

## Morphological Content and Chromatic Illumination in Belief Fixation

### 1. *Rationality and the Frame Problem*

The frame problem in cognitive science, which first arose in the early days of artificial-intelligence research, reveals aspects of human rationality that often have been overlooked both in philosophy and psychology. The frame problem pertains both to “theoretical” rationality concerning matters like belief fixation and to “practical” rationality concerning matters like action and planning—while also revealing how deeply intertwined are these two dimensions. These underappreciated aspects of rationality were nicely brought into view by Dennett (198?), an important and illuminating essay on the frame problem. As prelude and partial motivation for our subsequent discussion, we will begin by quoting Dennett at some length, commenting on quoted passages along the way. He starts the essay with the following parable:

Once upon a time there was a robot, named R1 by its creators. Its only task was to fend for itself. One day its designers arranged for it to learn that its spare battery, its precious energy supply, was locked in a room with time bomb set to go off soon. R1 located the room, and the key to the door, and formulated a plan to rescue its battery. There was a wagon in the room, and the battery was on the wagon, and R1 hypothesized that a certain action which it called PULLOUT(WAGON, ROOM) would result in the battery being removed from the room. Straightaway it acted, and did succeed in getting the battery out of the room before the bomb went off. Unfortunately, however, the bomb was also on the wagon. R1 *knew* that the bomb was on the wagon in the room, but didn't realize that pulling the wagon would bring the bomb out along with the battery. Poor R1 had missed that obvious implication of its planned act.

Back to the drawing board. “The solution is obvious,” said the designers. “Our next robot must be made to recognize not just the intended implications of its acts, but also implications about their side effects, by deducing these implications it uses in formulating its plans.” They called their next model, the robot-deducer, R1D1. They placed R1D1 in much the same predicament that R1 had succumbed to, and as it too hit upon the idea of PULLOUT(WAGON, ROOM) it began, as designed, to consider the implications of such a course of action. It had just finished deducing that pulling the wagon out of the room would not change the color of the room's walls, and was embarking on a proof of the further implication that pulling the wagon out would cause its wheels to turn more revolutions than there were wheels on the wagon—when the bomb exploded.

Back to the drawing board. “We must teach it the difference between relevant implications and irrelevant implications,” said the designers, “and teach it to ignore the irrelevant ones.” So they developed a method of tagging implications as either relevant or irrelevant to the project at hand, and installed the method in their next model, the robot-relevant-deducer, or R2D1 for short. When they subjected R2D1 to the test that had so unequivocally selected its ancestors for extinction, they were surprised to see it sitting, Hamlet-like,

outside the room containing the ticking bomb, the native hue of its resolution sicklied o'er with the pale cast of thought, as Shakespeare (and more recently Fodor) has aptly put it. "Do something!" they yelled at it. "I am," it retorted. "I'm busily ignoring some thousands of implications I have determined to be irrelevant. Just as soon as I find an irrelevant implication, I put it on the list of those I must ignore, and..." the bomb went off.

This parable illustrates nicely the extent to which ordinary human common-sensical planning and action depend on pertinent background information available to the human agent—information that somehow gets appreciated and accommodated, whether or not it becomes explicitly present in consciousness along the way. It also illustrates the extent to which planning and action typically are suffused with pertinent *expectations* about what will happen next if one acts a certain way, and about what *would* happen next if one were to act some other way (e.g., pulling the wagon out of the room without first removing the bomb)—again, whether or not those expectations become explicitly present in consciousness along the way. Such expectations are aptly considered aspects of theoretical rationality, and their presence is intimately bound up with practical rationality—whether or not they qualify as full-fledged *beliefs*. Concerning the robots in his parable, Dennett goes on to say this:

All these robots suffer from the *frame problem*. If there is ever to be a robot with the fabled perspicacity and real-time adroitness of R2D2, robot-designers must solve the frame problem.... [I]t is a new, deep epistemological problem—accessible in principle but unnoticed by generations of philosophers—brought to light by the novel methods of AI, and still far from being solved....

One utterly central—if not defining—feature of an intelligent being is that it can "look before it leaps." Better, it can *think* before it leaps. Intelligence is (at least partly) a matter of using well what you know—but for what? For improving the fidelity of your expectations about what is going to happen next, for planning, for considering courses of action, for framing further hypotheses with the aim of increasing the knowledge you will use in the future, so that you can preserve yourself, by letting your hypotheses die in your stead (as Sir Karl Popper once put it)....

But when we think before we leap, *how do we do it?*

So the frame problem, in short, is the problem of understanding how the human cognitive system manages to bring to bear, in real time, pertinent information it possesses in such a way as to effectively manage these kinds of theoretical-cum-practical tasks. The depth and difficulty of this problem is easy to overlook, precisely because the common-sense rationality that has proved so dauntingly difficult to engineer into robots is so easy and so natural for us humans. But evidently this ease and naturalness belies the complexity of what goes on outside of explicit conscious awareness. Dennett remarks:

The myth that each of us can observe our mental activities has prolonged the illusion that major progress could be made on the theory of thinking by simply reflecting carefully on our own cases. For some time now we have known better: we have conscious access to only the upper surface, as it were, of the multi-level system of information-processing that occurs in us. Nevertheless, the myth still claims its victims.

The analogy of the stage magician is particularly apt. One is not likely to make much progress is figuring out *how* the tricks are done by simply sitting attentively and watching like a hawk. Too much is going on out of sight. Better to face the fact that one must either rummage around backstage in the wings, hoping to disrupt the performance in telling ways; or, from one's armchair, think aprioristically about how the tricks *must* be done, given whatever is manifest about the constraints. The frame problem is then rather like the unsettling but familiar "discovery" that so far as armchair thought can determine, a certain trick we have just observed is flat impossible.

He offers the following mundane example, commenting on it in a way which, like his parable of the three robots, thematizes the richness of background knowledge that needs to be in play in ordinary planning and plan-execution—knowledge of the kind that ordinary common sense typically accommodates without its becoming explicitly present in consciousness. Immediately after the just-quoted passage, and continuing with the analogy of the stage magician, Dennett says:

Here is an example of the trick. Making a midnight snack. How is it that I can get myself a midnight snack? What could be simpler? I suspect that there is some leftover sliced turkey and mayonnaise in the fridge, and bread in the breadbox—a bottle of beer in the fridge as well. I realize that I can put these elements together, so I concoct a childishly simple plan: I'll just go and check out the fridge, get out the requisite materials, and make myself a sandwich, to be washed down with a beer. I'll need a knife, a plate, and a glass for the beer. I forthwith put the plan into action and it works! Big deal.

Now of course I couldn't do this without knowing a good deal—about bread, spreading mayonnaise, opening the fridge, the friction and inertia that will keep the turkey between the bread slices and the bread on the plate as I carry the plate over to the table beside my easy chair. I also need to know how to get the beer out of the bottle and into the glass.... [O]ne trivial thing I have to know is that when the beer gets into the glass it is no longer in the bottle, and that if I'm holding the mayonnaise jar in my left hand I cannot also be spreading the mayonnaise with the knife in my left hand. Perhaps these are straightforward implications—instantiations—of some more fundamental things...such as, perhaps, the fact that if something is one location it isn't also in another, different location; of the fact that two things can't be in the same place at the same time; or the fact that situations change as the result of actions....

Such utterly banal facts escape our notice as we act and plan, and it not surprising that philosophers, thinking...*introspectively*, should have overlooked them. But if one...just thinks...about the purely informational demands of the task—what *must* be known by any entity that can perform this task—these banal bits of knowledge rise to our attention. We can easily satisfy ourselves that no agent that did not *in some sense* have the benefit of the information (that beer in the bottle is not in the glass, etc.) could perform such a simple task. It is one of the chief methodological beauties of AI that it makes one be a phenomenologist in this improved way.... [O]ne reasons about what

the agents must “know” or figure out *unconsciously or consciously* in order to perform in various ways.

He also stresses, rightly, the important role that doxastic states of *expectation* play in the forming and executing of plans—again, largely without becoming explicitly present in consciousness (not, anyway, as long they do not get thwarted). They are important, *inter alia*, because of how they contribute to the intelligent flexibility with which human agents execute their plans and intentions, adapting aptly to unanticipated contingencies. Continuing with the example of the midnight snack, he writes:

We assure ourselves of the intelligence of an agent by considering counterfactuals: if I had been told that the turkey was poisoned, or the beer explosive, or the plate dirty, or the knife too fragile to spread mayonnaise, would I have acted as I did? If I were a stupid “automaton”... I might infelicitously “go through the motions” of making a midnight snack oblivious to the recalcitrant features of the environment. But in fact, my midnight-snack-making behavior is multifariously sensitive to current and background information about the situation.... [A]n intelligent agent must engage in swift information-sensitive “planning” which has the effect of producing reliable but not foolproof expectations of the effects of its actions. That these expectations are normally in force in intelligent creatures is testified to by the startled reactions they exhibit when their expectations are thwarted. This suggests a graphic way of characterizing the minimal goal that can spawn the frame problem: we want a midnight-snack-making robot to be “surprised” by the trick plate, the unspreadable concrete mayonnaise, the fact that we’ve glued the glass to the shelf. To be surprised to you have to have expected something else, and in order to have expected the right something else, you have to have *and use* a lot of information about things in the world.

In sum, Dennett’s above-quoted observations about the frame problem, largely focused on his parable of the three robots and his example of the midnight snack, make vivid several crucial points that will be presupposed in our discussion below. First, rational agents deploy—and must deploy—extensive amounts of pertinent background information in carrying out even quite mundane everyday activities. Second, such background information typically gets accommodated without becoming explicitly present in consciousness—a fact that has often led philosophers (and psychologists too) to seriously under-appreciate the complexity and subtlety of common-sense rationality. Third, “practical” rationality is much more suffused with elements of “theoretical” rationality than is often realized—viz., *expectations*, which (like background information) typically are operative in ongoing practical agency without becoming explicitly present in consciousness. (Expectations are doxastic, surely, whether or not they always qualify as *beliefs*; and, being doxastic, they fall under the rubric of “theoretical” rationality.)

## **2. *Morals of the Frame Problem, I: Belief Fixation as Abductively Inferential and as Non-Modular***

Belief fixation is the generation of beliefs and the maintenance of existing beliefs. (What will be said in this section about belief fixation applies, *mutatis mutandis*,

to expectation fixation too, regardless whether or not one counts all expectations as full-fledged beliefs.) Rational belief-fixation is the generation and maintenance of beliefs under circumstances in which those beliefs (i) are evidentially well supported by one's available evidence, and (ii) arise *because* they are thus evidentially supported. Thus, rational belief-fixation is, in an important sense, *inferential*: it is a process whereby beliefs are generated on the basis of evidence, possessed and deployed by the cognitive system, that supports them. And in general, rational belief-fixation also is *abductively* inferential, since normally the pertinent kind of evidential support is abductive rather than demonstrative.

This being so, the frame problem arises in a stark way for belief fixation—a point that long was rightly and emphatically stressed by another influential philosopher who, like Dennett, kept a close eye on pertinent developments in cognitive science—the late Jerry Fodor. Here are some pithy remarks he made at the beginning of an essay of his own on the frame problem:

There are, it seems to me, two interesting ideas about modularity. The first is the idea that some of our cognitive faculties are modular. The second is the idea that some of our cognitive faculties are not.

By a modular cognitive faculty, I mean—for present purposes—one that is 'informationally encapsulated'. By an informationally encapsulated cognitive faculty, I mean one that has access, in the course of its computations, to less than all of the information at the disposal of the organism whose cognitive faculty it is. So, for example, I think that the persistence of the Muller-Lyer illusion in spite of one's knowledge that it *is* an illusion strongly suggests that some of the cognitive mechanisms that mediate visual size perception must be informationally encapsulated....

It's worth emphasizing a sense in which modular cognitive processing is *ipso facto* irrational. After all, by definition modular processing means arriving at conclusions by attending to less than all of the evidence that is relevant and available. And ignoring relevant and available evidence is, notoriously, a technique of belief fixation that will get you into trouble in the long run. Informational encapsulation is economical; it buys speed and the reduction of computational load by, in effect, drastically delimiting the database that's brought to bear in problem solving. But the price of economy is warrant. The more encapsulated the cognitive mechanisms that mediate the fixation of your beliefs, the worse is your evidence for the beliefs that you have. And...the worse your evidence for your beliefs is, the less the likelihood that your beliefs are true....

[R]ational processes have their debilities too.... If, for example, you undertake to consider a nonarbitrary sample of the available and relevant evidence before you opt for a belief, *you have the problem of when the evidence you have looked at is enough*. You have, that is to say, Hamlet's problem: when to stop thinking.

The frame problem is just Hamlet's problem viewed from an engineer's perspective.... What is a nonarbitrary strategy for delimiting the evidence that should be searched in rational belief fixation? I don't know how to answer this

question. If I did, I'd have solved the frame problem and I'd be rich and famous.

In other writings Fodor elaborated at length on these thoughts, in three respects. First, he argued that typically the degree of evidential support that a potential belief possesses, relative to one's total available evidence, (i) is non-demonstrative and abductive (rather than being a matter of logical entailment by the evidence), and (ii) is highly and multifariously *holistic* (as are evidential abductive-support relations in science). Like confirmation of hypotheses in science, non-demonstrative evidential support for potential beliefs typically exhibits two holistic aspects that Fodor called the "isotropic" feature and the "Quineian" feature. In Fodor (1983) he wrote:

By saying that confirmation is isotropic, I mean that the facts relevant to the confirmation of a scientific hypothesis may be drawn from anywhere in the field of previously established empirical (or, of course, demonstrative) truths. Crudely: everything that the scientist knows is, in principle, relevant to determining what else he ought to believe.... (1983, p. 105)

By saying that scientific confirmation is Quineian, I mean that the degree of confirmation assigned to any given hypothesis is sensitive to properties of the entire belief system; as it were, the shape of our whole science bears on the epistemic status of each scientific hypothesis (1983, p. 107).

Isotropy brings in the whole current belief system: any bit of actual or potential information from any portion of the belief system might, in some circumstances, be evidentially relevant to any other. Being Quineian makes confirmation holistic in a deeper way: confirmation depends upon "such considerations as simplicity, plausibility, and conservatism" (Fodor, 1983, p. 108), which are determined by the global *structure* of the whole of the current belief system and of potential successor systems.

Second, he urged that human belief fixation is often quite rational, thereby operating in a non-modular manner that successfully brings to bear the holistic, abductive, evidential support for a given potential belief that is actually possessed by the cognitive system. (Likewise for expectation fixation, as it figures in planning and in guidance of action.)

Third, he bemoaned the fact that a tractable computational solution to "Hamlet's problem," for creatures who have as much available evidence to draw upon as do real humans, looks unattainable. Hence his ongoing, and increasingly pessimistic, negative attitude about the prospects for what he called the "computational theory of mind" (CTM). In Fodor (1983) he wrote, concerning processes like belief fixation and planning:

The difficulties we encounter when we try to construct theories of [these kinds of] processes are just the sort we would expect to encounter if such processes are, in essential respects, Quineian/isotropic.... The crux in the construction of such theories is that there seems to be no way to delimit the sorts of informational resources which may affect, or be affected by, central processes

of problem-solving. We can't, that is to say, plausibly view the fixation of belief as effected by computations over bounded, local information structures. A graphic example of this sort of difficulty arises in AI, where it has come to be known as the "frame problem" (i.e., the problem of putting a "frame" around the set of beliefs that may need to be revised in light of specified newly available information) (Fodor, 1983, pp. 112-3).

In a more recent essay, he expressed his ongoing deep pessimism as follows (with our interpolations about what, in context, he clearly had in mind):

Computational nativism [i.e., the CTM] is clearly the best theory of the cognitive mind that anyone has thought of so far (vastly better than, for example, the associationistic empiricism that is the main alternative); and there may indeed to aspects of cognition [viz., informationally encapsulated ones, such as the cognitive mechanisms that mediate visual size perception] about which computational nativism has got the story more or less right. But it's nonetheless quite plausible that computational nativism is, in large part [viz., vis-à-vis informationally unencapsulated cognitive faculties like planning and belief fixation], not true.

The problem is not nativism (which asserts that human cognitive mechanisms deploy a significant amount of information that is possessed innately), but rather the idea that human cognition is *computation*, i.e., manipulation of explicit mental representations in accordance with rules of the kind that could constitute a computer program. Tractable computational processes require modularity—an informationally encapsulated data base—whereas rational belief fixation is highly holistic and non-modular. (The case against the CTM is further articulated in Horgan and Tienson 1994, 1996.)

### ***3. Morals of the Frame Problem, II: Belief Fixation as Non-Computational and as Essentially Morphological***

The human cognitive system somehow accomplishes highly holistic, highly non-modular, belief fixation and expectation fixation. How might it do so? Dennett's analogy of the stage magician is very pertinent here. If one supposes that the trick is accomplished by means of tractable computation that somehow deploys what Fodor called "a nonarbitrary strategy for delimiting the evidence that should be searched in rational belief fixation," then one will find it difficult to avoid concluding (as Dennett put it) "that so far as armchair thought can determine, a certain trick we have just observed [viz., the holistically rational fixation of beliefs and expectations] is flat impossible."

The way forward is to follow Dennett's two-pronged advice: "either rummage around backstage in the wings, hoping to disrupt the performance in telling ways; or, from one's armchair, think aprioristically about how the tricks *must* be done, given whatever is manifest about the constraints." This was done by Horgan and Tienson (1994, 1996). Aprioristically, they argued that highly holistic, highly non-modular, cognitive processing must operate in such a way that much pertinent background information gets accommodated *automatically and without being rendered explicit*, by virtue of the standing structure (the morphology) of the cognitive system. They

called information that gets implicitly accommodated this way “morphological content.”

Rummaging backstage in the wings, at a time when neural-network models had come to prominence in cognitive science, Horgan and Tienson argued that abductively holistic cognitive processing should be understood not as the computational manipulation of explicit mental representations of all the evidentially pertinent information, but rather in terms of the kind of mathematics that goes naturally with neural-network modeling—viz., dynamical systems theory. Morphological content is embodied in the topological contours of the  $(n+1)$ -dimensional “activation landscape” of an  $n$ -node neural network—with the activation-level of each node in the network being a dimension in the network’s  $n$ -dimensional “state space,” and with “downhill” on the landscape being the direction of time.<sup>1</sup> Temporal trajectories along the activation landscape from one activation-state to another, as determined largely by the morphological content embodied in the topological contours on the landscape in the vicinity of the initial activation-state, constitute transitions from one occurrent cognitive state to another.

Horgan and Tienson argued that two apparent morals emerge from the fact that cognitive processes like belief fixation are richly holistic and highly non-modular. First, such processes are too subtle and too complex to constitute representation-level computation, i.e., manipulation and transformation of mental representations in accordance with programmable rules that advert to the content-encoding structure of those representations.<sup>2</sup> Second, these processes depend essentially and heavily upon morphological content, i.e., information that is accommodated automatically by such processes without being rendered explicit along the way. Although the first of these two apparent morals will not figure prominently in our subsequent discussion in this paper, then second one will. We will be focusing on two interrelated matters: (1) the question of how the abductively inferential, essentially morphological, nature of rational belief-fixation is related to the conscious, experiential, aspects of episodes of belief fixation, and (2) the question of the potential implications for epistemology of the contention that rational belief-fixation typically is both abductively inferential and essentially morphological.

#### **4. *Morphological Content and Consciousness, I: The EOC Assumption***

Rational belief-fixation is abductively inferential: rationally generated beliefs are ones that are well-supported—typically non-demonstratively—by evidence one possesses. But since processes of rational belief-fixation depend heavily upon essentially morphological aspects, these processes do not constitute *fully explicit* inferences; i.e., they do not constitute inferences in which all pertinent evidence, and all of the holistic, Quineian, respects in which the evidence combines to abductively support the given belief, get explicitly represented during the process belief fixation.

Prima facie, it seems quite plausible that content that gets implicitly and automatically accommodated during cognitive processing, rather than getting explicitly represented along the way via an occurrent mental state with that very content, cannot be an aspect of conscious experience. Rather, it might seem, any thought-content that is present in consciousness must be the explicit content of occurrent mental state, rather than morphological content. This is not to say that all



occurrent mental states must be conscious; for, some might be entirely unconscious. Rather, the claim is that being explicitly represented by an occurrent mental state is a *necessary condition* (albeit not a sufficient condition) for an item of content to be present in consciousness. We will call this claim the *explicit occurrent content assumption* (for short, the EOC assumption).

We next set forth and compare several epistemological positions that presuppose the EOC assumption. Thereafter we will argue that the EOC assumption, despite being initially plausible, is false; and we will consider the apparent import of its falsity for epistemology.

## 5. *Doxastic Justification, I: Three Alternative Positions Each Embracing the EOC Assumption*

If one accepts (1) that rational belief-fixation typically is abductively inferential in a highly holistic, isotropic-cum-Quineian way, (2) that such belief fixation depends heavily upon morphological content, and (3) that the EOC assumption is correct, then what might all this mean concerning the nature of doxastic justification? We will canvass several alternative potential answers to this question.

First is what we will call *non-experiential reliabilism*. This view holds that the key constitutive feature of a justified belief is the fact that it was produced by a reliable belief-forming process. Although the source of this reliability is the sensitivity of the belief-fixation process to holistic evidential support that is possessed by the cognitive system, such evidence-sensitivity is not itself constitutive of doxastic justification; rather, reliability of the operative belief-fixation process is what's constitutive. Moreover, given the EOC assumption and the fact that belief fixation typically is abductively rational in a way that depends heavily upon morphological content, conscious experience is comparatively unimportant in epistemic justification: far too much of the abductive belief-fixing conscious processing takes place outside of consciousness. Hence, this form of reliabilism is aptly called "non-experiential."<sup>3</sup>

Second is what we will call *non-experiential evidentialism*. This view holds instead that the key constitutive feature of a justified belief is the fact that it was produced, by a holistically evidence-sensitive belief-fixation process, in response to evidence possessed by the cognitive agent that constitutes strong evidential support for the belief. Although a process of this kind is apt to be reliable, its reliability is a byproduct the constitutive feature of doxastic justification, rather than being constitutive itself; what's constitutive is evidence-sensitive processing in response to good evidence one actually possesses. But this view nonetheless is similar to non-experiential reliabilism in this respect: conscious experience is comparatively unimportant in epistemic justification, because far too much of the abductive belief-fixing processing takes place outside of consciousness. Hence, this form of evidentialism is aptly called "non-experiential."

Third is what we will call *gut-feeling experiential evidentialism*. This view treats certain conscious epistemic "gut feelings"—certain conscious experiences as-of a given proposition's *seeming to be true*—as the key constitutive feature of doxastic justification. (In order to be at all plausible, this view presumably would need to

impose certain constraints on the kinds of epistemic gut feelings that would qualify as justificatory, as distinct from other kinds that would not qualify.) The view also embraces the EOC assumption, and so also embraces the claim that much of the abductively inferential belief-fixing processing that generates the pertinent epistemic feelings takes place outside of consciousness. This view therefore denies that such unconscious processing is itself constitutive of epistemic justification; instead, such processing is treated as being merely the causal source of the epistemic gut feelings which themselves are what is really constitutive.

Each of these three views denies, in one way or another, that the conscious experience and the essentially morphological aspects of belief fixation are *both* key constitutive aspects of epistemic justification. The first two positions relegate conscious experience to a peripheral role, constitutively—claiming that very large portions of the crucial constitutive action occurs entirely outside of consciousness. The third position, by contrast, treats certain kinds of conscious gut feelings as crucially constitutive of doxastic justification, while claiming that the evidence-sensitive inferential processing underlying the pertinent gut feelings, to the (very great) extent that this processing operates morphologically, plays no constitutive justificatory role.

None of these views, we contend, is either theoretically satisfying or phenomenologically plausible. Conscious experience plays a crucial and central role, constitutively, in doxastic justification; but so do the essentially morphological aspects of abductively inferential belief-fixation processes. The key to having it both ways is to jettison the EOC assumption.

## **6. *Morphological Content and Consciousness, II: Chromatic Illumination***

In philosophy of mind it has been very common to divide intentional mental states into two categories, which often are regarded (if only tacitly) as both exclusive and exhaustive: *occurrent* and *dispositional*. Typically, the notion of a dispositional mental state with content *C* is understood as being a disposition to go into an occurrent mental state with content *C* under suitable elicitation conditions—where the elicitation circumstances don't constitute a reason not previously possessed to have the pertinent propositional attitude vis-à-vis *C*, but rather just prompt an occurrent instance of the state that was “already there” dispositionally. For instance, if you are asked whether elephants can fly, this question is apt to prompt in you an occurrent belief that elephants cannot fly—a belief that you already possessed dispositionally, whether or not it has ever previously become occurrent in your mental life. (It is a nice philosophical question how to distinguish a dispositional belief that *C* and a mere disposition *to believe* that *C*; for a discussion of this question, see Audi XXYY.) Dispositional mental states, as thus typically construed, are not and cannot be directly efficacious *qua* dispositional; rather, the pertinent disposition first must be triggered by an elicitation circumstance that initiates a causal process which generates an occurrent instance of the pertinent mental state—and this occurrent intentional mental state now can be efficacious in one's subsequent cognition and subsequent behavior.

We will call dispositional mental states, as thus construed, *canonically* dispositional. And we will say, of a canonically dispositional mental state with content *C* that has been triggered by an elicitation circumstance into an occurrent

mental state with content *C*, *mediately efficacious*; the idea is a canonically dispositional intentional mental state with content *C* can be efficacious only indirectly, via the mediation of a prompt-elicited occurrent intentional mental state with content *C*.

Consider now what we earlier called the EOC assumption, which asserts that a necessary condition for a content *C* to be present in consciousness is to be experiencing an occurrent, conscious, intentional mental state whose explicit content is *C*. If one supposes (if only tacitly) that the categories of occurrent and categorically dispositional are exclusive and exhaustive regarding intentional mental states, then the EOC assumption is apt to seem obviously and trivially true.

However, a third way to be in an intentional mental state, distinct from being in it either occurrently or canonically dispositionally, is to be in it morphologically. Although morphological content in a cognitive system is a dispositional matter, the pertinent dispositions differ crucially from those that accrue to canonically dispositional mental states: morphological-content dispositions can be *immediately efficacious*. I.e., an occurrent cognitive process can accommodate morphological content *C*—can evolve, non-accidentally by virtue of the cognitive system's morphology, in a way that is sensitive to *C*—without *C* ever getting explicitly represented during the process.<sup>4</sup>

Given that a cognitive agent's possession of morphological content *C* can be immediately efficacious in this way with respect to occurrent, conscious, cognitive processes, it is entirely possible that such immediate efficacy can figure importantly in the specific phenomenological character of those occurrent processes. There could well be “something that it's like”—some specific phenomenological aspect(s)—distinctive of conscious cognitive processing in which background morphological content *C* is being immediately efficacious. This would be the “what it's like” of *implicitly appreciating C consciously without C being explicitly represented in consciousness*.

In Horgan and Potrč (2010), we argued that occurrent conscious experience quite often has such phenomenological aspects, which we labeled *chromatic illumination*. This expression is inspired by a visual metaphor, which we find helpful in conveying the nature of such phenomenology. Consider a visual scene that is illuminated in certain ways by light sources that are not themselves visible (from the observer's perspective) within the scene, and that significantly affect the overall look of the scene. Think, for instance, of the famous 1892 painting by Toulouse Lautrec, “At the Moulin Rouge,” which hangs in the Art Institute of Chicago. Various figures in the painted scene are illuminated in strikingly different ways—e.g., the women more prominently than the men, one woman by lighting to the left but outside the scene, another woman by lighting from the lower right but outside the scene, a peculiar light-induced greenish tint to some of the illuminated faces that blends with the greens in the background of the scene, etc. The presumptive sources of these distinctive features—lighting of various kinds at various positions in the presumptive wider environment, producing light with various different chromatic characteristics—are not present in the visible scene. They are not explicitly represented. Nonetheless, they are *implicitly* present in the scene anyway, in the ways that the figures in the scene are chromatically illuminated by those presumptive light-sources.

The directly visible scene presented in a painting—or in a photograph, or on a stage—can be taken as a metaphorical stand-in for one’s present occurrent conscious experience, which includes the various specific ways that items explicitly represented in the scene are illuminated. By contrast, the out-of-view sources of the visible scene’s various aspects of illumination can be taken as a metaphorical stand-in for what is implicitly present in consciousness—present not by being explicitly presented or depicted, but rather by virtue of how it affects the character of what *is* there explicitly. These effects, within occurrent consciousness, of information not explicitly present in consciousness, are what we call chromatic illumination.

In Horgan and Potrč (2010) we argued, largely by appeal to introspective phenomenological reflection, that conscious experience frequently exhibits chromatic illumination by background information—information that thus is implicitly appreciated experientially without being explicitly present in consciousness. We argued in part, in a way we now briefly repeat here, by invoking a familiar type of experience: understanding a joke. (An advantage of focusing on a non-sensory cognitive phenomenon is that it is appropriately analogous, so we will maintain, to belief formation.)

For specificity, consider a cartoon by Danny Shanahan that appeared in the *New Yorker* magazine roughly twenty years ago.<sup>5</sup> (We will describe it rather than exhibiting it—which underscores the fact that pertinent aspects of joke-getting normally do not depend essentially on specific sense modalities.) Two guys are sitting in a bar, and the one who is speaking looks unhappy and exasperated. He says to the other guy, “What I don’t, like, get is how she, like, figured out I was, like, having an affair with the, like, babysitter.”

Getting this joke, or any joke, is an instantaneous experience. In that instant, normally one needs to appreciate quite a wide range of pertinent background information; also, one needs to appreciate why and how all this information combines, holistically, to constitute an instance of funniness. Each item of that information must be appreciated, together with the ways that all the pertinent items interact with one another to make for funniness—for, otherwise one would not be understanding the joke.

For any reasonably clever joke, typically it is possible to elaborate at some length upon the various items of background information all of which are pertinent. In the example at hand, for instance, probably the most salient such item—the item that one would mention first, in seeking to explain the joke to someone who does not yet get it—is that the practice of persistently inserting the word ‘like’ into one’s spoken English is a distinctive feature of way *teenage girls* talked in the U.S., twenty years ago. (Supposedly the practice was initiated by teenage females in the San Fernando Valley, adjacent to Los Angeles—so-called “valley girls,” who are sources of numerous youthful fads in the U.S. Nowadays, regrettably, such ‘like’-talk is much more prevalent, so much more so that contemporary American undergraduates frequently don’t easily get the joke.)

But that is just the beginning, in terms of pertinent background information that must be grasped in the instant in order to get the joke. Additional such items of information can be made salient by posing suitable questions—questions that might be asked by someone who is persistently deficient at understanding jokes. Whose

children did this babysitter care for? Roughly how old is the babysitter? Who is the “she” who figured out that the guy is having an affair with the babysitter? How did that person react, upon learning this fact? Why is the guy upset that she figured it out and reacted that way? Why doesn’t he himself understand how she figured it out? And, of course: What’s so funny about all this? One could write a monograph about the pertinent background information and its holistic relevance. Likewise for most any even moderately subtle joke.

In the instant of joke-getting, very little of this kind of information seems to be explicitly present in consciousness. Nonetheless, all of it is being *appreciated* in one’s conscious experience, because otherwise one would not get the joke. Thus, the information is *implicitly* present in consciousness, by way of the chromatic illumination it exerts upon the phenomenological character of one’s overall synchronic experience. Furthermore, the “getting it” aspect of experience is not some generic feature, such as experiencing oneself laughing or inclined to laugh (perhaps without knowing why). Rather, it is quite content-specific: some particular item(s) of explicitly conscious content (in this case, what the guy at the bar is saying, and his obvious consternation) is appreciated as funny by virtue of how those explicitly conscious items are relevantly interconnected with a rich body of pertinent background information. Thus, all those *specific* items of background information are implicitly present in the conscious joke-getting experience, by virtue of the *specific* way that the experience chromatically illuminates the phenomenological character of the experience.

Admittedly, it is an empirical hypothesis that joke-getting has the features lately described. But it is a very plausible hypothesis, for at least two mutually supporting reasons. First, it conforms with the deliverances of introspection: when one attends introspectively to one’s joke-getting experience (albeit perhaps retrospectively via one’s memory of the recent past), it does not seem that all that background information is explicitly present in consciousness; yet, it also seems, upon reflection, that one would not have understood the joke without consciously *appreciating* all that information and its holistic humor-relevance. Second, there are strong theoretical grounds—really the same grounds as in the case of belief fixation—for the claim that the cognitive processes that underwrite joke-getting are essentially morphological: viz., joke-getting exhibits very similar kinds of holistic, Quineian-cum-isotropic, dependence on background information. This means that it is not tractably possible for joke-getting cognition to proceed by deploying explicit conscious representations of all pertinent items of information and their holistic, Quineian-ishly funnymaking, interconnections—especially not in the very short time it normally takes to get a joke.

## 7. *Chromatic Illuminated Epistemic Seemings*

A chromatically illuminated occurrent mental state is itself rife with dispositional potential, in ways directly connected to the morphological content doing the chromatic illuminating. One kind of looming potentiality, for example, is the capacity to spontaneously arrive at consciously explicit answers to suitable background-probing questions—in the case of the occurrent state of getting the babysitter joke, for instance, questions like “Roughly how old is the babysitter?”;

“Who is the “she” who figured out that the guy is having the affair with babysitter?”; etc. The presence of such looming potentialities normally is itself a matter of chromatic illumination, rather than something that is explicitly represented in consciousness. Moreover, this aspect of chromatic illumination is part of the phenomenological content-specificity of one’s experienced mirthfulness, as distinct from the phenomenology of finding oneself laughing with no clue about why. And even when *some* pertinent background information becomes explicitly present in consciousness, triggered by the joke-getting experience in combination with some other occurrent mental state (e.g., occurrently considering a probe question), normally much else that one is appreciating remains only implicitly present in consciousness in the form of chromatic illumination. In explaining the joke to someone who doesn’t yet get it, for example, normally one would bring explicitly to mind, and articulate, only a small part of the pertinent background information—e.g., the fact that twenty years ago, “like”-talk in English was employed primarily by teenage girls. That explicit information will trigger a joke-getting experience in one’s interlocutor only if the interlocutor now undergoes an occurrent conscious mental state that is suitably chromatically illuminated by the various other pertinent items of background information and also is suitably chromatically illuminated by how funniness supervenes, in a holistically Quineian way, on the full body of pertinent information.

Similar observations apply, *mutatis mutandis*, to the phenomenology of occurrent conscious beliefs. Normally, an occurrent belief that *P* is not experienced as evidentially un-tethered—as arising, for no apparent reason, “out of the blue.” (That would be the analogue of finding oneself laughing without a clue about why.) On the contrary, normally an occurrent belief that *P* is accompanied by an occurrent “epistemic seeming”—a phenomenological aspect as-of *P*’s *seeming true in virtue of specific evidence one possesses*. Often this aspect will itself be present in consciousness only implicitly in the form of chromatic illumination, rather than being present as explicit conscious belief alongside the conscious belief that *P*. And although some pertinent items of evidence might be explicitly present in consciousness alongside the occurrent belief itself, much of the isotropic, Quineian, evidential support for *P* that one possesses will be consciously appreciated only implicitly as chromatic illumination. Such chromatic illumination will endow the conscious belief-state with various looming potentialities—often including the capacity to bring explicitly into consciousness, piecemeal, certain specific components of the overall evidential support for *P* that is being appreciated implicitly and chromatically. Normally, however, even when certain evidentially pertinent items of information are explicitly present in consciousness and are explicitly being regarded as evidence for *P*, one’s consciously so regarding them will take place against a rich background of further evidentially pertinent information that remains only implicitly conscious in the form of chromatic illumination. Typically, therefore, articulating “the” justification of one’s belief that *P* will be similar to articulating “the” explanation of a joke: what comes to mind explicitly is only part of the story, and one’s appreciation of the (partial) relevance of this part of the story involves a rich further background of holistic, isotropic and Quineian, considerations that are consciously present only implicitly in the form of chromatic illumination.

## 8. *Doxastic Justification, II: Chromatic-Experiential Evidentialism and Chromatic-Experiential Hierarchical Pluralism*

At the close of Section 5 we urged each of the following claims: first, conscious experience plays a crucial and central role, constitutively, in doxastic justification; but second, the same goes for the essentially morphological aspects of abductively inferential belief-fixation processes. The key to reconciling these two claims is to embrace the further claim that the conscious phenomenological character of an occurrent belief that *P* can include, in the form of chromatic illumination, implicit appreciation of the holistic, isotropic/Quineian, evidential support that one possesses for *P*. The resulting position, which we will call *chromatic-experiential evidentialism*, asserts that the fundamental constitutive features of a doxastically justified belief that *P* are these: (1) one possess adequately strong holistic evidential support for *P*, (2) one experiences an epistemic seeming-that-*P* that is chromatically illuminated by this evidential support, and (3) one forms the occurrent belief that *P* because of this chromatically illuminated epistemic seeming.

Chromatic-experiential evidentialism, we submit, is much preferable to any of the three accounts of doxastic justification that we canvassed in Section 5. Those accounts, despite their differences from one another, are phenomenologically implausible in essentially the same way: they treat belief fixation as occurring very largely “out of the blue” in a merely gut-feeling way, as far as consciousness is concerned—rather than arising in a way that embodies implicit conscious appreciation of the pertinent evidential support that one possesses. Conscious experience is indeed constitutively centrally important for doxastic justification, but it is important precisely because it includes conscious epistemic seemings that, far from being blind to the essentially morphological aspects of holistic, isotropic/Quineian, evidential support, are richly chromatically illuminated by those very aspects.

We also maintain that a fully adequate account of doxastic justification should be somewhat pluralistic, urging the pertinence of hierarchy of goals and subgoals that are constitutively connected to belief (Henderson, Horgan, Potrč, and Tierney 2017; Horgan, Potrč, and Strahovnik, forthcoming). The hierarchy can be briefly and approximately characterized as follows—with level 4 being the principal constitutive *telos* of belief, and each successive level constituting an epistemic agent’s constitutive best means toward the goal at the level immediately above it:

Level 4: *Reliable veridicality*. Believing only what is true, by deploying a reliable process of belief fixation.

Level 3: *Objective rationality*. Believing only what is objectively highly likely to be true (relative to one’s total available evidence), by appreciating the objective import of one’s evidence and being gripped by belief through that appreciation. (This is one’s constitutively best available primary sub-goal as means toward reliable veridicality.)

Level 2: *Sensibility-based subjective rationality*. Believing only what is subjectively likely to be true according to one’s own deep epistemic sensibility (relative to one’s total available evidence), by appreciating the import of one’s evidence (according to one’s deep epistemic sensibility) and being gripped by belief through that appreciation. (This

is one's constitutively best available secondary sub-goal as means toward objective rationality.)

Level 1: *Experiential subjective rationality.* Believing only what accords with one's epistemic seemings that meet certain specific phenomenological criteria, by doing so on the basis of those seemings. (This is one's constitutively best available immediate sub-goal as means toward sensibility-based subjective rationality.)

On this pluralistic approach, normative evaluation of an epistemic agent's belief that *P*, in terms of the notion of doxastic justification, can be primarily focused on any specific level in this hierarchy, or on several together—depending on one's evaluative purposes in a specific context. Also, in order for an epistemic agent's belief that *P* to be doxastically justified relative any given level of the constitutive hierarchy, it also must be doxastically justified relative to each lower level.

Within this pluralistic framework, chromatic-experiential evidentialism would apply most directly to level-3 normative evaluation, since level 3 is the optimal kind of epistemic rationality to which an epistemic agent can aspire.<sup>6</sup> However, chromatically illuminated epistemic seemings would figure importantly vis-à-vis levels 1 and 2 as well—even though one's actual epistemic seemings sometimes might reflect a subjective epistemic sensibility that are imperfectly aligned with objective relations of evidential support, sometimes might fail even to accord with one's own subjective sensibility, and sometimes (as with many sincere climate-change deniers) might so ill-accord with the objective rational import of one's available evidence as to be epistemically benighted.

---

<sup>1</sup> The *physical* morphology that is characterized by this high-dimensional dynamical system is not the intrinsic physical structure of the neural network or of its successive activation-states. Rather, it is the structure of the network's overall *physical potentiality profile*—the full body of ways it would evolve from any given activation-state to subsequent activation-states. (This profile is subserved, of course, by the intrinsic physical structure of the network.) One can think of such potentiality profile as the locus of nature's "engineering design" for intelligent cognizers. For elaboration of how a goddess named Eva—a thought-experimental stand-in for evolution—might implement this design-project, thereby harnessing the enormously rich internal structure possessed by certain dynamical systems and by the physical potentiality profiles they characterize, see Horgan and Tienson (1992) and Horgan (2012).

<sup>2</sup> Horgan and Tienson did not deny, however, that intelligent human cognition needs to deploy occurrent mental representations some of which have language-like structure. On the contrary, they argued in favor of this claim; see Chapter Chapter 5 of Horgan and Tienson (1996), entitled "Why There Still Has to Be a Language of Thought, and What That Means," and Horgan (2010). Jerry Fodor and Zenon Pylyshyn (XXXX) had claimed that insofar as cognitive science invokes language-like mental representations that are subjected to structure-sensitive processing, neural-network aspects of cognitive-science models cannot be anything more than an alternative "implementation architecture" for the classical, computational, conception of cognition. But this reflected a failure to appreciate the possibility that processes like holistically abductive belief-fixation are too subtle and complex to constitute representation-level computation—as indeed is so, argued Horgan and Tienson.

<sup>3</sup> Differing versions of non-experiential reliabilism are possible, depending on how one construes the pertinent notion of reliability. One familiar construal is reliability in the cognitive agent's actual global environment. But Henderson and Horgan (refs) propose an alternative construal, which has the advantage *inter alia* of evading the so-called "new evil demon problem" for standard reliabilism that was posed by Cohen and Lehrer (ref). Henderson and Horgan's proposal is to construe the pertinent kind of reliability as what they call "trans-global reliability under suitable modulational control." More on this matter in Section XX below.

<sup>4</sup> In describing morphological content as a matter of "immediately efficacious" dispositions, we do not mean to deny that some occurrent event must trigger such a disposition. Rather, the point is this: in a cognitive system with morphological content *C*, once some occurrent event occurs which initiates a mental process to which *C* is relevant, the process evolves in a *C*-appropriate manner without deploying an explicit representation of *C* along the way.



---

<sup>5</sup> This specific cartoon, despite being culturally somewhat dated, illustrates especially well the apparent richness and extensiveness of chromatic illumination in joke-understanding cognition. We authors are constantly on the lookout for a more recent joke or cartoon that could serve this purpose equally well. Suggestions would be gratefully received.

<sup>6</sup> How does reliability of one's belief-forming processes figure in doxastic justification? First, although there are indeed reliability-involving normative desiderata concerning belief fixation, we maintain that the most fundamental one—the one linked constitutively to the truth goal—is what Henderson and Horgan call *transglobal* reliability, viz., reliability across a wide range of epistemically possible global environments. (The belief-forming processes of your brain-in-a-vat experiential duplicate are no less transglobally reliable than are yours, even though these processes are not reliable within the envatted brain's own actual global environment.) Second, transglobal reliability is itself so tightly intertwined conceptually with objective epistemic rationality—i.e., with likely truth, given one's available evidence—that the two cannot come apart. See Horgan and Henderson (2001, 2006, 2007, 2011 Chapters 3-5 and Section 7.1), Henderson, Horgan, and Potrč (2007).