

MEDITATION, ENACTIVISM AND
INTROSPECTIVE TRAINING

by

MICHAEL DAVID ROBERTS

A thesis submitted to the University of Birmingham for the degree of DOCTOR
OF PHILOSOPHY

Department of Philosophy
School of Philosophy, Theology and Religion
College of Arts and Law
University of Birmingham
September 2018

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

This PhD thesis concerns introspective approaches to the study of the mind. Across three standalone papers, I examine the significance of introspective data and advise on appropriate kinds of training for the production of such data. An overview document first introduces major themes, methods and arguments of the thesis. Paper 1 then begins the argumentative work, interrogating the constraining function of introspection in cognitive science. Here, I evaluate “enactivist” claims about the significance of introspection, clarifying central enactivist suggestions to draw out the broader importance of introspection in science and philosophy. Paper 2 then examines the proposed employment of Buddhist meditation practices in the production of rigorous introspective data. I defend such proposals against concerns that meditators yield ungeneralizable data, given the transformative character of these attention-training techniques. I argue that some meditation-trained transformations are actually epistemically-beneficial, undermining popular associations between transformation and “distortion”. Paper 3 then reviews difficulties involved in integrating meditative training into research. I emphasise the importance of specific contextual supports to meditation as critical ingredients of introspective proficiency, showing how difficulties replicating these threaten to limit the scope of meditation’s scientific benefits. I layout methodological responses to this that can maximise meditation’s positive impact going forwards.

For my parents

Acknowledgements

This thesis owes much to many. My biggest thanks go to my supervisors. I am deeply grateful for the support of Scott Sturgeon, whose enviable well of illustrative metaphors proved an excellent catalyst to my research, and whose enthusiastic direction has helped me to turn small kernels of intuition into well-formed arguments. I am also thankful for the patience and generosity of Maja Spener, whose meticulous and challenging advice has transformed my writing style no end, greatly improving the clarity of the ideas expressed here. My gratitude likewise goes to those who have read drafts of my work—including Susanna Siegel, Christian Coseru and Parimal Patil—and to the University of Birmingham postgraduate philosophy community, especially Ray Schuur, Stewart Yarlett, Casey Elliott and Jen Elliott. This research would not have been possible without the generous funding of the Arts and Humanities Research Council and the Royal Institute of Philosophy. Final thanks go to my family for their unending kindness and encouragement.

List of Contents

Thesis Overview.....	1
Themes, Aims and Background	1
Paper Summaries.....	5
Paper 1 – “Phenomenological Constraints: A Problem for Radical Enactivism”.....	5
Paper 2 – “Meditation and Introspection: Insight Through Transformation”	8
Paper 3 – “Introspective Training: A Broader Path?”	10
Conclusions and Future Directions:	12
Take Introspection Seriously	12
Develop “Transformational” Methods of Inquiry.....	13
Transformation and Insight are Co-supportive	14
Phenomenological Constraints: A Problem for Radical Enactivism	17
1 Introductory Comments	17
2 Basic Features of Enactivism	21
3 The Structural Resemblance Constraint.....	22
4 Reduction and Identity	27
5 Enactivist Phenomenology and its Consequences	33
6 Radical Enactivism and SRC	37
7 Radical Enactivism: Correcting Inconsistencies	41
8 Questions Going Forwards and Neurophenomenology	49
Meditation and Introspection: Insight through Transformation.....	56
Introduction.....	56
1 Meditation: Proposals and Objections	62
1.1 Buddhist Roots	62
1.2 Scientific Uses: The Meditative Turn in Cognitive Science	65
1.3 Objections	68
2 Meditation and the Cognitive Psychology of Attention.....	72
2.1 Attentional Skill in Focused-Attention	72
2.2 The Science of Top-Down Attentional Control	75
2.3 Focused-Attention in Focus	79
2.4 Effects and Benefits of Focused Introspective-Attention	82
2.4.1 Phenomenological Effects.....	82
2.4.2 Clarity Considered	84

2.4.3	Introspective Value	86
2.5	A Foundation to Build On.....	94
3	Further Distortion Concerns.....	95
3.1	Objectifying the Subjective.....	95
3.2	Stilling the Stream.....	97
3.3	Complex Experiences, Intentionality and Different Probes.....	99
3.4	Increased Richness, Increased Detail	103
3.5	Conceptual Tainting	104
3.6	A Route to Handling Distortion	106
4	Using Attentional Skill: Pedagogical and Scientific Considerations.....	108
4.1	Refining Concentration	109
4.2	Appropriate Labels: Broad, Simple and Brief	115
4.3	Piecemeal Progress.....	119
4.4	Concentration as a Preliminary Factor	127
4.5	Ground-Clearing	131
5	Conclusions and Future Directions	137
	Introspective Training: A Broader Path?	141
	Introduction.....	141
1	Meditation and Introspective Training.....	145
1.1	The Buddhist Context	145
1.2	Meditation and Cognitive Scientific Methods	152
2	The Significance of Context.....	156
3	A Broader Path.....	160
3.1	Standalone benefits	160
3.1.1	Meditative Training.....	160
3.1.2	Ethical Training.....	162
3.1.3	Intellectual Training	163
3.2	Supportive benefits.....	164
3.2.1	Pre-empting Difficulty	165
3.2.2	Developing a Practical Toolkit.....	166
4	Methodological Repercussions	170
4.1	Summary Findings: A Dilemma	170
4.2	A Methodological Choice	172
4.2.1	A Broader Path.....	173
4.2.2	A Narrow Path.....	176

4.3 Weighing the Choice.....	178
5 Conclusions and Routes Forward.....	183
References.....	188

List of Figures and Tables

Figures

Paper 1:

Fig. i. A sensorimotor loop (arrows show direction of influence) P. 21

Tables

Paper 2:

Fig. i. Meditation Schematics P. 64

Paper 3:

Fig. i. Meditation Schematics P. 148

Abbreviations from the Pali Canon

AN Anguttara Nikaya

MN Majjhima Nikaya

SN Samyutta Nikaya

Note on Referencing

Individual Discourses from the Anguttara Nikaya and Samyutta Nikaya are referenced by division number and sutta number. For example, AN 3:101-102 indicates Discourses #101-102 of Division #3 of the Anguttara Nikaya.

Discourses the Majjhima Nikaya are referenced by sutta number. For example, MN 20 indicates Discourse #20 of the Majjhima Nikaya.

Discourses and translations are sourced from the following works (with specific sources indicated in text)

AN / MN Bodhi, B. (2005) *In the Buddha's Words: An Anthology of Discourses from the Pali Canon*. Somerville, MA: Wisdom Publications.

SN Access to Insight (ed.) (2013) Samyutta Nikaya. *Access to Insight (BCBS Edition)*. Retrieved August 8 2018. Available

<http://www.accesstoinsight.org/tipitaka/sn>

Thiradhammo, A. (2014) *Working with the Five Hindrances*. Belsay, UK: Aruno Publications.

Thesis Overview

Themes, Aims and Background

This collection of papers examines introspective approaches to the study of the mind – their significance, their utility, and the kinds of training required to pursue them. By “introspective approaches”, I refer to those means of interrogating consciousness that use subjective reports as principal methodological tools, employing judgements about the mind made “from the inside”, as it is sometimes put (Spener, 2015, p. 300). In a series of standalone works, this PhD pursues two principal themes on this topic.

First, I evaluate the warrant for affording subjective reports a place of significance in the study of the mind by engaging with contemporary arguments emerging within the *enactive approach* to cognitive science and philosophy (paper 1). Drawing influence from continental phenomenology, Buddhist thought, and non-representational developments in cognitive science, enactivists account for the mind in terms of embodied interactions between organisms and their environments (see Varela, Thompson and Rosch, 1991/2017; Stewart, Gapenne and Di Paolo, 2014), and they distinguish themselves as especially vocal about the methodological importance and advantages of introspective methods (e.g. Colombetti, 2014; Thompson, 2015; Thompson and Cosmelli, 2011). The first aim of this PhD is to unpack these phenomenological commitments of enactivists using more conventional theoretical material from analytic philosophy of mind, helping to draw out both the broader significance of subjective reports and the way they are held to motivate enactivist theories.

Second, I critically examine one form of mental training identified as a suitable support for introspective endeavours – Buddhist meditation (papers 2 & 3). Meditation comprises a set

of attentional training practices, which have long been held to sensitize their practitioners to experience within the Buddhist traditions (Thompson, 2015; Lutz et al., 2015). Such practices are increasingly touted as valuable resources for introspectively-minded researchers, with both enactivist and non-enactivist theorists pushing them as means of training the gestures and attitudes required for rigorous introspective methods (e.g. Colombetti, 2014, chpt. 6; Varela, Thompson and Rosch, 1991/2017; Wallace, 1999, 2009, 2012). The second central aim of this PhD is to critically assess these proposals, offering a more balanced account of the methodological role such practices can play within science and philosophy, through a more careful and in-depth engagement with Buddhist theoretical and pedagogical literature.

Pursuing these objectives has made it necessary to engage some very different areas of thought, each of which has grappled with theoretical and methodological questions about introspection in very different ways. And the attempt to synthesize ideas from such disparate traditions faced immediate challenges. First, it was important to address my materials in the depth necessary to avoid over-simplification, whilst resisting overly technical presentations that would distract from central arguments and make it difficult to bring ideas into dialogue. For clarity, my engagement with the enactive approach has thus been restricted largely to enactive accounts of perceptual experience. Moreover, I have focused on externalist leanings present here—the tendency to view perception as something extending constitutively beyond the subject and into the environment—avoiding many other conditions that enactivists hold critical, including biological autonomy and self-organization (see Stewart, Gapenne and Di Paolo, 2014; Barandiaran, 2016).

The need for a well-defined and manageable set of raw materials also presented a particular challenge for my engagement with Buddhist meditation. Research in this area typically proceeds from the fields of Buddhist studies or religious studies, which favour rather

extended and scholastic approaches to their theme. More recently, meditation has also become a topic of interest in cognitive psychology (see e.g. Jha, Krompinger and Baime, 2007; Hölzel et al., 2011; Lutz et al., 2008; 2015) which, in contrast, inclines towards specialised examinations of the effects of practice upon our neural architecture, being rather light in engagement with the web of traditions and beliefs that give such practices their meaning, function and character. This left relatively few models for analytic philosophical engagements with meditation that could draw out its essential features in a manner sensitive to surrounding theoretical frameworks, without getting bogged down in scholarly detail.¹

Given the need for a sophisticated yet *accessible* presentation of Buddhist thought, I have prioritised source materials that are relatively unambiguous, generally held as plainly representative of major themes of Buddhist thought, and which present less formidable interpretative challenges to the interdisciplinary minded philosophical writer (and reader). I have focused largely upon theoretical proposals about meditation and its function in the early Pali Canon (in particular, discourses known as the *suttas*), whose contents are fairly well unified and generally held closest to the teachings of the historical Buddha himself (Albahari, 2006, p. 1; Gethin, 1998, p. 1). This is to present what Gethin (p. 3) refers to as some “foundations” of Buddhist thought – ideas and practices that are fundamental to, and largely consistent across, the major Buddhist schools. I pair this with an engagement of the pedagogical literature from the contemporary “Insight” meditation movement, itself drawing principal inspiration from early Buddhist discourses and the more conservative Theravāda tradition (Sharf, 2015, pp. 472-473; Cousins, 1996)

Given my selective approach to the Buddhist literature on meditation, it must be acknowledged that my broad references here to “Buddhist meditation” and “Buddhist thought”

¹ For some good exceptions to this, see Albahari (2006), Thompson (2015) and Wallace (1999, 2008, 2012)

will conceal many significant idiosyncrasies in the way that meditation is employed and elaborated across the different Buddhist traditions, not to mention some controversies over even the earliest Buddhist views of meditation (see Griffiths, 1981). Nonetheless, this is a necessary strategic manoeuvre, highlighting one orthodox interpretation of meditation's essential character. And it should be understood as one of several plausible means of unpacking the fundamentals of Buddhist meditation, upon which the different traditions of Buddhism later alter, extend and elaborate to produce their own idiosyncratic meditative programmes.

It has also been important to the success of these works to remain sensitive to the different contexts of the ideas I engage. In the enactive approach and broader western traditions of inquiry, interest with introspective methods is motivated largely by an epistemological drive to understand the mind and its components. In the contemplative traditions of Buddhism this is present too. However, such an endeavour is always traditionally paired with, and indeed founded upon, a more central interest in the soteriological function that introspection can serve – a principal function of introspective inquiry here is not knowledge itself, but to acquire kinds of knowledge that are conducive to *liberation* of the self (and the other) from suffering. The Buddhist student aims to familiarise themselves with the mental conditions productive of suffering through meditation, precisely *so that these can be overcome*. Being mindful of this broader goal is essential, for one must be careful to note that the practice of meditation ultimately favours transformation of the individual's psyche, over the accumulation of knowledge. This has been increasingly stressed by Thompson as a critical and often neglected point, when it comes to understanding both the personal and epistemological utility of meditation practices (Varela, Thompson and Rosch, 1991/2017, pp. xxiv-xxvi). In fact, this very tension serves as a basis for one of the primary objections to the use of meditation in

science, which I engage in paper 2, a nuanced treatment of which is fundamental to any good account of meditation's value here.

Another aim of this work, however, has been to show that the transformative and the epistemological orientations of meditation practice can actually sit well together – they can be co-supportive. And in paper 3, I argue that an understanding of this mutually-beneficial relationship is critical for assessing the scope of meditation's benefits to science and philosophy, and for advising on how that can be maximised. In this way, I hope to better stress how the introspective and the soteriological components of meditative practice need not be seen as problematic competitors, but can be beneficial partners.

Paper Summaries

Paper 1 – “Phenomenological Constraints: A Problem for Radical Enactivism”

In paper 1, published in *Phenomenology and the Cognitive Sciences* (Roberts, 2018b), I seek to clarify and support the important role assigned to subjective reports in the construction of early enactivist ideas. Such an undertaking helps me both to crystallise some broader important points about the constraining function of introspection in science and philosophy and to advise how enactivism can go forwards in a manner congruent with its early commitments.

Thompson (2011) notes that it is central to enactivist thought that ‘experience is not an epiphenomenal side issue, but central to any understanding of the mind and needs to be investigated in a careful phenomenological manner’ (p. 20). This enactivist foregrounding of introspective (or phenomenological) methods reflects a much broader rehabilitation of subjective reports into cognitive science. Increasingly, it has been recognised that introspection is an indispensable source of, and primary means of characterising, the explanatory targets that

constrain and guide the explanatory endeavours of scientists and philosophers (Chalmers, 1999; Jack and Roepstorff, 2003, 2004; Kriegel, 2015, p. 21; Thompson, 2015). And, in this paper, I delineate the precise manner in which this constraint is understood by enactivist thinkers. I argue that a majority of enactivists are committed to something I call the “structural resemblance constraint” (SRC). SRC stipulates that sub-personal models or explanations of mental phenomena, in science and philosophy, ought to implicate certain global structures that resemble structural features of experience described from the first-person perspective.

This commitment to the importance of structural resemblances between first and third-person accounts is critical to wider enactivist suggestions that descriptions of experience can actually help researchers in the *identification* of biological processes important to experience. This larger project is known as ‘neurophenomenology’ (Varela, 1996). Holding structural resemblances as a necessary feature of satisfying accounts of the mental enables researchers to *search for* biological processes displaying structural properties (particularly temporal properties) resembling those featuring in subjective accounts and thereby implicate such discovered bodily processes as central biological underpinnings of the mental phenomenon in question (see e.g. Lutz et al., 2002).

In this broader neurophenomenological method, subjective reports used to guide investigation are then to be refined and precisified with the help of empirical data, which can suggest alterations to subjective accounts and further, more detailed structures to seek out and investigate introspectively. In this way, subjective reports can be honed in a way that catalyzes further empirical findings, generating an ongoing, back-and-forth influence between the first and third-person. Francisco Varela (1996) argued that such a project of ‘mutual constraints’ could help researchers gradually narrow in upon a more complete and detailed account of

conscious experience that was better able to address the purported ‘explanatory gap’ between the physical and the mental.

In my paper, I focus upon enactivist treatments of perception to illustrate the first stage in this enactivist methodology. I show that enactivists pedestalize phenomenological reports which construe perception as a type of *felt bodily engagement with the environment* and use this in combination with SRC to favour sub-personal accounts of perception spanning brain, body and environment. I support the enactivist endorsement of SRC by introducing more familiar explanatory constraints delineated in mainstream philosophy of science. In so doing, I seek to make this often-misunderstood commitment of enactivist thought more intelligible to analytically-minded thinkers and better concretize the important constraining role of introspection in scientific and philosophical theorizing.

With SRC unpacked and motivated, I then demonstrate how this constraint poses a problem for one recent version of enactivism – the ‘Radical Enactivism’ of Hutto and Myin (2013, 2017). Reviewing Radical Enactivism’s claims about the appropriate conditions upon explanations of the mental, along with their purported openness to *internalist* accounts of perceptual experience, reveals inconsistencies in their views about introspective constraints, and a failure to abide by SRC. I show that, in order to retain the most explanatorily interesting feature of the enactivist project – its ability to *do better justice* to the phenomenology of experience, and thereby make progress against the explanatory gap – Hutto and Myin ought to reject internalism (according to which, perceptual experience can be exhaustively explained in terms of processes internal to the subject). Only this way can enactivism be armoured from a theoretical fall back into less interesting and less satisfying territory. This is to use Hutto and Myin’s ideas as an instrument with which to draw out some important points about the status, function and value of introspective methods in science and philosophy.

Paper 2 – “Meditation and Introspection: Insight Through Transformation”

If introspective reports have the power to constrain and guide research into the mind, this begs the question: which ones? Which kinds of reports have the rigor, precision and reliability required to act as suitable subjective data? And what are the criteria for ascertaining this? Given the growing acceptance of the inevitability and advantages of introspection, these questions have been subject to much interest in the past two decades (see Bayne and Spener, 2010; Depraz, Varela and Vermersch, 2003; Spener, 2015; Shear and Varela, 1999). There is growing consensus that the reports we prioritise should emerge from subjects possessing particular mental capacities that help them to suspend assumptions about their experiences and set aside distractions, so as to be sensitive to the raw materials of experience itself (see Colombetti, 2014, chpt. 6; Thompson, 2015). Likewise, subjects should be able to mobilise a *passive* form of attention that avoids the disturbance of experience prior to its description (Thompson, Lutz and Cosmelli, 2005, pp. 69-75).

It’s been variously suggested that subjects meeting such criteria are those trained in the phenomenological reduction (Gallagher and Zahavi, 2008; Thompson, 2007), in rigorous psychological methods (Hurlburt and Akhter, 2006; Schwitzgebel, 2004) or, as concerns me here, in *Buddhist meditation*. A growing number of researchers hold up Buddhist meditation as an especially valuable form of training in the capacities underlying reliable introspection (Colombetti, 2014; Depraz, Varela and Vermersch, 2003; Thompson, 2015; Wallace, 1999), given that the practice develops abilities to manipulate the attentional systems and mobilise a relatively “bare attention” that eliminates all that is strictly unnecessary for pure observation (Thompson, Lutz and Cosmelli, 2005, pp. 69-75).

It is this “meditative turn” in cognitive science that I evaluate in papers 2 and 3. The first of these papers, “Meditation and Introspection: Insight through Transformation” aims to

support and refine the turn to meditation by responding to the most influential objection against it. This objection relates that meditative methods actually involve the *transformation* of experience, rather than its straightforward illumination, meaning that they will not produce generalisable data that reflect the mind of untrained and inattentive subjects (see Dreyfus, 1993; Colombetti, 2014, chpt. 6). I suggest that such an objection is founded on the *distortion assumption*: the implicit belief that transformation of the mind necessarily produces a false or misleading version of the mind's natural or pre-transformed state and should therefore be avoided in any investigative methods.

The distortion assumption is false. And it runs directly counter to Buddhist contemplative theory, which relates that some insights actually come not merely in spite, but, *in virtue* of particular kinds of experiential transformation. I motivate this position in paper 2, arguing that there are *beneficial* kinds of transformation available through meditation that can be exploited in introspective methods. Doing this, I seek to sketch a more nuanced understanding of the way that transformation and insight can sit together in the study of the mind, breaking the spell binding transformation to distortion, and promoting a more careful understanding of the way that meditation can help us to investigate experience.

To achieve the above, I bring theoretical material from the attention sciences into dialogue with pedagogical and phenomenological literature on meditation internal to the Buddhist tradition, both of which have been relatively neglected in considerations of this issue. I show that theoretical models in cognitive psychology implicate certain phenomenological changes possible through meditation that are actually of *epistemic benefit* rather than of detriment. And I then use the attention sciences and the Buddhist pedagogical-phenomenological literature together to show how these can be exploited in a way that minimises, or accounts for, more problematic “distortions” that might also be introduced

through meditative training. This strategy thus employs both third and first-person considerations to delineate the kinds of experiential transformation possible through meditation, and to reveal some that can do the kind of work necessary to support a broad range of subjective reports that can further guide the investigation of human experience. In this way, I mirror the ‘mutual constraints’ method favoured in neurophenomenology, using both first and third person data in reciprocal fashion, in the service of an increasingly detailed understanding of the transformational features of meditation practice, wherein both introspective and empirical data can fall into an appropriately detailed equilibrium, illuminating the practice in more depth.

By providing a clearer account of the nature and function of transformation in meditation, I am then able to make some specific recommendations about the methodological employment of meditative training – recommendations which are more sensitive to the traditional understanding and employment of such practices and which are more resistant to objections.

Paper 3 – “Introspective Training: A Broader Path?”

The final paper in this collection, “Introspective Training: A Broader Path?” is the most forward-looking of the group, critiquing the scope of benefit realistically available through employing meditation in cognitive science. In this paper, I target Francisco Varela’s (1996) suggestion that, not merely experimental subjects, but *researchers themselves* should become proficient in the kinds of gestures necessary to sensitize them to experience (p. 346). Such proposals have been taken up and elaborated by more recent thinkers (e.g. Colombetti, 2014, pp. 149-155; Desbordes and Negi, 2013; Kordeš and Markič, 2016) who suggest that

meditative practices can help researchers productively guide their own research efforts, as well as help researchers to understand and direct the introspective reports of experimental subjects.

Once more, my critique of these proposals centres upon understanding meditation as part of a broader transformational and soteriological project. For, when we fully appreciate this broader goal, we can see some major obstacles that threaten to drastically limit the advantages advertised by enthusiasts for meditation in cognitive science. This problem is well highlighted by looking to recent scholarly debates over psychotherapeutic applications of “mindfulness” meditation practice. Here, it’s well-appreciated that a primary therapeutic aim of the practice is to eliminate certain afflictive contents from the practitioner’s mental landscape – primarily the experiential manifestations of *craving*, which are considered chief causes of human suffering in early Buddhism. Looking to such debates, we see a growing recognition that this process of eliminating craving is hampered when meditation practice is plucked from its traditional environment of beliefs, practices and rituals (Dreyfus, 2011; Kirmayer, 2015; Sharf, 2015).

Importantly, as I show in paper 2, one of the primary *epistemic* benefits of meditation comes through this very transformation referenced above, for experiences of (even very subtle) craving are traditionally considered one of the central components of mental distraction, destabilising attention, and are therefore central targets for removal in meditation practice in order to reach deeper truths. An implication of this is that the introspective and the transformational-therapeutic components of meditation are not two separate enterprises but overlapping endeavours. As such, the broader contextual elements of Buddhist practice, traditionally held important to secure meditation’s therapeutic value, are also important in securing its introspective value. Accordingly, I argue that any realistic assessment of the practice’s value within cognitive science needs to be attentive to the broader supports beginning

to feature in psychotherapeutic discussions, supports including: intellectual analysis; practices for cultivating “wholesome” mental qualities; and ethical practices.

Again then, we see a major theme of these papers manifest – the transformative and the epistemic elements of meditation practice cannot be easily separated. And this presents something of a methodological problem, for the kinds of transformation required to support introspective endeavours are crucially supported by a web of practices seemingly alien to the scientific mindset and overtly very difficult to incorporate into any research paradigm. The remainder of this paper outlines the different ways one might respond to this dilemma, before weighing these options and suggesting plausible routes forward. All in all, the above problem is not a devastating one. Certainly, awareness of the importance of context helps to temper over-drawn expectations about the benefits meditation can afford science, but it also reveals important aspects of the path to introspective proficiency – underappreciated in many other traditions of introspective inquiry. Knowledge of these factors is critical in devising future methodologies that can maximise the impact that meditation can make within the field.

Conclusions and Future Directions:

From this research, three major conclusions emerge, based on which I also make some concrete recommendations about the appropriate methodologies to employ in cognitive science going forwards.

Take Introspection Seriously

Firstly, these collected works further impress the growing calls to treat introspection seriously in cognitive science. Satisfying accounts of the mind require us to “do justice” to what experience is like from the inside and to be open to the idea that introspection might actually

be used as a guide to the identification of important biological processes underlying conscious awareness. This means that more should be done to familiarise researchers, in science and philosophy, with the broad and definitive phenomenological dimensions of the different types of experience that they are studying. As well as concerning themselves with topics like the epistemic status of imaginations, delusions and emotions, or highly abstract models of consciousness, it would be valuable for researchers to have a broad understanding of at least the higher-level phenomenological features characterising these different phenomena. What is it about these experiences that distinguishes them *first-personally* from each other? What is it that characterises an experience of kind X, and thereby forms a central constraining factor in any account of kind X? To help here, researchers might, on the one hand, be encouraged to develop a close personal acquaintance with experience. On the other hand, they might become more familiar with the phenomenological analyses of those appropriately trained in introspection. Whether the broad phenomenological insights one uses to guide one's research come from personal expertise or from trained and rigorous traditions of introspective inquiry, keeping in touch with the actual dynamics of experience from the first-person perspective is critical if we wish to provide explanations of the mental that really hit the mark. Good research often demands more than a pre-occupation with theoretical models – models that can quickly come apart from experience and lead one astray – it demands a thorough understanding of experience's first-personal dimensions.

Develop “Transformational” Methods of Inquiry

If first-person methods are important for taking cognitive science forwards, then concerted effort should also be put into their improvement and into constructing methodologies that best utilise those mental capacities that are most important. In this light, the second overarching

message of these collected papers emerges. For, having reviewed the way that meditation practices yield their insights, we see that one powerful way to understand experience is to engage in its *transformation*.

Much historical research on consciousness has been informed by untrained or relatively inattentive introspection, for fear of distorting targets during their investigation. We might call these “preservational” methods of inquiry. Having shown, however, that transformation itself can be of epistemic benefit, I hope to have demonstrated that there is no need for consciousness research to so restrict itself. Rather, there should be an openness to more overtly “transformational” methods of first-person inquiry, and in developing the methodologies that can exploit such things. In paper 2, I have given some broad indications about how these might look, guided in large part by a turn to the pedagogical literature on meditation. Yet, much more effort is needed to develop these into something that can be operationalised. And this can be done with the help of further examination of meditative texts, some directions for which I have indicated. Too often, instructional texts have been neglected in the turn to meditation, and correcting this promises to help us use contemplative practice to greater effect in the future.

Transformation and Insight are Co-supportive

One important condition for developing effective methodologies for the use of meditation is to recognise that the transformational and the epistemological aspects of the practice can actually be co-supportive. This is the final take-home message of these papers. Historical biases against the transformation of experience during introspective inquiry have obscured the way that it, and indeed underlying attitudes *aiming at transformation*, can actually work as a partner to our epistemological concerns. In Buddhist thought, insight is a condition for transformation – one must understand the nature of one’s problems in order to resolve them. However,

transformation is also a condition for insight – one must quieten and still the mind in order to reveal important things about it. And this co-supportive relationship between insight and transformation is something that can be made more of in future introspective methods.

This speaks to the broader context in which researchers pursue their work and allows for some comments on the underlying motives possible within science. If the epistemological (insight-oriented) and the soteriological (transformation-oriented) objectives of meditation practice are co-supporting, this means that many different motives can produce good research. It reveals that the desire to pursue *useful* knowledge, of the kind that can benefit self and other, can actually be a condition for some kinds of knowledge to be acquired. The broader desire to transform ourselves and others for the better is considered a condition for the removal of afflictive mental conditions in the Buddhist traditions, and the removal of these is precisely what can help to put us in better touch with our experience and so guide our explanations of it. An appreciation of this reinforces the possibility of what Thompson (2016) calls ‘a new way to relate science and what many people like to call spirituality’ (pp. 931-932), wherein we recognise that the knowledge-oriented practices of science and the transformation-oriented practices of the contemplative traditions can work together, in a mutually enriching fashion, to understand and enact the processes by which society might be improved (see pp. 931-932; Kordeš and Markič, 2016, p. 166).

[this page is left intentionally blank]

Phenomenological Constraints: A Problem for Radical Enactivism

[First published (2018) *Phenomenology and the Cognitive Sciences*, 17 (2), pp. 375-399]

ABSTRACT: This paper does two things. Firstly, it clarifies the way that phenomenological data is meant to constrain cognitive science according to enactivist thinkers. Secondly, it points to inconsistencies in the ‘Radical Enactivist’ handling of this issue, so as to explicate the commitments that enactivists need to make in order to tackle the explanatory gap. I begin by sketching the basic features of enactivism in sections 1-2, focusing upon enactive accounts of perception. I suggest that enactivist ideas here rely heavily upon the endorsement of a particular explanatory constraint that I call the *structural resemblance constraint* (SRC), according to which the structure of our phenomenology ought to be mirrored in our cognitive science. Sections 3-5 delineate the nature of, and commitment to, SRC amongst enactivists, showing SRC’s warrant and implications. The paper then turns to Hutto and Myin’s (2013) handling of SRC in sections 6-7, highlighting irregularities within their programme for Radical Enactivism on this issue. Despite seeming to favour SRC, I argue that Radical Enactivism’s purported compatibility with the narrow (brain-bound) supervenience of perceptual experience is in fact inconsistent with SRC, given Hutto and Myin’s phenomenological commitments. I argue that enactivists more broadly ought to resist such a concessionary position if they wish to tackle the explanatory gap, for it is primarily the abidance to SRC that ensures progress is made here. Section 8 then concludes the paper with a series of open questions to enactivists, inviting further justification of the manner in which they apply SRC.

1 Introductory Comments

The last two decades have seen steadily growing interest in the enactive approach within cognitive science and philosophy (e.g., Colombetti, 2014; Noë, 2004; Stewart, 2014; Thompson, 2007). Originally articulated as an approach to understanding and studying the mind, early enactivism stressed that cognition is best conceived as a form of *embodied action* (Varela, Thompson and Rosch 1991/2017, p. 172) and made two central claims: firstly,

cognition is ontologically continuous with the basic life-regulation processes of living organisms (Thompson, 2007); secondly, cognition's nature is strongly dependent upon the embodiment of the organism in question, and thus the kinds of sensorimotor interaction with the environment that it can maintain (Varela, Thompson and Rosch, 1991/2017, p. 173). As recently noted (De Jesus, 2016), contemporary enactivist theorising has centred largely upon how best to systematise and develop the approach, leading to the crystallisation of a number of competing versions of enactivism that afford varying degrees of significance to the claims set out in *The Embodied Mind* (Varela, Thompson and Rosch, 1991/2017).¹ This diversity has understandably been accompanied by controversy over the authenticity of these different approaches to early enactivist ideas (see Di Paolo, Rohde and De Jaegher, 2014). In this paper, however, I am not much concerned with these theoretical details. I am interested rather in the *means* by which the basic enactivist picture was originally derived and the broad repercussions that this has for the future development of enactivist thought. Specifically, I shall be concerned with the suggestion that enactivism is a thesis arrived at principally through *phenomenological reflection* (Thompson and Cosmelli, 2011, p. 165).

The founders of enactivism note that their ideas were chiefly derived from the phenomenological considerations and analyses of Merleau-Ponty (1945/1962), which stress the embodied and interactive qualities of experience, and they mark *The Embodied Mind* (Varela, Thompson and Rosch, 1991/2017) as an extension of his research project (p. lxi). This text thus sets out a bare-bones model for how the cognitive sciences can proceed in a way that is informed by Merleau-Ponty's phenomenological insights, namely, one which focuses

¹ The main contenders here are *autopoietic enactivism* (e.g., Colombetti, 2014; Di Paolo, Rohde and De Jaegher, 2014; Thompson, 2007), *sensorimotor enactivism* (e.g., O'Regan and Noë, 2001; O'Regan, 2011; Degenaar and O'Regan, 2015; Hurley, 1998) and *radical enactivism* (Hutto and Myin, 2013, 2017). These terms are due to Hutto and Myin (2013). See Hutto and Myin (2013, pp. 23-36) for a good outline of the principal differences between these contenders.

cognitive science's empirical lens upon the appropriate horizons: the embodied interactions maintained between organisms and their environments. Applying this line of thinking to *perceptual* cognition, Thompson and Cosmelli note that they 'use phenomenological considerations about perceptual experience to constrain how [they] think about the subpersonal mechanisms of perception' (2011, p. 165) and they recommend that cognitive science more broadly take this approach, interrogating subpersonal mechanisms extending across brain, body and environment.²

This paper will be concerned firstly with two important questions that arise in response to this strategy, especially amongst those unfamiliar with the continental literature from which enactivism springs: (1) Why ought we to think that phenomenological considerations should be able to influence our cognitive science? (2) In precisely what manner should this occur? An initial aim of this paper is to answer these in detail, for I believe that the motivation for Thompson and Cosmelli's stated explanatory strategy remains unfortunately opaque in the literature, particularly to thinkers working exclusively in the analytic tradition, who have inherited a long-standing suspicion of the utility of first-person considerations when it comes to theorising about the mind.³ Without clarity here, enactivists have been left open to sustained critique over the powerful influence that it affords to phenomenological data (e.g. Rupert, 2009, 2015) and a persistent misunderstanding of basic enactivist motivations.

I will attempt to rectify this situation in the first half of this paper, unpacking the reasoning behind Thompson and Cosmelli's approach at length. I will begin by sketching some minimal details of enactivism in section 2, before moving on in section 3 to introduce a

² Those researchers who have taken up this model now make up what is called *enactive cognitive science*. See Di Paolo, Rohde and De Jaeger (2014) for more on this.

³ For a good recent example of such suspicion, see Schwitzgebel (2008). For a diagnosis of how this suspicion came to be commonplace, see Hatfield (2005).

particular explanatory constraint that the majority of enactivists have picked up, which I shall argue can help us to understand why enactivists make the claims that they do. I call this the *structural resemblance constraint* (SRC). I shall outline the justification for SRC in section 4, where I will mark SRC as an attempt to motivate identity claims between the mental and the physical, and as itself supported by standalone considerations from the philosophy of science which are familiar to analytic thinkers. In section 5, I shall then crystallise the manner in which enactivists have made use of SRC in order to arrive at their views on perception.

The second aim of this paper will then be critical, for the reliance upon SRC within the enactivist camp is not without its problems. In sections 6-7, I will reveal this difficulty in the “Radical Enactivism” of Hutto and Myin (2013). Despite implicitly supporting SRC as a means of motivating identity claims, I believe that their programme for Radical Enactivism contains claims which threaten to undermine SRC completely. The problem here is Radical Enactivism’s stated consistency with the narrow (brain-bound) supervenience of perceptual experience. I shall argue that, given Hutto and Myin’s phenomenological commitments, SRC is inconsistent with narrow supervenience, and that one of the two must be rejected.

Focusing on the case of Radical Enactivism this way will allow me to arrive at a broader and more significant point. I will outline that the endorsement of SRC is critical to the claim that enactivism can tackle the explanatory gap. As such, I will argue that enactivists must avoid being concessionary on the issue of supervenience, and *reject* the possibility of (perception’s) narrow supervenience, if they wish to preserve enactivism’s progress against the explanatory gap and promote its further bridging. The remainder of the paper (section 8) will then offer some questions for enactivists concerning their application of SRC.

2 Basic Features of Enactivism

To narrow the subject matter, I focus here upon enactivism's construal of a single subset of (human) cognition: perception.⁴ Speaking very broadly, enactivists claim that perception is a way of acting in the world, wherein a meaningful perceptual world is constituted at least partly by the specific activities the organism performs (Di Paolo, Rohde and De Jaeger, 2014, pp. 39-40; Noë, 2004). Enactivists suggest that perception consists not merely in *receiving* the world, but also in *choosing* which parts of the world one will be affected by, which aspects of the world are brought into an organism's sensorimotor loop (Stewart, 2014, p. 3), shown in *Fig. i*.




Fig. i. A sensorimotor loop (arrows show direction of influence)

To explain, the embodied activity of the organism (O) is said to determine the aspects of the environment (E) that will impinge upon it, and thus its next sensory inputs, while those sensory inputs are used to guide the actions that are needed to sustain things within this loop, themselves required to maintain the organism's viability (Stewart, 2014, pp. 1-4; Thompson, 2007). A

⁴ My treatment of perception takes it to be an essentially *conscious* process. Though the term is also used to refer to unconscious aspects of sense-making, it here concerns only those mental states/processes that there is "something it is like" for the subject to undergo. In the remainder of the paper, I will thus equivocate between perception and perceptual experience, as is common amongst enactivists. I will also use the term 'process' rather than 'state', given the enactivist understanding of perception as something which takes time to happen.

central claim of enactivism is therefore that *both* of these sensory and motor processes go into perception.

For present purposes, the details of this process are insignificant. We need only observe that the circular phenomenon outlined in *Fig. i.* is christened as the ‘*enaction*’ or ‘*bringing forth*’ of a world (Varela, Thompson and Rosch, 1991/2017, pp. 205-206), that is, the *allowing of something to present itself* or to show up for the organism. Given this account of perception, it is crucially important for the enactivist that one be in constant contact with the environment *out of which* a world is enacted. Accordingly, it has become common more recently to regard the enactive understanding of perception as something not merely active and embodied, but also *interactive*. Thompson and Cosmelli have thus qualified in a more recent article that ‘to perceive is to be in an interactive relationship with the world’ (2011, p. 165; see also Noë, 2004; Pepper, 2014; Ward, 2012).

Now, this is an understanding that, as I have said, is purported to arise from phenomenological considerations about the nature of perceptual experience. I suggest that enactivist motivations to assert this understanding depend not merely upon a certain Merleau-Pontian phenomenological picture, but also upon the concomitant endorsement of a particular explanatory constraint that I call the *structural resemblance constraint* (SRC).

3 The Structural Resemblance Constraint

SRC is a constraint upon explanations of *conscious/phenomenal* mental processes (those that there is *something it is like* to undergo) within cognitive science. Before outlining SRC, it will be useful to outline two broad distinctions that can help identify the type of explanation that SRC targets.

Firstly, we have a distinction between *horizontal explanations* and *vertical explanations* (Cummins, 1983). *Horizontal explanations* seek to account for something's occurrence by proffering a sequence of distinct *preceding* events (Drayson, 2012; Bermúdez, 2006). For example, we might seek to explain the occurrence of a tornado by reference to a series of pressure fronts and air moisture levels converging over time. This form of explanation thus thematises *why or how something was brought about*. Contrastingly, we have *vertical explanations* which offer some set of constituent parts and relations that collaboratively make up or constitute some phenomenon, thereby accounting for its features and causal powers (Drayson, 2012, p. 2). For example, we might explain the properties and causal powers of a tornado by referencing the real-time internal rotation dynamics of the air it contains. This form of explanation thus thematises *what something is made of*.

Both of these forms of explanation are found within cognitive science. For example, a clinical psychologist might explain some emotion by citing a person's previous attitudes and negative experiences - a horizontal explanation. Contrastingly, a cognitive neuroscientist might seek to explain the same emotion by relating the set of neuronal processes and relations sufficient to constitute it - a vertical explanation.

With this distinction on the table, I can state that it is solely *vertical explanations* that concern SRC. Furthermore, it is *complete*, rather than *partial*, vertical explanations which SRC targets - those which seek to proffer the necessary *and sufficient* constituent parts and relations of some phenomenon. A second distinction helps to further clarify the kind of vertical explanations of interest here. This is a distinction between *personal-level* and *subpersonal-level* explanations.

A *personal-level* explanation will invoke that which we ordinarily ascribe to *whole persons* (e.g. beliefs, desires and experiences) in its explanatory narrative. It will speak in

personal terms. This is the mainstay of folk psychology, where we might explain some individual's behaviour, say, by invoking their preceding desires. A *subpersonal-level* explanation, meanwhile, will invoke things that we do *not* ordinarily ascribe to *whole persons*. It will speak in *subpersonal* terms taken from the larger domain of inquiry from which the explanation emerges. For example, a neurological explanation of some mental process will speak in the terms of neurological theories, such as “neuronal activations” or “long-range synchrony”, while a quantum-mechanical theory will use very different terms.

While I am concerned here solely with vertical explanations, it is perfectly possible for these to proceed at either the personal or subpersonal-level. For example, Introspectionist thinkers (e.g. Titchener, 1899) have offered accounts of mental processes as being built out of simpler experiential *elements* or *atoms*, thereby giving vertical explanations in personal-level terms.⁵ However, these kinds of explanation (both vertical *and* personal-level) are not typically considered the task of contemporary cognitive science, and I set these aside here too. Cognitive science is routinely taken to account for the personal-level in terms of the subpersonal-level, and it is this kind of *subpersonal vertical explanation* that SRC concerns. This clears the space to offer its full outline.

The Structural Resemblance Constraint (SRC): An explanation of some conscious mental process X (which is vertical and subpersonal) should reveal a strong *structural resemblance* between (a) the combined constituent parts and relations that it invokes and (b) X as it is best characterised phenomenologically.

⁵ See Kriegel (2015) for a contemporary exploration of related ideas.

This needs unpacking. Firstly, a structural resemblance holds iff two relata possess some shared set of structural features, i.e. iff a mapping of one domain onto the other preserves at least some of its structural relations. If two relata share such structural features, it will be possible to abstract away from the two to reveal a structural description (a formal model) that is equally accurate in describing both (Bayne, 2004, p. 360). The limit case of this is an *isomorphism*, where the mapping is both one-to-one and *fully* structure preserving.

Given that a set of shared structural features can be larger or smaller, structural resemblance occurs on a spectrum. For example, suppose we have three pencils, all of which possess the same overall macro-shape (all are equally long, pointed cylinders). The first two will be more strongly structurally resemblant to each other than they will be to the third, if these two are both traditional wooden pencils with a common simple internal make-up, while the third is a mechanical pencil with its own more complex internal organisation.

Secondly, as to precisely *how strong* a resemblance SRC demands, we can think of this as a question of degree; the stronger the resemblance revealed, the more satisfactory the explanation will be, having taken all other factors into account. I will return to this important issue later.

Finally, relata (b) in SRC refers only to the entity as invoked and characterised *intrinsically* by our best phenomenological accounts, rather than in terms of the causal role it plays within a larger system, or the relations it holds to other mental processes. Regarding what our best phenomenological accounts might be, and where they might come from, I will tackle these issues in subsequent sections.

Now, SRC is a weaker formulation of the structural *isomorphism* constraint picked up upon in the work of Lutz (2002), Metzinger (2000, pp. 67-69), Varela (1997) and Wheeler (2005, p. 133, pp. 225-236, 2013, p. 147), thus allowing for *degrees* of explanatory value to

emerge, rather than proposing an all-or-nothing bargain.⁶ I also believe that SRC is implicitly endorsed in the work of a majority of self-avowed “enactivist” thinkers, as I shall go some way to demonstrate. I am in favour of such endorsement, for I believe that SRC is valuable as an explanatory constraint. However, I think the reasons for this are not always transparent in the literature at present. As such, I want to offer a standalone justification for SRC here. This manner of selling SRC will show it to tap directly into more mainstream explanatory concerns in the philosophy of science and mind, thereby attempting to make enactivism more comprehensible to a broader swathe of thinkers and, simultaneously, opening up a larger store of existing philosophical material for enactivists to use in support of their views. In brief, I wish to implicate SRC as no more than a specific application of something more widely discussed: a constraint upon *reduction* in the sciences. Doing this will allow me to make two points in support of SRC: (A) SRC mirrors a specific criterion for ontological reduction in the sciences that serves to increase the *intelligibility* of making identity claims, and thus (B) SRC assuages broader extant philosophical concerns about subpersonal accounts of conscious mental processes. (A) is the central point, while (B) shows the benefit of (A) in the realm of cognitive science.⁷ Let us take (A) first.

⁶ Though these thinkers speak in terms of ‘isomorphism’ rather than ‘resemblance’, I suspect that most of them would be willing to grant the value of this weaker formulation. Also worth noting here, while it is common to invoke dynamical-systems-theory (DST) as providing an appropriate kind of abstract model (a *mathematical* model, in this case) with which to capture a structural resemblance (e.g. Lutz, 2002), SRC is itself more general and can allow for less technical models to be utilised. We can therefore think that a structural resemblance holds to the extent the *any* broadly-construed abstract structural model can be applied to both relata.

⁷ Some enactivists may immediately balk at the sight of (A), most obviously those expounding avowedly “non-reductive” theories of mind (e.g. Colombetti, 2014; Thompson, 2007). I address this issue in f13.

4 Reduction and Identity

Before beginning, an account of reduction must be offered. A loose sense of the idea as given in the *Stanford Encyclopedia of Philosophy* will be sufficient here:

The term ‘reduction’ as used in philosophy expresses the idea that if an entity x reduces to an entity y then y is in a sense *prior to x* , is *more basic than x* , is such that x *fully depends upon* it or is *constituted by it*. Saying that x reduces to y typically implies that x is *nothing more than y* or *nothing over and above y* (Van Riel and Van Gulick, 2016, ‘Scientific Reduction’, para 1, emphasis in original)

What then are some reasonable constraints upon a reduction? I will look here at some recent comments on this issue in Hutto and Myin’s *Radicalizing Enactivism* (2013). By examining the criteria invoked here, which Hutto and Myin themselves endorse, we shall see that SRC offers merely a domain-specific application of one such criterion. Consequently, by interrogating the reasons given in favour of such a criterion, we shall have reasons in favour of SRC.⁸

Hutto and Myin (2013) first note that contemporary philosophy of science no longer holds inter-theoretic reductions to require strict *a priori deduction* of the truths of one domain from the truths of another (p. 175), as is sometimes stipulated concerning the mental and its relation to the physical (e.g. Chalmers, 2010, p. 244). This is considered unnecessarily demanding and not to reflect the standards typically accepted in science, given that it does not allow for the possibility of theory correction (Bickle, 1998; Churchland, 1989; Schaffner, 1993). Instead, Hutto and Myin (2013) highlight Churchland’s ‘New Wave’ criteria for

⁸ As to the significance of Hutto and Myin’s *endorsement* of such criteria, this will be returned to in the second half of the paper.

reduction, noting that a reduction should rather preserve the ‘*image* of the higher-order theory’ in the lower-order theory on offer, and thus its explanatory and predictive resources (pp. 175-6; Churchland, 1989, pp. 49-53).⁹ The higher order here is the *reduced* theory, while the lower-order is the *reducing* theory. Now, given that the preservation of a theory’s explanatory and predictive resources might reasonably be thought to be entailed by the preservation of its *image*, I will focus purely on this latter criterion here. What then does the preservation of an *image* entail?

To make things simpler (and for consistency with the characterisation of reduction given at the beginning of this section) I address this question from the perspective of *ontological reduction*, rather than theory reduction. This is to target the reduction of individual objects or properties referenced within particular theories, rather than whole theories themselves. There is no significant difference in the reductive criteria here; one aims to preserve the *image* of these, in a lower-order explanation. Let us flesh out how this might happen then.

When reducing some higher-level object, a reductive theory will first need to invoke its own specific (i) ontology and (ii) dynamics, i.e. it will offer (i) some lower-level objects and (ii) some account of how these things are disposed to behave (that is to say, it affords them certain properties). It will then cook up an account of the relationship between those objects which can be seen to yield some *macroscopic properties* that mirror the properties of the higher-level entity being targeted for explanation. A macroscopic property here refers to some property instantiated at a scale larger than that of the lower-level objects themselves. Put

⁹ Referring to the higher-order theory as the ‘older theory’, and the lower-order as the ‘newer theory’, Churchland states: ‘A reduction consists in the deduction, within T_N [the newer theory], not of T_O [the older theory] itself, but rather of a roughly equipotent *image* of T_O , an image still expressed in vocabulary proper to T_N [...] the older theory, accordingly, is never deduced; it is just the target of relevantly adequate *mimicry*.’ (1989, p. 49).

differently, it is a property that occurs at the *system-level* rather than the component-level. (If it spans the *whole* system, it is called a *global property*). Let us then take an example of the above procedure.

When reducing human biology to cell biology, one might take a human heart (a specific higher-level object, with properties disposing it to behave a certain way). One might then, from within the lower-level theory, outline how assemblies of muscle cells are disposed to coordinate over time via electrical signals so as to create periodically fluctuating global pressure levels across the assembly, a disposition which will itself be seen to mirror the heart's disposition to expand and contract (or beat). Here one is thus invoking some lower-level ontology and dynamics, i.e. some set of cells which are disposed to behave in particular ways, and concocting an account of their relationship that can be seen to yield some *macroscopic property* that is a *mirror image* of higher-level properties of the human heart.¹⁰

We can see that the production of such an image is equivalent to the demonstration of a strong *structural resemblance* between the two relata, for a mirror image can be subject to a relatively detailed abstract level of structural description that is equally applicable to that which it mirrors.¹¹ It is for this reason that I claim SRC is a domain specific application of this criterion for reduction. If we want, we can thus substitute the vocabulary of reduction into SRC and say that (a) the combined constituent parts and relations invoked by our cognitive science, should comprise a system that instantiates *macroscopic properties* displaying an imagistic relation to the properties invoked in relatum (b) of SRC.

¹⁰ I am here narrating only how *intrinsic features/properties* of the higher-level can be mirrored. Churchland himself talks of mirroring both intrinsic features/properties *and roles/functions* within a larger system (1989, p. 52). I do not focus on these latter two here, given that replicating something's intrinsic features will also enable it to play the same causal/functional role (see Drayson, 2012, p. 2).

¹¹ Indeed, for something to count as a mirror, it must replicate a fairly *large set* of structural features. The less it preserves, the more *distorting* the mirror and the nearer it approaches a *loss of mirror status*, and thus its capacity to present an *image*.

Returning to relatum (b) here, it is important to remember that the target to be mirrored by our cognitive science is always some *conscious process*. This means firstly that the relationship between the constituent parts invoked will need to be a *dynamic* one (given its being a *process*, i.e. something which changes over time) and, secondly, that the target will be provided and characterised *as a target* through phenomenological reflection (given its being a *conscious process*). It is on account of this second and more significant point that SRC invokes *phenomenological characterisations* in its outline; phenomenological theory is the higher-level theoretical domain which constructs the targets that cognitive science is to mirror.¹²

Now that we have all this on the table, we can ask: what is the benefit of SRC? Here we can reintroduce points (A) and (B) from before:

(A) SRC mirrors a specific criterion for ontological reduction in the sciences that serves to increase the *intelligibility* of making identity claims.¹³

(B) SRC assuages broader extant philosophical concerns about subpersonal accounts of conscious mental processes.

¹² It will be noted that I have outlined/compared only the detail for a *retentive reduction* in this section, where that which is reduced *retains its conceptual validity*, as when one reduces the biological to the chemical (see Bickle, 1992). I do this on account of our explanatory target in SRC being the output of our *best* phenomenology, which I here assume to be accurate enough to resist any eventual elimination. A retentive reduction is thus set in contrast to a *replacement reduction*, which *eliminates* that which it reduces, e.g. in the case of phlogiston theory. I will say more on this important issue of phenomenological accuracy in sections 5, 6 and 8.

¹³ It is worth emphasising here that I am only suggesting that SRC mirrors *one criterion* for reduction in point (A), i.e. *that which I believe to be critically important for motivating identity claims*. One might reasonably argue that this is a necessary but *insufficient* condition for reduction, and that there are additional reasons that the enactivist will wish to construe their approach as non-reductive. See, for example, Thompson (2007) for reasons in favour of a non-reductive construal of enactivism, and Bayne (2004, pp. 258-259) for a contrasting appraisal.

Point (A) is fairly simple; the production of an *image* of the higher-level, within the conceptually richer lower-level theory, motivates us to believe that both areas of inquiry have indeed managed to latch onto the very same worldly entity (see Churchland, 1989, p. 52; Metzinger, 2000, pp. 67-69). Put differently, the structural resemblance that has been revealed gives us a strong reason to think that the two levels of description are in fact describing *one and the same thing* in approximately correct ways. Of course, identity relations demand a *complete isomorphism* between two relata. But given that we ought not expect our science to quickly arrive at some “complete and final theory” of the mental, nor ought we to expect even our best phenomenology to be completely immune from any error, a strong structural similarity still provides strong reasons to claim identity (see Churchland, 1989, p. 52; Metzinger, 2000). Similarly, it seems that the stronger the structural resemblance revealed, the more intelligible it is to make such claims. Strong resemblance relations give a good vertical explanation what Gillett calls ‘ontologically unifying power’ in this way (2016, pp. 225-7). What precisely ought one’s identity claims to feature? The set of constituent parts and relations themselves, or the macroscopic features that these display? The answer to this will depend upon one’s position on multiple realization. Opting for the former would be to equivocate between a relation of constitution and identity, and to deny any metaphysical difference between the whole and its appropriately arranged parts. Opting for the latter and identifying the mental with (physically-realized) *macroscopic* processes would be to make room for multiple realizability (see Papineau, 1998, pp. 377-378). I leave this choice to others, given that it has no bearing on the value of SRC as such.

From (A), we can arrive at the second and more pointed advantage of SRC: (B) SRC assuages broader extant philosophical concerns about subpersonal accounts of conscious mental processes. Specifically, SRC respects the intuition that our subpersonal scientific

accounts of the mental and our phenomenological accounts should in some sense *fit together*. This is a criterion that McDowell (1994) has previously offered, recently emphasised by Wheeler (2005, 2013). McDowell (1994) notes that our scientific accounts should not be ‘phenomenologically off-key’; they ought not seek to identify the mental with things that do not fit onto (or *match the pitch of*, to extend the musical metaphor) our best phenomenological accounts. What then would it mean to be phenomenologically *on-key*? I think we can see that this is precisely what one is doing when seeking to satisfy Churchland’s criteria, which I have shown SRC to be a domain specific version of. Specifically, it is by so constructing an *image* or ‘analogue’ (Churchland, 1989, p. 52) of that which is postulated at the higher (phenomenological) level, *within* the lower (scientific) level, that we match its pitch. Such an imagistic/analogic relation between the two ensures a phenomenologically *on-key* account.¹⁴ Without such a relation, our cognitive science remains phenomenologically impoverished and we are left feeling that something is missing.

Having outlined SRC then, we are now in a position to return to the questions raised at the end of section 1: (1) Why ought we to think that phenomenological considerations should be able to influence our cognitive science? And (2) in precisely what manner should this occur? We can answer (1) by saying that a scientific account can better motivate identity claims by *doing justice* to the phenomenology (and thus being constrained by it). We can answer (2) by

¹⁴ The distinction between these two roughly matches McDowell’s distinction between a *constitutive understanding* and an *enabling understanding*. A constitutive understanding, Wheeler (2013) notes, ‘concerns the identification, articulation and clarification of the conditions that determine what it is for a phenomenon to be the phenomenon that it is’ (pp. 142-143). Rather confusingly here, a McDowellian constitutive understanding is thus a case of clarifying what the target for explanation actually is, the domain of phenomenology in the current context. An *enabling understanding* concerns ‘the causal elements, along with the organization of, and the systematic causal interactions between, those elements, that together make intelligible to us how a phenomenon of the kind could be realized or generated in a world like ours’ (p. 143); this is the domain of cognitive science here. A good enabling understanding can therefore make the constituted object intelligible by offering an account of causal elements and relations that can be seen to realize (rather than generate, in this context) *analogic* macroscopic processes.

saying that cognitive science can do *justice* to the phenomenology by demonstrating a strong structural resemblance to that which phenomenological reflection sets up.

Let us now return to the issue of perception then, as we began, and apply SRC here. Doing this, we see that we are immediately faced with an important question: what *is* the appropriate phenomenological picture of perception to be mirrored? SRC makes the answer to this critical. So let us see the enactivist response. Doing so will enable us to understand precisely how the enactive account of perception is derived.

5 Enactivist Phenomenology and its Consequences

Central to enactivism is the claim that, *phenomenologically*, perceptual experience is an episode of real-time interaction with the environment (Noë, 2004; Thompson and Cosmelli, 2011; Ward, 2012). Interactivity is purportedly invoked in our most rigorous and accurate phenomenological accounts, themselves held to emerge from the methods of (amongst other things) continental phenomenology and Buddhist mindfulness practice (Colombetti, 2014; Depraz, Varela and Vermersch, 2003; Shear and Varela, 1999; Varela, Thompson and Rosch, 1991/2017). Of course, enactivists list numerous additional phenomenological properties of perception—for example, that one half of such interaction comprises a *felt-body*, that acts as a subjective and affectively charged pole of such interaction (see Colombetti, 2014; Thompson, 2007)—but it is this broad phenomenological structure of *interaction with an environment* (i.e. with something *external* to, or transcending, the subject of perception) that I wish to focus on to serve the purposes of the present argument. Let us entertain some example descriptions then, where this picture emerges.

O'Regan and Noë (2001) exemplify the enactivist stance using the 'feel of a Porsche', an experience that they say is naturally expressed, *phenomenologically*, in terms of the

particular ‘Porsche-like give and take’ between us, the car and the environment (pp. 79-80). When asked to describe such an experience, we talk of Porsche-specific responses to our actions. We talk of *the way the Porsche handles*, the fact that it accelerates quickly when we press on the accelerator for example, pulling us swiftly away with it, and more general facts about both (i) *our own activity* and (ii) *the way the Porsche responds to us*. What we have here is *synchronic interaction*, between subject and object.¹⁵

Let’s take another more concrete example: the feeling of *hardness* associated with perceiving a table by touch. This might be described by the enactivist as a sense of *resistance to one’s action*. No matter how hard one pushes against the table, the table pushes back. It does not give way. The feeling of hardness should be described in terms of the refusal of that object to be perturbed by the subject, and its exertion of force back upon the subject in response to attempts to distort it.

Enactivists would also wish to extend this kind of phenomenological account to more difficult perceptual experiences, *the apprehension of colour* for example. An enactivist might here wish to highlight the artist Kandinsky’s description of the phenomenological difference between yellow and blue: ‘if two circles are drawn painted respectively yellow and blue, brief concentration will reveal in the yellow a spreading movement out from the centre, and a noticeable approach to the spectator. The blue, on the other hand, moves in upon itself, like a snail retreating into its shell, and draws away from the spectator.’ (1914/1977, pp. 36-37)

¹⁵ It’s important to emphasise that this description concerns the feel *of a Porsche*, not the feel *of the experience*. It has been common in the consciousness literature to sometimes speak of experiences themselves having “feels”, or “feeling a certain way to a subject”, creating additional complications in explanation (see Hacker, 2002; Janzen, 2008, chpt. 2). I am generally suspicious of such reference to second-order feels, however, and take all phenomenological descriptions in this paper to be direct accounts of the experiences themselves (i.e. of their qualities).

Certainly, one might dispute such phenomenological claims—which is an important point we shall return to later—though stating them overtly here should make crystal-clear the further conclusion that enactivists draw about perception.¹⁶ Given enactivists’ widespread commitment to both (i) SRC and (ii) interactive phenomenology, an interactively structured cognitive-science account fits better than a non-interactive one. We have no more than an inference to the best explanation (Rupert, 2015, p. 157).

Of course, shared interactivity alone gives only a fairly weak structural resemblance without detail regarding the substructures contained within each pole of interaction, or the kind of interaction involved.¹⁷ Convincing accounts of perception should aim to mirror as many phenomenological features as possible in order to strengthen the resemblance. However, enactivists think that their explanations can also account for just as broad a swathe of these additional features as can competing accounts of perception (which tend to miss interactivity), thus making the resemblance that enactive scientific accounts display stronger, and their explanations more satisfying.

Now that we have these central enactivist commitments on the table—a commitment to (i) SRC and (ii) interactive phenomenology—we can better understand why enactivists make the claims they do about perception. For example, Thompson and Cosmelli remark that ‘given

¹⁶ In fact, we shall see that it is crucial for enactivists to hold fast to the principle that phenomenological descriptions *can* be open to error, for such a proposal will undergird their strategies for addressing the explanatory gap (see sections 6 and 8). I am in general agreement with such a position on fallibility. However, it should be noted that, while some stress the pervasive unreliability of introspective judgement across a broad swathe of areas (e.g. Schwitzgebel, 2008), the arguments in the remainder of this paper rely mostly on only a narrow range being held questionable. We might want to remain cautious, for example, about questioning the authority of a subject’s judgements concerning the *occurrence* of particular mental states. Yet, I think that we are certainly justified in holding fallible judgements about *that in virtue of which* such a state has a particular content/character (i.e. the structural qualities that, for instance, *make* an experience an experience of hardness). And fallibility concerning these latter types of judgement is the important thing here.

¹⁷ I here, and henceforward, use “interaction” as shorthand for temporally-extended, synchronic interaction between an organism and aspects of the environment that are external to the subject of perception.

this [interactive] conception of perceptual experience, we can't specify the mechanisms of perception only in terms of what goes on in the brain without including the body and its dynamic sensorimotor coupling with the environment' (2011, p. 165). This claim can be understood by noting that it rests on an implicit endorsement of SRC. The authors see a cognitive-scientific explanation yielding global processes labelled in terms of "dynamic sensorimotor coupling" as the most viable means of *mirroring* the interactivity of the phenomenology.¹⁸

Taking SRC as an implicit commitment also enables us to understand similar claims made by Ward (2012, p. 734), who notes that enactivists' phenomenological commitments make it 'natural to adopt a *complementary* conception at the subpersonal level [...] it is natural for them to resist a subpersonal conception restricted to the neural activity of the organism engaged in such interaction [for example].' I suggest that the "natural" pull of such manoeuvres stems from a tacit commitment to SRC, which I have shown to be supported by standalone considerations from the philosophy of science.¹⁹

When entertaining Ward's remark, it is also important to appreciate that combining SRC with an enactivist phenomenology allows enactivism to remain fully compatible with the continued reference to *neural activity* or *neural representations* in cognitive scientific accounts of perception. The point is merely that an enactivist will not seek to *identify* perception with such things. Any cognitive scientific account will always invoke manifold component parts *involved* in the realization of some mental process, which needn't do justice to the

¹⁸ A similar example of such an approach to explanation, which talks in terms reminiscent of dynamic sensorimotor coupling, is Bruineberg and Rietveld's 'Radical Embodied Cognitive Science' (2014). This seeks to utilise existing cognitive scientific concepts, such as "affordances", "sensory feedback", "prediction error" and "free-energy" etc., in order to offer explanations yielding macroscopic processes labelled as "self-organizing brain-body-environment systems", mirroring the phenomenology at issue.

¹⁹ I shall return to Ward's own reasons for this conclusion in section 7, when they will become important for appraising enactivism's capacity to address the explanatory gap.

phenomenology, and which can be studied as important standalone phenomena (Ward, 2012, p. 733). It is perfectly possible therefore for neural representations, for example, to play a role (i.e. at the component level) in the realization of macroscopic interactivity. This is an important nuance recently stressed by Wheeler (2013).²⁰

Now that the nature of SRC has been delineated then, along with its warrant and implications, I wish to approach more critical points. I am concerned specifically with the ‘Radical Enactivism’ of Hutto and Myin (2013) which, despite seeming to implicitly endorse SRC, makes other claims that I think withdraw from its full implications and threaten to undermine the enactivist project. The problematic claim here is the supposed compatibility of Radical Enactivism with an entirely *narrow* (i.e. brain-bound) supervenience base for perceptual experience. To draw out this problem in the next section, I will first sketch the nature of Hutto and Myin’s endorsement of SRC and will then outline the problematic claims. This will enable me to bring out the broader point that enactivists must resist such a manoeuvre, if they wish to maintain that enactivism can address the explanatory gap.

6 Radical Enactivism and SRC

Hutto and Myin (2013) have recently set out ‘Radical Enactivism’. This is a form of enactivism which *begins with the science*, rather than the phenomenology, and applies considerations about the individuation of cognitive systems in order to recommend an interactively structured cognitive science of the mind. For Hutto and Myin, philosophical critique of our cognitive science can demonstrate that (i) cognitive systems *do not always terminate at the skin*, and (ii) cognitive science need not always utilise the notion of internal ‘content’ (i.e. representations

²⁰ Though see Gallagher and Zahavi (2008, pp. 93-94) for more sceptical thoughts on the plausibility of the notion of “representation” at this componential level.

that are subject to accuracy conditions) in order to satisfactorily explain the mind.²¹ Radical Enactivism therefore falls broadly into line with the central enactivist emphasis upon the need to understand and conceive of the mind in terms of world-involving interaction (p. 4). In this section, I wish to look at some of the claims that Hutto and Myin make about the philosophical implications of such ideas. I focus specifically upon the conclusion of their radical enactivist manifesto, for it is here that we can see their implicit endorsement of SRC.

At the end of their programme for Radical Enactivism, Hutto and Myin come to address the question of the ‘explanatory gap’ (Levine, 1983). The explanatory gap is our purported inability to provide or reveal an *intelligible link*, or a *conceptual overlap*, between the physical and the mental, with our explanations (Davies, 2008). Hutto and Myin note that any account which fails to do this will leave only *unintuitive identity claims* (2013, chpt. 8). The authors also note that they too are out to endorse a strict identity thesis and suggest that ‘belief in such identities can and should be motivated’ (p. 157), thus expressing the belief that the explanatory gap can be closed. To do this, they suggest that most philosophers must go through a process of ‘reconceiving’ (or, more appropriately, *fixing*) the phenomenological accounts that feature in their identity claims (p. 176). They state:

Like other enactivists, we reject the standard ways of characterising the ‘phenomeno’ side of phenomeno-physical identity claims. The difficulty with other existing conceptions of phenomenal properties is that their advocates are wedded to confused pictures of what is to be identified, when they imagine the relata to be qualia and brain states (p. 176)

‘Qualia’ is a notoriously complex term, denoting a particular conception of phenomenal properties that usually ascribes to them at least some of the following set of (second-order)

²¹ It is the emphatic rejection of the idea the cognition necessarily involves content that is said to afford this version of enactivism its purported radicalism (Hutto and Myin, 2013, p. 8).

properties: *intrinsically qualitative, private, ineffable, incorrigible, atomic* etc. (p. 156; see Dennett, 1993). It is unnecessary to go through the entire list here though. What the authors seem to be particularly concerned with in this context is the suggestion that phenomenal properties are both *ineffable* and *non-dynamic*. ‘As long as this picture [of the phenomenal] remains in play,’ they note, ‘there can be no progress in understanding how phenomenality intelligibly relates to, or might be instantiated in, nature’ (p. 157). They suggest instead that we ought to invoke phenomenological accounts that are ‘in tune with’ (ibid.) the way that we naturally speak about our experiences. They remark, ‘[n]aturally occurring ‘what it feels like’ illocutions take activities as their natural objects. When we describe phenomenal experience, we cannot help but mention environment-involving interactions’ (p. 177). As such, we see the same phenomenological picture recommended as was outlined in section 5. What’s more, it is important to stress that we are asked to reconceive our phenomenology in such an interactive manner *because these accounts are supposedly more accurate*, not simply because they are natural, nor because they help us argumentatively. But what argumentative function does this “revised” (or fixed) phenomenology serve? How does it help us work against the explanatory gap?

Hutto and Myin continue, in what is I think is the crucial comment, stating that ‘the plausibility of the proposed identities [between the physical and the phenomenal] looks entirely different, and far less contrived, if it is assumed that the phenomenal character of experiences must, ultimately, be understood by appealing to interactions between experiencers and aspects of their environment’ (2013, pp. 176-177). This seems to me to be nothing other than a commitment to SRC; we can lessen contrivances, and therefore motivate identity claims, by making our relata appropriately resemblant in structure. Having recommended an interactively shaped cognitive scientific account of perceptual experience, they suggest an interactively

shaped phenomenological one with which to pair it. Hutto and Myin remark that they ‘foreground the ways in which environment-involving activities are required for understanding and conceiving of phenomenality [and] abandon attempts to explain phenomeno-physical identities in deductive terms for attempts to motivate belief in such identities by reminding us of our common ways of thinking and talking about phenomenal experience.’ (2013, p. 177) Here, we therefore see an attempt to provide conceptual overlap between the mental and physical via a revisionary stance on phenomenal properties.

We can also see here that Hutto and Myin are therefore abiding by SRC from the *reverse direction* to Thompson and Cosmelli (2011). Having spent their outline of Radical Enactivism arguing for the need to approach the mind *subpersonally* (in cognitive science) in terms of world-involving interaction—using metaphysical arguments about the individuation of systems and what they call the ‘Hard Problem of Content’ (p. xv)—they *then* invoke interactive phenomenology *post-hoc* as a means to generate structural resemblance and validate such a subpersonal framing of perception.²² Yet in so doing, they seem to be promoting SRC in essentially the same manner. We can construe this kind of argumentative manoeuvre as a direct counter to the concerns of Rupert (2015), for example. Evaluating the strength of enactivist theories, Rupert draws an unfavourable assessment of their supposed progress against the explanatory gap through invoking subpersonal interactions, asking ‘[h]ow could that stuff—interaction with the environment, for example—*be* phenomenal experience? It doesn’t *seem* at all like consciousness’ (2015, p. 161).²³ Hutto and Myin would be apt to respond here by saying

²² Clearly, Hutto and Myin are invoking a different kind of phenomenological source to that of Thompson, for example, who invokes the results of more disciplined kinds of investigation (e.g. 2007, 2015; Thompson and Cosmelli 2011; Varela, Thompson and Rosch, 1991/2017). See section 8 for more on this issue.

²³ Despite being unclear in this quote, the enactivist need not construe all forms of experience in terms of interaction. I am concerned here only with perception, which is argued to have such character, but other experiences (e.g. pain, or emotion) need not.

‘yes, it doesn’t *seem* like experience, because you are working with a confused phenomenology’. Rupert is unpersuaded because he seems to take the phenomenal properties of perception to be non-dynamic (i.e. non-interactive) while, for the enactivist, dynamism is *characteristic* - it’s dynamism all the way down. Rupert thus perceives an explanatory gap to remain because of his conception of the phenomenal. Radical Enactivism seeks to work against the explanatory gap by offering a *re-conception* of the phenomenal.²⁴ So far, so good.²⁵

Yet all is not well with Radical Enactivism. My concern here is Hutto and Myin’s insistence that Radical Enactivism is compatible with the *narrow supervenience* of perceptual experience. Examining this now can help bring out a larger point about the commitments enactivists must make in order to retain the full philosophical force of their ideas.

7 Radical Enactivism: Correcting Inconsistencies

When assessing the explanatory benefits of enactivism’s proposal to “go wide” with perceptual experience— i.e. to appeal to world-involving processes of interaction in subpersonal (cognitive scientific) explanation—Hutto and Myin state that ‘[t]he explanatory value of this move is not undermined, even if it is accepted that the supervenience basis of phenomenality

²⁴ In this sense then, Radical Enactivism joins other enactivists in rejecting any *metaphysical* gap between the mental and the physical (see e.g. Noë and O’Regan, 2002). It maintains the existence only of an *epistemic* gap, occurring at the level of individuals, which can be appropriately addressed (i.e. narrowed) by providing those individuals with better phenomenology.

²⁵ Rupert has a further counter to enactivist suggestions here, arguing that structure is tangential to the issue at hand. Linking the explanatory gap to the hard problem (given that bridging the explanatory gap would solve the hard problem) he states, ‘the hard problem is that of explaining *the nature of things on which structure is imposed*, not the structure itself’ (2015, p. 161, emphasis added). As such, given that enactivist phenomenology only speaks in terms of structured dynamics, they can only leave the gap as large as ever. Bayne (2004, p. 361) also posits this same predicament. In response, we might ask: what relevant things *could* we say about the ‘nature’ of conscious processes beyond their structure and causal efficacy? It is unclear what is being demanded here. It seems equivalent to asking a scientist to explain the intrinsic nature of some physical entity beyond its compositional structure and causal powers. Perhaps there is something else important here, but I do not see it.

is wholly brain-bound’ (p. xix).²⁶ Though they do not endorse the narrow-supervenience thesis, they are keen to suggest that their Radical Enactivism is *compatible* with it. This kind of move goes directly counter to positions held by many other contemporary enactivists.

A majority of prominent enactivists support the Extended Conscious Mind Thesis (ECM): the material (constitutive) supervenience base of some kinds of experience extends into the environment (Hurley, 1998; Noë, 2004; Pepper, 2014; Thompson and Cosmelli, 2011; Ward, 2012). The typical candidate here is *perceptual experience* (Pepper, 2014, p. 99). ECM about perceptual experience is thus a commitment to a *wide* supervenience base.²⁷

Hutto and Myin wish to retreat from this. Instead, they offer a metaphysical thesis restricted to ‘basic cognition’. Basic cognition refers to foundational forms of cognition, including such things as cricket phonotaxis (pp. 42-43), but also relatively simple aspects of human cognition. It is defined as ‘mental activity that exhibits intentional-directedness, but it doesn’t necessarily imply phenomenality [i.e. phenomenal properties]’ (p. x). They claim that such basic cognition is ‘extensive’, i.e. ‘fundamentally, constitutively already world-involving’ (p. 137). Yet they stress that this does not entail a commitment on their part to *conscious* mental activity (such as perception) being itself extensive.

This theoretical retreat seems highly problematic to me. I think that there is good reason that enactivists tend towards ECM, for I suspect that the ‘explanatory value’ of enactivism is indeed put in danger by Hutto and Myin’s stance here. What I find problematic is not so much Hutto and Myin’s claim about basic cognition, it is *the claim that they can avoid ECM* (a thesis

²⁶ Though this quote is ambiguous, an interpretation of it as thematising *perceptual* phenomenality/experience, rather than only *some* (other) kinds of (non-perceptual) phenomenality/experience, is motivated by its later evaluation in the context of ‘ordinary on-line perceiving’ (Hutto and Myin 2013, p. 161).

²⁷ Note that on the current rendering a rejection of narrow supervenience does not entail a commitment to ECM. ECM concerns the environment, while narrow supervenience concerns only the brain. One could argue that perceptual experience supervenes on brain and (non-neural) body, but not the environment. However, given that Hutto and Myin’s phenomenology is interactive, it is ECM that I focus on here.

of wide supervenience), in conjunction with other claims they have made about enactivism's explanatory value. As such, we must first understand what they mean by 'explanatory value'.

This begins to be revealed in the below quote:

It remains an open question whether the strategy of “going wide” for explanatory purposes implies that the minimal supervenience base for phenomenality is like-wise (i.e. extensive). Although *a full and satisfying understanding* of phenomenality cannot be achieved without going wide, it is compatible with the supervenience base for phenomenality remaining entirely confined to the brain (2013, pp. 157-158, emphasis added)

We see here that it is *understanding* that is at stake. Going wide provides full and satisfying understanding; this is its explanatory value. Thus, enactivism's explanatory value is preserved if it can offer such a full and satisfying account. However, it seems that the authors' attempts to make Radical Enactivism compatible with narrow supervenience invoke *two different kinds of understanding*, which are themselves inconsistent. Only one of these is full and satisfying. Furthermore, this kind is *incompatible* with narrow supervenience. My suggestion is therefore that Hutto and Myin cannot claim to provide both of these forms of understanding; they must choose between them based upon the philosophical work that they want Radical Enactivism to do. Let me unpack this.

On the one hand, Hutto and Myin suggest that “going wide” provides a *contextual* understanding. They note that ‘the strategy of going wide is necessary when trying to understand the phenomenality of experience *in a wider context*’ (p. 165, emphasis added). This is to say, such a cognitive-scientific story can tell us about the specific sensorimotor patterns of interaction between organisms and environments that are *ordinarily needed in place* for such experiences to occur (Ibid.), and which perceptual experience has been historically tied to. This is to suggest that “going wide” with our cognitive science can give us the typical *contextual*

conditions of perceptual experience, whether these be real-time causal conditions, or historical conditions, without taking these to be constitutive of perception, or ‘part of its metaphysical essence’ (Ibid.). Hutto and Myin rightly note that such contextual understanding is useful on many counts, especially if we wish to understand how such experience arises in nature, or the conditions under which it can be most easily generated (pp. 163-4).

This first kind of understanding is invoked while Hutto and Myin entertain challenges to ECM. They believe that the most powerful way to problematise ECM is the ‘Argument from Shared Phenomenality’ according to which very similar, if not identical experience can occur in both the presence *and the absence* of organism-environment interaction. For example, we might consider extravagant Swampman thought experiments, where the neural activity of a perceiver is brought into being from nowhere (Davidson, 1987), or ‘direct neural manipulations of envatted brains’ (Hutto and Myin 2013, pp. 161-2). The suggestion here is that we can coherently imagine the same kind of experience being preserved in the absence of real-time environmental interaction, with appropriate metaphysical conclusions then being drawn.

Alternatively, the Argument from Shared Phenomenality can be made by referencing more mundane cases of supposed experiential overlap in the apparent absence of environmental interaction, such as ‘individuals dreaming, imagining, and suffering from “locked in syndrome”’ (ibid). We might also imagine such things as subjectively indistinguishable hallucinations. Broadly, all such arguments rest upon motivating the possibility of shared phenomenality across cases of both online perceiving, where real-time environmental interaction is in place, and (often less epistemically ‘good’) cases where that interaction is absent. Hutto and Myin are suspicious of such arguments, yet want Radical Enactivism to remain standing even if these go through (pp. 161-3). They thus use the hypothetical soundness of such arguments to motivate invoking the above kind of *contextual understanding*.

On the other hand, however, we have a second form of understanding offered by Hutto and Myin: a *constitutive understanding*. We have already seen that the authors are keen to motivate identity claims between the mental and the physical (p. 157), and suggest that this can be done by bringing our cognitive science and phenomenology into a relation of structural resemblance. However, this can only motivate identity claims if we take the science to be giving us a *constitutive account* of a mental process, rather than a contextual account of how it came to be. Returning to an earlier distinction, this means we must be offering a *vertical* rather than a *horizontal* explanation. As such, Hutto and Myin here tie the explanatory value of enactive cognitive science directly to its offering (i) a constitutive, vertical account that (ii) *abides by SRC*. In this case then, they cannot divorce the explanatory value of enactivism from metaphysical claims about relations of supervenience (i.e. *constitutive* supervenience, here)

Additionally, I have shown SRC to be a domain-specific application of Churchland's criterion for reduction. And this is a criterion which seems to satisfy Hutto and Myin. Indeed, they themselves promote this conception of reduction as a means to motivate identity claims, stating that we should bring our explanatory criteria 'in line with this more liberal thinking about how identities can be established' (p. 176), further motivating the thought that this is indeed the kind of understanding they are keen to provide, and again suggesting a supervenience entailment.

Having got these two kinds of understanding on the table, it is first immediately clear that they are incompatible. *Either* one thinks that "going wide" with our cognitive science is telling us about the causal/historical context of perceptual experience *or* one thinks that it is telling us about its constitution and thus its supervenience base; one cannot claim both from one explanation. So we have a curious oscillation here, on Hutto and Myin's part, between two

different and incompatible forms of understanding.²⁸ Which one ought they then to favour? It seems clear to me that, given Hutto and Myin's claims about the explanatory gap, they ought to favour dropping the claims about contextual understanding. Why so? Because it is only within the realm of a constitutive understanding that we can hope to bridge the explanatory gap. Tackling the gap requires accounting for *what the mental is* in terms of the physical, not accounting for how it came about. It requires demonstrating conceptual overlap between things that are *already proposed as candidates for identity*. And this is indeed what Radical Enactivism sometimes seems to aim at in its demonstration of structural resemblance between specific "interactive" phenomenological construals of perception, and cognitive-scientific renderings of these same things that manifest an analogic interactive structure at the global-level. What's more, it seems to be in this engagement with the explanatory gap that the explanatory value of enactivism is properly secured. Hutto and Myin want to offer explanations that provide "full and satisfying" understandings. If they do not address the explanatory gap, they do not do this.

A second point can be made against the idea that enactivist cognitive science should be construed as offering a merely contextual account. By definition, a contextual account will invoke interactive subpersonal processes that are *not themselves constitutive* of some mental process. A proposed contextual account of this kind would then necessarily admit the possibility of some *additional* constitutive subpersonal story about perceptual experience that *could not feature subject-object interaction* (given that interaction is meant to *frame* this, and cannot therefore figure *within* it). Yet, given Hutto and Myin's emphasis upon an interactive phenomenology, this would dissolve the central resemblance that they elsewhere invoke in support of their ideas (see Ward, 2012, pp. 734-5). So not only does their promotion of this

²⁸ See Froese (2014) for a similar assessment.

kind of contextual understanding fail to bridge the explanatory gap, it actively works against the way they elsewhere suggest closing it.

This allows me to make a broader point than about enactivism in general. If enactivists wish to maintain that they *can* address the explanatory gap, they simply cannot be concessionary in the way that Hutto and Myin suggest. If enactivism is allowed to be compatible with narrow supervenience, *in spite of insistence upon an interactively-structured phenomenology of perception*, it will need to drop any endorsement of SRC, and be content to offer more limited kinds of understanding.²⁹ This, however, would be to disarm enactivism of its major weapon against the explanatory gap. Perhaps this is an appealing strategy for some, though I cannot imagine it would satisfy Hutto and Myin, given their otherwise valuable remarks on this issue.³⁰

A final way to motivate the above choice amongst the broader enactivist community (opting for a *constitutive* construal of enactive cognitive science) is to note that it helps enactivists retain their contention that phenomenology can have some bite. We can stress here that there would be no reason at all to endorse SRC, as a means of addressing the explanatory gap, unless one believed that one's phenomenological accounts were themselves *accurate* in some sense, that is to say, unless one believed that they truthfully revealed *something* about the

²⁹ Of course, opting for this more ambitious kind of understanding means that enactivists will need to offer ways of countering the Argument from Shared Phenomenality. They might take, for example, the same line as disjunctivists on this issue, denying that veridical perception and illusion/hallucination are of a common kind. See also Thompson and Cosmelli (2011) and Hutto and Myin (2013, pp. 161-162) for further suggestions on this issue.

³⁰ There is an additional argument against SRC which may further motivate some enactivists to drop it. This is due to Wheeler (2015) who notes that a constraint along SRC lines contradicts the occasional assertions of sensorimotor enactivists that real-time interaction is not strictly *necessary* for perceptual experience. Instead, the thought goes that one only needs a practical mastery of the laws of sensorimotor contingencies that *would* obtain were one to engage in certain motor behaviours (i.e. the possession of certain practical *sensorimotor knowledge*) (pp. 169-173; see O'Regan, 2011). Given Hutto and Myin's deeply sceptical stance on the minimal necessity of such knowledge however (2013, pp. 24-32), I think they are unlikely to find such an argument convincing. See also Ward (2012, f7) for a response to Wheeler along these same lines.

subpersonal make-up of perception. This point is made by Noë (2007) who states that, if one wishes to avoid the ‘epistemic isolation’ of phenomenology, one must be prepared to say that phenomenology makes a *theoretical commitment*, which is accordingly subject to pressure from science (p. 232). After making the phenomenological claim that perceptual experience is ‘world-involving’, he states that ‘if it turns out that it is possible for me to have an experience of the same kind as a visual experience of a pencil in the absence of a pencil, then it turns out that I am wrong about my phenomenology itself [...] [for] I take a stand on the theoretical question when I take a stand on the phenomenology’ (2007, p. 236).

For these reasons, Ward (2012) argues that once enactivists have an interactive phenomenology in play this makes it natural to resist any narrow subpersonal construal of perceptual experience (and in turn favour ECM), for ‘the picture of the subpersonal that would result from doing so threatens to make it unintelligible how their conception of things at the personal level could be correct’ (p. 735). And surely enactivists like Hutto and Myin *do* want to maintain that their proposed phenomenology is correct. If this were not so, it would sit very oddly with their emphasis upon the importance of *reconceiving* our phenomenology. Why would it matter how we conceive of the phenomenology if its descriptions could be so readily bypassed?

Noë uses the above thoughts to highlight what he thinks is the contrast between continental phenomenology as a disciplined examination of lived experience and the “introspective” tradition of phenomenological reflection, which too often conceives it simply as a free-standing exercise - a ‘descriptive preliminary to theorising’ (p. 232) that metaphysics then comes to supersede. And it is occasionally this kind of free-standing conception that Hutto and Myin seem to be in danger of falling into when they back-slide into claims about contextual understanding.

It must thus be emphasised that Noë's contention that phenomenology makes a theoretical commitment is a *condition* for endorsing SRC. Unless one holds one's phenomenology to make claims about the nature of the world, one has no reason for thinking that phenomenology should have any influence when it comes to cognitive science. As such, if the enactive project has any place for phenomenology it must, as Noë remarks, conceive phenomenology to be 'concerned with nature itself' not only with 'how things seem' (p. 234).

This brings to an end the substantive claims I wish to make in this paper. I now wish to conclude on a more cautious note. For though I think that SRC is warranted, and that the way it is wielded by enactivists suggests a promising route forwards, it must be admitted that a number of significant concerns remain to be addressed here.

8 Questions Going Forwards and Neurophenomenology

First and foremost, given that SRC affords phenomenology an ability to constrain cognitive science, one must be careful to ensure it is being constrained by the *right phenomenology*. Recapitulating an earlier worry: who is to say that an interactive phenomenological account of perception is the correct one? One might claim that phenomenological methods reveal perceptual experiences to be no more than mental paint (Block, 1990) or unequivocally representational in nature (Chalmers, 2010, p. 334). This is an area of high controversy and complex debate, where arguments require the extended justification of phenomenological methods that can work not merely at the broad level, but also at the finer level, to isolate precisely what is central to perceptual experience from what is perhaps tangential or contextual.³¹ The drawing of these lines is no easy task.

³¹ See Bayne and Spener (2010) for a more detailed sketch of this debate as manifest across the study of a wider array of our experiences.

Enactivists contend that certain individuals are well equipped to perform this function, which does not in itself seem an unreasonable suggestion. But what are the best sources for maximally accurate and rigorous phenomenology? Certainly, suspicions are likely to be raised by Hutto and Myin's invocation of naïve or "natural" introspective judgement as an appropriate source of phenomenological constraint. This is highly unlikely to persuade without a great deal of argument that naïve descriptions of experience somehow gain access to some wellspring of phenomenological truth that abstract philosophising manages to miss.³²

Perhaps the methods of continental phenomenology or Buddhist mindfulness meditation are more appropriate than here. Yet these are subject to their own criticisms. In particular, there are concerns that such methods, as forms of mental training, also *change* a subject's experience, making it problematic to generalise from the phenomenological reports of such persons (e.g. Rupert, 2015, p. 172, f9). A number of recent arguments have sought to assuage such worries, suggesting that transformation of the mind through training in no way rules out gaining insight into its pre-transformed character. Colombetti (2014, p. 156) for example, agrees with Gallagher and Zahavi (2008) that such training always 'involves [both] a gain and a loss' (p. 63), bringing out certain aspects of experience, ordinarily implicit and hidden, at the expense of distorting other aspects of that experience. However, further arguments can be offered to explicate the scope of such gains and the significance of such losses, given that these methods are playing such an important role in the theoretical work being produced.

³² This is not to say that this is not a potentially fruitful avenue to pursue. Given that such conceptions (folk and philosophical) are neither univocal nor atemporal, sociological and historical investigations of the diverse, shifting understandings of human experience (and indeed "the physical") will likely prove valuable tools for addressing the explanatory gap. This would be to venture into the realm of *experimental philosophy*.

In addition, further material can be offered by enactivists to crystallize precisely *how* we can determine the accuracy of first-person methods. Existing options here include: (i) intersubjective validation (Gallagher and Zahavi, 2008) (ii) introspection-reliant abilities (Spener, 2015) and (iii) subpersonal constraints (e.g. Varela, 1996), where third-personal data is also used as a restraint upon our phenomenology. While I do not wish to arbitrate between these, I will end with a few remarks on the final of this trio of options, for this possibility is particularly important to keep in mind for those who might worry that SRC affords phenomenology an unreasonably strong influence.

It must be emphasised here that SRC does *not* leave our phenomenology immune from constraints working in the opposite direction, namely from the scientific to the phenomenological. In fact, SRC leaves open that our *best* phenomenological characterisations might themselves need to be determined with the help of subpersonal data. I will thus end here with some brief comments on the *neurophenomenological movement*, often allied to enactivism, which has begun to delineate a way of combining SRC with this opposite direction of constraint, resulting in a method of so-called ‘mutual constraints’ (Varela, 1996).

Thompson, Lutz and Cosmelli tell us that ‘[t]he [first] working hypothesis of neurophenomenology in an experimental context’, which should be familiar given the material covered in this paper, ‘is that phenomenologically precise first-person data produced by employing first-person/second-person methods provide strong constraints on the analysis and interpretation of the physiological processes relevant to consciousness’ (2005, pp. 46-47). In short, neurophenomenology *assumes* the value of SRC and looks for subpersonal physiological dynamics showing structural resemblance to aspects of the phenomenology. Doing so, it aims to reveal previously unrecognised physiological dynamics that are critical to the mental process being thematised. In a famous example, Lutz et al. (2002) examined the experience wherein a

subject gradually perceives a three-dimensional figure emerging from an auto-stereogram, showing that subtle transitions between various stages of this experience (described from the first-person) were mirrored by transitions into distinct stages of “phase synchrony” across different regions of the brain.³³

It is particularly important to note here though that neurophenomenology was also described as a method that would employ constraints in the reverse direction, i.e. it would also use *subpersonal data* to help evaluate the phenomenological descriptions on offer and to make amendments to invoked phenomenology (Varela, 1996). Thompson, Lutz and Cosmelli (2005) note that this data can ‘provoke revisions and refinements of the phenomenological accounts’ (p. 47). For example, should a phenomenological report suggest structural features that seem strongly discordant with those found in the subpersonal dynamics, the subject might be asked to re-examine this aspect of their experience. Attending more intimately to this part of experience might then yield different phenomenological reports. However, Thompson, Lutz and Cosmelli go on to note that appropriately discordant third-person data can also give the subject a means of latching onto ‘previously inaccessible’ aspects of their experience (ibid.). It can give the subject *something to look for*. When (or perhaps, if) this something is found and examined in more detail from the first-person perspective, this might allow *richer phenomenological details* about this part of the experience to emerge, that can then be used, in

³³ Of course, this experiment demonstrates only a weak resemblance between phenomenology and *neural dynamics*, capturing a shared *temporal structure*. To strengthen structural resemblance between phenomenology and the subpersonal, enactivists believe we will need to extend our constitution base beyond the brain. This is something that Colombetti (2014) notes in her programme for affective science (the study of feeling/emotion): ‘observing brain activity provides only a partial glance into the organismic activities that enact lived experience [...] This is also why neurophenomenology, which has so far limited its third-person methods to brain activity, ought to include recordings of bodily activity—thus becoming what I shall call [...] *neuro-physio-phenomenology*’ (p. 143). And in the case of perceptual experience, if we buy an interactive phenomenology, a constitution base will need to include the organism’s environment, delineated in terms akin to the ‘Radical Embodied Neuroscience’ of Bruineberg and Rietveld (2014).

turn, to discover (i.e. look for) relevant (structurally-resemblant) finer-grained subpersonal dynamics. Thus, by combining the two directions of influence, we are supposed to use phenomenological and scientific data as ‘reciprocal constraints’ to help us improve each one, and home in upon what is central (Varela, 1996, p. 343). Engaging in this neurophenomenological back-and-forth is thus intended to catalyse the production of ever stronger structural resemblances, which Varela himself saw as part of a ‘methodological solution to the hard problem’ (ibid.).

Unfortunately, this reverse constraint from science to phenomenology is seldom evidenced in neurophenomenological experiments, and far more evidence of the efficacy of the method of mutual constraints is needed. Furthermore, even with a method of reciprocal constraints, we have a problem, for it leaves us uncertain precisely *when* we are supposed to reverse the direction of constraint. In other words, how long do we allow our phenomenology to hold sway as a constraint upon our science, instead of applying the constraint in the reverse direction?³⁴ We must remember that resemblance can be generated by modifying *either* the phenomenal or the subpersonal relata, and thus be prepared to do either one.³⁵

This relates to another problem that enactivists more broadly must deal with. Namely, enactivism’s preferred method of operationalising SRC—neurophenomenology—seems to construe third-person constraints as acting primarily as *refinements* to the phenomenology, i.e. as suggesting alterations *within* their interactive phenomenological picture, rather than allowing any major overhaul of this interactive conception, thus yielding a rather one-sided rendering of the mutual constraint method. Perhaps the neurophenomenologist is warranted in

³⁴ Thanks to Mike Wheeler (personal correspondence) for this point.

³⁵ It seems to be for this reason that Noë asserts that ‘phenomenology doesn’t *fix* the outcome of important theoretical investigations [...] phenomenology *bears on such disputes* without fixing them, just as other information about how things are may bear.’ (2007, p. 237, emphasis added)

replying to the effect that “look, we need to begin somewhere and, given that we’re interested in conscious processes, and that these only emerge *as targets* from phenomenological reflection, we should give this a certain priority in characterising the targets.” Nevertheless, such reply leaves open important questions about the precise range of judgements for which phenomenological reflection will have priority, and more must be said about how to non-arbitrarily carve off such a region.

Finally, given any such neurophenomenological commitment to the authority of phenomenology in such domains (i.e. concerning the broad shape of our experiences), it is evident that something distinct from third-personal data will be needed to ascertain the validity of such first-person descriptions. This might concern the methodological criteria for good phenomenological reflection itself, which is something that other neurophenomenology enthusiasts have picked up on recently (see e.g. Bitbol and Petitmengin, 2016; Petitmengin, 2009; Petitmengin and Bitbol, 2009), though this will need fleshing out more fully to satisfy sceptics.

These are all questions that the contemporary enactivist must address if they wish to continue making use of SRC in the manner that they have done so far. But, however these questions are answered, I believe that the enactivist promotion of SRC as a standalone constraint is a reasonable one. What’s more, I hope that SRC’s outline here, if it has any take-home message, demonstrates the indispensability of phenomenology to cognitive science and philosophy, and the importance of doing it well.

[this page is left intentionally blank]

Meditation and Introspection: Insight through Transformation

ABSTRACT: This paper sketches how transformation of the mind through Buddhist meditation practice can support introspective investigations of experience in science. Rebuffing conventional associations between transformation and distortion, it carves out a space for *epistemically-beneficial* transformations. §1 first introduces meditation’s place within Buddhist thought, outlining traditional claims that the practice cultivates attentional gestures important for interrogating the mind. It then outlines proposed uses of these practices within science, before introducing worries over their utility. Such worries propose that meditative gestures transform and thereby *distort* the mind, making resultant introspective judgements unrepresentative of untrained or inattentive experience. The remainder of the paper combats these worries using material from two distinct fields. §2 introduces literature from the cognitive psychology of attention to sketch a first-pass account of how meditative transformations might be of benefit. It argues that converging models of attention here can precisify the phenomenological changes available through meditative training, such that their epistemic merits can be better evaluated. I identify one kind of meditation practice as training a form of top-down attentional control. And using cognitive psychological models of this capacity, I argue that it can (i) *accentuate* and (ii) *isolate* particular features of experience, to our epistemic advantage. §3 outlines some more challenging, distortive dangers surrounding the introspective use of top-down attentional control, showing how it can be *misappropriated* to yield genuinely unrepresentative accounts of experience. Responding to these, §4 brings the attention literature into dialogue with the pedagogical literature on meditation practice to show *how to use* this attentional faculty appropriately in introspective investigations, addressing such dangers. This allows me to conclude in §5 with some comments on prudent approaches to introspective inquiry within science.

Introduction

The past two decades have seen a resurgent interest in introspective approaches to understanding the mind (e.g. Chalmers, 2004; Gallagher and Zahavi, 2008; Kriegel, 2015; Thompson, 2007; Shear and Varela, 1999). By “introspective” approaches (or methods), I mean those approaches employing subjective reports as principal methodological tools

(Overgaard, Gallagher and Ramsøy, 2008, pp. 100-102), prioritising the *first-person access* through which subjects come to make judgements about the mind “from the inside” (Spener, 2015, p. 300).¹ Cutting the cloth broadly, such judgements might concern occurrent/conscious aspects of the mind—those there is “something it is like” to undergo, such as emotions—or they might concern non-occurrent/unconscious aspects, including dispositions and attitudes like background beliefs. Only the former concern me here.

Though there is a long and influential tradition stressing the fallibility of introspective methods, it’s now widely recognized that this does not render them useless (e.g. Bayne, 2015; Bayne and Spener, 2010; Kriegel, 2015, p. 23).² One finds increasing agreement that subjective reports serve an essential “target-setting” role in consciousness science, being a principal source of (and a means of characterising) the targets that constrain and guide the explanatory endeavours of philosophers and scientists (Chalmers, 1999; Jack and Roepstorff, 2003, 2004; Kriegel, 2015, pp. 18-21; Thompson, 2007, chpt. 1 & 2, 2015). On a theoretical note, good descriptions afford experiences with broad, definitive properties that should be taken seriously by theoreticians, whose accounts should aim *ceteris paribus* to do justice to these properties (see McDowell, 1994; Roberts, 2018b; Ward, 2012; Wheeler, 2005, p. 133, pp. 225-236).³ On a narrower empirical note, so-called “neurophenomenologists” use subjective reports to help reveal important biological (especially neurological) underpinnings of experience. By prompting detailed accounts of the internal structure of experiences, and then searching for

¹ I use the terms “introspective” and “subjective” interchangeably to describe such methods throughout. One also finds reference to “phenomenological” or “first-person” methods in the literature. I take all such terms to be univocal here, avoiding theoretical assumptions about the mind-world relationship (c.f. Zahavi, 2007).

² For recent examples of caution over the use of first-person methods see Dennett (1993), Rupert (2015), Schwitzgebel (2008, 2011) and Spener (2011).

³ This is not the place to interrogate what “doing justice” would look like. For thoughts on this, see Roberts (2018b) and McDowell (1994). Neither should this claim be read to assert that good introspective accounts are indefeasible – only that they should be constraining factors upon our explanations.

analogous structure in third-person data (i.e. data often dismissed as mere “noise”), researchers can highlight bodily processes previously unrecognised as important to the experience in question (see e.g. Dor-Ziderman et al., 2013; Lutz et al., 2002; Reinerman-Jones et al., 2013).⁴

As receptivity to introspective data has been building, many have stressed the importance of sound methods for producing it. At the turn of the twentieth-century, Chalmers identified the development of rigorous and systematic methods for bringing experience to report as the ‘greatest challenge’ facing consciousness science (1999; see also Frith, 2002). Many have since taken up this challenge, re-energizing the enterprise of the early Introspective Psychologists.⁵ In this climate there has been an increasing turn to Buddhist attention-regulation practices, i.e., meditations (e.g. Colombetti, 2014; Depraz, Varela and Vermersch, 2003; Kordeš and Markič, 2016; Thompson, 2007, 2009, 2015; Varela, Thompson and Rosch 1991/2017), which are my concern here.

Meditation has long been employed by Buddhist contemplatives to gain intimacy with experience. Amongst contemporary researchers too, it’s now held to train several attentional gestures important for introspective methods. Meditation is said to develop capacities to put aside distractions and assumptions during introspection, along with an ability to sustain attention upon present experience without slipping into evaluative or judgemental narratives or thoughts about past and future (see Colombetti, 2014, chpt. 6; Thompson, 2009). In general, meditation is thought to support a “bare attention”, or “passive observational stance”, unobtrusive enough to avoid disturbing target experiences or colouring their description with theoretical preconceptions (Thompson, Lutz and Cosmelli, 2005, pp. 69-75).

⁴ For reasons that structural analogy is thought important here as a methodological and explanatory constraint, see Roberts (2018b), Ward (2012), Thompson and Cosmelli (2011).

⁵ See Boring (1945) and Spener (2018) for good accounts of these earlier attempts at systematising introspective methods.

Despite this enthusiasm, there remains continued disquiet over such proposals. There is a recalcitrant worry that the attentional gestures trained in meditation don't simply help to reveal the mind; they actually transform it. They are accused of yielding different kinds of experience rather than illuminating existing ones and thereby of re-shaping and potentially “distorting” experience in the attempt to characterise it (see Dreyfus, 1993; Colombetti, 2014, p. 150; Fox et al., 2012, pp. 6-8; Thompson, Lutz and Cosmelli, 2005, p. 72; McAuliffe, 2018, p. 239).⁶ In brief, it remains controversial how practices seemingly supporting the transformation of experience can yield generalizable data (see Froese, Gould and Barrett, 2011, pp. 264-265).

Many attempts to rescue the “meditative turn” from these worries have contented to highlight (promised or actual) empirical progress through similar methods as counter-evidence (e.g. Colombetti, 2014, chpt. 6; Thompson, 2015, pp. 56-57) or have proceeded on a phenomenological basis that makes introspective claims about the nature of “bare attention” itself, seeking to downplay problematic experiential changes here (e.g. Kordeš and Demšar, 2018; Petitmengin and Bitbol, 2009, pp. 372-381; Bitbol and Petitmengin, 2013; Depraz, Varela and Vermersch, 2003; Thompson, Lutz and Cosmelli, 2005, pp. 72-73). Yet, there has been relatively little attempt to combat a central assumption underlying concerns over meditative methods – the assumption that experiential transformations within introspective investigations are necessarily harmful to the process of understanding the mind and should in principle be avoided.⁷ We can call this the “distortion assumption”, in that it equates mental

⁶ Those *au fait* with the history of introspective methods will sense familiar territory here. Such concerns resemble older worries about ‘observational distortion’, long troubling the methods of the psychological and phenomenological traditions (see Petitmengin and Bitbol, 2009). I revisit this similarity in §3.

⁷ I do not mean to suggest that this assumption is unquestioned or unchallenged in the existing literature (examples of which will be evidenced later), merely that there have been no extended attempts to dismantle it.

transformation (in this circumstance) with the production of a false or misleading (i.e. “distorted”) picture of our natural mental processes.

I suggest that the distortion assumption is false. Mental transformations need not be harmful to the process of understanding the mind. And I wish to argue this conclusion in the first part of the paper by unpacking a central yet neglected aspect of Buddhist contemplative theory. In classical meditation literature, much of the practice’s epistemic value is held to arrive not in spite experiential transformation, but *in virtue* of it. The insights of “insight-meditation” are said to demand a background in attentional practices self-consciously aimed at developing capacities to transform the mind (Gethin, 1998, pp. 174-177; Williams and Tribe, 2003, pp. 81-84; Wallace, 1999, pp. 175-180). In this way some transformations are held not only unproblematic but *beneficial*. I will unpack this proposal in the paper’s first half, seeking to free discussions about meditation from the insidious effects of the distortion assumption. This way, one can then construct a more judicious response to concerns over meditation’s scientific utility, which I shall detail in the remainder of the paper.

To achieve my first aim, I shall employ some conceptual resources increasingly recognised as powerful ways to illuminate the meditative enterprise: those of the cognitive psychology of attention. It’s surprising, given meditation’s growing study as an “attention regulation practice” (see Lutz et al., 2015), that little use has been made of contemporary attention science to address concerns over the generalizability of meditative insights. Looking to cognitive psychology allows one to pin-point some well-defined attentional capacities trained and improved in meditation. And contemporary models of these capacities can yield an increasingly precise understanding of their phenomenological effects. Knowing these effects more comprehensively, one can better gauge whether training subjects to become more skilled in these capacities might be of use or of detriment. In this first section, I shall thus employ the

attention sciences to reveal two different kinds of *epistemically beneficial transformation* that meditation makes possible through attentional training.

In this way, the attention sciences can make important aspects of Buddhist theory more intelligible and undermine the distortion assumption without relying solely upon phenomenological claims. Equally though, they also help bring into focus some more appropriate and cutting worries about meditative methods, which I shall turn to in the second half of the paper. The attention literature highlights additional and more problematic kinds of transformation possible through meditation-trained attentional skill. And their possibility means one must be very careful about *how such skill is used* in introspective investigation, for there are many ways it can be misappropriated to yield genuinely unrepresentative accounts of human experience.

In the final section of the paper, I shall turn to a second neglected field to advise how these pitfalls can be best avoided: the pedagogical literature on meditation. In such literature, one finds a rich vein of instructional commentary delineating how to use attentional skills appropriately in epistemically-oriented “insight meditation” practices. This has been significantly under-appreciated in analytic treatments of this topic. Interrogating practice instructions in Buddhist texts—looking at how the contemplative quest actually proceeds—allows for the prescription of, for instance: the kinds of introspective target for which specific attentional capacities are appropriate; the manner in which these capacities should be utilised; and, when they should be transcended. Through continual dialogue with the attention literature, I shall show that the actual pragmatics of meditation instruction show how to minimise these more pressing dangers in contemporary scientific environments. In tandem, the two fields can help us approach more detailed methodological protocols for the effective employment of meditative training in introspective methods. Moreover, they will help to showcase and

crystallise a number of different ways that experiential transformations can be incorporated into introspective methods – the varying senses in which transformation and insight can happily sit together.

To begin, I must do some groundwork, delineating the basic character of meditation and the reasons it has been considered valuable to cognitive science.

1 Meditation: Proposals and Objections

1.1 Buddhist Roots

Unpacking meditation requires saying a little about the broader Buddhist context from which these practices emerge. Speaking of “Buddhism”, in this broad sense, obscures many significant differences across the various Buddhist traditions. Nonetheless, it is sufficient for current purposes, which require me to convey only some “foundations” of the Buddhist traditions – a term Gethin (1998) coins for those fundamental ideas and practices present in early Buddhism that are largely shared or assumed by its varied, later manifestations (p. 3).

It is first critical to emphasise Buddhism's essentially soteriological orientation – its fundamental concern with suffering and liberation from suffering. Suffering is believed to be rooted in craving and its cessation in the abandonment of craving, which is to be achieved by following the Eightfold Path.⁸ Meditation practices should be understood in this context – they are a collection of attention regulation techniques prescribed on the Eightfold Path, aimed at removing craving (see Lutz et al., 2008, 2015). Meditation redirects attention to particular objects, in particular ways, to achieve this.

⁸ Some traditions identify *ignorance* as the central condition for suffering and that which is to be overcome (see Lusthaus, 2003, p. 243). Largely though, these hold ignorance central *because it underpins craving*; one craves because one doesn't know enough – a relationship I interrogate shortly.

Though the exact meaning of ‘craving’ (Pali: *taṇhā*) is complex (see Peacock, 2008), it can be glossed here as desire that has become *obsessional*, such that the object of one’s desire (which could be a material object, person, experience, or usually an abstract idea) has taken on an aura of necessity.⁹ Craving can take two principle forms. One can crave to attain things or to escape things. In either case, the attitude requires no mere preference, but the felt framing of some object as imperative to one’s continued identity or existence (see Lusthaus, 2003, p. 61; Teasdale and Chaskalson, 2011, pp. 94-100). Traditionally, this is considered to manifest in such things as lust, anger, worry and doubt which are collected (with others) under the heading of the ‘Hindrances’ in early Pali discourses (Thiradhammo, 2014, pp. 17-21).

A central proposal concerning craving’s removal is wonderfully simple: one should become *familiar* with the operation of the Hindrances and the way they produce harm (AN 3:101-102 in Bodhi, 2005, pp. 192-193; Thiradhammo, 2014, pp. 17-19). Doing so, one can curtail harmful ways of being, and prioritise beneficial ones. Regrettably, such familiarity is difficult to attain. The mind of the beginner is held unpliant – so stormy and chaotic as to obscure the mechanics of craving and suffering in a tangle of activity and distraction (Thiradhammo, 2014, p. 31; Gethin, 1998, pp. 174-177). Put simply, unprepared investigation usually results in a swift transition from a receptive attitude to a *reactive* one. The student will get quickly dragged into long chains of evaluative response and distraction—thinking “I shouldn’t be feeling that”, “this means I’m not making progress” or “I’ll never be a good meditator”—thereby disrupting calm observation. Or they will simply get overwhelmed by the magnitude of mental activity. This is where meditation practices come to the fore; particular

⁹ Parenthesised italics henceforth give original Pali terminology, Pali being the language in which Buddhist ideas were first committed to the page and the language of the early Theravāda tradition.

kinds of meditation train specific attentional gestures held important for investigating experience.

Two principal kinds of meditation are outlined in *Fig. i.* below.¹⁰

Focused-Attention	Open-Monitoring
1. Hold attention upon a designated object	1. Be open and attentive to all contents arising in the stream of experience, moment by moment
2. Notice distractions that drag one away from the object	2. Notice reactivity to mental contents, or past and future narratives, that make one lose touch with experience
3. Release distraction	3. Release reactivity
4. Return attention to the object	

Fig. i. Meditation schematics.

Focused-Attention practices are simpler. The student chooses a single object on which to meditate and tries to sustain undivided attention upon it, cycling through the above four steps repeatedly.¹¹ The object here might be an external object (such as a pebble, mandala, or coloured disc) or internal object (such as the breath, a body part, feeling, or mental image). Internal objects can therefore be ostensibly physical or mental, with the term “object” designating something which one sets oneself opposed and attentive to, not something necessarily “out there” in the world (see Thompson, 2007, p. 23). Given that my interests are

¹⁰ These outlines are adapted from Lutz et al. (2008). See Gethin (1998, chpt. 7) for a broader account of both practices. The categorical terms used here are not native to Buddhist traditions; they are neologisms introduced in contemporary scholarly and scientific literature to better categorise diverse practices from across the Buddhist world (Thompson, 2015, pp. 51-52).

¹¹ Certain factors complicate this account of Focused-Attention practice. For example, traditionally it’s common to *shift* object once a certain degree of concentration is reached (see Gethin, 1998, pp. 181-184; Dunne, 2011, p. 80; Shankman, 2008, pp. 57-59). Nonetheless, this simple outline is sufficient for here.

with introspection here, I concern myself henceforth only with the training of attention to mental objects (i.e. experiences).

Focused-Attention is prescribed primarily for its ability to train a number of gestures held important to the investigation of experience – those comprising the steps of practice. By bringing the mind back repeatedly to a single aspect of experience, Focused-Attention practices are traditionally said to cultivate the ability to direct attention to a specific point and hold it there (steps 4 and 1, respectively) (see Gethin, 1998, p. 176; Davis and Thompson, 2013, p. 592). Relatedly, it's held to help practitioners notice and dissolve distractions more quickly and more easily (steps 2 and 3) (Lippelt, Hommel and Colzato, 2014). Open-Monitoring practice, meanwhile, focuses chiefly upon supercharging the latter two capacities, with distractions conceived here as *reactivity*, or anything that drags one away from calm observation, manifesting a desire for things to be otherwise. It also trains sensitivity to a broader range of mental activity in step 1, where attentional focus de-emphasised. Collectively, the attentional gestures honed in these two practices are believed important for the investigation of experience (itself traditionally occurring in “insight” or *vipassanā* practices) and contemporary cognitive science has begun to pick up on this.

1.2 Scientific Uses: The Meditative Turn in Cognitive Science

A growing number of researchers in cognitive science are seeing meditation as a means to support a general introspective proficiency that can be of scientific benefit. Many mark meditators' abilities to hold attention upon prescribed targets (experiences) as facilitating more accurate and detailed subjective reports (Thompson, 2009, p. 189; Colombetti, 2014, chpt. 6). Much is also made of related capacities to release and suspend habitual concerns, judgements and assumptions (Ibid. ; Thompson, 2015, pp. 52-53). These might include theoretical or

normative judgements about experience, or simply other ongoing mental activity (passing thoughts, sensations, feelings) irrelevant to one's experimental interests, which Wallace (1999) speaks of as general mental "excitation" (pp. 176-177).

Without such capacities, it's thought that subjects will be in danger of: expressing prior beliefs or expectations about the way their experiences unfold (Colombetti 2014, p. 157; Schooler and Schreiber, 2004, p. 33); slipping into explanatory rather than descriptive narratives (Bitbol and Petitmengin, 2013, pp. 271-273; Hurlburt, 2009; Petitmengin, 2006, p. 235); mistaking judgements for components of target experiences themselves (Colombetti, 2014, p. 157); or simply losing touch with the present stream of experience (Wallace, 1999; Thompson, 2009, pp. 188-189).

Taken together, the attentional gestures trained in meditation are held to underpin a "bare attention" or "receptive openness", held important for the description of experience. Thompson, Lutz and Cosmelli (2005) note that 'bare attention means noticing, witnessing, or being present to what is happening in one's experience, without explanation or judgement' (p. 70), and others have held this up as a holy-grail of introspective methods. Colombetti (2014) considers it as central to describing the phenomenological micro-dynamics of emotions, and important for identifying the unique biological markers of emotion types (p. 151). Such descriptions require putting aside preconceptions and normative evaluations (e.g. strategies for dealing with an emotion, or judgements about its appropriateness) such that one can focus on the lived-character of emotion itself (chpt. 6). More generally, bare-attention is held central to the phenomenological epoché, where habitual concerns and assumptions are set aside to interrogate one's openness to the world in experience (Thompson, Lutz and Cosmelli, 2005, p. 71).

Integral to the meditative turn, however, is the assumption that meditators are not merely good at describing their own experiences, but that they can facilitate better conclusions about experience *in general* (Varela, Thompson and Rosch, 1991/2017; Thompson, 2007, 2015). Our cognitive science isn't especially concerned with truths about the individual experiences of meditators, but truths about particular inter-subjectively shared *types* of experience – the defining properties of “anger”, “shame” or “pain”, say – that are distinctive of the human mind. It's proposed that we can achieve this by submitting the judgements of trained subjects (i.e. meditators) to a generalisation process (see Thompson, Lutz and Cosmelli, 2005, sec. 6; Gallagher and Zahavi, 2008, chpt. 2). First, we assess what's common to judgements about tokens of the same type within the individual. This yields *intra-subjective* truths about these types. Then we compare and corroborate such judgements across many individuals to yield *inter-subjective* truths. These latter kinds might express the first-personal character of the various emotions (precisely *what it is like* to experience joy, say, in contrast to excitement) or the defining properties of the various perceptual modalities (e.g. what it is like to see rather than imagine) – general truths known as ‘invariant’ or ‘essential structures’ of experience in the phenomenological tradition (Ibid., p. 28). In this sense then, the important claim of the meditative turn is not that meditation has individual introspective benefits (affording better judgements about one's own mental states), but that it has *scientific and philosophical benefits*, supporting more rigorous phenomenological accounts of the types of experience that our (scientific and philosophical) theories of mind are concerned with, which can then be used to guide and constrain research.

1.3 Objections

The above proposals are not without their critics. Many resist such “revelatory” views of meditation, offering more problematic assessments of the capacities that meditation trains. By far the most prominent criticism here relates that, while such capacities might sensitize meditators to *their own* experiences—that is, their attentive and stable experiences—they will also make such experiences importantly different to those of non-meditators, or experiences immediately prior to the deployment of attention (see Colombetti, 2014, p. 150, pp. 155-158; Dreyfus, 1993; Froese, Gould and Barrett, 2011, pp. 264-265; McAuliffe, 2018, p. 239; Thompson, 2015, pp. 56-57; Thompson, Lutz and Cosmelli, 2005, pp. 72-73; Shear and Varela, 1999, p. 13). From such a perspective, it would be a mistake to generalise from meditators’ judgements to broader truths about the human mind, throwing into doubt meditation’s value for the broader scientific objectives noted above.

As Colombetti (2014) notes, objections of this kind tend to slide between two formulations. One can object that the transformative character of the attentional capacities trained in meditation undermines the practice’s ability to yield insights into (i) the natural, inattentive and *untrained* experiences of non-meditators, or (ii) the mind as it was *prior* to deliberate attentional gestures being deployed, often called “pre-reflective” or “lived experience” (pp. 155-157). In this paper, I take the above formulations together, under the broader issue of whether meditation is able to illuminate *lived experience*.

This line of criticism is certainly not new. It coarsely re-capitulates a central concern apparent in much older discussions over the appropriate experimental methods for psychology, occurring during the hey-day of Introspective Psychology and its competitors. Though these were broad-ranging and nuanced (something I’ll later unpack in more detail), a central topic of concern here related how methods of investigation employing purposeful, *directed* attention

could transform and thereby distort the experiences of interest to science (see Bitbol and Petitmengin, 2013; Spener, 2018). In light of such concerns, many attempts to employ introspective methods in psychology sought to devise maximally inattentive and what we might call “preservational” (as opposed to “transformational”) methods, which could probe experience without inducing changes to its content (see Spener, 2018, pp. 156-166).

Given the continuing influence of these concerns, and enduring associations between transformation and distortion, many have begun to hold up open and undirected forms of meditation as most valuable to science (e.g. Bitbol and Petitmengin, 2013; Depraz, Varela and Vermersch, 2003, chpt. 1.2; Petitmengin and Bitbol, 2009; Thompson, 2015). For instance, Petitmengin and Bitbol (2009) identify abilities to *direct* and *sustain* attention, trained in Focused-Attention practice, as most overtly transformative and problematic in character. They focus instead upon promoting maximally passive, higher-level Open-Monitoring meditations, which they suggest train an effortless, receptive and broad-scoped form of attention that sensitizes one to experience without the problems associated with directed attention (p. 378).

In spite of this, skills in directing and sustaining attention are considered essential components of introspective proficiency in the contemplative traditions. Focused-Attention practices are usually treated as precursors to Open-Monitoring types and are framed as important foundations for more advanced insight-practices (Gethin, 1998, chpt. 7; Thompson, 2015, p. 52). Accordingly, others have started to sketch plausible benefits to introspective methods available through these more overtly transformative skills in directed attention, emphasised in Focused-Attention practices (e.g. Davis and Thompson, 2015). Given their important status in the Buddhist tradition, and their relative neglect in recent literature, Focused-Attention practices warrant keener attention. And in the remainder of this paper, I

shall sketch how the capacities they train can play a significant and beneficial role in first-person scientific methods.

As indicated earlier, concerns about the value of directed attention (and by extension meditation) rest largely upon the *distortion assumption* – the implicit belief that transformation of the mind necessarily produces a false or misleading, i.e. “distorted”, version of the mind’s natural or pre-transformed state. I am not the first to recognise this as lying at the heart of many objections to the scientific use of trained introspectors. In their introduction to neurophenomenology, Thomson, Lutz and Cosmelli (2005) counter an analogous objection by explicitly critiquing the distortion assumption. Rather than seeking to downplay the transformative effects of the investigatory methods they propose, they suggest that some kinds of transformation can actually be of epistemic benefit. The authors marshal Husserlian phenomenological claims to the effect that skilled methods of investigation can help to bring out features of pre-reflective experience more distinctly (pp. 72-73). Similarly, Colombetti (2014, chpt. 6) endorses Gallagher and Zahavi’s (2008) proposal that a skilled method of self-observation is able to ‘disclos[e], disentangl[e], explicat[e] and articulat[e] [...] components and structures which were contained implicitly in the lived experience’ (p. 63).

Phenomenological considerations certainly have a place to play in these debates (see Roberts, 2018b). However, arguments seeking to *justify* introspective methods—to establish their epistemic credentials—solely by appeal to introspection itself (trained or otherwise) are likely to arouse suspicion (see Bayne, 2015, p. 5). Furthermore, such arguments have not gone far to alter popular beliefs on this issue. Colombetti (2014) identifies naïve associations between transformation and distortion as recalcitrant sources of armchair criticism levelled against her recent programme for the use of introspective methods in emotion science (pp. 155-158).

Perhaps most importantly though, claims of the above kind are so broad as to leave it unclear precisely *how* transformation might be of benefit. What, for instance does it mean to suggest that some phenomenological transformations can “disclose”, “explicate” or “articulate” its implicit features? Many such terms do little more than state the desired conclusion, employing epistemically-loaded vocabulary to describe the transformations involved. They don’t tell a story about the *kinds* of transformation that can be seen to secure these epistemic functions.

Given the above shortcomings, I will attempt a slightly different track here. I suggest that a useful way to make progress involves turning to the attention sciences. Attention science allows us to re-conceptualise the contentious gestures developed in Focused-Attention meditation (the directing and sustaining of attention) in terms of a specific and well-defined attentional capacity. It also offers detailed subpersonal models of this capacity, accounts that should constrain how we think about the kinds of transformation it involves. I suggest that looking more carefully into this reveals specific, *epistemically-beneficial kinds*, fitting the mould of those referenced by the above phenomenological thinkers.

My approach here is common in contemporary philosophy of mind. Increasingly, researchers seek to clarify the character of mental phenomena by appeal to scientific models. Examining how broad phenomenon can actually be *implemented* at the physiological level has refined how we think about their finer phenomenological details. For instance, the discovery of a relatively narrow bottleneck in visual processing has been taken to favour less “rich” conceptions of visual phenomenology (Noë, Pessoa and Thompson, 2000). Gallagher and Francesconi (2012, p. 6), meanwhile, note that the unearthing of common neural structures active in the pre-frontal cortex during both performed and perceived action has favoured early Husserlian conceptions of social cognition as involving the “mirroring” of other’s kinaesthetic

sensations in one's own body, rather than mere intellectual inference (see also Gallagher, 2005). My attempt to pursue an analogous approach in the case of meditative transformations builds upon preliminary work by Davis and Thompson (2013, 2015), who seek to bring the attention and meditation literatures into dialogue to pursue a genuinely 'cross-cultural' cognitive science, better able to reveal the transformative capacities of the human mind, along with their epistemic merits. Davis and Thompson's own interest is with epistemic improvement in general though, rather than a narrower concern with introspection as is mine here.

Again, I stress that my initial turn to the attention literature does not reflect a belief that phenomenological considerations have no place here. On the one hand, this argumentative strategy is simply pragmatic, recognising the explanatory persuasiveness of empirical considerations in contemporary philosophy of mind. More importantly though, it is best understood in terms of a larger project of "mutual constraints", wherein we recognise that *both* first and third-person data should inform our views about the mind and can be used in reciprocally-influencing fashion, over time, to generate increasingly precise understandings of experience (see Gallagher, 1997; Gallagher and Zahavi, 2008, pp. 32-33; Varela, 1996), which is what I shall move towards in section 4. When it comes to issues of transformation and distortion, insufficient work has thus far probed the third-personal side of this bargain, so this is where I begin.

2 Meditation and the Cognitive Psychology of Attention

2.1 Attentional Skill in Focused-Attention

To distinguish the relevant attentional capacity trained in Focused-Attention meditations, it's helpful to go into further detail about the experience of undertaking the practice. As outlined

in *Fig. i.*, distinctive of Focused-Attention is the repeated return of attention to a single (mental or physical) object. This is typically done by mentally rehearsing a suitable *label* in step 4.¹² This can be the object's name, or a strongly associated term. If the object is a candle flame, the student might sub-vocalise "flame" or "seeing". If the object is the experience of the breath, they can silently remark "breathing". This act of labelling facilitates a return to the object, and with constant iteration it is said to yield a sense of becoming increasingly held or "tethered" to that object (Gethin, 1998, pp. 176-181). With enough training, it's proposed that meditators can even "fix" their mind upon an object in a "one-pointed concentration" (Ibid., p. 181), where no re-orientation is needed at all. In this case, they will remain solely in step 1 of the practice. At this point, we can introduce some considerations from the attention sciences to precisify the attentional capacity trained through the above labelling technique.

Within cognitive psychology, the employment of labels to redirect attention is referred to as an instance of *top-down (attentional) control*. Generally, the term "top-down" is used in the field when attention is internally guided by a subject's prior knowledge, plans and goals (Katsuki and Costantinidis, 2014, p. 509), with the canonical guide being conceptually-formed *intentions* to attend to a specific location or object.¹³ It's also important to note that attention can be classed as "top-down" whether one attends to the objects of the world (i.e. with perceptual attention) or to the objects of experience (i.e. with introspective attention) so long as it fulfils the above condition (see Wu, 2014, p. 255).¹⁴ Top-down attention comprises one of

¹² For a good account of the importance of label use, glossed as "applied and sustained thought" or "initial and sustained mental application" (*vitakka-vicāra*) in early Buddhist discourses, see Shankman (2008, pp. 38-40).

¹³ This is admittedly a loose characterisation of top-down control. For instance, the "guidance" of attention can be understood in a number of different ways, and to occur across a number of varying timescales. I might in a sense "guide" my attention to my alarm clock by setting it to sound at 7am. For present purposes though, we should take top-down attentional control to occur in those cases where a subject's attentional targets are determined in *direct/immediate consequence* of some independent psychological state of the subject.

¹⁴ I say more on the relationship between perceptual and introspective attention throughout §2-3.

two broad kinds of attention distinguished in cognitive psychology. The second kind is “bottom-up” attention, where attention is guided by external factors, including specific and salient properties of attended objects themselves, as when a close flying pigeon inadvertently attracts and holds one’s gaze (see Corbetta and Shulman, 2002).

Ordinarily, top-down attention is initiated voluntarily, with the recollection of specific intentions that express goals of the subject. On account of these properties, top-down attention is also sometimes called *voluntary*, *goal-driven* or *endogenous attention* and set in contrast to *involuntary*, *stimulus-driven* or *exogenous* characterisations of bottom-up attention (Pinto et al., 2013; but see Wu, 2014, pp. 29-38).¹⁵ One should also note that top-down attention captures both cases where (i) a subject purposefully *shifts* attention (e.g. from the book they’re reading to their emotional state), and (ii) a subject *sustains* attention in virtue of (or under the “control” of) some independent psychological state. For instance, deliberately holding attention upon a specific part of the visual field is also a case of top-down (controlled) attention. Here, we have a case where top-down attentional control extends beyond a single attentional shift (see Pinto et al., 2013, p. 2).¹⁶ Returning to Focused-Attention practice itself, we see that both such kinds are being mobilised – in the repeated *return* of attention to a specified object in step 4 and the *holding* of attention in step 1. For clarity’s sake, it’s helpful here to tie these two together under a broad skill; Focused-Attention practice is distinctive in training a kind of *top-down*

¹⁵ For a more detailed consideration of these two forms of attention, including the difficulties inherent in such simple distinctions, see Wu (2014, chpt. 1). Unlike Wu, I use the terms “top-down” and “controlled” as equivalent here.

¹⁶ Given that top-down attention *usually* extends beyond a single attentional re-orientation, Pinto et al. (2013) note that it’s common to equivocate between top-down attention and *sustained attention* (p. 2). This is unwise, given that not all top-down attention need be sustained, and not all sustained attention need be top-down (see Lutz et al., 2008, pp. 166-167). I introduce the possibility of a bottom-up, sustained attention in §4.1.

(attentional) control. And we can conceive this skill rather broadly going forwards in terms of an ability to *control attention according to one's goals*.¹⁷

There is growing empirical work to suggest that top-down control is indeed improved through continued Focused-Attention practice (Jha, Krompinger and Baime, 2007; Lutz et al., 2008; Lutz et al., 2009; MacLean et al., 2010). Practiced meditators display faster responses to cue directions and a larger differentiation between the neural responses to different cues, indicating faster or more efficient top-down control in instances of re-orienting (Kerr et al., 2011) along with an ability to sustain attention upon specified objects for longer (Carter et al., 2005). Other, bottom-up attentional capacities are trained in the meditative enterprise (which I shall comment upon later), but it is this top-down capacity that I am interested in primarily here, given that its exercise is most susceptible to controversial and potentially distortive transformations of experience. With this skill identified, it's now possible to approach a better understanding of its phenomenological effects by looking to some dominant models of top-down control in the sciences. This will reveal specific experiential transformations possible through such skill that can be of benefit to our first-person sciences of mind, bolstering the utility of Focused-Attention meditation in such an enterprise.

2.2 The Science of Top-Down Attentional Control

Following Davis and Thompson (2015), we can look to two prominent models of top-down attentional control in the attention literature to bring out its transformational benefits. These models converge upon a similar sub-personal account of the phenomenon and reflect a broad consensus in the field. Each unpacks top-down attentional control in the context of perceptual

¹⁷ In line with f11, this should be taken to cover cases where “control” occurs in relatively direct or immediate consequence of one's intentions.

attention. But given the character of these models, we will later see that they are also relevant to introspective attention.

A first popular theory understands top-down attentional control in terms of the influence of “control-sets” (or “task-sets”) (Duncan and Humphreys, 1989; Wolfe and Horowitz, 2004; Olivers et al., 2011). Call this the “control-set” model. Control-sets are conceived as templates that specify a particular collection of stimuli for further processing and *ex hypothesi* attentional selection (Grubert et al., 2017, p. 843). From this perspective, one exercises top-down attentional control whenever a control-set is activated and biases cognitive processing in favour of its specified stimuli, prioritising incoming data that match those one would expect from those stimuli over other kinds.

Davis and Thompson (2015) note that *working memory* is central to the activation and maintenance of control sets (p. 50; see also Fockert et al., 2001) – a form of short-term memory that transiently holds task-relevant information in an accessible state (Fougnie, 2008). This operates when one holds travel directions in mind for a car journey, or steps in a cooking recipe, while such tasks are performed. It’s believed that control-sets can be activated by holding appropriate conceptual or linguistic representations within working memory (Davis and Thompson, 2015, pp. 50-51). This proposal is nicely motivated by Jones et al. (2010). They show that directing subjects to attend to specific body parts (the hands in this instance) using visual cue words (e.g., showing them the word “left hand” or “right hand”) leads to increased neuronal activity in brain regions associated with these body parts and an improved capacity to detect stimulation applied to these areas (see also Schubert et al., 2006, 2008, 2009). Such directions are thought to transfer relevant conceptual representations to working memory, wherein Davis and Thompson explain that they control attention by mobilising control-sets and

‘directly amplifying early sensory responses to stimuli in this area of the body and inhibiting responses to other areas’ (2015, p. 50).

This model of top-down attentional control is largely consistent with Desimone and Duncan’s (1995) ‘biased competition’ model of attention. Desimone and Duncan posit that early sensory signals in the nervous system—whether these represent objects out in the world or inside our body—are always in competition with one another for further, limited processing resources. Such competition is said to be perpetually “biased” by both other psychological states of the subject and the character of sensory signals themselves. For example, competition can be biased by both the relatively high-level occupants of working memory (e.g. the words “left hand”), and the lower-level salience of sensory stimuli (e.g. the amplitude of pain signals in the left hand). From this perspective, both such factors are always at play in determining which data will be subject to further processing and their objects thereby attended. Whether we describe attention as “top-down” or “bottom-up” from this perspective will thus depend upon the relative balance between competing high-level and low-level forces. Again though, this model conceives top-down attentional control in terms of the prioritisation of local and specific sensory data, through the influence of high-level psychological states, including working memory.

In the attention literature, this local prioritisation of sensory data is said to raise something called *phasic alertness* (Davis and Thompson, 2015, pp. 52-53; Sturm and Willmes, 2001). Phasic alertness denotes a subject’s local or task-specific *sensitivity* to particular stimuli, or stimuli in specified sensory fields, and it is typically indexed by an increased ability to detect and report upon those stimuli. For instance, phasic alertness rises when one watches out for a tennis ball to be served or listens for the starting gun in a sprint-race. In each case, top-down influences raise phasic alertness in the sense that they devote greater cognitive resources to

selected data than competing data, and thereby make the subject more sensitive to correspondent stimuli.

To appreciate the significance of these phasic alertness increases through top-down attentional control, it's important to note a growing consensus that top-down influences raise phasic alertness by causing local fluctuations in a foundational field of *tonic alertness* (Britton et al., 2014; Raz and Buhle, 2006). Tonic alertness denotes one's wider sensitivity to stimuli across the breadth of the sensory and interoceptive fields ("interoception" referring to awareness of the body's internal state, including activity in the muscles, joints and viscera). And research suggests that overall activity levels in the tonic alertness system correlate with a subject's "degree" or "level" of consciousness, with a deficit in the alertness system characterising sleepiness, fatigue and an inability to focus (Britton et al., 2014, pp. 65-66). This link between phasic and tonic alertness can help us understand the nature of top-down attentional control more clearly, for it suggests that top-down attentional control can be understood in terms of the localised heightening of conscious awareness.

As with many proposals in consciousness science, these links between alertness and consciousness are controversial. Longstanding debates over putative distinctions between *phenomenal consciousness* and *access consciousness* question, for instance: how far conscious awareness depends upon the allocation of processing resources to incoming data; which systems those data need access to; and the importance of abilities to report upon relevant stimuli (see Block, 2005; 2011; Cohen and Dennett, 2011; Taylor, 2013). It's not possible to enter these debates usefully here. For present purposes, I assume minimally that conscious awareness is something that occurs in degrees (on a spectrum of strength), and that the *degree* to which sensory data is prioritised and processed by the cognitive systems, and thereby made at least more *accessible* for report, will correlate with the level of consciousness of those data (see

Dennett, 1993; Thompson, 2007, pp. 262-264). From this perspective, top-down attentional control should induce localised *increases* in the levels of conscious awareness, whatever the prior level of awareness happened to be (including no awareness at all). There are many ways to understand what is meant by an increased “level” of conscious awareness (see Bayne, Hohwy and Owen, 2016). A common and reasonable interpretation, however, suggests that an increased level of consciousness, or being “more aware”, involves simply *being aware of more* (Ibid., p. 407). From this perspective, increasing phasic alertness through top-down attentional control should make the subject aware of *more incoming data in selected regions* than prior to the exercise of this capacity.

With these theoretical points on the table, it’s now possible to return to Focused-Attention practice and unpack what’s happening here in greater precision, first at the sub-personal level and then at the personal level.

2.3 Focused-Attention in Focus

In Focused-Attention practice, through repetition of object-appropriate labels the student is (*ex hypothesi*) transferring (or strengthening) a representation of their meditation object to (or within) working memory. As the meditator repeats the term “breathing”, they mobilise an attentional control-set that prioritises the processing of sensory data from bodily structures relevant to the breath, increasing phasic alertness in these regions. Given limited cognitive resources, acts of prioritisation mean that the meditator is also diverting cognitive resources away from competing data, or inhibiting these signals (Thompson and Davis 2015, p. 52). For instance, sustained visual focus upon a candle through appropriate labels (“seeing” or “light”) will prioritise sensory data associated with these labels and inhibit sensory data from elsewhere, including that pertaining to visceral activity, pains or the breath. It is these dual sub-personal

effects of label application that underpin top-down attentional control and thus the ability to return (step 4) and hold the mind (step 1) to the meditative object. And it is these effects that are being trained in Focused-Attention practice. Top-down influence comes in varying strengths or “weights” (see Balluch and Itti, 2010; Goldstone, 1998). So, the more prominently that meditators can make labels feature in working memory, the more quickly/efficiently they should induce phasic alertness changes in cases of re-orienting, and the firmer their minds will be fixed upon the object in step 1, where these phasic alertness changes are sustained.¹⁸

At this point, it’s important to recall that Focused-Attention practices can take either physical or mental objects as targets. In this way they can improve top-down control in cases of either perceptual attention (including interoceptive attention) or introspective attention. Repeated practice upon *one’s experience* of the breath will lead to more skilled top-down attention *to experience*. For the purposes of this paper, I’ll use the term “focused introspective-attention” to denote this particular *introspective* top-down attentional faculty improved in Focused-Attention practice. Focused introspective-attention denotes internally guided (or “controlled”) introspective attention, wherein attention is focused upon a specific mental target by the influence of some independent psychological state of the subject. Moreover, when I speak of a “skill” in focused introspective-attention, I refer to the capacity to guide/control introspective attention *according to one’s goals*, wherein one’s specific targets are in line with one’s intentions. Given that my concern in this paper is with introspection, it is only this form of top-down attentional control (focused introspective-attention) that will concern me henceforth.

¹⁸ The sense of becoming increasingly “fixed” to a meditation object might therefore be linked to improved re-orienting through an ability to effectively mobilise *strong* biases that themselves allow for more sustained holding of the object as they are retained in working memory. Alternatively, it might be explained in terms of a gradual strengthening of working memory representations through the continual repetition of labels each time the mind wanders.

Of course, it's controversial what the difference between perceptual and introspective attention is, or whether there is any difference at all (see Wu, 2014, pp. 253-267). Is there, for instance, any difference between attending to the physical processes of breathing versus the mental sensations of the breath, besides the kinds of judgements that these cases inform? It doesn't seem so. Is the kind of attention one mobilises in cases of introspection then the same as the kind in perception? Or are its *objects* the same? Given that this is a controversial topic, I avoid any equivocation between these two kinds of attention here. Theoretical complications arise when addressing their relation. These would take us too far off-topic here, but I shall return to them in §3.3. For now, it's important to acknowledge only two things.

First, Focused-Attention practices do mobilise a *kind* of introspective attention (top-down, controlled introspective attention) – a kind natural to everyday life, as when we purposefully direct attention to bodily feelings, emotions, pains and hunger pangs, etc. This is what I'm calling focused introspective-attention. Second, focused introspective-attention is mobilised in Focused-Attention practice by employing labels bearing conceptual contents of the same kinds believed to bias the processing of sensory data in models of top-down *perceptual* attention. Given that these labels bear specific, empirically-grounded powers in explanatory models of perceptual attention, we should expect them to have largely equivalent powers (i.e. biasing effects upon the allocation of sensory processing resources) in the introspective case. There's no reason to think that a subject's having an introspective target will somehow strip these conceptual contents of their causal powers or change those powers in any significant manner. Indeed, experimental work on meditators, demonstrating an increased proficiency bringing about the noted subpersonal effects underlying top-down perceptual attention, has used practitioners self-consciously trained in attending to mental rather than physical targets (see Kerr et al., 2011) suggesting an equivalent subpersonal biasing effect.

Employing labels like “breath” or “body” will therefore affect our cognitive systems in the shape outlined, regardless of whether our assumed target is perceptual (physical) or introspective (mental). In either case, conceptual contents will prioritise information about sensory stimuli in the relevant regions of the body and increase phasic alertness in these areas. Similarly, employing the label “seeing” to direct oneself to visual *experience* will exert the same biases as if using this label to direct oneself to the *objects seen*. And as I’ve intimated in §2.2, these rises in phasic alertness have important phenomenological effects, to which we can now turn.

2.4 Effects and Benefits of Focused Introspective-Attention

2.4.1 Phenomenological Effects

By the above models, we’d expect focused introspective-attention to first heighten levels of experiential *richness*, *sharpness* or *granularity* in attended regions (see Davis and Thompson 2015, p. 50; Farb et al., 2015, p. 15; Nielsen and Kaszniak, 2006; Teper, Segal and Inzlich, 2013). Deploying this faculty will heighten phasic alertness, causing localised increases in conscious awareness. And I’ve claimed that this entails the subject *becoming aware of more incoming data* or increasing the number of “grains” composing this region of experience. For this reason, focused introspective-attention shouldn’t be conceived in terms of what Wu (2014) calls “direct models” of introspection (pp. 256-267). According to these models, introspective attention simply “embeds” existing aspects of experience without affecting them. It turns the inner eye without disturbing what it sees. If there is such a kind of introspective attention, focused introspective-attention doesn’t fit the bill.

Importantly, the above models show that focused introspective-attention also involves inhibiting sensory signals from other stimuli (i.e. divert processing resources away from them) thereby decreasing the richness of these aspects of experience. Focused introspective-attention thus alters *overall ratios* of richness across experience. Such alterations are part of what it is to re-direct attention to experience, and they are preserved in sustained attention upon it. Of course, increases in richness mean that aspects of experience stabilised upon will now be different to before. However, it's important here that the Buddhist traditions consider this specific difference to bring advantages, rather than solely problems, in the quest to understand the mind's nature. In Tibetan Buddhism, the degree of mental richness is called *gsal cha*, often translated as "clarity" (or sometimes "vividness", see Thompson, 2015, Wallace, 1999). In the context of Focused-Attention practice, clarity is used to track the subjective richness of the particular aspect of experience focused upon. However, it can also be used in the context of more open (i.e. Open-Monitoring) practices to track the richness of the whole experiential field (Thompson, 2015, p. 76). And Gethin (1998) notes that the major traditions of Buddhism consider the manipulation of mental clarity as a central condition upon which Insight (*paññā*) is built – insight into both (i) the nature of one's own mind and (ii) general laws governing the nature of the human mind, the world and their relation (pp. 174-176; see also Thompson, 2015, p. 76; Wallace, 1999).

I suggest that it is mental clarity (as I shall designate this phenomenological condition henceforth) that lies at the heart of the revelation-versus-distortion debate. The key question is whether changes in clarity through focused introspective-attention must always be avoided when investigating experience, or whether they might be exploited for epistemic advantage, as is the contemplative posit. This latter possibility can be unpacked by first further probing what's involved in clarity changes.

2.4.2 Clarity Considered

The first observation to make here is that increases in clarity should *fill out* or *saturate* target aspects of experience with some mental contents of equivalent low-level kinds to those present within those aspect's pre-attended forms. When introspective attention is turned to the sensations of breathing (via the label "breathing"), this will prioritise incoming data pertaining to the bodily activity associated with the breath. And, given that such data-sets are constrained by the actual bodily activity taking place at the moment of attending, sub-personal selective processes will gather, and raise to awareness, *at least some data* of equivalent low-level kinds to those already being processed and contributing to pre-attentive experience. For instance, they will raise to awareness more data pertaining to the particular dynamics (i.e. the form and intensity of activity) of those body parts involved in the breath. In this way, they increase the quantity of "grains" bearing such contents within an experiential target, filling out or saturating the target.

When making this point, one must be sensitive to so-called "refrigerator light" concerns. Just as there is no light in the refrigerator before it is opened, one might worry that there is no occurrent experience in play before focused introspective-attention is mobilised; direction to previously unattended aspects of the mind might thus simply *create* the experience it purports to clarify and fill out (Block, 2007, p. 489; Jaynes, 1976, p. 23). However, we can here return to the scientific models introduced in §2.2 to mellow this concern.

As specified, both reviewed models of top-down attentional control (including focused introspective-attention) understand its effects in terms of rises in *phasic alertness*. And phasic alertness changes themselves are thought to be explicable in terms of localised increases in a more basic tonic alertness field. Importantly, there are growing suggestions that this foundational alertness system is the minimal system sufficient for our being consciously aware.

An influential theory developed Parvizi and Damasio (2001) designates it the “core consciousness” system (see also Bossier, Jonker and Treur, 2008). Tononi and Edelman (1998) mark it as the “dynamic core” of consciousness (see also Edelman and Tononi, 2000; Edelman, Gally and Baars, 2011). And Davis and Thompson (2015) frame it as the “ground-floor” of consciousness” (p. 49). Moreover, it’s believed that the tonic alertness system, centred anatomically around the thalamus and brainstem, is capable of acting *independently* of top-down selective processes dependent upon higher cortical regions of the brain. This makes top-down processes strictly unnecessary for experience. Rather, top-down forces are said to mould experience by habitually manipulating the workings of the more basic system.

Davis and Thompson (2015) outline this relationship as follows. The tonic alertness system is responsible for their being “something-it-is-like” for the subject. Top-down forces (when active) then manipulate this basic field to determine precisely *which* sensory data one is most responsive to and, thereby, precisely *what* it is like for the subject (p. 49; see also Searle, 2000) – they mould, punctuate and locally concentrate a basic and self-standing field of awareness. And the independence of this tonic system from top-down selective influences means that we are always attending and responding to far more stimuli than are specified by top-down mechanisms, meaning that there should be genuine instances of ongoing (“bottom-up”) experiences that predate such attentional shifts, themselves illuminable through introspective-focus.

The second phenomenological claim we can make about alterations in clarity levels through focused introspective-attention is motivated by the inhibitory effects underlying top-down control in our reviewed models. Improvements to focused introspective-attention will improve one’s ability to inhibit stimuli outside of one’s attentional-set. And this should have *isolatory* effects upon experience. Inhibition of extraneous data, decreasing phasic alertness in

regions outside of one's focus, will isolate target aspects of experience from their surroundings. For instance, focus upon the experience of the breath will inhibit bodily signals underpinning the experiences of pain – it will at least partially suppress these background or peripheral aspects accompanying the experience of breathing, so as to isolate the sensations of breathing themselves. For this reason, Focused-Attention practices are traditionally considered temporary means of *suppressing* the Hindrances, prior to their later elimination (Gethin, 1998, p. 175; Shankman, 2008, p. 92; Thiradhammo, 2014, pp. 22).

With these observations about clarity-transformations laid down, it's now possible to crystallise their epistemic/introspective value.

2.4.3 Introspective Value

The value of the above phenomenological changes can be demonstrated with the help of two analogies. First, take the filling-out of experiences available through increased mental clarity.

Let us imagine an artist, Gina, who has been commissioned to create a pictorial mosaic. Gina has begun work, but she knows that her client is impatient and will come to view the mosaic before its completion. With this in mind, she has sparsely filled each section of her mosaic with a uniformly-distributed set of tiles, each matching the designated section-colour, to give her client a sense of the mosaic's form prior to its completion. Now, let's suppose that the impatient client, after his viewing, decides to observe Gina complete the work, gradually adding in more of the coloured tiles to each section. As this happens, the more vivid the mosaic will become, the more prominently its features will stand out to the client, and the easier it will be to discern what the mosaic depicts.

This is how we can think about the first benefit of focused introspective-attention. Experiences that are vague, dull and imprecise can become fuller, sharper, more complete, and

their properties can accordingly become more salient. Through heightening levels of clarity one is, *as least in part*, increasing the concentration of componential features that afford an experience of its higher-level properties. In this way, one widens the supervenience base of those properties so as to bring out these features of experience more fully. For instance, focused introspective-attention upon feelings of discomfort in one's head will increase the quantity of interoceptive data from the head making it to one's awareness, allowing one to better discern whether one's discomfort is merely a kind of pressure, or an experience of pain. By increasing the supervenience base of these higher-level properties, one might also get a better sense of the character of the lower-levels themselves, better impressing the phenomenological micro-dynamics *underlying* one's discomfort or pain. Replacing a patchy and sparse awareness of these subtle sensations with a fuller and richer kind might allow one to determine, for instance, whether one's pains are composed of dull, aching sensations, or sharp stinging sensations. We can think of both these effects in terms of the *accentuation* of experiential properties through the filling out of experience.

Of course, not all properties of experience are accentuated. Some are changed in the process; the experience becomes sharper, richer, and its form becomes more definite than before, further implications of which shall be broached in §3. Yet, it is through the above accentuatory changes that we can achieve what Gallagher and Zahavi (2008) speak of as the "disclosure" or "articulation" of structures contained within lived-experience.

This accentuatory effect, where properties are made salient through *filling out*, should be distinguished from what we might call the "intensification" of experience. Intensification is closer to the "accenting" of musical performance, where a note is emphasised by increasing its dynamic. Rather, if one wanted to bring the musical analogy to the introspective case, the accentuatory effect is closer to moving from (i) a note played on a solo clarinet, to (ii) a note

played on many clarinets at once. In accentuation, one is *widening* the supervenience base of higher-level properties; in accenting, one is changing the character of (i.e. intensifying) the individual components of the supervenience base themselves.¹⁹

The benefit of the second above noted phenomenological effect of focused introspective-attention—the *isolation* of particular aspects of experience—is well revealed with another analogy. Here, we can imagine an ethologist, Joanna, walking through a forest and looking for wildlife. At some point, Joanna catches sight of an animal moving through the trees and wants to take notes about the animal (its species and behaviour, say). It is prudent here for Joanna to *stop walking* and stand still. This way, she creates a background of stillness against which the animal's character and activity becomes more apparent. Contrarily, if everything is moving in her visual field, accurate discernment is significantly more challenging. This is how we can conceive the introspective benefits of phenomenological isolation.

In the introspective case, one aims not to make judgements about the world, but about one's own mental states, though ideas about epistemic benefit run much the same. With lots of

¹⁹ I don't suggest here that intensification never occurs in acts of focused introspective-attention. In fact, empirical work suggests that top-down perceptual attention to specific properties (e.g. apparent size and spatial distance in visual attention) *can indeed* intensify one's experience of those properties (see Carrasco, Fuller and Ling 2008), so we should think that focused introspective-attention is open to the same effects. It also needs admitting that this distinction between clarity and intensity is a grey one. It may be that the filling out of experience *itself* increases the intensity of an experience, just as playing a quiet note on many clarinets, in a sense, increases the volume of the music. In the present narrative though, I assume that there are genuine instances of increased saturation of experience as *distinct* from increased intensity of experience, through focused introspective-attention. Potential concurrent changes in the intensity of experience (in *particular* instances) won't prove impediments for introspective methods utilising focused introspective-attention, for reasons that will become apparent in §4. Moreover, it's worth noting that the argument I make here could run a similar way by identifying intensity, rather than clarity, as the epistemically beneficial transformation of focused introspective-attention. Just as a *louder* drum is more likely to reveal certain qualities of its sound (it's timbre, texture, etc.) than a quiet one, increased intensity of experience itself plausibly affords epistemic benefits during introspective investigations. On account of this, though I run the argument in terms of clarity here, one might substitute this for intensity, with appropriate narrative adjustments, to reach an equivalent conclusion.

experiential activity outside one's focus, this tends to call attention towards itself. Superfluous data exerts an insidious detracting effect upon introspective judgement, over-shadowing our targets and exerting a *cognitive drain* upon introspective capacities. When such activity is dissolved, introspective judgements can benefit accordingly (Gethin, 2004, pp. 207-8; Colombetti, 2014, p. 147). And it is this isolatory effect of introspective-focus that underpins the beneficial “disentanglement” of experience during rigorous self-observation, earlier proposed by Gallagher and Zahavi (2008).

Confusingly, some more contemporary Buddhist literature refers to this disentanglement of experience again in terms of increased “clarity” (see Frønsdal, 2005). So conceived, clarity would denote the degree to which some aspect of experience emerges *without competition for attention*—without the typical morass of extra experiential phenomena that challenge it for our concern—rather than in a particularly vivid manner. To avoid confusion here, I'll avoid using “clarity” in this sense. However, it's worth emphasising that the models of top-down attentional control reviewed in §2.2 make sense of this dual usage. The prioritisation of some sensory data (underpinning vividness and accentuation) is going to *require* (given limited processing resources) the inhibition of others (underpinning isolation and disentanglement). Thus, the two senses of clarity are really like two faces of the same coin. To illustrate their scientific benefits more concretely though, it's useful to entertain some more example cases, in which we first see improvements to a subject's judgements about *their own* lived-experiences.

First, take the recent enthusiasm for “mindful eating” in clinical (Kristeller and Hallett, 1999; Kristeller and Wolever, 2011) and non-clinical contexts (e.g. Albers, 2012; Bays, 2012). This approach to eating promotes a healthier relationship with food through more dedicated attention to the feelings one has surrounding food. It aims to help people eat only what is really

needed by their body and will most satisfy them. This goal is supported by giving participants forms of attentional training, including mindfulness-based eating awareness training (MB-EAT) (Kristeller and Wolever, 2011), where subjects practice deploying top-down attention to hunger and satiety cues and their emotions surrounding food. More general Focused-Attention practices are also used (Kristeller and Hallett, 1999).

There are good indications that these approaches are effective (Godfrey, Gallo and Afari, 2015; Katterman et al., 2014). And it's theorised that this is underpinned, at least in part, by improved introspective judgements about hunger and emotional experiences. In light of the above models, one way to unpack this benefit (and one endorsed amongst mindful eating theorists) is as follows: directing careful focused introspective-attention to one's hunger sensations fills out our experience with more interoceptive data of the kinds that were previously underpinning our hunger experience (Kristeller and Hallett 1999, p. 358). By accentuating the properties of occurrent hunger experiences this way, and simultaneously disentangling them from similar experiences (e.g. body-based emotional signals), focused introspective-attention is proposed to improve judgements about *levels* and *kinds* of hunger that have been in play, and helps subjects to distinguish between hunger-cues and emotional cues that are often mistaken for genuine hunger signals and precipitate unnecessary eating (Hill, Craighead and Safer, 2011, p. 2; Kristeller and Wolever 2011, pp. 50-51).²⁰ To the extent one

²⁰ The theoretical foundation of the mindful eating movement relies upon a correlation or correspondence between how hungry one *feels* and how hungry one actually *is* (i.e. the kinds of "objective" bodily activity that are induced when the body is in need of sustenance). This seems right to assume as the default relationship (see Spener, 2015, p. 311). This correspondence would also explain slippage in the literature between speaking of improvements to (i) judgements about physiological signals themselves and (ii) judgements about our *experience* of these signals (see e.g. Kristeller and Wolever, 2011). So long as the correlation holds, an improvement in either one should make for an improvement in the other. For a more developed attempt to investigate introspective proficiency by appeal to the successful exercise of skills, see Spener (2015) on "introspection-reliant abilities".

is *skilled* in focused introspective-attention then, the easier it should be for the subject to remain in touch with the ups and downs of their pre-reflective hunger experiences.²¹

For a second example, we can turn to a common practice taught to novice meditators – the “body scan”. This is often given as a preliminary to other meditations and requires the subject to direct attention through different regions of bodily sensation – usually travelling upwards from the feelings in the feet, to sensations in the head. Here, one aim is to become more aware of subtle bodily sensations and feelings of discomfort, including feelings of tension in the face, shoulders or neck. As students move attention in this manner, they are encouraged to release particular tensions to prepare the ground for other practices. And at the conclusion of the practice students report a sense of increased overall bodily ease compared to pre-practice levels, something also reflected by physiological measures (Ditto, Eclache and Goldman, 2006).

One way to understand this example is to posit that focused introspective-attention helps the student discern subtle properties of the background state of bodily experience that they had brought with them to the meditation session. The student systematically disentangles and accentuates various factors that had been contributing to this background experience. In so doing, they are not merely prompted into reflective awareness *that* an experience of a certain kind was ongoing, they are also better able to discern those aspects of the body in which their

²¹ Other effects of attention training might also play an explanatory role here. For example, emphasis upon *acceptance* in mindfulness-practice is also something much emphasised in MB-EAT (Godfrey, Callo and Afari, 2015). Further work can disentangle the relative importance of these different factors in the efficacy of such programs. Though, see Teper et al. (2013) for the suggestion that acceptance and introspective proficiency are mutually reinforcing.

pre-reflective state of discomfort (for instance) was grounded and ameliorate these with appropriate bodily adjustments.²²

During the body-scan the student might also be instructed to look out for experiences of “feeling tone” – the expression used in classical Buddhist texts for the bodily sense of pleasure, displeasure or neutrality arising in response to passing mental contents (e.g. thoughts and imaginings) (see Thiradhammo, 2014, pp. 79-80). Focused introspective-attention to such experiences is then made the central aim of a practice known as “mindfulness of feeling tone”. Here, one can again theorise that focused introspective-attention through appropriate intentions (e.g. subvocalizing “feeling”) devotes greater cognitive resources to those interoceptive stimuli underpinning experiences of feeling-tone. This way, it can help students register the particular hedonic tones that accompany their perceptions/imaginings of particular objects. By accentuating their properties, and disentangling them from the larger experiential whole, focused introspective-attention can help students to become more aware of their habitual reactive tendencies, and thence look out for how their mind picks up and runs with these basic feelings into more complex emotional reactions (see Thiradhammo, 2014, pp. 66-81).

In each of these three examples then, focused introspective-attention supports a better understanding of the subject’s *own* lived-experiences. It induces phenomenological effects (accentuation and isolation) that allow for better discernment of properties possessed by the pre-transformed and pre-attentive experiential landscape. Subjects are able to make *retrospective judgements* about their inattentive experiences by extrapolating back from features of the attended state. One example has targeted experiences of habitual and normally unattended reactive tendencies. And as indicated earlier, intimacy with such customary

²² Certainly, there are also non-introspective factors supporting this increased sense of bodily wellbeing. The diversion of cognitive resources *away* from other aspects of experience that sustain feelings of unease (e.g. unhealthy narratives) will also help the body to relax. I say more on this phenomenon in §3-4.

reactivity is important from Buddhist perspectives; if a person is unaware of the involvement of these tendencies in the perpetuation of their suffering, they will be poorly placed to tackle that suffering.

In addition, the above insights about personal experience ground the second epistemic pay-off of a skill in focused introspective-attention, and the one of concern in the present paper. From some personal insights, it will be possible to generalise to the broader phenomenological truths of concern to our scientific and philosophical theorising. For example, Nyanaponika (2015) talks of top-down, label-use as a means to help students ‘dispel the illusion that mental processes are compact [and] [...] discern their specific nature or characteristics’ (pp. 80-81; see also Thiradhammo, 2014, pp. 79-81). From this perspective, bringing experiences into relief through top-down attentional control helps deconstruct complex and compact phenomena, such that their defining features become more apparent – broad truths about experience in general. Taking bodily experiences of desire as our archetype, this can proceed by first bringing attention to many different cases of desire, to reveal both the number and characteristic features of different *intra-subjective* types, including the precise bodily sensations underpinning them. These results could then be corroborated across many individuals (e.g. through discussions with other meditators and teachers, the reading of Buddhist texts, or strict methods of “intersubjective validation” in continental phenomenology) to reveal the *general* subjective nature of desire, or of hedonic tone, for instance – the kinds of truth in which our science can deal. In this way, experiential transformation can play a role in illuminating the general nature of lived experience.

2.5 A Foundation to Build On

The above has sketched how experiential transformations of the kind practised in Focused-Attention might be able to illuminate lived-experience. I've identified specific transformations induced by a trained skill in focused introspective-attention capable of supporting the quest for *general* truths—truths about the human mind, not merely one's own mind—of interest to scientists and philosophers. To do this, I've argued that converging models of attentional top-down control suggest that focused introspective-attention heightens phasic alertness in targeted regions of the mind, whilst quietening other stimuli contributing to experience. This implicated certain phenomenological changes—*accentuation* and *isolation* of experiential properties—that could be induced on demand to illuminate the mind.

In so mobilising the resources of attention science, I hope to have made a basic premise of Buddhist thought more intelligible—that we can approach truths about experience through its transformation—thereby undercutting the distortion assumption (the assumption that methods involving transformation will necessarily produce misleading or misrepresentative accounts). Nonetheless, the attention literature can do much more than simply make meditative methods plausible; it can actually make them more robust to further criticism and help flesh them out more carefully. After all, so far, we have only some suitable foundations upon which an account of meditation's utility can be built. And further use of the attention literature can extend these foundations to show how such transformations might be sensibly exploited within first-person methods. Specifically, the models of top-down attentional control reviewed here help to unveil some better-founded, distortion-oriented concerns about the use of meditation in science, which themselves need addressing if we are to approach a rigorous and practical methodology for the employment of contemplative practice in first-person methods.

In the next section, I'll show how the attention sciences help us rework the distortion assumption into more biting distortion-oriented concerns. For, though such models demonstrate the *availability* of epistemic benefit through a skill in focused introspective-attention, they also spark concern that this skill can be misappropriated – that it might be wielded unwisely in first-person methods to yield genuinely misrepresentative accounts of lived experience. As I'll outline, they implicate several *additional* phenomenological transformations available through focused introspective-attention (some of which have already begun to rear their heads), which have more problematic effects upon experience. This reveals numerous pitfalls available through the imprudent use of focused introspective-attention, which one will need to be sensitive to when devising methods of introspective investigation.

Many of these dangers have already been alluded to in the philosophical, phenomenological and psychological literatures, and I shall relate them to their historical forebears when possible. Nevertheless, as we shall see later, the attention literature doesn't simply reinforce these older concerns; it actually helps to identify the crux of the problems – to distinguish the *roots* of these dangers. And once this is done, we will see in §4 that these models leave theoretical space for the avoidance of such problems, leaving an important place for top-down attentional skills in the investigation of the mind.

3 Further Distortion Concerns

3.1 Objectifying the Subjective

A first, and more challenging, distortion-oriented worry concerning introspective methods employing focused introspective-attention was well-captured by the neo-Kantian thinker Paul Natorp. Natorp (1912) noted how purposeful attention to experience could transform the

subjective, something *identified* with, into something set apart from oneself. ‘[O]ne apparently never grasps the subjective, as such’ Natorp states, ‘[... o]n the contrary, in order to grasp it scientifically, one is forced to strip it of its subjective character’ (p. 103, cited in Zahavi, 2003, p. 157; see also Petitmengin and Bitbol, 2009, p. 366, 377).

Take the feeling of tiredness following a night of poor sleep. Throughout the day, this sense of tiredness is largely “lived through”, with the objects of one’s attention being the contents of the world itself. The feeling of tiredness might influence the way that we orient ourselves towards the world, but it is not usually something we are directed towards. Yet, when the subject deliberately attends to the tiredness, the experience is transformed from something lived through *onto* the world (something subjective), into something to which they are now opposed. In so making the tiredness an “object” of attention, we therefore introduce some novel volitional or agentic component to experience (a substitute “subjective” component) that restructures the conscious landscape and allows the tiredness to become an “object”.

Prima facie, this concern is especially pertinent in the present context. By the earlier models, acts of top-down attentional control over experience require working memory to be loaded with conceptual representations that can mobilise appropriate control-sets. Consequently, these will usually be dependent upon deliberate acts of sub-vocalized intending (labelling), introducing the above sense of positionality to experience. So conceived, the benefits of focused introspective-attention appear to require, at least in the normal case, the generation of new experiential and relational properties of the kind that Natorp describes. This kind of transformation is not a revelatory one; it does not itself *bring out* features of lived-experience in any obvious manner. Rather, it is the addition of something alien to lived experience which can be thought to either “distort” the overall landscape of the mind, or

potentially overcloud those aspects of the mind we are interested in (given that addition is the very opposite of the *isolation* posited to have epistemic benefits).

Worries of this sort can be strengthened by highlighting contemporary portrayals of meditation as a means to gain increasing “detachment” from experience (see e.g. Sujîva, 2000, p. 179; Nyanaponika, 2015, p. 89). One seems to be replacing ordinary immersion in experience with something peculiarly distanced. Through constant repetition of labels, and the development of greater skill in focused introspective-attention, one might even worry that we are *worsening* regular objectification problems attached to deliberate attending. Proficiency in focused introspective-attention therefore seems at best to involve, and at worst to exacerbate, the novel experiential and relational properties bestowed upon introspective targets in the regular case.

3.2 Stilling the Stream

Along with objectifying experience, Natorp (1912) noted how deliberate attention to the mind could petrify or deaden its natural, flowing character. He spoke of this as ‘killing subjectivity in order to dissect it’ (p. 102). One is forced, he says, ‘to artificially still and interrupt the continuous stream of becoming, which surely is how inner life presents itself, to isolate the individual finding, to fixate it with the isolation in mind, to sterilize it, like the anatomist does with his specimen’ (pp. 101-102). Here, Natorp echoes William James’ suggestion that attempts to investigate experience were akin to ‘seizing a spinning top to catch its motion’ (1890/2007, p. 24).²³ Both articulations of this concern aptly capture the problems underlined

²³ Petitmengin and Bitbol (2009, pp. 366-367) give a good survey of other ways this concern has been elaborated.

by the above models of top-down attentional control, for its underlying *inhibitory* character means that many natural elements of experience are in danger of being lost to the introspector.

On the one hand, focused introspective-attention forestalls the natural tendency of the mind to switch between varied mental contents (sensations, volitions, thoughts, images, etc.). This suggests that it is of limited benefit if one wishes to illuminate spontaneous patterns of activity that manifest across the breadth of the mind. On the other hand, even narrow-scale dynamism can be undermined, given that some of these broader elements may in fact be helping to retain the shape, character or flow of those aspects we turn towards. Wundt (1897) and James (1890/2007, pp. 243-245) believed this made it impossible to learn about the natural flow of human thought through directed attention (see Petitmengin and Bitbol, 2009, p. 366), given that thought's character is often dependent upon being in the background and emerging in involuntary response to other mental contents. Similarly, Brentano (1874/1995, p. 30) suggested that deliberate attention to one's anger would "diminish" the anger itself. Our models of top-down attentional control underline this point, suggesting that turning attention to the intentional act of anger requires diverting resources from the perceptual or imagined object sustaining the anger in its original form. One can thus disarm the emotion of its object and slow the "spinning-top" of emotion.²⁴

In light of the above then, one might fear that the benefits of focused introspective-attention will come at the expense of "deadening" experience, taking it further away from the experience of life "as lived" by untrained persons. And again, one might be concerned that this deadening effect is simply exacerbated by meditative training in such things.

²⁴ See Spener (2018) for an account of how this worry informed the experimental methods of thinkers in the early to mid-twentieth century, particularly those of Introspectionist and later Gestalt Psychologists.

3.3 Complex Experiences, Intentionality and Different Probes

The next distortive danger arises when investigating more complex experiences than most of those referenced above. Perhaps focused introspective-attention has some relatively straightforward advantages when it comes to simple bodily experiences like hunger and hedonic tone. But other experiences like emotions will have complex internal dynamics that present more difficult challenges for investigation. For one thing, unlike simple interoceptive experiences, emotions have an *intentional structure*, of which different aspects will be open to investigation.²⁵ It will be important to distinguish how these different aspects can be probed, and to avoid running them together when gathering introspective data.

This difficulty is underlined by the above scientific models, which emphasise that the phenomenological effects of focused introspective-attention will be heavily dependent upon the character of the labels/intentions used to direct attention, not to mention the fact that focused introspective-attention also *inhibits* that which is outside its focus. Insensitivity to these facts, and failure to discern appropriate probes, could result in conflation between reports about different aspects of experience, or the neglect of important parts of an experience being targeted and resultant theoretical overemphasis on merely some of its features. These dangers are well illustrated by the case of emotional experience.

Emotions are widely thought to possess an intentional structure (Kind, 2013, p. 117; Goldie, 2002). By this, we mean that they are *directed towards* or *aim at* something. As with other intentional states, that which they aim at (some person, object, state of affairs, e.g.) is their “intentional object”, and their specific manner of directing us towards that object is their

²⁵ See Dahlstrom (2014, pp. 149-153) for a good discussion of interoception and intentionality.

“intentional act” (or “intentional mode”).²⁶ In the case of fearing a particular person, the intentional object is the person feared, and the manner by which one is directed towards them—fearfully—is the intentional act, which will have its own characteristics that distinguish it from other emotions, such as being *lovingly* or *angrily* directed towards that person. This structure yields a complication for the probing of emotion that’s revealed by considering an instruction often put to meditation students.

Take the request to “observe the Hindrances that drag one away from the present moment”. The Hindrances, recall, are episodes of obsessional desire or craving—the felt necessity for particular things to be other than they are—that form a central introspective target of Buddhist contemplative programmes. And they can be probed in several different ways. For instance, observation might target the *affective dimension* of craving – the bodily feelings of lust, or aversion say – that is, the intentional act. Contrastingly, it might target *what* one craves or “feels towards” (Goldie, 2002, p. 241), as when asked to “confront one’s fears”, i.e. the *intentional object*. This might be a specific future event, with particular features that explain one’s fear, discernment of which might involve attention to more cognitive dimensions, like thoughts or mental images. Perhaps there are also kinds of probes able to target the entire emotional complex at once, including both act and object.

Engaging focused introspective-attention to emotions will need to be sensitive to the possibility of probing in these multiple ways. One must ascertain both how and whether each kind is initiated. What will be the appropriate labels for directing us to the intentional act? And how will these be distinguished from those that thematise the intentional object? More broadly,

²⁶ Intentionality theorists also speak of intentional “contents” as distinct from “objects”, which designate *that which one attributes* to the object (e.g. dangerousness in the case of fear). For more on this distinction, see Crane (2000, pp. 51-53). I avoid such talk, given that the act-object distinction is sufficient to motivate the concern related here.

any systematic introspective employment of focused introspective-attention should proceed by first contemplating the complexity of one's targets and whether they are open to such differential probing.

Take another common target of introspection: perceptual experience. Can one attend to the intentional act here? To the act of seeing a cup, say? Some claim not, arguing that any attempt to do so leads one to "look through" the act (of seeing) to the presented object itself (the cup). This is the "transparency observation" about perception.²⁷ Advocates of transparency propose that introspective attention to perceptual experiences can only attend to the same object (or features) as perceptual attention – it will entail seeing through to the object itself (Grice, 2002, p. 45; Harman, 1997, p. 667; Wu, 2014, pp. 257-262).²⁸

For some such thinkers, the introspection of perceptual experience will be conceived akin to the 'Transparency Model' entertained by Wu (2014, pp. 258-267) or Dretske's (1995) 'deferred perception' account. In such accounts, the introspection of perceptual experiences proceeds by simply applying psychological instead of empirical concepts in the course of ordinary perceiving. I introspect a perceptual experience of X by first perceptually selecting some object (X) and then biasing my judgements in favour of the relevant psychological concepts, i.e., speaking in terms of *how X looks* rather than *what X is*. Schwitzgebel (2012)

²⁷ Note that transparency advocates needn't claim that introspective knowledge is exhausted by knowledge about the (intentional) objects of experience presented. The intentional act or mode of a perceptual state, i.e. the perceptual modality (vision, olfaction etc.) *through which* that object is presented, clearly transcends this kind of knowledge, despite being open to introspective judgement (Crane, 2000, pp. 59-60; Thompson 2007, p. 285). The transparency observation as construed here is a mere phenomenological claim that attention to experience can retain only the perceptual object as *thematic* (that which is attended to). It says nothing about the extension of the judgements we can make subsequent to this.

²⁸ Whether one wishes to describe the resultant introspective attention here as attention to the (external) object or the *intentional* object will depend upon whether one favours internalist or externalist accounts of perception's intentional object, as well as broader issues about the objects of illusion and hallucination. See Crane (2000, pp. 55-58) for more on this issue. This point is not significant here. All that is important is the claim that any attempt to attend to the *act* of perception itself will fail.

marks introspection of perception so conceived as ‘perception with a twist’ (p. 35) – it is the pairing of perceptual attention with a novel introspective or psychological attitude.

For transparency theorists then, focused introspective-attention upon perceptual experience can bring introspective advantage only by bringing clarity to the intentional object of perception. It can tell us more about the experience only by accentuating and isolating what the experience is *about*, something that remains a proper part of the experience’s phenomenal character (see Wu, 2014, p. 258). For instance, from this perspective, mobilising or heightening focused introspective-attention to tactile sensations of a table can only reveal what those sensations are about - whether that table is felt *as* something hard, or smooth, or rubbery, or greasy. Unlike in the case of emotion, focused introspective-attention would not be able to illuminate (in any direct manner, at least) features of the *act* of experiencing – the *act of feeling* the table, independent of *what* is felt.

However, the transparency observation is controversial. Thompson (2007) suggests that, while attention to perceptual experience *usually* looks through to the (intentional) object, there is a way by which we can (‘with effort’) attend to the act of perception (p. 284). He calls this the ‘moderate transparency thesis’. For Thompson, it’s possible to attend to an experience’s ‘subjective features’. An experience’s subjective features are not qualities of the object (as are attended to in perception) but qualities of the way that said object is brought into view (pp. 285-287). And through such attention, he suggests that features of experience on the side of the intentional act, which usually remain implicit or latent, can be made explicit and available for phenomenological consideration (p. 287). I will return to this complex issue in §4. For now, it’s sufficient to note that, if Thompson is correct, the use of focused introspective-attention in the investigation of perceptual experience will also need to be sensitive to the possibility that

different aspects of perception's intentional structure might be interrogated with its aid. And this will require identifying appropriate labels for initiating the respective probes.

3.4 Increased Richness, Increased Detail

The fourth danger surrounding the use of focused introspective-attention is that its capacity to increase the granularity of experience can deceive us into thinking lived experience contains much more than is really the case. For example, when we shift attention to a peripheral aspect of our visual experience, we find it rich in colour and detail. However, empirical work shows an inability to accurately report upon colours in the periphery (Ferree and Rand, 1919; Moreland, Jameson and Hurvich, 1972). We also know that subjects can fail to detect overt and incongruent phenomena, if primed to be selective of (focus upon) only certain features in a scene; Simons and Chabris' (1999) famous "invisible gorilla" experiment showed that subjects often failed to recognise a man in a gorilla suit walking through a ring of people passing a basketball, if asked to count the number of passes made. This supports a phenomenon of "inattention blindness" to many features of the world (see Mack and Rock, 1998)

In §2, I suggested that "refrigerator light" objections (per which, *all details* of experience are new ones) are too extreme, if offered as default objections to acts of bringing the pre-reflective to attention. Nonetheless, we'd be naïve to think that genuine cases of attending to (conscious) pre-reflective experience will "fill out" that experience only with those data that were already present. Not only can focused introspective-attention add very overt things to experience, of the sort missed in inattention blindness experiments, but the increased granularity that it affords will also capture subtler detail and nuances. For example, focused introspective-attention to gustatory experiences in "mindful eating" programmes may support awareness of subtler flavours and details in one's food, enriching the eating experience in a

manner that underpins the increased enjoyment of food reported in such programmes (Hong, Lishner and Han, 2014).

Relatedly, one should be careful of thinking that the prioritisation (and raising to awareness) of even equivalent sensory data through focused introspective-attention will be entirely homogenous. This process may well differentially prioritise stimuli from the original set. This means that the increased richness of experience can also reset the balance of features in experience, causing it to display a novel and perhaps more intricate structure. For example, attending to the taste of sweetness in one's coffee, might prioritise specific sweetness flavours in the coffee over others, rebalancing the original ratio and turning the experience into a *particular type* of sweetness, not the vague and diffuse taste one had previously. In these cases, focused introspective-attention is introducing non-trivial novelty to the experience, rather than unobtrusively "accentuating" its existing features, through clarity increases.

3.5 Conceptual Tainting

A final danger worth mentioning here concerns possible *conceptual distortions* introduced to experience through focused introspective-attention. A growing body of research argues that experience is subject to widespread "cognitive penetration" (see Zeimbekis and Raftopoulos, 2015), as when a subject's irrational belief that person X is angry with them might cause them to experience person X's expression as more "angry-looking" (Brogaard and Chomanski, 2015, p. 472). In a similar way, we might worry that attempts to direct introspective attention by mobilising conceptual representations generate something capable of infecting the experience itself.²⁹ These concepts might introduce novel conceptual content to experience. Mobilising the

²⁹ See Brogaard and Chomanski (2015, pp. 470-472) for thoughts on the relation between these two kinds of case.

concept *CANDLE* to direct attention to one's visual experience of a particular object (i.e. a candle) might itself introduce the representational content *CANDLE* to the experience, where it was previously absent. Or, for the less representationally-oriented, these concepts might yield new "gestalts", wherein non-conceptual content is restructured into new figure-background relations.³⁰ If so, the very method used to examine experience is pre-determining what one finds there.

These concerns also extend to theoretical concepts. It's suggested that even the philosophical schema one brings to introspection can taint what one discovers through it. Firth (1949) proposed this as one means to explain disputes over sense-datum theories in the early twentieth century, questioning whether 'underlying prejudices' at play might 'prevent many people [...] from examining perceptual consciousness with complete objectivity' (p. 452). One way to interpret this is to say that philosophical conceptions about experience might lead subjects to initiate different kinds of probe, that accordingly transform experience differently. Those believing perception to be entirely transparent may have been targeting only its objects, while those favouring less transparent conceptions may have been sensitive to and picking up upon *other* aspects of perceptual experience, explaining divergences between the two groups (pp. 462-463; see also Spener, 2018, pp. 153-156). This returns us to the concerns raised in §3.3, to which we can now add the possibility that even different *philosophical conceptions* of experience might induce different kinds of probing, yielding different results.

One finds similar issues discussed in epistemological debates internal to the field of Buddhist Studies. Thompson notes that it remains 'an open and interesting question' in the field whether meditative experience informs, *or is informed by*, Buddhist philosophical ideas (Thompson, Varela and Rosch, 1991/2017, p. xxiii). One can well argue that canonical

³⁰ See Siegel (2006) for more detail on the difference between these two accounts.

Buddhist texts and meditation manuals might themselves shape the experiences of meditators. For instance, a philosophical schema favouring the *discreteness* of experience might promote a means of access that yields an experiential landscape mirroring such a schema, rather than revealing the mind as it is independent of such concepts (see also Thompson, 2015, pp. 56-57). Some even argue that attentional skills might be trained *precisely to alter* experience, such that it better accords with doctrinal truths. Sharf (1995) outlines how the historical assignment of genuine knowledge to purported cases of meditative insight ‘often require[d] the complicity of spiritual exegetes [...] called upon to attest to the orthodoxy of one’s meditative accomplishment’ (p. 270). He suggests that Buddhist meditation might better be considered a ‘script for performance’, or ‘*ritualization* of experience’ (p. 269), serving to legitimate the doctrine in traditional scripture and preserve a certain unity amongst the tradition.

3.6 A Route to Handling Distortion

The revamped distortion-oriented concerns reviewed above present a more formidable challenge to the use of meditation in the study of the mind. They do not simply equate the transformation of experience with its distortion but highlight *specific kinds of change* fostered by a skill in focused introspective-attention (trained in Focused-Attention practice) that are counterproductive and can promote a misleading picture of lived experience. In sum, they suggest focused introspective-attention brings a swathe of dangers along with its benefits. Nonetheless, I suggest that none of these dangers is severe enough to warrant relinquishing such a skill. Rather, one can retain an epistemically-beneficial place for it (and thus for Focused-Attention practice) in the study of the mind, so long as one is sensitive to *how this skill is used* within our introspective endeavours. It must be used in a way that *exploits* the epistemically-beneficial transformations of focused introspective-attention whilst either

minimising or accounting for those more misleading and deceptive transformations noted above.

As I've emphasised, contemplative theory retains an important role for top-down skills. In the final section, I'd thus like to bring the attention literature into dialogue with the pedagogical literature on meditation to advise how to use focused introspective-attention prudently. Here, I'll look to the instructions for "insight" practices, where top-down attentional skills are utilised for epistemic benefit. Doing this, we will see that the actual specifics of meditation instruction are such as to side-step, minimise or address many of the dangers just outlined. They will reveal: the kinds of introspective target for which focused introspective-attention is appropriate; the manner in which it needs to be employed; and the point at which it needs transcending.

Moreover, I shall show that these pedagogical suggestions are fully consistent with the models of top-down attentional control reviewed above – such models leave theoretical room for (i.e. they can explanatorily capture) the less problematic ways of utilising focused introspective-attention indicated in the instructional literature. Delineating this can therefore help us flesh out a more careful approach to the use of focused introspective-attention, that can be replicated in contemporary scientific contexts. Along the way, we shall also acquire a more nuanced sense of the way that transformation of the mind can be handled within introspective methods, enabling us to distinguish several different ways in which transformation and insight can sit together.

4 Using Attentional Skill: Pedagogical and Scientific Considerations

My turn to the pedagogical literature on meditation centres upon the contemporary Insight Meditation Movement and closely-related “Burmese style” *Vipassanā* tradition, rooted in the teachings of Mahasi Sayadaw (1904-1982). The central elements of this twentieth-century Theravāda meditation “revival” draw directly from canonical Theravāda material, especially the ‘Discourse on the Establishment of Mindfulness’ (*Satipatthāna Sutta*), and the commentarial material of Buddhaghosa, centring on his ‘Path of Purification’ (*Visuddhimagga*) (Sharf, 2015, pp. 472-473; Cousins, 1994).

Theravāda is typically considered the most conservative of the Buddhist traditions and closest in doctrine and practice to Early Buddhism (Gethin, 1998, p. 1). Focusing on contemporary renderings of Theravāda thought thus allows me to strike a balance between (i) thematising relatively “foundational” aspects of Buddhist thought, consistent with my earlier aims, and (ii) avoiding the need for heavy-duty exegetical work needed to unpack the nuances of classical meditation manuals. Moreover, these contemporary manuals place special emphasis upon the “insight” stage of contemplative practice (hence their name), where attentional skills trained in Focused-Attention and Open-Monitoring are put to use for epistemic benefit. They also contain a wealth of nuanced student-centred, pedagogical advice that is de-emphasised in classical texts at the expense of the aesthetics of presentation and structure.

This thematic choice is therefore pragmatic, serving to ease exposition. As such, it must be acknowledged to arrive at the expense of strict representativeness to the very earliest

Buddhist meditation instruction.³¹ Nonetheless, there is no need to privilege older traditions when looking for “authentic” instruction. As with other later manifestations of Buddhism, the Insight Movement is rooted in the classical canon and is one of many forms of a *living* tradition that attempts to present those foundations in a manner appropriate to its environment. It can be considered here as just one pragmatic model for insight-oriented practice – one amongst many.

4.1 Refining Concentration

The first noteworthy feature of Insight Meditation manuals is their emphasis upon the possibility of *refining* the concentration that is induced by focused introspective-attention. Contrasting the relatively coarse-grained western commentary on this topic, the Buddhist traditions have historically emphasised a spectrum of increasingly pure kinds of concentration, not all of which suffer the above problems. In contemporary meditation manuals, one thus finds advice on how to “purify” concentrative experiences. This advice is sensitive to canonical distinctions between “right concentration” and “wrong concentration” and takes inspiration from descriptions of the *jhānas*—a series of increasingly refined concentrative states—catalogued in texts like Buddhaghosa’s *Path of Purification*.

As teacher Sujîva (2000) notes, “‘concentration’ actually covers a wide range of experience’ (p. 143). Concentration is said to denote the “holding” of an object. And Sujîva distinguishes a number of increasingly refined kinds of holding. He notes that the aim in Focused-Attention practice is to hold ‘without clinging defilements’ (p. 145) - ‘[n]ot with obsession, not with anger not with greed but with clear awareness’ (p. 163). Thus, there are forms of concentration to be avoided in insight practice – primarily those *forceful* types

³¹ There is significant debate over how representative these contemporary meditation revivals are of canonical instruction. For nuanced accounts of the Insight Movement’s relation to older material, see Cousins (1994) and Sharf (2015).

manifesting feelings of necessity for some outcome and distinctive of craving (*taṇhā*). For example, impatiently slamming one's attention back upon one's meditation object, having been distracted for the hundredth time, is a case of unrefined, impure or “wrong” concentration.

It's common for meditators to note that concentration becomes increasingly easy with practice, requiring less pronounced forms of effort and intention, and this is accompanied by a lessening sense of positionality against one's object (Lutz et al., 2008; Wallace, 1999). The observation that feelings of positionality begin to dissipate, as intentions become less pronounced, is consistent with early Theravāda theories of mind. Such accounts hold the mental factor of craving, which is believed to accompany the majority of our intentions, not merely productive of suffering (*dukkha*), but also a critical component of the sense-of-self, to which this feeling of positionality or being *over-and-against experience* is intrinsic (Albahari, 2006, p. 27, pp. 61-63).³²

These points suggest that the first distortion concern noted above—the seemingly fundamental sense of positionality that accompanies concentrative experiences induced by focused introspective-attention—may actually latch onto particularities in the way that concentration is initialised or sustained, rather than something intrinsic to it. Certain properties that we take as necessary may be common but inessential extras that can taint a narrower phenomenon. It is for this reason that Buddhist meditation is often highlighted as mobilising a “bare attention” (Colombetti, 2014, p. 156) wherein the attentional systems have been purged of elements not strictly proper to them. These things include particular volitional and affective

³² It's important to recall that craving (*taṇhā*) is something held to manifest most often on a very subtle level in contemplative thought, rather than covering merely the kinds of overt longing we commonly associate with it. It is this commitment to the subtlety of craving that allows it to play an important role in Theravāda accounts of the sense-of-self. Moreover, one should distinguish the sense of positionality from the *perspectival quality* inherent to many experiences. While perceptual experience, for instance, has an inherent perspectival dimension, wherein objects are presented from a particular spatio-temporal location (or “point of view”), this is distinct from the positionality felt as we are set *against* perceptual experience itself (see Albahari, pp. 6-21).

tones that accompany many intentions used to shift or sustain focus, as when one anxiously attempts to hold the mind upon an object as a means to block out thoughts concerning some recent traumatic event.

Focused-Attention is thus a pragmatic approach for *refining* concentration so less distortive kinds can be mobilized. In fact, many traditions of Buddhist literature reference completely “non-dual” forms of concentration, where the sense of opposition to an object disappears entirely (see Dunne, 2011). Describing concentrative practice upon a *kasina* – a mental image of a “circle of light” – Sujîva (2000) notes that:

There will come a time when there is unification—the mind and the circle of light are one. That moment, when one does not seem to be able to differentiate between the two, is what we call samadhi. It is a kind of absorption. As long as one is still consciously knowing and differentiating at that very moment, is it still access concentration (the degree of concentration traditionally held necessary for “insight” practices) [...] There is no subject-object differentiation at the moment of samadhi (p. 149, parenthesis added)

In samadhi, the sense of positionality associated with craving is entirely absent – here there is no sense of a subject “doing” the concentrating. And this possibility is consistent with the above models of top-down attentional control, if they are unpacked with sufficient care. The establishment of concentration, through focused introspective-attention, is tied here to the holding of conceptual representations in working memory. But while the effortful and intentional rehearsal of appropriate labels assists in this process, there is nothing about such models that require their incessant employment. Appropriate representations can remain in working memory for significant amounts of time without rehearsal (Ericsson and Kintsch, 1995) allowing one to deploy focused introspective-attention without the constant need for linguistic commentary that induces feelings of positionality.

This possibility is also outlined in meditation manuals devoted to the establishment of the *jhānas*. Here, it's suggested that coarser kinds of concentration proceed through the factors of *vitakka* and *vicāra*, with the two usually being introduced together as the amalgam *vitakka-vicāra* roughly translated as “applied and sustained thought” or “initial and sustained mental application” (Shankman, 2008, p. 39). Here, concentration is maintained through effortful label use. But Shankman (2008) outlines that, as one moves through the *jhānas*, these are left behind:

As concentration deepens, the mind becomes more still. The mind in the second *jhāna* is free from discursive thought [...] *vitakka-vicāra* [...] drops away in the deeper levels of *samādhi*. Concentration has been sufficiently strengthened so that it need not be tethered to an object by the factors of *vitakka* and *vicāra*, since it naturally remains steady [...]. At this stage the awareness remains stable and unbroken. The *Samṇamṇḍikā Sutta* states that wholesome intentions, a form of mental activity, cease without remainder with the subsiding of *vitakka-vicāra* upon entering the second *jhāna* (pp. 44-45).

This suggests that, though applied thought is usually important for the *development* of concentration, it is not essential to concentrated states themselves. The subject can eventually engage a more refined focus upon experience that proceeds without the feelings of positionality associated with intentional thought. Moreover, given that this form of concentration is said to remain naturally steady, despite the fact that representations in working-memory are thought to decay over time (Barrouillet et al., 2017; Lemaire and Portrat, 2018), these are likely forms of introspective concentration that do not fit under the “top-down” umbrella at all. They would be analogous to cases of involuntary perceptual “absorption” in some natural scene, as when the beauty of a sunset keeps one glued to it, or cases where a particularly interesting stimulus holds our attention independent of any intention or agenda of our own. These are *bottom-up* forms of sustained attention, where one’s attentional targets are not strictly “controlled” by

one's intentions in a direct sense. While it's not possible to go into detail about these states here, Lutz et al. (2008) offer a useful "dynamicist" account of such experiences:

[high-level] meditation states might not be best understood as top-down influence in a classical neuroanatomical sense, but rather as dynamical global states that, in virtue of their dynamical equilibrium, can influence the processing of the brain from moment to moment [...]. In this alternative "dynamicist" view of top-down control, spatio-temporal trajectories of neural activity emerge from complex non-linear neural interactions following rules of dynamical theory [...] In this view, the brain goes through a succession of large-scale brain states, with each state becoming the source of top-down influences for the subsequent state (p. 167)

By Lutz et al.'s (2008) model, the maintenance of these maximally pure forms of concentration need not be explained by any *additional* psychological state, meaning that these states are not instances of top-down attentional control as understood in the present paper. Rather, the experience *sustains itself* when isolated from perturbing factors, in virtue of the self-organizing properties of these biological systems (see Thompson, 2007, chpt. 3).³³ In respect of this, it's plausible that Focused-Attention practice not only trains a more refined kind of focused introspective-attention, but that such top-down states themselves can lead to bottom-up forms of sustained and introspective attention, which might also be exploited within science.³⁴ The description of these two possibilities within the meditative literature, as well as their theoretical consistency with the contemporary attention science in question, therefore works against the first distortion concern noted in §3. Both motivate thinking that the alien positionality

³³ The importance of isolation from other influences here explains why we find it so hard to become immersed in something when we have a lot on our minds. Becoming properly absorbed in nature, for instance, usually requires that we've somehow managed to set aside our habitual list of obligations.

³⁴ Such advanced non-positional kinds of concentration are usually *not* considered requirements for insight practice though (Sujīva, 2000, p. 149; Gethin, 1998, chpt. 7; but see Bronkhorst, 1993).

associated with focused introspective-attention can be lessened by appropriate training in these two faculties.

At this point, it's also worth noting something rarely commented on in responses to this objection – the sense of positionality is not, in fact, alien to all lived experience. Often, our experience is one of being overtly positioned *against the world*. We are not always immersed in the world and its objects, as objections of kind outlined in §3.1 intimate, but often felt to be detached from and standing against those objects. This phenomenological sense of opposition is thus a feature of some lived experiences, which can be captured by focused introspective-attention. If transparency theorists are correct in suggesting that the attention mobilised when introspecting perceptual experience in a top-down fashion is nothing other than top-down *perceptual attention*, then we are merely replicating the phenomenological dimensions characterising top-down perceptual attention, rather than “distorting” lived-experience. Even if the two forms of attention are different, both will sometimes involve a dimension of positionality (irrespective of what one is positioned against), mellowing the concerns about the unrepresentativeness of reports gathered through top-down introspective skills.

In spite of this qualifier though, certain features of experience will be overshadowed if this positional attitude is all one can bring to bear in investigations, for the reasons reviewed above. And this section can be concluded with some concrete recommendations for introspective methods motivated by prior considerations. First, the examination of subtle features of experience would be best to proceed with prior competence in more advanced concentrative states (whether refined top-down kinds or bottom-up kinds). Feelings of positionality are unlikely to undermine attempts to distinguish salient differences between coarse emotional states like anger or fear. However, they are much more likely to subvert phenomenological investigations of the micro-dynamics of implicit bias, fleeting associations,

or the fast-changing temporal properties of perceptual phenomena. Without a capacity to rid the mind of things inessential to these targets, introspective reports of such properties may be tainted or hampered by features not strictly proper to them. More broadly, what distinguishes the two kinds of case above is the relative difference in prominence, intensity or salience between the experiential properties we are investigating, and those we are generating in order to target such things (feelings of positionality, volition, intention etc.). Subtle aspects of experience will require a subtlety of concentration so as not to be overshadowed by more overt features of consciousness.

At this point, we can also distinguish the first broad way that transformation can be handled appropriately within first-person methods. As in the above mobilisation of concentration, one can work to (1) mobilise beneficial transformations, whilst *eliminating* detrimental transformations (novel and overt positionality, in this case) at the very earliest stage of investigation. This is the simplest and most intuitive way to approach transformation.

4.2 Appropriate Labels: Broad, Simple and Brief

Despite the above problems associated with the use of labels to direct attention, these will remain important within investigations of coarser aspects of the mind, and in building up to more refined top-down and bottom-up concentrative states. The Insight Meditation literature accordingly devotes much instruction to the appropriate kinds of label for initialising focused introspective-attention, where it is known as the practice of “mental noting”. What one sees here is the predominance of labels that are (conceptually) broad, simple and brief.

Firstly, labels used to direct attention tend to express maximally broad contents. Those used include: “feeling” or “sitting” in mindfulness of the body; “seeing”, “touching” or “hearing” in mindfulness of the “sense-bases” (sensory fields); and “breathing”, “in”, “out”,

“rising”, “falling” to hold the mind upon the breath (Sujîva, 2000, pp. 28-41; see also Gethin, 2015, pp. 28-30). More specific labels are sometimes used *after* this initial stage of attentional application (e.g. one can register thoughts as “remembering”, “planning”), though these tend to be used to help *register* what has emerged, rather than to *direct* the mind; their function is to ingrain conceptual insights, rather than make insights possible. One might think of this difference in terms of canonical distinctions between *sati* (mindfulness) and *sampajañña* (clear comprehension). Nyanaponika (1988, p. 46) notes that the former concerns the attentional *holding* of the object, while ‘[c]lear comprehension is the right knowledge (*Ñāṇa*) or wisdom (*Paññā*) based upon right attentiveness (*sati*).’ (see also Dreyfus, 2011, pp. 49-50)

Relatedly, practitioners are cautioned to avoid specificity in their labelling when that endangers artificially preserving some aspect of experience. Sujîva (2000), discussing attention to the sensations of breathing, states that even ‘the “rising” or “falling” [is] not constant and it may disappear while one is watching it’ (p. 31). Warning against the unnatural preservation of these aspects of experience, he notes:

The labelling can be said to point to a window to which we direct our mindfulness. It helps us hold our mind to the meditation object and thereby, develops the concentration which sees, through mindfulness and bare attention, the realities that occur there. Here, we cannot choose what we see; we only direct our mindfulness to the “window” and observe whatever arises (p. 30)

As Davis and Thompson (2015) note, such labelling is also used more informally through the day during Vipassanā meditation retreats; they suggest that the aim here is, again, not primarily to describe, but to *hold* the mind to the present (p. 51).

In either situation, mental noting will likewise avoid complex concepts likely to prompt further reflection by the student and will tend towards brevity, rarely extending beyond two

words and normally just one. This helps avoid what's known as "conceptual proliferation" (*papañca*) in Early Buddhism – the snowballing of thoughts and evaluations about experience (Ñānananda, 1971/1997, pp. 4-5). More complex and extended notes are likely to evolve from *directive* aids to *descriptive* or *discursive* thought. They will take a direction of their own, pulling one away from experience or leading to affective reactions and thus "impure" forms of concentration.

With these points noted, one can establish some clearer guidelines for the use of labels within first-person scientific methods. First, focused introspective-attention should be initialised with the help of labels that are simple and brief. This will be critical in supporting what's referenced in the phenomenological literature as a 'receptive openness' to experience. Colombetti (2014) elaborates this as 'a *passive-observational stance* towards one's mental life', which is '[not] inquisitive, judgemental [or] actively discriminating' (p. 149, see also p. 156; Thompson, Lutz and Cosmelli, 2005, p. 70). Simple and brief labels will support this stance, forestalling discursive thought that can stir the mind into distraction and impure concentration.

Second, one should err in favour of broad and neutral labels, relative to one's introspective targets. This will help (at the subpersonal level) to mobilise attentional control-sets that exert maximally homogenous increases in alertness and (at the personal level) to minimise the possibility of conceptual distortions of experience or the introduction of novel content. Relatedly, the pedagogical literature suggests that one should be keenly aware of the move from directive labelling to descriptive labelling, which will be more specific and more likely to 'fulfil' the experience in a certain way, stabilising particular and more specific properties in a more artificial manner (see Depraz, Varela and Vermersch, 2003, p. 71; Petitmengin, 2007, p. 74). This distinction between directive and descriptive language use is

rarely made in the broader introspective literature, but it can help to minimise dangers introduced by rushing into specific, complex and extended description too quickly.

Note though that these are all strategies for the *minimisation* rather than elimination of conceptually-induced novelties in experience. For instance, the enriching of an experience through top-down processes, bringing detail and precision, may lead to increased specificity or determinacy of that experience, *however broad* the label employed. A vague and diffuse sense of enjoyment might be transformed into a new and particular kind of enjoyment with attention (see Colombetti, 2009) or reveal subtler bodily stimuli than were apparent before.³⁵ Nonetheless, though we can only seek to *minimise* rather than eliminate such transformations, this causes no intractable problems for those first-person methods open to such issues. One reason for this is that problematic novelties can be dealt with during the *latter stages* of first-person methods, if we know of their likely occurrence. Though this point might not feature prominently in debates within the Buddhist traditions, it forms a central concern of other traditions of phenomenological inquiry, which introduce additional procedures after individual introspective reports have been made, and which are important partners to the use of meditation-trained skills in the study of the mind. This is a point not stressed often enough in response to distortion concerns, which we can unpack with a specific example, and which can help distinguish the second broad way that transformation can be handled.

Take attempts to establish what uniquely identifies the lived experience of desire – some general characteristic of the mind, common across many persons. Good introspective

³⁵ Nanay (2009) and Stazicker (2011) have proposed determinacy increases as a necessary feature of visual attention. They employ the determinable/determinate distinction to make this point, where determinates of any determinable are conceived as *more precise way of being that determinable*, as scarlet is a determinate of red. From this perspective, top-down attention will always lead to increases in specificity of visual experience, which we might think will apply in the case of focused introspective-attention upon visual experience too. For a counter-argument, see Wu (2014, p. 125).

methods aimed at this will always involve the employment of *additional processes* after token descriptions have been solicited from individual subjects. These processes are employed precisely to correct for some of the problematic idiosyncrasies that will be formed in the individual introspective act. The phenomenological tradition stresses the importance of “eidetic reduction” and “intersubjective validation” here. In the eidetic reduction, the subject engages in a form of imaginative variation, to isolate what is essential to experiences of a certain type from what is inessential or ephemeral. Intersubjective validation then compares, contrasts and corroborates these results across many individuals. These then are two attempts to *filter out* some of the idiosyncrasies that might be introduced at the individual level of self-observation by acts of attention to experience, and which can permeate individual descriptions. The problem of distortion seems so large often on account of forgetting this important step in first-person methods.

This enables us to distinguish the second broad way that transformation can be handled. We can (2) mobilise *beneficial* transformations, whilst *minimising* detrimental transformations and then (in latter stages) *filtering out* the novelties infecting descriptions of token experiences. In this sense then, it doesn’t matter so much whether descriptions about particular instances of experience are in fact misleading or erroneous with respect to particular features of experience. Cognitive science does not care in large part about the experiences of individuals, it is concerned with general truths—the “invariant” features of experience—that can be approached by correcting for more distortive transformations in the process of their revelation.

4.3 Piecemeal Progress

As noted in §3.3, introspective methods need to be sensitive to the fact that some experiences will be complex targets, with numerous moving components, and open to variable probing. An

implication of this is that it will be important to gauge how to consistently target these different aspects with focused introspective-attention, so as to avoid theoretical overemphasis upon some, or the conflation between reports of different kinds. Turning to the Insight Meditation literature is again beneficial here, particularly in its approach to investigating the *emotions*.

Emotions are primary means by which the Hindrances manifest. Given that these are chief concerns of Buddhist practitioners, one would expect the Buddhist tradition to have much to say about the appropriate means to investigate emotions. And they do. A first thing of note in the Insight literature is the recurrent suggestion that the intentional “act” of emotions is best investigated through the body. Contemporary teacher Thiradhammo (2014) well exemplifies this in the following practice instructions:

Trying to be aware of [the Hindrances] at the very beginning of practice, it is easy to be pulled into them [i.e. into the story about their object] or caught in doubt about them: ‘What am I actually looking at?’ However, if you have a very good grounding in *awareness of the body*, you can always relate back to it, or cross-reference it: ‘What is the condition of the body? Is it lacking in energy? Or has it got too much energy?’ Through the body you are able to recognize: ‘Oh, there is lethargy’, or ‘there is restlessness’. Thus you can generate greater awareness of the Hindrances through *awareness of their expression in the condition* of the body (pp. 24-25, emphasis and parentheses added).

This proposal can be unpacked by noting some features of contemporary emotion theory. There’s a consensus amongst emotion theorists that bodily sensations play a pivotal role in the experience of affect and emotion (Damasio, 1999, 2003; Pollatos and Schandry, 2008; Prinz, 2004; Seth, 2013; Whiting, 2011). Different emotions are known to correlate with different kinds of bodily activity (Nummenmaa et al., 2014), and it’s thought that emotions constitutively involve an awareness of some such activity (Colombetti, 2014; Prinz, 2004; Whiting, 2011). From this theoretical standpoint, the prioritisation of interoceptive data (i.e.

data about bodily features, such as the muscles, skin and organs) through top-down attention to bodily experience (or the body itself) should help to illuminate certain characteristics of emotion – namely the intentional *act* of emotions, wherein we feel towards particular things.

Expressing this point, Davis and Thompson (2015) suggest that mindfulness-practice ‘may result in increased awareness of one’s emotional state by virtue of increased conscious experience of interoceptive changes involved in one’s physiological reactions’ (p. 55; see also Colombetti, 2011, p. 302; Farb, Segal and Anderson, 2013; Hölzel et al., 2011; Sze et al., 2010; Teper, Segal and Inzlicht, 2013). This increased awareness might concern the simple occurrence of some undefined form of emotional reactivity or it might concern the type of emotion taking place. The authors actually attribute this possibility to longer-term increases in bodily awareness available through mindfulness-meditation, which I shall review in §4.5. Yet, such increases can also be induced on demand using focused introspective-attention with body-centric labels (e.g. “body”, “feeling” or “abdomen”), as in Thiradhammo’s (2014) above account. This will both accentuate the bodily properties of the emotional act and isolate them from distractors, to our epistemic advantage.

Such a strategy exemplifies a more general benefit of labelling noted by Nyanaponika (2015). He speaks of labels as means of ‘singling out the separate strands forming [the] intricate tissues [of experience]’ (p. 76) such that they can be better investigated and catalogued. In the case of emotion, focused introspective-attention to bodily experience, through appropriate labels, works to accentuate and isolate the bodily expression of the intentional act from the emotion’s more cognitive elements (or *that which* the emotion is directed towards). This way, we can learn certain distinctive things about the different types of emotional experience, irrespective of their objects. For instance, we can reveal the precise bodily structures involved in any particular emotional type, as well as the spectrum of sub-varieties of those emotions that

are possible. For instance, attention to the body during episodes of anger can reveal kinaesthetic sensations preparing the arms and hands for movement as distinctive features of this emotional type (see Nummenmaa et al., 2014).

This can be done whilst acknowledging that we don't learn *everything* about emotion through attention to the body and that this act has additional and more distortive effects upon our experience, which will need accounting for. For instance, Buddhist contemplatives are well aware that directing the mind to the intentional act of emotion will sap some of its intensity. Mirroring Brentano, Nyanaponika (1988) notes that direction to a bodily state of anger about a disturbing noise works to dissolve that anger by 'diverting attention' away from the noise that fuels the anger (p. 72). Indeed, this forms a central strategy for suppressing the Hindrances, an understanding of which is held important for the practitioner to develop (Thiradhammo, 2014, p. 27). Yet, these facts do not mean that one can learn nothing about the intentional act and the larger emotional structure here. In the case of emotion, attention to bodily experience does not immediately overturn the comportment of the body, and thus the emotion it expresses. If it did, the soteriological project would be an easy one!

Furthermore, some of the seeming "losses" involved in any individual probe are merely temporary and recoverable. Note how, often, reflective awareness of the simple fact that one is emotionally reacting *at all* is a condition for revealing the intentional objects of emotion. It can prompt a search for what one is reacting to, where one directs attention to "mental objects".³⁶ For example, while cycling to work, there may be many fleeting emotional states occurring in my background experience as I focus upon the road. Many of these will be outside of my reflective awareness, and at the end of the bike-ride I may be unable to report either their

³⁶ This is practiced in "mindfulness of the objects of clinging" where one gets more of a sense of the things one is pre-occupied with, rather than the manner of being pre-occupied itself. See Thiradhammo (2014, pp. 44-45) for a good personal account of how this might proceed.

occurrence or their intentional objects. If, however, I can gain reflective awareness of an emotional episode by tapping into the body at regular intervals during my bike-ride, this can set off a recognitional alarm bell (e.g. “I’m feeling worried”), which can prompt me to seek out the intentional object of that experience. This is possible given that the act of diverting attention to the body (i.e. the feelings of worry) rarely completely disarms my emotion of its intentional object or prevents it from returning and rumbling around in my mind. The very recalcitrance of the Hindrances is what armours top-down investigations of experience against some of these criticisms.

Altogether, we see that the meditative assembly of introspective knowledge is self-consciously presented as proceeding in a gradual, accumulative manner. It is sensitive to more problematic transformations induced by any individual step in this process and works around these, over time, to get a sense of the “lived experience” of the mediator. All things considered then, from this review of contemplative approaches to the investigation of emotion, three points emerge as “take home” messages about the appropriate ways to use focused introspective-attention.

Firstly, the literature helpfully emphasises that top-down introspective inquiry is often an *extended* and *piecemeal* affair (i.e. when one’s introspective targets are complex). There is no reason to stipulate the engagement of a single introspective act able to simultaneously illuminate everything perfectly clearly.³⁷ The long historical engagement of Buddhist contemplatives with experience has led to an understanding of its complexities, and a complementary *series* of different acts by which those complexities can be illuminated. So long as one is sensitive to the different effects of each probe, one can proceed safely.

³⁷ Petitmengin (2006, pp. 237-238) also makes this point, but of what she calls ‘retrospective evocation’, that is, re-enacting different aspects of experience in memory.

Secondly, this helps us identify two further ways that transformation can be handled in the investigation of experience. Already we have seen that it is possible to exploit positive transformation whilst (1) *eliminating* distortion during the act of probing or (2) *filtering out* distortive idiosyncrasies that have come to infect reports in later stages. In the above examination of emotion, two more possibilities emerge. We can instruct introspectors to conduct a series of probes, asking them in each case to (3) *ignore* those aspects of experience that we know will be distorted by probing at the point of report (for example the intensity and “directed” nature of emotion) and to direct themselves only to those properties likely to be most reflective of lived experience (e.g. the broad bodily components of emotional kinds). This way, so long as one proceeds cognisant of the full set of phenomenological effects of focused introspective-attention, one can avoid being misled by them. For, just as turning on a light to investigate which of our relatives is occupying a dark room *changes* them, in the sense that it will raise their temperature ever so slightly, it does not change them with respect to the property we are interested in.³⁸ This approach is distinguished from *filtering out* (method (2)) in that it forestalls such distortions from even contributing to judgements about lived experience in the first place.

In contrast, it is also possible to (4) *derive knowledge* about the properties of lived experience from distorted properties themselves when employing top-down introspective attention. In the above analogy, we could also derive some knowledge of the person’s temperature so long as we know either roughly how the light affects them (i.e. that it *raises* their temperature) or indeed how strongly (i.e. *how much* it raises their temperature). In the introspective case, knowing that the strength of an emotion is likely to decrease (*ceteris paribus*) upon bringing it into attention, we can infer certain things about the pre-reflective

³⁸ Thanks to Scott Sturgeon (personal correspondence) for this analogy.

intensity of this emotion too – that it was *no more than* intensity-level X. In this way, we can arrive upon an increasingly accurate and complete understanding of our own lived experiences, while acknowledging that some such aspects of this picture are approached indirectly.

The third and final take-home point from the Insight Meditation literature’s treatment of emotion is its emphasis upon *the body* as the route into the intentional act of experience (see also Depraz, Varela and Vermersch, 2003, p. 36). I’ve illustrated how this proceeds in the case of emotion, though we might also use it to inform discussions about perception. Earlier, I noted how debates concerning transparency (§3.3) and sense-data (§3.5) might be explicable in terms of different ways of probing experience, rather than direct conceptual “distortions” of the experiences in question. It may be that transparency theorists, for instance, hold the view about perception that they do precisely because they are unaware of the appropriate means to draw attention to the intentional act of perception, while the meditation literature helps unveil the body as key here.

This possibility is further motivated by returning to Thompson’s (2007) comments on perception. Earlier, I noted Thompson’s proposal that perception always presents more than the sensory qualities of the world; it presents also *subjective features* that it is possible to attend to. Thompson describes this in terms of the “self-presentation” involved in perception, which I understand here to mean the presentation of *aspects of the self*, as opposed to the world.³⁹ Drawing upon Husserl and Merleau-Ponty, he elaborates that self-presentation is made up of a ‘pre-reflective bodily self-consciousness’ (p. 265). For example, the presentation of the table’s

³⁹ This should be understood differently to other uses of the term “self-presentation” (or “self-representation”) which take the “self-(re)presenting” quality of perception to designate the implicit awareness *that* one is conscious of something, accompanying at least some of our experiences (see Rosenthal, 2002, p. 409; Zahavi, 2014, p. 15). I say more on this feature of the mind in §4.5. See Coseru (2009) for an understanding of self-presentation closer to my own and Thompson’s (2007) interpretation, which Coseru attributes to Buddhist thinker Dignāga.

hardness is said to involve more than just registration of the table's rigidity, or the pressure it exerts upon me; it also involves a subtle awareness of my own corporeal engagement with the table, wherein I actively attempt to put pressure on and manipulate the table itself. From this perspective, awareness of the body is also a constitutive part of perceptual experiences. And it seems reasonable that we might heighten awareness of the intentional act of perceptual experiences by bringing attention to their *bodily features*, analogously to the way emotional experiences can be illuminated.

There is an important dis-analogy though between emotional and perceptual experiences. When attending to the bodily aspects of emotion, the emotion is temporarily sustained in the absence of its intentional object. However, it seems difficult if not impossible to turn away entirely from the objects of perception whilst sustaining those perceptions themselves. Thompson thus cautions that the appropriate means of attending to the intentional act here involves 'not [...] turning our attention away from what that experience is *of* (that is, the intentional object)' (2007, p. 285) but by engaging an *extra* form of attention atop this basic kind.

Thompson offers only a short and speculative account of this dual attentional stance. He describes it to involve 'direct[ing] our attention to the appearance of the object [...] *while vigilantly keeping in mind that appearances are objective correlates of subjective intentional states*' (p. 287). Here then we seem to have a balance of two things in play. One must hold the object in place, whilst "keeping in mind" the corporeal, active and subjective features of perception. Thompson marks such "keeping in mind" as involving a kind of "cognitive attention". This suggests a new kind of top-down attention being employed once concentration upon the object itself has become steady, but one which must somehow accentuate one's bodily

involvement in the perceptual act without biasing against incoming sensory data (at least in any significant sense).

I must admit that this gesture remains obscure to me. It is difficult to understand both the practicalities of its performance, and how it can be made consistent with the models of top-down attentional control reviewed here. Nonetheless, relating these two presents an interesting avenue for future investigation. And I suggest that further engagement with the contemplative literature might assist here. Manuals devoted to *kasina* practice, for instance, involve detailed accounts of the steps by which one turns attention from the intentional object of perception to the intentional act itself (see Wallace, 1999). And further engagement with these texts may better unpack the nature and practicalities of such attentional gestures.

To summarise this third take-home point then, we can say that the Insight Meditation literature reinforces the idea that proficiency in heightening bodily awareness may be an important condition for proficiency in introspective endeavours and gives useful models for moving between act and object.

4.4 Concentration as a Preliminary Factor

Despite all the above qualifications about the appropriate use of focused introspective-attention, the Insight Meditation literature also emphasises its ultimate limitations within introspective endeavours. Some kinds of insight are simply unsuited to, and even hampered by, the deployment top-down skills, requiring other kinds of attention. This reflects a broader tendency within the literature to regard concentration as an ultimately *preliminary* factor, which is to eventually be de-emphasised in insight practices.

Focused introspective-attention's insufficiency lies in the narrowness of the experiences it promotes. Its tendency to inhibit things outside its scope means that it can

remove surrounding components of experience that retain the shape and dynamism of that which is attended to. Similarly, inhibitory effects mean that such attention can ultimately obscure broader patterns of activity important for understanding the *causes and conditions* of suffering. The wider aim of insight practice is to develop and deploy “wise attention” (*yoniso manasikāra*) that discerns events in terms of the Four Noble Truths. This entails understanding more than just when craving and suffering are present, and their intrinsic properties; it demands familiarity with: how they came about; what makes them disappear; and what keeps them from returning (Thiradhammo, 2014, p. 26). Discernment of these patterns is best served by a kind of attention that is sensitive to such broad and diverse features of mental activity.

For these reasons Sujîva (2000) notes that the development of Right Concentration—that is, appropriate concentration for insight practice—does not mean pushing for extreme concentration at the expense of all else. Rather, it means developing a certain degree of proficiency in top-down attentional skill and then *de-emphasising* holding the mind in place, in favour of an interest or curiosity about the place one has taken up (pp. 145-147). Put differently, one employs top-down attention to first bring the mind to rest on a certain point. One then deploys a less selective attention that can be sensitive to whatever emerges within and *around* one’s target, yielding a broadening in the range of experience (see Thompson 2015, p. 52).

Here, the subject is moving to a more distinctively *bottom-up* form of introspective attention, whose targets are determined (at least, most significantly) by the intrinsic features of experience. Nonetheless, it seems that, here, the practitioner relies upon some of the ‘natural steadiness’ of higher-level concentrative states to hold the mind around a particular point. There is likely some residual, “hangover” effect of the self-organizing and self-sustaining properties of concentrative states (see §4.1) temporarily retained in this new state, allowing

one's attention to centre upon and around a certain point, despite being open to broader patterns of activity. In this way, it will be possible to both *direct* one's focus whilst being sensitive to broader features of experience. This step approximates the "loosening" of attention that Petitmengin and Bitbol (2009) identify as an important stage in introspective investigations, to be entered after attention has been deliberately re-directed. They note:

[u]nlike Focused-Attention [i.e. focused introspective-attention], which is narrow, concentrated on a particular content, this attention is panoramic, peripheral, open on a vast area. This diffuse attention is however very fine, and sensitive to the most subtle changes. Several people have described this openness to us as a subtle shift of the area usually perceived as the centre of attention towards the back of the skull, or from the head down into the body (p. 378).⁴⁰

In recommending this eventual move to bottom-up forms of introspective attention, the Insight literature reinforces the idea that first-person methods can't rely solely on focused introspective-attention; a skill here needs to be paired with other capacities. While the ability to generate refined forms of concentration (including top-down and possible bottom-up kinds) is important, top-down, selective approaches to the investigation of experience should eventually be de-emphasised in favour of a more natural, unbiased and open curiosity, that better retains the mind's breadth and dynamism. A skill in focused introspective-attention is therefore a *condition* for investigation, whose exercise is eventually to be overtaken.

In this next stage of Insight practice, subjects rely upon a more general and involuntary sensitivity to experience, trained in Open-Monitoring practices (see §1.1). This involuntary

⁴⁰ This deliberate "loosening" of attention also allows pre-reflective aspects of bodily experience to become more salient, despite their not being strictly focused, potentially offering another route into the intentional act of perception – when attention is not exclusively focused upon the perceptual object, it makes room for stimuli underpinning the perceptual act to become more conscious (see Petitmengin and Petitmengin, 2009, pp. 377-381).

sensitivity is usually illustrated by describing cases of its absence. A much referenced and relatable example is Armstrong’s truck driver. Here, Armstrong (1980, pp. 59-60) relates the story of a long-distance truck driver who, at some point in their travels, “comes to”, realising that they have for the past while been driving without being aware of what they were doing, perhaps lost in some other thoughts. Here, Armstrong thinks we have sufficient reason to believe the driver had perceptual consciousness of stimuli required to drive; nonetheless, they were lacking some awareness *of* this consciousness – which Armstrong calls ‘introspective consciousness’. It is this introspective consciousness that returns when the driver “comes to”.

In the scientific literature, Thompson (2015, p. 52) notes that this phenomenon is captured under the concept of *meta-awareness*, which he glosses as ‘awareness of awareness’. Though meta-awareness covers a varied and often conflicting set of capacities within cognitive psychology, we can conceive it rather broadly here in terms of one’s automatic epistemic sensitivity to the contents of one’s mind – something that needn’t be prompted by deliberately “turning inwards”.⁴¹ And it is meta-awareness that comes to the fore in insight meditation once concentration has been built up to a suitable degree. Though, it must be emphasised that one first needs a *degree of proficiency* in top-down skills before this stage can be reached. Without skills in focused introspective-attention, the mind quickly runs off into different territory

⁴¹ Somewhat confusingly, “meta-awareness” is sometimes used to describe only those occasions where one gains awareness of mental contents by such *deliberate turns inwards* (e.g. Chin and Schooler, 2009). This makes its use rather awkward here, for, though such deliberate introspective gestures are performed at the outset of Open-Monitoring meditation (i.e. one turns towards the whole of the experiential field), the practice aims to make one’s sensitivity to experience more automatic and passive, rather than something that needs to be actively engaged, with this automatic capacity the primary target for improvement here (see Lutz et al. 2015, p. 640). An alternative term for this might be “inner perception”, which Spener (2018) notes was coined by Brentano to describe the ‘fairly automatic and passive awareness one has of one’s own conscious experience, as one goes along in the world in an ordinary manner’ (see Spener, 2018, p. 159) and set in deliberate contrast to “self-observation”, where attention is actively and deliberately turned towards the mental (i.e. focused introspective-attention). For consistency with other meditation research though, I continue with the term “meta-awareness” here.

entirely. There will be no possibility to build up to more naturally stable states, nor therefore to direct one's inquiry at particular targets, leaving investigation a superficial and haphazard affair. So understood, focused introspective-attention is best conceived as an essential *preliminary* skill.

Yet there is an even more preliminary role played by this faculty, whose illumination shall be my final aim here. §4.5 will show how top focused introspective-attention can yield an extra type of epistemically-beneficial transformation in the longer-term. And this very transformation actually serves to improve the capacity turned to—meta-awareness—once focused introspective-attention has done its work.

4.5 Ground-Clearing

A final important characteristic of the Insight Meditation literature is its emphasis upon the preliminary “ground-clearing” function of top-down attentional control. In contemporary references to meditation, it is rarely emphasised that top-down attentional skills are used also in contemplative programmes to set the appropriate conditions *within which* to investigate experience, rather than merely being ways to probe experience itself. As Gethin (1998) notes, Buddhist meditation regimes tend to be framed largely as a two-stage procedure:

This then is the basic theory of Buddhist meditation stated in the terms of the oldest texts. While later schools and traditions may change and adapt the terminology used, while they may elaborate the stages and techniques in a number of different ways, while they may give distinctive technical accounts of the content of the knowledge gained [...] the basic principle for the most part holds good: one stills and clears the mind and *then* turns it towards investigation and insight (Gethin, 1998, p. 176, emphasis added)

Gethin makes it clear here that investigation is preceded by a ground-clearing procedure of “stilling” and “clearing”, which usually occurs through concentrative (Focused-Attention) practices. Even in the contemporary Insight Meditation literature, where investigation and insight are prioritised, there is emphasis upon the supportive benefits of prior grounding in concentrative practices and the according development of the *jhānas* (see Sujîva, 2000, pp. 228-230; Nyanaponika, 1988, p. 62). Sujîva marks the ability to enter states of high concentration as an advantage (p. 230). While Nyanaponika (1988) notes that concentrative approaches to mindfulness of breathing can be used as ‘a prelude to other exercises’ (p. 62). Here top-down attention is something used not *within* the investigation, but as an important prior. How then does this yield epistemic benefits?

The broad proposal here is that repeated returning of the mind to a single object pacifies the mind of its habitual busyness; it generates states of relative quietude or calm that give traditional Focused-Attention practices their framing as “tranquillity” (*samatha*) practices (Williams and Tribe, 2003, pp. 81-82; Gethin, 2004, p. 207). Importantly, this quietude is not something that immediately disappears once Focused-Attention practice is left behind. It seeps over into the post-concentrated state, producing a general or broad-scoped state of quietude.

This post-concentrated quietude can be understood by returning to earlier talk of *clarity*. In one of its Buddhist guises, clarity designates the emergence of aspects of experience without competition for attention. This can occur in a narrow sense, as when what one concentrates on emerges in relative isolation from distractors. But it can also occur in a broader fashion, where there is a more *general sparseness* to the mind. Though the mind is no longer “one-pointed” here, being populated by a broader variety of mental contents, there is nonetheless less going on in general, making for less competition amongst possible targets of introspection. And note that this supports the flip-side of mental clarity – vividness. With less mental elaboration, the

mind's resources are distributed over a smaller range of phenomena, making one's experience of those phenomena richer.

With mental quietude achieved, mental contents emerging into the stream of consciousness will be easier to discern, both in the deliberate, top-down introspective probing of experience, and through more automatic meta-awareness capacities. On the latter note, Markovic and Thompson (2016) outline how 'meta-awareness requires maintaining openness to experience and overriding one's habitual tendency for conceptual elaboration' (p. 92). This makes focused introspective-attention a key support for meta-awareness, for it can create a state where there is an ongoing disposition towards less elaborate processing of mental contents, allowing meta-awareness capacities to function more effectively. So understood, mental quietude does not *improve* meta-awareness capacities themselves; rather, it is the environment in which they function best.

Davis and Thompson (2015) link this idea to the earlier models of attention. They marshal evidence to suggest that raising phasic alertness—the region-specific sensitivity to stimuli, underpinning the local accentuation and isolation in acts of top-down attention—also raises *tonic alertness* in the longer term (Jha, Krompinger, and Baime, 2007; Robertson et al., 1998). Tonic alertness, recall, designates a person's broader degree of sensitivity to stimuli across the entire spectrum of sensory (including interoceptive) modalities. And we'd expect increases here to yield phenomenological changes akin to the broader-scale clarity characterising post-concentrated states. Less elaborate kinds of mental activity in the sparse post-concentrated landscape—especially, fewer *conceptual* dealings with experience (see Thompson, 2015, pp. 51-52)—allow for more cognitive resources to be devoted to remaining stimuli across the range of sensory (including interoceptive) fields, thereby accentuating what is left in experience, for the reasons reviewed in §2.4.3. Moreover, the phenomenological

effects of tonic-alertness increases do not suffer from many of the problems associated with focused introspective-attention (i.e. those underpinned by *phasic* alertness increases), for they are more uniformly distributed across experience as a whole.

Importantly, these increases in tonic alertness are not simply residual effects that persist temporarily post-practice. They gradually come to permeate the everyday life of the meditator, becoming a more stable “trait” (Lippelt, Hommel and Colzato, 2014, p. 3; Kilken et al., 2015). This means that their supports to meta-awareness can extend here too. Davis and Thompson (2015) note that the reduction in conceptual elaboration, and consequent accentuation of the features of bodily experience (through raising tonic-alertness), can support awareness of body-based emotional reactivity in everyday life (p. 53). This happens without the need to go looking for such experiences, as one would when using focused introspective-attention as a probe, but thanks to its longer-term effects on the functioning of meta-awareness.

Classical Buddhist texts explain this transition from raised *phasic* to tonic alertness in terms of the “karmic arc”. By removing some of the mental agitations (or karmic “seeds”) in concentration practice, say particular Hindrances like “sensual desire”, one pre-empts their future effects (their karmic “fruits”). By framing things as necessities, as the Hindrances do, individuals push themselves into discursive planning to attain those things (see Nyanaponika 2015, pp. 92-4). And this proliferation of discursive activity not only co-opts cognitive resources that could be spent in awareness, it promotes further Hindrances that do likewise, given the potential (and likelihood) for such plans to be frustrated. As the mind is less plagued by agitations then, Nyanaponika notes that ‘the centrifugal forces of mind, making for mental

distraction, will peter out’ (p. 95), creating a state of broad quietude for introspective investigations to flourish.⁴²

We can summarise the general point here by saying that top-down skills help to create an appropriate *environment* for Insight practices to occur – whether this is done deliberately just prior to investigation or more organically over the long-term. Experiences arising (or even deliberately precipitated) in this environment won’t be *reacted to* in the habitual way, creating a sparse, yet well-punctuated landscape, that is more amenable to description. In contemporary terminology, this broad-scale clarity can be distinguished as an “operational condition” for introspection – a condition under which introspective judgements tend to come out good or accurate (Goldman, 2004, p. 14; Spener, 2015, p. 303, 316). This is not too far from what is already suggested in some recent treatments of introspection. Spener (2015), for example, identifies a common set of conditions contributing to introspective accuracy, including the subject being ‘alert, not distracted, not under the influence of drugs’ (p. 316). Buddhist contemplative programmes merely systematise methods for producing conditions of alertness and non-distraction at the broadest-scale (in the clarity that spills over from Focused-Attention practice), so that it can be exploited in introspective endeavours. Thus, renowned Insight teacher Ajahn Chah has remarked that ‘the deeper the calm, the deeper the insight’ (cited in Thiradhammo, 2014, p. 42)

One might object that conducting introspective methods within this environment brings its own dangers of unrepresentativeness, offering a variant of the ‘stilling the stream’ objection from §3.2. Recall Brentano’s claim that deliberate attention to emotions serves to “deaden” them. We might worry that the prospects of studying “real” anger, say, (of the intense and

⁴² For a more elaborate account of the contours of the karmic arc, as detailed in early Buddhist Abhidhamma texts, see Lusthaus (2003, chpts. 9-10).

raging kind that actually drives our behaviours) is in fact *worsened* by the state of quietude traditionally cultivated as a precursor to Insight practice. This is an important objection, and we must be careful to acknowledge this and similar differences between the naïve introspector and the trained meditator. Nonetheless, so long as we are again sensitive to differences between the two (which Froese et al. (2011) remind us that meditators are well aware of (p. 265) and are used to incentivise practice), we will be able to exploit “revelatory” differences, whilst either minimising or accounting for more distortive differences in the four ways already outlined.⁴³

We can also re-impress the difficulty of the soteriological project here. Practitioners will testify that meditative practice does not rapidly banish ordinary kinds of emotional reactivity from their existence. There may well be *less* instances of these, but only the most idealistic conception of the contemplative project will posit their complete disappearance. The very recalcitrance of the Hindrances ensures that important features of experience will therefore be shared (e.g. the emotional kinds noted in §4.3) and interrogable through the above methods. Yes, there will be increases in granularity here. Yes, there will be different levels of intensity when it comes to emotional experience. But one can be sensitive to these. And more productively, one can then focus on the ways that meditative quietude might “prime” the mind to be more susceptible to the shorter-term probing of focused introspective-attention, rather than those ways that it takes the mind further away from that of the non-meditator.

In conclusion then, two lessons can be drawn from the Insight meditation literature concerning the “ground-clearing” benefits of focused introspective-attention. On the one hand, practices devoted to ground-clearing can be employed immediately prior to introspective

⁴³ It’s also worth mentioning that meditators’ journeys through to a different “default state” are likely to make them more familiar with the things that they have relinquished.

investigations to induce appropriate operational conditions. On the other hand, they can be used as standalone practices to exert longer-term trait changes to tonic alertness levels, and thereby the clarity of experience, which can be epistemically exploited without effort. This further secures the importance of top-down attentional skill within the introspectors toolkit.

5 Conclusions and Future Directions

In the above, I've sought to persuade that the transformative character of meditation practice can be an advantage to introspective methods within science. I've shown, in §2, that Focused-Attention meditations train a skill in focused introspective-attention, itself capable of inducing *epistemically beneficial* transformations to experience. Noting the many dangers surrounding such skill in §3, I've shown in §4 that a turn to the pedagogical literature on meditation, paired with proper scrutiny of models of top-down attentional control, reveals how these difficulties can be dealt with. In this way, I've argued that focused introspective-attention can play an important role in unveiling invariant features of human experience, particularly those of the emotions and affective states. In this way, it is an important component in the investigative repertoire of the proficient introspector. Moreover, it does this not in spite, but *in virtue* of its transformative qualities.

Undertaking the above, I've thus supported and clarified the broader Buddhist posit that transformation of the mind can actually be exploited, rather than avoided, in the mind's investigation, helping to break the popular spell that binds change to distortion. On this point, we can agree with Boyde Henry Bode, who noted over a century ago that '[t]he proper test for a sound introspection is not the degree of change which it introduces, but *the kind*.' (Bode, 1913, p. 88). So long as we remain sensitive to the *kind* of changes induced then by the skills we train, we will be able to use meditative methods effectively. In this work, I've identified the

various kinds of change induced by top-down introspective attention, and distinguished four ways that such changes can be handled for epistemic benefit:

- (1) Induce beneficial kinds and *eliminate* distortive kinds
- (2) Induce beneficial kinds and *filter out* the effects of distortive kinds upon reports
- (3) Induce beneficial kinds and *ignore* distortive kinds at the point of report
- (4) Induce beneficial kinds and *derive knowledge* from distortive kinds

These points also motivate some broad comments about the possible future directions of consciousness science, which I shall end on.

Scientific investigations of consciousness have typically sought to minimise change when it comes to the formation of introspective judgements and have thereby disparaged methods that seem open to it. This has meant a reliance on relatively off-the-cuff and often retrospective reports that attempt to minimise potential distortions by using inattentive and relatively untrained subjects in distant and cautious methods. Yet this bias against careful and attentive introspection has also led to a situation where the raw materials of our science – the introspective reports that shape our explanatory targets – are criticised on account of being both massively unreliable and inconsistent on one hand (Schwitzgebel, 2008), and hopelessly lacking in detail on the other (Chalmers, 1999). Chalmers notes that we've tended to capture only 'gross and simple features of conscious experience' with our descriptions or have ended up '[employing] language which is obviously course-grained and imprecise' (p. 10), talking relatively uninformatively in terms of 'an experience *of red*, [or] *of a horizontal line*' (emphasis added). Especially lacking is detail about the phenomenal character of such experience, which

can express precisely *what* “it is like” to have an experience of anger, of red, of pain, or of even of thought.

In carving out an epistemic role for experiential transformation, we see instead that first-person methods need not be restricted to what one might call “preservational” types – those that aim to keep experience wholly intact. And the Buddhist traditions provide excellent models for these alternative transformational approaches. Of course, the refinement of introspective methods using Buddhist insights should also acknowledge that the Buddhist path is primarily a *soteriological* rather than epistemological one. When push comes to shove it will favour transformation over and above knowledge (i.e. even if transformations are distortive in nature). This means one must remain cautious when learning from contemplative approaches. Nonetheless, one should note that the path to reducing suffering in the Buddhist tradition is explicitly said to depend upon properly grasping the nature of things, articulated herein as “seeing and knowing how things really are” (*yathābhūtañāṇadassana*) (Davis and Thompson, 2015, p. 43, 56). This means that the soteriological and the epistemological elements of meditation practice intersect in a way that science can use to its advantage.

[this page is left intentionally blank]

Introspective Training: A Broader Path?

ABSTRACT: This paper critiques contemporary proposals to employ Buddhist meditation practices within cognitive science as forms of introspective training. I draw out under-appreciated complexities in the Buddhist path to introspective proficiency and reflect upon the bearing these might have for future incorporation of meditation into scientific research programmes. §1 delineates existing suggestions to utilise meditation within science. §2 then casts doubt on these proposals by invoking recent worries over the utility of meditation that has been “stripped for export” and practised in isolation from traditional ethical and philosophical frameworks. I illustrate a prevailing consensus in scholarly, pedagogical and psychotherapeutic circles that the therapeutic value of meditation practice suffers when undertaken without these traditional supports. I demonstrate that a parallel problem emerges for meditation’s introspective value and thereby its utility for cognitive science. §3 surveys this broader web of supportive practices, fleshing out a “broader path” to introspective skill within Buddhism that is much more demanding than is often credited. §4 then crystallises a dilemma that is revealed for proponents of meditation in science, and highlights a methodological choice needing to be made between promoting broad or narrow paths of training, each of which affords science different kinds of benefit. I comment on how to make this choice effectively. In §5, I conclude with some comments on the how science might make best use of meditative practice going forwards.

Introduction

Across a wide range of Buddhist traditions, the practice of meditation is held central to the development of introspective skill and self-knowledge. Such attentional training has long been employed to promote understanding of the mind and the individual’s relationship to the world – supporting kinds of “wisdom” that can liberate people from suffering (Dreyfus, 2011; Gethin, 1998, chpt. 7; Nyanaponika, 1988). Unsurprisingly, most contemporary scientific interest has targeted the latter therapeutic potential of these practices (see Keng, Smosku and Robins, 2011; Khoury et al., 2011). Yet, an important line of research has begun to unpack the broader advantages of the introspective skill developed in meditation, traditionally said to underpin such therapeutic functions. It is this broader issue that I address here.

In particular, this paper shall critique the proposal that meditation’s introspective benefits—its capacity to promote intimacy, at the individual level, with one’s mental life—can allow it to play an important methodological role in the sciences of mind. Rather than treating meditation as an interesting *object* of study, many suggest that one can employ meditation as a *tool* to direct and improve scientific inquiry itself (e.g. Depraz, Varela and Vermersch, 2003; Thompson, 2009, 2015; Varela, Thompson and Rosch, 1991/2017; Wallace, 1999, 2009).

Such proposals arise in a cognitive-scientific climate increasingly friendly to “subjective” or “introspective” methods of data production – methods prioritising that special form of first-person access where we come to make judgements about our minds “from the inside” (Spener, 2015, p. 300).¹ There is growing acknowledgement that such methods serve, at minimum, an important “target-setting” role in the scientific probing of the conscious mind, yielding data that scientists use to guide and constrain their research efforts (e.g. Bitbol and Petitmengin, 2017; Kriegel, 2015, pp. 18-21; Thompson, Lutz and Cosmelli, 2005).² Moreover, it’s thought that the more accurate and detailed one’s characterisation of such targets, the likelier one is to uncover their biological underpinnings, fostering a growing appreciation for especially *skilled* introspection (see Chalmers, 1999; Colombetti, 2014; Thompson, 2015, p. 77). It’s even proposed that researchers can benefit from *personal* introspective skill, affording them an intimacy with subjective experience able to productively inform their own research enterprises (Colombetti, 2014, chpt. 6; Levit Binun and Tarrasch, 2014; Varela, 1996).

This recent focus upon the merits of introspection has catalysed the re-development of systematic methods for bringing experience to report (e.g. Depraz, Varela and Vermersch,

¹ I employ the terms “subjective” and “introspective” interchangeably throughout. In the literature, one also sees talk of “phenomenological” or “first-person” methods. All such terms are held univocal here, avoiding theoretical suppositions about the relationship between the mind and world (c.f. Zahavi, 2007).

² Given my focus here, all further use of the term “introspection” targets only the process of arriving at judgements about *conscious* aspects of the mind.

2003; Froese, Gould and Seth, 2011; Hurlburt, 2009).³ And many such methods call for specific attentional skills to be employed – skills of the kind that Buddhist meditation is thought to develop (Thompson, 2009, 2015, chpt. 2; Colombetti, 2014, chpt 6.). With increased awareness of this link, meditation has been seen as one such means of facilitating more rigorous introspective methods, with central focus upon two categories of practice known in the sciences as “Focused-Attention” and “Open-Monitoring” meditation (Lutz et al., 2008).

Such suggestions are not without their critics. Most prominently, many worry that such training regimes actually transform and *distort* the experience they are supposed to sensitize us to, yielding ungeneralizable data (see Colombetti 2014, pp. 155-158; Kordeš and Markič, 2016, pp. 159-161). I think this worry can be handled though, and that enough ink has already been spilt doing so (see Bitbol and Petitmengin, 2013, Colombetti, 2014, pp. 155-158; Roberts, 2018a; Thompson, Lutz and Cosmelli, 2005, pp. 72-73). In this paper, I thus wish to draw attention to another under-appreciated factor that ought to give further pause for thought over the above proposals.

The critique I offer here takes inspiration from contemporary debates within scholarly and clinical communities over the importance of *contextual supports* within psychotherapeutic applications of mindfulness-meditation – a practice incorporating (but arguably exceeding) the two kinds of meditation noted above.⁴ In recent years, many have stressed the dangers of a common assumption that mindfulness can be plucked from its traditional environment of Buddhist beliefs, practices and rituals, with little harm to its efficacy (see e.g. Kirmayer, 2015; Kordeš and Markič, 2016; Murphy, 2017). There is growing recognition that many desired

³ For accounts of earlier attempts to institutionalise rigorous first-person methods within Introspective Psychology, see Boring (1953, esp. pp. 171-172), Danzinger (1980) and Schwitzgebel (2004).

⁴ It is notoriously difficult to pin down an uncontroversial definition of “mindfulness” and I won’t attempt to do so here. See Williams and Kabat-Zinn (2013) for more on this issue.

therapeutic outcomes of mindfulness are more difficult to secure without presenting these basic attentional techniques in a manner which retains or replicates certain traditional underpinnings that provide the context for practice, including: intellectual engagements; practices for cultivating wholesome mental qualities; and lifestyle realignments to incorporate more ethical behaviours.

I shall suggest in this paper that sensitivity to contextual factors is not merely important for psychotherapeutic communities, but also for those advocating meditative training as a methodological boon to cognitive science. My first aim will be to argue this point by demonstrating that many of the alleged “effects” making meditation therapeutically valuable are also critical to its introspective utility. This has important repercussions. It implicates a *broader path* to introspective skill within the Buddhist traditions, which I shall spend some time unpacking. It likewise entails that realistic assessments of meditation’s value for science require giving attention to these broader supports now featuring in psychotherapeutic discussions. And I shall devote the paper’s final section to such an assessment, revealing important limiting factors upon meditation’s benefits to science, which will need acknowledgment.

Echoing Dreyfus’s (2011) similarly critical motives, these points are intended not as a kind of scholarly ‘got you’ (p. 46), but primarily to highlight routes forward. Increased awareness of traditional contextual supports has already helped to mould expectations within psychotherapy about the therapeutic scope of meditation, generating important dialogue over how best to maximise that scope (see Kabat-Zinn, 2011; Kirmayer, 2015). This awareness can now spread to the cognitive scientific arena, for similar benefit. For instance, sketching the broader path to introspective skill highlights *additional* practices one might consider as methodological supports for research. It also reveals a methodological choice that will be faced

by those pushing for meditation’s employment in science – a choice between promoting (i) a “broader path” of training, incorporating some such additional practices, with more significant individual introspective benefits or (ii) a “narrow path” of training, with less significant individual benefits, but wider appeal, given its less challenging character. Moreover, I shall show that each option has different consequences, fostering skills useful for different kinds of questions in cognitive science.

Here then is an outline of how I shall proceed. §1 first presents Buddhist claims about the introspective benefits of meditation. It then delineates suggestions to employ meditation within science, centring upon the idea that meditation might be used as a form of researcher training. §2 then begins the argumentative work. It unpacks the close relationship between meditation’s introspective and therapeutic functions in order to argue that the contextual difficulties noted in psychotherapeutic circles are also relevant within cognitive-scientific applications of meditation. In §3, I give a brief and unapologetically selective account of the *broader path* to introspective skill that this implicates. While §4 draws out repercussions for the use of meditation in science and suggests plausible routes forward.

1 Meditation and Introspective Training

1.1 The Buddhist Context

The traditions comprising “Buddhism” are numerous and varied. And as with many introductions to meditation, I shall here employ some broad brushstrokes when speaking of these traditions and meditation’s place within them – an approach that hides countless internal disagreements and nuances in major theoretical positions. Nonetheless, given my present aim to highlight the widespread commitment across these traditions to a broader framework for the

development of introspective skill, it seems wise to begin with relatively broad points. The following thus offers some of what Gethin (1998) terms the “foundations” of Buddhist thought surrounding meditation – ‘those fundamental ideas and practices that constitute something of a common heritage shared by the different traditions of Buddhism that exist in the world today [... which] are all, in one way or another, assumed by and known to all Buddhism.’ (p. 3)

The first thing to note here is meditation’s subservient status to the broader *soteriological* aims that characterise Buddhist thought. At root, the ideas and practices of Buddhism are uncontroversially concerned with *suffering* and its overcoming. The historical Buddha is held to have discovered the origin of suffering and the means to escape it; suffering arises from *craving*, while its dissolution comes in the abandonment of craving (see SN 56.11 in Thanissaro, 1993). Literally “thirst” (Pali: *taṇhā*), craving should be understood as more than simply desire, but desire that has taken on a compulsive quality (see Lusthaus, 2003, p. 61).⁵ The object of desire—which might be an object, person, pleasurable experience, or even an idea—is herein framed as a *requirement* rather than a mere preference. Teasdale and Chaskalson (2011) note that ‘this compulsion is reflected in our felt experience and inner language which are dominated by a sense of must, should, ought, have to, need to, if only’ (p. 94). We might crave to *attain* things or to *escape* from things, with craving traditionally considered manifest in experiences like greed, lust, hatred and aversion. These are variously collected (with others) under the category of the ‘defilements’ (Gethin, 1998, p. 175), ‘fetters’ (Bodhi, 2000, pp. 1565-66) or ‘Hindrances’ (Thiradhammo, 2014), as is the term I shall go forwards with.

How then is one to abandon craving? A fundamental proposal here is straight-forward: one must first come to understand the operation of the Hindrances and the way they produce

⁵ All future parenthesised italics indicate original Pali terminology.

harm (AN 3:101-102 in Bodhi, 2005, pp. 192-193; Thiradhammo, 2014, p. 17-19). One must *see the way in which one suffers*.⁶ Unfortunately, such vision is held difficult to attain. The untrained mind is believed so turbulent as to obscure the patterns productive of suffering within a whirlwind of activity and distraction. Given such a predicament, relevant knowledge is said to require that one first generate certain conducive mental conditions. And this is the task for which meditation is held central.

Though later traditions variously adapt the meditative enterprise, Gethin (1998) notes that several principles of Early Buddhist meditation theory tend to ‘hold good’, highlighting the central significance of “stilling” and “clearing” the mind in meditation (pp. 174-176). These acts generate two conditions—*stability* and *clarity*—widely acknowledged in Buddhism as pillars atop which relevant introspective knowledge can be built; the mind should be stable, and the mind should be clear. In contemporary terms, we might call these two “operational conditions” for introspective judgement, i.e., conditions under which introspective judgements tend to come out accurate (see Goldman, 2004; Spener, 2015, p. 316).

Stability concerns the degree to which the mind is fixed upon a particular object, be that object mental or physical. The benefits of being able to cultivate high stability are relatively intuitive; good introspective judgements require that one can stay *on target*. Training is thus needed in *directing* and *holding* attention upon appropriate mental regions or specific mental features, without continual lapses into distraction (e.g., into concern with the past or future) (see Davis and Thompson, 2013, p. 592; Shankman, 2008, p. xvi; Wallace, 1999;).

Clarity is more complex, as this condition tends to be unpacked in two distinct fashions, tracking two ways that we typically make use of the term. In its first sense, clarity is the degree

⁶ This quest to understand the workings of the Hindrances forms part of the more fundamental aim to understand all phenomena in terms of the Three Marks: “suffering”, “impermanence” and “not-self” (see Gethin, 1998, pp. 174-176).

to which some aspect of the mind manifests as *sharp* or *vivid*. Conventionally, it's understood that the more vivid the experience, the less likely it will be to pass unnoticed or poorly apprehended (Davis and Thompson, 2015, pp. 50-54; Wallace, 1999, p. 177;). A more vivid emotional experience, for instance, is thought to be more easily detected, and to make more salient the emotion's type, internal structure and the parts of the body involved or presented therein (see Teper, Segal and Inzlicht, 2013; Farb, Segal and Anderson, 2015).

In its second sense, clarity maps the degree to which some aspect of the mind manifests *without competition for attention* – without the usual mass of accompanying (mental) phenomena that challenge it for our concern (Catherine, 2013, p. 21). In the same way that a person in a crowd will be more difficult to track and describe than if alone, so parts of experience will be harder to discern and interrogate if crowded by other activity. It is therefore useful to create what we might call a “background of stillness”. Without this, there is a tendency for much of that surrounding activity to call attention towards itself, exerting an insidious detracting effect upon our judgements.⁷ To the extent that superfluous mental activity subsides then, introspective judgements can be improved (see Gethin, 2004). This sense of clarity is closely linked to the first, for the removal of resource-sapping distractions outside one's focus allows for resources to be *diverted* into one's object of focus itself, which increases the vividness of that aspect of experience (see Thompson, 2015, p. 76). In this way, the two senses of clarity are really like two sides of the same coin.

I hold the above claims reasonable, in at least qualified form (see Roberts, 2018a) and won't defend them in the present paper. Rather, I shall interrogate proposals that assume

⁷ One might worry that this second sense of clarity is dangerously close to stability. Yet, it's important to note that the two can diverge. Stability refers merely to the degree of unvarying focus. And though it's true that phenomena that challenge for our attention have a tendency to *then* destabilise the mind, this needn't always happen. It's perfectly possible to be subject to subtle distracting influences from outside (e.g. from the periphery of one's vision) without one's attention being completely dislodged.

stability and clarity can be of epistemic benefit. Certainly, it remains contentious both how conducive such conditions are to introspection and whether the judgements that arise from them will be generalisable. Nonetheless, their theoretical plausibility and growing empirical support (see Fox et al., 2012; Thompson, 2015, p. 57) make their repercussions well worthy of consideration. One such repercussion is that introspective capacities might be *improved* by training people to cultivate such conditions. And this is where meditation practices become relevant.

Fig. i. gives a schematic outline of two central types of Buddhist meditation practice.⁸

Focused-Attention	Open-Monitoring
1. Hold attention upon a designated object	1. Be open and attentive to all contents arising in the stream of experience, moment by moment
2. Notice distractions that drag one away from the object	2. Notice <i>reactivity</i> to mental contents, or past and future narratives, that make one lose touch with experience
3. Release distraction	3. Release reactivity
4. Return attention to the object	

Fig. i. Meditation Schematics.

Focused-Attention is the simplest form of meditation, appearing as *samatha* or “calm-abiding” in Early Buddhism. In this practice, the student chooses a simple object on which to meditate, such as the breath, a pebble, or a mental image of the Buddha, and cycles repeatedly through the above four steps. An object of Focused-Attention might be ostensibly physical or mental,

⁸ The categorical terms used in *Fig. i.* are not native to Buddhism; they have been formulated in contemporary scholarly and scientific literature as neologisms to better categorise diverse practices from across the Buddhist schools (Thompson, 2015, p. 51). Though this simple division suffices here, see Lutz et al. (2015) for more recent and more nuanced approaches to categorisation.

internal or external, with “object” simply designating something one is opposed to (see Thompson, 2007).⁹

Contrastingly, in Open-Monitoring practice, the student has no set object of focus. Instead they engage a calm, non-judgemental attention, which simply registers whatever emerges in experience—be that sights, sounds, thoughts, or feelings, etc.—sometimes applying simple labels to arising contents. Thompson (2015) notes that the goal here is to practise monitoring ‘without getting caught up in cognitive and emotional reaction’ to what arises (pp. 52-53). Whenever one does get caught up, one registers that too and returns to a state of open awareness.

The basic route by which these practices can improve introspection is relatively simple. With practice, students become more proficient at undertaking the attentional gestures comprising the steps of the schematic. And each of these gestures promotes one or both of the conditions believed conducive to introspection.

Firstly, practice in Focused-Attention allows one to gradually hold attention upon an object for longer (Carter et al. 2005; MacLean et al., 2010) – it makes one more proficient at undertaking step 1, itself clearly conducive to stability. Likewise, meditators develop skills in recognising and releasing distractions (Ainsworth et al., 2013; Menezes et al., 2013; Tang et al., 2007), becoming more proficient at steps 2 and 3, such that less time is spent away from the object. Stability is also promoted by an improved ability to re-direct attention according to one’s intentions (step 4) once concentration has lapsed (Jha, Krompinger and Baime, 2007).

Improvements in steps 2 and 3 of Focused-Attention practice also promote clarity. With practice, it’s reported that *peripheral* distractors (antithetical to clarity in its second sense) can

⁹ Certain factors complicate this account of Focused-Attention practice. Traditionally it’s common to *shift* object once a certain degree of concentration is reached (see Gethin, 1998, pp. 181-184; Shankman, 2008, pp. 57-59). Nonetheless, this simple outline is sufficient here.

be noticed and released before they completely dislodge one's focus (see Lutz et al., 2015; Thompson, 2015, pp. 51-52). And, as noted above, such capacities support clarity's flipside (vividness), as more cognitive resources can then be channelled into the object in step 1, increasing the subjective vividness of this aspect of experience.

Open-Monitoring practices are more oriented towards clarity. Steps 2 and 3 are aimed at supercharging clarity by regulating ever subtler reactivity to arising mental contents. Disengagement from overt reactive tendencies is thought to become faster and easier for proficient meditators (Britton et al., 2014, p. 73; Hasenkamp et al., 2012) meaning more subtle reactivity can then be tackled, while the overall level of reactivity also eventually lessens (Lutz et al. 2008, pp. 2-3). In the context of Open-Monitoring practice, mental contents antithetical to clarity ought not be conceived as focal or peripheral distractors, they are instead *anything* that calls one away from non-judgemental open-awareness, and thus acts as a competitor to the calm, observing faculty. Relinquishing these "judgemental" detriments to awareness, it's thought that the subject can 'free up resources' (Britton et al. 2014, p. 73) for use elsewhere, enabling them to increase the vividness of the stream of experience monitored in step 1.

To summarise then, Focused-Attention and Open-Monitoring strengthen attentional gestures supporting stability and clarity. These conditions (*ex hypothesi*) provide an environment conducive to introspective knowledge. Moreover, they are cultivated in meditation for epistemic gain – they promote the "wisdom" essential to the larger Buddhist project of tackling suffering. As Dreyfus (2011) notes, the point of meditation is not simply to *attain* a state of stability and clarity, 'but to *use* this state to gain a deeper understanding of the

changing nature of one’s bodily and mental states so as to free our minds from the habits and tendencies that bind us to suffering’ (p. 52, emphasis added).¹⁰

Now, given increased demands for introspective skill within cognitive science, and growing awareness of meditation in the west, many have begun to speculate over the potential uses of meditation within science itself.

1.2 Meditation and Cognitive Scientific Methods

Proposals concerning the methodological employment of meditation are most overt in the “neurophenomenological” research program, which proposes that skilled introspective descriptions can be used to ‘guide and shape’ scientific inquiry (Thompson, 2007, p. 329; see also Lutz and Thompson 2003, pp. 31-33; Varela, 1996). Neurophenomenological experiments proceed by first offering subjects a degree of introspective training and then soliciting descriptions about the *structure* of certain experiences, particularly their temporal structure (how they unfold over time). With such accounts in hand, scientists can then *search for* structurally analogous patterns in collected third-person data. In this manner, neurophenomenologists seek to reveal underlying biological processes previously unrecognised as important to those types of experience (Colombetti, 2014, p. 147; Lutz and Thompson 2003, pp. 31-33).¹¹ Phenomena investigated this way include: perceptual fusion

¹⁰ Framing meditation as aimed principally at understanding or “insight” (*paññā*) is one of two inconsistent models in early Indian texts. Griffiths (1981) outlines a tension between this framing and another in which meditation aims primarily at tranquillity (*samādhi*). This model de-emphasises the need for understanding in favour of purifying the mind of ever subtler activity. I work with the insight model here.

¹¹ For some reasons that structural analogy is thought important, see Roberts (2018b). Neurophenomenology also seeks to apply constraints in the reverse direction. Third-person data should itself be used to ‘refine and revise’ phenomenological analyses, along with giving subjects new things to look for. In tandem, these are held to yield a method of ‘reciprocal constraints’ (Varela, 1996), better able to address the “explanatory gap”.

(Lutz et al., 2002); the sense-of-self (Don-Ziderman et al., 2013); and feelings of awe and wonder (Reinerman-Jones et al., 2013).

Most neurophenomenological studies have not employed meditators (but see Don-Ziderman et al., 2013; Garrison et al., 2013a, 2013b), instead employing subjects minimally trained in attentional gestures resembling those trained in meditation, or in different capacities held important for introspection like descriptive or categorical skills (e.g. Lutz et al., 2002). Nonetheless, meditation is consistently recommended as a valuable source of training here (Thompson, 2007; Thompson, 2015; Thompson, Lutz and Cosmelli, 2005; Varela, Thompson and Rosch 1991/2017). Indeed, it has been proffered as an ideal means to achieve the ‘bracketing’ of ideas, assumptions and judgements—the epoché—held important for investigating experience in the phenomenological tradition (Thompson, Lutz and Cosmelli, 2005, pp. 70-71). The meditative traditions have thus been held to fill instructional holes left by phenomenology concerning the actual *practicalities* of deploying introspectively conducive conditions in a neurophenomenological context.

Most radically though, Varela (1996) suggested that any science capable of tackling consciousness needed introspective training to become internal common practice not merely upon experimental subjects but also *amongst researchers themselves* (see also Gallagher and Zahavi 2008, p. 33). He spoke of the need to ‘build a sustained tradition of phenomenological examination [demanding] [...] a re-learning and a mastery of the skill of phenomenological description’ (p. 346). Furthermore, he noted that:

[This] proposal implies that every good student of cognitive science who is also interested in issues at the level of mental experience, must inescapably attain a level of mastery in phenomenological examination in order to work seriously with first-person accounts (p. 347)

Varela is unequivocal here about the significance of introspective proficiency. And in the same work he identifies the meditative traditions of Buddhism as a chief support in gaining such mastery (p. 341, 346). This is the proposal of principal interest to me here.

Varela's ideas have been pushed recently, with theorists proposing meditative training *within* the research community itself. For instance, Colombetti's (2014) "enactive" manifesto for emotion science identifies meditation as a means to train researchers in the attentional gestures underpinning skilled "self-observation" (i.e. introspection) (p. 149; see also Desbordes and Negi, 2013; Kordeš and Markič, 2016). On the one hand, it's thought that such skills can help researchers to understand the reports of their experimental subjects and guide those subjects to previously unnoticed dimensions of their experience while gathering subjective data (Colombetti, 2014, pp. 149-155). More broadly though, the proposal seems to be that introspective skill can help to fruitfully guide researchers' own empirical endeavours, preventing them from building poor phenomenological assumptions into their research agendas or experimental designs and thus pursuing philosophical or empirical dead ends (see Kramer and Bitbol, 2014; Metzinger, 2016).

To appreciate this latter idea, one can make analogy to other fields of study. Consider how uncontroversial it is to suggest that the committed linguist or anthropologist immerse themselves in their topics of study, by experiencing a language or culture first-hand.¹² Doing so, mistakes based upon received wisdom are minimized. Individuals keep in touch with the real subject matter of their research, rather than being sucked in by theoretical models that are apt to be taken more importantly than (or confused with) the phenomena under question – a mistake Korzybski (1933/1993) christened as *mistaking the map for the territory* – or which can swiftly diverge from what is supposed to be measured by these things without recognition.

¹² Thanks to Scott Sturgeon (personal correspondence) for this analogy.

Varela's (1996) proposal can be read as suggesting that consciousness researchers might benefit in an analogous fashion through introspective training.

A proposal resembling this comes from Gallagher (2010), who suggests that we use skilled introspective insights to inform researchers' experimental designs – a method he calls “front-loading”. Doing so reduces the likelihood that third-person measurements will confound different experiential features and properties, or lump distinct dimensions of experience together under a single kind and provide misleading or unedifying data (pp. 26-29). Gallagher recommends drawing such insights from external authorities, though there is no reason that these could not come also from the introspectively-proficient researcher themselves. Either way, Colombetti (2014) holds the failure to incorporate introspective proficiency—the institutional acceptance of what she calls the “just-take-a-look” attitude to introspection (p. 148)—as a central factor underlying the failure to identify unique biological signatures for different kinds of emotional experiences (pp. 160-161, p. 170), and she identifies meditation practices as one means to reinvigorate emotion science accordingly. Relatedly, Desbordes and Negi (2013) have called for ‘increasing the pool of individuals trained in both contemplative practice and modern science’ (p. 2). And these are more than just ideas; they are already being put into action.

Large organisations such as the Mind and Life Institute are bringing together researchers and meditators for exchange of ideas and practices. One also sees specialised events like “Towards a Mindful Cognitive Science”, a week-long researcher retreat tailored to provide Buddhist meditation training to consciousness researchers (see Kramer and Bitbol, 2014). And one finds attempts to enculture the new generation, with the “Mental Autonomy Project”, setup by Thomas Metzinger, offering PhD scholarships in philosophy of mind and cognitive science that incorporate at least 1200 compulsory hours of meditation training (see Metzinger, 2016).

In sum, we are witnessing the emergence of a serious research programme aimed at utilising Buddhist meditation as a means to introspective proficiency and cognitive scientific progress.

I must first state my general support for these kinds of project; they seem valuable attempts to drive science forwards. Nonetheless, I see a significant obstacle that will be faced here, which I shall unpack in the remainder of the paper. This obstacle is revealed by examining how meditation practices are actually presented in Buddhist contemplative literature, for, when one investigates the way that introspective skill is approached here, one finds that the prescription is much more demanding than often acknowledged. It is not a simple case of thirty minutes a day on the meditation cushion; meditation traditionally comes with a whole set of additional practices held up as vital contextual supports. As I shall lay out, these contextual factors seem especially difficult to replicate in scientific contexts, casting doubt upon the realistic *scope* of the benefits that meditation practices can offer the scientific enterprise.

2 The Significance of Context

We can motivate the above concerns by looking to parallel discussions about the psychotherapeutic value of meditation. Since Jon Kabat-Zinn's introduction of mindfulness-meditation into clinical psychotherapy in the 1970s, interest in its therapeutic benefits has snowballed (see Kabat-Zinn, 2003). Yet, attempts to pin down mechanisms of benefit here have led some to stress the difficulties inherent in extricating such practices from their Buddhist homelands and applying them in alien arenas.

The emerging concern is well expressed by Kordeš and Markič (2016), who note that most of those responsible for applying mindfulness-meditation in therapeutic settings 'seem to believe that this practice will be effective, even though it has been separated from the context of Buddhist practice and belief' (p. 161). In contrast, many have begun to emphasise that

meditation is traditionally only one part of a much wider program of individual transformation—The Noble Eightfold Path—which includes practices for transforming doxastic structures, social attitudes and over-arching existential goals (see Farias and Wikholm, 2015; Kirmayer, 2016; Murphy, 2017; Ozawa-de Silva, 2016). And, O’Brien (2017) notes that ‘[w]ithin the context of Buddhist practice, all parts of the Path support and affect all other parts of the Path [...] [so] [f]rom a Buddhist perspective, when mindfulness is practiced in isolation of the rest of the Path it becomes something different from Buddhist mindfulness’ (Ibid.).

The remainder of the Path is therefore not only a standalone means of tackling craving and suffering; it supports the pursuit of this goal *within* meditation too. Growing sensitivity to this within psychotherapy has led some to voice concerns that the therapeutic value of mindfulness practice is endangered or limited by failing to replicate its traditional context. Wilks (2014) notes this as a live debate in Buddhist journals and amongst Buddhist teachers. She relates how mindfulness-based therapies ‘are said to be diluting the Dharma, watering down the radical teachings of the Buddha into some sort of “Dharma-lite” and offering a “one-fold path”, without reference to the other limbs of the eight-fold path.’ Now, this consideration is all well and good, but what relevance has it to meditation’s *introspective* value and the present proposals?

To understand the significance of the above, one can return to steps 2 and 3 of the Focused-Attention practices and ask: what precisely *are* the distractions that the practitioner is attempting to recognise and set aside here to cultivate appropriate “operational conditions” for introspection? Those who have attempted such practice quickly realise that, in the majority of cases, these distractors are nothing other than the Hindrances themselves - the experiential manifestations of craving. They are: concerns about the past or future; strong urges to get up and do something different; simmering resentments about recent personal injustices.

Accordingly, Nyanaponika (1988, p. 33) outlines the meditative development of concentration (stability) as a procedure of gradually removing the Hindrances, while Catherine (2013) marks the Hindrances as obstacles to ‘clear seeing’ (p. 21, see also Shaw, 2014, chpt. 3; Thiradhammo, 2014, pp. 17-18).

In the above, we can see that the distractors one aims to set aside in Focused-Attention (and the “reactivity” in Open-Monitoring)—whose dissolution cultivates conditions of clarity and stability—often *are* the very things targeted for removal in therapeutic contexts.¹³ So understood, the introspective and therapeutic quests are *not two distinct enterprises*. What seems at first a purely therapeutic benefit of meditation can be an introspective benefit too. This means that difficulties in the one enterprise will mean difficulties in the other. And if the Path is important for ensuring meditation’s therapeutic value, it will also be important for ensuring its introspective value. In consequence, the importance and viability of such factors will need considering not merely in western psychotherapeutic communities (where the issue is already difficult) but also within scientific circles proposing to utilise meditation, where this problem might appear even more acute.

This overlap between the introspective and therapeutic benefits of meditation might cause confusion at first. For it was suggested in §1.1 that the meditator was aiming to remove certain distractors in practice to better *understand* the nature and function of the Hindrances (an introspective benefit), in virtue of which they could *then be eliminated* (a therapeutic benefit). The above points, however, make it seem that this procedure is being short-circuited – meditation involves also the removal of craving straight off the bat. If both these things are

¹³ For this reason, a verse in the Samyutta Nikaya relates that the Hindrances ‘produce blindness, *cause lack of vision and ignorance*, obstruct wisdom, associate with distress and are not conducive to awakening’ (SN 97, in Thiradhammo, 2014, p. 17).

true, then the practitioner would appear to be removing craving, *in order to understand craving*. This is indeed the proposal.

As Gethin (1998) notes, the meditative path aims to make practitioners aware of *gradually more subtle* forms of craving, which is said to involve the progressive removal of overt kinds so more elusive kinds can become apparent (pp. 174-181, p. 198). Thiradhammo relates that subjects need to be free of some aspects of the craving mind to properly register others (2014, p. 31). Comprehending subtler aversions to particular situations, or people, for instance, is only possible when the mind is free from raging anger. There is thus a reflexive (i.e. feedback) relationship posited between introspective proficiency and Hindrance removal. The abandonment of some Hindrances promotes greater awareness of others, which in turn supports the “wisdom” that can lead to further reductions in the Hindrances (and thus further insights to emerge). This constant interplay between the introspective and the therapeutic in meditation only further impresses the seeming significance of the contextual problem that will be faced by those promoting meditative training in science. To the extent that the therapeutic effects of meditation are left unsupported by the remainder of the Path, this reflexive relationship will also be curtailed, and resultant introspective benefits will be less pronounced.

To understand the significance of these concerns, it’s essential to delve further into the contextual factors increasingly recognised as important to meditation within psychotherapy. In the next section, I’ll thus give a brief overview of such broader practices aimed at supporting the meditative removal of the Hindrances. This will flesh out a “broader path” to introspective skill, as I shall call it, beyond the two central practices that usually get the headlines. By interrogating the broader path’s character, one can then better judge how realistic it will be to provide such supports in science. And in examining the nature of its supportive function, one

can better gauge the degree of introspective benefit in danger of forfeiture if this broader web of practice is not replicated.

3 A Broader Path

Across the Buddhist traditions, an enormous range of practices can be found for working with the Hindrances. Nonetheless, they are essentially divisible into three categories—meditative, ethical and intellectual—reflecting (roughly) a traditional three-fold partition of the eightfold-path. Though vast amounts could be said about each, there is space here only for a cursory survey which can prompt deeper investigation by others.

The practices I shall outline address the Hindrances in two ways: (i) they act as *standalone* means of dissolving the Hindrances (ii) they *support the meditative procedure* of tackling the Hindrances in Focused-Attention and Open-Monitoring practice. I shall spend more concerted time on the latter. But to comprehend this, it is first essential to unpack the former.

3.1 Standalone benefits

3.1.1 Meditative Training

§1.1 illustrated two central meditative practices for working with the Hindrances (Focused-Attention and Open-Monitoring). Yet, most Buddhist traditions prescribe additional meditations serving this same function. The most prominent practices involve the deliberate cultivation of certain virtuous emotional attitudes or “mental qualities”, with four qualities in particular becoming central. These are known as the four “Divine Abodes” (*Brahma-vihāra*) or “Immeasurables” (*appamaññā*) – (1) loving-kindness (*mettā*); (2) compassion (*karuṇā*); (3)

sympathetic-joy (*muditā*); and (4) equanimity (*upekkhā*). Meditations devoted to the Immeasurables involve the deliberate regulation of attention—bringing to mind and focusing upon particular phrases, people and scenarios, often in a very imaginative fashion—to yield increasingly strong manifestations of these four qualities.

The first and most noticeable benefit of such practice is that the Immeasurables can be then deployed *in situ* to dissolve occurrent Hindrances. They can act as immediate counteractive measures, thanks to being attitudinally *opposite* in nature and therefore incompatible with the continued existence of a target Hindrance. For example, the individual who suffers from the Hindrance of ill-will is advised to cultivate loving-kindness or compassion, such that the occurrent mode of appraisal manifest in ill-will is broken. *The Path to Purification (Visuddhimagga)* for example, a prominent fifth-century meditation manual, describes loving-kindness as a “solvent” that can “melt” anger (Buddharakkhita, 1995). In this manner then, operational conditions for introspective judgement can be promoted on the fly by mobilising mental qualities capable of stabilising the ‘churning’ of the mind overrun with Hindrances.

Secondly, and of greater importance, the Immeasurables are employed to reduce the overall degree of self-centredness underpinning potential *future* cravings. They transform one’s *disposition* to crave. By cultivating virtuous qualities, the student is said to witness the meaning these bring to their lives and gradually become more content with a life based upon selfless attitudes (see Brewer et al., 2013). This satisfaction is supported by the positive feedback received from the “ethical” actions that the Immeasurables promote. For this reason, the Immeasurables have been rather awkwardly referred to in science as ‘Ethical Enhancement’ practices (Vago and Silbersweig, 2012) – they motivate the broadly ethical conduct that can further support Hindrance removal, to which I now turn.

3.1.2 Ethical Training

Occupying the Right Conduct (*sila*) division of the Path, one finds Right Speech, Right Action and Right Livelihood, prescribing models for broadly “ethical” behaviour. Such normative models might reference generosity, truthfulness and even the importance of vocational decisions. Importantly, they are not merely intended to support social harmony, but also to yield beneficial changes in an individual’s psychological landscape. As Huxter (2015) notes:

Acting with ethical integrity includes making the choice to be harmless in our action. When we are not intentionally harming ourselves or others, it is likely that our minds will not be plagued with hatred, guilt or fear. When we avoid harmful speech, actions and occupations, our conscience is more likely to be clear and our minds more easily able to focus on the immediate experience of life (p. 38)

Here, we see the clear linkage between ethical conduct and the factors of stability and clarity that emerge through dissolving the Hindrances. Though the different Buddhist schools flesh out these ethical prescriptions in inexpressibly varied detail, the basic relationship between ethical conduct and Hindrance dissolution is clear enough. Given greater awareness of this relationship, Western psychotherapeutic communities have increased the roll-out of compassion-based practices alongside or within mindfulness-based therapeutic programs (Wilks, 2014), facilitating the behavioural conduct traditionally underpinning mindfulness practice and *extending* the therapeutic effects available through narrower attentional techniques alone. However, awareness of this relationship is minimal amongst those advocating Buddhist contemplative methods within cognitive science. Bar a few notable exceptions (e.g. Wallace, 2004, chpt. 3; Depraz, Varela and Vermersch, 2003), much focus on the introspective benefits of Buddhist methods has been blind to the importance of ethics in this endeavour.

3.1.3 Intellectual Training

The final set of practices appropriate for tackling the Hindrances are *intellectual* in nature, involving careful reflection upon Buddhist teachings. Themes of such reflection might include ethical or metaphysical principles extracted from foundational canonical scriptures or broader Buddhist literature. For example, a Theravāda student can reflect upon the “three marks of existence” (*tilakkhaṇa*) noted in the *Dhammapada* – three ‘marks’ characterising all phenomena: impermanence (*annica*), being subject to inevitable dissolution; not-self (*anattā*), not identifiable as “me” or “mine”; and suffering (*dukkha*), being a potential source of unsatisfactoriness or pain. Sustained reflection upon these works to reduce levels of craving for obvious reasons – the holding of a particular thing as imperative to one’s existence (i.e. craving) will be weakened by properly grasping that both you and it will inevitably turn to dust (Gethin, 1998, p. 70).

While one should acknowledge that many Buddhist traditions emphasise the dangers of intellectual reflection (e.g. forms of Chan and Zen), holding the ‘Right Wisdom’ of the Path to arrive primarily through other (i.e. meditative) routes, this is unrepresentative of Buddhism as a whole, particularly in its Indo-Tibetan forms (see Williams, 2008, pp. 45-51). Generally, students are encouraged to reflect upon Buddhist teachings, remain sensitive to their relevance in life and meditation, and firmly sear them into their minds.

Early Theravāda schools emphasise the study of the “Higher Teachings” (*Abhidhamma*) alongside meditation (Sharf, 2015, p. 474). Likewise, early Mahāyāna traditions, despite shunning the scholasticism of the Abhidhamma schools, also encouraged intellectual and philosophical reflection. Williams (2008) notes that the ‘Perfection of Wisdom’ (*Prajñāpāramitā*) texts of such traditions, despite highlighting the centrality of non-conceptual forms of wisdom, are not intended as prescriptions for training. Rather he states

that such training involves initially ‘*extending* the analysis’ (p. 55, emphasis in original). Thus, ‘to conclude that wisdom for the Perfection of Wisdom is the result of simply cutting discursive thought [would] be a historical and religious error’ (ibid.). One finds a similar emphasis upon “discursive meditations”, more akin to the “meditations” of the West, in the Mādhyamika schools of India (see Williams 2008, pp. 79-81) and later Tibetan traditions, wherein a distinction is made between ‘resting meditation’ and ‘analytical meditation’, the latter emphasising the important role of ‘questioning’ and ‘inquiring’ (Thangru, 2004).

In sum, intellectual reflection plays an integral role in developing the wisdom that can overcome craving. And when this is considered in combination with the remainder of the Path, we see that the corpus of Buddhist methods for generating conditions conducive to introspection is a broad, varied and demanding prescription, which encompasses meditation but is not exhausted by it.

Having now unpacked the above standalone power of the broader path against the Hindrances, it’s now possible to comment on how this path can support the removal of the Hindrances *within* meditation.

3.2 Supportive benefits

There are two ways that the broader path’s supportive relationship to meditation can manifest. Firstly, it can reduce the number and strength of the Hindrances the student brings to meditation thanks to dispositional transformations, making for easier progression to more advanced stages where more refined attentional gestures are practiced. Call this *pre-empting difficulty*. Secondly, the broader path gives the student a set of strategies and techniques that can be deployed in situ at points of difficulty to catalyse the practice. Call this *developing a practical toolkit*.

3.2.1 Pre-empting Difficulty

The importance of entering meditation with prior immersion in the broader path is much commented on in contemplative circles as a way to forestall obstacles. For instance, Wallace highlights the prudence of reducing self-centredness through ethical action and practising the Immeasurables as a precursor to Focused-Attention meditation (2004, p. 73; see also Huxter, 2015, pp. 33-41). With these under one's belt, one will face *fewer* Hindrances in meditation, or *less intense* Hindrances, easier to put aside, supporting awareness (and thence removal) of deeper and more subtle cravings and distractions (Wallace, 2004, p. 73; Thiradhammo, 2014, p. 41). In this way, the student gets the biggest distractors out of the way ahead of time, speeding up practice from afar, meaning they can devote more time to practising the recognition of subtler Hindrances.

Moreover, the practice of this latter capacity is one of the really unique things about Focused-Attention and Open-Monitoring practices. The broader path is well suited to tackling overt Hindrances. And more advanced stages of meditation are the principal environments suited to sensitizing oneself to and tackling subtler attachments of the mind. And it is really *this* task that the student wants to be engaging in meditation. Such a feat thus requires work to pre-empt the more overt Hindrances before taking to the cushion. Without such background of preparation, Wallace (2004) quips that he 'could spend a whole year in a *samatha* [Focused-Attention] retreat and achieve nothing but frustration.' (p. 73).

The situation is similar with intellectual practices. Periods of reflection, in which one conceptually grasps why meditation is needed, for instance, result in less of the Hindrance of doubt in the practice and swifter progress. Without a framework of reflection that ingrains the reasons to meditate, and the prior benefits it has brought, one will be faced with recalcitrant doubt in practice sessions. Nyanaponika thus speaks of 'firm conviction concerning the

Buddha, Dhamma (his teachings) and Sangha (the spiritual community)’ as conducive to meditation, itself achieved through ‘knowledge of the Buddhist scriptures’, ‘[a]sking questions about them’ and ‘keen investigation into reality’ (1994, parentheses added). Similarly, reflection upon metaphysical teachings about Impermanence will transform one’s disposition to bring a mind full of raging desire to meditation, which would take a long time to pacify and prevent the practise of more advanced skills.

In short, both intellectual and ethical practices provide a key dispositional support to the efficacy of the meditative enterprise, without which one’s progress will be limited.

3.2.2 Developing a Practical Toolkit

Not all difficulties in meditation can be pre-empted. Thus, the broader path also helps the student to construct and hone a *practical toolkit* to be brought to meditation, which can tackle difficulties as they arise. This comprises a set of robust mental strategies and gestures that can be deployed *in situ* to support the student’s progress.

To illustrate this point, consider an all too common experience for novice meditators: sitting down to Focused-Attention practice, only to be faced with a recalcitrant Hindrance (say, resentment at some recent event) that simply cannot be put aside and which threatens to endure for the entire meditation period. Here, the student cannot perform step 3 (“release distraction”) effectively. And it is in these scenarios where the prior development of a practical toolkit proves it worth.

For the student who has done their “homework” and practised the Immeasurables, these attitudes become tools available within practice – they can be mobilised *in situ* to tackle particularly stubborn Hindrances. For instance, in an early discourse titled ‘The Removal of Distracting Thoughts’ (MN 20 in Bodhi, 2005, pp. 275-278), it’s recommended that the student

deliberately turn their mind to ‘some other sign connected with what is wholesome’ to counteract ill-will when it arises during meditation, which can entail thinking of some topic conducive to loving-kindness (*mettā*). This can accelerate the removal of stubborn Hindrances, progression through the steps of practice, and thus advancement to more advanced stages.¹⁴ Without the capacity to easily generate such states, they won’t be available to call upon, and the student may become stuck with specific distractions, unable to stabilise the mind. In this manner, the toolkit helps students to *actually undertake the steps prescribed* in meditation instruction and thus to advance onto more refined levels of practice, to the benefit outlined in §3.2.1.

The situation is again analogous in the case of intellectual practices. Kuan (2008) highlights how early mindfulness texts prescribe the effortful ‘recollection’ of inspiring subjects, themes or persons as either (i) tone-setting preliminaries to practice, or (ii) tools to be used within meditation itself (pp. 52-56, 62-65). For instance, the ‘Ten Recollections’, delineates a list of ten themes for reflection that can be used in either manner. Here, students ‘call to mind’ aspects of the Buddha’s personality, inspiring Buddhist deities, and memorised extracts of Buddhist teaching (e.g., on the inevitability of death). This proceeds through effortful discursive thought and constructive imagination of example cases.

¹⁴ On account of their “Hindrance-burning” qualities, some positive mental qualities are considered important building blocks to construct in developing the early stages of *jhāna* (a series of ever more refined concentrations). Here, one deliberately cultivates such qualities not in response to token Hindrances, but as an ongoing affective background to support practice (see Shankman, 2008, pp. 32-42). One might object here that, unless these are swiftly dropped after “burning away” the Hindrances, they are merely replacing those with further potentially distracting mental processes. In fact, meditative theorists suggest that positive mental qualities can actually support the lower-degrees of concentration—for instance, the feeling of joy can help one stay on target by ‘preventing boredom’ (Shankman, 2008, p. 201)—because they are fundamentally different in kind from the Hindrances. The Hindrances themselves have an inherent *yearning* or *driving* character, seeking what is not present and are thus especially responsible for ‘churning up’ the mind in an introspectively detrimental manner.

In their first function, these recollections are used immediately prior to meditation as “ground-clearing” tools that sweep away difficult states of mind or set an appropriate motivational tone. In their latter function, they are integrated more minimally within the meditation itself, being deployed at points of difficulty. Kuan (2008) marks the recollection of efficacious themes as an example of the ‘deliberate forming of conceptions’ recommended in early discourses (p. 52). For instance, though a student engaged in mindfulness of the body will usually simply re-orient attention to some prescribed bodily region when a Hindrance is registered, ‘The Removal of Distracting Thoughts’ also instructs students to remind themselves of the disadvantageous consequences of indulging the Hindrances when meeting a stubborn Hindrance that will not be dissolved by ordinary means. Similarly, the call for the meditator to direct his mind to ‘some other sign connected with what is wholesome’, in this discourse, might involve reciting a particularly meaningful phrase from a teaching.

These examples touch upon a larger point made by Chambers, Gullone and Allen (2009) about the broader significance of intellectual reflection in therapeutic applications of mindfulness-meditation. The authors note that attempts to study mindfulness under the extant category of ‘emotional regulation’ practices have sometimes ignored distinctive intellectual features deviating from familiar forms like cognitive behavioural therapy (CBT). In particular, the authors diagnose an under-appreciation of the rich metaphysical theory historically underpinning the practice (p. 567). Here, students are taught to conceive of their emotions under a philosophical framework which stresses their transience, reducibility, and lack of ownership by any independent “self” who must feel responsible for them. With such a framework in place—i.e., with these ideas seared into the subject’s mind through analytic

reflection—emotional disturbances can be *philosophically* tackled, as teachings are recollected in meditative practice.¹⁵

For example, when a feeling of ill-will arises during practice, it will be common for novice meditators to then feel bad or guilty about having felt this – the experience will be taken as a sign of failed progress, of general bad character, or of being the kind of person unable to benefit from practice. In this way, rather than simply dropping the Hindrance, the student will set off a cascade of further distractions. In contrast, for a practitioner well-grounded in Buddhist theory about the lack of any ultimate “self” who must control and feel guilty about emotions, such cascades can be forestalled. In bringing to mind such teachings at the point of emotional reaction, further reactivity of a *second-order* nature can be pre-empted. It’s unfair to say that the significance of this broader theoretical framework isn’t appreciated amongst mindfulness-based therapists (see Wilks, 2014), though it’s important not to ignore the philosophical depth of this structure in Buddhism, which can be sanitized when the practices are brought under western theoretical frameworks.

This completes my outline of the factors held important for tackling the Hindrances. As argued in §2, it consequently reveals numerous additional factors of neglected importance in the cultivation of introspective proficiency. With these on the table, it’s possible to re-address the proposed employment of meditation in science, better equipped to comment upon the scope of prevailing strategies here.

¹⁵ One might worry that such metaphysical dealings might “cognitively penetrate” experience, biasing what is found through such methods. This is an important point, but I think one that can be handled. I address this issue in Roberts (2018a). See also Thompson (2015, pp. 56-7) and Colombetti (2014, pp. 155-158)

4 Methodological Repercussions

4.1 Summary Findings: A Dilemma

To recapitulate briefly, the turn to Buddhist meditation in cognitive science grew out of renewed enthusiasm for subjective reports and rising awareness that rigorous reports demanded subjects possess certain cognitive skills. These skills could ensure the appropriate “operational conditions” prior to introspective judgement. Varela, Thompson and Rosch (1991/2017) heralded the Buddhist contemplative traditions as one source of such skills – they were said to offer a storehouse of rigorous techniques for attentional and attitudinal development that would enable practitioners to perform the kind of attentional gestures necessary to gain reliable access to experience.

Delving further into this Buddhist treasure-trove however, prompted by contemporary psychotherapeutic discussions, has revealed it to contain a wealth of practices important to introspective skill, stretching far beyond the central meditative practices typically held up as valuable in science. This itself reveals valuable practices that have been neglected in the quest for introspective proficiency. But it also shows that the Buddhist programme for introspective development is *far more demanding* than is often credited in cognitive science.

Moreover, §3.2 has suggested that the supports of these broader elements are critical for advanced introspective skills. Without such conditions in place, students are apt to remain stuck in the early stages of practice, rarely able to reach the degree of mental quietude enabling the practice of more refined attentional gestures. These underpin the most nuanced introspective skills, involving the recognition of subtler tendencies of the mind. What then does this mean for the proposed use of meditative techniques in science?

Most significantly, the benefits accrued through meditative training in soteriological environments will be rather different to those available in science, if the latter does away with meditation's conventional accompaniments. Many of the traditionally declared "effects" of meditation depend upon the *holistic* nature of Buddhist practice. So, attempts to utilise mere elements in science for epistemic benefit will have different consequences. Furthermore, having surveyed the character of this broader path to introspective proficiency, one will be struck by its seeming incongruity with the methods and practices of science. Considered from a researcher's perspective, it's not merely the breadth of the Path that is demanding, but the components themselves. Prescriptions for ethical behaviour and compassion practice lie far beyond the ordinary remit of the scientific community. Indeed, the contextual problem appears even more acute than in psychotherapeutic contexts.

Within psychotherapy, the contextual problem refers to the diminishing therapeutic benefits available for the patient as the holistic nature of Buddhist practice is narrowed to mindfulness-meditation alone. In this environment, it's difficult enough to encourage take-up of contextual supports. Nonetheless, there is at least some precedent within psychotherapy for prescribing interventions that reach far into the daily lives of those they will benefit. Within a scientific paradigm, the proposal tackled here identifies researchers themselves as the targets of benefit, and this merely multiplies the contextual problem. There is much less acceptance in science of such invasive methods for cultivating professional skill. Researcher training is kept distinct and isolated from the broader kinds of "life-practices" prescribed on the broader path. Yet, the more one narrows this path to make it accessible to scientists, the less pronounced the individual benefits will be. In this way, one is left with a dilemma, knowing that many of the most important introspective skills available through meditation practice require pursuit of a

wider program of training, but also that this programme appears over-demanding for many (if not most) in the scientific community.

However, I do not think that this dilemma presents insurmountable problems for the proposals at issue here. In fact, I suggest that an appropriate response to such concerns mirrors that emerging within psychotherapy, where revelations of an analogous dilemma have prompted clearer reflection upon the manner in which meditation practice is presented. The relevant question facing enthusiasts for meditation in science is: how can one respond to contextual concerns, *using* knowledge of the above dilemma, to maximise meditation's benefits to the scientific enterprise as a whole? It is, after all, this wider benefit that is the ultimate concern of those advocating meditative training regimes, over and above specific benefits to individuals. With this question in mind, I suggest that the above considerations reveal a methodological choice facing those advocating meditation within science, whose options will have different consequences for the way cognitive science develops.

4.2 A Methodological Choice

For those advocating meditative training in science, there is a decision to be made over whether this training should prioritise (i) the *degree of benefit* that it makes available to individual researchers, or, (ii) the *accessibility* of the training itself. One must choose between promoting either (A) a *broader path* aimed at higher-level introspective skills for individuals, finding less take-up, or (B) a *narrow path* aimed at making lower-level skills more widely available. This is no black-and-white choice, of course. But presenting it in such stark terms here helps to draw out its character and consequences.

4.2.1 A Broader Path

The first option involves prioritising the degree of benefit made available to individual researchers. This entails promoting introspective training that replicates at least some of the contextual supports revealed as important in §3. This can be described as the more “receptive” option – it responds to revealed difficulties involved in generating introspective skill by welcoming a broader set of practices into programmes of training. From this perspective, the exposition of a larger-than-meditative framework for introspective skill fleshes out valuable details – underappreciated pieces of the puzzle – unknown or under-represented in other traditions of introspective reflection.

In concrete terms, this option can involve promoting institutional programmes of training that offer the core attentional practices (Focused-Attention and Open-Monitoring meditation) as important centre-pieces, with a requirement for some minimum number of elements of the broader path taken upon alongside. These extras could be chosen based upon specific interests of individuals. Such broader practices needn’t even be presented in a particularly novel manner. They might be framed similarly to existing “non-academic” supports to study offered in institutional settings. For instance, existing welfare provisions encouraging an appropriately calm and concentrated mind for study could be extended and emphasised as *important supports* to researcher training itself, cultivating the appropriately “primed” mindset to bring to other aspects of introspective training. Similar re-framings could be given to extra-curricular programmes for developing certain pro-social attitudes and communicative virtues under the broader banner of “citizenship skills”, which already find favour as important components of professional education in universities. In this way, promoting the broader path needn’t require proposing drastic institutional changes or any demanding commitment to a “Buddhist” way of living.

Alternatively, promoting the broader path could involve pushing something very much akin to Buddhist contemplative training. One might develop more specialist institutions – or centres within existing ones – self-consciously oriented towards contemplative approaches to consciousness science and more overt about the necessary supports to practice and their resemblance to Buddhist regimes. These could offer a wider gamut of supports in a more rigorous manner. Such setups would likely appeal only to those resonating strongly with the broader orientations of the Buddhist path, but they would generate the most advanced introspective proficiency.

What then are some likely benefits of promoting the broader path? Here one can point to the criticality of high-level introspective skill for advancing particular research enterprises in science and philosophy. Such skills will be most advantageous for well-established areas of research and those areas of science targeting phenomena with particularly fine-grained temporal structure. Developed sub-fields like perceptual science already have a reasonably developed phenomenology guiding their investigations. Any phenomenological stimulation of such research is thus more likely to require a sensitivity to very subtle experiential distinctions.

One component of this might be sensitivity to fine-grained *temporal* properties of experience. This would allow subjects to report more rigorously on the rapid evolution of experience – important in the study of experiences that unfurl over small time-scales. Take phenomena like perceptual pop-out, object-perception, depth-perception and indeed temporal perception. The subpersonal processes underpinning these happen at lightening pace. So, if there is any hope of *phenomenological* intimations about these, of the kind that can further illuminate our understanding of them, it is likely that only advanced phenomenological

proficiency will provide them. Accordingly, the broader path would seem important for appropriate phenomenological “front-loading” of experimental design here.¹⁶

Perhaps more importantly though, it is through the broader path that one might make substantive progress on the *philosophical* problems surrounding human experience. Nagel (1972) argued that addressing the “problem of consciousness” would require significant improvement in abilities to describe the subjective character of experience (p. 437, pp. 448-449; see also Chalmers, 1999). Only with an appropriately detailed picture of our explanandum, the thought goes, can we properly understand how it might relate to, or be instantiated in, the physical world. This requires introspective accounts not merely of a metaphorical nature (“X is like a cold winter’s day”), or about mental specific contents (“X is an experience of red”), but about such things as “*structural features* of perception” (pp. 449, emphasis added) which are more congruent with the explanations offered by “physical” theories. Put differently, the suggestion is that more fine-grained, literal descriptions might reveal those properties *in virtue of which* certain experiences have the contents or character that they do, further narrowing the aspects of experience seemingly “left out” by physical accounts. Indeed, it seems to be this very issue that Varela (1996) had in mind with the promotion of contemplative training, framing it as an important means of closing the “explanatory gap”.

For these reasons, broader-path enthusiasts can argue that insufficient emphasis upon contextual factors is likely to *waste* the potential of Buddhist contemplative training regimes. To argue for the broader path, they can emphasise how narrower paths will forego the possible

¹⁶ One might object that Lutz et al.’s (2002) famous study of perceptual fusion produced important insights into the neural structures underpinning the phenomenon despite employing only minimally trained subjects, and that high-level training is thus unwarranted in such cases. Yet, one should note that, though *subjects* were relatively untrained here, the *experimental design* was critically and self-consciously informed by the immersion of co-author Francisco Varela in the intricate phenomenological literature of the Buddhist Abhidhamma (see Thompson, 2015, pp. 41-45).

superiority of the contemplative traditions over *other* regimes of introspective training, such as those found in the continental phenomenological tradition. And in so doing, one may be passing up the chance to make really significant empirical and philosophical progress.

4.2.2 A Narrow Path

The second option for meditation enthusiasts is to prioritise the *accessibility* of contemplatively-inspired introspective training. Here, one promotes a *narrow* path of practice, akin to that already advocated, being prepared to sacrifice some of the introspective proficiency available through Buddhist methods to broaden appeal. In concrete terms, taking this option means pushing for greater institutional availability of the core meditative elements of Buddhist practice (Focused-Attention and Open-Monitoring). One can advocate for these to be routinely integrated into existing programmes of researcher training, even if only as optional components. What then are the plausible benefits to science of this alternative?

Aiming at a more limited lower-level improvement in introspective skill across the field at large is no insignificant development. Most obviously, such skills can support areas of research that have made scant use of phenomenological reports as methodological guides. As Colombetti (2014) notes, emotion science is a field ripe for benefit from such researcher training, with introspective reports utilised here being very minimal and coarse (pp. 44-49). Even small introspective improvements will be useful here. The same goes for research avenues in their infancy. Take the recent explosion of interest in implicit bias (e.g. Brownstein and Saul, 2016; Greenwald and Krieger, 2006; Holroyd, 2012). There was an initial tendency to regard this phenomenon as entirely unconscious, yet there are growing suggestions that it actually operates – at least on some occasions – on a conscious level, with a subtlety which regular persons are often insensitive to (see Gawronski, Hoffman and Wilbur, 2006; Hahn et

al., 2014). If so, this presents an opportunity for introspective training to help guide study into the cognitive architecture underpinning implicit bias and the search for methods to counteract it. The lack of developed phenomenology in this arena means that even small introspective improvements (supported by a narrower path of training) should benefit here.

Several arguments can be deployed to further the case for narrow training. Let us start with the most obvious. One might argue that the Buddhist route to high-level introspective proficiency is not merely demanding, but *overdemanding*, in a scientific context. Even if one is persuaded that the broader path will have significant benefits for individuals, one can argue that these won't be easily attained by scientific researchers themselves. By this reasoning, the best way to benefit cognitive science involves expending energy upon making smaller, less conspicuous practices more widely available, rather than wasting time encouraging things that are unlikely to find fertile soil. Indeed, Varela's original proposals for incorporating meditative training have already been challenged as unrealistic (see Braddock, 2001), so to go even further might be framed as severely optimistic.

Relatedly, one can emphasise the difficulties involved in operationalising elements of the broader path in any way that meets the levels of standardisation and rigour demanded in science. How exactly could one regularise practices of ethical behaviour or compassion? Even for those unconvinced of the necessity of this (of which I am one), such demands will prove difficult barriers to overcome in any attempt to put the broader path into practice.

Perhaps most significant of all though, one can point to seemingly critical differences in the ultimate character of the two paradigms that one is attempting to integrate here. In an excellent comparative piece, Patrik (1994) claims that there is a fundamental divergence between the phenomenological method of inquiry as instigated by Husserl and the contemplative methods of Buddhism. The former, she notes, is at root *knowledge-driven*, while

the latter is *liberation-driven*, i.e., aimed at the release of beings from suffering (see also Kordeš and Markič, 2016, pp. 162-165). Though phenomenologically-guided science targets different kinds of knowledge to Husserl, an equivalent contrast in underlying motive is apparent. This difference is important, for it suggests that the catalytic potential of contextual supports, in the meditative development of introspective skills, will depend upon a more foundational resonance with Buddhism's soteriological objectives. Let me elaborate this.

As I've tried to emphasise in §3, the ease with which one can put aside the Hindrances in meditation (and thus distractions in introspective methods) will be determined by such things as the degree to which one's underlying dispositions are ones of self-centredness or compassion and the level of contentment achieved through living an ethical life. The broader path helps to cultivate and strengthen these very things. Yet, sceptics of the path's utility in science might claim that a further, more basic, attitude will be required to underpin these – a commitment to some loftier *soteriological* project. Even if these contextual supports are replicated, one might think that their efficacy will be weakened without some larger soteriological purpose in mind that really commits the practitioner to their effective practice. The problem facing proponents of the broader path is thus perhaps not merely its demandingness, but that it will lose some of its power if one only has in mind *epistemic gain*.¹⁷

4.3 Weighing the Choice

§4.2 has outlined two responses to contextual concerns given in §3 – two possible means of employing meditation in science, along with their plausible consequences and some arguments for each. When deliberating between these options, the first thing to note is that this is clearly

¹⁷ From this perspective, we might think that it would be just as productive to advocate the training of Buddhist contemplatives *in science*, rather than vice-versa. For more on this possibility, see Desbordes and Negi (2013). I return to this issue of underlying motives at the paper's conclusion.

not an either-or choice. Some combination of both is possible, and I think advisable, if one wishes to maximise meditation's benefits to the field at large.

First, there is good reason for caution over broader introspective training regimes. When the broader path to introspective skill (as specified in the Buddhist traditions) is fully illuminated, it is revealed as of such a character as to deter large swathes of people that would otherwise be keen on introspective improvement. This will ensure that its influence could only be rather limited in the current climate, arriving through a small subset of individuals. Given that this considerably narrows its plausible benefits to the field as a whole, it seems sensible not to afford it too much emphasis in the promotion of meditative training.

In contrast, the relatively anodyne nature of the narrow path is likely to find wider appeal. And its ability to spread lower-level introspective skill more broadly is more likely to have significant impacts, particularly in the short term. Given that the current environment of cognitive science is one of almost *no introspective skill at all*, a strong catalytic effect upon research is plausible here in a relatively short timescale. Furthermore, gathering stronger evidence of methodological benefits through meditative training seems an important precondition for advocacy of the broader path to find favour. If more minimal kinds of contemplative training can be seen as advantageous, it will be much easier to push for their extension. Consequently, it seems wise to put most concerted efforts into promoting these narrow forms. Though the "stripped down" character of these programmes will need to be acknowledged to armour them the criticism of Buddhist scholars, and indeed from overdrawn expectations, this narrow approach to training seems best placed to benefit cognitive science in its current state.

Moreover, insufficient attempts to operationalise this more conservative strategy means it has not yet been possible to ascertain the level of introspective benefit that it can have.

Bringing meditation into the sciences is not simply to strip it of context; it is to give it a *new* context. One uses it here to study novel targets and to address questions for which there is no precedent in Buddhist thought. For this reason, one should not rule out that minimal regimes of introspective training could have unexpectedly significant benefits, without the contextual factors that prove important for the *particular* insights that Buddhist programmes aim at.¹⁸

Perhaps the most significant factor favouring the narrow path, however, concerns the knock-on benefits that will be felt in the research community through more widespread introspective proficiency amongst researchers. For instance, it's been suggested that philosophy of mind has something of a 'quality control' problem when it comes to the introspective guidance of theoretical work; bad theories are produced atop of bad phenomenology (Hutto and Myin, 2013, pp. 176-178; Noë, 2004, chpt. 2) – theories which then gain unwarranted pre-eminence and sap others' time. By my estimations, the best way to forestall fruitless research is not to pedestalise a small cadre of introspective authorities, to the exclusion and suspicion of the majority, but to spread introspective skills more widely. Without this, the sheer numbers of people engaged in different research paradigms will influence people's views on the cogency and value of those paradigms, even if they are doing little more than creating more work for themselves. In combination then, I think the above points support the case for advocates of meditative training in science to channel their efforts principally into establishing narrow forms of training.

¹⁸ One might extend such an argument by questioning the validity of *general metrics* of introspective proficiency, that err in favour of the broader path. For example, Fox et al.'s (2012) study of correlations between subjective and objective measures of tactile intensity showed *no improvements at all* in novice meditators (<28hrs). Improvements were only seen in experts (>7000hrs) who'd had contextual supports in place. However, attempts to extrapolate from this single measure to conclusions about the importance of contextual supports for *general* introspective proficiency may well be problematic.

This being said, I'm unpersuaded that the broader path should be dismissed off-hand. Encouraging at least some uptake of this different approach ensures that the full range of introspective skills available through Buddhist methods can help guide scientific research. One might argue that this could be achieved by simply adverting to external introspective authorities, as has been done in some prior neurophenomenological experiments (e.g. Lutz et al. 2002). Yet, one should again emphasise the novelty of the context here - meditation is being employed for novel questions in science, and important insights may be lost if scientists rely solely on existing and external introspective authorities, who have (in the Buddhist case) been guided only towards very specific aspects of the mind.

Additionally, I think that enthusiasts for the broader approach can question how much appeal for contemplative methods will be lost by replicating some of their traditional breadth. Though it's unequivocal that broader methods will find *less* appeal amongst the scientific community, one should be careful not to be bluntly dismissive of potential interest. For one thing, this underestimates the level of fervour that many have for their research enterprises.

Take perhaps the most prominent exemplar of this more immersive approach, Alan Wallace, who has set forth a manifesto for a 'contemplative science' (2009, 2012), combining first and third-person methods to tackle important questions about the aetiology and treatment of human suffering. Wallace's own research practice is informed by extensive experience in diverse contemplative methods, mirroring others like Francisco Varela and Evan Thompson, who have produced significant empirical and theoretical work off the back of personal immersion in the contemplative traditions. The presence of such figures in the scientific community suggests that swift dismissal of the broader path's appeal might be short-sighted. It may fail to appreciate what a future science of mind *could* look like. And greater awareness

of the importance of meditation's contextual supports might only increase the number of those interested in engaging such things.

In making this point though, one must be careful to acknowledge that some aspects of the traditional supportive web seem largely anathema to the scientific mindset, placing fundamental limitations upon meditation's possible integration. Take, for instance, the attitude of extreme worldly pessimism or even depression that Obeyesekere (1985) argues best represents the central mindset held to support meditation in Sri Lankan Buddhism. This is a long way from the chipper, empirical outlook that the Buddha is sometimes argued to have endorsed, itself an interpretation heavily criticised by Lopez (2012), but often used to reconcile the Buddhist and scientific enterprises. Or consider the commitment to re-incarnation in the supportive philosophical framework delineated in canonical Buddhist texts. It's highly unlikely that either of these two things will find fertile soil in science.¹⁹

Still, one can also counter-balance this difficulty by mirroring existing attempts to secularise aspects of Buddhist practice. Advocates for the broader path might claim that it won't be necessary to *directly mirror* many traditional supportive factors. In fact, it is likely possible to retain many of their essential supportive mechanisms whilst stripping them of aspects problematic to contemporary minds. For instance, I'm inclined to think that commitments to re-incarnation serve no essential supportive function in the meditative enterprise. What is important is a certain systematic intellectual framework that serves to temper the cravings one faces in practice. This might (though needn't) involve a revisionary interpretation of re-incarnation doctrine, which stresses the interdependence of persons and the inescapable impact of one actions upon one's future and the future of others.

¹⁹ Thanks to Ian Kidd (personal correspondence) for prompting these two points.

The precise level of advocacy one should engage for the broader path depends upon how plausible this secularisation process is, which can be better ascertained by examining existing attempts to do so in psychotherapy and indeed theoretical work on the plausibility of what Flanagan (2011) calls “naturalising” Buddhism. Yet, I think it holds sufficient promise as to bolster the idea that traditional supportive factors to meditation may indeed be replicable in science without too drastic a drop in their appeal. In tandem then, the plausibility of these approaches finding some fertile ground in science, as well as the importance of retaining the full spectrum of introspective benefits available through Buddhist methods, makes it important to retain a stream of advocacy for broader paths too. It would be both over pessimistic and short-sighted to dismiss such approaches completely.

To summarise then, in this section I’ve recommended that narrow forms of training be favoured in the push for contemplatively-inspired introspective development within science. Yet, I’ve also argued that one should retain a place for the advocacy of broader methods, encouraging awareness of the significance of contextual supports and the fact that narrow training will have different consequences and different kinds of benefit to more traditional forms. This is a pragmatic strategy, sensitive to the current climate in cognitive science, but also respectful of the possible directions that cognitive science might take in the future.

5 Conclusions and Routes Forward

In this paper, I’ve argued that proponents for the methodological employment of meditation in science, as a form of researcher training, should be aware of a number of contextual factors conventionally held up as important supports. I’ve motivated this by drawing attention to an analogous situation meeting proponents for meditation in psychotherapy, where it’s increasingly recognised that broader aspects of Buddhist practice play an important supportive

role in meditation's therapeutic function. In fact, I've motivated thinking that these two enterprises are much closer together than is typically acknowledged.

In many ways, the current conception of Buddhist contemplative practice amongst cognitive scientists resembles that popular in the Indian Mahāmudrā traditions. These favoured drastically streamlined conceptions of the Path, emphasising only a handful of basic meditative practices at the expense of the broader ethical and intellectual web articulated by prior Buddhist schools. Dunne (2011) notes how later Tibetan Mahāmudrā teachers made a deliberate attempt to remind people of context in their teachings, out of concern that important factors in the meditative enterprise were being forgotten (pp. 84-86). I hope to have done something of the same here. A similar process of reconsidering these factors has already begun in psychotherapeutic communities, and cognitive science should follow suit if it is to maximise the benefits available to it.

Having reviewed the character of these contextual supports, I've delineated two possible responses to this review. One emphasises broader regimes of training, replicating traditional supportive factors for maximum introspective proficiency. The other emphasises the accessibility of meditative training and is content to acknowledge more restricted benefits to individuals taking them up. I've drawn out plausible benefits of each form, before coming down principally in favour of the narrow option.

In making this suggestion though, it's important to acknowledge that the issues are not clear-cut. While I've favoured greater emphasis upon narrower forms of training, uncertainty over the plausibility of "secularising" the broader path leaves it unclear precisely how much that path should be promoted, and future work on this can help to fine-tune the balance of training regimes appropriately. Of course, this does not mean we should hold back in trying either. As with all attempts to institutionalise new methodological orientations to research, they

shall begin life insufficiently justified and requiring constant feedback for their optimisation – this is how progress is made. As with any research paradigm, the proof of the pudding will be in the eating. So, we’d do well to pay careful attention to the way this one develops. Keeping a finger on the pulse, sensitive to precisely how much empirical or theoretical success stems from individual projects, and how much of the broader path is required or encouraged in each such venture, the methodological value of these different paths can be played out in their explanatory and predictive fruits.²⁰ We transform the question into an empirical one.²¹

It is perhaps not important to replicate all components of the path to secure its supportive function, especially in unfamiliar territory where it is being employed for novel purposes. So, paying attention to how the employment of meditation proceeds can help us work out *which* bits of the path are most critical. Can one dial back some aspects of the broader path and get similar or identical benefits? Is there a minimum set of practice needed to reach a relative “expertise” beyond which advantage drops sharply or completely? Is there a *threshold level* of engagement in the broader path for securing introspective benefits?²²

As I’ve indicated, to some extent this experiment is already in progress, with a steadily growing stream of researchers beginning to implement contemplative practice into their lives and work. Moreover, I’ve argued in favour of at least *some* attempts to replicate semblances of the broader path of meditative supports in continued enterprises of this sort. As such, I end here by remarking on one factor that I think will be critical in determining how successful these broader enterprises can be, knowledge of which the enthusiast for the broader path can utilise

²⁰ Thanks for Scott Sturgeon (personal correspondence) for these points.

²¹ To see an example of this strategy, see Levit-Binun and Tarrasch’s (2014) preliminary study examining the effects of integrating meditative methods on student enthusiasm and progress in undergraduate neuroscience courses.

²² One might also seek answers to these questions by looking to parallel attempts to ascertain the importance of the specific contextual supports within psychotherapy.

to maximise that success. This concerns something much more fundamental, already alluded to by those favouring more cautious approaches: the *ultimate motivations* of those involved.

As noted in §4.2, the broader path is likely to find most appeal and commitment amongst those who resonate with its broader soteriological aim – the reduction of suffering for self and other. The degree to which such practices can efficaciously inform cognitive science will thus depend upon the personal motivations of the individuals making up the discipline. And though it is right to stress this as a reason for hesitation over the promotion of broader paths, Kordeš and Markič (2016) rightly caution that it would be unfair to hold soteriological orientations necessarily orthogonal to the scientific enterprise.

They note that such beliefs stem from an ‘idealized’ vision of science, in which research springs from ‘pure curiosity’. In contrast, they suggest that many researchers are driven by the desire to produce *useful* knowledge and also ‘existential commitments’ (pp. 164-165). The latter concerns deep-rooted authenticity considerations, through which we seek to act (and cultivate ourselves to act) in ways which accord with our own nature. In this way, research in cognitive science can be construed as a kind of self-exploration, through which we come to understand ourselves in ways which help us live. In respect of this, I conclude by suggesting that enthusiasts for the broader path would do well to put concerted effort into making the scientific community more aware of its reasons for research. It seems important to remind that community of the above forms of motivation, which might actually complement and support, rather than simply compete with, more epistemological motives. Herein, there is no unattractive Buddhist proselytizing. There is only a reminder that *many incentives spur good research* and an emphasis upon a source of motivation that some may have unwillingly detached from and wish to re-connect with. Through such prompts, we open a position in motivational space wherein the practices of the contemplative traditions might find broader appeal, where they

might be engaged for the right reasons, and where their value to cognitive science can accordingly be maximised. This seems a valuable attempt to drive cognitive-science forwards. And it's worth remembering with Colombetti (2014, p. 158) that, even if it ultimately comes up short, it is likely to have good side-effects.

References

- Ainsworth, B., Eddershaw, R., Meron, D., Baldwin, D.S. and Garner, M. (2013) The effect of focused attention and open monitoring meditation on attention network function in healthy volunteers. *Psychiatry Research*, 210 (3), pp. 1226-1231.
- Albahari, M. (2006) *Analytical Buddhism: The Two-Tiered Illusion of Self*. Basingstoke, UK: Palgrave Macmillan.
- Albers, S. (2012) *Eating Mindfully*. Oakland, CA: New Harbinger Publications.
- Armstrong, D.M. (1980) *The Nature of Mind and Other Essays*. Ithaca, NY: Cornell University Press.
- Baluch, F. and Itti, L. (2010) Training top-down attention improves performance on a triple-conjunction search task. *PloS one*, 5 (2), e9127. doi: 10.1371/journal.pone.0009127.
- Barandiaran, X.E. (2017) Autonomy and enactivism: Towards a theory of sensorimotor autonomous agency. *Topoi*, 36 (3), pp. 409-430.
- Barrouillet, P., Uittenhove, K., Lucidi, A. and Langerock, N. (2017) On the sources of forgetting in working memory: The test of competing hypotheses. *The Quarterly Journal of Experimental Psychology*, 71 (8), pp. 1714-1733.
- Bayne, T. (2004) Closing the gap? Some questions for neurophenomenology. *Phenomenology and the Cognitive Sciences*, 3 (4), pp. 349-364.
- Bayne, T. (2015) 'Introspective Insecurity', in T. Metzinger and J. M. Windt (eds) *Open MIND: 3(T)*, Frankfurt am Main: MIND Group. doi: 10.15502/9783958570214.
- Bayne, T. and Spener, M. (2010) Introspective humility. *Philosophical Issues*, 20 (1), pp. 1-22.
- Bayne, T., Hohwy, J. and Owen, A.M. (2016) Are there levels of consciousness? *Trends in Cognitive Sciences*, 20 (6), pp. 405-413.
- Bays, J. (2009) *Mindful Eating: A Guide to Rediscovering a Healthy and Joyful Relationship with Food*. Boston: Shambhala Publications.
- Bermúdez, J.L. (2006) *Philosophy of Psychology: Contemporary Readings*. New York: Routledge.
- Bickle, J. (1992) Mental Anomaly and the New Mind-Brain Reductionism. *Philosophy of Science*, 59, pp. 217-230.

- Bickle, J. (1998) *Psychoneural Reduction: The New Wave*. Cambridge, MA: MIT Press.
- Bitbol, M. and Petitmengin, C. (2013) A Defense of Introspection From Within. *Constructivist Foundations*, 8 (3), pp. 269-279.
- Bitbol, M. and Petitmengin, C. (2016) On the possibility and reality of introspection. *Mind and Matter*, 14 (1), pp. 51-75.
- Bitbol M. and Petitmengin C. (2017) 'Neurophenomenology and the microphenomenological interview', in Velmans, M. and Schneider, S. (eds) *The Blackwell Companion to Consciousness*. 2nd Edition. Chichester: Wiley & Sons, pp. 726–739.
- Block, N. (1990) 'Inverted Earth', in Tomberlin, J. (ed) *Philosophical Perspectives*, 4. Atascadero, CA: Ridgeview, pp. 53-79.
- Block, N. (2005) Two neural correlates of consciousness. *Trends in Cognitive Sciences*, 9 (2), pp. 46-52.
- Block, N. (2007) Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30 (5-6), pp. 481-499.
- Block, N. (2011) Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, 15 (12), pp. 567-575.
- Bode, B.H. (1913) The method of introspection. *The Journal of Philosophy, Psychology and Scientific Methods*, 10 (4), pp. 85-91.
- Bodhi, B. (2000) *The Connected Discourses of the Buddha: A Translation of the Samyutta Nikaya*. Somerville, MA: Wisdom Publications.
- Bodhi, B. (2005) *In the Buddha's Words: An Anthology of Discourses from the Pali Canon*. Somerville, MA: Wisdom Publications.
- Boring, E. (1953) A history of introspection. *Psychological Bulletin*, 50 (3), pp. 169-189.
- Bosse, T., Jonker, C.M. and Treur, J. (2008) Formalisation of Damasio's theory of emotion, feeling and core consciousness. *Consciousness and Cognition*, 17 (1), pp. 94-113.
- Braddock, G. (2001) Beyond reflection in naturalized phenomenology. *Journal of Consciousness Studies*, 8, pp. 3-16.
- Brentano F. (1874/1995) *Psychology from an Empirical Viewpoint*. Rancurello, A.C., Terrell, T.D., and McAlister, L.L. (trans.) London: Routledge.
- Brewer, J.A., Davis, J.H. and Goldstein, J. (2013) Why is it so hard to pay attention, or is it? Mindfulness, the factors of awakening and reward-based learning. *Mindfulness*, 4 (1), pp. 75-80.

- Britton, W.B., Lindahl, J.R., Cahn, B.R., Davis, J.H. and Goldman, R.E. (2014) Awakening is not a metaphor: the effects of Buddhist meditation practices on basic wakefulness. *Annals of the New York Academy of Sciences*, 1307 (1), pp. 64-81.
- Brogaard, B. and Chomanski, B. (2015) Cognitive penetrability and high-level properties in perception: Unrelated phenomena? *Pacific Philosophical Quarterly*, 96 (4), pp. 469-486.
- Bronkhorst, J. (1993) *The Two Traditions of Meditation in Ancient India*. Delhi: Motilal Banarsidass.
- Brownstein, M. and Saul, J.M. (eds) (2016) *Implicit Bias and Philosophy* (Vol. 2). Oxford: Oxford University Press.
- Bruineberg, J. and Rietveld, E. (2014) Self-organization, free energy minimization, and optimal grip on a field of affordances. *Frontiers in Human Neuroscience*, 8, 599, pp. 1-14.
- Buddharakkhita, A. (1995) Metta: The Philosophy and Practice of Universal Love, *Access to Insight (BCBS Edition)* Available at <http://www.accesstoinsight.org/lib/authors/buddharakkhita/wheel365.html> (accessed 19 Aug 2018).
- Carrasco, M., Fuller, S. and Ling, S. (2008) Transient attention does increase perceived contrast of suprathreshold stimuli: A reply to Prinzmetal, Long, and Leonhardt. *Perception and Psychophysics*, 70 (7), pp. 1151-1164.
- Carter, O.L., Presti D.E., Callistemon C., Ungerer Y., Liu G.B. and Pettigrew, J.D. (2005) Meditation alters perceptual rivalry in Tibetan Buddhist monks. *Current Biology* 15, no. 11, R412-R413.
- Chalmers, D. (1999) First-person methods in the science of consciousness. *Consciousness Bulletin*, Fall, pp. 8-11.
- Chalmers, D. (2004) How can we construct a science of consciousness? *The Cognitive Neurosciences III*, pp. 1111-1119.
- Chalmers, D. (2010) *The Character of Consciousness*. New York: Oxford University Press.
- Chin J.M. and Schooler J.W. (2009) 'Meta-Awareness', in W.P. Banks (ed) *Encyclopedia of Consciousness (Volume 2)*. Oxford: Elsevier, pp. 33-41.
- Churchland, P.S. (1989) *Neurophilosophy: Toward a Unified Science of the Mind-brain*. Cambridge, MA: MIT Press.
- Cohen, M.A. and Dennett, D.C. (2011) Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, 15 (8), pp. 358-364.

- Colombetti, G. (2009) What language does to feelings. *Journal of Consciousness Studies*, 16 (9), pp. 4-26.
- Colombetti, G. (2014) *The Feeling Body: Affective Science meets the Enactive Mind*. Cambridge, MA: MIT Press.
- Corbetta, M. and Shulman, G.L. (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3 (3), pp. 201-215.
- Coseru, C. (2009) Naturalism and intentionality: A Buddhist epistemological approach. *Asian Philosophy*, 19 (3), pp. 239-264.
- Cousins, L.S. (1996) 'The origins of insight meditation', in Skorupski, T. (ed) *The Buddhist Forum*, Vol. 4, 1994-1996, Tring, UK: The Institute of Buddhist Studies, pp. 35-58.
- Crane, T. (2000) Introspection, intentionality, and the transparency of experience. *Philosophical Topics*, 28 (2), pp. 49-67.
- Cummins, R.C. (1983) *The Nature of Psychological Explanation*. Cambridge, MA: MIT Press.
- Dahlstrom, D.O. (2015) 'Interoception and self-awareness: An exploration of interoceptive phenomenology', in Dahlstrom, D.O., Elpidorou A. and Hopp W. (eds) *Philosophy of Mind and Phenomenology: Conceptual and Empirical Approaches*. New York: Routledge, pp. 141-164.
- Damasio, A.R. (1999) *The Feeling of What Happens: Body and Emotion in the Making of Consciousness*. New York: Harcourt.
- Damasio, A.R. (2003) *Looking for Spinoza: Joy, Sorrow, and the Feeling Brain*. New York: Harcourt.
- Danziger, K. (1980) Wundt's psychological experiment in the light of his philosophy of science. *Psychological Research*, 42 (1-2), pp. 109-122.
- Davidson, D. (1987) Knowing One's Own Mind. *Proceedings of the American Philosophical Association*, 61, pp. 441-458.
- Davies, M. (2008) 'Consciousness and explanation', in Weiskrantz, L. and Davies, M. (eds) *Frontiers of Consciousness*. Oxford: Oxford University Press, pp. 1-54.
- Davis, J.H. and Thompson, E. (2013) 'From the five aggregates to phenomenal consciousness: Toward a cross-cultural cognitive science', in Emmanuel, S.M. (ed) *A Companion to Buddhist Philosophy*. Chichester, UK: Wiley, pp. 585-597.

- Davis, J.H. and Thompson, E. (2015) 'Developing attention and decreasing affective bias', in Brown, K.W., Cresswell, J.D. and Ryan, R.M. (eds) *Handbook of Mindfulness. Theory, Research, and Practice*. New York: Guildford, pp. 42-62.
- De Jesus, P. (2016) Autopoietic enactivism, phenomenology and the deep continuity between life and mind. *Phenomenology and the Cognitive Sciences*, 15 (2), pp. 265-289.
- Degenaar, J. and O'Regan, J.K. (2015) Sensorimotor theory and enactivism. *Topoi*, doi: 10.1007/s11245-015-9338-z.
- Dennett, D.C. (1993) *Consciousness Explained*. London: Penguin.
- Depraz, N., Varela, F.J. and Vermersch, P. (2003) *On Becoming Aware: A Pragmatics of Experiencing* (Vol. 43). Amsterdam: John Benjamins Publishing.
- Desbordes, G. and Negi, L.T. (2013) A new era for mind studies: training investigators in both scientific and contemplative methods of inquiry. *Frontiers in Human Neuroscience*, 7, 741, doi: 10.3389/fnhum.2013.00741.
- Desimone, R. and Duncan, J. (1995) Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18 (1), pp. 193-222.
- Di Paolo, E. (2009) Extended life. *Topoi*, 28 (1), pp. 9-21.
- Di Paolo, E., Rohde, M. and De Jaegher, H. (2014) 'Horizons for the enactive mind', in Stewart, J. Gapenne, O. and Di Paolo, E. (eds) *Enaction: Toward a New Paradigm for Cognitive Science*. Cambridge MA: MIT Press, pp. 33-87.
- Ditto, B., Eclache, M. and Goldman, N. (2006) Short-term autonomic and cardiovascular effects of mindfulness body scan meditation. *Annals of Behavioral Medicine*, 32 (3), pp. 227-234.
- Dor-Ziderman, Y., Berkovich-Ohana, A., Glicksohn, J. and Goldstein, A. (2013) Mindfulness-induced selflessness: a MEG neurophenomenological study. *Frontiers in Human Neuroscience*, 7, 582, doi: 10.3389/fnhum.2013.00582.
- Drayson, Z. (2012) The uses and abuses of the personal/subpersonal distinction. *Philosophical Perspectives*, 26 (1), pp. 1-18.
- Dretske, F. (1995) *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- Dreyfus, G. (2011) Is mindfulness present-centred and non-judgmental? A discussion of the cognitive dimensions of mindfulness. *Contemporary Buddhism*, 12 (1), pp. 41-54.
- Dreyfus, H.L. (1993) Review of 'The Embodied Mind'. *Mind*, 102 (407), pp. 542-546.

- Duncan, J. and Humphreys, G.W. (1989) Visual search and stimulus similarity. *Psychological review*, 96 (3), p. 433-458.
- Dunne, J. (2011) Toward an understanding of non-dual mindfulness. *Contemporary Buddhism*, 12 (1), pp. 71-88.
- Edelman, G. M. and Tononi, G. (2000) *A Universe of Consciousness: How Matter Becomes Imagination*. New York, NY: Basic Books.
- Edelman, G.M., Gally, J.A. and Baars, B.J. (2011) Biology of consciousness. *Frontiers in Psychology*, 2, doi: 10.3389/fpsyg.2011.00004.
- Ericsson, K.A. and Kintsch, W. (1995) Long-term working memory. *Psychological Review*, 102 (2), p. 211-245.
- Farb, N., Daubenmier, J., Price, C.J., Gard, T., Kerr, C., Dunn, B.D., Klein, A.C., Paulus, M.P. and Mehling, W.E. (2015) Interoception, contemplative practice, and health. *Frontiers in Psychology*, 6, 763, doi: 10.3389/fpsyg.2015.00763.
- Farb, N.A., Segal, Z.V. and Anderson, A.K. (2013) Mindfulness meditation training alters cortical representations of interoceptive attention. *Social Cognitive and Affective Neuroscience*, 8 (1), pp. 15-26.
- Farb, N.A., Segal, Z.V. and Anderson, A.K. (2013) Mindfulness meditation training alters cortical representations of interoceptive attention. *Social cognitive and Affective Neuroscience*, 8 (1), pp. 15-26.
- Farias, M. and Wikholm, C. (2015) *The Buddha Pill: Can Meditation Actually Change You?*. London: Watkins Media Limited.
- Ferree, C.E. and Rand, G. (1919) Chromatic thresholds of sensation from center to periphery of the retina and their bearing on color theory. *Psychological Review*, 26, pp. 16-41.
- Firth, R. (1949) Sense-data and the Percept Theory. *Mind*, 58 (232), pp. 434-465.
- Flanagan, O. (2011) *The Bodhisattva's Brain: Buddhism Naturalized*. Cambridge, MA: MIT Press.
- Fockert, J.W. de., Rees, G., Frith, C.D. and Lavie, N. (2001) The role of working memory in visual selective attention. *Science*, 291 (5509), pp. 1803-1806.
- Fougnie, D. (2008) 'The Relationship between Attention and Working Memory', in Johansen, N.B. (ed) *New Research on Short-Term Memory*. New York: Nova Science Publishers, pp. 1-45.

- Fox, K.C., Zakarauskas, P., Dixon, M., Ellamil, M., Thompson, E. and Christoff, K. (2012) Meditation experience predicts introspective accuracy. *PloS one*, 7(9): e45370, pp. 1-9. doi: 10.1371/journal.pone.0045370.
- Frith, C. (2002) How can we share experiences? *Trends in Cognitive Sciences*, 6 (9), p. 374.
- Froese, T. (2014) Review of ‘Radicalizing Enactivism: Basic Minds Without Content’. *The Journal of Mind and Behavior*, 173 (4), pp. 466-500.
- Froese, T., Gould, C. and Barrett, A. (2011) Re-viewing from Within: A Commentary on First-and-Second-Person Methods in the Science of Consciousness. *Constructivist Foundations*, 6 (2) pp. 254-269.
- Froese, T., Gould, C. and Seth, A.K. (2011) Validating and calibrating first-and second-person methods in the science of consciousness. *Journal of Consciousness Studies*, 18 (2), pp. 38-64.
- Fronsdal, G. (2005) Hindrances to Clear Seeing, *Insight Meditation Centre*, Available at <https://www.insightmeditationcenter.org/books-articles/articles/the-five-hindrances-handouts/hindrances-to-clear-seeing/> (accessed 14 Sep 2018)
- Gallagher, S. (1997) Mutual enlightenment: Recent phenomenology in cognitive science. *Journal of Consciousness Studies*, 4 (3), pp. 195-214.
- Gallagher, S. (2005) Phenomenological contributions to a theory of social cognition. *Husserl Studies*, 21 (2), pp. 95-110.
- Gallagher, S. (2010) ‘Phenomenology and non-reductionist cognitive science’, in Gallagher S. and Schmicking D. (eds) *Handbook of Phenomenology and Cognitive Science*. Berlin: Springer, pp. 21-34.
- Gallagher, S. and Francesconi, D. (2012) Teaching phenomenology to qualitative researchers, cognitive scientists, and phenomenologists. *Indo-Pacific Journal of Phenomenology*, 12 (sup3), pp. 1-10.
- Gallagher, S. and Zahavi, D. (2008) *The Phenomenological Mind*. New York: Routledge.
- Garrison, K., Santoyo, J., Davis, J., Thornhill, T., Kerr, C. and Brewer, J. (2013a) Effortless awareness: using real time neurofeedback to investigate correlates of posterior cingulate cortex activity in meditators' self-report. *Frontiers in Human Neuroscience*, 7, 440, doi: 10.3389/fnhum.2013.00440.
- Garrison, K., Scheinost, D., Worhunsky, P.D., Elwafi, H.M., Thornhill, T.A., Thompson, E., Saron, C., Desbordes, G., Kober, H., Hampson, M. and Gray, J.R., (2013b). Real-time

fMRI links subjective experience with brain activity during focused attention. *Neuroimage*, 81, pp. 110-118.

- Gawronski, B., Hofmann, W. and Wilbur, C.J. (2006) Are “implicit” attitudes unconscious? *Consciousness and Cognition*, 15 (3), pp. 485-499.
- Gethin, R. (1998) *The Foundations of Buddhism*. Oxford: Oxford University Press.
- Gethin, R. (2004) ‘On the Practice of Buddhist Meditation According to the Pali Nikāyas and Exegetical Sources’, in Bechert, H. and Gombrich, R. (eds) *Buddhismus in Geschichte und Gegenwart* 9. Hamburg: University of Hamburg, pp. 201-221.
- Gethin, R. (2015) ‘Buddhist Conceptualizations of Mindfulness’, in Brown, K.W., Cresswell, J.D. and Ryan R.M. (eds) *Handbook of Mindfulness: Theory, Research, and Practice*. New York: Guildford Press, pp. 9-41.
- Gillett, C. (2016) ‘The metaphysics of science, nature and the rules of engagement’, in Aizawa, K. and Gillett, C. (eds) *Scientific Composition and Metaphysical Ground*. London: Palgrave MacMillan, pp. 205-248.
- Godfrey, K.M., Gallo, L.C. and Afari, N. (2015) Mindfulness-based interventions for binge eating: a systematic review and meta-analysis. *Journal of Behavioral Medicine*, 38 (2), pp. 348-362.
- Goldie, P. (2002) Emotions, feelings and intentionality. *Phenomenology and the Cognitive Sciences*, 1 (3), pp. 235-254.
- Goldman, A. (2004) ‘Epistemology and the Evidential Status of Introspective Reports’, in Jack, A. and Roepstorff, A. (eds) *Trusting The Subject? Vol. 2*. Thorverton, UK: Imprint Academic, pp. 1-16.
- Goldstone, R.L. (1998) Perceptual learning. *Annual Review of Psychology*, 49 (1), pp. 585-612.
- Greenwald, A.G. and Krieger, L.H. (2006) Implicit bias: Scientific foundations. *California Law Review*, 94 (4), pp. 945-967.
- Grice, H.P. (2002) ‘Some remarks about the senses’, in Noë, A. and Thompson, E. (eds) *Vision and Mind: Readings in the Philosophy of Perception*, Cambridge, MA: MIT Press, pp. 34-55.
- Griffiths, P. (1981) Concentration or insight: The problematic of Theravāda Buddhist meditation-theory. *Journal of the American Academy of Religion*, 49 (4), pp. 605-624.

- Grubert, A., Fahrenfort, J., Olivers, C.N. and Eimer, M. (2017) Rapid top-down control over template-guided attention shifts to multiple objects. *Neuroimage*, 146, pp. 843-858.
- Hacker, P. (2002) Is there anything it is like to be a bat? *Philosophy*, 77 (2), pp. 157-174.
- Hahn, A., Judd, C.M., Hirsh, H.K. and Blair, I.V. (2014) Awareness of implicit attitudes. *Journal of Experimental Psychology: General*, 143 (3), pp. 1369-1392.
- Harman, G. (1997) 'The intrinsic quality of experience', in Block, N.J., Flanagan, O.J., and Güzeldere, G. (eds) *The Nature of Consciousness: Philosophical Debates*. Cambridge, MA: MIT Press, pp. 663-676.
- Hatfield, G. (2005) 'Introspective evidence in psychology', in Achinstein P. (ed) *Scientific Evidence: Philosophical Theories and Applications*. Baltimore: John Hopkins University Press, pp. 259-286.
- Hill, D.M., Craighead, L.W. and Safer, D.L. (2011) Appetite-focused dialectical behavior therapy for the treatment of binge eating with purging: A preliminary trial. *International Journal of Eating Disorders*, 44 (3), pp. 249-261.
- Holroyd, J. (2012) Responsibility for implicit bias. *Journal of Social Philosophy*, 43 (3), pp. 274-306.
- Hölzel, B.K., Lazar, S.W., Gard, T., Schuman-Olivier, Z., Vago, D.R. and Ott, U. (2011) How does mindfulness meditation work? Proposing mechanisms of action from a conceptual and neural perspective. *Perspectives on Psychological Science*, 6 (6), pp. 537-559.
- Hong, P.Y., Lishner, D.A., and Han, K.H. (2014) Mindfulness and eating: An experiment examining the effect of mindful raisin eating on the enjoyment of sampled food. *Mindfulness*, 5 (1), pp. 80-87.
- Hurlburt, R.T. (2009) Iteratively apprehending pristine experience. *Journal of Consciousness Studies*, 16 (10-1), pp. 156-188.
- Hurlburt, R.T. and Akhter, S.A. (2006) The descriptive experience sampling method. *Phenomenology and the Cognitive Sciences*, 5 (3-4), pp. 271-301.
- Hurley, S.L. (1998) *Consciousness in Action*. Cambridge, MA: Harvard University Press.
- Hutto, D.D. and Myin, E. (2013) *Radicalizing Enactivism: Basic Minds Without Content*. Cambridge, MA: MIT Press
- Hutto, D.D. and Myin, E. (2017) *Evolving Enactivism: Basic Minds Meet Content*. Cambridge, MA: MIT Press.

- Huxter, M. (2017) 'Mindfulness and the Buddha's noble eightfold path', in Shonin, E., Van Gordon, W. and Singh, N.N. (eds) *Buddhist Foundations of Mindfulness*. Cham, CH: Springer, pp. 29-53.
- Jack, A.I. and Roepstorff, A. (2003) *Trusting the Subject I*. Special Issue of *Journal of Consciousness Studies*, 10, 9-10.
- Jack, A.I. and Roepstorff, A. (2004) *Trusting the Subject II*. Special Issue of *Journal of Consciousness Studies*, 11, 7-8.
- James W. (1890/2007) *The Principles of Psychology, Volume 1*. New York: Cosimo.
- Janzen, G. (2008) *The Reflexive Nature of Consciousness* (Vol. 72). Amsterdam: John Benjamins Publishing.
- Jaynes, J. (1976) *The Origin of Consciousness In the Breakdown of the Bicameral Mind*. Boston, MA: Houghton Mifflin.
- Jha, A.P., Krompinger, J. and Baime, M.J. (2007) Mindfulness training modifies subsystems of attention. *Cognitive, Affective and Behavioral Neuroscience*, 7 (2), pp. 109–119.
- Jones, S.R., Kerr, C.E., Wan, Q., Pritchett, D.L., Hämäläinen, M. and Moore, C.I. (2010) Cued spatial attention drives functionally relevant modulation of the mu rhythm in primary somatosensory cortex. *Journal of Neuroscience*, 30 (41), pp. 13760-13765.
- Kabat-Zinn, J. (2003) Mindfulness-based interventions in context: past, present, and future. *Clinical Psychology: Science and Practice*, 10 (2), pp. 144-156.
- Kabat-Zinn, J. (2011) Some reflections on the origins of MBSR, skillful means, and the trouble with maps. *Contemporary Buddhism*, 12 (01), pp. 281-306.
- Kandinsky, W. (1914/1977) *Concerning the Spiritual in Art*. New York: Dover Publications.
- Katsuki, F. and Constantinidis, C. (2014) Bottom-up and top-down attention: different processes and overlapping neural systems. *The Neuroscientist*, 20 (5), pp. 509-521.
- Katterman, S.N., Kleinman, B.M., Hood, M.M., Nackers, L.M. and Corsica, J.A. (2014) Mindfulness meditation as an intervention for binge eating, emotional eating, and weight loss: a systematic review. *Eating Behaviors*, 15 (2), pp. 197-204.
- Keng, S.L., Smoski, M.J. and Robins, C.J. (2011) Effects of mindfulness on psychological health: A review of empirical studies. *Clinical Psychology Review*, 31 (6), pp. 1041-1056.
- Kerr, C.E., Jones, S.R., Wan, Q., Pritchett, D.L., Wasserman, R.H., Wexler, A., Villanueva, J.J., Shaw, J.R., Lazar, S.W., Kaptchuk, T.J. and Littenberg, R. (2011) Effects of

mindfulness meditation training on anticipatory alpha modulation in primary somatosensory cortex. *Brain Research Bulletin*, 85 (3-4), pp. 96-103.

- Khoury, B., Lecomte, T., Fortin, G., Masse, M., Therien, P., Bouchard, V., Chapleau, M.A., Paquin, K. and Hofmann, S.G. (2013) Mindfulness-based therapy: a comprehensive meta-analysis. *Clinical Psychology Review*, 33 (6), pp. 763-771.
- Kiken, L.G., Garland, E.L., Bluth, K., Palsson, O.S. and Gaylord, S.A. (2015) From a state to a trait: trajectories of state mindfulness in meditation during intervention predict changes in trait mindfulness. *Personality and Individual Differences*, 81, pp. 41-46.
- Kind, A. (2013) 'The case against representationalism about moods', in Kriegel, U. (ed) *Current Controversies in the Philosophy of Mind*. New York: Routledge, pp. 113-134.
- Kirmayer, L.J. (2015) Mindfulness in cultural context. *Transcultural Psychiatry*, 52 (4), pp. 447-469.
- Kordeš U. and Markič O. (2016). Parallels between mindfulness and first-person research into consciousness. *Asian Studies*, 4 (2), pp. 153-168.
- Kordeš, U. and Demšar, E. (2018) Excavating belief about past experience: experiential dynamics of the reflective act. *Constructivist Foundations*, 13 (2), pp. 219-229.
- Kordeš, U. and Markič, O. (2016) Parallels between Mindfulness and First-person Research into Consciousness. *Asian Studies*, 4 (2), pp. 153-168.
- Korzybski, A. (1933/1993) *Science and Sanity: An Introduction to Non-Aristotelian Systems and General Semantics*, 5th Edition. Englewood, NJ: The International Non-Aristotelian Library/Institute of General Semantics.
- Kramer, G. and Bitbol, M. (2014) *Toward a Mindful Cognitive Science*. Available <http://www.mindfulnessitalia.it/sys/wp-content/uploads/2014/06/Towards-a-Mindful-Cognitive-Science-2014-final.pdf> (accessed 18 Sep 2018)
- Kriegel, U. (2015) *The Varieties of Consciousness*. Oxford: Oxford University Press.
- Kristeller, J.L., and Hallett, C.B. (1999) An exploratory study of a meditation-based intervention for binge eating disorder. *Journal of Health Psychology*, 4, pp. 357-363.
- Kristeller, J.L. and Wolever, R.Q. (2011) Mindfulness-based eating awareness training for treating binge eating disorder: the conceptual foundation. *Eating disorders*, 19 (1), pp. 49-61.

- Lemaire, B. and Portrat, S. (2018) A Computational Model of Working Memory Integrating Time-Based Decay and Interference. *Frontiers in Psychology*, 9, doi: 10.3389/fpsyg.2018.00416.
- Levine, J. (1983) Materialism and qualia: The explanatory gap. *Pacific philosophical quarterly*, 64 (4), pp. 354-361.
- Levit Binnun, N. and Tarrasch, R. (2014) Relation between contemplative exercises and an enriched psychology students' experience in a neuroscience course. *Frontiers in Psychology*, 5, 1296, pp. 1-16.
- Lippelt, D.P., Hommel, B. and Colzato, L.S. (2014) Focused attention, open monitoring and loving kindness meditation: Effects on attention, conflict monitoring, and creativity – A review. *Frontiers in Psychology*, 5, 1083, doi: 10.3389/fpsyg.2014.01083.
- Lopez, D.S. (2012) *The Scientific Buddha: His Short and Happy Life*. New Haven, CT: Yale University Press.
- Lusthaus, D. (2003) *Buddhist Phenomenology: A Philosophical Investigation of Yogacara Buddhism and the Ch'eng Wei-shih Lun*. New York: Routledge.
- Lutz, A. (2002) Toward a neurophenomenology as an account of generative passages: A first empirical case study. *Phenomenology and the Cognitive Sciences*, 1 (2), pp. 133-167.
- Lutz, A. and Thompson, E. (2003) Neurophenomenology integrating subjective experience and brain dynamics in the neuroscience of consciousness. *Journal of Consciousness Studies*, 10 (9-10), pp. 31-52.
- Lutz, A., Jha, A.P., Dunne, J.D. and Saron, C.D. (2015) Investigating the phenomenological matrix of mindfulness-related practices from a neurocognitive perspective. *American Psychologist*, 70 (7), pp. 632-658.
- Lutz, A., Lachaux, J.P., Martinerie, J. and Varela, F.J. (2002) Guiding the study of brain dynamics by using first-person data: synchrony patterns correlate with ongoing conscious states during a simple visual task. *Proceedings of the National Academy of Sciences*, 99 (3), pp. 1586-1591.
- Lutz, A., Slagter, H.A., Dunne, J.D. and Davidson, R.J. (2008) Attention regulation and monitoring in meditation. *Trends in Cognitive Sciences*, 12 (4), pp. 163-169.
- Lutz, A., Slagter, H.A., Rawlings, N.B., Francis, A.D., Greischar, L.L. and Davidson, R.J. (2009) Mental training enhances attentional stability: neural and behavioral evidence. *Journal of Neuroscience*, 29 (42), pp. 13418-13427.

- Mack, A. and Rock, I. (1998) *Inattentional Blindness*. Cambridge, MA: MIT press.
- MacLean, K.A., Ferrer, E., Aichele, S.R., Bridwell, D.A., Zanesco, A.P., Jacobs, T.L., King, B.G., Rosenberg, E.L., Sahdra, B.K., Shaver, P.R., Wallace, B.A., Mangun, G.R. and Saron, C.D. (2010) Intensive meditation training improves perceptual discrimination and sustained attention. *Psychological Science*, 21 (6), pp. 829–839.
- Markovich, J. and Thompson, E. (2016) ‘Hypnosis and meditation: a neurophenomenological comparison’, in Raz, A. and Lifshitz, M. (eds) *Hypnosis and Meditation: Towards an Integrative Science of Conscious Planes*. Oxford: Oxford University Press, pp. 79-106.
- McAuliffe, A. (2018) Concerns with the Validity and Practicality of First-Person Data: Where Do We Go from Here? *Constructivist Foundations*, 13 (2), pp. 238-239.
- McDowell, J. (1994) The content of perceptual experience. *The Philosophical Quarterly*, 44 (175), pp. 190-205.
- Menezes, C.B., Paula Couto, M.C. de., Buratto, L.G., Erthal, F., Pereira, M.G. and Bizarro, L. (2013) The improvement of emotion and attention regulation after a 6-week training of focused meditation: a randomized controlled trial. *Evid. Based Complement. Alternat. Med.* 2013, pp. 1–11, doi: 10.1155/2013/984678.
- Merleau-Ponty, M. (1945/1962) *Phenomenology of Perception*. Smith, C. (trans), London: Routledge Press.
- Metzinger, T. (2000) *Neural Correlates of Consciousness: Empirical and Conceptual Questions*. Cambridge, MA: MIT Press.
- Metzinger, T. (2016) *The Mental Autonomy Project*. Available at <https://www.blogs.uni-mainz.de/fb05philosophie/files/2013/04/MAP.pdf> (accessed 18 Sep 2018).
- Moreland, J.D., Jameson, D. and Hurvich, L.M. (1972) ‘Peripheral colour vision’, in Jameson, D. and Hurvich, L.M. (eds) *Handbook of Sensory Physiology: Vol. VII/4. Visual Psychophysics*. New York: Springer, pp. 517–536.
- Murphy, J. (2017) Review of ‘The Buddha Pill: Can Meditation Change You?’. *Secularism and Nonreligion*, 6, p. 7, doi: 10.5334/snr.81.
- Ñānananda, B. (1971/1997) *Concept and Reality in Early Buddhist Thought: An Essay on Papañca and Papañca-Saññā-Saṅkhā*. Kandy, Sri Lanka: Buddhist Publication Society.
- Nanay, B. (2009) Attention and perceptual content. *Analysis*, 70 (2), pp. 263-270.
- Natorp, P. (1912) *Allgemeine Psychologie*. Tübingen: J.C.B. Mohr.

- Nielsen, L. and Kaszniak, A.W. (2006) Awareness of subtle emotional feelings: a comparison of long-term meditators and nonmeditators. *Emotion*, 6, pp. 392-405.
- Noë, A. (2004) *Action in Perception*. Cambridge, MA: MIT Press.
- Noë, A. (2007) The critique of pure phenomenology. *Phenomenology and the Cognitive Sciences*, 6 (1-2), pp. 231-245.
- Noë, A. and O'Regan, J.K. (2002) 'On the brain-basis of visual consciousness: A sensorimotor account', in Noë, A. and Thompson, E. (eds) *Vision and Mind: Selected Readings in the Philosophy of Perception*. Cambridge, MA: MIT Press, pp. 567-598.
- Noë, A., Pessoa, L. and Thompson, E. (2000) Beyond the grand illusion: What change blindness really teaches us about vision. *Visual Cognition*, 7 (1-3), pp. 93-106.
- Nummenmaa, L., Glerean, E., Hari, R. and Hietanen, J.K. (2014) Bodily maps of emotions. *Proceedings of the National Academy of Sciences*, 111 (2), pp. 646-651.
- Nyanaponika, T. (1988) *The Heart of Buddhist Meditation*. York Beach, Maine: Samuel Weiser.
- Nyanaponika, T. (1994) *The Five Mental Hindrances and Their Conquest: Selected Texts from the Pali Canon and Commentaries*. Nyanaponika T. (trans), *Access to Insight (BCBS Edition)*, Available at <http://www.accesstoinsight.org/lib/authors/nyanaponika/wheel026.html> (accessed 18 Aug, 2018).
- Nyanaponika, T. (2015) *The Vision of Dhamma*. 2nd edition, eBook, Onalaska, WA: BPS Pariyatti Editions. Available at https://store.pariyatti.org/Vision-of-Dhamma--PDF-eBook_p_2571.html (Downloaded 15 Sep, 2018).
- O'Brien B. (2017) *The Mindfulness Controversy – Buddhism vs Psychology?* Available at <https://www.thoughtco.com/the-mindfulness-controversy-449769> (accessed 19 Aug 2018).
- Obeyesekere, G. (1985) 'Depression, Buddhism, and the work of culture in Sri Lanka', in Kleinman, A. and Good, B. (eds) *Culture and Depression*. Berkeley: University of California Press, pp. 134-152.
- Olivers, C.N., Peters, J., Houtkamp, R. and Roelfsema, P.R. (2011) Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*, 15 (7), pp. 327-334.

- O'Regan, J.K. (2011) *Why Red Doesn't Sound Like a Bell: Understanding the Feel of Consciousness*. New York: Oxford University Press.
- O'Regan, J.K. and Noë, A. (2001) A sensorimotor account of vision and visual consciousness. *Behavioral and brain sciences*, 24 (05), pp. 939-973.
- Overgaard, M., Gallagher, S. and Ramsøy, T.Z. (2008) Integration of first-person methodologies in cognitive science. *Journal of Consciousness Studies*, 15, 5, pp. 100-120.
- Ozawa-de Silva, B. (2016) Contemplative science and secular ethics. *Religions*, 7 (8), 98. doi: 10.3390/rel7080098.
- Papineau, D. (1998) Mind the gap. *Nous*, 32 (S12), pp. 373-388.
- Parvizi, J. and Damasio, A. (2001) Consciousness and the brainstem. *Cognition*, 79 (1-2), pp. 135-160.
- Patrik, L.E. (1994) Phenomenological method and meditation. *Journal of Transpersonal Psychology*, 26, pp. 37-54.
- Peacock, J. (2008) Suffering in Mind: The Aetiology of suffering in early Buddhism. *Contemporary Buddhism*, 9 (2), pp. 209-226.
- Pepper, K. (2014) Do sensorimotor dynamics extend the conscious mind? *Adaptive Behavior*, 22 (2), pp. 99-108.
- Petitmengin, C. (2006) Describing one's subjective experience in the second person: An interview method for the science of consciousness. *Phenomenology and the Cognitive sciences*, 5 (3), pp. 229-269.
- Petitmengin, C. (2007) Towards the source of thought: The gestural and transmodal dimension of lived experience. *Journal of Consciousness Studies*, 14, pp. 54-82.
- Petitmengin, C. (2009) The validity of first-person descriptions as authenticity and coherence. *Journal of Consciousness Studies*, 16 (10-1), pp. 252-284.
- Petitmengin, C. and Bitbol, M. (2009) Listening from within. *Journal of Consciousness studies*, 16 (10-1), pp. 363-404.
- Pinto, Y., Leij, A.R. van der., Sligte, I.G., Lamme, V.A. and Scholte, H.S. (2013) Bottom-up and top-down attention are independent. *Journal of Vision*, 13 (3), pp. 1-14.
- Pollatos, O. and Schandry, R. (2008) Emotional processing and emotional memory are modulated by interoceptive awareness. *Cognition and Emotion*, 22 (2), pp. 272-287.

- Prinz, J. (2004) *Gut Reactions: A Perceptual Theory of Emotion*. New York: Oxford University Press.
- Raz, A. and Buhle, J. (2006) Typologies of attentional networks. *Nature Reviews Neuroscience*, 7 (5), pp. 367-369.
- Reinerman-Jones, L., Sollins, B., Gallagher, S. and Janz, B. (2013) Neurophenomenology: an integrated approach to exploring awe and wonder. *South African Journal of Philosophy*, 32 (4), pp. 295-309.
- Roberts, M. (2018a) Meditation and introspection: Insight through transformation. [unpublished]
- Roberts, M. (2018b) Phenomenological constraints: A problem for radical enactivism. *Phenomenology and the Cognitive Sciences*, 17 (2), pp. 375-399.
- Robertson, I.H., Mattingley, J.B., Rorden, C. and Driver, J. (1998) Phasic alerting of neglect patients overcomes their spatial deficit in visual awareness. *Nature*, 395, pp. 169–172.
- Rosenthal, D.M. (2002) ‘Explaining consciousness’, in Chalmers, D.J. (ed) *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press, pp. 406-421.
- Rupert, R.D. (2009) *Cognitive Systems and the Extended Mind*. New York: Oxford University Press.
- Rupert, R.D. (2015) Embodiment, consciousness, and neurophenomenology: Embodied cognitive science puts the (first) person in its place. *Journal of Consciousness Studies*, 22 (3-4), pp. 148-180.
- Schaffner, K.F. (1993) *Discovery and Explanation in Biology and Medicine*. Chicago: University of Chicago Press.
- Schooler, J. and Schreiber, C.A. (2004) Experience, meta-consciousness, and the paradox of introspection. *Journal of Consciousness Studies*, 11 (7-8), pp. 17-39.
- Schubert, R., Blankenburg, F., Lemm, S., Villringer, A. and Curio, G. (2006) Now you feel it—now you don't: ERP correlates of somatosensory awareness. *Psychophysiology*, 43 (1), pp. 31-40.
- Schubert, R., Haufe, S., Blankenburg, F., Villringer, A. and Curio, G. (2009) Now you'll feel it, now you won't: EEG rhythms predict the effectiveness of perceptual masking. *Journal of Cognitive Neuroscience*, 21 (12), pp. 2407-2419.

- Schubert, R., Ritter, P., Wüstenberg, T., Preuschhof, C., Curio, G., Sommer, W. and Villringer, A. (2008) Spatial attention related SEP amplitude modulations covary with BOLD signal in S1—a simultaneous EEG—fMRI study. *Cerebral Cortex*, 18 (11), pp. 2686-2700.
- Schwitzgebel, E. (2004) Introspective training apprehensively defended: Reflections on Titchener's lab manual. *Journal of Consciousness Studies*, 11 (7-8), pp. 58-76.
- Schwitzgebel, E. (2008) The unreliability of naive introspection. *Philosophical Review*, 117 (2), pp. 245-273.
- Schwitzgebel, E. (2011) *Perplexities of Consciousness*. Cambridge, MA: MIT Press.
- Schwitzgebel, E. (2012) 'Introspection, what?', in Smithies, D. and Stoljar, D. (eds) *Introspection and Consciousness*. New York: Oxford University Press, pp. 29-48.
- Schwitzgebel, E. (2014) 'Introspection', in Zalta, N. (ed) *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Retrieved from <https://plato.stanford.edu/archives/win2016/entries/introspection/> (accessed 19 Aug 2018).
- Searle, J.R. (2000) Consciousness. *Annual Review of Neuroscience*, 23, pp. 557–578.
- Seth, A.K. (2013) Interoceptive inference, emotion, and the embodied self. *Trends in Cognitive Sciences*, 17 (11), pp. 565-573.
- Shankman, R. (2008) *The Experience of Samadhi: An In-Depth Exploration of Buddhist Meditation*. Boston, MA: Shambala.
- Sharf, R.H. (1995) Buddhist modernism and the rhetoric of meditative experience. *Numen*, 42 (3), pp. 228-283.
- Sharf, R.H. (2015) Is mindfulness Buddhist? (and why it matters). *Transcultural Psychiatry*, 52 (4), pp. 470-484.
- Shaw, S. (2014) *The Spirit of Buddhist Meditation*. New Haven: Yale University Press.
- Shear, J. and Varela, F.J. (1999) (eds) *The View from Within: First-person Approaches to The Study of Consciousness*. Thorverton, UK: Imprint Academic
- Siegel, S. (2006) 'Which properties are represented in perception?', in Gendler, T.S. and Hawthorne, J. (eds) *Perceptual Experience*. Oxford: Clarendon Press, pp. 481-503.
- Simons, D.J. and Chabris, C.F. (1999) Gorillas in our midst: Sustained inattention blindness for dynamic events. *Perception*, 28 (9), pp. 1059-1074.

- Spener, M. (2011) 'Disagreement about cognitive phenomenology', in Bayne, T. and Montague, M. (eds) *Cognitive Phenomenology*. Oxford: Oxford University Press, pp. 268-284.
- Spener, M. (2015) Calibrating Introspection. *Philosophical Issues*, 25 (1), pp. 300-321.
- Spener, M. (2018) 'Introspecting in the 20th Century', in Kind, A. (ed) *Philosophy of Mind in the Twentieth and Twenty-First Centuries*. London: Routledge, pp. 148-174.
- Stazicker, J. (2011) Attention, visual consciousness and indeterminacy. *Mind and Language*, 26 (2), pp. 156-184.
- Stewart, J. (2014) 'Foundational issues in enaction as a paradigm for cognitive science', in Stewart J., Gapenne, O. and Di Paolo, E. (eds) *Enaction: Toward a New Paradigm for Cognitive Science*. Cambridge MA: MIT Press, pp. 1-31.
- Stewart, J., Gapenne, O. and Di Paolo, E. (2014) 'Introduction', in Stewart, J., Gapenne, O. and Di Paolo, E. (eds) *Enaction: Toward a New Paradigm for Cognitive Science*. Cambridge, MA: MIT Press, pp. viii-xvii.
- Sturm, W. and Willmes, K. (2001) On the functional neuroanatomy of intrinsic and phasic alertness. *Neuroimage*, 14 (1), S76-S84.
- Sujīva, V. (2000) *Essentials of Insight Meditation Practice*. Revised Edition, eBook, Petalang Jaya, Malaysia: Buddhist Wisdom Center. Available at https://www.buddhanet.net/pdf_file/essentials.pdf (Downloaded 15 Sep 2018).
- Sze, J.A., Gyurak, A., Yuan, J.W. and Levenson, R.W. (2010) Coherence between emotional experience and physiology: does body awareness training have an impact? *Emotion*, 10 (6), pp. 803-814.
- Tang, Y.Y., Ma, Y., Wang, J., Fan, Y., Feng, S., Lu, Q., Yu, Q., Sui, D., Rothbart, M.K., Fan, M. and Posner, M.I. (2007) Short-term meditation training improves attention and self-regulation. *Proceedings of the National Academy of Sciences*, 104 (43), pp. 17152-17156.
- Taylor, H. (2013) Is attention necessary and sufficient for phenomenal consciousness? *Journal of Consciousness Studies*, 20 (11-12), pp. 173-194.
- Teasdale, J.D. and Chaskalson, M. (2011) How does mindfulness transform suffering? II: the transformation of dukkha. *Contemporary Buddhism*, 12 (01), pp. 103-124.

- Teper, R., Segal, Z.V. and Inzlicht, M. (2013) Inside the mindful mind how mindfulness enhances emotion regulation through improvements in executive control. *Current Directions in Psychological Science*, 22 (6), pp. 449-454.
- Thangru Rinpoche (1994) Huckenpahler, V. (ed), Stern, M. and Kunsang, E.P. Martin, M. (trans), *The Seven Points of Mind Training*. Auckland, NZ: Zhyisil Chokyi Ghatsal Publications. Available at <http://www.rinpoche.com/teachings/sevenpoints.htm> (accessed 18 Aug 2018).
- Thanissaro B. (trans.) (1993) Dhammacakkappavattana Sutta: Setting in Motion the Wheel of Truth (SN 56.11) *Access to Insight (BCBS Edition)*. Retrieved August 8, 2018. Available <http://www.accesstoinsight.org/tipitaka/sn/sn56/sn56.011.than.html>.
- Thiradhammo, A. (2014) *Working with the Five Hindrances*. Belsay, UK: Aruno Publications.
- Thompson, E. (2007) *Mind in Life*. Cambridge, MA: Harvard University Press.
- Thompson, E. (2009) 'Contemplative neuroscience as an approach to volitional consciousness', in Murphy N. and Brown, W. S. (eds) *Downward Causation and the Neurobiology of Free Will*. Berlin Heidelberg: Springer, pp. 187-197.
- Thompson, E. (2011) Precipice of mind in life. *Journal of Consciousness Studies*, 18 (5-6), pp. 10-22.
- Thompson, E. (2015) *Waking, Dreaming, Being: Self and Consciousness in Neuroscience, Meditation, and Philosophy*. New York: Columbia University Press.
- Thompson, E. and Cosmelli, D. (2011) Brain in a vat or body in a world? Brainbound versus enactive views of experience. *Philosophical Topics*, Vol. 39 (1), pp. 163-180.
- Thompson, E., Lutz, A. and Cosmelli, D. (2005) 'Neurophenomenology: An Introduction for Neurophilosophers', in Brook, A. and Akins, K. (eds) *Cognition and the Brain: The Philosophy and Neuroscience Movement*. New York: Cambridge University Press, pp. 40-97.
- Titchener, E.B. (1899) *An Outline of Psychology*. New York: Macmillan.
- Tononi, G. and Edelman, G.M. (1998) Consciousness and complexity. *Science*, 282, pp. 1846-51.
- Tye, M. (2009) *Consciousness Revisited: Materialism Without Phenomenal Concepts*. Cambridge: MIT Press.

- Vago, D.R. and Silbersweig, D.A. (2012) Self-awareness, self-regulation, and self-transcendence (S-ART): a framework for understanding the neurobiological mechanisms of mindfulness. *Frontiers in Human Neuroscience*, 6, pp. 1-30.
- Van Riel, R. and Van Gulick, R. (2016) Scientific Reduction. *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), Zalta, E.N. (ed) Available at <http://plato.stanford.edu/archives/spr2016/entries/scientific-reduction/> (accessed 20 Feb 2017).
- Varela, F.J. (1996) Neurophenomenology: A methodological remedy for the hard problem. *Journal of Consciousness Studies*, 3 (4), pp. 330-349.
- Varela, F.J. (1997) The naturalization of phenomenology as the transcendence of nature: Searching for generative mutual constraints. *Alter: Revue de Phenomenologie*, 5, pp. 355–385.
- Varela, F.J., Thompson, E. and Rosch, E. (1991/2017) *The Embodied Mind: Cognitive Science and Human Experience*. Revised Edition. Cambridge, MA: MIT Press.
- Wallace, B.A. (1999) The Buddhist tradition of Samatha: Methods for refining and examining consciousness. *Journal of Consciousness Studies*, 6 (2-3), pp. 175-187.
- Wallace, B.A. (2004) *The Four Immeasurables: Practices to Open the Heart*. New York: Snow Lion Publications.
- Wallace, B.A. (2009) *Contemplative Science: Where Buddhism and Neuroscience Converge*. New York: Columbia University Press.
- Wallace, B.A. (2012) *Meditations of a Buddhist Skeptic*. New York: Columbia University Press.
- Ward, D. (2012) Enjoying the spread: Conscious externalism reconsidered. *Mind*, 121 (483), pp. 731-751.
- Wheeler, M. (2005) *Reconstructing the Cognitive World: The Next Step*. Cambridge, MA: MIT press.
- Wheeler, M. (2011) Mind in life or life in mind? Making sense of deep continuity. *Journal of Consciousness Studies*, 18 (5-6), pp. 148-168.
- Wheeler, M. (2013) Science friction: Phenomenology, naturalism and cognitive science. *Royal Institute of Philosophy Supplement*, 72, pp. 135-167.
- Wheeler, M. (2015) The Revolution will not be Optimised: Radical Enactivism, Extended Functionalism and the Extensive Mind. *Topoi*, doi: 10.1007/s11245-015-9356-x.

- Whiting, D. (2011) The feeling theory of emotions and the object-directed emotions. *European Journal of Philosophy*, 19 (2), pp. 281–303.
- Wilks, J. (2014) Secular Mindfulness: Potential & Pitfalls. *Insight Journal*. Available at <https://www.buddhistinquiry.org/article/secular-mindfulness-potential-pitfalls/> (accessed 18 Aug 2018).
- Williams, J.M.G. and Kabat-Zinn, J. (eds) (2013) *Mindfulness: Diverse Perspectives on its Meaning, Origins and Applications*. New York: Routledge.
- Williams, P. (2008). *Mahāyāna Buddhism: The Doctrinal Foundations*. 2nd Edition. New York: Routledge.
- Williams, P. and Tribe, A. (2003) *Buddhist Thought: A Complete Introduction to the Indian Tradition*. London: Routledge.
- Wolfe, J.M. and Horowitz, T.S. (2004) What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, 5 (6), pp. 495-501.
- Wu, W. (2014) *Attention*. London: Routledge.
- Wundt W. (1897) *Outlines of Psychology*. Judd, C.H. (trans), Leipzig: Wilhelm Engelmann.
- Zahavi, D. (2003) How to investigate subjectivity: Natorp and Heidegger on reflection. *Continental Philosophy Review*, 36 (2), pp. 155-176.
- Zahavi, D. (2007) Killing the straw man: Dennett and phenomenology. *Phenomenology and the Cognitive Sciences*, 6 (1-2), pp. 21-43.
- Zahavi, D. (2014) *Self and Other: Exploring Subjectivity, Empathy and Shame*. Oxford: Oxford University Press.
- Zeimbekis, J. and Raftopoulos, A. (2015) *The Cognitive Penetrability of Perception: New Philosophical Perspectives*. Oxford: Oxford University Press.