HIDING IN PLAIN SIGHT

Challenges unconscious mentality places on self-knowledge

By Stephania Donayre Pimentel (2-year MA)

Advisor: Tim Crane, PhD

Submitted to Central European University

In partial fulfilment of the requirements for the degree of Master of Arts in Philosophy

Vienna, Austria

2023

ABSTRACT

Self-knowledge refers to knowledge of one's own mind: what we feel, think, believe, desire… Since Descartes, it is commonplace to take self-knowledge as categorically distinct from our knowledge of the outside world. This distinctiveness is often put in terms of "privileged access": I can get to know my own mind in a way which no one else can. Indeed, privileged access has also been championed as the "mark of the mental": the defining criterion of mental phenomena.

It is my view that privileged access is not the mark of the mental, but only the mark of the *conscious*. This is because, first of all, the vast extent of our cognitive and affective processing occurs with little or no conscious effort. Secondly, many quintessential mental phenomena do not depend upon our continued awareness: I can continue to believe and desire something even when I am not conscious of it.

If this is right, the notion of privileged access fails to capture all that there is to to know. In particular, it fails to recognise unconscious mental phenomena as proper objects of self-knowledge. Bringing unconscious mental phenomena into the fold calls for a critical reevaluation of the aims and methods of attaining self-knowledge. It is this project that I undertake in this dissertation.

In chapter 1, I begin by proposing a basic ontological and epistemic distinction between mental events and mental states. In chapter 2, I build upon this schema by clarifying the nature of unconscious mental states. I go about this by answering two questions: (1) What makes unconscious states *mental*? and (2) What makes mental

states *unconscious*? In chapter 3, I bring this framework to bear on a philosophical debate concerning whether self-knowledge is a matter of discovery or «making up one's mind» (interpretivism). While siding mostly with interpretivists such as Crane and Farkas, I caution that we should not overestimate our capacities for articulation. I claim that any faithful representation of our unconscious states must accommodate or negotiate their inherent uncertainty and contradiction, ceding terrain for the mindful wisdom of self-ignorance.

AUTHOR'S DECLARATION

I, the undersigned, Stephania Donayre Pimentel, candidate for the MA degree in Philosophy declare herewith that the present thesis is exclusively my own work, based on my research and only such external information as properly credited in notes and bibliography. I declare that no unidentified and illegitimate use was made of the work of others, and no part of the thesis infringes on any person's or institution's copyright. I also declare that no part of the thesis has been submitted in this form to any other institution of higher education for an academic degree.

Vienna, Austria

2023.12.13

ACKNOWLEDGEMENTS

To all our misunderstandings and the misguided attempts to solve them.

TABLE OF CONTENTS

0. INTRODUCTION

This dissertation is about how we bring unconscious states to awareness in a way that the prospects and limitations of self-knowledge are properly accounted for. In analytic philosophy of mind, 'self-knowledge' can be an unfortunate term, as it might mislead us into thinking that I am referring to knowledge of the self and its nature— but this is not what I am doing. As a matter of fact, one of the implications of my work might be that there is not a unitary self as habitually conceived. But that is an orthogonal issue; this work is not about the metaphysics and epistemology of personal identity.

The notion of self-knowledge employed on this thesis is knowledge of one's own mental events and states: what we are feeling or thinking, or what we believe or desire. In western philosophy, this is traditionally conceived to be distinctive, hence we have a special name for it (unlike the knowledge of one's body, for contrast). At least since Descartes, it is commonplace to take self-knowledge (i.e., the knowledge of our own mind) as categorically distinct from our knowledge of the outside world. This distinctiveness is often put in terms of "privileged access": an umbrella term for the notion that one's is in a relatively safe position of authority to acquire, via introspection, knowledge about one's mind.

Central to some of the implications of this work, the notion of privileged access has been put forward as the "mark of the mental". This *epistemic* conception of mind is fundamentally shaped by Cartesian philosophy: "The nature of the human

mind, and how it is better known than the body" (Descartes, 1641/2017). Katalin Farkas (2009) says: "Privileged accessibility is the mark of the mental: it enables the subject to know its *subject matter* in a way that no-one else who is endowed with the same cognitive faculty, can" (emphasis added). Self-knowledge has been then often construed in terms of the idea of privileged access. While there is a clear tendency to consider it as a special category of knowledge, however, there is not nearly as much clarity about the objects of this knowing. The following question naturally arises: *What is there to know (*about one's mind)*?*

My first task in chapter 1 is to lay an account for that question. As such, I think it's important to disambiguate the different ways one can be said to know one's mind. On one hand, I want to further secure the right to say that we have privileged access: I certainly get to know my excitement or pain in a way no-one else can (although others can certainly know it too). This refers to the knowledge, I will posit, of our *mental events*— which involve experiences that are present, occurrent, and conscious. As experiences involve an awareness of themselves, attention to them affords us privileged access. It is in this very restricted or bracketed sense that I would characterise privileged access as the mark of the *conscious* (not the mental!): we can employ the metaphysics of conscious experiences to explain its first-person epistemology. Privileged access is, indeed, the mark of the conscious.

But in this work, I insist, as many philosophers do, that we have mental features beyond those that pass through the stream of consciousness. First, the vast extent of our cognitive and affective processing occurs with little or no conscious

effort; and second, we don't lose most of our beliefs and desires when our minds are occupied with something else: they persist, as standing *mental states*, predisposing us to act in predictable ways. Unlike phenomenal, conscious experiences, I think there is no privileged access to many or our dispositional states because, for a start, one can only pay attention to what is conscious.

A central goal of this thesis is then to change the terms of the self-knowledge debate. When we expand our conception of the mental beyond what is conscious, the quest for self-knowledge becomes much more puzzling: *how could one get to know something one is, by definition, unaware of?* Chapter 1 establishes that there are two different ways and prospects of knowing one's mind insofar we have different subject matters with different metaphysics and respective epistemologies.

As it will stand clear after the first section of this work, I am not too worried about the ontological and epistemic status of conscious episodes. I have tangible evidence of their reality: I am directly aware of them. What seems to me much more obscure, thus in need of further illumination, is the realm of unconscious standing states. Hence a second goal of this thesis will be to characterise some of the features of unconscious states as to be in a better position to assess our epistemological prospects (and limitations).

Attempting to characterise the nature of unconscious standing states has posed immense challenges, for there is the risk of trivialising it if we understand it simply as 'mentality without consciousness'. Although the sheer absence of consciousness does of course imply non-consciousness, it does not encompass the

explanatory scope of unconsciousness relevant to my ends. Instead, I am interested in mental activity that occurs beyond the reach of immediate awareness and may aid in explaining behavior from the subject's point of view. We do not want to limit ourselves to information-processing, but we want to include higher-order, personal mental phenomena such as drives, intentions, and judgements.

To cast some light into this subject matter, two queries then need to be addressed. The first question concerns what makes an unconscious state mental. The second question pertains to what makes a mental state unconscious in the intended, dispositional sense— in that positing unconscious mental states may help us explain various otherwise puzzling behaviors. The difficulty surrounding these questions, I believe, arises not only from their inherent complexity but also from the false starts in our exploration of them. This is the rationale behind my chosen historical methodology for chapter 2. By analysing some of the history of philosophy of mind, I will be aiming to provide insights into the structure of the (unconscious) mind itself, as well as the limitations of the conceptual foundations handed down to us. Chapter 2 highlights three moments in the history of 20th century philosophy of mind, and how theories of the unconscious may have developed differently had Freud's theories received more substantial attention within analytic philosophy of mind.

Indeed, the history of philosophy of mind could have been different (20th century). Freud's understanding of the unconscious constitutes a substantial departure from the negative approaches found in traditional philosophy of mind. He viewed the unconscious as an active and influential force on human behavior.

Importantly, while the repressed unconscious is commonly studied, it comprises only a narrow subset of unconscious mental states. Focusing solely on the repressed unconscious can be potentially deceptive and limits our understanding of the scope and importance of unconscious mental states, currently studied under the framework of the cognitive sciences. By clarifying the broader domain of the unconscious over the merely repressed, I seek to support the continuity between Freud's unconscious theory and the mental states and processed studied within cognitive science today, paving the way for a clearer, future understanding of their connection.

The third chapter of this work attempts to change the course of this history. Chapter 3 is concerned with the content of unconscious representation, and how it might differ from conscious representation. The traditional "propositional attitude" picture of both mental kinds of content will be rejected, where it is wrongly assumed that unconscious states represent in the way conscious ones do. This traditional view has deceptive consequences for the agenda of self-knowledge, as we may be misled into thinking that the contents of conscious judgements correspond with those found in unconscious states. As such, it will be argued that unconscious content lacks the specificity, determinacy, and consistency of conscious content— characteristics, I further speculate, that may be ultimately connected to our linguistic capacities.

Chapter 3 will defend and spell out Crane's 'worldview' conception of the unconscious: the subject's entire unconscious orientation toward the world, which thereby encompasses the individual's dispositions to act. While I will largely side with Crane and Farkas, on the basis of the multi-track, indeterminate, and holistic

nature of many dispositions, I will raise some doubts over the epistemic competency of interpretation under a worldview theory. According to their proponents, interpreting one's mental states can optimally resolve the indeterminacy and ambiguity of one's unconscious, resulting in conscious judgments that can settle matters from one's perspective. I claim that self-interpretation raises the same basic issue about the first-person access and authority as introspection does, as exposed in chapter 1. Indeed, what is self-interpretation but the self-attribution of beliefs and desires deploying introspection?

Notwithstanding the fact that a worldview theory will be favoured over a propositional attitude picture, both views will be exposed to share a normative aspiration that matters for self-knowledge. That is, both views attempt to determine with existing natural language items what is unconscious. A propositionalist suggests that there are determinate facts about unconscious mentality, and we just need to employ some method for discovery. A worldview 'interpretivist' will deny that there are such facts to be "found", but he does offer such a method of interpretation — albeit for 'transcendental articulation', not immanent' discovery.

I will resist the scope of the idea that a complete picture of unconscious content can be optimally articulated through language expression. Instead, I will propose a third way for complex cases, in which honouring indeterminacy appears as an epistemically virtuous practice, beyond interpretative decisions and propositional discoveries. Any faithful representation of our unconscious states must accommodate or negotiate their inherent uncertainty and contradiction, ceding terrain for the mindful wisdom of self-ignorance.

1. NECESSARY DISTINCTIONS: *What is there to know about one's mind?*

This chapter is dedicated to making the necessary distinctions about the ways one could go to know one's mind. A fundamental difference is established between how we attain knowledge of conscious and present experiences such as sensation and occurrent thinking episodes, on one hand, and our dispositions, such as beliefs and desires, on the other. For this purpose, I will deploy metaphysical[1] theses about mental events and states to shed light on their divergent epistemologies.

1.1 Metaphysical account: phenomenological vs. dispositional

Given my chosen methodology, it is necessary to delve deeper into an underlying ontology, in which [*see diagram, over*]: "mental universals" are general features that can be attributed to various mental particulars, and "mental properties" are synonymous with capacities and dispositions that, when exercised, manifest as mental events and states. The relationships between properties and their manifestations are of consciousness and intentionality. Crucially, consciousness is a feature applicable to mental events—- it is the "how it seems like" for them. This is the minimal conception of "phenomenal consciousness" (Block, 1995) which I will exclusively regard as worthy of the adjective "conscious" in this work, unless

[1] Not in a revisionary sense, but in the descriptive kind of sense. That is, here I am not looking for truths that are true by necessity, or trying to disambiguate 'appearance' from 'reality'. Here I am looking for the best way to think about things: the conceptual foundations for our thought.

otherwise specified[2]. Finally, I draw a fundamental distinction between episodic mental events which are happening at specific —albeit not precisely measurable— times, and those that do not. While mental events manifest at particular points in our lives (like '*falling* in love'), mental states persist throughout most of them (like '*being* in love').

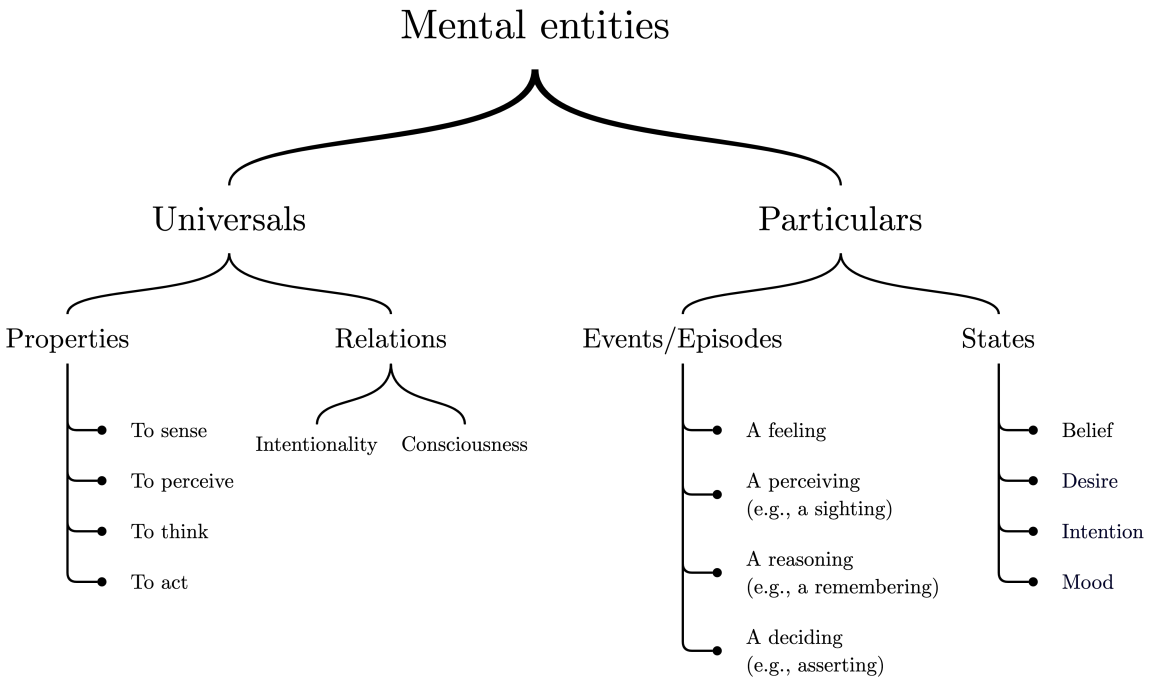

1.1.1 Mental events as experiences

A mental event is a discrete manifestation of a mental property. For example, from the capacity *to think*, episodes of thinking arise. By exercising the capacity of memory (i.e., *to remember*), an episode of remembering is instantiated: e.g., myself

[2] This phenomenal notion is in direct contrast with the one from "access-consciousness" (Block, 1995), which for my purposes is to be avoided— for reasons that will become clearer later in this investigation. The concept of "access-consciousness", on my view, may *smuggle* awareness into the notion of representation or information processing, creating further confusion. I see "access-consciousness" as a mongrel concept itself, as it includes both personal and sub-personal levels of cognition.

remembering the auspicious night I met my partner for the first time. Similarly, specific instances of feeling, seeing, thinking, remembering, deciding: all these are occurrences that have temporal parts that are streamed in consciousness— they are being experienced as such[3].

1.1.2 Mental states as dispositions

A state is a standing instantiation of a property over time. This notion contrasts with the one-time occurrences exemplified by mental events, such as decision-making. Consider intentions and ambitions for contrast, which involve the continuous and stable commitment to a particular goal or action. For instance, the state of admiration towards a role model persists over time, as opposed to a fleeting feeling of awe experienced during an inspiring performance. Mental states, as such, persist through changes in one's consciousness— and even in the absence of it.

Notice that the general term "thought", in a non-technical use, typically encompasses various events and states from both aforementioned categories. So the subset of acts of reasoning, counting, inferring, supposing, etc. as well as the subset of beliefs, hopes, intentions, etc. would be included in the broader set of "thoughts". This is thinking in a *broad sense*. For the technical purposes of my investigation, however, I will be using thinking in the *narrow sense (*unless otherwise specified*)*: i.e.

[3] One could point out that thinking doesn't seem to have the same (if any) phenomenal character as other kinds of episodes, like feelings. As such, one could be tempted to posit that the conscious character of thoughts differs from the conscious character of other kinds of episodes. This may be a relevant observation for other sorts of investigations, but it is not of concern here just yet. I will just note that conscious thinking, insofar as it is present to the mind, seems to be manifestly real or factual as it occurs in 'real-time': it occupies time by having temporal parts.

properly episodic acts of thinking that acts of remember that, judging that, wondering whether, and so on.

This clarification will prove most useful in our following chapter, where the content of the mental state 'belief' is posited primarily or totally unconscious. That is, if we accept a dispositional account of belief as taking the world to be (in a certain way), which governs one's behaviour, then most of what we believe cannot be items in the stream of consciousness. In contrast to thinking (in the narrow sense), my beliefs must persist in the absence of manifestation: I believe my name is Stephania, even when I am sleeping. Under this account,, beliefs as mental states are not essentially conscious.

This can be all a bit cumbersome, as the conscious acknowledgement of the truth of something is, in a non-technical sense, is also typically called "belief". But I will continue to insist on the occurrent thought/standing belief distinction, as it is predicated upon the unique functional roles that these entities play in our cognition. When we are "conscious" of beliefs, we are thinking: we are bringing a belief into consciousness in the form of thought, and it is only in that realm where we decide over what to do with them (e.g. judging, assenting, avowing). Yet believing, by itself, persists as an unconscious disposition to act.

1.2 Epistemological account: direct vs. inferential

Under our metaphysical account, we have two kinds of objects (of self-knowledge): mental events and mental states. These two are distinguished by their

ontological categories and bear different relations to consciousness. While mental events involve experiences—that is, they are present, occurrent, and conscious—, mental states include dispositions: they stand and persist over time, and as such are not necessarily present in the stream of consciousness. The nature of these two kinds of objects of knowledge puts us in two very different epistemic positions— our methods of access and hence prospects for achievement are not the same.

In the case of conscious and present experiences, the relation to consciousness is more direct— what it is like for us to have an experience (e.g a feeling, a sighting) is resolved by the way we are aware of an experience. As I will specify in the following section, this does not mean that the content of the judgement is the same as the content of inner awareness (that is, identical to the phenomenal character of the experience). What it means is the event is right there, "in front of our mind"— as we go through the experiences, they form the very vehicles of relevant, associated conscious judgements. As a result, the method of getting to know these events becomes a relatively undemanding task. For us to judge the character of a flavour, smell, or sensation, all that it takes is re-directing one's attention within. Introspection is illuminated under a "phenomenological light", so to speak.

According to some philosophers, if we accept that conscious experiences involve an awareness of themselves, attention to them can afford us "limited phenomenal infallibility", what they identify as leading to "phenomenal knowledge" (Horgan & Kriegel, 2007). For the purposes of my investigation, I believe it is not necessary to go as far as claiming infallibility— my commitment is limited to

privileged access. Indeed, I do not think grasping the nature of mental events is an infallible task (e.g., is this a feeling of excitement, or nervousness? Am I angry, or just hungry?), but this is not of my epistemic concern. Even if I am uncertain about the nature of my mental events, I still can "point" at them, and claim authority over the existence of whatever experience I am having.

In the case of standing states, the relation between us having a disposition and the way we are aware of them is more obscure. Getting to know many our beliefs and desires —think most dramatically, implicit bias and prejudice— is not as simple as re-directing one's attention within. Indeed, our awareness does not readily reveal the content of a standing state— but merely their conscious manifestations, if any. As a result, introspection loses its epistemic force: directing one's attention within does not reveal their contents. We do not have privileged access to our mental states, even when they remain present in our life through their manifestations. And very often, as we will see in the following chapters, as speculative explanations for our behavior.

1.2.1 Objective facts about subjectivity

Here I would like to establish a technical albeit pivotal disambiguation. Namely, I would like to further distinguish experiences themselves (e.g., redness, pain), with *judgements* about them (e.g., "*This* is how red looks like", "*This* is how pain feels like"). Technically, in this work, experiences (nor dispositions, for that matter) are not strictly the objects of knowledge, only facts —our judgements about them— are. In turn, experiences and dispositions are the objects of our judgements. I

would like to then reveal then the epistemological commitments of this work by a well-known vignette: let's disambiguate what's new for Mary after she leaves the lab (Jackson, 1982; *cf.* Robinson, 1982).

No-one (as far as I know) denies that Mary has a new experience. That is not up to controversy. What is up to controversy, I think, is first, whether this experience constitutes knowledge by itself (my answer would be no), and second, whether this experience can lead to knowledge (my answer would be yes). From the perspective of this work, experiences are not, intuitively, true nor false— they just *are*. These are ineffable, intrinsic, and private experiences arguably irreducible to the truth and falsity of judgements. But, as Wittgenstein (1953/2009) has pointed out, knowledge (and this includes self-knowledge) should be compatible with mistake: where one cannot be wrong, one cannot be right either.

To know, on my view, is to be in non-accidental contact with truths in the exercise of our rational capacities. In this sense, merely going through an experience does not call for reflective exercise. If it's an achievement, it's at best sensorial, unlikely cognitive— at least in the sense that matters for an epistemology based on truth and falsity. Experiences and dispositions are the objects of our judgements, but it's the judgement itself the object of knowledge, which is either true or false. (even though, as we have seen, going through experiences is usually accompanied by the acquisition of knowledge, since we have privileged access via introspection).

Merely experiencing something does not constitute knowledge. This, of course, does not preclude that having an experience wouldn't *enable* us to then gain

new knowledge (as previously stated, it often does), or even that there is some knowledge you can learn *only* if you go through certain experiences. Indeed, I exercise my reflective capacities and somewhat easily make a judgement about the experience (e.g., "Oh, so *this* is what red looks like"). From our vignette, Mary has learned an objective fact about subjectivity. She could even embed this content into conditionals, a mark of factuality: "if seeing red is like this, then it is not like that".

Notice Mary's epistemic progress. First, she experiences something new when she leaves the room and sees red, and then she learns something new when she reflects about it ("Aha! This is red!"). The result of this reflection is knowledge. Now, with this new truth, she can form new concepts, she gains the capacity to ask new questions, and to form new hypotheses (e.g., "Apples are red, are strawberries red too?"). One shall shall resist the temptation to identify an epistemic achievement from the conditions that enabled it. The experience of redness is not the subject matter of epistemology, or so I say. What is of epistemological concern are the judgements we form after we have reflected on our experiences.

1.3 Summary

We have two kinds of objects (of self-knowledge) associated to two different ontological categories which bear different relations to consciousness. Judgements about mental events are about experiences, which are present, occurrent, and conscious. Judgements about mental states are about dispositions, which stand and persist over time, and as such are not necessarily present in the stream of

consciousness. The nature of these two kinds of judgements, or objects of knowledge, puts us in two different epistemic positions. For mental events, introspection affords us privileged access to their content (i.e., experiences). As such, we produce generally reliable judgements over what these are like. For mental states, introspection loses its epistemic force: directing one's attention within does not reveal their contents (i.e., dispositions). As a result, there is not necessarily privileged access to many of our mental states.

The methods of access and hence prospects for self-knowledge are plural. There are at least two different ways and prospects of knowing oneself insofar we have different subject matters with different metaphysics and respective epistemologies. While we have tangible evidence of the reality of mental events — they're there, directly in front of our minds—, the realm of unconscious standing states is more obscure. Particularly, the inference of their content. As such, a second goal of this work will be to characterise some of the features of unconscious states as to be in a better position to assess our epistemological prospects (and limitations).

2. SUBJECT MATTER: *A historical approach*

One should exercise caution when determining what qualifies as an unconscious phenomenon, for there is the risk of trivialising it away. First, the mere absence of consciousness cannot be equated with the explanatory sense of unconsciousness relevant to this investigation. For instance, flowers do not "unconsciously" turn towards the sun, trees "unconsciously" shed their leaves in autumn, nor planetary bodies "unconsciously" follow elliptical orbits. These are examples of biological and physical phenomena devoid of mental activity, very much like the knee-jerk reflex or the withdrawal of a hand from a hot surface. Flowers, trees, planets, and hands are entities with no mental lives. Such phenomena can be classified as *non-conscious,* while mental phenomena that elude the possessor's awareness during their occurrence, even when prompted, can be designated as *unconscious.* For contrast, consider the case of someone showcasing mastery of complex patterns without realising them (e.g., grammatical rules, dance moves, gender norms): herein there is mental activity occurring beyond the reach of immediate awareness that may aid in explaining various behaviors.

My task in this chapter is then to clarify these ideas —that of unconscious mentality and its behavioural relevance—- in two questions that aim to elucidate my subject matter. First, what makes an unconscious state mental. And second, what makes the mental state unconscious in the intended sense: that is, in the sense that positing unconscious mental states may help us explaining various otherwise

puzzling behaviors from the point of view of the subject. The notions of *representation* and *dispositionality*, respectively, will help me delineating answers to these questions[4].

2.1 What makes an unconscious state *mental*?: REPRESENTATIONALISM

2.1.1 Decoupling cognition from consciousness allows us to isolate representation

At the turn of the 20th century, a common theme in a myriad of philosophical movements was consciousness as the central feature of the mental. Crane (2019) highlights at least three axes of discourse: perception and realism discussed by G.E. Moore and Russell, realism and pragmatism by William James and C.I. Lewis, and the phenomenological movement pioneered by Husserl. They have all emphasised the significance of understanding the mind in relation to consciousness, henceforth leaving behind the conceptual foundations for a mentality epitomised in sensation and perception.

At the time, consciousness and cognition[5] were thought to be phenomena that were closely related to one another— and should not be examined separately: consciousness is how one gets to know the world; it's the vehicle of one's knowing. Indeed, in early 20th-century analytic philosophy, the problem of consciousness revolved around the nature of our access to the external, mind-independent world. As such, inheriting an epistemic conception of mind is historically expected in

[4] The methodology of this chapter is largely historical. Many ideas developed here grew in a previous paper of mine (Donayre Pimentel, 2023) for a history course, "The Making of the Modern Mind". The instructor, Emese Lafferton, has kindly authorised me to use the material on this dissertation.

[5] For the purposes of this investigation, *cognitive* = *intentional* = *representational*, unless otherwise specified.

modernity. The emphasis on the centrality of consciousness was also shared by the nascent field of scientific psychology, which emphasised the need for a more systematic approach to study the contents of consciousness. Founders of psychology such as William James and Wilhelm Wundt saw consciousness as the primary subject matter of their discipline, as well as introspection as its appropriate method (James, 1890; Wundt, 1896).

Notwithstanding its introspectivist roots, throughout the 20th century, psychology turned increasingly preoccupied with establishing itself as a legitimately Galilean discipline. As such, behaviorism was seen as a methodological avenue that could provide the scientific respectability it sought. With its focus on the objective measurement of behavior as the only valid scientific method, other "undefinable" factors such as consciousness became inessential to behaviorally efficacious mental processes. This perspective represented a significant departure from the earlier emphasis on consciousness as the primary focus of psychology, and it would set the stage for subsequent developments in the field, such as the "cognitive revolution" of the 1950s and 1960s.

Crane's key insight (2019) is that the impact of behaviorism on the philosophy of mind in the 20th century did not stem from its plausibility as a substantive theory about the nature of mind and consciousness. Instead, its philosophical importance lies in the fact that it introduced and reinforced a view of mental life that could be largely independent of consciousness. For the purposes of an investigation into unconscious mental processes, I see this as the "bright side" of behaviourism: its

emphasis on observable behavior and its rejection of introspection as a valid source of knowledge helped to establish the idea that some behaviourally efficacious mental states could be understood without appealing to the contents of conscious experience.

In the late 20th century, then, consciousness is no longer deemed to be the fundamental characteristic of the mind. As we are now able to explain some behaviors without appealing to conscious states, cognition and consciousness can be decoupled—regarded as separate and distinguishable phenomena, as they largely remain today. This modern distinction between cognition and consciousness enables me, as a 21st-century theorist, to differentiate between conscious and unconscious mentality and hence propose an answer to the question of what makes an unconscious state mental: unconscious mental states *represent*. This stance is achieved by positing that consciousness is not essentially cognitive, intentional, or representational; and representations, intentionality, or cognition, are not essentially conscious.

2.1.2 Moving away from a negative conception of unconsciousness

The disjunction between consciousness and cognition can provide philosophers in the 21st century with a means of accounting for unconscious mentality in terms of representation. However, this notion of unconsciousness is not exhaustive, as it does not specify how it matters to behavior in the sense that matters

to this investigation. That, is, an account of standing states— as described in chapter 1, our second object of self-knowledge.

While one may grant that unconscious mental states are representational, numerous non-mental states also convey information about the external world. For instance, the states of one's kidneys may register the concentration of salt in the blood (Berger, 2014). Also, some "sub-personal" unconscious information-processing states (e.g., retinal states providing information about the properties of edges and colors) can be considered mental, on the usual understanding. But unlike these mere "sub-personal" unconscious information-processing states, standing unconscious states are "personal": they actually partly constitute the subject's point of view (e.g., their beliefs, their desires), and may be in some sense accessible to the subject.

Simply carrying information may not be enough to establish the kind of representation that is considered to be unique to the unconscious mentality— that which is relevant to the standing states we posit to explain our behavior. Therefore, representationalism needs to find a way to distinguish personal from subpersonal states, as well as to distinguish subpersonal mental states from other subpersonal states that carry information.

While understanding the mind in relation to consciousness did not explicitly deny the existence of unconscious mentality, it exerted a significant influence on how we conceive of it. Specifically, conscious cognition has been inherited as the template through which all forms of mentality, including the unconscious, are to be understood. This position leads to a narrow understanding of the unconscious that is

largely modelled upon conscious states. As a result, what is unconscious, is often reduced to a collection of subpersonal mental states at best, or even denied mentality at worst.

We shall call this the *negative* (in the sense of absence) approach to unconscious mentality. This perspective has, troublesome, far-reaching consequences to how we understand unconscious mentality. Under this outlook, unconscious processes and states are merely a consciousness-free version of the conscious mind. That is, the fundamental nature of representation is common to both conscious and unconscious states; the only difference is the absence of the quality that renders conscious states as conscious— their "light". Consequently, and potentially perniciously for the study of self-knowledge, the properties of conscious judgements are thought to correspond with those found in unconscious states. However, as we reviewed in chapter 1, there are compelling reasons to think otherwise; that is, that we are dealing with distinct objects of knowledge.

The history of analytic philosophy of mind has left us with dichotomies that largely neglect to consider what unconscious mentality is like. Even though the unconscious is increasingly being referred in various contexts to explain behavior, it is often conceptualised under the template of consciousness, in the way just described. As a result, today the major theories of consciousness place restrictions on the kind of unconscious mental states that there are. See in *Philosophy Compass*: "The nature of unconscious mentality depends on how we understand the nature of consciousness" (Berger, 2014), or more recently: "There is evidence that different types

of mental states can be unconscious, though that conclusion depends on the theory of consciousness assumed" (Dienes & Seth, 2022).

The preceding discussions underwrite the inadequacy of consciousness as the template through which all mental phenomena are understood– especially unconscious mental phenomena. To do justice to unconscious representation, and to properly investigate its relevance to behavior, we must abandon the negative definition of the unconscious. To this end, one shall often turn to examine influential movements outside philosophy; for example, psychoanalysis and the cognitive sciences. Given space constraints, I will focus in psychoanalysis in the section that follows— not without explicitly stating the continuity of these disciplines for reasons that will become clearer later.

2.2 What makes a mental state *unconscious*?: DISPOSITIONALITY

2.2.1 Freud: the repressed, the dynamic, and the wider realm of the unconscious

Sigmund Freud proposed the existence of unconscious mental states and processes as an explanation for a wide range of human behaviors (1917/1963; 1933/1964). Freud's primary argumentative strategy is an *inference to the best explanation:* he argued that myriad puzzling phenomena, from "the detritus of everyday life" (e.g., parapraxes, dreams, idiopathic pains) to associations of ideas[6] and inconsistent behavior (e.g., acting contrary to one's sincerely professed beliefs) would remain unintelligible if we take all mentality to be conscious. The postulation

[6] Seemingly inexplicable "ideas that come into our head we do not know from where" (Freud 1923/1961).

of unconscious processes, in turn, would render behavior to be in principle understandable. Most dramatically, Freud maintained that if the postulation of unconscious states leads to an "effective influence upon the course of conscious processes", then it provides "incontrovertible proof" of their existence (1915/1963). Freud took the ability of psychoanalysis to influence and treat psychological issues[7] as empirical support for the existence of unconscious mental states and processes.

The attribute of "incontrovertible" strikes as too strong. While Freud's postulation of unconscious mental states could potentially offer an explanation for some of the aforementioned phenomena, it is also possible to conceive of other theories that do not rely on them. For instance, the random flow of ideas in dreams could have a purely physical, or impersonal explanation (e.g., Hobson et al, 2000); or inconsistent behavior could be explained away through a conscious, mentalistic explanation— in terms of convictions and motives that are not grounded in the unconscious (e.g deflationist theories of self-deception; Galeotti, 2018). Additionally, explanations involving the unconscious might not always render a person's behavior comprehensible in terms of existing linguistic, or symbolic patterns. For instance, the application of a psychoanalytic approach to psychosomatic symptoms or neuroses may not distinctly clarify these occurrences in terms of causal explanations. Suffice to

[7] In this work, I do not intend to defend the therapeutic efficacy of psychoanalysis. Suffice to say for my purposes, Freud's reliance on therapeutic success as an argument echoes fallacies recurrent in both the history of religion and psychotherapy. Myriad narratives of success stem from theories that lack both robust theoretical foundations and empirical rigour— e.g. placebo effects often contribute to a patient's improvement solely through their perception of receiving care. These reservations certainly don't disprove of all Freud's assertions, but they are sobering: the complexity involved in determining causation in the psychological domain is certainly more challenging than Freud acknowledged.

say, appealing to unconscious mental states does not always make people's behavior more intelligible.

What Freud requires is evidence of behavioral patterns that could be explained by the content of the specific unconscious states he posits. Sometimes, he provides such evidence. In cases of "glove anaesthesia", for example, Freud and Breuer observed that hysterical patients were often "paralyzed" in areas not typically affected by paralysis, but instead in areas that were colloquially *thought to be* susceptible to paralysis (i.e., hands). Freud's explanation for this unusual pathology involved reference to the content of the belief: he argued that the content of the colloquialism was present in the patient's mind, but not consciously accessible. This highlights the theoretical centrality of unconscious content in explaining certain behaviors, such as this patient's paralytic symptom.

Freud's conception of the unconscious is often interpreted in the light of what he sometimes calls "the truly dynamic sense" (1915/1963), a specific type of unconscious mental states. This sense is idiosyncratic by which its unconscious processes are intellectually rudimentary, only subject to "the pleasure principle"—what Freud came to call the 'id' as a whole. It is also under this "truly dynamic sense" that Freud believed that such mental processes start as an unconscious ones, and only become conscious if they meet resistance or censorship: "every psychical act begins as an unconscious one, and…may remain so or go on developing into consciousness according as it meets resistance or not" (1915/1963).

Focusing on the repressed unconscious is customary, but potentially deceptive: this term comprises a narrow subset of unconscious mental states; it does not exhaust them. Repressed unconscious states are those unwanted thoughts which are often unreachable due to being actively blocked from consciousness by "psychological defences": e.g., projection, reaction-formation, rationalisation. Due to the active nature of these suppressing forces, Freud used the term "dynamically unconscious" states to refer to repressed unconscious states. While his work largely focuses on these —as they are thought to be the root of all psychopathology— it is important to note that the dynamic unconscious refers only to one category of unconscious mental state, rather than the totality of them. In fact, Freud makes a clear distinction between the dynamic unconscious and the wider realm of the unconscious: "Everything that is repressed must remain unconscious; but let us state at the very outset that the repressed does not cover everything that is unconscious. The unconscious has a wider compass: the repressed is a part of the unconscious, but not the whole of it". (Freud, 1900/1953). Not all unconscious mental states need be repressed.

2.2.2 The mark of the Freudian unconscious

If not repression, then what is the true mark of the Freudian unconscious? According to Wakefield's (2018) exegesis, the characteristic feature of unconscious mental states is their active functional role in cognition. He argues that it could be inferred that Freud considered a mental state to be 'active' at a particular moment only when its content engages with other contents in the mind, and this interaction is

partially observable in consciousness. Freud hints: "We call a process unconscious if we are obliged to assume that it is being *activated* at the moment, though at the moment we know nothing about it" (1923/1961).

Consider unexpectedly feeling an intense surge of fear when entering a particular place. This sudden emotional response could be explained by the interconnection between unconscious mental states (such as past traumatic experiences or repressed memories) and conscious mental events (the current perception of stimuli or sensory cues). Or imagine someone experiencing an unexplained sense of nostalgia upon encountering a familiar scent. This might result from the interplay between unconscious mental states, like buried memories, and conscious mental events, such as olfactory perceptions. Certain aspects —but not all — of this interaction would be observable in consciousness, and we could *infer* the presence of an 'active' unconscious mental state in Freud's framework.

This notion of 'background activity' emphasises the notion of a hidden, internal realm of mental happenings that operate somewhat independently of conscious awareness. The active contents from unconscious mental states remain beyond conscious perception, yet exert an impact on conscious experience. This notion underscores how there are some standing unconscious states that partly constitute the subject's point of view, and may be in some sense accessible to the subject (though if not specifically psychoanalytic, through systematic observation of behavior, personal history, and thought processes).

The aim of this section is not to defend any specific postulate of psychoanalytic practice. Rather, it seeks to emphasise two key aspects of Freud's theory as a philosophy of mind. First, Freud's understanding of the unconscious as an active and influential force on human behavior, which stands in contrast to negative approaches in traditional philosophy of mind. Had Freud's theories received more substantial attention within the history of analytic philosophy of mind, the field may have developed differently. Second, by clarifying the broader domain of the unconscious over the merely repressed, this section supports the continuity between Freud's unconscious theory and the positing of unconscious mental states and processes currently studied under the cognitive sciences today, paving the way for a clearer, future understanding of their connection.

For Freud's theories on the unconscious to be further substantiated, the best evidence would come from behaviors explicable by the specific contents of unconscious states. These are not thought be mere sub-personal information processes. Rather, they are thought to represent enduring, personal properties of the individual, shaping their perspective or worldview, and being potentially accessible to them in some manner. Notice he only reason Freud could find for mental content not becoming conscious is that it has been in this way censored or 'repressed'. As it stands, then, Freud's theory cannot account for other possible reasons why non-repressed material remains unconscious, such as a lack of appropriate causal or informational paths between unconscious content and introspective mechanisms. It is part of the ambitions of this thesis to shed light on these other possible reasons with an analysis of the content of unconscious representation in the following chapter.

3. CHARACTERISATION: *Unconscious Content*

If we follow philosophical convention of calling the way a state represents the world its *content*, then we can ask what is the difference between the content of unconscious and conscious mental states. In traditional philosophy of mind, a "propositional attitude" picture of both mental kinds of content has been given, where the same account of content is given for both what is conscious and unconscious. On this chapter, I reject this dominant picture as a way of introducing the idea that the manner in which unconscious states represent differs from how the conscious ones do. In particular, it will be argued that unconscious content lacks the specificity, determinacy, and consistency of conscious content. These are characteristics that are ultimately connected to our linguistic capacities.

The view I am defending in this chapter owes a lot to Daniel Dennett (1983), Eric Schwitzgebel (2010), and especially Tim Crane and Katalin Farkas (Crane 2017; Crane & Farkas 2022a, 2022b), particularly on their views about intentionality, belief-ascription, and standing states, respectively. Crane's proposal is that ascriptions of mental states are partial characterisations of what he calls a subject's 'Worldview'—their entire unconscious orientation towards reality, which forms the ground for all their cognitive, conative, and affective dispositions (2017). For the purposes of a thesis on self-knowledge, the "worldview theory" will appear in stark contrast to the "propositional theory", where states are propositionally represented (whether states are interpreted as 'sentences in the head' or as entirely determined dispositions). As

such, these theories will offer very different roadmaps for the quest of attaining self-knowledge.

While I will largely side with Crane and Farkas on their description of the multi-track, indeterminate, and holistic nature of dispositions, I will offer further epistemological constraints to the normative role of interpretation for knowing one's mental states. As it will be argued, interpretation is at the core of their worldview theory— its theorists remain optimistic that the complexity of content can be articulated through language expression. Crane (2017) says: "to model an aspect of a worldview with a proposition is to interpret it". However, self-interpretation raises the same basic issue about the first-person authority as introspection did in chapter 1. Indeed, what is self-interpretation but the self-attribution of beliefs and desires using introspection?

In this chapter, I will focus on the phenomenon of belief, but I think many of these claims would be applicable to other states, such as wants or desires. In what follows, I articulate some paradigmatic features of the representational content of conscious mental states (SPECIFICITY, DETERMINACY, and CONSISTENCY). These three properties are not shared by the representational content of unconscious mental states.

3.1 SPECIFICITY: propositions and behaviors

Consider the standard conception of belief as a propositional attitude. On this view, the representational content of a belief, whether as a persisting, unconscious

state, or an occurrent, conscious event is the same: a *proposition*— a bearer of a truth value. As such, beliefs are thought to be individuated by their propositional content: they are set apart by the propositions that specify how they take the world to be. The state of belief, following the ontology laid in chapter 1, would be an instance of a property (with numerically distinct states corresponding to separate instances), or a relation to a propositional content under the propositional attitude view.

My belief and your belief that I wrote this thesis would be treated as instances of the same property (to believe), or the same relation by the term "believes" to a propositional content (that I wrote this thesis). These are a couple of differentiated states, just as the belief that I wrote this piece is different from other of my beliefs like I am competent enough to finish it. The individuated states can be then expressed in myriad ways, like dispositions to utter certain words, to behave in a certain way, or to aim at truth.

This conception of belief raises several puzzles. First, if for each belief I have there is a correspondent proposition to which I am related, how many beliefs do I actually have? If a belief is truly a numerically individuated state, specified by its content, then it must be distinguishable from other beliefs and hence fundamentally countable. But this seems just plainly wrong— how shall one count all the "*P*'s" one takes to be true as one navigates the world as a whole? Second, can infants and animals have beliefs if they cannot entertain propositions? But they seem to build an unconscious worldview as they attempt to register the facts from their environment

at varying levels of accuracy. If we're asking such questions, we must have gone wrong somewhere in our propositional theorisation of belief…

One promising theoretical alternative is to distinguish between occurrent episodes of believing and *belief* as a dispositional state, as the ontology laid out in Chapter 1 allows us to do. This distinction is predicated upon their unique functional roles in our cognition. Indeed, one can strategically stipulate that the term 'belief' should be exclusively used in the *dispositional* sense, while 'think' (or more precisely, "judge" or "assert") should be reserved for the intended *occurrent* sense of believing. In the ontology presented in chapter 1, merely entertaining the thought that *P* is neither necessary nor sufficient for believing that *P* (even when the conscious thinking that *P* aligns with one's underlying belief). Beliefs are states, while thoughts are events. In this way, thinking-events are conscious, whereas doxastic-dispositions remain unconscious. A disposition cannot be called conscious if it persists over periods of unconsciousness (e.g., sleep and inattention). However, the *manifestation* of a disposition can be conscious, such as certain episodes of thinking.

In this latter framework, the distinct ways in which unconscious and conscious states represent become more evident. Even when we use the term "belief" in a strict dispositional sense, the possibilities of deception and plurality of behavior show how there is not a straightforward connection between acting and taking the world to be in a certain way, or believing something to be true— for example, that "skydiving is dangerous". Any type of action can be the product of disparate beliefs, given divergent motivations (e.g., my will to survive or perish), external

circumstances (place), and other relations to other mental states (desire for adventure). Unconscious content cannot be specified or individuated by a single behavior, as conscious content by single propositions can.

Here I am admittedly betting on a functionalist notion of belief as a mental state. That is, I am building the causation of belief into its very nature by defining it in terms of its causal or dispositional relations to action, perception, and other mental states (e.g., desire). Hence, what counts as behaving if we take the world to be in a certain way, or something to be true, cannot be easily specified. The very same belief can give rise to many different other states or actions, depending on which other mental states one has. Belief (like desire) is a 'multi-track' disposition.

3.2 DETERMINACY: attribution and models

The distinction between unconscious belief and conscious thought promises to reconcile ordinary belief ascriptions with the intricacies of individual's psychologies. Crane (2017) suggests that we consider propositional attitude attributions as models, as defined within the philosophy of science: "a concrete or abstract object used to draw attention to some features of a system under investigation, and to make the study of the system more tractable".

A model is a representation employed to highlight specific features of a target system, thereby enabling surrogate reasoning about that system. Models employ idealisation, abstraction, and simplification to emphasise certain features of the subject they represent. For example, some propositional models can focus on aspects

of your beliefs, and employing various types of propositions can highlight distinct features. Propositional attitude attributions emphasise (as well as suppress; Woods and Rosales 2010) salient aspects of one's mental states.

If a model is a representation used to highlight specific features of a target system, the target system in this context is a person's "worldview", according to Crane (2017): their entire unconscious representation of the world. The worldview embodies the subject's entire orientation toward the world, and thereby encompasses an individual's dispositions to act[8] . On this account, individual belief attributions do not correspond piecemeal to individual belief states. What belief ascriptions do is model —not fully determine— various aspects of this comprehensive structure. What counts as behaving if we take the world to be in a certain way, or something to be true, can be very well indeterminate.

The worldview as a whole has a representational nature— it is the unconscious framework itself. It encompasses all the inclinations associated to one's mental states, which are directed on the world. Indeed, in the case of belief, its "direction of fit" would be "mind-to-world" (Searle, 1983); while the case of desire, it would be "world-to-mind". Therefore, instead of determining a single disposition with a specific propositional content, the propositions employed in attributing beliefs are tools to represent facets of this representational structure.

[8] Crane introduces his worldview theory on his 2017 article "The Unity of Unconsciousness" and slightly modifies it on his 2022 article along Farkas: the worldview does not merely contain the subject's entire *doxastic* orientation toward the world, but also their affective, conative, etc. dispositions.

3.3 CONSISTENCY: the role of interpretation

Much of what we believe tends to be fragmentary, cloudy, undeveloped, and all the more inconsistent. The traditional model, which involves a single proposition matched with an individual belief state, cannot offer a psychological picture where this is possible. But our inner lives are full of the following examples:

Meet Tasneem, an avid food critic, and her best friend, Ohan, who aspires to be a chef. Tasneem genuinely cares for Ohan and supports their culinary endeavours. Ohan opens a restaurant, and Tasneem praises the dishes publicly, writes glowing reviews, and encourages others to try the food. Now, the question arises: Does she "really" believe that Ohan's cooking is that good? Or she has been "blinded" by the affection for her friend? Or does she not believe his cooking is good, but pretends it is out of loyalty and friendship?

It's a complex situation. Does Tasneem, or does she not, believe that Ohan is an exceptional chef? Tasneem's worldview encompasses a web of emotions, friendship, and support for Ohan, resulting in a set of dispositions that might not consistently align. Parsing her belief about the restaurant's culinary excellence is intricate. In such cases, providing a straightforward yes or no answer may prove elusive.

Certainly, not all cases conform to this pattern. Numerous facets of our worldview are uncomplicated, with dispositions amenable to efficient single-proposition modelling, like "I believe (*think*) it's raining outside". (However, even in these uncomplicated scenarios, the comprehensiveness of the unconscious ensures that there is no strict requirement for a single dispositional profile linked to each believed proposition— *maybe I'll grab an umbrella, maybe I just want to get wet*).

It is sure important to recognise that, in many cases, interpreting the worldview isn't very hard— we can get to know many of our beliefs and desires as we make a consistent connection between them and our behavior. Indeed, when we query our own thoughts, sometimes the answer comes effortlessly: "How to say a toast in Georgian?" It comes to mind: *Gaumarjos!* The information already stored, seamlessly resurfaces into consciousness, making my belief explicit and conscious. The precise mechanism behind this transition (that is, bringing what is in our unconscious to consciousness) is a difficult inquiry, but the potential capacity for such retrieval cannot be doubted.

In much of our inner lives, several elements consistently align: conscious thoughts, unconscious states, and behaviors harmoniously synchronise. We can consciously acknowledge the truth of "red traffic light = stop", our judgement conforms to the expected behavior linked to holding such a belief, and our underlying dispositions change as our interaction with the light changes to green. Further, some other of our dispositions can be understood in terms of the usual understanding of desires: for example, someone's longing for personal growth would align with the expected behaviors, thoughts, emotions, and intentions characteristic of someone with this aspiration.

But in many cases —especially central to people's lives—, these elements come apart. Our worldview and conduct are frequently a blend of equivocal commitments, vague desires, confusion, fear, and weakness of the will. I take that we are all a little bit like Tasneem when it comes to answering many complex and very personal

questions. And many of us try to resolve these perplexities and inconsistencies within ourselves, to ourselves. How might we—or ought we—approach this task? Crane and Farkas suggest that "we can benefit greatly from attributing a mental state*, as long as we make sure we present a complex enough picture to preserve *consistency* […] as long as one can present a *plausible* model, one can say that the subject has the mental state"[9] (2022a).

Interpretation is at the core of the matter here. Crane (2017) says: "to model an aspect of a worldview with a proposition is to interpret it". Making a belief conscious, then, would involve interpreting an aspect of one's worldview, where interpretation entails grasping a system—oneself—and making it consistent with the rest of our conscious judgements. From a first-person perspective, understanding one's beliefs would be an act of self-interpretation—uncovering an optimally consistent way of articulating or portraying one's beliefs. This process might involve "making up one's mind" or, in straight-forward cases, finding what is contained within it. Resolving ambiguity and vagueness in your worldview comes through *re*-discovering one's beliefs, leading to a conscious decision that brings clarity for us. From this perspective, the content of a belief emerges as it is investigated.

3.4 Roadmaps to self-knowledge: decisions, discoveries, and silences

The 'worldview theory' stands in stark contrast to the propositional representation of belief states, whether interpreted as 'sentences in the head' or as

[9] Redacted from "desire" while keeping the intended meaning. Emphases added.

entirely determined dispositions. As such, these offer two very different roadmaps for the quest of self-knowledge. The propositional framework implies that all facts about our beliefs are readily available for discovery. Under this account, determining 'what one thinks' is one thing, while 'making up one's mind' is another. Following from Crane's worldview theory, Crane and Farkas (2022a, 2022b) reject the sharpness of this distinction. According to them, uncovering one's beliefs can resolve the indeterminacy and ambiguity in one's worldview, resulting in a conscious judgment that settles matters from one's perspective.

Following Crane, I think that the picture of fully determinate mental states should be rejected, whether these are taken to be sentences in the head, or fully determinate dispositions— each of which corresponds to a single attributed mental state. Indeed, things like beliefs or desires should not be thought as determinate, discrete states; and our conscious judgements as the outputs of them— "in the way a printer outputs the text stored in a computer" (Crane & Farkas, 2022b). To be clear, the point it is not that there are no clear cases in which we can find out what we already believe (e.g., "I am a grad student", "*joie* means joy in French"); or that there are no straightforward cases in which we make up our minds (e.g., "I prefer pistachio over vanilla, please"). Crane and Farkas would say that the point is that there is no sharp distinction. The way I would rather put the point, however (for reasons that will become clear later), is that there are no obvious criteria for us to decide which is the case we have at hand.

One could agree with much of the worldview theory as a sketch of unconscious states, while disagreeing with the normative scope of interpretation in acquiring knowledge of them. I will refer to the latter as an 'interpretivist' stance about self-knowledge. Indeed, Crane and Farkas further their view normatively by claiming that, from a first-person perspective, understanding one's beliefs can be an act of self-interpretation—uncovering an *optimal* way of articulating or portraying one's beliefs (2022a, 2022b). This position puts them closer to a view of self-knowledge spoused by philosophers like Richard Moran (2001), where a complete characterisation of the first-person perspective requires bringing explicitly the subject as active participant. But doesn't self-interpretation raise the same basic issue about access in introspection, addressed in chapter 1?

Let's recall that what affords us privileged access to our occurrent mental events is the fact that these involve experiences, which are present, occurrent, and conscious. For these cases, introspection affords us privileged access to judging their content, our first object of knowledge. In contrast, in chapter 1 we have shown that there is no privileged access between us having a disposition and the way we are aware of them. Our awareness does not grant us access to the content of our standing states. As a result, the epistemic force that introspection has for events, it is lost for states: directing one's attention within with the goal of self-interpretation does not reveal the contents of our dispositions. Henceforth we concluded that the nature of these two kinds of objects of self-knowledge puts us in two different epistemic positions— one of relative lightness, and one of relative obscurity.

Even in less straightforward cases (like with Tasneem), a 'propositional picture of self-knowledge' would insist that there is a determinate fact of the matter about individuals' mental states– we could find out for sure if only we employed the right method. One can say this is a view of unconscious states as *discovery*. By contrast, an 'interpretivist' picture over these complex cases would claim that by the very act of bringing aspects of our belief to awareness, we actually resolve their vagueness. We *decide* to express ourselves in one way or the other in a linguistic medium— we articulate. For an interpretivist, the content of the worldview can be made into the content of a conscious judgment through verbal expression, either spoken aloud or internally. The verbal articulation of belief renders it more defined. In conscious thinking, 'the sharp delineation of determinacy in our verbal output replaces the vagueness of our convictions' (Dennett, 1978).

On my view, both propositional and interpretivist views are similarly normative in the sense that they both attempt to determine via linguistic articulation what is unconscious. A propositionalist suggests that there are determinate facts about unconscious mentality, and we just need to employ some method for discovery. An interpretivist denies that there are such facts to be "found", but he does offer such a method of interpretation— albeit for 'transcendental' articulation, not 'immanent' discovery. One may wonder, then, if there could be a third way for some cases, beyond decisions and discoveries, where one could honour indeterminacy— and claim *not knowing* as a legitimate epistemic practice.

3.5 Determining the indeterminate

In this chapter, I have exposed the shortcomings of the propositional attitude picture, and I have sided with a worldview theory in terms of the description of unconscious content. However, I insist that an interpretivist stance cannot be fully right (or tell the whole story) in all cases, at least from the point of view of self-knowledge. Indeed, self-interpretation raises the same basic issue about privileged access in introspection.

Recall that we have stated that what is unconscious is, often, rather indeterminate— certainly, we do not know how many beliefs and desires we have. Language, in contrast, follows determinacy: after all, one can only say one word at a time— specifying *this* belief, or *this* desire, attempting to make it all consistent as a whole. But in intricate scenarios, numerous belief-ascription questions may not have clear-cut yes/no responses. Perhaps, answers to complex questions such as Tasneem's should not be given by self-interpretation; that is, by determining what one believes, desires, etc. in verbal articulation. Perhaps answers to many of these questions need not to be provided, that is, one need not always to determine what is indeterminate— and instead can opt to honour indeterminacy.

3.5.1 The limits of interpretation

The Worldview theory holds that the ways that things are represented in the conscious and the unconscious mind are different. To model an aspect of the unconscious with a proposition, then, is to interpret it in consciousness— where

content appears relatively specific, with clear distinctions between beliefs, desires, intentions, and the like. The interpretivist would say, then, that the complexity of unconscious content can be optimally articulated through language expression this way, though interpretation.

I put the scope of this latter ambition into doubt. I do precisely because conscious determinacy is only linked to conscious content, and this link is indeed via our linguistic capacities. As such, the determinacy which we impose in our unconscious as we interpret it is a consequence of the determinacy of the language we use to describe it. At a closer examination, then, there appears to be a tension between attempting to model something with a tool that has some seemingly incompatible features with its object of judgement. One can very well be forcing linguistic forms into something that does not fit in a given linguistic shape. I take this to suggest that there must be some epistemic limitations to self-interpretation, in virtue of how the mind is thought to be structured (chapter 1).

Now this is all very abstract, but someone who is more cautious about self-interpretation would highlight the concrete dangers of *over*-interpreting in many complex cases, where seriously attempting to model with words —that is, to articulate— could distort the true nature of the person's unconscious. A familiar danger is confabulation, understood broadly as the phenomenon of confidently attributing oneself mental states that one does not actually possess, or did not occur. There are many theories of confabulation (Stammers & Bortolotti, 2020; Siderits, 2021), but for my purposes I would to highlight what I deem its centrally problematic epistemic

feature: the fact that they are post-hoc reconstructions, or rationalisations, that play a justificatory —versus a truly explanatory— role for one's behavior. If one often lacks privileged access to one's actual dispositions, then we should not be supposed to rely on plausible sounding, reconstructive reasons every time we want to "explain" why we behaved the way we did. Indeed, misattributions would result precisely from the subject's own strives to make sense of their puzzling actions in terms of pre-established mental states.

Of course, as previously acknowledged, there are many straightforward cases in which interpretation does offer an explanation for one's behavior that harmoniously aligns with the rest of our mental states and actions. But then the interpretivist bears the onus of offering criteria to distinguish between possible scenarios: i.e., what are the special conditions or circumstances in which an interpretation meets an explanatory versus a merely justificatory role? Could this burden ever be met? As it stands, I see no reason to think that mere introspection could enable us to detect whether a particular interpretation is confabulation or not. This brings me to a position where the epistemic competency of self-interpretation is put into doubt, or at the very least into a salutary place of probation.

It appears to me that a more faithful description of myriad complex behaviors involves a worldview that contains an extensive assortment of dispositions related to everything. These dispositions may not be internally consistent, and their piecemeal interpretation could be simply intractable. We should allow that answers to many questions may well be indeterminate — *"Why do you like her?", "How did you figure*

this out?", "Does Tasneem actually believe Ohan is an exceptional chef?, "What brings you to philosophy?". This need not result from the absence of a fact of the matter about my mental states. Rather, it may result from there being too many. Complexity, conflicting information, and confusion permeate the depths of the worldview. Any faithful representation of, let's say, belief, must accommodate or negotiate this immanent uncertainty and contradiction.

3.5.2 The limits of language

In previous sections, I have explored how various paradigmatic features of conscious mental content (SPECIFICITY, DETERMINACY, and CONSISTENCY) are not shared by the representational content of what is unconscious. I have also pointed that under a worldview theory, the determinacy which we impose on our unconscious —as we interpret it— is often a consequence of the determinacy of the language we use to describe it. In this last section, I will speculate on some related concerns related to this point, and I will suggest some topics for further research before drawing some general conclusions on an epilogue.

Per previous chapters, the content of unconscious states (e.g., dispositional belief) is categorically different from the content of correlated conscious episodes of thinking (e.g., judgment, assertion, professed belief). That is to say, the truth conditions for unconscious content and those of conscious content will differ. But for our aspirations at self-interpretation, this is all a bit daunting: content determinacy is said to be achieved through linguistic mechanisms— but we use language only when

we are conscious! If this is right, then language seems a properly determining factor only for the content of conscious judgements.

The "translation" of vast unconscious content into language as such would involve moulding it to conform to representations of a fundamentally different kind. This process, when carried out in complex cases, could distort the states of affairs in our minds. By pushing for an explanation in existing mentalistic idioms, in one direction or another, we may be failing to recognise the inarticulable nature of answers to many questions we deem as important. As it is unlikely that unconscious cognition conforms to the structure of propositionality, or fully determinate dispositions, I also take that it is unlikely to conform to the structure of language itself. Standing mental states themselves do not take place “in” language, but in a “deeper” conceptual, syntactic, or functional structure[10].

Future research will scrutinise the spectrum of the connection of language and thought (in the broadest sense), which I would say spans from the constitutive to epiphenomenal. Indeed, if content determinacy is to be achieved through linguistic mechanisms, the specifics of inner speech need to be further investigated. From the perspective of self-knowledge, a key feature of inner speech would be how it allows

[10] In the universe of cognitive science, this can be mentalese, language of thought, connectionist networks, Bayesian approaches, predictive frameworks… There are some existing, and there will be future options. One can remain agnostic of these specific vehicles. The point is that, when I sit and I judge; that is, self-interpret my beliefs and desires, what it is figuring is not my “beliefs” and “desires”— these are under a different medium of representation. We do not believe or desire “in” language. When we introspect, however, English and Spanish sentences are figuring again, but the believing and the desiring are happening in another medium of representation.

individuals interpret their experiences and dispositions. How is thinking carried out in inner speech, and why does this matter for self-interpretation and introspection?

Some view language as constituting mental acts, suggesting that inner speech is vital for self-interpretation and thought expression. This perspective would entail that language not only serves interpersonal communication but also shapes intrapersonal cognition, offering somewhat direct access to mental content, at least in some cases. In contrast, an epiphenomenal view would posit that language primarily converts thoughts for communication, while deeper cognitive processes occur largely independently of it. Here, language would serve as a tool for expressing thoughts rather than driving cognitive processes, offering no direct access to the mental content. As a result, inner speech's epistemic capacity for self-understanding would diminish, as language is seen as a vehicle for external expression rather than a direct pathway to the content of thought.

Propositional and interpretative perspectives presume that an underlying structure akin to (mentalistic) language exists —or optimally align— within the unconscious states they reference or seek to interpret. In the propositional case, it is expected that there is an unconscious standing state in the form "I believe that x" and that this coincides with the conscious recognition of the truth of x. In the interpretivist case, it is hoped that, after some reflection, Tasneem concludes that Ohan is a great cook, and this gives a far more determinate character to Tasneem's state of mind than it might previously have had. But taking this conscious judgement as the "last word" that settles the matter ignores many of other Tasneem's

dispositions that rather point towards different explanatory directions. Tasneem's worldview involves a vast collection of dispositions in relation to Ohan, some of which do not consistently align, which involve such vivid complexity that it might not be feasible to articulate them into a single word or sentence.

If there is no straightforward alignment between our mentalistic language and the actual psychological reality that validates their truth, then I challenge continued adherence to an articulative structure mirroring a belief/desire model. Interpretation will not honour an adequate explanation in cases where we have no clear intuitions, and embracing indeterminacy seems to me the most fitting approach to make sense of such instances.

Recall in chapter 2 that the only reason Freud could find for mental content not becoming conscious is that it has been in this way censored or 'repressed'. This work offers another more general reason: there seems to often be a lack of appropriate informational paths between unconscious content and introspective mechanisms. That is, a disconnect between the form that consciousness takes in introspection (i.e., linguistic) and the form responsible for taking the world to be in a certain way (more abstract structures such as connectionist networks, Bayesian systems, etc.). The actual mental states occur somewhere else, under a different medium of representation, calling for a different sort of language— that often times manifests as silence.

4. EPILOGUE: *Virtuous self-ignorance*

By writing this work on self-knowledge (and its challenges), I notice I have previously and largely started with the assumption that, in order to know ourselves, our various behaviors shall be given an etiological examination— at least in principle. That is, for us to know ourselves reasonably well, we are expected to know *why we do what we do.* As such, it is often insightful to speak of our behaviors (and others') as having certain causes, and there is some sense in which identifying these causes forms part of knowing ourselves.

'Causes' here is not an unambiguous term. Indeed, what a good etiological examination of a given behavior entails is not always clear. Think about the myriad and often incompatible ways one can theorise over questions such as '*What does Tasneem really think? Why did she take that path?*', '*Why didn't he come*?', '*Why did I say that?* ' What it is clear is that when humans try to do this —to explain a behavior to themselves, and to others—, they look for a set of psychological causes that are picked out by a particular mentalistic language that includes terms for the things we reviewed in chapter 1: events (perceiving, remembering, reasoning, deliberative thinking…) and states (beliefs, desires, intentions, goals…). People attempt to explain their behavior by appealing to these mentalistic notions.

By writing this work, I also notice that I have also initially assumed that this mentalistic language should have had at least some degree of similarity to a scientific idiom. Another way to put this commitment (arguably, a more controversial one) is

that I took references to mental states to be construed *realistically* (see below). Mental states are the *bases* of our dispositions to act and to have certain other mental features. In addition to the relevant dispositions, however, I impose certain constraints on the internal mechanisms and component structures that realises those dispositions. It is not surprise then that, from my perspective, it cannot all be up to interpretation.

Mental states seem to be causally efficacious. This causal efficacy presumably stems from some component structure which interacts in systematic (albeit complex and largely obscure) ways with the structures of other mental states. I take that distinct states have distinct physical realisations which cause distinct thoughts and behaviors, and I am aware this is a contentious programme. As such, for the most part, it is an assumption on my part to take representation thus construed "realistically", or physically realised in the brain. Those who would object to this assumption may instead read the various questions that this thesis raises *conditionally*: e.g., supposing that there really were representations physically realised in the brain, what difference (if any) would there be between those that are conscious and those that are unconscious? How do we get to know them? How are they related to language? How does this matter for self-knowledge?

Some philosophers have famously argued that mentalistic terms may not seamlessly fit into a scientific framework (e.g., Churchland 1981; Stich 1983). In fact, there is a sense in which much of the present work offers further reasons to doubt a smooth integration, in that it posits states and processes which representational

medium escape verbal articulation. However, as I have previously recognised, the project of self-knowledge as explored in this work does arise at least in part from the expectation to explain one's behavior by appealing to mentalistic terms. Seen from that perspective, this thesis offers a case study opportunity to further reflect on the extent to which a notion of self-knowledge —that which includes unconscious mentality— either challenges or broadens the frameworks developed by philosophy, psychology, psychoanalysis, the cognitive sciences, neuroscience.

The idea that one often cannot verbalise one's unconscious is an idea that is found among artists (e.g., Dali, 1942/1993), psychoanalysts (Jung, 1969/1981), mystical traditions (Watts, 1957), and scientists (Tenenbaum et al 2011) who seek different mediums of representation beyond language. However, it is not immediately clear how to state this thesis precisely. I have attempted to dissect this piece of wisdom within the framework of intentionality. I have described how unconscious mental content is often less specific, coherent, and determinate than conscious content— which may be very well linguistic content when we are talking about introspection and self-interpretation. As such, I take to have exposed a tension between our attempts to model something with a tool that has some seemingly incompatible features with its object of study. We may be trying to sculpt something with a knife that hasn't been sharpened for such a task.

What is unconscious is holistic, non-linguistic, mostly opaque to reason and deliberation. Yet the main way we use to explain ourselves (and others) involves a tool that has features of discreetness, determinacy, seriality. When we consciously

judge an unconscious state of ours, then, we may be thinking in an ill-posed medium. We may be exploiting certain properties of natural language and failing to acknowledge a lack of informational path, a leap in representation. Indeed, language is most likely not itself the form that mental states take; instead, these shall be often found on deeper associative (versus propositional) structures. I can take our unconscious psychological dispositions to be "realistically" representational, but claim that their structure do not neatly fall into the category of our commonsense notions of 'belief' or 'desire'.

This has all been a very abstract explanation over the limits of interpretation (that may very well be the limits of language), but we can appreciate this phenomenon concretely when we catch others (as it is too difficult within ourselves) confabulating. When we confabulate, illusions develop in thinking, as language is all we have to think with— and pay attention to. As such, in the governing egoism of the remaining mental states (e.g. our desires and motivations), the self-attribution of conscious mental states often take the form of confabulation. This is an immanent obstacle to the knowledge of ourselves.

Introspection (and self-interpretation) often delivers a psychologically compelling, but misleading, view of our mental life. And it is my hypothesis that this has a lot to do with the properties of language. Hence I advise future research in how thoughts are carried in introspection: i.e. if this form is linguistic (or it explores certain properties of natural language), then there must be something to further say

about how to connect with the deeper vehicles of our standing, non-linguistic unconscious mental states.

We delegitimise concepts for which no precise words exist, the uses of an identical word invites us to presume that we are dealing with identical concepts, and we reify abstract words as if they had referents in the real world. It is often the case that we feel we cannot verbalise how we ultimately get from one step in reasoning to the next, or how we reach a certain decision, but then after gaps in reasoning all seem perfectly plausible as they are supported by language: we are able to confabulate, rationalise, etc.

Natural language can obscure the understanding of some of our mental states and dispositions this way. It is often the case that even when some our mental states are indeterminate (i.e there is not a fact of the matter over we believe or want something), the determinacy of linguistic practices forces us to articulate it in a propositional fashion. "Why do you like her?", "How did you figure this out?", "Why did you move to Vienna?", "What brings you to Philosophy?" These are questions that call for modelling some of our unconscious attitudes through articulation. However, are these things we can know? Are these things we can talk about? Are they determinate enough for them to be articulated? Or we shall we remain virtuously ignorant of them?

This work lies in the intersection of philosophy of mind and epistemology. While philosophy of mind seems to me in principle descriptive, epistemology is often normative or prescriptive. Accordingly, I would argue that the fact that we cannot

often articulate our own unconscious mental states means that, in some thorny cases, we should not even try— and instead seek for other resources. It is the case that when we try to answer many complex questions, we are making something determinate which by nature it is not. That is, we are telling much more than we can know, and this is a *sui generis* form of falling in the realm of falsity. What I am suggesting, instead, is a *sui generis* form of virtuous ignorance.

There are cases in which precise information is practically or actually unknowable. In such contexts, individuals who take pretensions of certainty and self-assurance indicate incompetence, while voicing uncertainty implies expertise. I think that this is indeed the case for knowing many of our complex dispositions. Indeed, when we expand our conception of the mental beyond what is conscious, the quest for self-knowledge becomes much more challenging, ceding terrain for the mindful wisdom of self-ignorance. Sometimes, saying less than we can know may not signal further knowledge, but wisdom: virtuous ignorance.

5. References

Berger, J. (2014). Mental States, Conscious and Nonconscious. *Philosophy Compass*, *9*(6), 392–401. https://doi.org/10.1111/phc3.12140

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, *18*(2), 227–247. https://doi.org/10.1017/S0140525X00038188

Churchland, P. M. (1981). Eliminative Materialism and the Propositional Attitudes. *The Journal of Philosophy*, *78*(2), 67–90. https://doi.org/10.2307/2025900

Crane, T. (2017). The Unity of Unconsciousness. *Proceedings of the Aristotelian Society*, *117*(1), 1–21. https://doi.org/10.1093/arisoc/aox001

Crane, T. (2019). A Short History of the Philosophy of Consciousness in the Twentieth Century. In A. Kind (Ed.), *Philosophy of Mind in the Twentieth and Twenty-First Centuries: The History of the Philosophy of Mind* (Vol. 6). Routledge. https://www.taylorfrancis.com/chapters/edit/10.4324/9780429508127-4/

Crane, T., & Farkas, K. (2022a). Mental Fact and Mental Fiction. In T. Demeter, T. Parent, & A. Toon (Eds.), *Mental fictionalism* (pp. 303–319).

Crane, T., & Farkas, K. (2022b). The Limits of the Doxastic. In T. Crane & K. Farkas, *Oxford Studies in Philosophy of Mind Volume 2* (pp. 36–57). Oxford University Press. https://doi.org/10.1093/oso/9780192856685.003.0002

Dali, S. (1993). *The Secret Life of Salvador Dalí* (Reprint edition). Dover Publications.

Dennett, D. C. (1978). *Brainstorms: Philosophical essays on mind and psychology* (1st ed). Bradford Books.

Dennett, D. C. (1983). Beyond Belief. In A. Woodfield (Ed.), *Thought & Object: Essays On Intentionality*. Oxford University Press. https://doi.org/10.2307/2218897

Descartes, R. (1641/2017). Second Meditation: The nature of the human mind, and how it is better known than the body. In J. Bennett (Trans.), *Meditations on First Philosophy in which are demonstrated the existence of God and the distinction between the human soul and body* (pp. 3–8).

Dienes, Z., & Seth, A. K. (2022). Conscious and Unconscious Mental States. In *Oxford Research Encyclopedia of Psychology*. https://doi.org/10.1093/acrefore/9780190236557.013.804

Donayre Pimentel, S. (2023). The Unconscious Mind in Analytic Philosophy of Mind: Tracing the False Starts of the 20th Century Thought on the Nature of Unconscious Mentality. Ms, Central European University.

Farkas, K. (2008). Phenomenal Intentionality Without Compromise. *The Monist*, *91*(2), 273–293.

Freud, S. (1900/1953). The Interpretation of Dreams (First Part). In J. Strachey, A. Freud, C. L. Rothgeb, & A. Richards (Eds.), *The standard edition of the complete psychological works of Sigmund Freud: Vol. IV*. Hogarth Press.

Freud, S. (1915/1963). Introductory Lectures on Psycho-Analysis (Parts I & II). In J. Strachey, A. Freud, C. L. Rothgeb, & A. Richards (Eds.), *The standard edition of the complete psychological works of Sigmund Freud*: Vol. XV. Hogarth Press.

Freud, S. (1917/1963). Introductory Lectures on Psycho-Analysis (Part III). In J. Strachey, A. Freud, C. L. Rothgeb, & A. Richards (Eds.), *The standard edition of the complete psychological works of Sigmund Freud*: Vol. XVI. Hogarth Press.

Freud, S. (1923/1961). The Ego & the Id. In J. Strachey, A. Freud, C. L. Rothgeb, & A. Richards (Eds.), *The standard edition of the complete psychological works of Sigmund Freud*: Vol. XIX. Hogarth Press.

Freud, S. (1936/1964). New Introductory Lectures on Psycho-Analysis & Other Works. In J. Strachey, A. Freud, C. L. Rothgeb, & A. Richards (Eds.), *The standard edition of the complete psychological works of Sigmund Freud*: Vol. XXII. Hogarth Press.

Galeotti, A. E. (2018). Investigating Self-Deception. In *Political Self-Deception* (1st ed.). Cambridge University Press. https://doi.org/10.1017/9781108529242

Hobson, J. A., Pace-Schott, E. F., & Stickgold, R. (2000). Dreaming and the brain: Toward a cognitive neuroscience of conscious states. *The Behavioral and Brain Sciences*, *23*(6), 793–842. https://doi.org/10.1017/s0140525x00003976

Horgan, T., & Kriegel, U. (2007). Phenomenal Epistemology: What Is Consciousness That We May Know It so Well? *Philosophical Issues*, *17*(1), 123–144. https://doi.org/10.1111/j.1533-6077.2007.00126.x

Jackson, F. (1982). Epiphenomenal Qualia. *The Philosophical Quarterly (1950-)*, *32*(127), 127–136. https://doi.org/10.2307/2960077

James, W. (1890). The principles of psychology (Vol. 1). Henry Holt and Company.

Jung, C. G. (1969/1981). *The Archetypes and The Collective Unconscious* (R. F. C. Hull, Trans.; 2nd edition). Princeton University Press.

Kihlstrom, J. F. (2000). The psychological unconscious. In E. Tulving & F. I. M. Craik (Eds.), *The Oxford handbook of memory* (pp. 37-50). Oxford University Press.

Moran, R. (2001). *Authority and estrangement: An essay on self-knowledge*. Princeton University Press.

Robinson, H. (1982). *Matter and sense: A critique of contemporary materialism*. Cambridge University Press.

Schwitzgebel, Eric (2010) 'Acting contrary to our professed beliefs, or the gulf between occurrent judgment and dispositional belief ' Pacific Philosophical Quarterly 91: 531–553. https://doi.org/10.1111/j.1468-0114.2010.01381.x

Searle, J. R. (1983). *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press. https://doi.org/10.1017/CBO9781139173452

Siderits, M. (2021). Self-Knowledge and Non-Self. In M. Siderits, C. Keng, & J. Spackman (Eds.), *Buddhist Philosophy of Consciousness: Tradition and Dialogue* (pp. 189–208). Brill Rodopi.

Stammers, S., & Bortolotti, L. (2020). Introduction: Philosophical Perspectives on Confabulation. *Topoi*, 39(1), 115–119. https://doi.org/10.1007/s11245-019-09668-z

Stich, S. P. (1983). *From folk psychology to cognitive science: The case against belief*. MIT Press.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to Grow a Mind: Statistics, Structure, and Abstraction. *Science*, *331*(6022), 1279–1285. https://doi.org/10.1126/science.1192788

Wakefield, J. C. (2018). Freud's theory of unconscious mental states: A response to Schneider. *The Journal of the American Psychoanalytic Association*, 66(2), 313-331.

Watts, A. (1957). *The Way of Zen*. Vintage Books.

Wittgenstein, L. (1953/2009). Philosophische Untersuchungen: Philosophical investigations (G. E. M. Anscombe, P. M. S. Hacker, & J. Schulte, Trans.; Rev. 4th ed). Wiley-Blackwell.

Woods, J., & Rosales, A. (2010). Virtuous Distortion: Abstraction & Idealization in Model-Based Science. In L. Magnani, W. Carnielli, & C. Pizzi (Eds.), *Model-Based Reasoning in Science and Technology: Abduction, Logic, and Computational Discovery* (pp. 3–30). Springer. https://doi.org/10.1007/978-3-642-15223-8_1

Wundt, W. (1896). Grundzüge der physiologischen Psychologie [Principles of physiological psychology] (6th ed.). Wilhelm Engelmann.