In *Inference and Consciousness*, eds. Timothy Chan and Anders Nes, pp. 15-39, Routledge, 2020

PREPRINT

Unconscious Inference Theories of Cognitive Achievement

Kirk Ludwig and Wade Munroe

Philosophy Department

Indiana University

The Ethiops say that their gods are flat-nosed and black,
While the Thracians say that theirs have blue eyes and red hair.
Yet if cattle or horses or lions had hands and could draw,
And could sculpt like men, then the horses would draw their gods
Like horses, and cattle like cattle; and each they would shape
Bodies of gods in the likeness, each kind, of their own.
–Xenophanes

1. Introduction

What explains cognition and perception? What explains how the world looks to us, why we are subject to systematic illusions? What explains our capacity to speak and to understand language? The ultimate infrastructure for personal level cognition and perception lies in the physical construction of our bodies. But description at the level of fundamental physics doesn't promise much insight into the mechanisms of cognition and perception. We want rational insight, so to speak, into the infrastructure of cognitive achievement. We are tempted to seek an explanation not in terms of a family of concepts disjoint from those under which we bring the explananda but in terms of the same or allied concepts.

How does a research team solve a problem that none of its members is able to solve alone? It institutes a division of labor, in which different members of the team carry out different portions of the task, drawing on complementary skills and knowledge. When we have a specification of the division of labor, the subtasks and the processes by which they were carried out, and the organization of the team members that explains how their several contributions are combined into a complete solution, we understand the mechanism by which the research team solved the problem. This gives us rational insight into its cognitive achievement. This gives us one model for how to understand personal level cognitive achievement, namely, in terms of a division of labor into subtasks conceived of as problems to be solved at the subpersonal level. It promises rational insight into cognitive achievement, provided that the operation of these subpersonal processes can be characterized in terms of the same concepts as personal level cognitive operations, that is, in terms of rational problem solving or inferential reasoning, of some form or other. Since these processes are to explain personal level cognitive achievement, they are conceived of as being strictly subpersonal. Insofar, they are unconscious inference theories of personal level cognitive achievement.

This chapter considers the allure and prospects for unconscious inference theories of cognitive achievement (henceforth, 'UIT'). UITs explain various conscious, perceptual, and cognitive phenomena by postulating inference-like processes that operate over unconscious representational states. They subdivide into positions that hold (a) that these unconscious representational states and inference-like processes are in principle accessible to consciousness, and therefore are personal level states and processes, and (b) that they are strictly subpersonal and in principle inaccessible to consciousness (access or phenomenal (Block, 1995, 2002)). Our interest lies in the latter. UITs in category (b) subdivide into those that hold that the inference-like processes are (i) genuine inferences or (ii) not inferences but merely inference-like, inference facsimiles. We subdivide UITs in type (b)(ii) into those that hold inference facsimiles are defined over genuine representational states and those that hold they are defined over a theory-internal concept of representation. These divisions are represented in Figure 1.

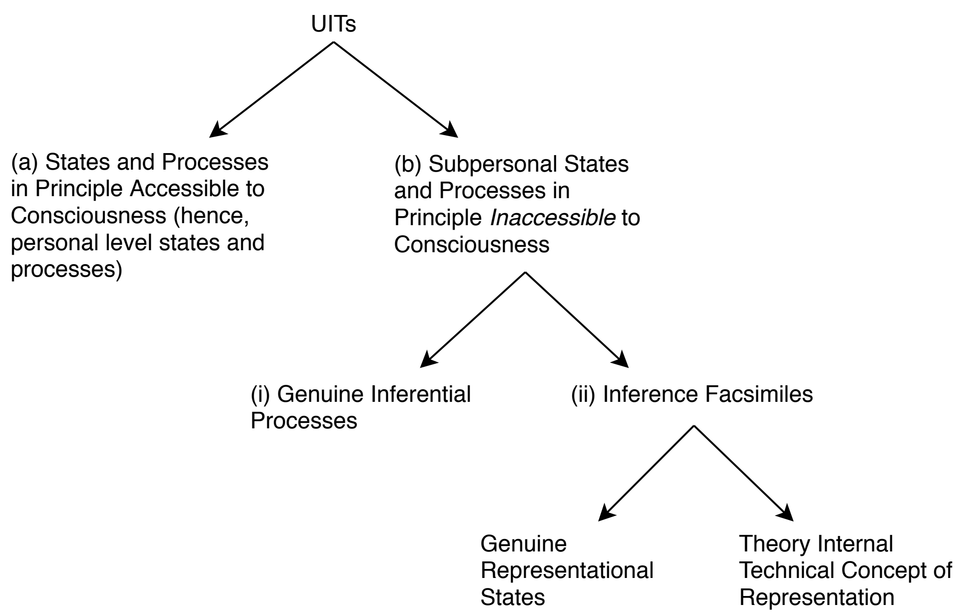

Figure 1: Unconscious Inference Theories

We argue that the only tenable UITs are ones that employ a theory internal technical notion of representation (lower right in Figure 1) but that once we give cash-value definitions of the relevant notions of representation and inference, it is difficult to see that much is left of the ordinary notion of representation. We suggest that the real value of talk of unconscious inferences lies in (a) their heuristic utility in helping us to make fruitful predictions, e.g., about illusions, and (b) their providing a high-level description of the functional organization of subpersonal faculties that makes clear how they equip an agent to navigate its environment and pursue its goals.[1]

In §2 we characterize the kinds of unconscious inference that we are concerned with. In §3 we review desiderata on what kinds of processes can count as inferences. In §4 we apply the desiderata to argue that there are no genuine modular subpersonal inferences. First, we argue that if they are inferences, they require a homunculus as their subject (§4.1). Next, we argue that the conditions required for this are not met by subpersonal modular capacities (§4.2). Finally, we argue that even waiving these points, UITs face a dilemma: they are committed either to an explanatory regress or the explanatory dispensability of unconscious inferences. In §5 we consider a retreat that merely requires inference facsimiles at the subpersonal level. We look at input-output representations (§5.1) and structural representations (§5.2) in Ramsey's sense (2007), and argue neither provide a genuine notion of representation suitable for use in a UIT. We then turn to cash-value definitions that make no pretense to connect with ordinary notions and suggest that they do not add new explanatory power, though talk of representations and inferences can play a useful heuristic role in theorizing about cognition and perception (§5.3). §6 summarizes.

2. Subpersonal Modular Inferences

Typically unconscious subpersonal inferences are treated as taking place in modular systems that serve narrowly defined functions. UITs of perceptual achievement (veridical representation of the environment) and linguistic understanding are paradigm examples. Although most of our discussion focuses on perception, the points carry over to other theories that treat information processing subsystems as inferential.[2]

While UITs of perceptual achievement have an ancient pedigree (Hatfield, 2002), contemporary theories trace their lineage back to Helmholtz (1867). Classic examples include

(Barlow, 1990; Brunswik, 1981; J. A. Fodor, 1983; R. L. Gregory, 1966; Richard L. Gregory, 1980, 1997; Irvin Rock, 1983; Irving Rock, 1984; Wandell, 1995). As Helmholtz puts it, the inferences that the perceptual system engage in

> … are in general not conscious, but rather unconscious. In their outcomes they are like inferences insofar as we from the observed effect on our senses arrive at an idea of the cause of this effect. This is so even though we always in fact only have direct access to events at the nerves, that is, we sense the effects, never the external objects. (1867, p. 430)

Similarly, according to Rock:

> Although perception is autonomous with respect to such higher mental faculties as are exhibited in conscious thought and in the use of conscious knowledge, I would still argue that it is intelligent. By calling perception "intelligent," I mean to say that it is based on such thought like mental processes as description, inference, and problem solving, although these processes are rapid-fire, unconscious, and nonverbal. "Description" implies, for example, that a perceptual property such as shape is the result of an abstract analysis of an object's geometrical configuration, including how it is oriented, in a form like that of a proposition, except that it is not couched in language. Such a description of a square, for example, might be "a figure with opposite sides equal and parallel and four right angles, the sides being horizontal and vertical in space." "Inference" implies that certain perceptual properties are computed from given sensory information using unconsciously known rules. For example, perceived size is inferred from the object's

> visual angle, its perceived distance, and the law that geometrical optics relating the visual angle to object distance. "Problem solving" implies a more creative process of arriving at a hypothesis concerning what object or event in the world the stimulus might represent and then determining whether the hypothesis accounts adequately for, and is supported adequately by, the stimulus. (1984, p. p. 234)

More recently, UITs have been given new life by the idea that the brain is a predictive engine, and, more specifically, a Bayesian reasoner, which engages in probabilistic inference about the hidden causes of sensory input with the goal of reducing sensory prediction errors (Clark, 2016; Hohwy, 2013).[3] As Hohwy puts it: "The brain infers the causes of its sensory input using Bayes' rule" (2013, p. p. 18).[4] According to Clark, "the predictive processing story, if correct, would rather directly underwrite the claim that the nervous system approximates a genuine version of Bayesian inference" (2016, p. p. 41). Rescorla notes that the inferences involved are strictly subpersonal:

> Perceptual processes are subpersonal and inaccessible to the thinker. There is no good sense in which the thinker herself, as opposed to her perceptual system, executes perceptual inferences. For instance, a normal perceiver simply sees a surface as having a certain colour. Even if she notices the light spectrum reaching her eye, as a painter might, she cannot access the perceptual system's inference from retinal stimulations to surface colour. (2015, p. p. 695)

At the level of our discussion, differences between classical and Bayesian inference theories will not be significant. The problem lies in the transference of concepts (i.e., the concepts of inference and representation) from one domain to another without taking seriously the conditions for their application. Details with respect to the nature and content of the postulated inferences make no difference.

Let's look at how a UIT might explain perceptual constancy. Perceptual constancies are described by a function that yields a constant value (perceptual representation) while its arguments (sensory stimuli) change. When the value is constant while inputs change, we have the representation of sameness of size, shape, color, etc. through changes of proximal stimulus. The UIT strategy is to explain how the perceptual system achieves constant representation of the relevant property by giving it knowledge of the function and knowledge of the appropriate arguments. For example, Emmert's Law states that the perceived linear size of an object is proportional to the product of its perceived distance and the angle subtended on the retina. There are analogues for constancy of represented shape through rotation relative to the observer, constancy of lightness and color through variations in illumination conditions, constancy of position relative to movement of the perceiver, and so on. The perceptual system gets information about, e.g., perceived distance (inferred from more basic cues) and the angle subtended by an object in the visual field and then infers using Emmert's Law a size which is to be represented in perceptual experience. Inferential mechanisms are also used to explain illusions. For example, in the Ponzo illusion illustrated in Figure 2, the black bars are the same length but the depth cues provided by the receding track generate a visual representation of the upper bar as longer than the lower.

Figure 2: Ponzo Illusion

For our purposes, the key features of the supposed inferential processes involved in UITs are that:

(1) they operate over representations that bear semantic relations to one another;
(2) they are modular in that they are relatively autonomous from personal level cognition, intention, belief, and reasoning;
(3) they are postulated to explain specific perceptual and cognitive capacities;
(4) their inputs are paradigmatically not personal level cognitive states but subpersonal representations, so that they are not conceived of simply as mediating personal level cognitive states as input and conscious output, and
(5) they are not in principle accessible to the person whose cognitive and perceptual capacities they subserve, because their functional role is to precede and to explain modifications of consciousness.

3. What are inferences?

In this section we review the central desiderata on an account of personal level inference (PLI). We identify features generally accepted as necessary for inference. This sets the stage for asking whether subpersonal modular processes subserving personal level cognition are plausibly thought of as inferential.

3. 1 Conditions on a successful account of PLI

PLI involves a transition from one set of propositional attitudes (e.g., beliefs, intentions, suppositions) to another. Inferences can terminate in a new attitude—a new intention in practical inference or belief in theoretical inference. Alternatively, they can terminate in an alteration of current attitudes—for example, an inference may result in relinquishing an intention or strengthening a belief. Attitudes whose contents support the attitudinal shift are *premise attitudes*, the new or altered attitudes that result are *conclusion attitudes*.

Practical and theoretical inferences are distinguished by the attitude types of their premise and conclusion states. For theoretical inference, the premise and conclusion attitudes have mind-to-world direction of fit. For practical inference, two kinds of attitudes are required, those with mind-to-world direction of fit (beliefs) and those with world-to-mind direction of fit (preferences). In practical inference, means-end beliefs provide premises, while preferences provide comparative evaluative judgments, such as that vanilla ice cream is better than chocolate.

An inference must involve states that have modes appropriate for its type. For example, a sequence of wishes whose contents are related by a logically valid argument form is not an inference, despite logical relations between their contents. We focus on theoretical rather than practical reasoning, in which transitions occur between states with mind-to-world direction of fit. This excludes treating modeless representations as figuring in inferential processes.

It is widely accepted that a successful account of PLI must satisfy the following three desiderata (P. Boghossian, 2014; Broome, 2013, 2014):[5]

(1) It must distinguish PLI from other types of transition in thought.

Not just any transition between propositional attitudes (even of the correct types) constitutes an inference. For example, an associative shift between two attitudes is not an inference. A change in attitudes must be caused 'in the right way' to constitute inference. Therefore, the fact that the attitude contents have logical relations between them suitable for epistemic support is not sufficient for a mental transition to constitute an inference. For example, although *A* follows from *A and B*, one might come to believe *A* from believing *A and B* by associating *A* with *B* rather than inferring it from *A and B*.

(2) It must allow that one can make a mistaken or non-normative inference, that is, one which still counts as an inference despite its not being a good inference.

Roughly, a mistaken inference is one in which the premise attitudes provide little or no support for the conclusion attitudes. People do not always reason correctly. An adequate analysis of PLI must allow that individuals can make inferences even when they commit the base-rate fallacy or affirm the consequent. So the fact that someone's attitude transitions do not conform to Bayesian norms on belief updating (e.g., conditionalization) does not entail that the person has failed to infer. Putting this together with (1), descriptively conforming to, e.g., Bayesian norms in attitude transitions is neither necessary nor sufficient for engaging in inference.

(3) It must explain how PLI is something we do and not merely something that happens to or within us.

We make inferences and *we* update our attitudes in doing so. An adequate account of PLI must allow for inference to be a controlled process as opposed to something that merely happens to a person.

Much of the contemporary literature on PLI focuses on what Paul Boghossian (2014) calls the Taking Condition (TC):

> [TC] Inferring necessarily involves the thinker taking her premises to support her conclusion and drawing the conclusion because of that fact.[6]

TC is regarded by many as central to satisfying conditions (1)-(3). TC distinguishes inference from other thought transitions by requiring that inference involves the reasoner *taking* her premises to support her conclusion and drawing the conclusion in virtue of this (desideratum 1). An associative transition does not depend on one's taking there to be an epistemic support relation between one's premises and conclusion. Furthermore, TC makes room for incorrectly taking one's premises to support one's conclusion. In consequence, one can count as inferring a conclusion even when it is not supported by one's premises (desideratum 2). Finally, because TC implies that a transition in thought only constitutes an inference when it is responsive to its subject taking the premises to support the conclusion, inference is properly something that can be attributed to the subject, rather than something that merely happens to or within her (desideratum 3).

TC requires clarification. A full account would need to explain:

(1) What *taking* one's premises to support one's conclusion consist in, e.g., whether taking is an intentional state like belief or intuition, or whether it is something like a disposition to judge that one's premises support one's conclusion,[7]

(2) What the content of the taking is,[8]

(3) What it is to *draw* one's conclusion in virtue of taking one's premises to support it, that is, what is required beyond the premise attitudes causing the conclusion attitudes.

Most contemporary work on inference either (i) focuses on answering one (or more) of the above questions or (ii) challenges TC as a necessary condition of inference. However, extant objections to TC assume that taking constitutes an intentional or representational state (McHugh & Way, 2016; Rosa, 2017; Wright, 2014). Thus, these are not objections to TC per se but to a certain way of explaining it.

The dominant account treats inference as a causal process that constitutes rule-following. For example, Boghossian (2014) adopts a rule-following account in which inferring just is following a rule of inference in moving from a set of premise attitudes to a conclusion attitude. Roughly, we can think of a rule as an instance of the following schema:

> If antecedent conditions, C, hold, then it is permitted/required to do/accept/believe A.[9]

In a rule-governed theoretical inference the antecedent conditions will be the possession of certain premise attitudes, and the rule will indicate which attitude(s) one is then permitted/required to adopt. Boghossian argues that in following a rule (rather than merely conforming to it) one takes the antecedent conditions as reason to perform the permitted/required action. So, if a person follows a rule of inference, nothing additional is needed for her to take her premises to support her conclusion. The taking falls out from the rule-following; no additional occurrent, intentional state constitutes the taking.

We treat TC as a fourth requirement on an adequate account of PLI, aimed in part at satisfying requirements (1)-(3). It expresses a relation of the reasoner to a process involving intentional states that makes sense of its being normatively appropriate from the point of view of the reasoner. We formulate this as follows:

> TC*: For a transition between premise states to conclusion states to count as an inference, their subject must take the premises to support the conclusion, whether this is a matter of an explicit intentional attitude with that content, the process being explained by the subject following a rule appropriate for the type of inference, or the subject being disposed to regard or act as if the inference is correct or justified.

We are deliberately non-specific to allow for various ways theorists have tried to make TC more precise. Notice that even if inference is construed as a matter of following a rule, when that is not understood in terms of a propositional attitude about the transition, since the rule is defined over the states involved in the transition, it requires the agent to be sensitive in its cognitive operations to the contents of those states.

4. Are there modular subpersonal inferential processes?

By subpersonal modular inferential processes (SMI) we mean processes which are genuinely inferential and not merely treated *as if* they were, or *modeled* by inferences (we return to 'as-if' talk in §5). We focus on the claim that there are modular processes involved in vision, language understanding, and other cognitive processes that (i) operate over states with intentional content and subserve personal level cognitive and perceptual processes by delivering appropriate personal level intentional states and (ii) constitute inferences on the part of the modular system itself as opposed to the personal level subject (PLS).[10]

In this section, we develop an argument against the plausibility and theoretical utility of SMI processes, construed as above.

(1) First, we argue that if inferential, SMI must be treated as inferences of a cognitive agent with propositional attitudes.

(2) Second, we argue that the conditions required to attribute SMI to subpersonal units conceived of as cognitive agents are not met.

(3) Third, we argue that even if conditions required to attribute SMI to subpersonal units were met, this would amount to a homuncular explanation of personal level cognitive achievement, and that, to avoid a regress, the explanation of the cognitive capacities of the homunculi would have to be given in different terms. Then the same style of explanation could be applied for personal level cognitive achievement, showing the homuncular explanation to be gratuitous. The explanation offered is therefore defective because it is in principle replaceable and there is no non-question-begging reason not to replace it at the first stage of explanation.

4.2 Modular Inferences Require Homunculi

If SMI are genuine inferences but not by the PLS, they must be inferences of a subpersonal level agent. We will say a *homunculus* is a subpersonal agent whose cognitive work subserves personal level cognitive achievements. Thus, if SMI are genuine inferences, they require homunculi.

First, for SMI to be genuine inferences, we must be able to make sense of their being taken to be correct by their subject. This requires us to conceive of their subject as taking intentional attitudes toward the inferences or at least to be disposed to take such attitudes toward them, or to be involved in rule-following of the sort that would support the idea that the subject takes a normative stance toward the relevant transitions. This is what it is for a cognitive agent to be making inferences. If the subject of SMI is not the PLS, it must be a homunculus.

Second, inferences involve transitions among propositional attitudes. Propositional attitudes have agents as their subjects. Moreover, SMI are theoretical inferences. Therefore, they require attitudes with mind-to-world direction of fit, that is, belief-like states.[11] The functional role of belief is to guide behavior, broadly construed, in the light of system goals. We can get a grip on states having mind-to-world direction of fit only if we are prepared to attribute to the system goals as well, and at least some form of rudimentary agency, in which its activities are directed in accordance with its beliefs and preferences. These are not personal level psychological states. They are not part of the psychological economy of the PLS. They therefore require a subpersonal agent.[12]

Third, the attribution of attitudes with contents presupposes that the concepts involved in the attitudes are possessed by their subject. Concept possession requires having the capacity to deploy concepts appropriately in relation to evidence and to reason in accordance with the requirements of their application conditions. Since these concepts are not and need not be possessed by the PLS in virtue of having the relevant modular capacity, as the case of children and non-human animals show, they are not concepts of the PLS—even if the PLS has the concepts independently. There must therefore be a distinct cognitive agent who possesses them.[13]

4.3 SMI are not Homuncular Inferences

The main argument against homuncular SMI is that attributions of inferential capacities require commitments that are not met by subpersonal processes subserving personal level cognition and perception.[14] We raise two problems, the holism of attitude attribution, and the holism of concept attribution.

First, inferences are not defined over representations but over attitudes with psychological modes appropriate for the forms of inference. In the case of theoretical inferences (about how things are) this requires a mode with mind-to-world direction of fit. But attitudes with mind-to-world direction of fit are part of a pattern that includes attitudes with world-to-mind direction of fit.

The reason that attribution of belief takes place in the context of attribution of desire and intention is that the canonical role of belief is to guide action in the light of preference. This is what gives us the idea that a state is a state *whose job it is* to represent something in the world as opposed to one that is merely lawfully correlated, like tree rings, with changes in the world. The difficulty with homunculi engaging in SMI is that there is no point in attributing to them any preferences or intentions, any more than there is to attributing preferences and intentions to trees. Trees do not engage in flexible goal directed behavior guided by representations of their environment. Neither do subpersonal cognitive faculties. We might attribute to a subpersonal module a *function*, relative to its contribution to cognition, but this is not to attribute a goal to the module itself. There is *no more* point to attributing goals to subpersonal modules that have functions subserving cognition and perception than to the heart or lungs or small intestines, all of which have biological functions as well.[15]

The second problem is connected. There are holistic constraints also on the attribution of concepts. The inferences typically attributed to modular faculties require sophisticated conceptual resources and reasoning capacities which there is no evidence that subpersonal modules possess, as opposed to those theorizing about them. This is one reason we do not want to attribute these inferences to PLS. The operations of subpersonal mechanisms subserving cognition are insensitive to whether the PLS possesses the competencies required by the concepts deployed in SMI. But then it is *even less plausible* to attribute these competencies to subpersonal agents that do not have the capacity to deploy these or even simpler concepts. Even as simple an inference as that involved in deploying Emmert's Law for linear size constancy requires geometrical concepts of angle, distance, size, and space, as well as the concept of equivalence and mathematical product—and this is just a beginning.[16]

Concept possession is constituted by competence in correct application. Vision theorists and linguists have these concepts because they can deploy them across different domains. Their attribution to theorists is supported by attribution of a range of supporting concepts, for vision theorists, of number, sum, cardinality, color, light, etc., and for linguists, of language, meaning, compositionality, rule, scope, domain, binding, etc. *None* of these *general* capacities can be attributed to subpersonal modules. The concepts attributed are only postulated to be deployed in a limited domain. No one thinks that the competencies required for possession of these concepts by the theorists who deploy them in describing SMI are possessed by subpersonal modules. But since the competencies are required to possess the concepts, the modules themselves do not possess the concepts. Therefore, they are not capable of performing inferences over contents involving them, since having the attitudes involved in the inference requires having the concepts they involve.

The attribution of SMI to subpersonal modules is a form of theoretical projection. If a vision scientist were to explain to someone how one might extract the relevant information present in, e.g., visual representation of the environment, from physical stimuli, given background knowledge of how the world works, she might use the sort of inferential account attributed to the visual system itself. Seeing the visual system as doing what the vision scientist is doing is supposed to make intelligible how the visual system does it: it does it just like that, like a vision scientist with tunnel vision, who cannot make any other inferences, who cannot think about anything in general, who cannot deploy the relevant concepts in any other domain. But the light at the end of the tunnel goes out when we see that having the concepts cannot be divorced from the general capacities that constitute competence in their deployment.

4.4 Explanatory Regress or Explanatory Dispensability

If we could find a subject for SMI, would we have an explanation of our cognitive achievements? The short answer is: Yes. But if we explain how cognition is possible in one agent by appeal to others engaging in cognition on its behalf, we have not explained how cognition as such is possible. It might be said that we can do better than this because we can explain the cognitive achievements of the subpersonal modules as well. But how? One might reapply the strategy of breaking the task down into subtasks performed by a second, deeper level of cognitive agents. Would this explain adequately the cognitive achievements of the first level of subpersonal cognitive agents? Yes … but only by postponing again the question of how cognition as such is possible.

One reply is homuncular functionalism (D. C. Dennett, 1978, p. p. 80): as we go down levels, we make an explanatory advance because the homunculi get successively dumber as they are given successively simpler inferences to make. [17] Yet even agents who are not very clever, if they are making even simple inferences, have to meet the holistic constraints on attitude and concept possession. Moreover, one needs to specify, at each level, exactly what the inferences are. It is unlikely that their content becomes less sophisticated as we go down the hierarchy, as they involve the concepts that theorists use to describe input and output. Consequently, simple inferences or not, we are postulating sophisticated cognitive agents.

One might reply that SMI were never intended to explain cognition as such. One could accept the force of the regress argument but still maintain that some explanatory progress has been made. This comes with a commitment to explain the subpersonal level cognition without adverting to further levels of subpersonal inferential processing. But now there is a dilemma. Suppose that putative SMI can be explained without appeal to further underlying subsubpersonal level inferences. We would then have an explanation of how cognition works which does not ultimately require appeal to cognitive operations. Why can't we apply this strategy at the first sublevel of processing? If we can, then the postulation of SMI is gratuitous because it is explanatorily dispensable. The only support that can be provided for it, since it is by hypothesis inaccessible to the PLS, is that if it were true, it would partially explain personal level cognition. Thus, if SMI are not explanatorily dispensable, they set us off on an explanatory regress. If they set us off on an explanatory regress, then they cannot provide an explanation of cognition as such. If the regress can be stopped, then SMI are explanatorily dispensable. Thus, SMI are either explanatorily dispensable or we cannot provide an explanation for cognition as such.

5. Inference facsimiles

Surely it is a mistake to suppose the sorts of inferences appealed to in UITs are intended to be inferences in the ordinary sense! Similarly, surely the representations over which unconscious inferences are defined were never intended to be ordinary representations. Thus, we do not need a subject of the inference who takes their premises to be support for their conclusions, and we do not need to worry about holistic constraints on attitude attribution or concept possession. From the standpoint of the working scientist, these criticisms are an example *par excellence* of the attempt to constrain the development of concepts appropriate for theoretical explanation to those developed in the domain of commonsense, which would frustrate the search for theoretical explanations across *every* domain in which science operates.

It is doubtful that all theorists who invoke unconscious inferences to explain cognitive achievements think of them as different from ordinary inferences in any respect other than being in principle unconscious and subserving personal level cognition and perception.[18] But a natural fallback is to suggest that the notions of inference and representation deployed in UITs should be understood in a different sense than the vernacular. On this view, to talk of "unconscious inferences" in the context of a UIT is not to talk about unconscious *inferences*, but about, as we will put it, unconscious inference facsimiles, like, in some respects, but not the same as, inferences. This leaves us with two questions. First, what is the content of such UITs, given their reliance on unconscious inference facsimiles, since we cannot rely on our antecedent understanding of 'inference' and 'content'? Second, in what does their theoretical utility lie: how are they to help us understand cognition and perception?

There is a hard and a soft line on the first question. The hard line maintains that while not subject to the usual holistic constraints, the states over which SMI are defined are genuine representations. The processes defined over them that subserve personal level cognition and perception are inferential insofar as the states in the processes bear semantic relations to one another that mimic inferential processes. Thus, the status of the processes as substantively inference-like depends on the states involved being genuine representations. The soft line relinquishes the idea that there need be anything of our antecedent notion of representation left and treats talk of representations as a proxy for something that could be explained without appeal to intentionality. We address the hard line first, then the soft line, which leads to the second question.

5.1 The Hard Line: IO-Representations

What constraints are there on genuine representation? Ramsey (2007) notes that it is not enough for the *theorist* to assign representational content to states, as when we treat voltage levels in transistors as representing 1 or 0. These make sense relative to tasks we design a machine to implement. It doesn't give the machine the task or intrinsic intentionality. The notions we appeal to, as Ramsey says, "must in *some* way be rooted in our ordinary conception of representation; otherwise, there would be little point in calling a neural or computational state a representation" (p. 25). But they can't be observer relative. We must make sense of the states having content *and* of their functioning as representations *for the system* containing them. Ramsey distinguishes two notions of representation for subpersonal cognitive processing that can be detached from propositional attitude psychology, input-output representation, or IO-representation (2007, sec. 3.2), and structural representation, or S-representation (2007, sec. 3.3). We deal with IO-representations in this section and S-representations in the next.

IO-representation applies to a system that already has representations as inputs and as outputs. If intervening processing can be explained by state transitions that, by an assignment of content to them, represent the process as involving semantically sanctioned transitions from input to output, then those internal states have IO-representations. Ramsey claims that IO-representations are genuine representations *for* the system, and so not merely observer relative. But there are two problems with this. First, there is no evident inconsistency in *denying* that the assignment of representations to internal states characterized neutrally captures something intrinsic to the system. For it is not inconsistent to claim that what mediates input and output is *simply causal-functional organization*. One might *stipulate* that *if* there is a semantic mapping, *then* mediating states are IO-representations. But this is a merely verbal maneuver and so not ampliative. Second, even if Ramsey were correct, IO-representation can't be applied to perceptual processing since it presupposes that both input and output independently have representational content—and although the output of perceptual systems is independently representational, the input is not (or not solely). Thus, this notion of representation for the perceptual system would rely on an observer relative assignment of representation to inputs to the perceptual system. This would give the intervening states IO-representational content, but they would be observer relative as well, and not representations *for* the system.

5.2 The Hard Line: S-Representations

S-representations are structural representations. The idea is illustrated by a map. Points on the map correspond to areas on the terrain (within a margin of error). The distance between points corresponds to the distance between the areas on the terrain. When we use a map, we exploit what we know about its structure and the mapping function to learn about the terrain. Put most generally, it is the idea of a modeling system consisting of elements (*m*-elements) and relations between them which are isomorphic to a target system with its elements (*t*-elements) and its relations in the sense that there is a mapping from m-elements to t-elements, and a mapping of *m*-relations to *t*-relations, such that for any *m*-relation, *r,* relating a sequence, *s,* of *m*-elements, the image of *r* in the target system relates the image of the sequence of *m*-elements in the target system. The image of a relation or element of the modeling system in the target system is what it is mapped onto. The idea is that subpersonal processes may extract information from models in this sense to guide what representations are produced at the personal level.[19]

The difficulty lies in the idea that a subpersonal process *extracts* information to *guide* a process. What gives substance to this idea? It is not that there is an isomorphism between elements of the system and something outside it. Isomorphism is not representation. The cars in one row in a parking lot may be isomorphic with those in the next. But neither row represents the other. Isomorphism is symmetric, representation is not. When we use a map to locate a restaurant, we are the ones who, by exploiting what we know of its structure and relation to a city, use it as a representation. This presupposes an agent who uses it as a representation for a purpose. For subpersonal processes hypothesized to exploit S-representations, however, there is no one who uses them. By hypothesis the PLS does not use them. And having set aside the appeal to subpersonal agents as implausible, unsupported, and pointless, there is no one else to use them either. We are left with a structure isomorphic with some bit of reality that plays a causal role in the production of appropriate personal level representations. We could define this as a S-representation! But, again, this is not ampliative. The adoption of the language of representation may give the impression of depth of explanation. But the definition accomplishes only an abbreviation.[20]

5.3 The Soft Line

This leads us from the hard line to the soft line. The hard line maintained that the states over which SMI are defined are genuine representations. The soft line treats talk of representations as a proxy for something that entails no commitment to genuine intentionality. But then why bother? The answer is that even if talk of representations and unconscious inferences plays no fundamental explanatory role, it can play a useful heuristic role. But what does it come to and how could it play a heuristic role?

The answer is that, to borrow an apt expression from Brunswik (1956, p. p. 141), even if to different purpose,

> if we can see processes subserving perception and cognition as ratiomorphic, we gain insight into *how* they perform a function serving personal level perception and cognition.

Processes subserving perception are ratiomorphic if (i) they have the formal or functional features of a bit of reasoning, under an appropriate mapping, but not the semantic content, and (ii) the processes having that structure, in the organism's environment, yield a largely effective updating of perceptual representations of its environment. This removes commitment to the processes involving genuine representations. But it keeps everything that is important for understanding. More precisely:

> A process is (thinly) *ratiomorphic* iff there is an isomorphism from its causal-functional structure to a system of rules and representations which shows how, from input described in a certain way, the system generates appropriate personal level representational output, where appropriateness is judged in terms of its general usefulness in guiding the consuming system's cognition and action, given its goals and purposes.

Why is it useful to think of subpersonal processes as ratiomorphic? There are at least four connected reasons.[21] (i) First, it gives us a way of thinking about the causal-functional structure of a process in terms of a familiar conceptual scheme with which we have great facility. It provides, in Egan's terms, a "function-theoretic" characterization of a mechanism subserving perceptual and cognitive capacities.[22] (ii) Second, it is an aid to discovery because it aids in thinking about the perceptual system from the design perspective. Thinking of processes as ratiomorphic helps us to see how the system could be structured to produce appropriately and dynamically changing output in response to stimuli given the environment and history of the organism's interaction with it. (iii) Third, it is an aid in making predictions because it involves adopting the intentional stance toward a subsystem, conceived as an oddly limited reasoner.[23] (iv) Finally, as Egan notes, it can "serve as a temporary placeholder for an incompletely developed computational theory of a cognitive capacity and so guide the discovery of mechanisms underlying the capacity" (2018, p. 13). This shows in what sense the function-theoretic structure identified is explanatory: it quantifies over realizations that implement it, providing insight into what the actual realizers contribute to the functioning of the system of which they are a part.

Cognition and perception are subserved by subpersonal processes. There are constraints on those processes given that they are supposed to deliver to the PLS, for the most part, accurate representations of the immediate environment. In the case of perception, this requires a causal-functional organization of the system that generates at the output a perceptual representation whose intrinsic nature reflects in its structure (even if structure alone is not sufficient for representation) a structure of similar complexity in the world (like a map and what it maps). The question which needs an answer is how the structure of the one is transmitted to the other.

When it is a design problem, we know what the target is, we know the nature of the environment, and we know what the input is. We can then seek to construct a mechanism that exploits structure in the input to transform it into the output we want, given the environment and a history of interaction with it. We understand how the system goes from physical input to a representation of the environment when we have an account of a mechanism that generates it from structure in the input. Since the input inevitably underdetermines the appropriate output, part of what we want insight into is how the system is structured to yield from input appropriate output. This requires something to be supplied by the mechanism that constrains the relation of input to output in a way that is sensitive to what is likely to be producing the input given the environment. This is what makes it apt for description as ratiomorphic.

If we think of the task as assigned to a person who has knowledge of general features of the environment and how the system is situated in it, and then is given knowledge of the input, we can think of an inferential process that would generate an appropriate output representation. This gives us a description of a functional-causal organization that will do the job. And if we implement the design in a physical system, then we will have an explanation of how that system does the job (assuming we have representations as output). What is crucial for understanding how the job is accomplished is not that there be representations and rules of inference in the system itself but only that its structure be isomorphic to a system of representations and rules of inference. Thus, thinking of the process as ratiomorphic (seeking to see it under a mapping) helps both (i) to formulate hypotheses about the functional-causal organization of a subsystem and (ii) to grasp it.

Once we have a hypothesis about a ratiomorphic structure, (iii) it can help us to make predictions. For example, from Emmert's Law we can predict that manipulating depth information will yield incorrect representations of object size, as is born out, for example, in the Ames Room Illusion. Thinking of the process as ratiomorphic makes the prediction particularly vivid because we think of someone deducing from incorrect premises a conclusion that follows from it. Conversely, thinking of illusions as generated by ratiomorphic processes provides additional clues to the structure of those processes. For example, the Muller-Lyer illusion, the Ponzo illusion (Figure 1), and the moon illusion provide clues to the functional-causal structure of the visual system, which we can seek to make intelligible from the design perspective, which encourages looking for ratiomorphic processes in the system.

Finally, what a hypothesis about ratiomorphic structure gives us is an account of the causal-functional structure of a mechanism relating input to the perceptual system (or its subsystems) and output, which explains, given ceteris paribus laws connecting features of the environment with input, why for the most part the output is appropriate for the organism. The causal-functional structure is a kind of mechanism sketch. The sketch is filled in by finding a realization of it in a lower level description of the system, and ultimately a description in terms of the neurophysiology of the brain. Thus, (iv) the hypothesis guides investigations into more detailed mechanisms underlying the functional relationship between environment, input, brain mechanism, and perceptual representation.

When the ratiomorphic approach is appropriate, representational talk doesn't add anything to our understanding of the nature of the process as such. Yet it does give us insight. It provides insight both into the causal-functional organization of the system that does the causal-structural translation job and into how it is fitted for the job that it does. The assignments make *perspicuous* to us how the system preserves or generates or selects *relevant* causal-structural information. It makes clear to us how it subserves a function for the system we explain in terms of goals or representations—but crucially it does so without our having to take seriously the idea that the mechanisms themselves have representational content.[24] In this sense the role is heuristic.

6. Conclusion

Serious use of the terms 'inference', 'content', 'representation', and 'concept' must pay attention to their application conditions or supply operational definitions. Attention to their application conditions makes clear that modular systems subserving personal level cognition do not engage in inferences, they do not involve, except in their output, representations, and they, as opposed to the system which they subserve, do not possess concepts. It is natural to respond by declaring that philosophy should not attempt to put a priori constraints on the development of theoretical concepts in the pursuit of scientific understanding. But that's not the point. If words are being used in their usual sense, we must respect their application conditions. If new theoretical concepts are being deployed, we must make clear what their nature is. When we provide operationalized definitions, it becomes clear that talk of inference and so on, is basically unrelated to the ordinary personal level notions, and supplies no explanatory power over what can be said without appeal to them, though the vocabulary retains a heuristic function.

REFERENCES

Barlow, H. B. (1990). Conditions for versatile learning, Helmholtz's unconscious inference, and the task of percetption. *Vision Research, 30*, 1561-1571.

Bennett, M. R., & Hacker, P. M. S. (2003). *Philosophical Foundations of Neuroscience*. Malden, MA: Blackwell Pub.

Block, N. (1995). On a confusion about a function of consciousness. *Brain and Behavioral Sciences, 18*(2), 227-–247.

Block, N. (2002). Some concepts of consciousness. In D. Chalmers (Ed.), *Philosophy of Mind: Classical and Contemporary Readings* (pp. 206-219).

Boghossian, P. (2014). What is Inference? *Philosophical Studies, 169*(1), 1-18.

Boghossian, P. A. (2008). Epistemic Rules. *Journal of Philosophy, 105*(9), 472-500.

Broome, J. (2013). *Rationality Through Reasoning*: Wiley-Blackwell.

Broome, J. (2014). Comments on Boghossian. *Philosophical Studies, 169*(1), 19-25.

Brunswik, E. (1956). *Perception adn the representative design of psychological experiments*. Berkeley, CA: University of California Press.

Brunswik, E. (1981). Levels of Perceptual Organization. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization* (pp. 255-278). Hillsdale, N.J.: Lawrence Erlbaum Associates.

Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge,: M.I.T. Press.

Chomsky, N. (1988). *Language and Problems of Knowledge: The Managua lectures*. Cambridge, Mass.: MIT Press.

Chudnoff, E. (2014). The Rational Roles of Intuition. In A. Booth & D. Rowbottom (Eds.), *Intuitions*: Oxford University Press.

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

Dennett, D. (2009). Intentional Systems Theory. In A. Beckermann, B. P. McLaughlin, & S. Walter (Eds.), *The Oxford Handbook of Philosophy of Mind* (pp. 339-350). Oxford: Oxford University Press.

Dennett, D. C. (1978). *Brainstorms: Philosophical Essays on Mind and Psychology* (1st ed.). Montgomery, Vt.: Bradford Books.

Dogramaci, S. (2013). Intuitions for inferences. *Philosophical Studies, 165*(2), 371-399.

Egan, F. (2010). Computational models: a modest role for content. *Studies in History and Philosophy of Science Part A, 41*(3), 253-259.

Egan, F. (2012). Representationalism. In E. Margolis, R. Samuels, & S. Stich (Eds.), *The Oxford Handbook of Philosophy and Cognitive Science*: Oxford University Press.

Egan, F. (2013). How to think about mental content. *Philosophical Studies*(1), 1-21.

Egan, F. (2017). Function-Theoretic Explanation and the Search for Neural Mechanisms. In D. M. Kaplan (Ed.), *Explanation and Integration in Mind and Brain Science* (pp. 145-163): Oxford University Press.

Egan, F. (2018). *A Deflationary Account of Mental Representation (draft)*.

Evans, G. (1982). *The Varieties of Reference*. Oxford: Clarendon Press.

Fodor, J. (1984). Observation Reconsidered. *Philosophy of Science, 51*(1), 23-43.

Fodor, J. A. (1983). *Modularity of Mind: An Essay on Faculty Psychology*. Cambridge, Mass.: MIT Press.

Frege, G. (1979). Logic. Posthumous writings: Oxford, England: Blackwell.

Gregory, R. L. (1966). *Eye and brain: the psychology of seeing*. New York,: McGraw-Hill.

Gregory, R. L. (1980). Perceptions as Hypotheses. *Philosophical Transactions of the Royal Society B: Biological Sciences, 290*(1038), 181-197.

Gregory, R. L. (1997). Knowledge in Perception and Illusion. *Philosophical Transactions of the Royal Society B: Biological Sciences, 352*, 1121-1128.

Hatfield, G. (2002). Perception as Unconscious Inference. In D. Heyer & R. Mausfeld (Eds.), *Perception and the Physical World: Psychological and Philosophical Issues in Perception* (pp. 113--143): John Wiley and Sons.

Helmholtz, H. v. (1867). *Handbuch der physiologischen Optik*. Leipzig,: Voss.

Hlobil, U. (2014). Against Boghossian, Wright and Broome on inference. *Philosophical Studies, 167*(2), 419-429.

Hlobil, U. (2016). What Is Inference?

Hohwy, J. (2013). *The Predictive Mind*. Oxford: Oxford University Press.

Hornsby, J. (2000). Personal and sub-personal; A defence of Dennett's early distinction. *Philosophical Explorations, 3*(1), 6-24. doi:10.1080/13869790008520978

Jackendoff, R. (1994). *Patterns in the Mind: Language and Human Nature*. New York: Basic Books.

Lakoff, G., & Mark, J. (1999). *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. New York: Basic Books.

Larson, R. K., & Segal, G. (1995). *Knowledge of meaning: an introduction to semantic theory*. Cambridge, Mass.: MIT Press.

Ludwig, K. (1996). Explaining Why Things Look the Way They Do. In K. Akins (Ed.), *Perception*. Oxford: Oxford University Press.

McHugh, C., & Way, J. (2016). Against the Taking Condition. *Philosophical Issues, 26*(1), 314-331.

Nes, A. (2016). The Sense of Natural Meaning in Conscious Inference. In T. Breyer & C. Gutland (Eds.), *Phenomenology of Thinking* (pp. 97-115): Routledge.

Neta, R. (2013). What is an inference? *Philosophical Issues, 23*(1), 388-407.

Orlandi, N. (2011a). Ambiguous figures and representationalism. *Phenomenology and the Cognitive Sciences, 10*(3), 307-323.

Orlandi, N. (2011b). The innocent eye: Seeing-as without concepts. *American Philosophical Quarterly, 48*(1), 17.

Orlandi, N. (2012). Embedded seeing-as: Multi-stable visual perception without interpretation. *Philosophical Psychology, 25*(4), 1-19.

Orlandi, N. (2013). Embedded Seeing: Vision in the Natural World. *Noûs, 47*(4), 727-747.

Orlandi, N. (2014). *The Innocent Eye: Why Vision is Not a Cognitive Process*: Oup Usa.

Orlandi, N. (2016). Bayesian Perception Is Ecological Perception. *Philosophical Topics, 44*(2), 327-351.

Ramsey, W. M. (2007). *Representation Reconsidered*. New York: Cambridge University Press.

Rescorla, M. (2015). Bayesian Perceptual Psychology. In M. Matthen (Ed.), *The Oxford Handbook of Philosophy of Perception* (pp. 694-716). Oxford: Oxford University Press.

Rock, I. (1983). *The Logic of Perception*. Cambridge, Mass.: MIT Press.

Rock, I. (1984). *Perception*. New York: Scientific American Books, Inc.

Rosa, L. (2017). Reasoning without regress. *Synthese*.

Valaris, M. (2014). Reasoning and Regress. *Mind, 123*(489), 101-127.

Valaris, M. (2017). What reasoning might be. *Synthese, 194*(6).

Wandell, B. A. (1995). *Foundations of Vision*. Sunderland, Mass.: Sinauer Associates.

Wright, C. (2014). Comment on Paul Boghossian, "What is inference". *Philosophical Studies, 1*(1), 1-11.

ENDNOTES

[1] Nico Orlandi has developed a critique of inferentialist or constructivist accounts of perceptual accomplishment in a series of papers and a recent book (2011a, 2011b, 2012, 2013; 2014, 2016). Orlandi argues that inferentialism is not *the best explanation* of the success of the visual system. Orlandi advocates an embedded view (EV) of perception. "According to EV, the visual system has physical features that make it act in a lawful manner. We should refrain from thinking of such features as representing anything. The features are biases shaped by environmental contingencies in the evolutionary past *and in the present*, and we can appeal to such contingencies to explain what we see" (2014, p. 57). We focus on whether the conditions for attributing inferences to subpersonal systems can be met in the first place, but we argue that there is heuristic value in inferential talk because identifies causal-functional structure that helps explain successful representation. We suggest that this construal of inference talk converges with the approach that Orlandi recommends on empirical grounds.

[2] For example, unconscious inference theories of linguistic cognition look back to Chomsky's work (1965, 1988) on the structure of the language faculty. Though Chomsky has claimed that it is a misreading to attribute to him a UIT, his followers have embraced it:

> … the unconsciousness of mental grammar is still more radical than Freud's notion of the unconscious: mental grammar isn't available to consciousness under any conditions, therapeutic or otherwise. (Jackendoff, 1994, p. p. 9)

> The cognitive unconscious is the massive portion of the iceberg that lies below the surface, below the visible tip that is consciousness. It consists of all those mental operations that structure and make possible all conscious experience, including the understanding and use of language … it is completely and irrevocably inaccessible to direct conscious introspection. (Lakoff & Mark, 1999, p. p. 103)

UITs have been extended to unconscious processing of semantic rules for interpretation as well at the level of LF (see also (Larson & Segal, 1995)). On this view, language processing involves a faculty that possesses innate knowledge of grammatical principles and principles of interpretation which are applied both to input when a child is learning a first language and in language processing subsequently.

[3] Orlandi (2014, 2016) argues that the Bayesian approach is better characterized as an ecological approach than as an inferential theory.

[4] Hohwy represents the error minimization theory being a successor of inference theories that stretch back to Helmholtz which differs in its account of the inferences involved. Hohwy thinks it obvious that there is unconscious inferential processing: “We can in fact engage in such inference, since we can perceive” (2013, p. 14), as if perception could not occur when proximal stimuli underdetermine distal causes without *inferential* processes being involved. The interesting question, on his view, is the kind of inference.

[5] We don’t claim these desiderata are exhaustive, only that they are of central importance. See Hlobil (2014, 2016) for discussion.

[6] Boghossian notes a historical precedent for the Taking Condition in Frege, who claims, “[t]o make a judgment because we are cognizant of other truths as providing a justification for it is known as inferring~~.~~” (1979, p. 3). See Hlobil (2016) for more on historical antecedents.

[7] Neta (2013) argues that taking is a judgement, Valaris (2014, 2017) that it is a belief, and Chudnuff (2014) and Dogramaci (2013) that it is an intuition or intellectual seeming. In contrast, Hlobil (2016) and Boghossian (2014) (and arguably Broome, 2013) deny that it is an intentional state.

[8] Nes (2016) claims that in inferring some proposition, *p*, from some set of propositions, *Q*, one has “the sense” that *Q means that p* where ‘means’ is taken to be natural meaning in Grice’s sense of the term. Broome (2013) claims that the content of taking is that *the contents of one’s premise attitudes imply the contents of one’s conclusion attitude*. Valaris (2017) argues that the content of taking is that *the contents of one’s conclusion attitude follows from the content of one’s premise attitudes*, where taking consists in *realizing* that all (relevant) possibilities that make one’s premise attitudes true, make one’s conclusion attitude true. Finally, Neta (2013) claims that the content of a taking state is that *one’s premise attitudes propositionally justify one’s conclusion attitude.*

[9] See Boghossian (2008) on epistemic rules.

[10] Modular inferences are subpersonal. It is prima facie consistent with their being unconscious and modular that they are inferences made *by* the person in whom they take place. We do not address this view here because (i) it is implausible and (ii) it is unlikely that those advocating for an unconscious inference theory of perceptual achievement attribute, e.g., the inferences supposedly being drawn by the visual system to the agent herself as subject. For example, many

of the inferences would involve concepts that the PLS (which may be a child or non-linguistic animal) clearly does not possess.

[11] Many theorists attribute genuine beliefs to subsystems in the brain. For example, “It asserts that at some level of description all creatures of the same phenotype share the same prior beliefs about what their sensory input should be ….” (Hohwy, 2013, p. p. 86). Here the sensory inputs are not conscious level but stimulus at the sensory surfaces, of which the PLS is ignorant.

[12] Could one argue that the goals are not goals of the module but of the system that contains it? (Thanks to Anders Nes for this question.) First, since the behavior being guided is that of the subsystem, the goals are directed at what the subsystem does rather than the containing system, and so are at the wrong level to be personal level goals. Second, as we note next, the concepts deployed in the subsystem representations cannot be generally assumed to be available to the containing system. These concepts will be involved in goal specification for the subsystem as well.

[13] Are we overlooking the possibility of non-conceptual content? Non-conceptual content has been claimed for perceptual experiences, but we are not here entertaining views that attribute perceptual experiences to subpersonal modules. However, subpersonal computations have also been said to have non-conceptual content because they traffic in representations whose correctness conditions would be specified using concepts the PLS does not possess (Evans, 1982, pp. p. 104, n 22). However, we are here concerned with the view that SMI involve genuine inferences of just the sort that occur in the PLS except for being subpersonal. (See the quotations in section 2 and notes 11 and 18 in this connection.) We will consider fallback positions in

section 5, where we conclude that there is no case for a subsystem having states that are representations *for* the subsystem itself rather than a grid projected onto it by the theorist.

[14] We pass over some problems related to conclusions and premises of SMI. The conclusion is in a different subject than the premises, is an experience rather than belief, and contains more information than the premises. The first is the most serious problem because there is no one subject to take the premises to support the conclusion. For the premises, how are the general principles, some of which are not innate, learned by subsystems, if they do not have access to information possessed by the PLS, and how do subsystems learn of what is going on at the sensory surfaces? Magic?

[15] On this topic, it is useful to note a feature of Bayesian models of perception. The Bayesian inference from perceptual input to, e.g., shape, yields a probability distribution, but perception is determinate. This is usually handled by invoking a utility function, which may be task dependent, that reflects the penalty for making a mistake (rather than just choosing the hypothesis with the highest posterior probability). The determinate output is the one that maximizes expected utility. However, first, this undercuts the idea that an *inference* is being made to what the environment is like. If you accept Pascal's Wager, you are not inferring that God exists, but reaching the practical judgment that belief in God maximizes expected utility. Second, whose utility? Not the perceptual system, but the PLS, since it is potential harm or benefit to the whole system that is taken into account. But then we have an action with no proper agent.

[16] The prediction error minimization project treats the perceptual system as a hierarchy of levels at each of which inferences are performed. At the lowest level it treats inputs to the inferences as

involving information about physical stimulation of the sensory surfaces. “The brain does have access to the sensory data that impinges on it” (Hohwy 2013, p. 50). The brain is also said to engage not just in first-order Bayesian reasoning but in “second order statistics that optimizes its precision expectations,” which is a matter of “perceptual inference about perceptual inference” (p. 66). This involves more conceptual sophistication than most people possess. It is a good thing the brain is smarter than the person it serves. Notably, the more sophisticated the theorist becomes, the more sophisticated the brain is said to be. The history of inferential accounts, which have become more and more sophisticated over time, suggests that the inferences lie in the eye of the beholder.

[17] Dennett’s proposal was bound up with his advocacy of the Intentional Stance as foundational in understanding propositional attitude attributions. See note 23 in this connection. See also Hornsby (2000, p. sec. 4) for how Dennett’s development of intentional systems theory led him away from an early strict division between personal level attributions of psychological states and subpersonal mechanisms.

[18] See the quotation from Rock in section 2, and Fodor (1984): “… what mediates perception is an inference from effect to causes. The sort of mentation required for perception is thus not different in *kind*—though no doubt it differs a lot in conscious accessibility—from what goes on in Sherlock Holmes’s head when he infers the identity of the criminal from a stray cigar band and a hair or two” (p. 31). In this connection see also the discussion in (Bennett & Hacker, 2003, pp. 23-33).

[19] This is very much the idea in the prediction error minimization account that attributes Bayesian reasoning to the perceptual system. The brain has a model of the environment which is

used to make a prediction and revised to minimize error between the prediction and environment (Hohwy 2013, ch. 2).

[20] Ramsey defends S-representations against the related charge that isomorphism is promiscuous by arguing that "components of the model become representations when the isomorphism is exploited in the execution of surrogative problem-solving" (2007, p. 96). One might think this solves the problem just outlined. But the question is still how to make sense of its being *exploited* in any sense other than playing a causal role in producing appropriate personal level representations, or of *problem solving* going on in any sense other than that an appropriate personal level representation results. Repeating a question begging description is not an argument.

[21] See Ludwig (1996, p. sec. 7). Frances Egan's (2010, 2012, 2013, 2017) two-part pragmatic deflationary account of representations in cognitive neuroscience separates mathematical from cognitive content in computational accounts of cognitive function. Our discussion focuses on what Egan calls the intentional gloss. The mathematical function gets into the picture only as more detailed mechanisms for realization of the "inferential processes" are proposed.

[22] Egan introduces this term (2013, 2017) to characterize a neural mechanism as computing a mathematical function, but it applies equally well to inferential theories at a higher level of functional organization.

[23] We treat the intentional stance as a matter of treating a system as-if it had intentional states. Dennett's intentional systems theory holds that the distinction between as-if intentionality and original intentionality is ill-motivated (2009). We reject this. The concept of the intentional stance presupposes intentionality since it presupposes an intentional agent who takes it up. If

intentional systems theory maintains that a system has genuine intentional states iff someone can usefully take the intentional stance toward it, it makes the explanans presupposes an understanding of the explanandum. Thus, the truth of the biconditional itself has to be settled on the basis of an independent analysis of intentionality. For further critical discussion, see (Bennett & Hacker, 2003, p. appendix 1).

[24] Assigning representational contents to states is analogous to assigning numbers to physical magnitudes like mass, energy, and momentum. We use the numbers and their structure to keep track of relations among the states that we assign them to. Similarly, to treat a state, say, as representing 1, or an edge, is to keep track of its role in the system, relative to a systematic assignment of contents to states and semantic relations to transitions.