of the very same. We utilise axiomatic generalisations unquestioningly, and it is just this unquestioning acceptance of what we do that can be analysed at a phenomenological level. Dennett calls this approach "sophisticated aprioristic anthropology of folk (naïve) psychology" (ibid.) (although, one wonders whether this was done tongue firmly planted in cheek).

In a similar move, Hutto (2008a: 2, 26-32) has expressed concern with the conflation of having an intuitive grasp of "beliefs" and understanding what it is to actually act for a reason[64]. He points to the six month gap between children assimilating an ability to attribute desires and their capacity to pass any false belief experiment (Hutto 2008a: 135-136; similarly alluded to by Ratcliffe 2007b: 17). He points out that children's ability to attribute desires to others, for example, arises long before an adequate grasp of what they entail via culture's "linguistically scaffolded mentalistic" framework[65] (more of which to follow below) (Hutto 2008a: 135). In other words, as has been alluded to above, the narratives children are faced with in their daily lives prime them for a culturally-mediated understanding of beliefs/desires and uncountable other mental states, and the interrelations of these, for the purpose of reason explanation. In line with this, and to pour even more cold water on the interpretations of the false belief experiments, it has been found that there are a vast plethora of factors which enable earlier capability in conceptual understanding in children. These can range from parental involvement, to size of family, and even, as alluded to earlier, the skewing effect of experimenting in atypical social environments such as during the false belief task (Dunn 1991; Surian & Leslie 1999). All of these factors produce an effect on the results of the false belief task experiment through unforeseen means by the original experimenters. The implications of this for theoretical FP is that if these experiments are not strengthening our understanding of FP in its canonical guise, then the belief/desire attributions to children appear to be *mis*attributed.

---

inner surface of the straw. And yet our minds still resist the scientific explanation given due to our axiomatic understanding of how things "should" behave. The liquid "shouldn't" act against gravity according to our folk understanding.

[64] He, quite rightly in my view, states that to understand an action as performed for a particular reason, it is not enough to appeal to a singular propositional attitude. We are in need of very many of them, not simply a lone desire, which all act together in "motivational crime" (2008a: 26).

[65] For instance, in one study, data gathered from conversational interaction between children showed that there is a propensity to "not use such terms as intend to, on purpose, or mean to until about 3, 4, or 5 years [of age]" (Wellman & Phillips 2001: 130). In contradistinction to this, the desires of others are already alluded to at around 2 years of age, primarily through the use of the word "want" (ibid.).

To be clear, and of more interest to this thesis, I argue that our concepts of "belief" and "desire" do not correlate to delineated mental states. There are many instances of advocates of theoretical FP who tacitly endorse the reality of "beliefs" and "desires" as instantiated on the neuronal level. For example, Nichols and Stich (2003: 14) attempt to describe the mechanisms which underpin our FP by appealing to "belief" and "desire" boxes in a flow chart that purports to describe the architecture of the mind. In fact, these mental states are far more indeterminate than these proponents of traditional FP would have it. The psychological individuations which are the homogenised "beliefs" and "desires" that we use in talk of FP cannot easily pick out our reality. If we are within the same physical space as another individual, we do not (normally) explicitly assign a multitude of individuated mental states to that person such as "D believes that there is too much noise coming from down the corridor", or "D believes that there is a box of chocolates stashed in the cupboard". We do not attribute such mental states, as we are in a shared environment within which interpretation of the other occurs (Ratcliffe 2007b: 191). In other words, there is not always a need to attribute a belief or a desire to understand behaviour.

Furthermore, the "belief" we attribute to someone who is, for example, moving forward from a traffic sign does not easily track reality as their perception is intimately bound with their behaviour. We may attribute the belief "D believes the traffic light is green", but it is difficult to distinguish between the perception of the need to react, for instance, and the intuitive, phenomenal sensation of "understanding" its practical implication. Can it really be said that there is a "belief" in this instance, or are action and perception so intimately interlinked that assigning a belief becomes a redundant gesture after the fact[66]? The concept "belief" does not mirror a particular mental state, but rather speaks to an abstract method of referring to a range of mental events. We appear to be describing psychological states on the assumption that these explanations pick out the *"actual"* mental states in question. But these explanations are in fact given when we stop attempting to describe the actual behaviour of others, and resort to these general explanations in accordance with our shared cultural norms of understanding (ibid.: 195). These reason explanations make the behaviours of others understandable, without necessarily indicating *the* actual reason for their behaviour. As Ratcliffe (ibid.) puts it, the norms utilised for the purpose of the reason explanation of behaviour infer good reasoning on the part of the one who is acting, but this is not the same as saying that this is *a* or *the* reason for their behaviour. Furthermore, simply adding more such detail to our descriptions of a belief

[66] Chapter Three's introduction of the predictive capacity of the brain blurs the lines between action and perception even further.

or desire attributed to a conspecific does not elucidate the actual state of affairs at the neuronal level.

Arguably, subtle variations in psychological states allow for a plethora of potential explanations for given behaviour, but these would probably not consist in a particular state of canonical "belief" as postulated by FP. This results in the inevitable conclusion that there are no exemplar cases of belief, merely an uncountable number of mental states which fall *under* the incredibly broad blanket term of "belief" (ibid.: 187-190). Subtle distinctions can be explicated within brief narratives (which we all do) in order to elucidate a given state of affairs (a particular "belief" for example), rendering FP highly socially constructed. Although, I argue, not even these narratives that we tend to provide (once prompted) of particular "beliefs" or "desires" provide a description of the unitary mental state of the individual and, therefore, fail to denote a distinct psychological state of affairs. On top of this, the emotional influence of a scenario is intimately bound up with the current mental state of the individual, rendering a description of "belief" in every such scenario necessarily an attenuated version of reality (ibid.: 198). There is an acknowledgement within the literature that emotions are often experienced as intentional states, but *not* of the propositional variety (see, for example, Chapters Four and Five in Solomon (2003) as well as the essays in Solomon (2004)). Within these emotional states, our bodily dispositions, internal bodily milieu, externally-derived affect, and the like are incorporated into the matrix of our intentional attitudes. This does not leave much space for obviously individuated propositional thinking of the sort utilised in many explanations of behaviour in our folk psychology. Therefore, we cannot have a neatly-packaged, comprehensive, *accurate* account of behaviour without "emotional experiences, experiences more generally, perception of agency, gesture and intention[67], normatively configured situations, character, motivation, habitual tendencies and so forth" (Ratcliffe 2007b: 199), and this seems an impossible endeavour for the "folk" to attempt. To explicate as accurate a portrayal of reason explanation as possible, an individuated reason in the form of a propositional attitude is not enough to account for the agent's behaviour, hence the need for the incorporation of the above-quoted aspects within the reason explanation. This leads me to endorse a view of belief states as comprising indeterminate dispositional states. We cannot account for more than "macro-reasons"[68] in our explanation of even our own behaviour, from within our own "privileged" phenomenological vantage point.

[67] In the cognitively loaded sense.
[68] These can be defined as the canonical exemplars of beliefs and desires we all use to describe behaviour.

Dennett has argued similarly in his assertion that there is "no fact of the matter" (1991a) in our cognitive deliberations, and by extension others' ability to interpret the same. His view of our consciousness as a multitude of "drafts" informs this position, whereby our phenomenal experience is constituted by narratives of mental states that compete to "win out", with no clear singular "I" to be found within this interaction except through hindsight (ibid.). This results in a view of consciousness that is fuzzy, opening up space for a multitude of competing dispositions. If true, this also leads to a non-propositional view of our mental states, undermining the efficacy of FP as propositional attitude psychology (as accurately depicting reality). Problematising the singular "I" that has clear reasons for its behaviour problematises the distinction between competing propositional attitudinal states. These critiques serve to elucidate the validity of doing away with the propositional attitudes as accurate portrayals of our mental states, and I shall go further in this regard by introducing the concepts of predictive processing and teleosemiotics in Chapters Three and Four respectively in order to further erode their standing. For now, I shall make a brief segue to a theory on how narratives potentially provide the necessary propositions that we utilise in our FP, before concluding the chapter.

4. The Narrative Practice Hypothesis

In this next section, I want to build upon what has been said concerning the need to de-essentialise the theory of propositional attitudes with an extension of Hutto's theories. I will be dealing specifically with his Narrative Practice Hypothesis (NPH) (2004; 2007a; 2007b; 2008a; 2008b) or, more specifically, what it entails and how this potentially elucidates our FP within our structured but ever-shifting contextual background. Its purpose is to introduce a plausible alternative to theories of innate propositional attitudes and to explain how it is that they may arise if they are not innate. My own thoughts on the validity of this view shall follow, and I shall remain neutral on its explanatory power up until that point.

Hutto strongly believes that what accounts for our impressive ability to implement our familiar folk psychological attributions is the reinforcement of *narratives* during our ontological development (Hutto 2007a; 2007b; 2008a: 23-40; 2008b: 177). The foundation necessary for attributing reasons for actions does rest on the basic capacities already available to us. Yet, children do not function in their early months of social interaction by attributing any beliefs or desires to an external agent's actions (Hutto 2008a: 135). Only over time do they start to utilise the now-familiar propositional attitudes in order to navigate their interpersonal niche. These

propositional attitudes are learnt through a step-by-step process whereby subsequent mastery of the syntactic (and eventually semantic) components of their mother tongue are assimilated into their future interpersonal cognition (Hutto 2008a: 27; 2008b: 182). Hutto states that it is not the inherent capacities that are missing in early childhood, but the necessary, for lack of a better word, *rules* of their co-integration, which are learnt through social interaction (Hutto 2008a: 27). Yet, any resultant competency in the assigning of propositional attitudes to external agents does not necessarily translate to actual *knowing* in the sense of knowing how they can be recombinantly constructed to formulate an understanding of reasons (ibid.: 28). In other words, it does not necessarily denote the existence of a fully-functioning folk psychologist, in the sense that an individual can understand, interpret and manipulate the norms which govern how the propositional attitudes help understand reasons. Naturally, the normative interrelations of these attitudes must be grasped and understood. As Hutto states,

> [w]hat children are missing, on acquiring a practical grasp of the concept of belief, is not therefore another ingredient needed for baking the folk psychological cake – rather it is the instructions for mixing all the ingredients properly in the making of *many* such cakes. (ibid.) [emphasis added].

The proposal for how these norms arise from a child's ontogenetic development is simple enough. Hutto introduces the NPH in order to elucidate this transferral of folk psychological skills. In short, it is the direct experience of narratives (stories) about people who are clearly presented in a manner whereby there are *reasons* for their actions. These are what Hutto calls "folk psychological narratives" (Hutto 2008a: 28; 2008b: 178). As Goldie (2004: 115) puts it, the narratives are constructed in order for "…the audience or the reader to make sense of the thoughts, feelings, and actions of those people who are internal to the narrative".

What should be taken away from this is that the above narratives are what provide the necessary ontogenetic preparation for the elucidation of reasons for actions. These narratives are provided as "exemplars" (Hutto 2007b: 53, 2008a: 28) which nurture an inherent comprehension of the normative components of folk psychological understanding. Children acquire an understanding of how the propositional attitudes can be recombinantly formulated in order to explicate the requisite belief, desire, and so forth, along with the emotions that accompany them. Furthermore, these narratives provide the contextual background, which presents the myriad ways in which an agent's reasons can be shaped and guided by the immediate context, the

agent's historical backdrop, and other things[69] (ibid.). This type of co-cognitive[70] development also appears to occur at the requisite point within ontogeny to lay the framework for more advanced folk psychological ascriptions (Hutto 2008a: 28-29).

The manner in which these folk narratives can perform such a function is due to their existence as objects of *public* consumption that represent complex social interactions (Hutto 2007b: 67, 2008a: 29). As clarified earlier, these narratives act as exemplars, but in addition, they are the requisite apparatus needed for reinforced learning over time. Children learn the interrelated rules and frameworks that arise from the continual reinforced use of established folk psychological attitudes, and are steered toward an emphasis on beliefs and desires[71] as typical instantiations of propositional attitude ascriptions (Hutto 2007b: 55, 2008a: 29, 2008b: 178). The propositional attitudes are introduced along with the methods of their interrelation, and in this way children are taught to perpetually recognise the way in which beliefs and desires seem to influence actions[72] (Hutto 2007b: 55-56, 2008a: 29). As should be clear, these principles of interrelation are not conveyed in an easy list of rules, but reiterated through continued learning.

As an important aside, this would bias any facets of explanation within the folk psychological framework toward already-used and already-understood uses of the propositional terms within a context-bound culture, as well as bias the ways in which these terms are thought to interrelate (Hutto 2008b: 182-183). It is far easier to make use of already-proven methods for social navigation than to develop detailed, fine-grained reason explanations for any behaviour[73]. For simple reason explanation, to understand what the specifics of what an agent's particular beliefs and desires are, how they change over time, as well as the contextual information necessary for this understanding, not to mention the *why* of the above, is provided for through these exemplar narratives (Hutto 2008a: 29). In other words, this is the methodology through which children acquire the more nuanced folk psychological understanding necessary for understanding reasons, and what they need to fine-tune when attempting to make sense of the above. This would naturally extend to explanations of their own behaviour, when needed (ibid.). As simply

[69] The ways in which children are given the requisite support to flesh out their reason-explanations can be found in Chapter Ten of Hutto's (2008a).
[70] By "co-cognitive" I refer to the ability of an individual to incorporate the inferred thoughts of a conspecific within her cognition, thereby being able to predict the conspecific's behaviour.
[71] It should be remembered that beliefs and desires do not occupy exclusive space within the pantheon of propositional attitudes, but it is clear that they are what we could call the exemplars of the exemplars.
[72] The ways in which beliefs and desires allegedly lead to action is perhaps what can be seen as the central thesis of folk psychology.
[73] Chapter Three's focus on the predictive capacity of the brain provides a window into how the brain focuses its energies on these already-extant social grooves on the macroscale.

as this story is presented, it is a compelling explanation for our early acquisition and understanding of the necessary fundamentals of our folk psychology, for how they are applied, as well as of the norms inherent in our social understanding. All of this is supplied without any overt rules describing how our folk psychologising ought to work. Instead these tacit fundamentals of our FP are acquired through repeated engagements in extended narratives depicting the contexts in which they function. The variations within these narratives convey a sense of what needs to be accounted for when explaining the reasons of conspecifics, as well as their own. This is why Hutto states:

> In this way children learn which kinds of factors must be taken into account and adjusted for when it comes to making sense of the stories that others tell about the reasons why they acted, as well as learning what needs mentioning when providing their own accounts. It is in this way and in this sense that children acquire an understanding of the core structure of folk psychology, its governing norms, and guidance on its practical application. (ibid.)

Fundamentally, the tacit norms within behavioural understanding are learnt through the above narratives (Hutto 2008a: 32; Hutto 2008b: 179). If Hutto is right, a theory of our FP must account for the application of these norms within social reason explanation insofar as these are acquired tacitly via learnt narratives. We need to work out a view that can incorporate our *practically-applicable* knowledge into the frameworks as described by the majority of FP theorists and their ilk (Hutto 2008b: 185). We must not lose sight of the fact that folk psychologising is primarily a practical endeavour, constantly and consistently called-upon to provide contextual understanding across multiple scales of interpretation. There is relatively little mention among FP theorists of just how children are able to acquire the practical applicability of these rules in their day-to-day existence (Hutto 2008a: 34). Adding up the ingredients of belief/desire and the assumption that the person is not irrational is hardly enough when the person's psychology, history, and context are vastly important for a suitably comprehensive explanation of their behaviour (ibid.: 34). On Hutto's view, FP narratives achieve this because they depict more than the aforementioned "idealised rational agents", but also *non-ideal* rational agents. They depict how both rational and irrational behaviour is engendered within the confines of pertinent contextual factors (ibid.: 35).

Intuitively, Hutto seems right. When we attempt to provide a reason for a person's actions, we rarely (I would contend never) offer anything more than a truncated view of her social reality. Instead, we have an embodied understanding of the social context which is then neatly

packaged in "just so" tales for public consumption. But these "just so" tales do provide a streamlined explanation of behaviour which gets us by in our social environments. On this view, any examples from folk psychology utilising beliefs and desires are "just so" tales. "S believes that p" is effectively a social imposition. There is more to the tale than what has been included in any "just so" story, and we generally only apply reason explanation on a "need-to-know basis" (Hutto 2008a: 36). Attention to learned behavioural patterns usually provide enough sense for guiding our immediate reactions and we accomplish much of our our social interactions without the need for explicit explanations of behaviour, which prompts Hutto (ibid.) to declare that in our day-to-day social environment we "do not need to interpret what others are doing or why by asking about their reasons".

As a result of the above, treating FP as an *accurate* theory of how our social minds work should be discouraged due to "the bad effects it has had (and continues to have) on the imaginations of many philosophers, psychologists, and others working on this topic" (Hutto 2008a: 32).

As compelling as this view on the acquisition of our folk psychological abilities is, I remain unconvinced that this much emphasis should be placed on narratives alone[74]. Hutto does not deny, and indeed has advocated (Hutto & Myin 2017: 253; Chapter Six), that there are other factors at play when we explain behaviour. Hutto is correct in that coarse-grained intentional patterns can be discerned within our interpersonal narratives, but it is far more complicated than this theory would seem to suggest. He attempts to stay in the middle-ground with iterations of "belief" and "desire" that are not too explanatorily broad or narrow, making our use of these folk conceptions permissible, if factually inaccurate. However, one wonders whether this is truly informative or another "just so" tale that does not quite capture the psychological sophistication that we all appear to employ on a daily basis. Nevertheless, Hutto makes a convincing argument up to the point where we *should* invoke the NPH. There could just as easily be a space for it alongside a broader and potentially all-encompassing framework for the acquisition of our folk psychological ability. It does, however, appear that it should be utilised more as a complementary conceptual tool which explains the base acquisition thereof. It is for this reason that a theory of a more generalised "enculturation" in Chapter Four is introduced below to provide a broader framework for the acquisition of our FP abilities. The NPH was introduced to lay the groundwork for a plausible, non-nativist theory as to how propositional

[74] Ratcliffe (2007b: 214-216; 253, footnote 16) also has reservations concerning the all-explanatory nature of the NPH, despite their similar projects.

attitudes are actually introduced into our lives, and narratives appear to go a long way in terms of explaining their origins.

5. Conclusion

This chapter served to introduce the field of folk psychology by elucidating how it is thought (by many) to operate within our cognitive make-up. The notion of propositional attitudes as atomistic constituents of our cognition as well as their centrality to our FP was expanded upon. Special reference was made to Fodorian FP, which is a straightforward form of propositional thinking, in order to introduce the disjunction problem as a problem that stunts our intuitions regarding adaptational theories of cognition[75]. Thereafter, Churchland's eliminativism toward propositional attitudes in the form of vector processing was introduced. The introduction of these two theorists also served to highlight the poles of the debate, although neither position was advocated for. Fodorian propositional attitudes were argued to be not as fundamental in terms of our cognitive functioning as some assume, while the outright elimination of propositional attitudes was also discouraged. The core issues with the assumptions inherent in theorising about FP were then expanded upon in order to problematise talk around FP as propositional thinking in general. The false belief task was used as an emblematic example of biases introduced during theorising of this ilk. This was further utilised to endorse a de-essentialising of propositional attitudes as physically realised in the brain. A window into how propositional attitudes possibly get introduced into our discourse was given in the form of the NPH. It was explained that the introduction of narratives during ontogeny could potentially be a substantial building block for our propositional thinking, as opposed to some form of ToM innateness which Fodor (1995) would endorse.

After all of the above, an important consequence to consider is the possible evolutionary benefits of folk psychological ability. One can endorse a view of FP which depicts it as an end-product of evolutionary development, but one in which both learning and pre-existing cognitive architecture constrains it toward a particular form. This view of cognition and FP is predominantly what I shall be working with throughout the rest of this thesis. The view of propositional attitudes as demarcated and analysable on the neuronal level does not appear to align with evolutionary reality. Also, disjunction "problems", whereby it is alleged that the acquisition of mental contents cannot obtain without a suitable demarcation of the necessary accuracy conditions, do not appear to be as intractable as they are taken to be, as I shall argue

[75] The significance of this point will become apparent in Chapter Four.

for in the following chapters. These issues point away from the neuronal reality of propositional content, and instead toward the social learning of propositional concepts as useful heuristics for behaviour. In line with Humphrey's (1976) idea that what early mankind was predominantly concerned with was the social understanding of conspecific behaviour, as well as Dennett's (1991a; 1995; 2017) views of how language shaped the very consciousness we all know and love, the purpose of evolving folk psychological abilities lies in the selection pressures of social attunement. Therefore, we can plausibly depict FP in an "as if" fashion, while problematising the assumption that it is actual propositional thinking which is instantiated within the brain on a fine-grained level. FP, as Ratcliffe (2007b: 23) puts it, is

> an abstraction from social life that is misleading in various respects and has *no psychological reality*. At best, it is a convenient way of talking in certain areas of philosophy, which has become an entrenched and misguided philosophical institution [emphasis added].

Ratcliffe argues that this is a direct result of the mechanistic naturalism found within the analytic and scientific traditions (ibid.: 231-243). These research paradigms have sought to explicate a narrative of folk understanding that presuppose biased notions of what it is to understand conspecifics. I agree with this analysis, but we must be careful. For to deny "mechanistic naturalism" on his construal is to ignore the mechanistic (but intractable) nature of even the social world[76]. Saying it is this view of the world which is at fault in our assessment is not the same as saying it is this view of the world which is *incorrect in general*. We can analyse how our internal processes function through a mechanistic naturalism in order to explicate the true nature of our mental states, regardless of the indeterminacy inherent within them. Arch-scientism has perhaps not always helped our interpretations of the crushingly complex world both outside and within us, but difficulties and over-simplifications do not necessarily entail that a view is incorrect. Endorsing methodological naturalism, which is a form of naturalism that is congruent with the methods and aims of the sciences, is a way to cut through these challenges. Methodological naturalism allows us to suspend our disbelief when faced with seemingly intractable problems and work toward cogent explanations. We need not worry about finding determinate mental states in our cognitive make-up if that is where the science lies. As I hope to show in the subsequent chapters, by incorporating the latest relevant scientific research into our conception of FP, we will arrive at a more nuanced understanding

[76] By the mechanistic nature of the social world, I mean simply that there are "rules of engagement" which constrain behaviour, which in turn affect the formulation of our folk psychological ability.

that shines more of a light on the nature of our mental states through systematic scientific analyses. This shall also highlight that our phenomenal experience is not as rich in representational content as we believe it to be, undermining the traditional conceptions of FP as consisting of determinant mental states. To achieve this, we shall first look at a rising discipline within the realm of cognitive system dynamics in the following chapter: that of the Bayesian nature of the brain.

Chapter Three: The Mind as Predictive Processor

1. Introduction

In this chapter, I shall introduce recent literature surrounding dynamic system theory in cognitive science, which concerns the fine-grained cognitive processing that underlies our cognition. This will subsequently inform the theme introduced in Chapter Four of this thesis. My focus in this chapter will primarily be an account of the neuroscientific theories of hierarchical predictive processing[77] as described by Clark (2012; 2013a; 2013b; 2013c; 2015; 2016), with much seminal work done by Friston and colleagues (Friston 2005, 2010; Friston & Stephan 2007) as well as Hohwy (2008, 2013). [78] The reason for introducing predictive processing (PP) here is due to its rising popularity within the neurosciences and cognitive sciences, forcing us to take its implications seriously. There is a growing consensus that the PP approach could be the conceptual tool that could unify many aspects of consciousness and/or cognition. For example, it could unify cognition, action, and perception to such an extent that the conceptual joints meld together, problematising old distinctions. It could furthermore unite differing research paradigms under the same PP umbrella, bringing not only different empirical sub-disciplines together, but also uniting theoretic-philosophical approaches with the empirical. This chapter will first describe the predictive processing paradigm, elucidating it as a form of Bayesian processing and describing its general foundational characteristics in this light. Then I shall relate these Bayesian roots to the functioning of the brain, introducing the concepts of active inference and free energy. Thereafter, I shall look at the results that these features would have on our conception of the mental in that the Bayesian model of the brain inverts our understanding of the brain as a passive receiver of environmental information, thereby upsetting our previously conceived folk notions of ourselves. Therefore, the *implications* of this view of the mind/brain in terms of FP will be elucidated and critiqued.

[77] Originally formulated explicitly in Friston (2005), Lee & Mumford (2003), and Rao & Ballard (1999).
[78] While I align myself closer to Clark, Hohwy's view is highly similar, with only the glaring exception of the solipsistic nature of our predictive brains standing out (Hohwy 2008, 2013).

2. Bayesian Prediction and its Relation to the Mind

2.1 Bayesian Prediction and Predictive Processing

PP in its current formulation appears to be a promising conceptual tool to unify cognition, action, and perception under a common banner. To be clear, it is a computational tool, and is utilised in an abstract sense to capture the underlying neurophysiological processes of the brain[79]. These computations are thought to predominantly feature below the levels of phenomenal experience, although I contend that these foundational computations could be applied at higher levels of personal phenomenality. I argue that these underlying computations *could* upscale to folk conceptions of reality, in the form of behavioural reason-giving and other phenomena (more to come further below).

A core tenet of PP can be found in the writings of Kant's *Critique of Pure Reason* (1998: 110-111) wherein he describes the role that our faculties of intuition (our *Anschauungsvermögen*) play in the *formulation* of our sense-data. In other words, our sense-data is not simply "given" to perception, but at least partly constituted by it. This is not the received view of classical scientific conceptions of how our brain is thought to process information—that of the "passive, stimulus-driven" (Engel *et al*. 2001: 704) variety[80]. PP takes further inspiration from Helmholtz (1985), who developed Kant's line of reasoning into the notion that perception is more of a subconscious process of inference. He stated that the effects of sensory reality are only perceived through the constraining actions of the mind, rather than by simply taking up the sensory reality itself. In line with this, PP depicts perception as far more actively constructive than first thought. It is also far more selective in that higher processing can influence and alter the sensory intake, and thereby utilise what is "given" (or "accepted") to inform predictions relating to potential (future) sensory impingements (Swanson 2016: 9). This is a predominantly top-down depiction of cortical processing, wherein prior experiences of the environment inform cognition. Consequently, these prior experiences are given pre-eminence and are all-pervasive in perception. The standard formulation of PP depicts these prior experiences of the environment as modelled statistical estimations in some form within the brain (Clark 2016: 21). These estimations are functionally similar to "representations" of the environment,

[79] There have, however, been theories of just how these computations can be actualised in the brain's substrate (Engel *et al*. 2001; Friston 2005). Furthermore, the reality of Bayesian implementation in the brain has been explored to ascertain whether Bayesian frameworks are actually being instantiated within the brain (Maloney & Mamassian 2009; Maloney & Zhang 2010).

[80] This "passive, stimulus-driven" conception of brain processing is effectively that of the brain taking up information and inducing its cognitive functions purely through external stimuli.

although I shall problematise this view in the following chapter. Furthermore, these "representations" are continuously updated as a result of any detected misalignment between predictions of potential sensory uptake and what is actually received. This happens in service of minimising errors in future predictions by the "representational" models of the brain, thereby adhering to the norms engendered by Bayesian inference (Hohwy 2008: 17), more detail of which is to follow in the next section. Finally, and this truly begins to blur the lines of perception and cognition, these predictive "representations" are depicted as primarily "action-oriented" (Clark 2013a: 185) in nature. This is all in the service of "predictive control" (Wiese & Metzinger 2017: 4; Seth 2015) for successful navigation of the body through the environment. Many believe (controversially) that this means that action, so conceived, could be more pertinent than perception for the survival of the organism, with perception relegated to the (still significant) role of acquiring information about the environment, as opposed to instigating the body to sample the environment in the first place. Environmental information is used in service of action on a contextually-sensitive scale. Effectively, the information which is incorporated into our cognitive deliberations is pertinent only insofar as it instigates effective action for environmental survival. But, before going into more detail concerning all of these features of PP, we need to look into PP's relation to Bayesian inference.

Bayesian inference is closely tied to Bayes' Theorem[81], which is a statistical formula that aids in prediction of outcomes where there is a degree of uncertainty (Hohwy 2008: 17). This formula incorporates considerations of conditional events that have a bearing on the outcome. These events are the plethora of factors pertinent to whether a particular predictive outcome would occur. One consideration is the demarcated conditional probability[82] that is used to predict the possibility of an event actually occurring if it is taken that another is, in fact, true (ibid.: 18). The general idea is to feed data into the equation to predict the degree of certainty with which outcomes of a specific set of events or actions may occur. As an example, if an agent experiences uncertainty as to which set of assigned hypotheses holds true, the agent can assign differing hypotheses concerning the outcome with degrees of probability. These degrees of probability are effectively degrees of belief in a particular outcome. The result of this leaves the agent with an increased degree of certainty with which to act upon the world as it becomes

[81] It is named after Thomas Bayes (1701-1761), who was an English mathematician (as well as theological philosopher). His original formulation has developed significantly since the 18th century.
[82] This simply refers to the likelihood of an event occurring if another event has already occurred.

clearer, in theory, which new data needs to be incorporated into the working hypotheses[83] (ibid.: 16-17). In Bayesian terms, this leads the agent to incorporate this latest data into an *a posteriori* (or posterior) computation by conjoining an *a priori* (or prior) computation with the probability (ibid.: 17). This prior acts as the data already available to the agent, while the probability pertains to the relation of the data within a target domain with the new target domain's data. This all serves to elucidate the probability of a state of affairs occurring after new data has been incorporated into the prior. This, in turn, can be used to diminish uncertainty, updating the prior with more precise estimates of probability within the target domain (ibid.). The concept used within the literature to capture this process is labelled "Bayesian inference", as mentioned above. The inference between the prior and the posterior is key in the understanding of Bayesian modelling of this sort.

If we were to map this form of inference onto the mind, we are provided with a powerful tool for not only probability estimation, but for modelling *perception in general*. Bayesian prediction, generalised to Bayesian inference, can be postulated as a fundamental component of PP, which is thought to underpin perception. The priors that model the environment (also known as *percepts*), are the starting point from which posterior models of the environment can be inferred at the level of the sub-conscious (Friston 2005: 822; Hohwy 2008: 17-18; Clark 2013a: 183). Mapped onto the cognitive domain, the priors are analogous to neuronal weightings, which model the environment prior to sensory information being taken up. The posterior is analogous to neuronal weightings that model the environment *after* sensory information has been incorporated and has "informed" the prior. Before going into more detail, it should be noted that the assumption is that perception operates on a "need-to-know" basis, incorporating only pertinent data from the environment through Bayesian inference.

A lingering issue here for perception is that the data obtained from sensory stimuli is, as stated, *inferred* from the environment, leaving us with Hume's (2009) long-standing problem with inductive processes[84]. We cannot be certain of the validity of our perceptions of the environment, which results from such inferred data. Hence, our brains can only perform Bayesian inferences to the best possible approximation (Clark 2013a: 183). It is important to understand that, here, the *perception* of the environment does not necessarily entail an *accurate*

[83] As an example, one could infer the conditional probability of a patient's susceptibility to liver disease. If the patient is an alcoholic, this would be a pertinent bit of data which would inform the subsequent calculation. It would need to be calculated with a number of other variables (such as how many patients who are alcoholics tend to have liver disease etc.) in order to arrive at the best degree of probability.

[84] Hume (2009) famously argued that it is impossible for our causal inferences to be certain. This is known as an *inductive* inference whereby a conclusion is not fully supported by its premises.

portrayal of the environment. As noted by Wiese and Metzinger (2017: 4-5), Dennett has a humorous take on this age-old problem in his "robot control room" intuition pump, where you are asked to imagine yourself trapped within the control room of a robot (which is on board the robot):

> The robot inhabits a dangerous world, with many risks and opportunities. Its future lies in your hands, and so, of course, your own future as well depends on how successful you are in piloting your robot through the world. If it is destroyed, the electricity in this room will go out, there will be no more food in the fridge, and you will die. Good luck! (Dennett 2013b: 102).

On this construal of the mind, "we", operating within the limiting confines of the body, are provided only with an *indirect* access to the environment (in the case of the robot: through its sensor array). This implies that we are secluded from the objects of interest in our environment, furnishing us with the epistemological problem that the results of our actions can only be indirectly inferred[85]. Bayesian inference potentially provides us with a way to successfully elucidate our world and gain sufficient perspective thereof[86]. The epistemological concern here is known as the "inverse problem" (Harkness & Keshava 2017: 3; Wiese & Metzinger 2017: 5). The inferential certainty needed to derive a working (Bayesian-like) hypothesis to act on the environment must deal with a multitude of (i.e. noisy) signals from the senses. Unfortunately, these signals can only be found behind the veil of inference, or "Markov blanket"[87] (Clark 2013a: 184). It is an *inverse* problem, as the perceived sensory effects are used to infer concealed external sensory causes (think of Dennett's robot). In other words, the effects of sensory inputs are used to infer the causes of the sensory inputs, as opposed to passively receiving this information from its source. Bayesian inferential processes are ideally placed to explicate perceived sensory impressions from overdetermined (noisy) effects by appealing to the best inferred expectation (Burr 2017). These processes provide a way to overcome issues of multiple potential causes by assigning likelihoods to the occurrence of each cause. Furthermore, they could provide a neuro-computational solution to any disjunction

[85] Humanity "hallucinates at the world" in the words of Metzinger (2004: 52)
[86] For some preliminary examples it is worth reading Spratling's (2016) look at how this would work for the retina and in the cortex.
[87] The term is originally derived from Andrey Markov who pioneered the understanding of abstract systems which encode information in the form of data points. These values are used to derive a subsequent state within the system, thereby depicting a form of inference despite no actual memory being instantiated. Pearl (1988) introduced the concept of a Markov blanket to define the surrounding nodes within a network whose values can be used to infer a subsequent node's value by the property alone. The mathematical details are not pertinent in this treatment.

problems with regard to mental states[88] by simply appealing to the best possible inference, given the available sensory data, if we accept the theory of intentionality introduced in Chapter Four below. Briefly, the problem of fixing accuracy conditions in the first place arguably has it exactly backward, as probabilistic mental states infer only *broad* accuracy conditions, not *necessary* accuracy conditions. If one will recall the example of the duck-in-the-dark misrepresentation, our brains need not *actually* infer either the chicken or the duck-in-the-dark, but rather a cognitive imposition is placed upon what is perceived after the fact of perceiving. To be clearer, we only effectively "misrepresent" when we bring our cognitively-loaded intellect to bear on the case at hand. What is deemed an error at the probabilistic level is slowly weeded out over the course of development (whether we are speaking evolutionarily or ontogenetically). Bayesian inference within the brain is not about accuracy; it is about facilitating survival in the world. In Wiese and Metzinger's (2017: 6) terminology, Bayesian inference can be conceived of as a "probabilistic inverse mapping", whereby the myriad potential hidden causes are mapped onto a given sensory effect, and the most probable cause of this sensory effect is inferred from the mapping. This process leads to generative models[89] of the environment instantiated within the brain, which is central to PP (Clark 2016: 252-255). These generative models deal with generalised norms of mental state which are fixed over time (evolutionarily and ontogenetically). This form of development shall become clearer in Chapter Four; the next chapter introduces a potential causal history of mental states which can explain away the disjunction problem. Again, this helps resolve the disjunction problem by side-stepping the issue of stringent accuracy conditions, while still serving a richly perceptive cognition of the environment by utilising generative models which incorporate generalised norms of mental state.

Elaborating on generative models, Clark (2013a: 182) points out that an effective generative model of, for example, vision, would be one which attempts to account for the myriad ways "lower-level visual responses are generated by an interacting web of causes". In reality, this takes the form of top-down connective processing on multiple levels, resulting in a *hierarchical* processing system (Clark 2016: 19-21). The hierarchical system inferentially generates cognitive models of the environment, thereby delineating interactive causes from other environmental signals. Within this top-down hierarchical structure, models of the environment

[88] As a reminder, the disjunction problem brings to light the problem of distinguishing between causes that determine the content of mental states and causes that do not.
[89] These generative models are understood as a common distribution of a set of distributed variables.

are generated, predominantly utilising "higher-order" (prior) knowledge[90], which distinguishes this view of the mind from the passive cognitivist models (ibid.: 20).

The system therefore creates a "virtual" depiction of the incoming sensory signals through a cascade of multilevel processing. These depictions of top-down, hierarchical generative models, through these cascades of cortical processing, attempt to effectively predict the incoming sensory data, thus explaining this same data away (Clark 2013a: 182; Clark 2016: 37). What is meant by this is that if sensory data is successfully predicted by the model, it is not necessary to incorporate any alignments of the data with the predictive model (this would have no noticeable effect on the data). Only errors are propagated upward to inform the brain of its errors in prediction, thereby making any sensory matches with the model mostly redundant to the inner workings of the mind. This view has led some to reinterpret the brain's main task as approximately predicting the upstream (misaligned) sensory data rather than needing to reconstruct the actual sensory stimulation itself (Muckli 2010: 137; Rauss *et al*. 2011: 1249). This means that errors in prediction are "acquired" and then disseminated further up the hierarchical structure, all in order to influence future modelling of the incoming sensory signal for future tracking (Friston 2005, 2010; Friston & Stephan 2007; Hohwy *et al*. 2008; Lee & Mumford 2003; Rao & Ballard 1999, and many others). Furthermore, the precision estimates of this PP inhabit a functional role within the mind of the organism.

Effectively, the hierarchical predictive system is bidirectional in nature (Clark 2013a: 182-183). This form of prediction was first conceived of to make data compression more efficient in the processing of signals in computers (ibid.). For a basic example one can look at the processing of pixels within images: a single pixel can quite often be a predictor of the property of the pixels which surround it. Any difference in properties denotes a boundary between one pixel and that of another. Compression techniques therefore encode an image with only those properties that depart from predicted properties. Or, expressed differently, only a variation which is not "expected" is included in the image code, and it is only this variance between a predicted property and an actual property which is ultimately transferred[91]. This saves processing time and space when applied to processing models of this type[92]. It is just this form of processing which has been applied to neuronal PP to attempt to reconceptualise the way in

[90] Or "hyperpriors" if taken to the most general extreme (Hohwy *et al*. 2008: 691).
[91] This links with a minimal account of content (as an accurate portrayal of the environment) within basic perception, which shall preoccupy us in the next chapter.
[92] Prominent examples which have permeated our culture are lossless audio, motion-compressed video, and jpeg files.

which our minds work. Clark (2012, 2013a, 2013b, 2013c, 2015, 2016) has developed this further into a unified model of action and perception[93], whereby they are intimately linked to limiting errors in our predictive capacities through the "sculpting and selecting [of] sensory inputs" (Clark 2013a: 183). This particular aspect of his theory will be explored in more detail further below.

It should be noted that there is a vast difference between the base misalignment of incoming sensory data between environment and predicted model at a neuronal level; and agent-level, phenomenologically-accessible surprise. The literature distinguishes between these two highly interlinked phenomena, operating at differing levels within the brain. The terminology used for the implausible nature of incoming sensory signals is called "surprisal", which has been attributed to Tribus (1961). This is to distinguish it from the more molar-perspective sense of "surprise" that is familiar to us all. Prior models of the environment inform the brain as to the implausibility of given sensory stimulations, which then induces action on the part of the agent to reduce surprisal. This is achieved by altering the prior predictive model through highly efficient methods of "encoding" (Clark 2013a: 183). Importantly, the brain is not simply perceptually engaged whenever surprisal-inducing sensory input occurs, but under ecologically regular contexts there is always already the requisite priors active (ibid.). These priors influence cognitive processing of the sensory inputs immediately.

Interestingly, even illusions can be better explained with this interpretation of cognitive processing. One early prominent example has shown how illusions of motion could be the result of a "coherent computational strategy that is optimal under reasonable assumptions" rather than "sloppy computation" (Weiss *et al*. 2002: 603, see also Rescorla (2013)). The errors in visual perception, in this study, were found to be the natural result of (near) optimal Bayesian inference in light of available sensory data. Slower motions were more likely to be predicted, hence the illusion of slower object speeds as a natural consequence of the biasing effect of the percept (Weiss *et al*. 2002: 600-602). Hence, a recurring theme in the literature is that it is not about optimal prediction, but *near optimal* prediction of the uncertainties already available to our cognitive processes. Our cognitive processes are only able to predict utilising *a priori* information, which consequently biases us toward expected as opposed to accurate predictions of the environment. This does not bode well for Fodorian correctness conditions in our mental states.

[93] This was originally developed within the writings of Friston & Stephan (2007), Friston *et al*. (2009), Friston (2010), Brown *et al*. (2011), and others.

To reiterate, the "hidden causes" (Clark 2013a: 183) that impinge upon our senses can be explained away by utilising percepts. These percepts are utilised to form our subsequent prediction of that part of the environment which directly affects the sense modality in question, thereby, by extension, providing a prediction of the resultant effects in the particular modality. On the Bayesian predictive model, this is then contrasted with the *actual* sensory data uptake[94], and the discrepancy between the two (percept versus current sensory uptake), if any, is the prediction error which is fed upward so that the percept of the environment can be revised (Friston 2005: 821). The extent of the comparative discrepancy gives the mind an indication of the refinement that is needed of the predicted estimation of the environmental property under question. The subsequent revision of the discrepancy is known as *prediction error minimisation* (PEM) in the literature (Hohwy 2013: 41-55), with the idea being that our minds have evolved to seek minimisation of this error on a continually rolling basis. These fine-grained estimations of the environment are thus computationally incorporated into further predictions, juxtaposed with the continually-streaming modality signals, and thereby, all things being equal, prediction error is minimised[95] (ibid). Consequently, only once models of the environment have been suitably constructed from low-level sensory information[96] can the phenomenal character of our experience emerge. These phenomenal models emerge as a result of continual sensory input, informing our minds of the environment in which they are embedded (Clark 2013a: 183). Furthermore, these generative models of the environment are built from the past experiences of the individual. As individuals consistently take in sensory information from birth, it is not a case of "re-building" the environment with every conceivable experience, but rather altering via PEM the already-constructed world. Thus, Hohwy (2014: 4) states that priors come about not only through the immediate experience of the individual, but "through [past] experience, development and evolution"[97].

[94] Prediction error can be seen, in the words of Feldman & Friston (2010: 17), as a form of proxy for actual sensory data.

[95] Wiese and Metzinger (2017: 5) point out that it can occasionally be a *good* thing if prediction estimations are originally sub-optimal as to concealed causes. By continually updating the predictive models of our environment we can acquire a sense of "*confidence*" in these modelled representations of the hidden causes under question.

[96] This is predominantly in the form of information-as-covariance (Hutto 2017: 11), as shall be explained in the next chapter.

[97] This is also why Fristion *et al*.'s "dark room" problem (2012) – which asks why we wouldn't find a dark room somewhere and stay there to minimise prediction error (Mumford 1992; Sims 2017) – holds no weight. Prediction errors are bound up with prior predictions, which are in turn intimately linked with experience. We would expect to find a light source of some sort if we were truly faced with a dark room.

This black box[98] approach to cognition and perception (again, think of Dennett's robot above) leads to the seemingly solipsistic world that Hohwy (2012; 2013; 2017) endorses, forever wrapped in our Markov blankets. The lines of input/output that feed into and out of the brain must convey the most probable causes of impingement upon the senses despite this seemingly intractable barrier. All that the brain can discover, in any unequivocal sense, are its own fluctuations in its states, even if these have a rich phenomenal character (Clark 2013a: 183). Therefore, the brain only has direct access to these states, and nothing more. However, a way that the brain can link up with the world is through active motor directives and discover what new changes occur within these states as a result. On Clark's (ibid.) construal, this does not imply any "mapping relation" of these inner brain states to the environment. Rather, it implies an *inference* of the incoming data from the environment only on the basis of fluctuating brain states. The brain, as a hierarchical processor, uses top-down, generative environmental modelling to infer from its own internal states what the incoming sensory data entails. Or, as Clark (ibid.) interestingly puts it, an important task that the brain performs is that of inferring the "states of its own neuronal economy". The better the model of the environment, the better the inference, and the better the posterior probability of the current model.

To go into more detail, the way in which the proposed hierarchical predictive structure of the brain generates posterior models of the environment lies in the close interaction between the different levels of the cognitive structure. Priors at the "top" level can be utilised in a continually rolling guessing game to inform the priors "below". This takes the form of continuous "iterative estimation" (Dempster *et al*. 1977; Neal & Hinton 1998) which leads to a co-evolutionary interlinking of priors and the generated models of the environment, all in order to successfully predict the world (Clark 2013a: 183). A prior "higher" in the hierarchy places a constraint on the level "below", and these priors are constantly tuned by the incoming sensory data[99] (ibid.: 183-184). This attunement occurs at multiple levels within the hierarchy, while each level of priors attempts prediction of the level below through backward propagation (recurrent) connections. It is these backward connections which allow inference to occur, as fluctuations in the brain state are transferred higher up in the hierarchy as a new input. If there is no incongruity between the higher level priors and those below, then successful inference has been performed (Friston 2005: 817; Clark 2013a: 184). Only incongruities are fed

[98] This analogy refers to the opaque nature of the internal processes of the mind in juxtaposition to the (seemingly) clear inputs and outputs of the system.
[99] A prominent early example of this iteration of structural learning is that of Rao & Ballard's (1999) use of predictive coding within the visual cortex.

backward to a higher level to "inform" the priors of the error in prediction. Probabilistic models of the environment at higher levels are therefore updated continually on this basis, all in order to "explain away" any errors in prediction at lower levels (Friston 2005: 817-819).

It is thus only the backward connections that do the hard work, propagating the predictions to the higher levels, with the forward connections passing on only residual errors in prediction (Rao & Ballard 1999: 79; Clark 2013a: 183; 2016). The probabilistic model of the environment is thus updated, *predominantly* via backward propagation. The external environment is subsequently revealed to the brain via the gradual reducing of errors in prediction, which causes processional cascades within the brain's hierarchical structure[100]. The posterior probability of specific external environmental features are thus increased, and depicted as an internalised model (of some sort) within the brain (Clark 2013a: 184). It is this that allows us to infer what is behind the Markov blanket. Hence, we have a complete reversal of the usual paradigmatic view of the brain as a passive processor that incorporates external sensory impingements, processes this data, and then represents it to the mind in some fashion. Hohwy (2007: 320) stresses that the "top-down" predictive model is doing most of what we take to be "perceptual". The "bottom-up" informational signals provide more of a consistent recurrent feedback for the downward cascade of predictions, which still allows for heightened sensitivity toward environmental contexts. Gradually, the predictive model may be altered in this way, which has been depicted as a form of "reciprocal interaction between perception and learning" (Fletcher & Frith 2009: 53).

After all is said and done, the percepts that are provided from the continually updated perceptual Bayesian modelling of the environment can therefore be used to pierce the veil, thereby breaking through the solipsistic barrier after a fashion. However, in a monograph on PP, Hohwy argues that we are nothing more than "mental islands set over against the world, which is hidden behind the veil of sensory input" (2013: 258). We cannot delve behind the "evidentiary boundary" (Hohwy 2013, 2014). This solipsistic view of PP is perhaps a bit harsh, and Clark (2017) has argued that this seemingly intractable barrier is anything but. Clark points out that we are consistently re-configuring our mental boundaries in order to maintain successful adaptation with regards to the environment. In other words, he proposes a deep coupling of brain and world, which is in line with his view that our minds utilise the

[100] These depictions of cognition have been called "forest first, trees second" types (Friston 2005: 825; originally Hochstein & Ahissar 2002: 791). The general gist of the environment is depicted, with details following in its wake.

environment in order to function optimally[101]. In other words, our minds do construct our environment, but this ability is only afforded by very real attunement to sensory information. If our brains dealt in rich representations of the outer environment that stand in for the world, then it would be difficult to argue against charges of strong solipsism, but we need not view our inner neuronal weightings as directly mirroring the world. Hence, there is a sense in which our minds can predict the properties of objects within our environment *accurately enough* in that we are near optimal in our predictions.

The utilisation of these properties do not rely on truth relations of the sort that Fodor would endorse as necessary for accurate perception (as encountered in Chapter Two). Indeed, as Clark (2015a, 2015c, 2016, 2017) would agree, I propose that this tentative connection with the external properties of the world depicts brain and world as *more* closely entwined than traditional cognitivist theories. This view lies within the enactivist tradition. Enactivism, which is to be explored in more detail in the next chapter, is a view that depicts cognition as a result of dynamic processes of interaction between agents and their environment, blurring the line between perception and action more so than traditional cognitivist theories. As the section below will show, this opens up space for PP to be the underlying cognitive processing which enables the types of features that enactivism endorses. Note, however, that this does not necessarily speak to relations of accuracy and representation with regard to the environment, a view which Clark (2016: 288-291) shares to a certain extent.

2.2 Active Inference

This dynamic coupling between organism and environment has another dimension. One must keep in mind that the above is an explanation of a quite fine-grained, continually updating predictive "machine". The brain needs to formulate predictions in all sensory modalities at every moment of existence, dynamically updating its predictive models of the environment in order to account for the ever-changing world[102]. But, PEM must be seen primarily as a way to facilitate organismic survival by forming the basis upon which the agent may then subsequently *act* upon the world (Friston 2003; 2010; Friston *et al*. 2009; Friston *et al*. 2011). The internal environment of the body is also of high importance for the agent, and homeostasis[103] should be

[101] This is his Extended Mind Hypothesis (Clark & Chalmers 1998), which shall be explored in more detail in Chapter Four.
[102] This would include the agent's influence upon the sensory uptake, adding more layers to an already seemingly intractable computation.
[103] Homeostasis refers to the tendency toward equilibrium that physiological and biological systems exhibit.

considered as the primary concern of the agent/organism, as stability provides many affordances[104] within differing environments, while instability leads to low survival rates within even the most congenial of environments (Seth 2014b: 270-271; 2015: 9). Affordances have a close connection with an organism's *Umwelt*, which is the experienced world of the organism. The deep connection between the organism's body, environment and action is a key component of the PP paradigm, as perception of prediction errors in interoceptive, exteroceptive, and proprioceptive states lead to resultant action in the service of PEM. PP implies that the internalised mental states that serve to minimise prediction error in perception can be mapped, or overlap, with the mental states that serve action (Friston *et al*. 2010: 233).

In other words, this implication effectively binds perception to action in terms of not only the information processed, but in the very real sense that they are mutually supported by the *same* mental states. Those steeped in the literature will recognise this as a clear example of the ideomotor principle, which has been developed into the general proposal that neuronal states of causes derived from the environment map onto the neuronal states that enact preparation of action (Stock & Stock 2004: 176). In the language of PP, percepts and motor directives can function jointly as the same ideomotor mental state. This type of mental state is similar to the phenomenon which is closely tied to it: "mirror neurons". These "mirror neurons" are areas in the brain that have been found to fire not only when an action is performed, but when a suitably similar action is perceived when observing other organisms[105]. It is, in other words, an automatic and reflexive action.

Through PP, the ideomotor principle finds full expression in "active inference" as developed by Friston[106], which is neatly encapsulated in the following:

> Under active inference, there are no distinct sensory or motor representations, because proprioceptive predictions are sufficient to furnish motor control signals. […] They encode conditional expectations about hidden states in the world causing sensory data, while at the same time causing those states vicariously through action (Friston *et al*. 2011: 138).

Therefore, "representations", which entail both sensory and motor functions, serve a dual role in both inferring hidden causes within the environment, as well as causing the action to bring

[104] Recall, from Chapter One, that this term was originally coined by Gibson (1979) to refer to the positive and negative aspects of an organism's environment that matter to it.
[105] The literature concerning mirror neurons has somewhat overdetermined the range and extent of their impact on the mind, at least according to Sapolsky (2017: 535-542) and Hickok (2014).
[106] For example: Friston *et al*. (2009) and Friston *et al*. (2011).

these hidden causes to light. This results in a "circular causality" that "destroys conventional distinctions between sensory (consequence) and motor (cause) representations" (ibid.). Motor acts and perception are therefore intertwined, with the mental "representations" serving the dual role of being "both intentional and perceptual" (ibid.: 156). Motoric actions can be the basis from which "intentional" behaviour occurs through direct sampling of the environment caused by *direct* proprioceptive perception. Of interest to this thesis, this does not indicate that active inference appeals to "desired consequences" (ibid.: 157), but rather to prior learning and context-dependent inference. Any perceptual experience summarily invokes the percepts which, via active inference, conduct the agent/organism within and/or through its environment[107].

Active inference, moreover, can be elucidated as the computations that jointly underlie action and perception in order to minimise, at base, free energy (see below)[108] which in turn is all in the service of PEM. These computations are the Bayesian-optimal inferences that unify perception and action, thereby updating the ideomotor principle to a more nuanced understanding of fine-grained mental states. What this means is that active inference, so conceived, works toward equilibrium by bringing about changes in sensory states by directly acting upon the environment (Friston *et al*. 2012: 539). As our bodies are more at risk from unexpected and potentially harmful interoceptive states than unexpected and surprisal-inducing exteroceptive states, PP and active inference are a natural fit when explaining the homeostatic equilibrium of the organism. Seth (2015: 10) uses the example of declining sugar levels in the blood as a clear example of this. The percept of a yearning for sugar would result in prediction errors if there is a detection of deterioration of blood sugar levels. This leads to the body actively metabolising fat stores to minimise prediction error or through the macro-action of actively seeking and consuming more sugar.

This is why Clark (2013a: 185-186) has expanded his conception of PP to what he calls "action-oriented predictive processing" to drive the point home that this construal of PP (as a predictive coding variant on perception) encompasses action. It places action as predominant in these depictions of perception, as it is only through action that prediction error can be minimised (Friston 2009: 295), assuming a cogent environmental model exists and has been implemented.

[107] This also ties in with the idea of minimal content in the next chapter, if the percepts are minimal content models of the environment.

[108] This is enacted by a change in the internalised mental states of perception and those of sensory states by subsequent action on the world (Friston *et al*. 2012: 524). Take note that Friston *et al*. use active inference as a blanket term to refer to the intertwined inference of both action and perception of the agent. One could still make the case for separating the two.

On this understanding, action and perception are both instantiated using the very same processing capacities[109]. If perception is primarily concerned with matching incoming sensory signals with prior models of the environment in order to create the aforementioned posterior probability models, action utilises a similar strategy. The systems involved with errors in motor commands "self-suppress" (Friston 2003: 1349), not through neuronal facilitation, but through the movements that can change sensory signals. If your mind locks onto an environmental pattern, it causes a predictive cascade of what *should* be experienced as a result. As this predictive cascade progresses, motor action is provoked in order to successfully realise the aforesaid prediction (ibid.). This motor action could be very subtle, such as the above-mentioned maintenance of blood sugar levels, or movement of the entire body. In the literature, this means that action, in the form of control of motor commands, is mathematically the same as the Bayesian inference explained above[110] (Friston 2003, 2005, 2010; Friston *et al*. 2010; Friston et. al. 2011).

Action-oriented PP depicts motor commands as actively causing the continual sensory streams that our brains are in the business of predicting, leading to active inference (Clark 2013a: 186). Therefore, we are left with a view of action that depicts behavioural predictions as not only preceding sensation, but *determining* sensation. During active inference, the individual uses and/or moves their sensor arrays in order to seek, and as a result actively create, the sensory predictions that the brain expects in the first place (Friston 2009; Friston *et al*. 2010; Friston *et al*. 2011). Examples are the moving of eyes, or arms to manipulate the environment, in order to confirm an expected sensation. In the case of social interactions, we actively sample people's bodily language (gaze, expression and others) to confirm predictions. Action, therefore, fulfils our learnt expectations, depicting action and perception as intimately bound (Friston *et al*. 2009: 12). If active inference does not succeed in fulfilling a prediction, a cascade of prediction errors ensues, until the prediction is aligned with the individual's reality through further active inference (Friston 2010: 134). If this is not possible, the predictive model of the environment needs to be altered to fit the particular context. Again, this is closely tied to enactivism and its conception of direct coupling between body and environment, and may even inform the

[109] As Clark is constantly at pains to point out, this is the only game in town which could finally unify perception, action, and cognition (Clark 2013a; 2013b; 2016 and others).

[110] It is especially prevalent in "optimal feedback control theory" (Todorov 2009; Todorov & Jordan 2002), where a goal state is given, and Bayesian inference is utilised in order to discern the actions which will successfully achieve this goal. This is also prevalent and analogous to some literature on planning (Toussaint 2009), whereby a goal is given, and in order to plan the best course of action Bayesian inference is again used to infer the different states which will prove successful.

amendment to teleosemantic theories of content that shall be introduced in the subsequent chapter.

2.3 The Free Energy Principle

This is where an interesting element develops in the literature, intimated earlier, concerning the minimisation of free energy. This will be a recurring theme within the rest of this thesis as it is this very minimisation of free energy which I propose leads to minimal content in our cognition. This in turn shall inform my view of minimal instantiations of propositional attitudes in our folk psychological interactions. As described earlier, the problem for the agent/organism lies in the parsing of incomplete and noisy data into useable, *actionable* data. The agent/organism is not in a position whereby it has direct access to the probability distribution[111] that results from surprisal of the data within its environment.

The link between this fine-grained probability distribution and phenomenal perception can be found within the free-energy principle (FEP) developed by Friston and his colleagues (Friston 2003, 2008; Friston *et al*. 2006; Friston *et al*. 2008; Friston *et al.* 2009; Friston *et al*. 2010: 229-233, Friston *et al*. 2011: 139-144; Friston & Stephan 2007). The mapping of the expected outcome (i.e. the internal probability distribution) to the posterior (which has incorporated the hidden causes in the environment) needs to occur. But what also needs to occur is the changing of the sensory signals in order to facilitate lower surprise/surprisal *as efficiently as possible* (Wiese & Metzinger 2017: 12). This may appear, from the outset, to provide two sets of intractable problems: the mapping of the expected outcome to the posterior, as well as the minimisation of the surprise/surprisal of the sensory data (i.e. minimisation of prediction error). The FEP, as its name suggests, is concerned with the free energy within any entropic system, and in particular the *minimisation* of this free energy. This is *not* free *thermodynamic* energy, which is the energy available to the organism for successful navigation and action within the environment. In Friston's formulation, free energy on this construal is an information-theoretic quantity[112] that can be minimised in some fashion within the brain (Friston 2008, 2009, 2010)[113]. Information-theoretic quantities on this construal are simply the strength of the neuronal weightings that underlie the conveyance of pertinent data to the predictive

[111] This simply refers to the predictive model, which itself is a distributed model that infers the probability of a state of affairs.
[112] Informational theory deals with statistical analyses which trade noisy information for more refined statistical certainty.
[113] To do the topic justice, a more thorough explication of the FEP can be found within Bogacz (2015).

environmental models in the mind. This minimisation of free energy underlies *both* of the aforementioned operational problems concerning the mapping of the expected outcome to the posterior as well as PEM.

Therefore, we find a picture taking shape of the human organism as confined to its senses and actions but still able to find a way to provide as accurate a sensory prediction of the environment as possible. This is in order to satisfy the minimisation of free energy constraint, which the human body, in an attempt to survive in an entropic world, is in a constant battle to sustain. As Clark points out, the FEP is interesting even in isolation from PP, and could open the door to understanding "self-organisation within biological systems" (2013a: 187). As Friston (2010) has persuasively argued, the PEM theory can be depicted as a macro-consequence (in relation to the FEP) of the fundamental need in organisms to minimise free energy while existing in their environment[114]. Depicting free energy within the cognitive realm casts it as the difference between the way the world is and how it is "represented" (in an information-theoretic sense) within the mind (Clark 2013a: 186). In other words, in the information-theoretic sense, free energy is analogous to the energy needed to minimise prediction error. The better that an organism is able to predict the environment, the less thermodynamic energy is needed for its subsequent actions within said environment. Prediction error is effectively a stratagem to "report" the (information-theoretic) free energy (ibid.). Entropy can be depicted as "the long-term average of surprisal" (ibid.), in that the surprisal induced over time amounts to the entropy of the system, and any reduction of the free energy within the organism's cognitive modelling of the environment would result in an improvement of the environment's model. This leads to a reduction in prediction error, and therefore surprisal.

As Friston (2010: 133) has depicted it, the better the model of the environment, the better the organism can sustain its organisational structure through the minimisation of entropy. This would only be possible if the organism is able to resist an entropic increase within its neuro-cognitive system. Therefore, there would need to be a resistance of the second law of thermodynamics[115] over finite lengths of time through improvements of sensory modelling of the environment[116] (ibid.: 127-128). The minimisation of entropy, therefore, leads to a

[114] Or, to be exact, "an information-theoretic isomorph of thermodynamic free-energy" (Clark 2013a: 186).
[115] The simplest way to put this would be to say that entropy always increases within any closed system. Our bodies are open systems which exchange energy with the environment, which means that evolution has furnished us with the capability to mitigate immediate decay.
[116] Stated differently, the average free energy over a period of time is closely linked with adaptive fitness, or "free fitness" as it is known within evolutionary biology (Sella & Hirsh 2005).

minimisation of information-theoretic free energy. As a result, our brains are improved as predictors of sensory impingements, and therefore our chances at surviving within our environment are improved at the same time[117]. Therefore, the better we are at modelling our environment, the better chance we have at navigating it appropriately. If, as an abstract example, a predicted model of the environment elicits a response that accurately aligns with a threat, we are more than likely to avoid that threat than if there is significant surprisal between the model and the environment.

In line with the rest of this thesis, Friston (2009; 2010) has utilised the FEP to unite perception, action, and inference under a single stratagem of cognitive processing, as mentioned above. As an example, and of great interest to this thesis, I propose that we can actively strive to construct niches within our environment in order to minimise free energy and, by extension, prediction error. This can encompass the general social strategies we utilise in daily life, without a second thought, such as our use of belief/desire psychology. Under the FEP (if it is correct and it does unite perception and action in just such a way as to minimise free energy) human social interaction could be aimed at minimising free energy on the macro-scale. I propose that FP, as currently practiced, could be an unwittingly derived evolutionary adaptation toward minimisation of free energy on a bigger scale[118]. Propositional attitudes, if nothing else, pick out macro-reasons for behaviour in order to allow us to better model our environment for successful navigation thereof. Note that this does not entail an *accurate* portrayal of reality, as such, but merely those aspects of reality that *matter to the organism itself*. Socially constructed propositional attitudes within the context of humanity's folk psychological paradigm can be seen as just such a successful strategy for survival (in a broad sense) within the social environment. Also note that any future change in our FP could still adhere to the principles of energy minimisation, as this change could be a gradual shift in our understanding, rather than a wholesale and wasteful change[119].

[117] It should be noted that the FEP, according to Friston (2010), extends to all systematic structural organisation, from the entire morphological structure of the organism right on down to specific processes within the organism itself. Therefore, cognitive data processing is simply one element within this overarching principle.
[118] Hirsh *et al*. (2013) made a similar claim when they wrote that "personal narratives" are the molar-level integration of the underlying predictive processes of the mind (ibid.: 216-217).
[119] For a list of both the accepted and still contentious characteristics of PP, see Wiese and Metzinger (2017: 1).

3. Predictive Processing and Folk Psychology

But what room does the PP conception of mind leave for propositional attitudes and their interrelations? The two prominent philosophical adherents of PP have made mention of the negative impact this would have on our self-conception (Clark 2013a; 2016; Hohwy 2013: 2). Hohwy begins his monograph on PP with a capsule statement highlighting the fundamental changes that this view would afford us, and part of that radical shift is our very conception of "who we are" (Hohwy 2013: 2). Clark (2013a: 197) has similarly expressed the impact predictive modelling would have on "agent-level experience", potentially giving us a more coherent picture of our perception than "the basic framework of 'folk psychology'" (Clark 2016: 82). He also states that this would be a vindication of Churchland's neuro-computational elucidation of the inner workings of the mind (which the reader will remember was discussed in Chapter Two) (Clark 2013a: 197; 2016: 82). What Clark and Hohwy are alluding to in these statements are the more traditional iterations of FP as propositional attitude folk psychology. As I have expressed in the previous chapter, this is the predominant view of FP within philosophical literature on the topic, whether tacit or not. Therefore, looking at the impact that PP would have on our canonical depictions of beliefs and desires would be illuminating on this score.

3.1. The Impact of Predictive Processing on Canonical Beliefs

Firstly, it must be remembered that belief, as conventionally understood, is a propositional mental state linked to an attitude that bears an epistemic relation to the proposition where the proposition is regarded as true[120] (Dewhurst 2017: 4). This leads to interaction with other mental states, which create subsequent actions so as to act in accordance with this state of affairs (that of the proposition being true). This blanket statement concerning belief can still have a place within the PP conception of mind, as predictions generated by cortical processing could be interpreted as "beliefs" about the state of affairs in the world (ibid.). However, on the PP view, belief would be produced through the prior model, and subsequently juxtaposed to the world in order to produce the posterior probability model.

Friston has described beliefs as predictions in many of his works (Friston 2010; Hobson & Friston 2014), and so have Clark (2013a; 2016) and Hohwy (2013). While this may be an

[120] The mental state of belief is usually expressed in the typical formulation of a propositional attitude, such as "D believes that he has gone over this before".

innocent usage of the term "belief" within the PP framework, it is useful to separate the way "belief" is used within FP as traditionally conceived and the way it is used within the PP framework. As Dewhurst (2017: 4) points out, beliefs within traditional FP are usually treated as determinate mental states in the sense that either one believes what the mental state represents or one does not (as explained in Chapter Two). On the other hand, belief on the PP construal is a predominantly probabilistic disposition. For example, to believe that one is stuck in traffic or at a traffic light (from Bach & Dolan 2012: 573), the PP cognitive system would consign a probabilistic quantity to this state of affairs, which in turn provides impetus to act in the form of moving forward or not (although *when* the impetus occurs is of course problematised on this view). The example that Clark (2013a: 188; 2016: 41) uses is that the brain does not represent "CAT ON MAT", but instead a "conditional probability density function" that is indicative of the external reality (or at least indicative enough), taking what is given to the system.

It is clear from these examples that even writing, let alone thinking, about beliefs within PP forces upon us a radical shift in our conception of the interrelations of propositional attitudes and the mind. We can conceive of beliefs as probabilistic, similarly to Pettigrew (2015) who juxtaposes determinate belief with probabilistic belief and finds the former wanting. He sees a pluralistic take on beliefs to be more useful in terms of explicating the ways in which beliefs can be described on differing levels of analysis (ibid.: 201-203). We do use determinate beliefs in explaining behaviour, but these beliefs become increasingly indeterminate the further "down" the cognitive processes involved we go. Again, assigning an emblematic belief stating that one is stuck in traffic or at a traffic light as either being true or false is different to conceiving of predictive beliefs that contain finer details of the environment, those "such as edges and light gradients rather than the 'middle sized dry goods' that populate the folk ontology" (Dewhurst 2017: 4). Even the more determinate (but still importantly probabilistic) instantiations of belief within the PP system involve all the modalities, as well as somatic and emotional affectations, as described in Chapter Two with reference to beliefs and desires. These intermingle in ways that irreparably warp the typical conception of belief toward a strangely probabilistic distribution within the mind. Clark's conception of PP entails complexities that loop into mind and environment, thus rendering us incapable of adequately delineating canonical beliefs utilising the common conceptualisations within everyday communication (Clark 2016: 292).

In other words, the differing roles of our internal cognitive states do not correlate with typical conceptions of belief. We may express in language the notion that we are entertaining a belief, but our language is merely a way of fine-tuning, and hence manipulating in an artificial way, the precision gain on our predicted models of the world (Clark 2016: 284). Having said this, Clark (2013a: 199) argues that we must not discount the "causal potency of the folk-psychological constructs". These constructs are experienced and manipulated "just as surely as we encounter and model other constructs such as marriage, divorce, and taxes" (ibid.). Hohwy (2013: 60-61) similarly views beliefs as probabilistic, but also states that the interactions between the PP mind and the dynamic environment result in eventual regularities. These regularities, as more determinate mental states, are cognitively higher-order than the aforementioned probabilistic density functions and akin to our canonical conceptions of belief. Hohwy uses the example of a moving cat partially occluded by a picket fence to make his point (ibid.). Dewhurst (2017: 4-5) takes Hohwy to mean that the propositional belief "there is a cat moving behind a picket fence" is cashed out by both a coarse-grained prediction of the whole cat, as well as a finer-grained prediction of the parts of the cat seen through the fence over a period of time. The PP models therefore instantiate both the "diachronic" (ibid.: 5) slices of temporally immediate cat slices, as well as the higher-order "belief" of the existence of a whole cat. If this higher-order "belief" is suitably pertinent depending on the context (in this case: needing to have a conception of what is moving behind that fence), it is held at the conscious level. In other words, we should expect mental states which broadly resemble typical folk psychological mental states at higher levels of the predictive hierarchy, but these are still fundamentally probabilistic and too coarse to be determinate in the propositional sense (ibid.). I am partial to this view, as it encapsulates both the probabilistic and abstractly non-linguistic conception of belief, while allowing a seemingly determinate state of belief.

Predictions, through active inference, enable us to act upon our beliefs. In other words, the predictive models provide visceral, "off-line" norms for action. Clark (2016: 187) has elaborated that this is very similar to the Millikanesque representations of the Pushmi-Pullyu kind (Millikan 1995, 2005). These concepts and the lingering problems thereof shall be expanded upon in the next chapter, but briefly: Clark envisages predictive models as instantiating a form of content that is dually "descriptive as well as imperative" (Clark 2016: 187). This notion of representational content informs his view that the PP paradigm enables us to both model our environment as well as act upon it. The astute reader, as attested to by Dewhurst (2017: 5), will recognise that the line between beliefs and desires begin to blur. The

desire to act upon the world seems to be subsumed into the belief that the world is the way it is. This is why Hohwy has suggested a reconceptualisation of our concepts of "belief" and "perception" to something more akin to an all-encompassing "expectation" (Hohwy 2013: 70-73). Beliefs and desires, in other words, are done away with, and coalesce into one another as a singular cognitive state.

On this view of cognitive dynamics, it is difficult to place the traditional understanding of belief in relation to the fine-grained processing of PP. Ontologically-speaking, "belief" only exists at a far more molar level of observation. Yet, using the term "belief" to describe predictions does not necessarily delineate the typical propositional attitudes that inhabit our talk of FP. The determinate and linguistic nature of FP beliefs are starkly different to the probabilistic and fine-grained predictive models instantiated within the mind on this view (Dewhurst 2017: 5). At the very least, these predictive models are not propositional in any real sense. As alluded to above, these models may appear to coalesce, during higher-order processing, into something akin to our publicly-accepted notion of folk psychological belief, but this has a distinctly "as if" quality to it. As Chapter Four's depiction of the effect enculturation has on our beliefs and desires will show, we learn to attribute these beliefs and desires in a particular way through the effect of our ontogenetic development.

3.2. The Impact of Predictive Processing on Canonical Desires

Desires are similarly affected by the probabilistic framework of PP. Desires, as stated earlier, can be subsumed, along with beliefs, into a general expectational state. The action-oriented predictive qualities of this state do not simply model the environment, but stimulate the cognitive system to act within it through the motivational tool of active inference. This capability, that of stimulating the system into action, takes the place usually reserved for "desires" within our common understanding of FP (Dewhurst 2017: 5). A "representation" of what the predictive system depicts the environment to be like is generated, analogous to the "beliefs" of how it actually is, and subsequently action is entrained to bring these "desires" to fruition. The feeling of "desire" therefore features as a by-product of these sub-conscious processes, as the active minimisation of surprisal brings the body in line with the inner "representation" of what is expected[121]. It is an open question as to whether this phenomenal

[121] This is an all-pervasive phenomenon on this neurological framework. Everything from the surprisal induced by low glucose levels (instigating the chain of events that eventually lead you to eat) to the alignment of visual

by-product is either selected for via natural selection, or an arbitrary relic of the evolutionary past.

But this is a quite radical reversal of our conception of desire and action as seen through a traditionally folk psychological lens. Causation of action is reversed from a desire for a state of affairs to come about leading to behaviour, to an expected state of affairs being predictively generated, which then causes behaviour (via active inference). It appears, if we are to take the PP account seriously, that the phenomenological experience of desire is no more than the emotional and somatic instantiation of the predictive modelling process of a state of affairs in the world (ibid.: 6). Desire, as typically conceived, does not in any real sense inhabit a functional position which is causal in any manner. Said differently, active inference within our predictive minds "does not invoke any 'desired consequences'" (Friston *et al*. 2011: 157).

Desire, therefore, does not *lead to* behaviour in any obvious sense. Clark (2016: 129) endorses this view by reconceptualising desires (along with beliefs) as complex interrelations between predicted models and the world. Cognition and perception, as traditionally construed, become entwined in such a way as to become "fuzzy" (Clark 2013a: 190). Perception and desires have no clearly distinctive properties, and are rather thought to be variations within the cascading cortical processes, further influenced by the spatial and temporal differences depicted within the models themselves. Cognitive, more "overt" processing reacts to conceptualisations of the environment that are more abstract (and potentially culturally mediated). By contrast, perceptual, low-level models react to finer-grained levels of detail in the environment. It is the "precision-modulated […] interactions between these levels" (ibid.) which are the driving force of adaptation to the environment and intelligent behaviour. Thus, the world is perceived by this use of information, which is used to "explain away" driving sensory signals. Cognitive desires and perception are distinct only insofar as they are taken to be conceptual, rather than actual differing processes within the mind. The same internalised processes are utilised in their instantiation, and they are multiply-realised within the neuronal substrate (ibid.).

This is all very fine and well, but there is an existential issue with the aforementioned. This view comes at the cost of the conception of our personal agency and volition. Desires as consequences, as opposed to being the original instigators of behaviour, is a profoundly counterintuitive depiction of our experience. Here, active inference (depicted here as agency) is purely in the service of the PEM of the Bayesian system. It is not about *wanting* to do things,

sensory uptake with environmental expectations (the expectation that a child is cute leading to the confirmation thereof), can fall within the remit of PP.

instead: "…we *feel* like we want to do things *because* doing them will minimise prediction error" (Dewhurst 2017: 6). Hohwy (2013: 89) similarly states that it is not what the agent wills, so much as the process of PEM that instigates action. This would therefore not place our canonical folk conception of desire in significant interactions with action generating processes. The general predictions that instantiate action subsume folk psychological desires into themselves, and active inferences ensure that action is taken to make the predictions come to fruition.

4. Issues Going Forward

In summary, beliefs and desires of the traditionally folk psychological kind become something decidedly different on the PP picture. Firstly, both our beliefs and our desires are subsumed into a singular predictive state, which is action-oriented. Secondly, and more disconcerting in terms of our sense of personal agency, desire is not primary in our cognitive ontology, rather it is PEM which acts as the fundamental impetus for action. This is all very upsetting for those that insist on a FP predicated on semantic propositions as mental content. A move that these proponents can perform to undercut the impact of these theoretical impositions is to cede ground to the adherents of PP. Even if the PP picture were correct, the above does not entail a full picture of the way that we utilise folk psychology within our everyday interaction. What about the typical mentalistic states that can be attributed to ourselves and others? Are these theorists not performing some form of obfuscation, taking something away that exists in some obvious sense and has been useful in predicting mental states? Not exactly, since, as can be recalled from the previous chapter, contemporary theories may fail to vindicate folk psychological propositional attitudes in their canonical guise, but such propositional attitudes can be subsumed into a broader framework of learnt narrative competency. If these learnt narratives are what fix the propositional attitudes in behavioural explanations, we should be looking toward culturally-mediated influences on our folk psychological discourse, if we are to evaluate it. Arguably, our "beliefs" only appear determinate and linguistic because we have been taught to think of them as such.

The reader may also have picked up on elements of the disjunction problem within the Helmholtzian idea of multiple sensory causes potentially bringing about a particular effect, as discussed above. This issue, too, can be addressed. On the PP view, the predictive mind goes some way to explaining away the need for accuracy conditions, and instead advocates a mind

which collates *probability* over *accuracy* by "analysing" a gradient of likely sensory causes from a fixed history of generalised norm fixing. The existence of multiple causes potentially bringing about a particular effect are assigned weightings relative to each other, which subsequently attunes the mind to the most probabilistically pertinent cause. Again, this need not be accurate as such, but "accurate" *enough* for appropriate bodily response. There is nothing obviously contradictory about this beyond intuitive cognitivist bias, which means that the disjunction problem need not arise on this view.

The view of cognition that PP affords us could be an important component in answering the question of how it is possible that we acquire the information that fuels our perception and drives our action in the first place. In other words, how is it that our accuracy conditions are fixed such that our mental states can misrepresent? It may go some way to answering how the "spike trains"[122] within the brain are capable of "representing" knowledge as stored content for representation of the environment. It appears that it is only these "spike trains" that are available to the brain moment to moment after all. A way to alleviate these worries is to "go radical" with regards to representation.

To "go radical", as Hutto and Myin would put it (Hutto & Myin 2013a, 2017) in terms of representation in the mind, is to be eliminativist toward content. In this way, a view of the mind that can incorporate the positives of PP while doing away with the problematic elements of traditional FP, and which posits rich semantic mental content as foundational to our cognition, can be achieved. This account will be unpacked and expanded upon in the next chapter. Going radical with regards to representation in the mind is necessary as, although Hutto and Myin (2017: 57-66) find much to admire about PP, they balk at its insistence on representation as the hallmark of the mental. They see Clark's (2016) recent take on PP to be a cogent account which unifies the disparate phenomena of the mind into a coherent conceptual framework, partly because of its propositional attitude eliminativism (Hutto & Myin 2017: 58). But they have distinct misgivings with the lingering representational element of the PP paradigm, arguing for a minimal construal of content within the neuronal economy of the mind. This is due to the need to eliminate unnecessary processing in the brain, a view that is also adopted in this thesis. The admiration is similarly limited from Clark's side, as he considers it impossible to tell the tale of PP "in entirely *non*-representational terms" (Clark 2015c: 5, emphasis added). Clark does, however, have many affinities for the enactivist tradition with his previously formulated

[122] Action potential is the change in electrical potential of the axons within the brain, and spike trains are a temporal sequence of such action potentials.

Extended Mind Hypothesis (EMH) (Clark & Chalmers 1998), but this will be briefly expanded upon in the chapter to come. I will argue that a reconciliation of these related, yet marginally opposed viewpoints is needed in order to provide a more accurate theory of the inner workings of our cognitive processes, and by extension, of their effect on our folk psychological discourse.

In relation to the above, the gulf in understanding between our folk conception of ourselves and the scientific conception of the same are still poles apart. Sellars (1962) has termed these two spheres of discourse the "scientific" and "manifest" images. The manifest image denotes the day-to-day folk psychological understanding that we all employ[123], whereas the scientific image denotes the physical realm to which we have no direct cognitive access, such as atoms, molecules, the interaction of such, and everything in-between (such as the neuronal structures and their sub-personal processing currently under review). As Dennett has remarked (2013a: 209; 2017: 61-63), the role of the philosopher is to close the gap between these two images or, stated differently, to bridge the gap between surprise and surprisal[124]. This is why Dennett (2013a: 209-210) endorses the predictive processing paradigm, and goes on to draw an analogy between it and our folk predictions. He begins with what he has termed "Hume's Strange Inversion" (Dennett 2013a; 2013b; 2015; 2017): there seems to be something in the world known as causation, but we only see correlations, from which we *infer* causation, and these very correlations within the environment cause us to "see" the causation "out there" (Dennett 2013a: 210). For example, throwing a ball and seeing it cause the shattering of the window-glass elicits a sensation of perceiving the causation as out there in the environment (ibid.). But as Hume has long ago argued, this is actually an example of the mind's "great propensity to spread itself on external objects" (2009: 270). Instead, we are misrepresenting an inner sensation or, more correctly, an *anticipation*, for an external property of the world. In other words, we are misattributing internalised mental perception to particular occurrences in the environment, as a form of "user-illusion" (Dennett 2013a: 210). Hume (2009: 271) stated that we cannot seem to do away with this illusory sense of the reality of causation. Dennett (2013a: 210) states that this is still the contemporary case, as the cognitivist viewpoint is one in which perceiving representations of the world is primarily a case of receiving incoming sensory data from the outside[125].

[123] Or, as Dennett (2013a: 209) puts it: "[Our] everyday world of folk psychology, furnished with people and their experiences of all the middle-sized things that matter".
[124] Clark (2013a: 199) and Hirsh *et al*. (2013) also allude to Sellars with reference to PP.
[125] Here are a few other examples that he states are in need of "Strange Inversions": the intrinsic sweetness of sugary products; the intrinsic cuteness of babies; intrinsic sexiness of objects of lust; the funniness "in" the joke, and many others (Dennett 2013a: 210; 2015; originally Hurley *et al*. 2011).

Dennett proposes that to tie the manifest and scientific images together, we need to look to the predictive coding strategy (Dennett 2013a: 210). All organisms, from the single-celled up to humanity, have affordances that matter to their particular well-being, and which are pertinent to predictively discern (ibid.). Humans have a set of affordances that are vastly different from the simple homeostatic inclinations of a single-celled organism, including avoiding cars in the road, being able to open doors, and even being able to communicate socially. The Bayesian model sees the mind as deploying its resources toward doing as good a job as possible at "representing" the environmental features around us that matter to our future behavioural action (ibid.). We expect to not survive the impact of the car; we expect the door to open as we turn the handle; we expect certain social conventions to hold within social interaction. These expectations extend to our very selves, as predicting our behaviour is a good implicit survival technique when we are able to predict what we will think, do, and expect next (Clark 2013b: 240; Dennett 2013a: 210). As in the example of social conventions above, we expect the social conventions to hold, despite them not being intrinsic to the nature of any conversation. It is rather a felt dispositional projection of one's own expectations onto the conversation. It is not simply the existence of cues within the conversation, but an *expectation* of these cues that are projected onto it, providing the illusion of the intrinsic nature of these cues. Therefore, we "expect to expect" (Dennett 2013a: 210) the existence of these cues within the conversation. When these are proved correct, there are no prediction errors to propagate forward, which confirm the expected existence of the properties "inherent" within the conversation. It should be stressed that, as incorrect as this view of reality is from a certain vantage point, it is an evolutionary adaptation which has served us immensely well in surviving the vicissitudes of the world. My reading of this is that our nervous systems have evolved to the point of projecting properties onto the world that are not there, all in order to better serve our continual existence.

Clark himself has emphasised the disconnect between our folk conceptions of beliefs, expectations, fears and the like with the emergent picture of neuronal processing (Clark 1989; 2013a: 198-199). This disconnect forces us, he believes, to apply our latest scientific and conceptual understanding "to improve social relations and education, to increase human happiness, or to inform our responses to social problems" (Clark 2013a: 199). The coming together of these two depictions (the above-mentioned manifest and scientific images respectively) is a worthy project to undertake in order to achieve as much as possible from our emerging understanding of both. Extended to the realm of FP, we need to expose our conceptual failings between what we think we do and what the science is telling us. This is all

in order to expand not only the understanding of our world, but the understanding of our place within it. Conflating beliefs and desires can be an example of the gradual project of merging our manifest and scientific images. But, as Clark (ibid.) put it, our scientific conception of the world must also allow for the causal influence of the folk psychological framework.

5. Conclusion

This chapter has introduced a rising view within the cognitive sciences, that of PP. It depicts our minds as operating primarily through the Bayesian modelling of the environment. Perception, therefore, is potentially nothing more than an "explaining away" (Clark 2013a: 187) of incoming sensory data matched with hierarchical cascades of predictions at a multitude of temporal and spatial gradations. The predictions depict what the PP system has already gleaned from the environment (along with "doubts" within its own cognitive functioning). According to Clark (ibid.), this leaves us with a "theory-laden" view of perception whereby our priors (or hyperpriors) take prominence in our perceptions of incoming sensory data. Error signals, which are all that remain after the cancelling out of predictions and incoming sensory signals, are fed forward through the system[126]. Therefore, it is primarily the recurrent (or backward) connections that carry the pertinent informational burden. The unifying potential of this view was introduced in showing how PP within the brain is seemingly action-oriented toward environmental cues. This action-orientation elicits a dual role for our underlying neuronal architectures by seemingly incorporating them for both perception and action. It was further explained that this is all in the service of minimising free energy due to the energy constraints of the homeostatic system that is our brain and body. Thereafter, the effects of this view of our fine-grained processing on our canonical conceptions of "belief" and "desire" were discussed. The conclusions derived were deflationary and decidedly anti-realist as to what exists within our neuronal interactions at this level. Firstly, both beliefs and desires were found to be subsumed into a singular expectational state (Hohwy 2013: 70-73). Secondly, it was found that desire does not feature as primary in our cognitive ontology. It is more a felt disposition which is the result of the aforementioned expectational state. These have very clear implications for our understanding of our cognitive structures (as PP is counter-intuitive), as well as our conception of our own volition (which adds to the free will debate).

[126] To put a stop to "the gossiping", in the words of Friston (2005: 829).

It was proposed that the problems for us going forward with this view of brain functionality could be mitigated by endorsing a view of "representation" within the mind which is at the very least deflationary, opening the door for a radically enactive depiction of our inner cognitive processing. Furthermore, endorsing a research paradigm that emphasises a closing of the gap between our scientific and manifest images would go some way to mitigating any misgivings as to the decidedly counter-intuitive theory of our internal processing that PP provides. The next chapter will expand upon the radically enactive approach in order to further undermine traditional FP, by advocating a minimally-contentful depiction of our mental states. This shall advocate a "representation-less" view of our mental states, and incorporate PP in achieving this task. I shall start by introducing the literature on enactivism, and thereafter further problematising definitive content within our cognitive deliberations. This shall help in bringing PP and this radically enactive approach together, and in so doing, potentially sweep away misconceptions about the existence of propositional content.

Chapter Four: Teleosemiotics and the Minimal Turn

> "Conscious thought doesn't have much to do with this stuff. Doesn't have much to do with the way we live our lives, period, if you believe the psychologists. A bit of rationalization, most of it with hindsight. Put the rest down to hormonal drives, gene instinct, and pheromones for the fine-tuning. Sad, but true".
>
> - Richard Morgan (2002: [s.p.]).

1. Introduction

This chapter will move beyond PP to explicate a theory of mind which endorses minimal content-acquisition for the majority of our cognition, and by extension, for our day-to-day FP. As with PP, it is a view which does not have a use for unnecessary processing if there are more frugal ways to arrive at the same output. This view of minimal cognition is called radically enactive embodied cognition and is proposed by Hutto and Myin (Hutto 2011, 2015; Hutto & Myin 2013a, 2017; Myin & Hutto 2015). It employs a more direct fit between body and world to arrive at its radically embodied view. Fundamentally, it proposes an alternate framework for our general cognition in day-to-day existence than more mainstream theories of mental cognition and representation, such as the teleosemantic theories of content. Currently, these teleosemantic theories of content (Millikan 1984, 1993, 2004, 2005; Papineau 1987; Dretske 1988), to be explained below, serve as the dominant theories in which to ground intentionality, and by extension content, into a naturalistic worldview[127]. This means that this alternative view, as proposed by Hutto and Myin (2013a, 2017), offers a more naturalistic proposal in grounding our intentionality by taking a more scientifically viable, and therefore not exclusively philosophical, position. I shall go briefly into the problem of content within cognition, before describing the shift to more enactive theories of the mind and consciousness. After describing these foundations, I shall elucidate the particulars that make minimal cognition an appealing theory in juxtaposition to teleosemantic theories of content. I shall throughout show how this theory can capture the inner processes of our predictive minds, thereby linking its literature with that of PP**.** I shall then describe the functional role that social scaffolding potentially inhabits in our enactive and PP minds. This opens the door for an understanding of our general,

[127] Recall that naturalism in its most general sense advocates a congruency with our best sciences (and therefore precludes the supernatural from its theorising).

day-to-day internal mental processes as *necessarily* streamlined and frugal (but not without its shortcomings as a result). The reason for introducing this lies in the need to ground the enactive and the predictive capacities of the mind within our social realm. This is necessitated by the need to explain how it is that our folk psychological concepts arise from our cognitive architecture, as described in the foregoing chapters. I will show how it is possible for the features of our basic cognition to be separate from, but still inform, the fully representational content that FP trades in. In this way, FP is preserved as a "just-so" relic that is nevertheless parasitic upon, and is descriptive of, our molar-level social lives. This view of the mind is one that does not trade in rich content and representation at a neuronal level. Utilising the conceptual framework of this thesis, I shall end with an endorsement of a view of our FP as explicable in terms of the basic functionality of our embodied selves, which is a far cry from previously formulated theories along the lines of propositional attitude psychology.

2. The Myth of Content

To begin with, recall from Chapter Two that the way that our FP is allegedly instantiated at the neuronal level is separated from the processes we invoke once we seek to explain subsequent behaviour. It is only once we seek to *explain* that we delve into the mire of personal reasons, thereby invoking the now infamous belief/desire pairing that supposedly describes the causal chain of events which produce the resultant action. The content of these folk ascriptions, recall, are traditionally thought to be derived from our sense-data. Beliefs and desires in canonical FP are inherently content-involving. But if action and attendant processes occur below the level of overt thought, I argue that we need not involve content-inherent processes in our theories of cognition. As intractable as it is to unravel the implicit thinking we all employ, it is usually understood that to be able to employ a truly structured and *instrumental* thinking, we must be dealing with content-rich internal cognitive procedures (Hutto 2008a: 43). It is understood, as stipulated in Chapter Two, that these content-rich cognitive procedures are what constitute the propositional attitudes, the propositions that encapsulate our beliefs and desires, and which allow us to systematise a coherent framework of instrumental cognition (Hutto 2008a: 43; Hutto & Myin 2013a: 13-14, Hutto & Myin 2017: 1-2; Muller 2014: 171). These propositions are thought to be the constituents of our content-bearing states, which need to be reciprocally connected in order to obtain the sort of instrumental reasoning[128] needed for intentionality.

[128] Recall that instrumental reasoning denotes overt cognitive deliberation.

These propositional attitudes, traditionally, are seen as the framework that makes the apprehension of our environment possible in a way that provides an informationally-accurate and integrated understanding (Hutto 2008a: 2). This propositional content is then what enables us to cognate in an intentional way. But, to free ourselves from this increasingly tenuous cognitivist paradigm, there is a need for us to distinguish between content-rich intentional attitudes and those that do not adhere to this form of intention, namely, proposition-*less* attitudes. But it is often thought that if instrumental reasoning is to be instantiated, the requisite propositional attitudes must be sufficiently constructed in order to achieve this (ibid.: 44). Fodorian-style thinking on this score depicts holding particular thought content (the semantics) as a token which can be re-utilised within a syntactic chain of thought (Fodor 1994: 97; Hutto 2008a: 45). This is how we can acquire our different propositional attitudes within differing contexts, through the re-utilisation of these delineated semantic contents. Or so it goes. To be able to grasp a thought with particular propositional content enables instrumental reasoning, while communication is only possible through grasping the same content by more than one person (ibid.). On this view, in order to instantiate a thought with the requisite content one must have more than solely an intention toward a particular state of affairs: as I have mentioned above, we need to separate content-rich propositional attitudes from purely *intentional* attitudes[129]. More specifically, for reasons that shall become clear below, I shall argue for a form of minimally-contentful, directed behaviour which can be described as ur-intentionality; the "ur-" prefix denoting a more originary and foundational intentionality than that typically invoked in these discussions (see Hutto & Myin (2017), Chapter Five). Content-rich attitudes are those of the propositional variety, while non-vocalised responsive behaviour need not be instantiated by this variety of intentionality for it to do work, and by extension it does not need semantic[130] referential characteristics of the sort needed in traditional propositional folk theorising[131] (Hutto 2008a: 46). Simply appealing to the overt functionality of FP in social discourse need not speak to a fundamental cognitive role for propositional attitudes, as shall become clear below.

What basis does the above have in reality? Arguably, it is a very plausible model of our cognitive processes, in that it can account for their functioning in a very parsimonious manner.

[129] Or, more accurately, intentionally-directed dispositional states.
[130] Recall that mention of semantics refers to the representational content of the mental state.
[131] Evans' (1982) is the modern seminal text on perceptual states as primarily content-rich. It is very forthrightly stated that the subject's perceptual state has content which can only be represented as true or false (Evans 1982: 226-227). It is further elaborated that any content-rich internal state is intimately related to behaviour (ibid.). Furthermore, these internal states are created through perception of the environment (ibid.).

For starters, we should not think that incorporating sense-data, which seemingly informs the content of our thoughts, into our cognitive apparatus is in need of transduction[132] from its originary, non-conceptual form into a finalised conceptual vehicle. Furthermore, it is not necessary to incorporate this transduced content into further cognitive functioning. This is what Dennett has called "the myth of double transduction" (1996). He called it a myth, as many in the field still feel the tug of intuition toward a theory of the "encoding" of environmental information derived from the senses. This "encoding" is of sensory information into a different medium that is actually used within "true" consciousness. He has used this as a useful metaphor to undercut our intuitive understanding of mental representation[133].

It is difficult to pin down what extra explanatory work is being done when we invoke additional vehicles of content to explain the way that sense-data is represented to us (Hutto & Myin 2013b). It would seem that on this view the brain is doing more internal processing than is needed and, based on the preceding chapter's discussion of the notion of free energy minimisation, it would not make sense to evolve in such a way. PP theories' form of representation could perhaps be deflated to contain minimal content. Yet, we are still adding unnecessary processes to the brain's inner workings due to our need to incorporate semantic content at some level. This, to my eyes, is a direct result of the traditional analytic paradigm and, on a more fundamental level, the natural human inclination to humanise biological processes. Even teleofunctional theories of mind (introduced earlier in this thesis) endorse a view which is too dependent on information that affords a capacity for content to *represent* a state of affairs (Hutto 2008a: 47). Recall that teleosemantic theories attempt to explicate the notion that the norms of mental content, if they are accurate/true or not in virtue of their environmental causes, are fundamentally derivative of their *functional norms*[134]. Therefore, the accuracy or truth evaluability of mental content ultimately comes down to a biological form of functional normativity. But the degrees of fit which instantiate relations of a semantic nature, thereby guiding the organism toward a directed action, are thought to refer to semantically-defined vehicles within the brain and the features of the external environment which they represent (Hutto & Myin 2017: 43), and this also extends to the social environment. For me, there is no need for semantically-implicated vehicles of content.

[132] Transduction here refers to the change of one form of signal into another.
[133] For example, to assume that the information from the neuronal signals that are registering pain need to be *re*-encoded (transduced again) in order for us to properly cognate this pain, is a misunderstanding of how our mental processes work.
[134] The functional norms being those representational relations that have been fixed due to the particular evolutionary attunement of an intentional state toward distal stimuli.

Content-rich representation within the mind has often been automatically assumed, but now it appears that adherents are more defensive. For example, Colombo endorses a representational view akin to Clark and Toribio's (1994) wherein "representations are necessary to tackle representational-hungry" tasks (Colombo 2012: 9). Yet, notice how what is to be defended is assumed from the start. *Of course* one would need representations to explicate representation-hungry processes. It's the *existence* of representation-hungry processes in basic functional cognition that is in need of defending, not the form of the representations themselves. Hutto and Myin (2013b: 6-8) call this the "constitutive-explanatory fallacy", in that the biases inherent within theorising constitute and constrain the explanations derived therefrom. As a paradigm case for the invocation of representations within our basic cognitive processing, Colombo fails to introduce a single basis for preferring a representational approach to the mind, or even to explain how it adds any explanatory weight toward explaining "representation-hungry" processes, such as social attunement. No singular mental state accurately representing an external state of affairs is necessarily going to add anything over and above what a non-representational view would. Consider the following: Colombo states that neuronal information is transduced from sense-data and then "decoded" in the form of an "extraction of information […] from neuronal spiking" (ibid.: 7). This does not add anything more to the story. Effectively, Colombo does not base his explanation on an empirical basis, but instead relies primarily upon armchair philosophy to arrive at his conclusions. One can begin to see the issues at play here.

Representationalists need things in the mind which can account for, or stand in for, properties which are "abstract" and "absent" (Clark & Toribio 1994: 419). This would enable sensitive responsitivity to environmental stimulation, as these stand-ins represent the states of affairs and also somehow represent what is not present. The "absent" is that which is either previously experienced, or newly formulated within the mind. A certain form of sensitivity must be in place to account for that which is not present, but which nevertheless informs our perception of the state of affairs. Granted, this is an intractable issue regardless of how we think about the mind, as the difficulties of unravelling the complexity of the mind continue to sit with us. But this is an issue which we can empirically work toward solving if we ignore the distraction of the issues surrounding content and representation. Consider again: Colombo (ibid.: 6) states that neuronal firings correlate with environmental stimulation in a functional manner, and this functional relationship can be explicated as a form of code. But there is then quite a leap: Colombo (ibid.: 7) goes on to say that "*neural representations* can be said [to be] the

constituents of the neural code" (emphasis in original). Why are neuronal representations invoked at this point? For what purpose other than to shoehorn conceptual favourites into a preordained theoretical structure? We are even told by Colombo (ibid.: 20) that Bayesian predictive modelling of the sort introduced in the previous chapter *necessarily* entails representations of the content-rich kind. According to him, we can only account for our mind's ability to manipulate and control behaviour if representations are introduced, as representations are alleged to afford us the ability to "identify particular signals" (ibid.: 18). But there is current debate as to whether this is an accurate portrayal of just what "representation" means within the context of PP (eg. Gładziejewski 2016; Downey 2017; Dolega 2017). Secondly, he states that the explanatory power that the representations afford within cognitive science should also not be discounted. But an appeal to how scientists and philosophers have traditionally conceptualised their theoretical work is not a knockdown argument. In conclusion, to be naturalistic, one needs to adopt a rigorously scientific perspective, but in such a way that we are not restricted by traditional theorising on cognition and an over-emphasis on "representation" as the mark of the mental (see for example Ramsey (2014: 4-11) for a more extensive review). Effectively, taking representations "as *both* interesting explanatory constructs *and* as a necessary condition for a legitimate account [of the mark of the mental]" (ibid.: 8) comes across as circular.

3. An Enactive Turn (For The Better)

3.1 Theories More Radical

Before laying out a minimally-contentful picture of mental representation, a brief outline of the foundations of this view is necessary. Minimally-contentful theories of cognition stem from the "E-theories" of mind. Hutto and Myin (2017: 1-2) list the "embodied, enactive, extended, embedded, and ecological aspects of mind" to drive home the "E turn" within cognitive philosophy. These "E-theories" arose primarily to account for the fluid nature of cognition and to overcome the rigid constraints of utilising propositional processes to explain the mind (Sutton *et al*. 2011; Dreyfus 2014). These theories are canonically known as the "4E conception of the mind" (Rowlands 2010: 3, attributed to Shaun Gallagher), which refers to the mind's embodied, enactive, extended, and embedded aspects. Recall that enactivism is a catch-all term which refers to theories endorsing the view that cognition arises through the dynamic coupling between body and environment. This dynamism arises through the actions of an *active* body,

wherein reality is constructed through the body's selective interaction with the environment, and in the way that the environment acts on the body in turn. The "embodied" refers to the non-neuronal processes of the body which nevertheless influence the mind, while the "embedded" refers to the influencing effect of particular scaffolding within the environment, a typical example being the enhanced cognitive ability that language affords. In the absence of these scaffoldings, both bodily and environmental, the mind operates, at the very least, sub-optimally. There is also the notion of the extended mind, in terms of how the mental extends outward to incorporate the environment into its processing (more of which to follow).

Even though there has been an underlying acceptance of some form of propositional thinking in writings on the topic, there has also been a gradual shift toward other conceptions of cognition which can account for our quintessentially fast and seemingly streamlined thinking. The apparent dynamism of our minds is difficult to account for through brute, rigid propositionalism. Therefore, there has been a move toward the embodied, enactive, and extended forms of cognition, which make much of new empirical research into motor control functionality, mirror neurons, proprioception and the like[135]. In other words, the concepts of embodied, enactive, and extended cognition do not follow the contours of the traditional cognitivist paradigm.

E-approaches seek to move beyond the dominant cognitive theories which have been around since the middle of the previous century and which take it as fact that the mind is a representational and computational machine (Hutto & Myin 2017: 4). These two key concepts, representationalism and computationalism, are the backbone of the cognitivist approach, and in themselves rest on more fundamental methodological commitments. Firstly, in terms of mechanistic analysis, parts of the mind/brain are often analysed along with their interactive processes, and secondly, it is shown how these interact to create the target phenomena under analysis (Horst 2007: 16-17). Taking this further, it is assumed that the cognitive components of the brain, and the brain alone (Aizawa 2015: 2), are the foundation of intelligence. Linking with this thesis, the computational mechanisms which are thought to produce this behaviour entail brain-bound computational processes with inherently contentful mental representations. This neuronally-bound depiction of mind is quintessentially of an interiorised intellectualism,

[135] There are too many studies to count here, due to the vast proliferation of research on this theme. Research on motor circuitry reutilised for memory (Casasanto & Dijkstra 2010), mirroring within motor functions and emotions (Rizzolatti & Sinigaglia 2010), and many others have expanded our body of knowledge on the topic. These studies serve to highlight the interrelations of mind and environment, and hence depict them as not wholly distinct from each other.

where the "I" is paramount (Hutto & Myin 2017: 5). In contrast to this, E-approaches depict the cognitive realm as directly intertwined with bodily processes (embodied), the environment (enactive), and with non-neuronal objects *within* the environment (extended, embedded, ecological).

Not all theorists necessarily ascribe to every concept equally or even at all, with some even supporting fundamentally unaltered cognitivism with a few enactive or extended bells and whistles. The paradigmatic enactivist stance, for example, usually takes the form of some kind of embodied representationalism with inherent content (see Goldman & de Vignemont 2009; Goldman 2012, 2014; Gallese 2014). We could call these the conservative cohort of the embodied mind theorists. According to these theorists, the science points toward the re-deployment of embodied mental representations (representation of the body itself) as necessary for a vast number of cognitive processes. These embodied mental representations[136] "represent" information about the body in such a way so as to perform a useful function in cognitive processing. While having issues with some of the conclusions derived from research on embodied forms of cognition, Aizawa (2015: 4) puts it most clearly when he states that our cognitive processes are embodied in such a way that the underlying mechanisms of perception and action-response are *the very same for* conceptual reasoning. Therefore, the conservative branch of the E-theorist tree sees our mind as embodied in ways that extend beyond the confines of brain and skull, but claims that the ultimate processing of conceptual reasoning occurs via contentful handling of (semantic) representations.

The next and more radical step in extending the bounds of cognition appears in the writings of the extended mind theorists. They are more explicit in their views that cognition can be a temporally-extensive, non-neuronally-bound, embodied interaction with the environment and the objects within it. The most renowned version of this view is the extended mind hypothesis (EMH) as originally formulated by Clark and Chalmers (1998). On this view, the mechanisms which enable cognitive functioning extend beyond the body and into the environment, thereby offloading some of the cognitive burden during certain cognitive processing tasks (see also Clark 2008; Rowlands 2009; Wheeler 2010). The theorists who endorse this view place much emphasis on the extended capability that *tools* provide. These can consist of objects such as computers, bodily implants, or even the mundane in the case of the vast powers that words,

[136] For example, an inner "representation" of our bodily position in relation to the environment.

whether written or spoken, provide[137]. The general idea is that without complementary props (or extended cognitive vehicles) (Clark & Chalmers 1998: 8) in the environment, our minds are simply not capable of certain forms of enhanced cognition. These theorists usually advocate that causally effectual sources of information from these environmental props are what actually drive cognition (Aizawa 2014: 15). There is still a nod toward representationalism as inherently pertinent within cognitive processing, as any reading of Clark would attest to. But the strictures of traditional cognitivism have been relaxed to extend the boundary outward beyond the brain. However, there still appears to be a form of "intellectualised enactivism" (Hutto & Myin 2017: 61) at work here. For example, Clark often writes in a way that incorporates cognitivist references to "representations" (Clark 2013a; 2016 and others). This speaks to an enactivist who still believes that representation is necessary for the mental to do work. Clark quite clearly sees precision-weighting of estimates from our predictive minds (for example 2013a: 188) as the primary instigators of action, but these cannot be captured within "ordinary daily speech" (Clark 2015b: 5). What, then, can these precision-weightings be *of* if they perform the role of the contents of representations within the cognitive make-up of the brain? Clark proposes that the inferential processes that generate representations from within the "black box" of our mind are adequate to satisfy any philosophical qualms (2013a: 189). However, Clark does not see them as propositional in nature, so some adjustment needs to be made in order to tie this propositionless picture of the mind and his construal of PP together.

There has been another trend within the literature that moves even further away from the old cognitivist paradigm. This further move has coalesced into the sub-discipline of sensorimotor enactivism, first articulated within Noë's (2004). It is a thesis which emphasises fully embodied interactions with the environment in such a way as to blur the line between the two[138]. It differs from other forms of extended cognition in its rejection of functionalism, and especially the multiple realisability thesis. What this denotes is that it is only through our bodies that we can have our particular brand of cognition (Noë 2004: 25; 2009: 62), and any other substrate is

[137] Clark (2003), as well as Myin & Veldeman (2011), introduce abstract art as an unlikely yet perfect encapsulation of this thesis. They point to the work of van Leeuwen *et al*. (1999) who uses the example of sketchpad use within abstract art. External sketches, which are often used within abstract art, are necessary due to the difficult task of maintaining simultaneous interpretational images in the mind. These sketchpads therefore augment our limited processing capabilities.

[138] As a clear everyday example, consider the skilful riding of a bike (Degenaar & O'Regan 2017: 2). To successfully ride the bike, one must be able to keep one's balance despite mitigating circumstances (an upcoming bend, bumps in the road, excessive wind). This is an enactive process, as the "dynamic engagement with the environment" (ibid.) cannot be described in purely representational terms.

insufficient to account for our embodied interactions[139]. Further, it rejects the internal, richly-representational models of perception which are found in other views on cognition. Noë (2004: 216; 2009: 7) states (in an admittedly abstract fashion) that perception should not be thought of as something occurring within us, but rather something we actively do. So, while the neurons within the brain are necessary for perception to occur, it is only through the action of the individual that it can come about[140] (Noë 2004: 227; 2009: 47). Despite this radical shift in trying to explain cognitive processing, Noë's theory still holds onto the idea that rich content arises after the use and manipulation of sensorimotor information (Noë 2004: 17, although Noë's position has changed slightly since then, see 2009; 2012). On this formulation of sensorimotor enactivism rich content is still invoked, but it is *dependent upon* the utilisation of sensorimotor information. Therefore, there is still a reliance on representationalism within even basic cognition, despite this view's less conservative outlook than the extended mind theorists. While this need not be incorrect at face value, it does not go far enough according to the next theory to be reviewed.

There is one more step to take to arrive at the requisite platform for this thesis. This platform is what Hutto and Myin (2013a: 1; 2017: 9) have dubbed the "radically enactive, embodied account of cognition", or simply "REC". It takes a cue from sensorimotor enactivism in that cognition is still interpreted as an activity that individuals do, but it is instantiated through dynamic processes which loop back onto themselves while remaining attuned to the information (in an attenuated sense) derived from variables within the environment (Hutto & Myin 2017: 9). These informationally-attuned processes do not involve any content-inherent perception on a basic level (ibid.). Content, again, refers to correctness conditions which inform a state of affairs in the world as being accurately represented, while basic refers to the *fundamental functioning of our minds*, and not a simpler form of our brain processing. "Fundamental functioning" denotes what is left when content-rich social vehicles of thought, such as the propositional attitudes, are subtracted. In other words, this is what allegedly lies at the base of all of our cognition[141].

[139] One is wary of stating it in this way, as while it has the potentiality to be true for our particular cognitive processing, we should not think that this precludes other substrates from being able to instantiate a form of consciousness.
[140] Again, think of the example of riding the bike.
[141] This view of the mind rubs against the theories which posit an extended view of mind, but which still include a rudimentary form of content in their theories. Hohwy (2013) and Clark (2016) are examples of this within the PP view.

REC is a movement within enactivism which is primarily inspired by the latest in "robotics, dynamical systems theory, and ecological psychology and which finds philosophical support from the phenomenological, American naturalist, and Buddhist traditions of thought" (Hutto & Myin 2017: 9). While some of these elements are less scientifically helpful than others, they have provided inspiration for the general approach of the philosophers working in the field. As far as REC is concerned, it primarily falls under the umbrella of work begun with Varela *et al*. (1991) and their phenomenological approach to cognitive science inspired by the work of Merleau-Ponty (1962). Their project entailed a distinctly left-field approach by incorporating Eastern influences into their theorising which most have balked at, but their seminal text has inspired a proliferation of work in the field regardless. What all of the theories that have followed since have in common however, is that cognition is a situated[142], embodied, and enactive process, which also does not necessarily entail richness of semantic content.

There are a number of theorists in this sub-discipline who emphasise their own agendas, but they gravitate toward the same explanatory goals (Thompson 2007; Di Paolo 2009; Chemero 2009; Froese and Di Paolo 2011; Hutto & Myin 2013a, 2017; Bruineberg & Rietveld 2014). To reiterate, people who follow REC *do not* endorse content that relies on any form of "correctness condition". This usually amounts to talk of "representing" phenomena in a particular way so that there is the possibility of it *not* being so (Hutto & Myin 2017: 10). This is usually the quintessential view of content by philosophers of an analytic mindset that associate content with propositionality, whereby correctness is equated with accuracy, which in turn is synonymous with truth conditionality (ibid.: 10-11). For example, Brogaard (2014: 2) insists that perception is "accurate in virtue of some proposition p being true". The REC view opposes this. People who follow the REC path acknowledge the malleability of the concept of "content", and thus seek to relax the rigid strictures often tacitly accepted by philosophers of mind. This relaxation provides a notion of content which does not necessarily everywhere and always equate to a form of truth conditionality (Crane 2009; Burge 2010: 303). To be clear, this *does not* mean that REC endorses no content across the board (Hutto & Myin 2017: 11). This form of eliminativism would be *really* radical, even more so than REC. There are states of cognition which possess content in the rigid form of "correctness conditions", but it views these as not part of our basic functionality. Rather, it views these cognitive (truly propositional) states as a specific and late-developing ability enabled by cultural "scaffolding"

[142] Broadly-construed, this simply means that our brand of cognition is always found within particular contexts which in turn affect our theorising. Knowing is never seen as separated from our context.

(more of which to be explained below) (Hutto & Satne 2015: 527). Therefore, "basic minds" are our default, with semantically contentful minds denoting the "non-basic". Or, as Hutto and Myin (2017: 13) take it, basic minds are "phylogenetically and ontogenetically fundamental". There is nothing "basic" about basic minds beyond their lack of semantically-derived content (which does make one wonder why the terminology was chosen if it was inevitably going to lead to confusion). Content is parasitic upon our basic minds, and content is only introduced once social vehicles of content are introduced into the equation (which shall be looked at further below). REC, effectively, depicts our basic functionality as dispositionally directed toward the environment in a non-semantic way.

Fundamentally, REC finds its origins in ecological dynamics, which consists of the twin disciplines of ecological psychology along Gibsonian (1979) lines and dynamical systems theory. Gibsonian ecological psychology argues that organisms are tightly bound to their environment in such a way that perception is primarily in the service of action (ibid.). Perception, in other words, serves as the primary means for acquiring a practical grasp of the environment in a continually active manner. Furthermore, these directed and dynamic actions toward external phenomena do not require mental representations as mediators (ibid.: 279-280). Dynamical systems theory provides the mathematical capability for outlining this approach[143] (Hutto & Myin 2017: 9).

The basic understanding of REC (and other radically enactive theories), is that organisms' cognitive abilities evolve in context-sensitive capacities eliciting a *constrained* interaction between the body and environment over time[144] (ibid.: 22). These interactions are closely related to the *affordances* of the organism, being those things within the environment which matter most to it. It is this which enables and changes organisms' dispositional capacities. Neuronal weightings[145] are effected and changed, but this does not require positing a representational intermediary for causal power[146]. Being able to remember is simply a dispositional re-enactment of embodied capacities attuned to affordances, with minimal

[143] For example, Chemero (2009) has done an exemplary job of combining these two approaches to describe the efficacy of a radically enactive approach to mind in dissolving representationalism.
[144] Or, to acquire "a grip on the *patterns that matter* for the *interactions that matter*" (Clark 2015b: 5, emphasis in original).
[145] It is worth taking note of Kandel's (2001) seminal work on memory and the learning capacities of the brain in which he won a Nobel Prize. Kandel demonstrated (within simpler animal isomorphs) that it is not the properties of the neurons within the hippocampus which serve learning capabilities, but rather the *interaction between* these synaptic interconnections. Although, Kandel himself explicitly sought to reconcile this with representationalism in his writings (ibid.: 1035).
[146] Again: no double transduction in the brain.

content needed to derive an explanation for the dispositional processes of the organism (Ramsey 2007: 151-187). Barandiaran and Di Paolo (2014), in their genealogical review of the concept of "habit", point toward theorists who endorse the ways in which organisms enact an embodied responsitivity without any knowledge invoked during the process. Habits accrued through the lived experience of the organism can do the work needed for adequate responsitivity. No "thinking" in the full-blooded sense is occurring here at all (Rosenberg 2014: 25-27). To make this clearer: during the instantiation of environmental information within the brain, the function of differing regions within the brain which "represent" this information is dependent upon the moment that the information was instantiated. This is why memories are more of a "dispositional property" wherein the particular regions of the brain are "triggered by the right cue, in roughly the same pattern of activation they underwent during encoding" (De Brigard 2014: 169). Therefore it is the past lived experience of the organism which is paramount here, and not a focus on information as conveying contentful data.

But then what drives the notion that content is the harbinger of the meaningfully cognitive? It is seemingly primarily derived from the intuitive notion that *abstract* properties inherent within working memory must be coherently re-presented (Hutto & Myin 2017: 33). This intuition precludes the existence of minimally-contentful, basic minds. This persistent conception of the "mark of the cognitive" (ibid.: 13) insists upon vehicles of representation within the brain which carry experiential content, which are the necessary components for behavioural changes. The thinking is that abstract properties are transposed and replaced by these vehicles of thought in order to make them useable to the mind[147]. Clark himself states that it is necessary that "an inner item, pattern, or process whose content then corresponds to the abstract property" exists in order to be truly cognitive (Clark 1997: 167). The force of this intuition (perhaps driven primarily through our perceptual modality of sight) can overwhelmingly drive us toward the view that our brains deal with information processing and representational vehicles of thought (ibid.: 37). Of course, as discussed above, the crux of the issue is precisely how positing representations adds any explanatory weight over and above bodily dispositional learning through environmental engagement. The content/vehicle distinction in the classical cognitivist formulation falls away as a lack of content necessarily entails a lack of vehicles. Therefore, we need to step away from the Hard Problem of Content[148], as conceived of by Hutto and Myin

[147] A fairly (philosophically) recent primary text on the idea that representation within the brain is pertinent for fully-realised cognition is Clark & Toribio (1994).

[148] This is a play on Chalmers' (1995) notion of the Hard Problem of Consciousness. Chalmers posited that the actual problem, the really, truly, difficult problem of consciousness, lay in the explication of how we have our

(2013a, Chapter Four; 2017: 29**)**. A "gapless naturalistic account of cognition" (Hutto & Myin 2017: 41) is what is at stake here, and only depicting information in terms of covariance (see below) and the norms engendered by biological functionality seem to be able to provide this.

Granted, explanatory naturalism has notoriously found it difficult to account for the move from content-less to contentful cognition. Yet, many have proclaimed that content is not the problem it is purported to be: Chalmers (1996: 24) has spoken of content as not posing any mystery, while Strawson (1994: 44) and Miłkowski (2015) have similarly put forward that content does not pose any great problem for us. For these theorists, content should be invoked, but not only that — the problem of content has been solved. According to Miłkowski for instance, what has solved the problem of content is the explanatory power of teleological biological function as proposed by Millikan and others, which shall be explained further below (ibid.: 83). This does not appear to be as obvious as it has been made out to be, as it is my contention that content has not been suitably naturalised, which is a position I originally argued for when describing Fodor's views in Chapter Two. In light of this, if the framework of the paradigm restricts our deliberations, it is worth looking at our fundamental suppositions and altering them. We can walk away from the Hard Problem of Content by not assuming content from the outset. This is not a defeat, as some may surmise (Aizawa 2014, 2015), but a changing of the rules of engagement. The environment need not be accurately represented by cognition, but instead it should couple the organism with its environment in order "to stabilise appropriate coordinated patterns of behaviour" (Beer 2000: 97). Full-blown propositional attitudes will only follow in the wake of these basic forms of directed behaviour, once sufficient social scaffolding has been introduced.

While there have been many arguments for cognition along REC lines (also see Dreyfus (2014) for many examples), the depth of opposition to this view is vast. It is often argued that for something to be truly cognitive, it *must* equate with content as representational. We often lapse into talk of the representational ability of our minds. For example, it is thought that without representation we are "bereft of tools for explaining natural intelligence" (O'Brien & Opie 2009: 54). Similarly, REC has been labelled as more a theory of behaviour than illuminating any truly useful phenomena about the mind (Shapiro 2014; Roy 2015: 95). It is in this manner that representational views of cognition, even if more radical in the form of extended cognition,

phenomenal experiences (or qualia). The issues that plague cognitive and neuroscientists are the "easy problems" (ibid.). See Dennett (2005, Chapter Three) for why this is an unhelpful and distracting formulation of consciousness.

are argued to be more coherent approaches, without having to fully tackle the validity of a REC conception of mind. However, a better explanation is needed to show the worth of representationalism as a superior framework to REC-like cognitive theories, and this explanation should not be circularly touted as a necessary condition for cognition in the first place (Hutto & Myin 2017: 15-16).

3.2. An Amendment to Biosemantics

Arguably, to take a REC view of the mind, we are compelled to describe our particular brand of intentionality in another form, one which does not include content-rich representation. I do not believe that intentionality has been suitably naturalised in its attempt to explain the way in which content is allegedly instantiated at the neuronal level. To me, invoking content at this level is misleading, and we can go further to naturalise our understanding of our intentional states. We all know that biological organisms are attuned, or sensitive, to contributions from the contextual environment, even if it is difficult to extricate how some of these sensitivities are effected. It is these attuned sensitivities which allow for the identification and delineation of specific environmental discriminations that are of high significance for the organism in question (see Sterelny (2003), Chapter Two). Once these discriminations have impinged upon the sensory-intakes of the organism, a response is elicited in order to accommodate a successful synchronisation of action[149]. This array of processes is the end result of selective pressures within the environment, formed purely due to prerequisites for the continual survival of the organism[150] (Hutto 2008a: 50; Sterelny 2003, 2015). What one would perhaps notice is that there is no real need for the encroachment of content-involving actions for these creatures, *including* ourselves. This is an all-pervasive feature of our minds (except once content-rich public symbols are invoked), and introducing content at this level does no extra work. There is no need for the introduction of content-rich representations or the like at this stage of cognitive processing due to their lack of explanatory weight, and thus it cannot be said that it is these representations which elicit appropriate responses. To perceive through sensory attunement instantiates an immediate response which is characteristic of the context due to the aforementioned selective pressures (Hutto 2008a: 51). Affordances accrued through the

[149] One will note that PP depicts this sensory impingement and resultant action in a much more entwined fashion than explained here.
[150] Otherwise known as end-directed means.

evolutionary and ontogenetic development of the organism[151] provide perception "shortcuts" that bypass overt cognitive processing. Depending on the organism and its evolutionary environment, the differing adoptive approaches can be quite marked (again, see Sterelny (2003), Chapters Two and Three)[152]. Also, while the highly complex response-attunement of (for example) predator and prey or human social communication may be a world away from the modest, visceral actions/reactions of detection-response coordinations, these aforementioned are up-scaled versions of the very same basic processes (Hutto & Myin 2017: 13).

The character of the receptivity of these detection processes are explicable through the actions of what Hutto (2008a: 51) calls "indexically inspired Action Coordination Routines", or ACRs. These coordination routines are instantiated through the "scripted" configurations of recursive action, incorporating mechanisms that cooperate to achieve this end. These ACRs produce, after paradigmatic features are detected, changes within the body, which include proprioceptive, exteroceptive, and interoceptive alterations, resulting in a suitable paradigmatic action-response (ibid.). These detections are therefore responsive to sense-data and induce neuronal activation in the typical sense, but do not involve "contentful" responsitivity. Therefore, ACRs should be seen as minimally contentful action-response routines, but they are the base from which contentful comprehension flows[153]. In other words, what matters for my current argument is the realisation that these ACRs are, in Hutto's words, *indexical*, and without content (Hutto 2008a: 51). This indexicality refers to the ability of these ACRs to sensitively *indicate* phenomena in the environment despite shifting contexts. ACRs are the brain states which indicate, are sensitive to, and are therefore triggered by, particular environmental phenomena in particular environmental contexts. Furthermore, there is no need for transduction of this sense-data into another "content-bearing" state, and these minimally

[151] An everyday example would be the affordances that specific objects have for us: buttons can be pushed, levers can be pulled, wheels can be turned. Of course, affordances extend to much more than these obvious, everyday objects, and include direct attunement to the environment in general. These examples highlight instances of everyday interaction with the environment that do not necessarily require an overt thought such as "I believe that I should push this button". It does no extra work than what an implicit learnt mechanism of behaviour can attain.
[152] For example, in a series of experiments it was shown how capuchin monkeys are seemingly unable to utilise reasoning capabilities when given a choice of sticks to acquire peanuts out of a tube (Visalberghi & Limongelli 1995). Occasionally they would even use the tape that some sticks were bound with. In contradistinction to this, tamarins were shown to have the ability to select appropriate sticks for the job in a subsequent study (Byrne 2000). This need not point toward overt reasoning capability on the part of the tamarins, but the environmental space in which they have evolved has facilitated an ability to re-enact a skill-set across multiple domains.
[153] Dennett says something similar to this when explaining the affordances that all organisms have within their *Umwelt* that enable degrees of *competence*, but importantly not yet *comprehension* (Dennett 2017, Chapter Five). However, this was within the context of degrees of cognitive ability across species, and not differing levels of cognitive function within the same brain.

contentful states do not "represent" as cognitive thinking is often thought to. The environmental types (or kinds) that are relatively consistently identifiable provide the platform from which intentional behaviour is instantiated and at which it is directedly focused. It is perhaps here that the priors which inform predictive processing models[154] find their apparent ability to accurately enough guide the organism through the environment. At the very least, this indexicality has the potential to explain non-communicative cognitive processes far better and parsimoniously than needing to invoke actual understanding and comprehension on a conceptual level[155] (ibid.: 51-52). Bringing this into the PP fold, we can call them what Clark has termed "desert landscape" (2013a: 186, 200) depictions of environmental attunement, wherein

> proprioceptive prediction errors act directly as motor commands. On these models it is our expectations about the proprioceptive consequences of moving and acting that directly bring the moving and acting about. (ibid.: 186).

Furthermore, the responsive and intentional nature of non-communicative cognition could potentially be explained through the use of biological proper functions. These denote those biological attunements to the environment which have evolved through paradigmatic cases which have fixed their performative function over time (Millikan 1984: 17). Therefore, a good depiction of proper functions are those biological traits that organisms have accrued through the selective adaptation of evolution. There is a sense that there is a "proper function" that a particular trait has evolved to perform (through blind biological processes). This does extend further however, as even beliefs fall under this category, as the fixing of a belief is due to particular selective pressures within the environment (ibid.: 17-18). A more concrete explanation will occupy us further below in this chapter.

To determine proper function, as Hutto (2008a: 52) rightly states, we should investigate the difference between what the organism in question does with how it "ought" to behave. Cognitive mechanisms have evolved due to selection pressures toward a particular functionality, but if this functionality "misrepresents" a state of affairs leading to an inappropriate behaviour from the organism within a particular context, this does not negate its evolutionary function. It merely points to the fact that these mechanisms have not been optimally "designed" or that they have been deployed in the wrong context. Taking this

[154] As first introduced in Chapter Three.
[155] Millikan (2004: 211-220) also points to the important difference between thinking in terms of detached representations and "getting by" within the environment via sensitive indexical responding. This was in her juxtaposition between nonhuman and human modes of cognition, however.

perspective circumvents the issue surrounding proper functions and the disjunction problem, as it explains the potential for error within each bodily response. Fodorian issues regarding correctness conditions can be avoided if this approach is taken. Recall that the disjunction problem highlights the philosophical (and biological) issue of how to distinguish between causes that correctly determine the content of mental states and those that do not. To belabour a point, any mental states develop their "proper biological function", and therefore their accuracy conditions, according to a long historical process of "selection by consequences" (Hutto & Myin 2017: 105). To appreciate the accuracy conditions of a mental state, one cannot simply look at the current dispositional make-up of the mental state, but at how it became *historically fixed to perform such a function*. Advocating the need for synchronous accuracy conditions for each mental state, such that they fix the norms for each mental state about a particular thing in the world in order to be able to represent or misrepresent an object in the world, is an unnecessary theoretical imposition. Upscaling the biological proper function of mental states, we can begin to see space emerging for sufficient "norms" of function lying at the base of our intentionality which are fixed over time. To invoke accuracy conditions beyond this is to impose socio-cultural norms upon our base intentionality. The significance of undermining the disjunction problem lies in the fact that an unnecessary philosophical hang-up is discarded, forcing a more robustly biological grounding in the process. The evolutionarily-forged mechanisms were not created for optimum end results *at all times* but, rather, for optimum end results *for most of the time*, within environments suitably similar to that which selected for the mechanisms in the first place (ibid.: 105-106). Furthermore, the accuracy conditions may be upscaled to a generalised form of norm fixing, which is arguably all we need for the mental to do work. Talk of "accuracy conditions" in light of which a mental state can misrepresent something in the world does not make sense on the intentional scale.

The paradigm case of the issues that arise when attempting to naturalise representation in terms of adaptation is that of frogs and their fly-catching habits. In scientific experiments, these frogs are found to lash out for both flies as well as black dots. The epistemological issues surrounding what mental content is represented within the mind of the frogs has been a philosophical issue ever since (Ritchie 2008: 169). It follows the contours of the disjunction problem (from Chapter Two) in that philosophers have agonised over what can reliably be said about what is *actually* represented in the frog's brain when it mistakes a moving black dot for a fly. There is no suitably appropriate method for distinguishing between differing explanations which are equally likely to account for the semantic mental content of the frog. Any description of the

causal history of what created the attunement between mind and environment here would suffice. The evolutionary development of the frog would slowly fix the norms that matter to the frog in terms of base intentionality. To explain this causal history somewhat, we need to project our minds into deep time to see it more clearly. By way of example: a chemical differential between a unicellular organism and its environment may develop over time into more and more sophisticated attunements to the environment as it evolves. As time progresses, the "norms" of behaviour are fixed according to the homeostatic inclinations of the evolving organism. As evolution progresses, we eventually come to more complex creatures further up the phylogenetic tree which possess what we would call "minds". Along the way, these creatures have slowly and painstakingly acquired sensitive attunements to the environment which do not trade in accuracy conditions as such, but are nevertheless norms of biological behaviour attuned to the base biological needs of the organism. As for the frog, it merely continues in this vein, consistently striving to ingest the fly which it unwittingly "understands" as necessary for survival. We need not invoke accuracy conditions over and above what the frog has simply had hardwired into itself to survive.

The corollary to this, as the worry goes, is that such attunement is simply not fine-grained enough to provide the necessary (in other words unique) semantic mental content (Rosenberg 2014: 26). Only a subclass of the entire category of "moving black dots" need to obtain enough times for it to be a satisfactory evolutionary benefit for the frog. What is it then that the frog "represents" when confronted with a fly as opposed to a black pellet or a black dot on a screen? It, again, becomes an empirical issue whereby the intuitions of many a philosopher or scientist cannot reconcile the "root mismatch between representational error and failure of biological function" (Burge 2010: 301). My claim is that there is no further need for selective pressures to reinforce a more semantically fine-grained picture of the world, unless the failures really started to add up (Hutto 2008a: 68). It does not make sense to say that the frog "misrepresents" in the way that we mean it. We are imposing our own socio-culturally mediated accuracy conditions upon that which is more mentally autonomic than we suppose.

Proffered naturalised theories of representation which do not take into account evolutionary processes (broadly speaking) will always come up short, and these processes need not lead to a semantically fine-grained picture of the world. If one were to describe the *reasons* for an organism's behaviour, drawing from the informational content in the brain, it would not necessarily describe responses to specific environmental kinds (ibid.: 52-53). Again, we come back to the issues surrounding the "decoding" of content, or double transduction, and how no

re-interpretation of signals are needed for immediate and visceral response[156]. No internalised content-discrimination need be made on the back of internal content manipulation which states that this is indeed the current state of affairs. In other words, the cognitivist foundations of information processing and representation are potentially unnecessary for bodily action-responses. As a result, the disjunction problem, that of the problem of being able to accurately represent differing phenomena that elicit the same pattern of activation within the mind, is circumvented. In other words, as long as responsitivity to the environment is more or less consistently appropriate, then the disjunction problem is of no real import. Again, doing away with the disjunction problem weakens views which have strict conceptions of the correctness conditions of mental states, allowing for the more relaxed, embodied depiction of mental states advocated within this thesis.

In light of this, we are in a position to adaptationally explicate intentionality in a way which does away with issues of accuracy conditions and content-fixing. Typical philosophical theories which reflect on the biological nature of cognitive (semantic) norms derived from norms of function (or proper function) tend to formulate and expand on theoretical frameworks which rely on representational content (Millikan 1984, 1993, 2004, 2005; Papineau 1987; Dretske 1988). These are the teleosemantic theories of content, which attempt to naturalise the intentionality of brain states. Millikan's theories are perhaps most well-known in this regard. Her cleverly appropriated Pushmi-Pullyu "icons" (Millikan 1995, 2005) are depictions (of some form) of just what I have been speaking about in terms of the indexical ACRs. She still appeals to a rudimentary form of semanticism when she states that Gibsonian affordances are more than likely representational (ibid. 1995: 191, 2005: 174), however. These representational icons have a double purpose as they both lock onto a feature of the environment and, further, elicit an appropriate response to this tracked feature[157] ("pushed" from within and "pulled" from without[158]).

[156] Another well-known example which illustrates this in a clearer light is ring-dating the trunks of trees to determine their age. Their age may be represented from within *by inference* as extrapolated from our knowledge of how the seasons effect trunk growth, but this is purely content-less data which is not *used* in any assimilable way by the trees. Again, no translation from one state to another occurs, except through analogy.

[157] Or "indicative" and "imperative" features as they are sometimes called.

[158] For example, hens often emit a cry once a food source has been found, effectively letting chicks know there is food in a particular spot, while also imperatively driving the chicks to eat (Evans & Evans 1999). This denotes both the descriptive ("pulled") as well as the imperative ("pushed") aspects of a tracked feature.

Millikan's original agenda in the creation of her theory of content had more to do with the driving need to reconcile Wittgenstein's picture theory[159] with his later work (Millikan 2005: 77). In other words, although it was still an important component thereof, the primary example for naturalising content never had the *primary* intention of being biologically plausible in the first place! Natural selection was used primarily to satisfy the necessary correctness conditions which supervene on the dispositional state of the organism. Hence, she does not invoke her proper functions purely for biological necessity, but through a need to naturalise representational accuracy conditions[160] (ibid., see Chapter Four). For her, from the very basic level of environmental accuracy conditions we can then jump to language use amongst humans (ibid.: 87). If semantic rules govern the natural order of organismic responses, the semantics within language is thought to closely follow. Wittgenstein's later theories were, in her eyes, in need of alteration if they were to provide a more coherent outline for the normative features of content in biological organisms (ibid. 71-72). Therefore, the naturalisation of norms was needed to help her elucidate semantics within language-use. The only way for her to achieve this was to invoke natural selection as the over-arching rationale for maintaining correctness conditions, and by extension truth, through a distinct form of representation of the environment (ibid.: 87). Much like proper functions fix biological intentionality, forms of language are fixed with the environment in an intentional fashion (ibid.: 98). Of course, if we were to take this seriously, we run into problems immediately. To fix the representation of an object accurately enough is problematic for a teleosemantic theory, as the example of the frog striving for its black dot has shown. We are therefore still stuck with Fodor's disjunction problem.

But accuracy conditions need not matter if we move toward a more truly naturalistic approach that concerns itself with slow adaptational attunement to the environment. Moreover, there is something which can be salvaged from Millikan's teleosemantics: the relation between the neuronal state of the organism and the environmental features which it can denote (Pietroski 1992: 268; Hutto & Myin 2017: 114). We can still outline intentionality in terms of a natural relation between neuronal states and the environment, even if "true" representation is not instantiated (Godfrey-Smith 2006: 60). If we take the trusty example of the frog and its eating behaviour, we sidestep the issue of propositional representation if we acknowledge that its

[159] While the details of this need not divert us, it must be noted that Millikan even stated that she went for "picturing themes" in a paper on the topic (Millikan 2005: 77).
[160] Bee dances are her usual go-to in explaining her teleosemantic ideas. She states that bee dances have truth conditions wherein the rules which govern the nectar locations that they denote are fully semantic in nature. They are "designed to correspond to nectar locations as semantic rules" (Millikan 2005: 98).

behaviour has been hardwired by the incredibly extensive evolutionary processes of selection. We do not necessarily need representation to exist in the frog's brain to explain the naturalised behaviour of going for black dots and/or flies if this has been fixed through evolutionary processes. The *reason* for going for that black dot, which would denote a misrepresentation on the teleosemantic approach, is simply because this behaviour has enabled the continual propagation of the frog species in its *normal environment*. So we find that intentionality can still be saved, but propositionality of content can be done away with (Muller 2014: 158), as will be discussed below. I am more inclined to endorse this view, as no full-blown content, unless invoked in a weaker form, is explicitly necessary to explicate the internal processes of perception- and guidance-response. I argue that this upscales to human mental processes; reintroducing the chicken and duck-in-the-dark example from Chapter Two relating to the disjunction problem, we can effectively state that it does not matter if a particular mental state is undecidable with regards to being either a chicken or a duck-in-the-dark. Either can represent mental state M as long as we are suitably guided through the environment.

To replace Millikan's representational "pushmi-pullyu" proposal with a more plausible alternative, Hutto (2008a: 54) uses what he calls "Local Indexical Guides" (LIGs) to describe a process which is non-representational yet achieves successful action coordination in response to environmental features. These LIGs incorporate sensory intake into a vaguely structured order for the organism to be able to respond (ibid.). At this level, these LIGs do not exhibit any capability of incorporating practical know-how (of the overt propositional variety) into their make-up**,** but again, this is not needed.

An appropriate, but tired, example, is that of Millikan's bees, which she uses to explain the visceral responsitivity of organisms toward stimuli. Briefly: bees often perform dances which directly indicate the distance and direction of nectar for other bees to locate. These bee dances do not include semantic information explaining the location of nectar, but instead guide the organisms *directly* to their intended goal (Hutto 2008a: 54-55). This direct guidance occurs through indexical covariance with the environment, utilising the sun and the hive as a coordination mechanism. For example, the information conveyed by the bee dances cannot capture time and place, and they must alter their dances depending on where the sun is relative to the hive[161]. As these dances denote entirely specific information concerning the location of nectar, the bees can respond to this directly without needing further information. The dances

[161] I do not presume that this extends to all species of bee everywhere, but such simplification is necessary here.

do not "say" anything (i.e. they are not symbolic in the way that language typically works), but still point to a salient point of interest within the environment. Extending this line of reasoning, these vaguely structured directions could be related to the Bayesian models of the environment that are an important part of PP. We can see the beginnings of how these models can be the extension of these very low-level LIGs in action. It is possible that nested levels of LIGs could build upon each other and impart upon the organism, even humans, extensive cognitive capabilities when responding to environmental stimuli. By extension, if active inference is to be believed, these very LIGs could operate in such a way so as to instigate the action of the organism, *without* the need for any overt cognitive reasoning. This fundamental bodily responsitivity is arguably all that is needed for much of our more complex cognitive processing of a predictive nature. The environment could elicit behavioural response through direct neuronal linking to environmental cues (after sufficient evolutionary and ontogenetic development), such as immediate action-response to the imminent danger of an oncoming car. REC and PP could go hand-in-hand in their biologically-plausible explication of our cognitive mental states, eliminating semantic theories of representation in the process. As intimated above, a direct coupling between body and environment, thereby constituting minimal content, could be the basis for the active inference of the PP brain. One can see how this could undermine theories of neuronally-based, propositional FP.

Therefore, from the above, we should not assume that any successive action-response routine instantiated in an organism is primarily dependent on representational information being distilled from environmental features. These natural indicators from the environment do not provide representational content but, as explained above, they serve as a *direct* guide for subsequent action (Hutto 2015: 13). Again, one begins to see space opening up here for the applicability of PP within the context of mental "representation" and the way in which it is allegedly used by biological organisms. Active inference within PP is synonymous with this kind of direct embodied interaction. It is highly misrepresentative to state, without any further elaboration, that bodily responsitivity points to true representations of distal environmental cues (Hutto 2008a: 59). We can instead refer to a form of indexical content here. A further point that needs emphasising is that there is no sense in seeking to decompose an LIG into what amounts to imaginary constituent parts. Any change in an LIG would necessarily precipitate a change in resultant action from its end users. This leads Hutto to state that LIGs "guide action in an online and thoroughly context-bound way" (ibid.: 55), in the weaker sense of "online". Any significant change results in a failure to adhere to the LIG's "directive" (ibid. 55-56), but

this does not rule out accidents when these LIGs are operating normally (think of the frog striving for that black dot). Any failures can be explained by mitigating environmental circumstances which do not provide enough indexical information for predictively modelling the world. If predictive models are constituted by nested LIGs, it would necessitate highly context-bound predictive models of the environment. Also, with no determinate content, there is no way for our language to encapsulate the operations of these LIGs without resorting to some form of baseline presumption and "storytelling" of the sort associated with traditional FP[162].

Despite this inherently minimally contentful state of non-representation, we find that LIGs are all that are needed for perceptual reaction to have intentionality and even a normative basis (ibid.: 56). Normativity, here, refers to the baseline guidance that these LIGs exhibit for successful action-response to the environment. It manages to avoid the charge that teleosemantic theories face, which is that they do not provide an adequate conception of how the response of an organism can be misaligned to stimuli within the environment. It also, by extension, avoids the charge that such theories cannot elucidate how the mental can have truth-evaluable content to begin with (Hutto & Myin 2017: 43). This is because, as many argue, there is an important difference between "functioning properly (under the proper conditions) as an information carrier and getting things right (objective correctness or truth)" (Haugeland 1998: 309). Similarly, Stich has stated that natural selection cares more about "reproductive success" than truth (Stich 1990: 62). What teleosemantic theories fail to explain is the content of our representations. The method by which mental states represent a state of affairs is a notoriously intractable problem, and it has not been obviously solved of its disjunctive issues with teleosemantic theories of content. The solution on these theories has more been hinted at than thought through to its conclusion.

Arguably, then, LIGs are more suitable depictions of what actually enables the instantiation of even complex ACRs, as derived from the historic evolutionary environment of an organism (Hutto 2008a: 56). These succeed to sufficiently direct action-responses for organismic survival, without the need for any content, despite the sophisticated resultant actions of the organism. There need not be entirely accurate correspondence to external "facts" of the world when survival does not require accuracy-at-all-times (ibid.: 56-57). Again, using the frogs and their striving for the black dot as an example, accuracy need not come into the picture if the

[162] As outlined in Chapter Two, Hutto (Gallagher & Hutto 2008; Hutto 2004, 2007a, 2007b, 2008a, 2008b) has described the effect that narrative storytelling has on our ontogenetic cognitive development.

path of least resistance, that of simply striving for "black dot-like things", is enough for the animal to survive and propagate in its natural environment. Linking this with the previous chapter, it makes far more sense from an evolutionary perspective for minimally contentful attitudes of the individual organism to guide action, as this minimises the free energy of the system. This is why the concept of covariance is useful here, as covariant relationships are those that maintain their function, regardless of variable changes (within reason) (ibid.: 48-49). Tracking a distal environmental cue need not obtain in the sense that talk of "accuracy" and/or "truth" entails, but rather that the *functionality* of the covariant relationship obtains (Hutto & Myin 2017: 29). In this case, the sensory impingements would eventually trigger the requisite neuronal firings in such a way that the implicit detail derived from the environmental cue and its attendant neuronal populations is *secondary* to the teleofunction that it engenders. Therefore, function is what matters in covariant relationships, not accurate-at-all-times correspondence between environment and representation thereof, which diverges somewhat from the literature on proper functions. Teleosemantics (or biosemantics), as developed by theorists such as Dretske, Millikan, and the like, seeks to preserve a minimal role for semantic representation. But, as I hope is becoming clear by now, we need not invoke content-involving attitudes in order to be able to grasp the fairly basic directed intentionality of an organism. In other words, we need to appeal to a kind of ur-intentionality (again see Hutto & Myin (2017), Chapter Five). As alluded to above, this form of intentionality is a more semantically basic intentionality than that often involed in these debates, hence the "ur-" prefix. This ur-intentionality provides the necessary causal determination, without invoking proper functions, seeing as a semantically-free conceptualisation is enough to derive an understanding of our resultant action. What I have been explaining above concerning basic organismic response-routine is termed *biosemiotics* (Hutto 2008a: xiii, 57), or *teleosemiotics* (Hutto 2011: 335; Hutto & Myin 2013a, 2017), by Hutto, who wants to separate this reworked understanding of organismic response from the content-involving teleosemantic literature. In summary, he states:

> Organisms are informationally sensitive to and selectively end-directed at certain worldly features, objects, and states of affairs in a way that explains their success and failure on certain tasks. The intentionality they exhibit is an attitude of the whole organism expressed in their behavior; it is neither a property of the signs themselves nor of organismic inner states (Hutto 2008a: 57).

He is consistently at pains to emphasise that the relationship between environmental cues and the bodily responses elicited are of a covariant nature, dealing purely with information and not relations of truth or reference. This covariance entails the baseline coupling between our neurons and the environment, thereby attuning the mind toward appropriate cues within said environment. Therefore, the body responds effectively when faced with multiple opportunities for action within a particular context[163]. This is how teleosemiotics differs from teleosemantics, and which provides a theoretical advantage in its more biologically plausible alternative.

But what teleosemiotics lacks as a base theory are its ties with modern cognitive science. It can be proposed, however, that basic organismic response can be predicated upon PP's notion of active inference introduced in Chapter Three. Covariance could perhaps be indistinguishable from how PP is understood to relate the mind to the world. Active inference is an embodied response which could be depicted in terms that need not trade in truth-relations. Covariance relations enable suitable action responses at a fundamental level, as does active inference. Additionally, a bodily response which has been misdirected, which has exhibited inappropriate behaviour in relation to environmental cues, does not simply point to an "incorrect" internal representation that misses the "true state" of the environment (Hutto & Myin 2017: 111). It points more toward an incorrect *covariance* with the environment. In this way, PP and teleosemiotics may be similarly fundamental in their explanation of basic organismic responsitivity. The top-down model associated with PP can be what enables the covariant relationships endorsed by the teleosemiotic framework to have a grip on the world.

To express in language the set of conditional state of affairs that would need to hold for an organism to be successful in a specific action can mislead us as to our theorising on this level. To think of content that is non-conceptual is difficult enough, but we should be cognisant of whether we are expressing existant *true* content from the person's inner mental state, or expressing "faux" content derived from a fundamental environmental covariant relationship (ibid.). We often project thoughts into the heads of organisms to explain successful actions due to our extensive need to anthropomorphise internal processes and, to me, this has been extended

[163] Boxers are a clear example of this. They tend to position themselves and adjust their method of striking the punching bag relative to the space between them and the bag. Without realising that they are doing it, boxers learn to inhabit an "optimal metastable distance" from the bag which creates affordances for action (Bruineberg & Rietveld 2014: 10). This is a clear example of the ways in which experience and learning can alter base biological attunement to the environment, without overt propositional thinking.

to postulate inner (human) mental states which need not exist in the first place[164]. It is comfortable to posit organismic action in the form of propositional attitudes, but this is potentially a vital empirical error. To explain a state of affairs that would need to obtain for successful action-responses could equally well entail an organism not having any true intentional (fully cognitive) *comprehension* of the context in the way traditionally understood (ibid.: 116). Inferring organismic cognition from complex outward behaviour is not as straightforward as it would seem, and yet it could come down to the far more baseline, yet surprisingly elegant, solution of teleosemiotic, minimally contentful, bodily responses to environmental features. Active inference as proposed by the PP framework could potentially provide the bodily responses in terms of which this teleosemiotic account can ground its role in explanations of overt organismic behaviour.

Instead of having to unpack the inherent problems with needing to explain representational content gleaned from non-verbal actions, we can potentially sweep the problem away, and not simply under the rug, but out the door. This is to avoid, quite rightly, the Hard Problem of Content (Hutto & Myin 2013a, Chapter Four). To summarise the conclusions derived from this teleosemiotic approach, we can say that non-verbal action-responses can indeed be intentionally-directed (as well as *causally determined*) despite (1) not having content as traditionally construed; this directed intentionality is (2) a result of the entire organism's response to environmental cues; is (3) a result of the evolutionary history of the organism (Hutto 2008a: 59); and adding to this, (4) can explain the potential for error even if only LIGs are utilised to maintain action-responses, as cognitively-loaded intentionality is more than likely not instantiated at this level. Finally, (5) the top-down predictive models of PP can provide the necessary grip upon the world needed by covariant relationships.

I hope, by now, that it is clear that proponents of the cognitivist view could be highly misled due to, as mentioned above, our tendency to explicate behaviour in terms of concrete and representational mental states which "cause" the behaviour (Currie & Sterelny 2000: 155). Visceral organismic behavioural responses may be sensitive to information derived from the environment, but this responding does not entail the needed structure to integrate said information into the cognitive milieu beyond an in-built "scripted" pattern of response. This pattern of response could suitably be provided by the LIGs, with no need for modules

[164] Currie and Sterelny (2000: 155) make mention of anthropomorphism in belief ascription in their discussion of modularity of mind-reading. They use the example of inanimate objects which appear to possess purposive intentional movements, much like the cellular automata in Conway's Game of Life (originally in Gardner (1970)).

representing content denoting a desire or a belief (Hutto 2008a: 60). Thus, when we describe non-verbal behaviour as "intelligent", we need not justify it upon the back of a mental representation with inherent content (ibid.). We also need not believe that it is the existence of these recombinant representations, once amalgamated together, that provide the ability to formulate more sophisticated intentions and planning within non-verbal human cognition. As mentioned before, this form of sophisticated intentional thinking, propositional in nature, is arguably only introduced through the introduction of linguistically-scaffolded learning (ibid.: 60-61). Therefore, what has been described above can only be up-scaled to more sophisticated thinking in line with everyday human cognition once a form of *enculturation* has occurred. Further detail on this will occupy us further below.

I have been arguing that some if not most cognitive processes do not necessitate propositional thinking. Due to impingements on the senses having minimal informational content, as dictated by the teleosemiotic approach, there is no need for this information to be "attained", "transformed", and then "deployed" in subsequent operations. Positing semantic content does no additional work in explaining how cognition occurs. This should remind us of what Dennett proposes occurs in the brain, with our propositional style of thinking being more of a "user illusion" (2013b; 2017) coalescing into a fictional narrative of coherency. Our "thinking" selves are the result of those underlying processes which give rise to our emergent phenomenal consciousness. If we accept his theory, wherein the mind operates in a less determinate fashion than typically thought, we can incorporate a teleosemiotic approach to this view of the mind as a "fuzzy" processor[165]. Indeterminacy need not be an issue if the brain, and by extension the mind, operates on a need-to-know (or hardly-need-to-know?) basis with regards to its action-response routines. To elaborate, the brain needs to incorporate information from the environment only insofar as it is useful in tracking given events in said environment, and furthermore, the detail of this sense data need only be accurate in terms of covariance with the environment. This means that our experience of the world from within our phenomenological vantage point is always a truncated "representation" of the world. As a result, tracking sensitivity toward distal objects should be thought of in terms of *sensitivity* to *information*, not *acquisition* of *content* (Hutto 2008a: 61-62). Again, we need not assume double transduction of content into something "useable" if our responses are triggered by environmental cues, and we certainly need not assume a representation of the environmental milieu whereby the senses

[165] This denoting something which does not deal with computing accuracy so much as computing that which is salient for the present purposes of the organism. The "fuzziness" alludes to its indeterminacy.

incorporate content which is then transformed into a useable format, comprehensible to thought (as endorsed by Fodor).

We must not confuse ourselves by talking about information "contained" within a signal, instead of the more plausible explanation whereby information is frugally gleaned through the effects engendered by covariance. Along these lines, Hutto (ibid.: 63), in a move which has many affinities with the PP approach, makes the important and oft-overlooked point that we must not confuse ourselves by assuming that perception and action are always distinctive operations within the brain. Input-output mechanisms within the brain that rely on full-blown representation as their end product need not exist if immediate response can be initiated. Ratcliffe (2007b: 192) also makes the point that action and perception are inextricable[166], and we needlessly deconstruct seamless action along imaginary joints—content-involving and drive-motivating (very roughly: beliefs and desires). Again, if we go back to the PP framework and its central component of active inference, we can see that the current scientific literature supports this view with immediate action-response predicated on the inner predictive power of the brain.

The action-response routines instantiated by worldly cues do not rely on a "central observer" (Dennett 1991a) in order to function, as these routines would have evolved at a time which pre-dated any thinking, let alone belief, which itself requires language (if the above picture is correct). Basic action-response routines are, by extension, opaque to explication through language. They are instantiated without any representational thought, which serves the brain well, as it does not need to wait for the outputs of upper-level cognitive processes in immediate environmental tracking (Clark 2016: 179). As proposed in Chapter Three, the predictive capacity of the brain to process worldly cues is due to pre-modelling of the environment throughout the development of the organism, with only alterations to the environmental mapping that the brain already possesses. In a similar vein, Hutto almost has it right when he states that:

> [Action-response routines] kick in without any background thought. And, surely, this is a good thing. If creatures had to wait for the top-down verdicts of central cognition before responding in the simplest cases they would be effectively crippled. Indeed, without some such bottom-up responding to get a grip on worldly offerings, there would be no way for

[166] Ratcliffe (2007b: 192) asks us to imagine an explanation of someone playing tennis, who "believes" that the ball is approaching, and subsequently "desires" to return the ball. Can it really be said that these are two discrete preceding components of the resultant action? This would need to be interrogated a lot more before we could be so sure of our assumptions on this point.

> "central cognition" to judge, on reflection, which considerations are relevant to particular cases (Hutto 2008a: 63).

As he did not take into account the existence of PP when writing this, this is only partly correct. We cannot expect the brain to at all times wait for top-down responses before reacting on environmental cues, just as we need bottom-up signalling to alert the system to the same. But PP presents us with a view of the brain whereby *minimal* environmental cues are fed forward to impress upon the already extant cognitive model of the contextual environment. This means that it could fit the minimal content agenda that Hutto advocates on his teleosemiotic construal, while also allowing for the immediate, sensitive bodily responsitivity that needs to be accounted for. This could be a rare case of having our cake and eating it too, as there are important roles to play for both the top-down and the bottom-up responses of our brains, with top-down models of the environment being the predominant cause of action which provide the very shortcut that Hutto seeks. It is a neat package which could fit into his teleosemiotic and covariant picture of mind responsitivity rather well[167], along with tying it all into the literature on action-oriented PP.

But what does this mean for modular propositional thinking? LIGs, as contextually-bound bodily dispositions, are potentially all that are needed for even complex behaviour, even up to and including *social* behaviour. As expanded upon in Chapter Two, it has been shown just how little we *actually* use propositional thinking in our day-to-day living. Ratcliffe (2007a: 233-240, 2007b, Chapter Seven) goes further than this to show the lack of coherency in the general FP framework, elucidating quite starkly the proliferation of background assumptions we all employ on a daily basis (ibid.: 187-197). His description of the everyday cognitive capacities employed while enacting our FP abilities points to a lack of fundamental reasoning in our daily interactions. To reiterate what was said previously:

> The norms that feature in explanations of behaviour do, of course, include norms of good *reasoning*, in addition to situational norms. However, there is a difference between reasoning and having a reason (ibid.: 195).

[167] Although, not wanting to cede defeat, Hutto and Myin attempt to pour cold water on this "bootstrap hell" (Hutto 2017; Hutto & Myin 2017: 67-74**)**, despite their affinity for the PP paradigm in general. They use the term "bootstrap hell" to point toward the difficulty in constructing models of the environment from the ground up ("bootstrapping") through continuous sensitivity to representational content, which is what Clark proposes (2016: 19). Clark's distinctly more conservative approach fails to deal with the HPC (Hutto 2017: 10).

In other words, we are good at rationalising behaviour, but this does not mean that this picks out the main reason for a particular behaviour (refer to the example in footnote 172). Reading his work, I am drawn to propose that this insight could be extended to the forms of mental states that are picked out by much of our folk psychologising, with these mental states predominantly being *minimal* in content. The implication for this thesis being that, regardless of the result of a person's behaviour, there need not be any content-inherent (fully propositional) processes of reasoning that have brought this behaviour about. Ratcliffe supports Dennett's position on indeterminacy of beliefs[168] (ibid. : 205-211), which is a similar position to that which led Needham (1972: 125) to describe our use of the term "belief" as a "peg word" which takes the place of uncountable dispositions[169]. Through similar reasoning, Hutto (2008a: 60) denies that there need be an assumption of a basic architecture inherent to non-verbal action-response. Nichols and Stich (2003: 15) have described this basic architecture as an assumption made by theorists on the basis of the mind's representational states; i.e. that it contains the two distinct states of belief and desire. From what has been argued so far, it should be clear that I do not endorse this basic assumption either, primarily as a result of the picture of the indeterminate borders of our neuronal states, driven by probabilistic processing, proposed by PP. To simplify, if our neuronal states do not accurately represent the environment, then these neuronal states do not clearly relate to propositional beliefs and desires, and are more fundamentally dispositional in nature.

Taking a cue from Clark (2016), Dennett (1991a, 1995, 2013b, 2017), and others, it is clear that the ability to represent a coherent picture with regards to the world requires externally-derived props which have only recently been afforded to humanity in the form of our natural languages. While it is difficult to prove with any degree of accuracy, they propose that language is that which has afforded our minds an infinitely powerful tool to conceptualise and imagine our world, leading to the type of consciousness we all have. This brings us to the underlying point concerning the invocation of propositional attitudes within theories of FP. Traditionally, theoretical FP constructs a framework for embodied, non-linguistic action-response on the basis of our *linguistically-facilitated thinking* (Hutto 2008a: 73). It is thought that for non-linguistic action-responses to be content-involving, there must be roles played in our mental

[168] Dennett's (1987) theory of the "intentional stance" states that in order to read and/or predict behaviour, we take a detached stance toward a person or object rather than actually reading their mental states. This detached stance consists of *inferring* information about the person's or object's dispositional state. It is a form of neo-behaviourism.

[169] Needham (1972: 131) also states that belief "…can be so far translated away that the concept must disintegrate in correlation with the dispersal of its connotations". I am inclined to agree.

states that are similar to the roles that concepts, beliefs, names and others inhabit in our languages (ibid.: 74). It is assumed that there must be a counterpart which fulfils what amounts to the same function at this mental level, namely, a symbolic mental language, as proposed by Fodor with his LoT. As Hutto (ibid.) rightly states, the analogy falls apart here. It has been difficult enough to argue for a system of conventions that manages to "fix" name or concept use in natural language. And to utilise a similar strategy for what I argue to be a fine-grained, embodied, predictive "know-how" at a mental level appears to be fruitless. An appeal to direct functional processes holds far more explanatory power than traditionally-conceived ToMs. With a lack of a structured method of instantiation, the proposed mental symbols cannot co-refer in the way that names and concepts are capable of within language discourse, as Millikan (2000: 166) would also attest.

I propose that the "neat and tidy" approach of something like a LoT provides a dilemma which both over-complicates and over-simplifies what is occurring in our minds. The first horn of the dilemma is due to the over-complication derived by adding unnecessary functional processes to our minds which we need not have evolved a use for. Drawing from our discussion in the previous chapter concerning the minimisation of free energy, we must be wary of over-extending the bounds of what I call "least resistance processing". The frugality of the mind dictates that the path of least resistance will almost always be taken, arguably leading to covariant relationships between brain and environment, instead of necessarily contentful representation thereof. The second horn of the dilemma is due to the over-simplification of theories like the LoT. These theories appear to posit mental processes acting as decoupled instantiations of a sentential nature a little too easily, without any truly coherent framework on which to base these suppositions beyond armchair reasoning. We tend to gravitate toward explanations which make sense to our lived experience, which perhaps goes some way to explaining talk of not only a LoT, but the far more pervasive use of propositional attitudes in our discourse surrounding the mind. But if language-use is necessary for truly propositional thinking, perhaps we have not been utilising propositional attitudes for the purpose of practical reasoning as long as many have surmised within our evolutionary history. At least with regards to animal psychology, this is a consensus which appears to be growing, with Akins agreeing that much of the simple behavioural response of organisms could be explained by the operation of "narcissistic" systems (Akins 1996: 344-355), as opposed to representational concepts

within the mind[170]. Papineau (2003: 99) has stated that we cannot assume that more simplistic organismic behaviour takes the form of "means-end reasoning". Andrews (2008; 2012; 2015a; 2015b) has also recently advocated a view on FP highly similar to Hutto's from her perspective within the field of animal psychology. There is no empirical reason not to extend this approach to basic human minds. Non-linguistic (early) humans and some animals can be said to engage in a rudimentary form of FP, but this is not the propositional kind that we usually attribute to ourselves in our overt reasoning.

I am picking apart the use of propositional attitudes in explanations of non-linguistic action-responses in order to show that if we reject this, we are left with a far more useful and frugal explanation for behaviour. If we posit that the mind behaves in a manner that precludes the need for coherent propositional attitudes, we are some way to embedding our theory of mind into its evolutionary history. The content that informs representations of our environment serves no additional purpose in basic FP. But even if this is accepted, the teleosemiotic theory of action-response routines which downplays the propositional forms of thinking will have its detractors. What about the more sophisticated feats of FP that we are all capable of? To finalise this depiction of the non-propositional nature of the basic processing of our minds, and tying it up with its social aspect, this chapter will finish with the role that enculturation arguably plays in the development of full-blown propositional thinking. Denying propositional content in basic mental processes does not necessarily mean denying propositions across the board. It will be shown below how both the PP and the teleosemiotic processes discussed here underlie the eventual enculturated aspects of our learnt narrative competency.

4. Predictive Processing, Teleosemiotics, Social Scaffolding, and their Effect on Extended Folk Discourse

How does the preceding tie up more strongly with PP and its probabilistic account of the mind-world relation, as discussed in Chapter Three? Furthermore, while the preceding potentially explains indexical world responsitivity, how does it enable the fleshed-out propositional thinking that we all appear to employ? If the PP account of the mind is correct, if prediction error minimisation, and more fundamentally free energy minimisation, is the reality in need of

[170] An example of a "narcissistic" system can be that of our temperature gauging of the environment. Our thermoreceptors do not record the temperature (an accurate depiction of the environment), but what the temperature has evolutionary meant *for the organism* (via a covariant depiction of the environment).

explanation, we can see on the horizon a fundamental shift in our understanding of not only our allegedly propositionally-laden phenomenological experience, but also our social *co*-cognition. If PP were fundamental to our cognitive functioning, what does this say for our interaction with others who are operating on similar neuronal frameworks as ourselves? What does this mean with regards to our FP? These are questions too large for this dissertation alone, but it is worth formulating a general hypothesis of our cognitive functioning, incorporating a scaling up of perspectives from the fine-grained and the embodied theories (teleosemiotics and PP), to personal development (NPH), and beyond, to the co-cognition which the preceding arguably affords. This would naturally affect our thinking concerning FP, both in terms of how it is thought to function as a ToM as well as how we think when attributing reasons to the folk. These two barbs of folk psychological practice should be separated as there could still be a place for an aspect of FP within the PP fold if we acknowledge that FP, as understood by the *actual* folk, is not always *explicitly* formulated when attempting to explain behaviour (Dewhurst 2017: 10). Perhaps there is a place for a form of predictive FP in our cognitive makeup, but would this necessarily entail some form of propositional thinking?

While I have been deflationary with regards to folk psychology, one cannot deny the *existence* of folk psychology. However, it is its epistemological reality which concerns me. Taking the potential truth of both a teleosemiotic and a PP approach to cognition into account, we cannot easily maintain the view that our FP picks out a cognitive reality. We have seen that if we do not accept the reality of propositional attitudes as necessarily fundamental to our neuronal states, we begin to see that scientific (and/or philosophical) issues concerning content-acquisition, cognation of this content, and others, falls away. If the embodied mind evolutionarily acquires shortcuts (the frugal aspect of mind), assimilates them into its framework, and continues to utilise them in order to minimise free energy, many errors and environmental/social misfiring can be predicated on the activities and proclivities of the cleverly lazy brain (Clark 2015c: 9-12; 2016, Chapter Eight).

With the rise in popularity of PP theories, some theorists have begun to tackle the implications arising from the relation between PP and our socio-cultural environments. Fabry (2015, 2017, 2018) envisions complementary roles for PP and our developmental socio-cultural practices. More specifically she utilises the concept of *enculturation* as the primary complementary tool to the underlying PP framework, whereby socio-cultural practices support our predictive brains in particular ways. Furthermore, this enables our brains to be representational in the propositional sense. I see this as another, yet broader, formulation of what Hutto has written

concerning narratives and their fundamental role within ontogeny for the development of our folk psychological abilities. Enculturation, in other words, can be developed as a broader formulation of the NPH in that our socio-cultural practices inform our cognitive capacities. This extends to how propositional thinking becomes introduced into our cognitive abilities. Despite the fact that she does not focus on folk psychology in any overt way, I believe the foundation which she lays can be expressed as an up-scaled, general explanation for our folk psychological understanding. This would neatly complement (or, more accurately, lend empirical validity to) Hutto's views as well. Enculturation can be the means by which our brains are able to enhance our representational abilities.

Arguably, as the minimisation of prediction error and the ongoing processes of precision estimation are improved upon throughout an agent's life, the generated prediction models become more defined and therefore more accurate with regards to the agent's environment. We can therefore state that the ongoing process of prediction error minimisation delineates a vector through which the agent, throughout its existence, develops (Fabry 2017: 4). Fabry (ibid.) calls this continual updating of predictive models over the agent's lifespan, and primarily during ontogeny, the "predictive acquisition of cognitive capacities (PACC)". These capacities are what enable suitable and effective environmental interactions, and are formed through the frugal processing of the predictive brain. The computational system that the brain employs is an economical evolutionary trick which better serves the need for continually updating the predictive models of the environment. PACC and the acquisition of resultant cognitive practices are Fabry's primary focus, with cognitive practices denoting evolutionarily recent cognitive abilities such as reading and writing. Cognitive practices can also be defined by what Roepstorff *et al*. (2010: 1-2) have called "patterned practices". They are those practices which are utilised by many agents within our *shared* cognitive niches, the norms of language being a prime example. These practices are acquired and streamlined throughout the extended experience of the individual, a point often emphasised by Clark (2016). PACC, on the other hand, is achieved through enculturation specifically (Fabry 2017: 7). Enculturation is therefore the dynamic, tacit influence on neuronal plasticity that results from being embedded within a particular socio-cultural niche. This can include fine-grained changes in precision estimations as a result of what the agent is exposed to during ontogeny. Those who argue that enculturation is central to the acquisition of our higher cognitive abilities, such as Fabry (2015, 2017, 2018), further hypothesise that enculturated cognition requires us to commit to an understanding of the inner workings of our cognition as dependent upon the collaboration between our embodied

selves and our "cognitive niches". We can define the cognitive niche (somewhat pretentiously) as the contextual reality delineated by "the incrementally, trans-generationally structured socio-cultural environment that provides human organisms with epistemic resources for the completion of cognitive tasks" (Fabry 2017: 7-8). Epistemic resources are those that enable extended (and whichever other E-terms one wishes to include) cognition in the form of systems of representation and tools. Some simple examples would be the use of a calculator, or even the structured nature of schools.

The "patterned practices" mentioned above are objects of public consumption, much like Hutto's narratives[171]. "Patterned practices" can be seen as a generally defined concept under which the NPH can fall, as our cognitive practices are constrained and guided by inherent cognitive norms (see Menary (2007, 2010a, 2010b, 2013, 2015a, 2015b) for more on this view). What this means for the NPH is that particular folk narratives have been formulated for others to publicly consume and effectively communicate with. While more can be done to significantly align Hutto's theories with the literature on PP[172], we can see shades of it in Fabry's (2017: 8) formulation of the "normative aspect of enculturation", whereby the norms inherent within our cognitive practices are regulated by epistemically salient public markers (such as the communicative symbols of language). Therefore, public objects of consumption (such as symbols) are used to regulate and maintain the norms inherent within our cognitive practices. These norms are in a sense the "programming" that we acquire during development. Clark himself has pointed to the cultural processing of our cognitive practices within our socio-cultural environment, leading to his formulation of "scaffolded" dissipations of the cognitive load into culture[173] (Clark 1997: 179-180). Broadly speaking, the concept of socially-derived scaffolding denotes any set of "augmentations that allow us to achieve some goal that would otherwise be beyond us" (ibid.: 194-195). A simple example is that of using a pencil and ruler to draw an accurate line, but examples also extend to how people assist others (think of a child learning to walk). Therefore, social scaffolding can be seen as the structuring of our cognitive practices in a systematised procedure of "novice-expert interaction in the cognitive niche" (Fabry 2017: 8; see also Fabry's (2015) commentary on Menary's (2015a) paper on the topic). The "novice" in this interaction acquires the cognitive practices (such as learning a folk narrative), and the "expert" imparts the necessary skills within a particular domain (such as the

[171] Which were introduced at the end of Chapter Two.
[172] Although see Hutto (2017) and Hutto & Myin (2017: 57-66, 82-83, 150-163) for brief comments on the growing body of work supporting PP and its relation to their view on content.
[173] This is the broadly conceived functionality of his EMH in action (Clark & Chalmers 1998).

epistemic saliency of public symbols used within folk narratives)[174]. Fabry (2017: 8) calls this the "scaffolding aspect" of enculturation, and Hutto's notion of narrative development can quite easily be seen as a form of scaffolded learning in its depiction of the normative reinforcement of learnt folk psychological understanding. Hutto has indeed alluded to scaffolding in his writings (Hutto 2008a: 183, 224; Hutto & Myin 2013a, 2017, Chapter Six; Hutto & Satne 2015).

According to Hutto, human behavioural responses to the contextual environment occurs through non-accidental dispositional attunement toward distal cues, including others' intentional attitudes (Hutto & Myin 2017: 140). This view sees cognition as weaved into the patterned practices mentioned earlier. As Clark states:

> …higher cognition…is made possible only by our history of encounter with the increasingly exotic sensory flows created by our own culturally crafted 'designer' environments (Clark 2016: 138).

The effect that scaffolded learning and the norms that it engenders have on our cognitive practices can be profound. These cognitive norms can include the learnt targeting of specific intentional behaviour of others. As Sterelny (2003, Chapter Four) has argued, the ability to be attuned to others would require the gradual evolution of specific mental states in order to handle the complexity of social interactions. A clear example of this would be the learnt folk psychological dispositions which we all accrue over time. But this requires the ability to cope with multiple social triggers across extended temporal scales. On the view presented here, the social and physically embodied environment provide the norms which are subsequently assimilated to delineate "the content of environmental vehicles and how we manipulate them" (Menary 2010a: 229). With the help of cultural practices, our minds learn to attune to salient points of information gleaned from the environment, such as what a particular facial expression denotes, and what an appropriate response would be in light of this. Of course, if the teleosemiotic account of cognition is to be believed, it would render talk of content and vehicles superfluous to the baseline processing of the mind. Nevertheless, important for our purposes is that idea that these norms restrict our interpretation and subsequent manipulation of epistemically-salient environmental cues (Menary 2007: 136-141). Scaffolded learning enables the acquisition of such norms, which adjusts the ability of the agent's cognitive system to incorporate epistemic resources in service of these acquired cognitive practices (Menary 2010a;

[174] Sutton (2015) calls these developmental niches "cognitive ecologies".

2013). Interestingly, the interrelated components of cognitive norms are assimilated and adapted over time, resulting in these components shifting "from being entirely explicit and context free to being entirely implicit and embodied" (Menary & Kirchhoff 2013: 618).

This form of "content learning" occurs along a constrained pathway of development, which means that some capacities are only possible once others are already set up to enable them. For example, educational learning can set in place specific ideas which can be built upon to enable finer modes of cognition. This path-dependent learning is an important component of the aforementioned scaffolded learning. Clark (1997: 205) emphasises its importance by stating it implies that "you can't get everywhere from anywhere, and where you are now strongly constrains your potential future intellectual trajectories". Path-dependent learning, as the constraining effect on the learning process, provides the vector through which scaffolded learning unfolds. And, the inner processes of PP are ideally placed to provide support for this view of the complex interrelations of path-dependent learning (Clark 2016: 288). As Fabry (2017: 10) states, the systematised and temporally extended exposure to socio-culturally embedded epistemic resources during scaffolded learning (or in Hutto's more truncated paradigm: narrative learning) is corralled by the fine-tuning of precision estimates. On the PP view, the continual updating of predictive models of the environment is constrained by the particular cognitive resources available. These changes are shaped by the "functional biases" of the brain as well as body morphology (Fabry 2017: 7). For a clear and simple example, the physiological and anatomical structure of our arms and hands constrain the rolling prediction error minimisations necessary for proficient writing (Phillips *et al*. 2009: 585-586). The brain and body are only capable of fine-tuning its prediction within the straightjacketing effect of bodily limitations. Embodied active inference is therefore biased within these constraints, which complements the predictively-derived plastic changes within the brain to shape subsequent cognitive function acquisition. Likewise, within folk narrative learning, we acquire narrative competency only insofar as we are able to incorporate it into our existant cognitive makeup. This speaks to an embodied form of social cognitive capacity acquisition, which guides action in step with brain plasticity.

And what about the arrival of propositional-style thinking within our mental lives? These arrive on the scene through what Ramstead *et al.* (2016) call "cultural affordances". Cultural affordances are salient markers that help to predict the behaviour of others, and are therefore of vast import when navigating shared social spaces. Furthermore, during ontogenetic development, these cultural markers are *internalised* through repeated "social interactive

inference" (Metzinger 2017: 20). This points toward the possible gradual development toward more and more complex propositional-style thinking (Ramstead *et al*. 2016: 17) and the eventual formulation, usually around two to four years of age, of the autobiographical self[175] (Metzinger 2013: 5). This is in order to arrive at more fine-tuned models of the social environment, predicated on the successful navigation thereof. Only in this way is the much-vaunted content that is needed for a full-blown FP, in the form of propositional thinking, acquired. It is *this*, as Hutto and Myin have also pointed out above, that gives us any semblance of "full-blown" representation. Hence, we should see propositional attitudes as the enculturated artefacts of public discourse, and not instantiated in our neurons in their fully propositional form. These artefacts develop over time due to their useful benefits. PP and our teleosemiotic processes also have a profound effect on our culture, in that they shape our folk narratives, and by extension our propositional attributions (remember that our cognitive niches are also constrained by our inherent cognitive abilities). In turn, narratives serve to limit the behaviour of individuals "so as to make themselves more easily predictable [to] other agents" (Clark 2016: 286). This is important as it limits the otherwise debilitating plethora of potential causes from which behaviour could follow, providing other individuals with a foothold from which to infer actual causes. This does not speak to the need for ontological accuracy so much as pragmatic markers that enable successful inferences.

The picture painted above is fairly indistinguishable from Hutto's formulation of how narratives work within his NPH. There appears to be a growing consensus toward construing folk psychological discourse more holistically than as merely reflecting actual instantiated propositions within the brain. Naturalising folk psychological discourse by incorporating PP and teleosemiotic theories into our understanding of FP does not eliminate our practice of folk psychologising, as discussed in Chapter Two in relation to Churchland (1981). Rather, it entails the broadening of the traditional notion of our FP capacities as consisting in fundamentally content-rich representations at a neuronal level that cause us to act. Therefore, folk psychological discourse and the narratives which enable it could still be a fundamental component of a fleshed-out theory of cognition. The propositional thinking which we all employ is an important part of our reason explanation, and seems to have developed in such a way as to be parasitic upon our neural architecture, such as when we explicitly state "I believe that…." or "He desires that….". A tokening of linguistic propositions enables us to get a grip

[175] This more than likely contributes significantly to our inability to recall our lived experience at younger ages than this.

upon the world, but it puts a realist gloss on our actual neuronal interactions. As mentioned before, narrative and social scaffolding (or niche creation) are highly important aspects of the cognitive picture (Clark 2008; Clark 2016: 275-288), and folk psychological discourse can be depicted as a form of the very same (Dewhurst 2017: 9). This means that we can still find a use for it within a future theory of cognition. Folk psychological discourse, through the social norms which serve as the regulative principles for human behaviour, could also be understood as a macro version of active inference (ibid.). What we are left with is a more anaemic depiction of FP which can still act as a pragmatic framework for reason explanation, even if it does not depict the reality of our cognitive processing.

As Hutto and Myin (2017: 134) as well as Hutto and Satne (2015) note, the grasping of public symbols is necessary for proficient development and sensitivity toward practices of an intersubjective nature (from basic communication to the most complex of multi-agent projects). They also make the link between this and the growth of cognitive niches within humanity's sociocultural environments (Hutto & Myin 2017, Chapter Six). Sociocultural practices, as well as their continued maintenance, are the primary enablers for the prevalence of contentful cognition (Hutto & Myin 2017: 134). Basic minds[176], including our own, are only capable of specialised content-involving social interactivity once these sociocultural practices (such as propositional-style thinking) are incorporated into our cognitive make-up. As this sets contentful minds apart from other forms of minds, this truly does separate these minds from what has come before or since. They stress that this difference can be seen as one in *kind*, not degree, between humanity and the rest of the animal kingdom[177] (ibid.). Publicly available symbols supervene upon the basic structure of our minimally contentful minds[178]. Essentially, our basic minds utilise public epistemic markers to accentuate our basic functionality toward full-blown content-rich cognition. The details of how the mind acquires these abilities operationally is an open question, but one worth exploring.

From all of the above, one can see differing levels of functioning at work here. The fine-grained processing on the PP level, the teleosemiotic environmental attunement, and structured environmentally-discriminative path-dependant learning. I propose that the co-cognitive level of interrelated folk propositionality, as described above, is realised only after all of these

[176] Recall that a "basic mind" does not denote a "lesser mind".
[177] They in no way advocate a "special" place for humanity in the evolutionary game, but merely point to the unique shift in humanity's cognitive capacity as a wholly other form of consciousness.
[178] Something akin to what Dennett has been saying for quite some time (Dennett 1991a: 171-208, 1995, 2017). See Churchland (2002) for a brief critique of his views, and a rebuttal in Dennett's (2006).

components suitably interrelate. As the theory of PP has increased in empirical validity, we are in need of upscaled theories of its basic functionality at higher levels of cognitive functioning. The teleosemiotic account could explain how we acquire contextual features for our ever-updating (non-propositional) models of the environment. I believe that the underlying PP account of the brain, and by extension the embodied theories of the mind, could go even further to account for the behaviour and bodily dispositions of people in their social environments. Just how this would work is up for debate, but to elucidate how folk psychological interaction would work within the structure of PP would include a tale of abstract, high-level predictive models of behaviour which include the agent *as well as* the socio-cultural environment within which the behaviour occurs (Dewhurst 2017: 9). Dewhurst points to the supplementarity that particular learnt narratives[179] afford basic behavioural prediction (ibid.: 7), as does Fabry (2017: 10), as discussed above. This is more than the regular prediction of subsequent behaviour using *current* behavioural markers, but is instead a modelling of the context of the behaviour utilising already-learnt normative explanations of behaviour acquired over time. For example, knowing what a colleague's "belief" within a particular situation is, is informed by learnt information about past behaviour, which in turn attunes one's cognitive attention toward particular, salient behavioural markers. This differs from traditional FP, as behavioural regularities are all that is needed for what I would call an embodied folk understanding of the contextual environment. One can see how this can align with talk of PP and narrative structures within our FP. I call this iteration Dispositional Folk Psychology (DFP), as it moves beyond talk of overt mentalistic language toward a visceral, embodied understanding *that is nevertheless normative*. As described earlier, we navigate our everyday environment in an embodied fashion, but our embodied interaction with the environment also includes the social, which necessitates understanding norms of behaviour. This embodied interaction is a somewhat structured, predictive framework which does not incorporate overt language-use but nevertheless positions the embodied mind to react appropriately toward these distal social cues. Any learnt social cues, such as body language, are incorporated into our cognitive models, which allow us to minimise unnecessary overt processing of social cues, thereby allowing us to predict behaviour in a thoroughly immediate and visceral fashion. The shift toward implicit and embodied cognitive norms as discussed above (Menary & Kirchhoff 2013: 618) enables a

[179] Dewhurst does not use the word narratives in this particular context, but I argue that one can slip this term in and not lose anything significant from his original point. He does appear to state that there is more overt reasoning occurring during these everyday occurrences, more than I contend is possibly happening at any given moment, but any difference in our opinion is purely conjecture at this point.

use of cognitive resources which is frugal and adaptive to appropriate environmental stimulations.

In my eyes, DFP is what is left when the ascription of propositional attitudes are off the table, for most of our cognitive processes. One will recall that propositional attitudes as our primary means of folk understanding is an outmoded depiction of our neuronal reality. As DFP is a predominantly dispositional[180] experience of the environment, there is space here for an understanding of our base folk psychological behaviour as a minimally contentful form of prediction. What I am proposing is that DFP predominantly occurs throughout PP brains incorporating cognitive norms through teleosemiotic covariance[181]. This results in an updating of our predictive models while moving beyond an obvious need for vehicles of content. In addition, our predictive framework provides an embodied shortcut so as to operate according to the strictures of least resistance processing, restricted by the information-theoretic free energy of our cognitive system. These are not explanations or predictions as such, but more akin to "recognition-response patterns that generate embodied expectations" (Hutto 2007b: 44). In our continually evolving experience of social discourse, the consistent need for "online" reasoning when reacting or "understanding" social cues is done away with. Incorporating our knowledge of brain processing, it appears that consistent "online" reasoning would be far too burdensome a drain on our cognitive apparatus, and, as explained earlier, our brains have been evolutionarily jury-rigged along the path of least resistance. On this view, failure to predict behaviour is a prediction-error-in-waiting, which would be fed upward to update our looping predictive model of the social milieu. This need not be "content-involving", as an already extant predictive model of social understanding is partly acquired throughout the lived experience of the individual (which would include narrative learning). This model is already dispositionally embodied, but in need of occasional alignment in light of any failures of prediction through active inference. As DFP entails environmentally-coupled dispositions toward social cues, what can be seen as the "understanding" of behaviour could be more of a somatic, yet still emotionally affecting**,** experience, intimately associated with and elicited by the circumstances themselves. In other words, our experienced "intuitive" understanding of others can perhaps be seen as the emotional by-product of this immediate action-response. On top of this, folk psychological narratives inform the overt propositional attitude ascriptions which we apply after the fact. All of this brings our theories of FP in line with current theories

[180] Recall Ratcliffe's (2007b: 205-211) notion of beliefs as indeterminate dispositions.
[181] Covariance was described in detail in Section 3.2 of this chapter.

of cognitive processing, as discussed above, and hence makes it much more plausible than traditional theories based on armchair speculation constrained by linguistic commitments.

Regardless of which model we eventually arrive at, I argue that the most plausible theories of the operation of our brains lean toward a PP view of our internal cognitive processes. Simply put, if PP is on track to be the comprehensive answer to the above, its central notion of free energy minimisation should influence our embodied minds even more than is currently thought. The teleosemiotic account of our cognitive capacities which emphasises a propositionless, content-minimal mode of thinking, is an important upscaled expression of this frugal minimisation of energy. To function in the contextual environment, minimising free energy would make sense from an evolutionary perspective, which in turn leads to the minimally contentful responses of the embodied mind (as argued above). Therefore, PP and the teleosemiotic account of cognition go hand-in-hand as complementary theories that accommodate a fundamental biological need. Taking this further, enculturation (including the influences elucidated by the NPH) could be seen as an important component of our neural development, which extends into the embodied long-term existence of the agent. The cognitive load that differing types of narratives support, relates to a frugal use of cognitive processes which, again, minimises free energy for temporally-extended survival within the cognitive niche of the agent. This phenomenon supports the introduction of cognitive norms into the process, which provide further frugality in their ability to constrain behaviour along well-worn (socially-attuned) cognitive paths. This, I suggest, is the positive aspect of folk psychological discourse: its ability to guide us along paths already trodden, providing opportunity for free energy minimisation while "downloading" a form of coded norm programming. In other words, through social learning, folk terminology is gradually incorporated into our very baseline interactions with our environment, thereby bypassing to a great extent the need for overt folk reasoning. Therefore, propositional content predicated on truth evaluability at the neuronal level need not occur.

5. Conclusion

This chapter has introduced a minimal approach to representation within our cognitive processes, specifically when it comes to processes relating to FP, augmenting the PP picture of our minds presented in Chapter Three. Its purpose was to be deflationary with regards to mental content-acquisition and processing, and streamlining the already frugal interactions that a PP

theory of mind provides. A brief run-through of the literature was done, which served to highlight the shift toward minimal content when describing cognitive processes. This in turn made it possible to ground my affinity for non-representational cognition and the subsequent problematisation of the idea that semantic content is the mark of the truly cognitive. Thereafter, I introduced the teleosemiotic account of the covariant relationship between body and environment to lend theoretical weight to a more visceral blending of mind and environment than is typically the case. Thereafter, this discussion incorporated the literature on enculturation and scaffolded learning in order to construct a theory that joins FP, PP, and teleosemiotics. In conclusion, a view of frugal, least-resistance processing when it comes to our base folk psychologising was explicated to tie these differing but closely related fields together.

From the above it should be clear that I do not accept the reality of actual propositionally-laden folk explanations of behaviour, except in their minimal sense, and beyond their overt use in our linguistic practices. They are merely a means of what could be called "niche survival", where this niche happens to be the vast socio-cultural milieu which surrounds and envelops us. FP as propositional attitudinal thinking speaks to a truncated understanding of our minds, far too reliant on armchair reasoning to be helpful when it comes to explaining what occurs at the level of the mind/brain. It is easy to say that it is intuitively plausible that propositional attitudes are what lie behind our reason-giving capacities, but it can also quite easily be said that we usually do not utilise an overt reasoning process when we understand others' behaviour. In other words, it is highly plausible that the majority of our daily cognition does not take propositional form, due to our embodied and embedded existence within the environment. This is why I endorse a deflationary account of our FP. PP, and more importantly the FEP, points to a mind which takes shortcuts, prior to our ability to cognate reasons at all levels of interaction. This leads to a view of mind whereby a minimum of resources are utilised, and propositional reason-giving only occurs on a post hoc basis. Cultural norms, therefore, provide a vast cognitive resource which enables us to go "offline"[182] and achieve much more in our everyday reasoning than would be possible if we were to process the majority of our interpersonal experiences. We can view the norms of culture as providing an emancipation from unnecessary cognition in our ever-deepening social environment.

[182] By "offline" I mean in the sense of non-overt and subconscious reasoning.

Chapter Five: Conclusion

"Keep in mind that what seems like rationality is often just rationalization, playing catch-up with subterranean forces that we never suspect."

- Robert M Sapolsky (2017: 423)

1. Introduction

So much for this review of the current landscape of the philosophical and scientific fields of FP, PP, teleosemiotics, and the rest. But one may still feel as if there is a core element absent from this account of the mind and how it is said to process, represent, and instantiate folk psychological ascriptions. Even if all that has been covered is accepted, one may feel cheated, as if a sleight of hand has been performed. Not only have our folk psychological ascriptions been problematised, but a central cognitivist depiction of the mind has been challenged, as well as representational theories of the mind in general. Not only have we been jolted from our familiar folk thinking, but the very frameworks from which we tend to understand them have been subverted, and even supplanted. Where does this leave us regarding our self-conception? Are we poorer for this new view of ourselves? Well, perhaps, but perhaps not. This chapter will review the central themes of this thesis by summarising the take-home points from the analysis of folk psychological ascriptions in Chapter Two, before doing the same for the implications of PP for our FP theories. Thereafter, the teleosemiotic account of our basic cognition and what it says about representation and content in our cognition, and therefore our FP, will also be reviewed. To end this chapter, some issues that these theories hold for future research and conceptions of cognitive science, neuroscience, and philosophical depictions of our FP will be unpacked.

2. Folk Psychology Problematised

As seen in Chapter Two, traditionally-construed FP is far too constrained a theoretical imposition to be sustained as a workable theory of the mind. Putting aside the issues that arise for FP once the more plausible theories of PP and teleosemiotics come into play, the view of propositional attitudes as individuated components of our mentality is vastly untenable once their truncated natures are revealed. As was elucidated in that chapter, the reasoning behind

supposing the existence of full-blown propositional attitudes in behavioural reason explanation are not convincing enough to easily sway us any longer. It was shown that we arguably construct narratives entailing such attitudes in order to aid us in the navigation of our social environment, but this need not reflect reality to do its work. But Fodor's (1994) disjunction problem, as introduced in Chapter Two, is a reminder of the analytic worries that plague naturalised theories of cognition, wherein it is alleged that truth conditionality *must* hold *everywhere and always* for mental content to be useful or accurate. Propositional content is tied to accurate correspondence with the environment, leading to the issues surrounding the disjunction problem.

Nevertheless, I have argued against the assumption that propositional content is the mark of the cognitive, and have argued that this assumption needs to be done away with. As I have argued for covariant relationships which relate the mind to world, leading to a more holistic understanding of mental states, the disjunction problem falls away. It is a strange imposition to posit these abstract (yet allegedly rigid) propositional structures in the brain in order to account for our cognitive abilities. Not only this, but the insistence that it is *predominantly* these propositional attitudes that enable our truly human cognition, is a view that should be questioned. We are in need of a more comprehensive, empirically-informed theory of cognition before being able to say unequivocally that this is indeed the case. Often, as with many cognitive phenomena such as talk of "representation" and the like, propositional attitudes are assumed to be the end-goal of our theories of cognition. Instead of elucidating why we should incorporate propositional attitudes into our deliberations, they are smuggled in from the start (as argued by Ratcliffe 2007b). One can see why this presupposition could alter conceptual work on the target phenomena of scientific as well as philosophical theories of the mind. Moreover, such propositional thinking is seemingly introduced from an early age in the form of narratives (Hutto 2004, 2007a, 2007b, 2008a, 2008b), rendering a shift in perspective exceedingly difficult. As we have seen, these narratives serve as the exemplars from which co-cognition is effected and maintained across cultures, and they provide the means for reason explanation in our everyday FP. Therefore, our FP narratives provide the core abilities upon which our social understanding is based. As was seen in Chapter Two, Churchland (1981) has been a dominant theorist in the drive to eliminate propositional attitudes from our ontology, and he has done a commendable job. However, it is certainly time to move on from the all-or-nothing approach of true eliminativism, as propositional attitudes clearly occur in our conscious reasoning, regardless of their reality at a neuronal level. They may not be instantiated at this

level in any recognisable form, but they do seem to play some kind of role in cognition and, as argued above, they derive their efficacy from the social domain. They are cultural impositions upon the neuronal substrate.

3. Folk Psychology and Predictive Processing

Looking at folk psychological ascriptions through a plausible theory of our basic cognitive processes, namely predictive processing, also seems to suggest that propositional attitudes should be done away with, at least when it comes to most cognitive processes. There seems to be a basic incongruence between psychology of the propositional kind and the PP paradigm. On this view, the separation of beliefs and desires into their canonical distinctions are broken down and woven together on an active inference construal of cognitive processing. If predictions are action-oriented in the PP sense, we simply do not have beliefs and desires instantiated within the brain (Dewhurst 2017). What we do have in their place is a probabilistic model of the environment which is attuned to appropriate cues. Moreover, these proto-beliefs and -desires, as explained in Chapters Two and Three, are far too fine-grained to be picked out in publicly-accessible language (at least on most planes of the predictive hierarchy)[183]. These "beliefs" and "desires" are inter-woven and not clearly delineated and transparent enough for our ability to categorise them. A move we can perform to counteract the radical nature of this picture of FP would be to shift our understanding of what the folk have always meant when they have ascribed behaviour on the basis of propositional attitudes. It may be that canonical beliefs or desires were always meant to pick out the phenomena postulated by PP in some sense (Dewhurst 2017: 10). But this does not seem to track very well with what the folk have actually been doing. As Dewhurst (ibid.) points out, it is more plausible to say that beliefs and desires, as attributed by the folk, were always meant to pick out macro-reasons for behaviour. The way that I see it, these macro-reasons are evident on the broader scale of co-communicative understanding, yet they do not directly depict any corresponding cognitive processes, which can be disconcerting to our lived experience. Therefore, beliefs and desires are emergent properties once culture has had its effect on the brain. As described earlier in this thesis, FP is a fuzzy endeavour that nevertheless picks out framed "just so" narratives describing generalities about human behaviour for public consumption. There is simply no determinate

[183] Not to mention that on REC's teleosemiotic construal not much would be represented in any contentful sense in any case.

way of delineating an "actual", corresponding belief and desire as publicly understood at a cognitive level, but this, at the end of the day, is a "problem" more for our conceptions of who we are than any undermining of our ontological reality. We should take a Dennettian (1989) stance on our understanding of where propositional attitudes are situated, and instead position them (in a non-essential sense) somewhere between brain and environment (metaphorically speaking).

The incompatibility between the PP model and our public folk psychological ascriptions need not entail a drastic change in our *behaviour*, but rather a drastic change in the *conceptualisation of* our FP (Dewhurst 2017: 10). Our folk psychological ascriptions *need not necessarily* be congruent with reality in order to provide a useful elucidation of other peoples' thoughts. This means that, while FP as propositional attitude ascription is predominantly incompatible with PP, it is only the literalists such as Fodor and the like who will run into real problems as a result. If our FP and any theories which hang on it are set within boundaries demarcated by propositional attitudes, the PP framework undermines its efficacy as a coherent ToM. But this is also the case for Churchland's form of eliminativism which advocates that either FP be understood as an accurate portrayal of our internal cognitive processes, or it should be eliminated from our epistemological framework. It hardly needs saying anymore that this is not a necessary move in order to reach coherency in our scientific theorising. We could simply modify our understanding of what it is that we are picking out in the world when we folk psychologise, without entirely eliminating the terminology from public discourse. PP *could* seep into the mainstream and alter our understanding of what it is that we are *actually* describing when we speak of the behaviour of those around us as well as of ourselves. This would not be easy to guide in any controlled manner[184], and the need to do so is up for debate. But perhaps we can see this as more of an opportunity for a revising of our FP into a more scientifically cogent one. As Dewhurst puts it:

> We should aim to develop a novel conceptual taxonomy that more accurately reflects the structure of cognition and allows us to move beyond the limitations of folk psychological discourse. Understood in this way folk psychology could be used to identify interesting target phenomena and inspire scientific research […] but should not be used as a source of technical cognitive scientific concepts. (Dewhurst 2017: 11).

[184] Think of the notions of the Freudian unconscious and how it populated public discourse at the turn of the last century.

There have been a number of efforts to explicate a way of doing this, such as Anderson (2015), Klein (2012), Poldrack (2006; 2010) Price and Friston (2005), and others (in Dewhurst 2017: 11). These efforts have sought to revise the understanding of our cognition by amalgamating large sets of studies, all in order to segregate which processes are of functional relevance to our cognitive processes. From here, we can reconceptualise our cognitive processes (and by extension our FP) by incorporating terminology which is more appropriate[185]. By way of example, in order to incorporate our understanding of PP into FP, we would need to be aware of the way in which beliefs and desires are amalgamated into a singular predictive state. In addition, the differences in the interactive processes of such predictive states and those alleged to arise from propositional attitudes could be elucidated[186]. I am inclined to support this, although we must be careful if folk psychological terminology and concepts are used to "identify interesting target phenomena" (Dewhurst 2017: 11). This is because occasionally the long-standing usage of specific terminology may push us down conceptual avenues without recourse for revision. We would be philosophically stuck in our ways. The way to come out from under the yoke of decades of conceptual baggage is not always clear. It would not be as easy as changing the words we use ("belief" and "desire" etc.) to others (much like "surprise" and "surprisal" have been). I suppose, in any case, that FP as traditionally understood would still hang around, as the folk can still knowingly treat our coarse-grained interpretations in an "as if" sense. The entire edifice of FP, recognised as more than simply the attribution of snapshot mental states, could be subsumed into the PP framework in this way. Theoretical propositional attitude FP, however, would more than likely not survive the transition. The shift in our self-conception that PP would more than likely bring about would irrevocably alter how we conceptualise and understand FP.

4. Folk Psychology and Teleosemiotics

As for what the teleosemiotic account of cognitive processing gives us with regards to understanding our FP, we are in need of demarcating the functional processes of the brain as clearly as possible. I have given good reason to doubt the immediate assumption that the brain

[185] For example, Poldrack (2010) attempts to explicate the failings of attempting to map particular brain regions in terms of their particular functions. He instead advocates redefining the ontology of the brain (ibid.: 760).
[186] Predictive states, for example, could be vastly different to the alleged interactions between the neuronal states of propositional attitudes if predictive states cannot be clearly defined. We are in need of a different view of the imperative ("desire") and indicative ("belief") aspects of our neuronal states if they are indeed predictive states. Naturally, they would not be "belief" and "desire" traditionally conceived.

trades in representational "content" or "information" and the like. "Information" is a quite general term, but it is usually used to denote data of some variant that is "acquired" in some fashion, and *then* utilised within the brain. I introduced REC at this point to upset this commonly held view, as REC is a formulation of neurodynamic interaction within the brain that sees it as more "informationally sensitive" (Hutto & Myin 2017: 237) than informationally accurate. This view depicts neuronal pathways as influencing behaviour, but importantly *not* representing states of affairs in the world. Therefore, organismic behavioural responses are sustained through direct acquaintance with the environment, through the indexical ability of covariant relationships which calls into question the assumption that cognition is, at base, representational (ibid.: 29-31). The top-down models that are understood to operate within PP can be understood to afford a grip upon the world (via active inference) to make these covariant relationships do real cognitive work. It is this which fundamentally entwines the literature on REC and PP and explicates our basic cognitive functionality. As a result, we can effectively ignore the disjunction problem as a serious epistemic and metaphysical hindrance when we adopt this covariance model of basic organismic functionality. This is because it does not hold any great philosophical weight if we are dealing with information in a covariant sense as opposed to notions of accuracy. This is positive in that it frees us from philosophical baggage which does not align with evolutionary theory.

To move away from a contentful mind, and therefore the more traditional cognitivist thinking around propositional attitudes, we need to change some of our fundamental assumptions, as they place us at a considerable conceptual handicap. The repercussions of not adopting this approach is shown by Bechtel (2016), who assessed whether researchers have undergone studies[187] in such a way so as to assume the target phenomena of the study, in this case the representational abilities of the brain (ibid.: 3-7). Bechtel arrives at the conclusion that neuroscientists appear to be set on the idea that they are attempting to elucidate the form of representation within the brain, as well as the content itself (ibid.: 29-31). In the experiments, researchers attached electrodes to rats' brains to attempt to explain their spatial navigation in terms of "map-like" representations, which represent the environment in a contentful manner (ibid.: 29). In other words, the notion of contentful representation was foundational to the experiments. Bechtel does, however, believe that viewing content in this way is what enables such a rich body of work in the first place. This may be true, but of course there is more at

[187] In this case, studies on the neuronal representation that the place cells in the brain allegedly afford were analysed.

stake here than motivation. Accuracy of a fecund research paradigm relies on the lucidity of its conceptual premises, leading me to argue that the commitment to conceptual accuracy is of paramount importance. How can we say that our cognition is being elucidated if we are understanding the very nature of our research programs in a subtly flawed way? If we believe that robust representation is a necessary component of cognition, we are often going to be going up dead ends in future research on this score. REC appears to offer just the right frugal approach to research on cognition by avoiding adding redundant cognitive processes to the mix, and thus side-stepping the issue of content (and thus the HPC). Furthermore, any lingering worries about disjunction issues are similarly swept away.

If there is minimal content in the brain, then conventionally understood FP in terms of propositional attitude ascriptions becomes an increasingly untenable framework for understanding behaviour. If accuracy is not needed to be contextually-bound to the environment, the teleosemiotic account of intentionality provides a fairly novel tool to conceptualise the real-world sensitivity of our neuronal systems to said environment. The neuronal systems which instantiate navigation through our environment via their dynamic firings are all that are needed for appropriate (notice: not necessarily purely accurate) behaviour to be elicited. But these neuronal systems only exhibit their attuned environmental sensitivity once a particular context is experienced, thereby alluding to a correspondence between the neurons and the environment[188]. Notice, again, that this need not be representational, especially in the sense that propositional attitudes are thought to function. Therefore, what we are left with is a distinctly minimally contentful picture of our minds, and by extension our FP. But our cognition is fully realised, not only in its underlying processes being directly attuned to the environment, but also in its capacity to "loop" (Hutto & Myin 2017: 253) through culture, enriching the mind in such a way so as to provide the content that conceiving of propositional attitudes requires. In other words, through interaction with cultural artefacts of consumption, our sociocultural milieu provides some of our mental content, rather than the basic processing of our minds. The extensive, patterned practices of enculturation are the true harbingers of the truth conditional content needed by propositional attitudes, and it is only through our engagement with these practices that our content-enriched intentionality arises through a field of "cultural affordances" (Ramstead *et al*. 2016). Therefore, we have enculturation to thank for our seemingly "content-rich" cognitive representations. But this is a gloss on what our basic

[188] For example, place cells within the brain are believed to hardwire affordances within the environment to make the individual sensitive to future interactions when in a similar environment in the future (Ólafsdóttir *et al*. 2015: 1).

minds actually do, which is a direct coupling between environment and mind that trades on a form of cognitive "inertia" toward specific environmental stimuli. This "inertia" takes the place of truth conditional content, and serves exactly the same function in tuning the mind in such a way so as to be able to act upon the world. In other words, our minds are directed toward a field of affordances (Bruineberg & Rietveld 2014; Rietveld & Kiverstein 2014; Ramstead *et al*. 2016) not a field of content.

Lending weight to this idea, Hutchins (2005: 1558) has described how some experiments indicate that many concepts which "lack cultural coherence" simply cannot be grasped sufficiently within memory. For example, D'Andrade (1989) has shown that some inferential reasoning cannot easily be made unless tied to concrete concepts. In this case reasoning toward particular conclusions using simple premises with "x" and "y" as placeholders have a significant impact upon success rate. Substitute these placeholders with concrete concepts, and the rate of success increases exponentially. While a simple example, it is a window into the effect that shared models of the environment have as they reinforce the thinking and behaviour of individuals, over and above their base rudimentary cognition. Furthermore, as these cultural artefacts are embedded themselves within a vast mutually supportive network of other cultural artefacts, they have the effect of constraining themselves to a particular logical space. Hutchins (2013: 12-13) has also linked this concept with PP by alluding to the fact that cultural practices reduce entropy and increase the predictability of an individual's experience. Much more effort is needed to flesh out this aspect of social cognition.

What we are left with, once the layers of theory have been peeled away from talk surrounding FP, is something seemingly more anaemic than full-blown propositional attitude FP. But this anaemia only extends so far, as we must keep in mind the entire edifice of folk psychological discourse as a whole. The capacity for the utilisation of learnt narratives, predictions of behaviour, and the normative streamlining that enables these abilities, speak to a larger picture than a restricted FP model predicated on propositional attitudes (Dewhurst 2017: 10). Having said this, there is no need for accuracy in explicating the underlying components of our reason-explanation for behaviour when the general narrative frameworks successfully elucidate macro-reasons for behaviour at the coarse-grained level. Again, propositional thinking only shows itself at this level, but not necessarily further down into the underlying covariant relationships that hold between neurons and the environment. Therefore, there is an acquired capability that we utilise in our social "understanding", operating at the "basic" (but fundamental) level. This capability orientates us toward distal (in this case social) cues so as to

interpret them on a visceral level, all without full-blown reason explanations in the propositional sense. This latter reason explanation only arrives through overt cognition and co-communication, utilising public cultural symbols. Without this, our more generalised FP is dispositional in nature, leading to my formulation of DFP, as introduced toward the end of Chapter Four. This simply denotes the dispositional nature of how we react to social cues without overt reason explanation, which is significant insofar as it does not invoke propositional content in the traditional sense. Our social understanding is subsumed into our cognitive architecture, which enables a more streamlined and frugal use of cognitive resources. As stated earlier in this thesis, once propositional attitudes are off the table, this is what we are left with in our FP. To me, once the bigger picture is taken into account, there is no big loss here.

5. Future Endeavours

It would be remiss of me not to give a brief overview of the prevalent issues of this depiction of our minds as minimally-contentful prediction machines. We cannot supplant a decades-long view of the mind and expect no resistance and no repercussions. The conceptual leap alone is a shift in perspective many may not find easy to take, but luckily this is not something we need overly concern ourselves with. For example, if the PP framework is accurate, this reality may simply slowly seep into public discourse, or be co-opted when it is useful. This is a radical reconceptualisation of our underlying cognitive processing, but it is seemingly business as usual once out of the conceptual armchair and into the world. If anything, it is not clear what to make of this depiction of the mind as of yet. The science has not been settled significantly enough to say, so the ramifications could change by the time this is even read. But there is still a sense that if this is at least half the picture, our conception of our minds will never be the same.

The predominant ramifications perhaps lie in the effect that this would have on our theoretical propositional attitude FP, as both Clark (2016: 285-288) and Hohwy (2013: 2) have attested. A new FP taxonomy would need to encompass this conception of how our social understanding works in some form. The issues would perhaps sort themselves in time, but even using the word “prediction” usually denotes some form of agency on the part of the individual, so it would be a tough habit to break to make that cognitive shift. For now, we could focus on our folk psychological language use, or more accurately how precise it is in the elucidation of our

PP processes. This, of course, would only be necessary if the entire framework of PP is actually scientifically viable in the first place, but that is a purely empirical question. Some examples of neuronal dynamical studies that are seeking to eke out the details concerning motivation and desire at a fine-grained level are Friston *et al*. (2012), Gershman and Daw (2012), Bach and Dolan (2012), and Schwartenbeck *et al*. (2014). But what of the "higher" cognitive abilities displayed during social cognition and instrumental reasoning? These have only begun to be explored, and much interesting work can be found in Clark (2016)[189], Hohwy (2013) (especially in Chapters Nine to Twelve), Harrison *et al*. (2011), Hobson and Friston (2014), Moutoussis *et al*. (2014), and Seth (2014a). For example, Moutoussis *et al*. (2014: 10-11) cast "representations" within the mind during interpersonal interaction as a form of "belief" in a Pavlovian sense where social learning may skew us toward "beliefs irrespective of their consequences" (ibid.: 10) within certain contexts. Essentially, specific brain states are "taught" to be attuned to behaviour much in the same way that Pavlov's animals were behaviourally trained. Again, this is a social attunement not built upon sentential propositional attitudes, but upon generative models within the neuronal substrate. Studies such as these are the first steps to a more coherent picture of our neuronal as well as social functioning in terms of our belief/desire psychology.

The research looks bright, but a multitude of conceptual hurdles are still in the way. For instance: the extended battle over conceptual clarity concerning representation within the mind is far from over, with the promising but abrasive REC programme still waiting in the wings. Content may not be a problem if it is done away with but, as has been demonstrated, content is often introduced as the very mark of the cognitive. The REC approach is to ignore the problem and focus elsewhere, but it is clear that that is not viable if content is put on the proverbial pedestal and assumed to be pertinent to subsequent theorising. The problem of truth conditional content for propositional thinking may be swept aside if our minds are recognised as embodied and embedded within their environment in such a way so as to harbour covariant relationships to distal stimuli. The socio-cultural environment then provides the necessary accuracy-conditions as a narrative gloss over and above the actual neuronal reality. It is only in this sense that propositional thinking is "instantiated" within the brain. Again, there is still a further need to elaborate on how enculturation affects the brain at a concrete level, but this is a promising

[189] Especially in his discussions surrounding the possibility of generative models and how they potentially form a variant of simulation, enabling long-term planning in conjunction with social cognition. Moreover, there is discussion over how precision-weightings can be "self-manipulated" through "linguistically inflected reasoning" (Clark 2016: 300).

avenue of research. In short, the hurdles become wrinkles in the road if we simply shift our perspective toward these more radical theories, but this will require overcoming decades-worth of conceptual baggage. Academics themselves should at the very least apprise themselves of the latest science (and philosophy thereof) before theorising from their armchairs. Further research is needed in order to elucidate the bodily responsitivity toward environmental cues, while remaining neutral to what processes are being instantiated. Cao (2012), for example, conducted a study to figure out what is being said when specific parts of the brain are said to be in "contentful" communication with one another. Is there more explanatory weight that is borne by positing intermediary content? After analysing a diverse cast of functionally different processes and components within the brain, from neurotransmitters to networks of neurons, she arrives at the simple answer: not much. It is only once the organism as a whole is taken into account that semantics comes into the picture (ibid.: 70). I propose it is only in this light that philosophers can be useful in their contribution toward the endeavours of the cognitive sciences. If they are not "in the lab", they can certainly help clarify matters, building on what the science has revealed. The downside to this is a jockeying for position with regards to whose theory holds the best explanation. While I agree with Hutto and Myin in their belief that philosophers "are not mere onlookers" (Hutto & Myin 2017: 244), I am reticent to speculate just how much the majority of philosophical theorising in this field has been helpful. I am, however, confident that the avoidance of the HPC manages to sidestep otherwise "intractable theoretical mysteries" (ibid.).

A lesser issue in this regard (due to its waning popularity), is of the question of the role of theoretical propositional FP in its canonical guise. It is a research paradigm which is petering out, and perhaps this study can contribute to the slow death of the paradigm. Again, we shall continue to use folk terminology, but it is best if its technical usage falls away in time. It is much more plausible to view FP as trading in a priori modelling of the environment, which is altered on an ad hoc basis. Reason explanation is shown to be a retrospective instrumental thinking that affords us a grip on an otherwise cognitively fuzzy environment, providing platitudes that guide us along pathways that work in the macro-perspective of general social interaction. We will be left with talk of folk psychological propositional attitudes in our everyday speech, but hopefully within academic circles these will eventually turn out to be relics from a hard-to-shake-off jargon. A developing science needs to look to a future emancipated from our present conceptual shackles, and I hope that this thesis goes some way to showing the why and the how of achieving this.

BIBLIOGRAPHY

Adams, F & Aizawa, K. 2001. The bounds of cognition. *Philosophical Psychology*, 14(1): 43-64.

Aizawa, K. 2014. The enactivist revolution. *Avant (Torun)*, 5(2): 19-42.

Aizawa, K. 2015. What is this cognition that is supposed to be embodied?. *Philosophical Psychology*, 28(6): 755-775.

Akins K. 1996. Of sensory systems and the "aboutness" of mental states. *Journal of Philosophy*, 113: 337-372.

Anderson, ML. 2015. Mining the brain for a new taxonomy of the mind. *Philosophy Compass*, 10(1): 68-77.

Andrews, K & Gruen, L. 2014. Empathy in Other Apes, in *Empathy and Morality*, edited by HL Maibom. Oxford: Oxford University Press: 193-209.

Andrews, K. 2007. Critter Psychology: On the Possibility of Nonhuman Animal Folk Psychology, in *Folk Psychology Re-assessed*, edited by D Hutto and M Ratcliffe. New York: Springer: 191-209.

Andrews, K. 2008. It's in your nature: a pluralistic folk psychology. *Synthese*, 165: 13-29.

Andrews, K. 2012. *Do Apes Read Minds? Toward a New Folk Psychology* Cambridge: MIT Press.

Andrews, K. 2015a. *The Animal Mind: An Introduction to the Philosophy of Animal Cognition*. New York: Routledge.

Andrews, K. 2015b. The Folk Psychological Spiral: Explanation, Regulation, and Language. *The Southern Journal of Philosophy*, 53: 50-67.

Bach, DR & Dolan, RJ. 2012. Knowing how much you don't know: a neural organization of uncertainty estimates. *Nature Reviews Neuroscience*, 13(8): 572-586.

Baillargeon, R, Scott, R & He, Z. 2010. False-belief understanding in infants. *Trends in Cognitive Sciences*, 14: 110-118.

Barandiaran, X & Di Paolo, E. 2014. A genealogical map of the concept of habit. *Frontiers in Human Neuroscience*, 8: 522.

Baron-Cohen, S, Leslie, A & Frith, U. 1985. Does the autistic child have a "Theory of Mind"?. *Cognition*, 21: 37-46.

Barrett, LF & Bar, M. 2009. See it with feeling: affective predictions during object perception. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521): 1325-1334.

Bartsch, K & Wellman, H. 1995. *Children talk about the Mind*. Oxford: Oxford University Press.

Bashour, B & Muller, HD. (eds). *Contemporary Philosophical Naturalism and Its Implications*. London: Routledge.

Bateson, PPG & Hinde, RA. (eds). 1976. *Growing Points in Ethology*. Cambridge: Cambridge University Press.

Bechtel, W. 2016. Investigating neural representations: The tale of place cells. *Synthese*, 193: 1287-1321.

Beer, RD. 2000. Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4(3): 91-99.

Block, N & Kitcher, P. 2010. Misunderstanding Darwin: Natural Selection's Secular Critics Get it Wrong. *Boston Review*. [O]. Available: http://bostonreview.net/ned-block-philip-kitcher-misunderstanding-darwin-naturalselection [10 February 2018].

Bogacz, R. 2017. A tutorial on the free-energy framework for modelling perception and learning. *Journal of Mathematical Psychology*, 76(B): 198-211.

Botterill, G. 1996. Folk Psychology and Theoretical Status, in *Theories of Theories of Mind*, edited by P Carruthers and PK Smith. New York: Cambridge University Press: 105-118.

Branquinho, J. 2001. *The Foundations of Cognitive Science*. Oxford: Oxford University Press.

Brogaard, B. 2014. (ed). *Does Perception Have Content?*. Oxford: Oxford University Press.

Brook, A & Ross, D. (eds). 2002. *Daniel Dennett: Contemporary Philosophy in Focus*. Cambridge: Cambridge University Press.

Brown, H, Friston, KJ & Bestamnn, S. 2011. Active inference, attention and motor preparation. *Frontiers in Psychology*, 2: 218.

Brown, M, Salverda, AP, Dilley, LC & Tanenhaus, MK. 2011. Expectations from preceding prosody influence segmentation in online sentence processing. *Psychonomic Bulletin and Review*, 18:1189-1196.

Bruineberg, J & Rietveld, E. 2014. Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8: 1-14.

Bruner, J. 1990. *Acts of Meaning*. Cambridge: Harvard University Press.

Burge, T. 2010. *The Origins of Objectivity*. Oxford: Oxford University Press.

Burr, C. 2017. Embodied decisions and the predictive brain, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Byrne, RW & Whiten, A. (eds). 1988. *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans*. Oxford: Oxford University Press.

Byrne, RW. 2000. Evolution of Primate Cognition. *Cognitive Science*, 24(3): 543-570.

Campbell, M & O'Sullivan, M. (eds). 2014. *Wittgenstein and Perception*. New York: Routledge.

Cao, R. 2012. A teleosemantic approach to information in the brain. *Biology & Philosophy*, 27(1): 49-71.

Carruthers, P & Smith, PK. (eds). 1996. *Theories of Theories of Mind*. Cambridge: Cambridge University Press.

Carruthers, P. 2015. Two systems for mindreading?. *Review of Philosophy and Psychology*, 7(1): 141-162.

Casasanto, D & Dijkstra, K. 2010. Motor action and emotional memory. *Cognition*, 115(1): 179-185.

Chalmers, D. 1995. Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3): 200-219.

Chalmers, D. 1996. *The Conscious Mind*. Oxford: Oxford University Press.

Chemero, A. 2009. *Radical Embodied Cognitive Science*. Cambridge: MIT Press.

Churchland, PM & Churchland, PS. (eds). 1998. *On the Contrary: Critical Essays 1987-1997*. Cambridge: MIT Press.

Churchland, PM. 1981. Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, 78(2): 67-90.

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

Churchland, PM. 1993. State-space semantics and meaning holism. *Philosophy and Phenomenological Research*, 53: 667-672.

Churchland, PM. 1995. *The engine of reason, the seat of the soul: A philosophical journey into the brain*. Cambridge: MIT Press.

Churchland, PM. 1998a. Folk Psychology, in *On the Contrary: Critical Essays, 1987-1997*, edited by PM Churchland and PS Churchland. Cambridge: MIT Press: 3-15.

Churchland, PM. 1998b. Evaluating Our Self-Conception, in *On the Contrary: Critical Essays, 1987-1997*, edited by PM Churchland and PS Churchland. Cambridge: MIT Press: 25-38.

Churchland, PM. 1998c. Activation vectors versus propositional attitudes: How the brain represents reality, in *On the Contrary: Critical Essays, 1987-1997*, edited by PM Churchland and PS Churchland. Cambridge: MIT Press: 39-44.

Churchland, PM. 2002. Catching Consciousness in a Recurrent Net, in *Daniel Dennett: Contemporary Philosophy in Focus*, edited by A Brook and D Ross. Cambridge: Cambridge University Press: 64-81.

Churchland, PS, Ramachandran, VS & Sejnowski, TJ. 1994. A Critique of Pure Vision, in *Large-Scale Neuronal Theories of the Brain*, edited by C Koch, JL Davis, TJ Sejnowski and TA Poggio. Cambridge: MIT Press: 23-60.

Clark, A & Chalmers, D. 1998. The extended mind. *Analysis*, 58: 7-19.

Clark, A & Toribio, J. 1994. Doing without representing?. *Synthese*, 101(3): 401-431.

Clark, A. 1989. *Microcognition: Philosophy, cognitive science and parallel distributed processing*. Cambridge: MIT Press.

Clark, A. 1997. *Being there: Putting brain, body, and world together again.* Cambridge: MIT Press.

Clark, A. 2003. *Natural-Born Cyborgs: Minds, Technologies and the Future of Human Intelligence*. Oxford: Oxford University Press.

Clark, A. 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: Oxford University Press.

Clark, A. 2012. Dreaming the whole cat: Generative models, predictive processing, and the enactivist conception of perceptual experience. *Mind*, 121(483): 735-771.

Clark, A. 2013a. Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3): 181-204.

Clark, A. 2013b. The many faces of precision (Replies to commentaries on "Whatever next? Neural prediction, situated agents, and the future of cognitive science''). *Frontiers in Psychology*, 4: 270.

Clark, A. 2013c. Are we predictive engines? Perils, prospects, and the puzzle of the porous perceiver. *Behavioral and Brain Sciences*, 36(3): 233-253.

Clark, A. 2015a. Embodied prediction, in *Open MIND: 7*(T), edited by T Metzinger and JM Windt. Frankfurt am Main: MIND Group: [s.p.].

Clark, A. 2015b. Predicting peace: Reply to Madary, in *Open MIND: 7*(R), edited by T Metzinger and JM Windt. Frankfurt am Main: MIND Group: [s.p.].

Clark, A. 2015c. Radical predictive processing. *The Southern Journal of Philosophy*, 53: 3-27.

Clark, A. 2016. *Surfing uncertainty: Prediction, action, and the embodied mind.* Oxford: Oxford University Press.

Clark, A. 2017. How to Knit Your Own Markov Blanket: Resisting the Second Law with Metamorphic Minds, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Colodny, RG. (ed). 1962. *Frontiers of Science and Philosophy*. Pittsburgh: University of Pittsburgh Press.

Colombo, M. 2012. Explaining social norm compliance. A plea for neural representations. *Phenomenology and the Cognitive Science*s, 13(2): [s.p.].

Colombo, M. 2017. Social motivation in computational neuroscience (Or if brains are prediction machines then the Humean theory of motivation is false), in *Routledge Handbook of Philosophy of the Social Mind*, edited by J Kiverstein. New York: Routledge: 320-340.

Crane, T. 2009. Is perception a propositional attitude?. *Philosophical Quarterly*, 59: 452-469.

Currie, G, & Sterelny K. 2000. How to think about the modularity of mind-reading. *Philosophical Quarterly*, 50: 145-160.

Churchland, PM. 1988. Perceptual plasticity and theoretical neutrality: A reply to Jerry Fodor. *Philosophy of Science*, 55(2): 167-187.

Churchland, PM. 1989. Folk psychology and the explanation of human behavior. *Philosophical Perspectives*, 3: 225-241.

D'Andrade, R, 1989. Culturally based reasoning, in *Cognitions in Social Worlds*, edited by A Gellatly, D Rogers and J Sloboda. New York: McGraw-Hill: [s.p.].

Davies, M & Stone, T. (eds). 1995. *Mental Simulation*. Oxford: Blackwell.

Davies, PS. 2007. What Kind of Agent Are We? A Naturalistic Framework for the Study of Human Agency, in *Distributed Cognition and the Will: Individual Volition and Social Context*, edited by D Ross, D Spurrett and Harold Kincaid and GL Stephens. Cambridge: MIT Press: 39-60.

De Brigard, F. 2014. Is memory for remembering? Recollection as a form of episodic hypothetical thinking. *Synthese*, 191(2): 1-31.

Debruille, JB, Brodeur, MB & Franco Porras, C. 2012. N300 and social affordances: a study with a real person and a dummy as stimuli. *PLoS ONE,* 7(10): e47922.

Degenaar, J & O'Regan, K. 2017. Sensorimotor Theory and Enactivism. *Topoi*, 36(3): 393-407.

Dempster, AP, Laird, NM & Rubin, DB. 1977. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society,* Series B 39: 1-38.

Dennett, DC. 1987. *The Intentional Stance*. Cambridge: MIT Press.

Dennett, DC. 1991a. *Consciousness Explained*. Boston: Little, Brown and Company.

Dennett, DC. 1991b. Real Patterns. *Journal of Philosophy*, LXXXVIII: 27-51.

Dennett, DC. 1991c. Two Contrasts: Folk Craft versus Folk Science, and Belief versus Opinion, in *The Future of Folk Psychology: Intentionality and Cognitive Science*, edited by JD Greenwood. Cambridge: Cambridge University Press: 135-148.

Dennett, DC. 1995. *Darwin's Dangerous Idea*. New York: Simon & Schuster.

Dennett, DC. 1996. The Myth of Double Transduction. [O]. Available: http://ase.tufts.edu/cogstud/dennett/papers/transduc.htm [10 February 2018]

Dennett, DC. 2005a. *Sweet Dreams: Philosophical Obstacles to a Science of Consciousness*. Cambridge: MIT Press

Dennett, DC. 2005b. Two steps closer on consciousness, in *Paul Churchland (Contemporary Philosophy in Focus)*, edited by BL Keeley. Cambridge: Cambridge University Press: 193-209.

Dennett, DC. 2013a. Expecting ourselves to expect: The Bayesian brain as a projector. *Behavioral and Brain Sciences*, 36(3): 209-210

Dennett, DC. 2013b. *Intuition Pumps, And Other Tools For Thinking*. New York: W. W. Norton & Company.

Dennett, DC. 2015. Why and How Does Consciousness Seem the Way it Seems?, in *Open MIND: 10*(T), edited by T Metzinger and JM Windt. Frankfurt am Main: MIND Group: [s.p.].

Dennett, DC. 2017. *From Bacteria To Bach and Back: The Evolution of Minds*. New York: W. W. Norton & Company

Descartes, R. 2008. *Meditations on First Philosophy*. Translated by M Moriarty. Oxford: Oxford University Press.

De-Wit, L, Alexander, D, Ekroll, V & Wagemans, J. 2016. Is neuroimaging measuring information in the brain?. *Psychonomic Bulletin & Review*, 23(5): 1415-1428.

Dewhurst, J. 2017. Folk Psychology and the Bayesian Brain, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Di Paolo, E. 2009. Extended life. *Topoi*, 28(1): 9-21.

Dolega, K. 2017. Moderate Predictive Processing, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Downey, A. 2017. Radical Sensorimotor Enactivism & Predictive Processing. Providing a Conceptual Framework for the Scientific Study of Conscious Perception, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Dretske, F. 1988. *Explaining Behavior: Reasons in a World of Causes*. Cambridge: MIT Press.

Dreyfus, HL & Wrathall, M. (eds). 2014. *Skillful Coping: Essays on the Phenomenology of Everyday Perception and Action*. Oxford: Oxford University Press.

Dreyfus, HL. 2014 (2001). The primacy of phenomenology over logical analysis, in *Skillful Coping: Essays on the Phenomenology of Everyday Perception and Action*, edited by HL Dreyfus and M Wrathall. Oxford University Press: 146-167.

Dupré, J. 2012. *Processes of life: Essays in the philosophy of biology*. New York: Oxford University Press.

Dupré, J. 2014. A process ontology for biology. *Physiology News*, 100: 33-34.

Dunn J. 1991. Understanding others: Evidence from naturalistic studies of children, in *Natural Theories of Mind: Evolution, Development, and Simulation of Everyday Mindreading*, edited by A Whiten. Oxford: Blackwell: [s.p.].

Engel, AK, Fries, P & Singer, W. 2001. Dynamic predictions: Oscillations and synchrony in top-down processing. *Nature Review Neuroscience*, 2(10): 704-716.

Engel, AK, Friston, K & Kragic, D. 2013. Where is the action? The pragmatic turn in cognitive science. *Trends in Cognitive Science*, 17(5): 202-209.

Engel, AK, Maye, A, Kurthen, M & Konig, P. 2013. Where's the action? The pragmatic turn in cognitive science. *Trends in Cognitive Sciences*, 17(5): 202-208.

Evans, C & Evans, L. 1999. Chicken food calls are functionally referential. *Animal Behavior*, 58(2): 307-319.

Evans, G. 1982. *The Varieties of Reference*. Oxford: Oxford University Press.

Fabry, RE. 2015. Enriching the notion of enculturation: Cognitive integration, predictive processing, and the case of reading acquisition - A commentary on Richard Menary,

in *Open MIND: 25*(C), edited by T Metzinger and JM Windt. Frankfurt am Main: MIND Group: [s.p.].

Fabry, RE. 2017. Predictive Processing and Cognitive Development, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Fabry, RE. 2018. Enculturation and narrative practices. *Phenomenology and the Cognitive Sciences*, 17(5): 911-937.

Feldman, H & Friston, KJ. 2010. Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4: 215.

Fodor, JA & Lepore, E. 1992. *Holism: A Shopper's Guide*. Cambridge: Wiley-Blackwell.

Fodor, J & Piattelli-Palmarini, M. 2010a. "Misunderstanding Darwin": An Exchange. [O]. Available: http://bostonreview.net/BR35.2/darwin_exchange.php. [9 September 2019].

Fodor, JA & Piatelli-Palmarini, M. 2010b. *What Darwin Got Wrong*. New York: Farrar, Straus and Giroux.

Fodor, JA & Pylyshyn ZW. 1988. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28: 3-71.

Fodor, JA. 1975. *The Language of Thought*. Cambridge: Harvard University Press.

Fodor, JA. 1978. Propositional attitudes. *Monist*, 61: 501-523.

Fodor, JA. 1984. Semantics, Wisconsin Style. *Synthese*, 59: 231-250.

Fodor, JA. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge: MIT Press.

Fodor, JA. 1990. Information and Representation, in *Information, Language and Cognition*, edited by P Hanson. Vancouver: University of British Columbia Press: [s.p.].

Fodor, JA. 1994. *A Theory of Content and Other Essays*. Cambridge: MIT Press.

Fodor, JA. 1995. A theory of the child's Theory of Mind, in *Mental Simulation*, edited by M Davies and T Stone. Oxford: Blackwell: 283-296.

Fodor, JA. 1996a. Deconstructing Dennett's Darwin. *Mind & Language*, 11: 246-262

Fodor, JA. 1996b. On Swampkinds. *Mind & Language*, 11(1): 70-130.

Fordor, JA. 2003. *Hume Variations*. Oxford: Oxford University Press.

Fodor, JA. 2008a. Against Darwinism. *Mind & Language*, 23: 1-24.

Fodor, JA. 2008b. *LOT2: The Language of Thought Revisited*. Oxford: Clarendon Press.

Fletcher, P & Frith, C. 2009. Perceiving is believing: A Bayesian approach to explaining the positive symptoms of schizophrenia. *Nature Reviews: Neuroscience*, 10: 48-58.

Friston, KJ, Adams, RA, Perrinet, L & Breakspear, M. 2012. Perceptions as Hypotheses: Saccades as Experiments. *Frontiers in Psychology*, 3: 151.

Friston, KJ & Buzsáki, G. 2016. The functional anatomy of time: What and when in the brain. *Trends in Cognitive Sciences*, 20(7): 500-511.

Friston, KJ, Daunizeau, J & Kiebel, SJ. 2009. Reinforcement Learning or Active Inference?. *PLoS One*, 4(7): e6421.

Friston, KJ, Daunizeau, J, Kilner, J & Kiebel, SJ. 2010. Action and behavior: a free-energy formulation. *Biological Cybernetics*, 102(3): 227-260.

Friston, KJ & Kiebel, S. 2009a. Cortical circuits for perceptual inference. *Neural Networks: the official journal of the International Neural Network Society*, 22(8): 1093-1104.

Friston, KJ & Kiebel, S. 2009b. Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 364(1521): 1211-1221.

Friston, KJ, Kilner, J & Harrison, L. 2006. A free-energy principle for the brain. *Journal of Physiology, Paris*, 100(1-3): 70-87.

Friston, KJ, Mattout, J & Kilner, J. 2011. Action understanding and active inference. *Biological Cybernetics*, 104(1-2): 137-160.

Friston, KJ, Samothrakis, S & Montague, R. 2012. Active inference and agency: Optimal control without cost functions. *Biological Cybernetics*, 106(8): 523-541.

Friston, KJ & Stephan KE. 2007. Free-energy and the brain. *Synthese*, 159: 417-458.

Friston, KJ, Trujillo-Barreto, N, & Daunizeau, J. 2008. DEM: a variational treatment of dynamic systems. *NeuroImage*, 41(3): 849-885.

Friston, KJ. 2003. Learning and inference in the brain. *Neural Networks*, 16(9): 1325-1352.

Friston, KJ. 2005. A theory of cortical responses. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 360(1456): 815-36.

Friston, KJ. 2008. Hierarchical models in the brain. *PLoS Computational Biology*, x4(11):e1000211. [O]. Available: https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1000211 [20 December 2017].

Friston, KJ. 2009. The free-energy principle: A rough guide to the brain?. *Trends in Cognitive Sciences*, 13(7): 293-301.

Friston, KJ. 2010. The free-energy principle: a unified brain theory?. *Nature Review Neuroscience*, 11(2): 127-138.

Frith, U & Happé, F. 1999. 'Theory of Mind and Self-Consciousness: What is it like to be Autistic?'. *Mind & Language*, 14: 1-22.

Froese, T & Paolo, EA. 2011. The enactive approach: Theoretical sketches from cell to society. *Pragmatics & Cognition*, 19(1): 1-36.

Gallagher, S & Hutto, DD. 2008. Understanding others through primary interaction and narrative practice, in *The shared mind: Perspectives on intersubjectivity*, edited by J Zlatev, TP Racine, C Sinha and E Itkonen: Amsterdam: John Benjamins:17-38.

Gallagher, S. 2001. The Practice of Mind: Theory, Simulation, or Interaction?. *Journal of Consciousness Studies*, 8(5-7): 83-107.

Gallese, V. 2014. Bodily selves in relation: Embodied simulation as second-person perspective on intersubjectivity. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 369(1644): 20130177.

Gardner, M. 1970. Mathematical Games: The fantastic combinations of John Conway's new solitaire game "Life". *Scientific American*, 223: 120-123.

Gazzaniga, M. (ed). 2009. *The cognitive neurosciences*. Cambridge: MIT Press.

Gellatly, A, Rogers, D & Sloboda, J. (eds). 1989. *Cognition in Social Worlds*. New York: McGraw-Hill.

Gershman, SJ & Daw, ND. 2012. Perception, action and utility: The tangled skein, in *Principles of brain dynamics: Global state interactions*, edited by MI Rabinovich, KJ Friston and P Varona. Cambridge: MIT Press: 293-312.

Gibson, JJ. 1979. *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin

Gładziejewski, P. 2016. Predictive coding and representationalism. *Synthese*, 193(2): 559-582.

Godfrey-Smith, P. 2006. Mental representation, naturalism, and teleosemantics, in *Teleosemantics: New Philosophical Essays*, edited by G Macdonald and D Papineau. London: Oxford University Press: [s.p].

Goldie, P. 2004. *On Personality*. London: Routledge

Goldman, A & de Vignemont, F. 2009. Is social cognition embodied?. *Trends in Cognitive Sciences*, 13: 154-159.

Goldman, AI. 1989. Interpretation Psychologized. *Mind & Language*, 4(3): 161-185.

Goldman, AI. 2012. A moderate approach to embodied cognitive science. *Review of Philosophy and Psychology*, 3(1): 71-88.

Goldman, AI. 2014. The bodily formats approach to embodied cognition, in *Current Controversies in the Philosophy of Mind*, edited by U Kriegel. London: Routledge: 91-108.

Gorea, A. (ed). 1991. *Representations of Vision: Trends and Tacit Assumptions in Vision Research*. Cambridge: Cambridge University Press.

Greene, J. 2015. *Moral Tribes: Emotion, Reason, and the Gap Between Us and Them*. Atlantic Books: London.

Greenwood, JD. (ed). 1991. *The Future of Folk Psychology: Intentionality and Cognitive Science*. Cambridge: Cambridge University Press.

Grzankowski, A. 2012. Not All Attitudes are Propositional. *European Journal of Philosophy*, 3: 374-391.

Guttenplan, S. (ed). 1994. *A Companion to Philosophy of Mind*. Oxford: Blackwell.

Harkness, DL & Keshava, A. 2017. Moving from the What to the How and Where – Bayesian Models and Predictive Processing, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Harnad, S. 1990. The symbol grounding problem. *Physica D*, 42: 335-346.

Harrison, LM, Bestmann, S, Rosa, MJ, Penny, W & Green, GGR. 2011. Time scales of representation in the human brain: weighing past information to predict future events. *Frontiers in Human Neuroscience***,** 5(37). [O]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3084444/ [15 July 2018].

Haugeland, J. 1998. *Having Thought: Essays in the Metaphysics of Mind*. Cambridge: Harvard University Press.

Hayes, P. 1978. *Naive Physics 1: Ontology for Liquids*. Working Paper 35, Institut pour les Etudes Semantiques et Cognitives, Université de Génevè. [O]. Available: http://www.issco.unige.ch/working-papers/Hayes-1978-35.pdf [26 August 2016].

Hayes, P. 1980. The Naive Physics Manifesto, in *Expert Systems in the Microelectronic Age*, edited by D Michie. Edinburgh: University of Edinburgh Press: [s.p].

Hickok, G. 2014. *The Myth of Mirror Neurons: The Real Neuroscience of Communication and Cognition*. New York: Norton.

Hirsch, JB, Mar, RA & Peterson, JB. 2013 Personal narratives as the highest level of cognitive integration. *Behavioral and Brain Sciences*, 36(3): 216-217.

Hobson, JA. & Friston, KJ. 2014. Consciousness, dreams, and inference: The Cartesian theatre revisited. *Journal of Consciousness Studie*s, 21(1-2): 6-32.

Hochstein, S & Ahissar, M. 2002. View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, 36(5): 791-804.

Hohwy, J, Roepstorff, A & Friston, KJ. 2008. Predictive coding explains binocular rivalry: An epistemological review. *Cognition*, 108(3): 687-701.

Hohwy, J. 2007. Functional Integration and the mind. *Synthese*, 159(3): 315-328.

Hohwy, J. 2012. Attention and conscious perception in the hypothesis testing brain. *Frontiers in Psychology*, 3. [O]. Available: https://dx.doi.org/10.3389/fpsyg.2012.00096 [9 September 2018].

Hohwy, J. 2013. *The predictive mind.* Oxford: Oxford University Press.

Hohwy, J. 2014. The self-evidencing brain. *Noûs*, 50(2): 259-285.

Hohwy, J. 2017. How to Entrain Your Evil Demon, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Horst, S. 2007. *Beyond Reduction: Philosophy of Mind and Post-Reductionist Philosophy of Science*. Oxford: Oxford University Press.

Hume, D. 2009. *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*. Auckland: The Floating Press

Humphrey, N. 1976. 'The Social Function of the Intellect', in *Growing Points in Ethology*, edited by PPG Bateson and RA Hinde. Cambridge: Cambridge University Press: 303-317.

Hurley, M, Dennett, DC & Adams, RB. 2011. *Inside jokes: Using humor to reverse-engineer the mind*. Cambridge: MIT Press.

Hutchins, E. 2005. Material anchors for conceptual blends. *Journal of Pragmatics*, 37: 1555-1577.

Hutchins, E. 2013.The cultural ecosystem of human cognition. *Philosophical Psychology*, 27(1): 43-49.

Hutto, DD, Kirchhoff, MD & Myin, E. 2014. Extensive enactivism: Why keep it all in?. *Frontiers in Human Neuroscience*, 8: 706.

Hutto, DD & Myin, E. 2013a. *Radicalizing Enactivism: Basic Minds without Content*. Cambridge: MIT Press.

Hutto, DD & Myin, E. 2013b. Neural representations not needed: No more pleas, please. *Phenomenology and the Cognitive Sciences*, 13(2): 241-256.

Hutto, DD & Myin, E. 2017. *Evolving Enactivism: Basic Minds Meet Content*. Cambridge: MIT Press.

Hutto, D & Ratcliffe, M. (eds). 2007. *Folk Psychology Re-assessed*. New York: Springer

Hutto, DD & Sanchez-Garcia, R. 2015. Choking RECtified: Embodied Expertise beyond Dreyfus. *Phenomenology and the Cognitive Sciences*, 14(2): 309-331.

Hutto, DD & Satne, G. 2015. The natural origins of content. *Philosophia*, 43(3): 521-536.

Hutto, DD. 1999. *The Presence of Mind*. Amsterdam: John Benjamins Publishing Company.

Hutto, DD. 2004. The limits of spectatorial folk psychology. *Mind & Language*, 19: 548-573.

Hutto, DD. 2005. Knowing what? Radical versus conservative enactivism *Phenomenology and the Cognitive Sciences*, 4(4): 389-405.

Hutto, DD. 2006a. Both Bradley and biology: Reply to Rudd, in *Radical Enactivism: Intentionality, Phenomenology and Narrative: Focus on the Philosophy of Daniel D. Hutto*, edited by R Menary. Amsterdam: John Benjamins Publishing Company: 81-106.

Hutto, DD. 2006b. Unprincipled engagements: Emotional experience, expression and response, in *Radical Enactivism: Intentionality, Phenomenology and Narrative: Focus on the Philosophy of Daniel D. Hutto*, edited by R Menary. Amsterdam: John Benjamins Publishing Company: 13-38.

Hutto, DD. 2007a. Folk Psychology without Theory or Simulation, in *Folk Psychology Re-assessed*, edited by D Hutto and M Ratcliffe. New York: Springer: 115-135.

Hutto, DD. 2007b. The Narrative Practice Hypothesis: Origins and applications of folk psychology. *Royal Institute of Philosophy Supplement*, 60: 43-68.

Hutto, DD. 2008a. *Folk Psychological Narratives: The Sociocultural Basis of Understanding Reasons*. Cambridge: MIT Press.

Hutto, DD. 2008b. The Narrative Practice Hypothesis: clarifications and implications. *Philosophical Explorations*, 11(3): 175-192.

Hutto, DD. 2011. Elementary mind minding, enactivist-style, in *Joint Attention: New Developments in Philosophy, Psychology, and Neuroscience*, edited by A Seemann. Cambridge: MIT Press: 307-341

Hutto, DD. 2013a. Enactivism: From a Wittgensteinian point of view. *American Philosophical Quarterly*, 50(3): 281-302.

Hutto, DD. 2013b. Exorcising action oriented representations: Ridding cognitive science of its Nazgul. *Adaptive Behavior*, 21(1): 142-150.

Hutto, DD. 2013c. Fictionalism about folk psychology. *Monist*, 96(4):585-607.

Hutto, DD. 2014. Contentless perceiving: The very idea, in *Wittgenstein and Perception*, edited by M Campbell and M O'Sullivan. London: Routledge: 63-83.

Hutto, DD. 2015. Basic social cognition without mindreading: Minding minds without attributing contents. *Synthese*, 194(3): 1-20.

Hutto, DD. 2016. Narrative self-shaping: A modest proposal. *Phenomenology and the Cognitive Sciences*, 15: 21-41.

Hutto, DD. 2017. Getting into predictive processing's great guessing game: Bootstrap heaven or hell? *Synthese*, 195(6): 2445-2458.

Jordan, JS. (ed). 1999. *Modeling Consciousness across the Discipline*. Maryland: University Press of America.

Jordan, MI. (ed). 1998. *Learning in graphical models*. Dordrecht: Kluwer Academic Publishers.

Joyce, J. 2008. Sv: "Bayes' Theorem". *The Stanford Encyclopedia of Philosophy*. [O]. Available: http://plato.stanford.edu/archives/fall2008/entries/bayes- theorem/ [23 August 2016].

Kahneman, D. 2011.Thinking, Fast and Slow. New York: Farrar, Straus and Giroux.

Kandel, E. 2001. The molecular biology of memory storage: A dialogue between genes and synapses. *Science*, 294: 1030-1038.

Kant, I. 1998. *Critique of Pure Reason*. Translated by P Guyer and AW Wood. Cambridge: Cambridge University Press.

Keeley, BL. (ed). 2005. *Paul Churchland (Contemporary Philosophy in Focus)*. Cambridge: Cambridge University Press.

Kinzler, KD, Dupoux, E & Spelke, ES. 2007. The native language of social cognition. *Proceedings of the National Academy of Sciences*, 104(30): 12577-12580.

Kiverstein, J & Rietveld, E. 2015. The primacy of skilled intentionality: On Hutto and Satne's The Natural Origins of Content. *Philosophia*, 43: 701-721.

Kiverstein, J. 2016. *The Routledge Handbook of Philosophy of the Social Mind*. London: Routledge.

Klein, C. 2012. Cognitive ontology and region-versus network-oriented analyses. *Philosophy of Science*, 79(5): 952-960.

Koch, C, Davis, JL, Sejnowski, TJ & Poggio, TA. (eds). 1994. *Large-Scale Neuronal Theories of the Brain*. Cambridge: MIT Press.

Lee, TS & Mumford, D. 2003. Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 20(7): 1434-1448.

Lepore, E & Pylyshyn, Z. (eds). *What is cognitive science?*. Oxford: Blackwell.

Leslie, A, Friedman, O & German, TP. 2004. Core mechanisms in 'theory of mind'. *Trends in Cognitive Sciences*, 8: 528-533.

Lewis, D, 1966. An Argument for the Identity Theory. *Journal of Philosophy*, 63: 17-25.

Lewis, D. 1970. How to Define Theoretical Terms. *Journal of Philosophy*, 67: 427-446.

Lewis, D. 1972. Psychophysical and Theoretical Identifications. *Australasian Journal of Philosophy*, 50: 249-258.

Lewis, D. 1994. Reduction of Mind, in *A Companion to Philosophy of Mind*, edited by S Guttenplan. Oxford: Blackwell: 412-431.

Mahajan, N & Wynn, K. 2012. Origins of "Us" versus "Them": Prelinguistic infants prefer similar others. *Cognition*, 124(2): 227-233.

Maibom, HL. (ed). 2014. *Empathy and Morality*. Oxford: Oxford University Press.

Malle, B, Moses, LJ & Baldwin, DA. (eds). 2001. *Intentions and Intentionality*. Cambridge: MIT Press.

Malle, B. 2004. *How the Mind Explains Behavior: Folk Explanations, Meaning, and Social Interaction*. Cambridge: MIT Press.

Maloney, LT & Mamassian, P. 2009. Bayesian decision theory as a model of visual perception: Testing Bayesian transfer. *Visual Neuroscience*, 26: 147-155.

Maloney, LT & Zhang, H. 2010. Decision-theoretic models of visual perception and action. *Vision Research*, 50: 2362-2374.

Manzotti, R. 2011. *Situated Aesthetics: Art beyond the Skin*. Exeter: Imprint Academic.

Matthen, M. (ed). 2013. *Oxford Handbook of the Philosophy of Perception*. Oxford: Oxford University Press.

Menary, R & Kirchhoff, M. 2013. Cognitive transformations and extended expertise. *Educational Philosophy and Theory*, 46(6): 610-623.

Menary, R. (ed). 2006. *Radical Enactivism: Intentionality, Phenomenology and Narrative: Focus on the Philosophy of Daniel D. Hutto*. Amsterdam: John Benjamins Publishing Company.

Menary, R. 2007. *Cognitive integration: Mind and cognition unbounded*. New York: Palgrave Macmillan.

Menary, R. 2010a. Dimensions of mind. *Phenomenology and the Cognitive Sciences*, 9(4): 561-578.

Menary, R. (ed). 2010b. *The Extended Mind*. Cambridge: MIT Press.

Menary, R. 2012. Cognitive practices and cognitive character. *Philosophical Explorations*, 15(2): 147-164.

Menary, R. 2013. The enculturated hand, in *The hand, an organ of the mind: What the manual tells the mental*, edited by Z Radman. Cambridge: MIT Press: 349-367.

Menary, R. 2014. Neural plasticity, neuronal recycling and niche construction. *Mind & Language*, 29(3): 286-303.

Menary, R. 2015a. Mathematical cognition: A case of enculturation, in *Open MIND: 25*(T), edited by T Metzinger and JM Windt. Frankfurt am Main: MIND Group: [s.p.].

Menary, R. 2015b. What? Now. Predictive coding and enculturation: A Reply to Regina E. Fabry, in *Open MIND: 25*(R), edited by T Metzinger and JM Windt. Frankfurt am Main: MIND Group: [s.p.].

Menary, R. 2016. Pragmatism and the pragmatic turn in cognitive science, in *Where is the action? The pragmatic turn in cognitive science*, edited by AK Engel, K Friston and D Kragic. Cambridge: MIT Press: 219-237.

Merleau-Ponty, M. 1962. *Phenomenology of Perception.* Translated by C Smith. London: Routledge.

Metzinger, T & Wiese, W. (eds). 2017. *Philosophy and predictive processing*. Frankfurt am Main: MIND Group.

Metzinger, T & Windt, JM. (eds). 2015. *Open MIND.* Frankfurt am Main: MIND Group.

Metzinger, T. 2004. *Being no one: The self-model theory of subjectivity*. Cambridge: MIT Press.

Metzinger, T. 2013. The myth of cognitive agency: Subpersonal thinking as a cyclically recurring loss of mental autonomy. *Frontiers in Psychology*, 4: 931.

Metzinger, T. 2017. The Problem of Mental Action - Predictive Control without Sensory Sheets, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Michie, D. (ed). 1980. *Expert Systems in the Microelectronic Age*. Edinburgh: University of Edinburgh Press.

Millikan, RG. 1984. *Language, Thought, and Other Biological Categories*. Cambridge: MIT Press.

Millikan, RG. 1993. *White Queen Psychology and Other Essays for Alice*. Cambridge: MIT Press.

Millikan, RG. 1995. Pushmi-Pullyu Representations. *Philosophical Perspectives*, 9:185-200.

Millikan, RG. 2000. *On Clear and Confused Ideas: An Essay about Substance Concepts*. Cambridge: Cambridge University Press.

Millikan, RG. 2004. *Varieties of Meaning: The 2002 Jean Nicod Lectures*. Cambridge: MIT Press.

Millikan, RG. 2005. *Language: A Biological Model*. Oxford: Oxford University Press.

Miłkowski, M. 2015. The hard problem of content: Solved (long ago). *Studies in Logic, Grammar and Rhetoric*, 41(1): 73-88.

Mölder, B. 2016. Mind and Folk Psychology: A Partial Introduction. *Studia Philosophica Estonica*, 9(1): 1-21.

Morgan, R. 2002. *Altered Carbon*. London: Orion.

Moutoussis, M, Fearon, P, El-Deredy, W, Dolan, RJ. & Friston, KJ. 2014. Bayesian inferences about the self (and others): A review. *Consciousness and Cognition*, 25: 67-76.

Moyal-Sharrock, D, Munz, V & Coliva, A. (eds). 2015. *Mind, Language, and Action: Proceedings of the 36th Wittgenstein Symposium*. Berlin: De Gruyter.

Muckli, L. 2010. What are we missing here? Brain imaging evidence for higher cognitive functions in primary visual cortex V1. *International Journal of Imaging Systems Technology*, 20: 131-139.

Muller, HD. 2014. Naturalism and intentionality, in *Contemporary Philosophical Naturalism and Its Implications*, edited by B Bashour and HD Muller. London: Routledge: 155-181.

Mumford, D. 1992. On the computational architecture of the neocortex. II: The role of cortico-cortical loops. *Biological Cybernetics*, 66(3): 241-251.

Myin, E & Veldeman, J. 2011. Externalism, mind and art, *Situated Aesthetics: Art beyond the Skin*, edited by R Manzotti. Exeter: Imprint Academic: 51-85.

Myin, E & Hutto, DD. 2015. REC: Just radical enough. *Studies in Logic, Grammar and Rhetoric*, 41(1): 61-71.

Neal, RM. & Hinton, G. 1998. A view of the EM algorithm that justifies incremental, sparse, and other variants, in *Learning in graphical models*, edited by MI Jordan. Dordrecht: Kluwer Academic Publishers: 355-368.

Needham, R. 1972. *Belief, Language, and Experience*. Oxford: Blackwell.

Nichols, S & Stich, S. 2003. *Mindreading: An Integrated Account of Pretence, Self-Awareness and Understanding of Other Minds*. Oxford: Oxford University Press.

Noë, A. 2004. *Action in Perception*. Cambridge: MIT Press.

Noë, A. 2009. *Out of Our Heads*. New York: Hill and Wang.

Noë, A. 2012. *Varieties of Presence*. Cambridge: Harvard University Press.

O'Brien, G & Opie, J. 2009. The role of representation in computation. *Cognitive Processing*, 10(1): 53-62.

O'Brien, G & Opie, J. 2015. Intentionality lite or analog content?. *Philosophia*, 43(3): 723-730.

Olafsdottir, F, Barry, C, Saleem, A, Hassabis, D & Spiers, H. 2015. Hippocampal place cells construct reward related sequences through unexplored space. *eLife*, 4:e06063. [O]. Available: https://elifesciences.org/articles/06063 [25 September 2018].

Papineau, D. 1987. *Reality and Representation*. Oxford: Oxford University Press.

Papineau D. 2003. *The Roots of Reason: Philosophical Essays on Rationality, Evolution and Probability*. Oxford: Oxford University Press.

Papineau, D. 2010. Review of *What Darwin Got Wrong*, by Jerry Fodor and Massimo Piattelli-Palmarini. *Prospect*, 168: 83-84.

Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Francisco: Morgan Kaufmann Publishers.

Pettigrew, R. 2015. Pluralism About Belief States. *Aristotelian Society Supplementary Volume*, 89(1): 187-204.

Phillips, JG, Ogeil, RP & Best, C. 2009. Motor constancy and the upsizing of handwriting. *Human Movement Science*, 28(5): 578-587.

Pietroski, M. 1992. Intentionality and teleological error. *Pacific Philosophical Quarterly*, 73: 267-282.

Poldrack, RA. 2006. Can cognitive processes be inferred from neuroimaging data?. *Trends in Cognitive Sciences*, 10(2): 59-63.

Poldrack, RA. 2010. Mapping mental function to brain structure: How can cognitive neuroimaging succeed?. *Perspectives on Psychological Science*, 5(6): 753-761.

Price, CJ & Friston, KJ. 2005. Functional ontologies for cognition: The systematic definition of structure and function. *Cognitive Neuropsychology*, 22(3-4): 262-275.

Price, H. 2013. *Expressivism, Pragmatism and Representationalism*. Cambridge: Cambridge University Press.

Pylyshyn, Z. 1999. What's in your mind, in *What is cognitive science?*, edited by E Lepore and Z Pylyshyn. Oxford: Blackwell: 1-25.

Rabinovich, MI, Friston, KJ & Varona, P. 2012. *Principles of brain dynamics: Global state interactions*. Cambridge: MIT Press.

Radman, Z. (ed). 2013. *The hand, an organ of the mind: What the manual tells the mental*. Cambridge: MIT Press.

Ramsey, WM. 2007. *Representation Reconsidered*. Cambridge: Cambridge University Press.

Ramsey, WM. 2014. Must cognition be representational?. *Synthese*, 194(11): [s.p.].

Ramstead, MJD, Veissiere, SPL & Kirmayer, LJ. 2016. Cultural affordances: Scaffolding local worlds through shared intentionality and regimes of attention. *Frontiers in Psychology*, 7: 1090.

Rao, RPN & Ballard, DH. 1999. Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1): 79-87.

Rauss, K, Schwartz, S & Pourtois, G. 2011. Top-down effects on early visual processing in humans: A predictive coding framework. *Neuroscience and Biobehavioral Reviews*, 35(5): 1237-1253.

Ratcliffe, M. 2007a. From Folk Psychology to Commonsense, in *Folk Psychology Re-assessed*, edited by D Hutto and M Ratcliffe. New York: Springer: 223-243.

Ratcliffe, M. 2007b. *Rethinking Commonsense Psychology: A Critique of Folk Psychology, Theory of Mind and Simulation*. London: Palgrave Macmillan.

Ratcliffe, M. 2008. Farewell to Folk Psychology: A Response to Hutto. *International Journal of Philosophical Studies*, 16:3: 445-451.

Rescorla, M. 2013. Bayesian perceptual psychology, in *Oxford Handbook of the Philosophy of Perception*, edited by M Matthen. Oxford: Oxford University Press: 694-716.

Rietveld, E & Kiverstein, J. 2014. A rich landscape of affordances. *Ecological Psychology*, 26(4): 325-352.

Ritchie, J. 2008. *Understanding Naturalism*. Cape Town: Acumen.

Rizzolatti, G & Sinigaglia, C. 2010. The functional role of the parieto-frontal mirror circuit: Interpretations and misinterpretations. *Nature Reviews: Neuroscience*, 11: 264-274.

Roepstorff, A, Niewöhner, J & Beck, S. 2010. Enculturing brains through patterned practices. *Neural Networks*, 23(8-9): 1051-1059.

Rosenberg, A. 2013. How Jerry Fodor slid down the slippery slope to anti-Darwinism, and how we can avoid the same fate. *European Journal of Philosophy of Science*, 3: 1-17.

Rosenberg, A. 2014. Disenchanted naturalism, in *Contemporary Philosophical Naturalism and Its Implications*, edited by B Bashour and HD Muller. London: Routledge: 17-36.

Ross, D, Spurrett, D, Kincaid, H & Stephens, GL. (eds). 2007. *Distributed Cognition and the Will Individual Volition and Social Context*. Cambridge: MIT Press.

Rowlands, M. 2009. The extended mind. *Zygon*, 44: 628-641.

Rowlands, M. 2010. *The New Science of the Mind: From Extended Mind to Embodied Phenomenology*. Cambridge: MIT Press.

Roy, JM. 2015. Anti-Cartesianism and anti-Brentanism: The problem of anti-representationalist intentionalism. *Southern Journal of Philosophy*, 53: 90-125.

Russon, A, Bard, K & Taylor Parker, S. 1995. *Reaching into Thought: The Minds of the Great Apes*. Cambridge: Cambridge University Press.

Ryle, G. 2009. *The Concept of Mind*. London: Routledge.

Sapolsky, R. 2017. *Behave: The Biology of Humans at Our Best and Worst*. London: Vintage

Scholl, B & Leslie, A. 1999. Modularity, development, and 'Theory of Mind'. *Mind & Language*, 14: 131-153.

Schwartenbeck, P, Fitzgerald, THB, Mathys, C, Dolan, R & Friston, KJ. 2014. The Dopaminergic Midbrain Encodes the Expected Certainty about Desired Outcomes. *Cerebral Cortex*, 25: 3434-2445.

Searle, J. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press.

Seemann, A. (ed). 2011. *Joint Attention: New Developments in Philosophy, Psychology, and Neuroscience*. Cambridge: MIT Press

Segal, G. 1996. The modularity of theory of mind, in *Theories of theories in mind*, edited by P Carruthers and P Smith. Cambridge: Cambridge University Press: 141-157.

Sella, G & Hirsh, AE. 2005. The application of statistical physics to evolutionary biology. *Proceedings of the National Academy of Sciences of the United States of America*, 102(27): 9541-9546.

Sellars, W. 1963. *Science, Perception and Reality*. London: Routledge.

Seth, AK. 2014a. A predictive processing theory of sensorimotor contingencies: Explaining the puzzle of perceptual presence and its absence in synaesthesia. *Cognitive Neuroscience*, 5(2): 97-118

Seth, AK. 2014b. Interoceptive inference: From decisionmaking to organism integrity. *Trends in Cognitive Sciences*, 18(6): 270-271.

Seth, AK. 2015. The Cybernetic Bayesian Brain: From Interoceptive Inference to Sensorimotor Contingencies, in *Open MIND: 35*(T), edited by T Metzinger and JM Windt. Frankfurt am Main: MIND Group: [s.p.].

Shapiro, L. 2014. Review: Radicalizing Enactivism: Basic Minds without Content, by Daniel D. Hutto and Erik Myin. *Mind*, 123(489): 213-220.

Shuler, MG & Bear, MF. 2006. Reward timing in the primary visual cortex. *Science*, 311(5767): 1606-1609.

Sims, A. 2017. The Problems with Prediction: The Dark Room Problem and the Scope Dispute, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Sober, E. 2010. Natural Selection, Causality, and Laws: What Fodor and Piatelli-Palmarini Got Wrong. *Philosophy of Science*, 77(4): 594-607.

Solomon, RC. 2003. *Not Passion's Slave: Emotions and Choice*. Oxford: Oxford University Press.

Solomon, RC. (ed). 2004. *Thinking about Feeling: Contemporary Philosophers on Emotions*. Oxford: Oxford University Press.

Spratling, MW. 2016. A review of predictive coding algorithms. *Brain and cognition*, 112: 92-97.

Sterelny K. 2003. *Thought in a Hostile World*. Oxford: Blackwell.

Sterelny, K. 2015. Content, control and display: The natural origins of content. *Philosophia*, 43: 549-564.

Stich, S & Nichols, S. 2003. Folk Psychology, in *The Blackwell Guide to Philosophy of Mind*. Oxford: Blackwell: 235-255

Stich, S & Ravenscroft, I. 1994. What is folk psychology?. *Cognition*, 50: 447-468.

Stich, S & Warfield, TA. (eds). 2003. *The Blackwell Guide to Philosophy of Mind*. Oxford: Blackwell.

Stich, S. 1983. *From Folk Psychology to Cognitive Science: The Case Against Belief*. Cambridge: MIT Press.

Stich, S. 1990. *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge: MIT Press.

Stock, A. & Stock, C. 2004. A short history of ideo-motor action. *Psychological Research*, 68: 176-188.

Strawson, G. 1994. *Mental Reality*. Cambridge: MIT Press.

Surian, L & Leslie, AM. 1999. Competence and performance in false belief understanding: A comparison of autistic and normal 3-year-old children. *British Journal of Developmental Psychology*, 17: 141-155.

Sutton, J, McIlwain, D, Christensen, W & Geeves, A. 2011. Applying intelligence to the reflexes: Embodied skills and habits between Dreyfus and Descartes. *Journal of the British Society for Phenomenology*, 42(1): 78-103.

Sutton, J. 2015. Remembering as public practice: Wittgenstein, memory, and Distributed Cognitive Ecologies, in *Mind, Language, and Action: Proceedings of the 36th Wittgenstein Symposium*, edited by D Moyal-Sharrock, V Munz, and A Coliva. Berlin: De Gruyter: 409-444.

Swanson, LR. 2016. The predictive processing paradigm has roots in Kant. *Frontiers in Systems Neuroscience*, 10: 79.

Thompson, E. 2007. *Mind in Life: Biology, Phenomenology, and the Sciences of the Mind*. Cambridge: Harvard University Press.

Todorov, E & Jordan, MI. 2002. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11): 1226-1235.

Todorov, E. 2009. Parallels between sensory and motor information processing, in *The cognitive neurosciences*, edited by M Gazzaniga. Cambridge: MIT Press: 613-624.

Toussaint, M. 2009. Probabilistic inference as a model of planned behavior. *Künstliche Intelligenz*, 23: 23-29.

Tribus, M. 1961. *Thermodynamics and thermostatics: An introduction to energy, information and states of matter, with engineering applications*. New York: D Van Nostrand.

Ullman, S. 1991. Tacit assumptions in the computational study of vision, in *Representations of Vision: Trends and Tacit Assumptions in Vision Research*, edited by A Gorea. Cambridge: Cambridge University Press: 305-317.

Van Leeuwen, C, Verstijnen, IM & Hekkert, P. 1999. Common unconscious dynamics underlie uncommon conscious effect: A case study in the iterative nature of perception and creation, in *Modeling Consciousness across the Disciplines*, edited by JS Jordan. Maryland: University Press of America: [s.p.].

Varela, FJ, Thompson, E & Rosch, E. 1991. *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge: MIT Press.

Visalberghi, E & Limongelli, L. 1995. Acting and Understanding: Tool Use Revisited through the Minds of the Capuchin Monkeys, in *Reaching into Thought: The Minds of the Great Apes*, edited by A Russon, K Bard and S Taylor Parker. Cambridge: Cambridge University Press: 57-79.

Von Helmholtz, H. 1985. *Helmholtz's treatise on physiological optics*. Birmingham: Gryphon Editions.

Weiss, Y, Simoncelli, EP. & Adelson, EH. 2002. Motion illusions as optimal percepts. *Nature Neuroscience*, 5(6): 598-604.

Wellman, HM & Phillips, A. 2001. Developing intentional understandings, in *Intentions and Intentionality*, edited by B Malle, LJ Moses and DA Baldwin. Cambridge: MIT Press: 125-148.

Wheeler, M. 2010. In defense of extended functionalism, in *The Extended Mind*, edited by R Menary. Cambridge: MIT Press: 245-270.

Whiten, A & Byrne, R. (eds). 1997. *Machiavellian Intelligence II: Extensions and Evaluations*. Cambridge: Cambridge University Press.

Whiten, A. (ed). 1991. *Natural Theories of Mind*. Oxford: Blackwell.

Wiese, W & Metzinger, T. 2017. Vanilla PP for Philosophers: A Primer on Predictive Processing, in *Philosophy and Predictive Processing*, edited by T Metzinger and W Wiese. Frankfurt am Main: MIND Group: [s.p.].

Wiese, W. 2016. Action is enabled by systematic misrepresentations. *Erkenntnis*, 82(6): [s.p.].

Wimmer, H. and Perner, J. 1983. 'Beliefs about Beliefs: Representation and Constraining Function of Wrong Beliefs in Young Children's Understanding of Deception'. *Cognition*, 13: 103-128.

Zlatev, J, Racine, TP, Sinha, C & Itkonen, E. (eds). 2008. *The shared mind: Perspectives on intersubjectivity*. Amsterdam: John Benjamins Publishing Company.