

The Last Magic Show: A Blind Brain Theory of the Appearance of Consciousness

Abstract: According to the latest estimates, the human brain performs some 38 000 trillion operations per second. When you compare this to the amount of information that reaches conscious awareness, the disproportion becomes nothing short of remarkable. What are the consequences of this radical informatic asymmetry? The Blind Brain Theory of the Appearance of Consciousness (BBT) represents an attempt to 'explain away' several of the most perplexing features of consciousness in terms of information loss and depletion. The first-person perspective, it argues, is the expression of the kinds and quantities of information that, for a variety of structural and developmental reasons, cannot be accessed by the 'conscious brain.' Puzzles as profound and persistent as the now, personal identity, conscious unity, and most troubling of all, intentionality, could very well be kinds of illusions foisted on conscious awareness by different versions of the informatic limitation expressed, for instance, in the boundary of your visual field. By explaining away these phenomena, BBT separates the question of consciousness from the question of *how consciousness appears*, and so drastically narrows the so-called explanatory gap. If true, it considerably 'softens' the hard problem. But at what cost?

*How could they see anything else if they were prevented from moving
their heads all their lives?*

—Plato, *The Republic*

Introduction: The *Problematic Problem*

How many puzzles whisper and cajole and actively seduce their would-be solvers? How many problems *own* the intellect that would overcome them?

Consciousness is the riddle that offers its own solutions. We now know that many of the intuitions it provides are deceptive through and through, and we have our suspicions about many others. The obvious problem is that these faulty intuitions constitute the very explananda of consciousness. If consciousness *as it appears* is fundamentally deceptive, we are faced with the troubling possibility that we quite simply will not recognize the consciousness that science explains. It could be the case that the 'facts of our deception' will simply fall out of any correct theory of consciousness. But it could also be the case that a supplementary theory is required—a theory of the *appearance* of consciousness.

The central assumption of the present paper is that any final theory of consciousness will involve some account of multimodal neural information integration.¹ Consciousness is the product of a Recursive System (RS) of some kind, an evolutionary twist that allows the human brain to factor *its own operations* into its environmental estimations and interventions. Thinking through the constraints faced by any such system, I will argue, provides a parsimonious way to understand why consciousness *appears* the way it does. The ability of the brain to 'see itself' is severely restricted. Once we appreciate the way *limits on*

¹ The differentiation and integration that so fundamentally characterize conscious awareness necessitate some system accessing multiple sources of information gleaned from the greater brain. This assumption presently motivates much of the work in consciousness research, including Tononi's Information Integration Theory of Consciousness (2012) and Edelman's Dynamic Core Hypothesis (2005). The RS as proposed here is an idealization meant to draw out structural consequences perhaps belonging to any such system.

recursive information access are expressed in conscious experience, traditionally intractable first-person perspectival features such as the now, personal identity, and the unity of consciousness can be ‘explained away,’ thus closing, to some extent, the so-called explanatory gap.

The Blind Brain Theory of the Appearance of Consciousness (BBT) is an account of how an embedded, recursive information integration system might produce the peculiar structural characteristics we associate with the first-person perspective. In a sense, it argues that consciousness is so confusing because it literally *is* a kind of confusion.² Our brain is *almost* entirely blind to itself, and it is this interval between ‘almost’ and ‘entirely’ wherein our experience of consciousness resides.

The Facts of Informatic Asymmetry

There can be little doubt that the ‘self-conscious brain’ is myopic in the extreme when it comes to the greater brain. Profound *informatic asymmetry* characterizes the relationship between the brain and human consciousness, a dramatic quantitative disproportion between the information actually processed by the brain and the information that finds its way to consciousness. Short of isolating the dynamic processes of consciousness within the greater brain, we really have no reliable way to quantify the amount of information that makes it to consciousness. Inspired by cybernetics and information theory, a number of researchers made attempts in the 1950's and early 1960's, arriving at numbers that range from less than 3 to no more than 50 bits per second—almost preposterously low (Norretranders, 1999). More recent research on attentional capacity, though not concerned with quantifying ‘mental workload’ in information theoretic terms, seems to confirm these early findings (Marois and Ivanoff, 2006). Assuming that this research only reflects one aspect of the overall ‘bandwidth of consciousness,’ we can still presume that whatever number researchers ultimately derive will be surprisingly low. Either way, the gulf between the 7 numbers we can generally keep in our working memory and the estimated 38 000 trillion operations per second (38 *petaflops*) equivalent processing power (Greenemeier, 2009) possessed by the average human brain is *boggling* to say the least.³

Consider the relative informatic poverty of experiences like pain, or anger, or even insight. Our affects only provide the smokiest grasp of the brain's inner-workings. Even the apparent richness of our sensorium belies an informatic abyss. Not only does the ‘transparency of experience’ blot out all awareness of sensory processing, phenomena such as change blindness and inattention blindness show that much of the informatic richness we attribute to our perceptual awareness is assumptive. We generally don't possess the information we think we do!

One need only ask, What is your brain doing *now*? to appreciate the vertiginous extent of informatic asymmetry. For my own part, I can adduce some neuroanatomical terms, talk about gross changes in activation as tracked in various neuroimaging studies, and so on. But when it comes to the nitty gritty, I really don't have the faintest clue—and neither does any neuroscientist living. Given this, you might expect a large and vibrant literature on the topic. But with the possible exception of Daniel Dennett (1987, 1991, 2005), no theorist has ever considered the possibility that this fact, as glaring as it is, could inform, let alone provide, a *general* account of (what we call) consciousness.

This oversight becomes all the more curious when one considers that radical informatic asymmetry is pretty much what we should expect given the structure and natural history of the brain. At

² In this respect, the Blind Brain Theory seems to require the kind of collapse of perception and cognition found in Andy Clark's Action-Oriented Predictive Processing Model (2012) in addition to recursive integration.

³ ‘Information,’ of course, is a notoriously nebulous concept. Rather than feign any definitive understanding, I will simply use the term in the brute sense of ‘systematic differences,’ and ‘processing’ as ‘systematic differences making systematic differences.’ The question of the *semantics* of these systematic differences has to be bracketed for reasons we shall soon see. The idea here is simply to get a certain theoretical gestalt off the ground.

some point in our recent evolutionary past, perhaps coeval with the development of language,⁴ the human brain became more and more *recursive*, which is to say, more and more able to factor its own processes into its environmental interventions. Many different evolutionary fables may be told here, but the important thing (to stipulate at the very least) is that some twist of recursive information integration, by degrees or by leaps, led to human consciousness. Somehow, the brain developed the capacity to ‘see itself,’ more or less.

It is important to realize the peculiarity of the system we’re discussing here. The RS qua neural information processor is ‘open’ insofar as information passes through it the same as any other neural system. The RS qua ‘consciousness generator,’ however, is ‘closed,’ insofar as *only recursively integrated information reaches conscious awareness*. Given the developmental gradient of evolution, we can presume a gradual increase in capacity, with the selection of more comprehensive sourcing and greater processing power culminating in the consciousness we possess today.

There is no reason, however, to think that consciousness as we experience it represents anything more than one possible configuration, a mere point in the space of all possible consciousnesses. So if consciousness began as something dim and dismal, the product of some primitive precursor to the RS, how far has it progressed? The answer to this question, it seems, depends on the adequacy and extent of the RS—the way it is structurally and functionally ‘positioned’ vis a vis the greater brain. If the only information to reach consciousness is information that is recursively integrated, then the adequacy of consciousness depends, in no small measure, on the kinds of developmental and structural constraints confronting the RS. And we have good reason, I think, to believe these are (and were) quite severe.

There’s the issue of evolutionary *youth*, for one. Even if we were to date the beginning of modern consciousness as far back as, say, the development of hand-axes, that would only mean some 1.4 million years of evolutionary ‘tuning.’ By contrast, the brain’s ability to access and process external environmental information is the product of hundreds of millions of years of natural selection. In all likelihood, the RS is an assemblage of ‘kluges,’ the slapdash result of haphazard mutations that produced some kind of reproductive benefit (Marcus, 2008).

There’s its *frame*, for another. As far as informatic environments go, perhaps nothing known is more complicated than the human brain. Not only is it a mechanism with some 100 billion parts possessing trillions of interconnections, it continually rewires itself over time. The complexities involved are so astronomical that we literally cannot *imagine* the neural processes underwriting the comprehension of the word ‘imagine.’ Recently, the National Academy of Engineering named reverse-engineering the brain one of its Grand Challenges: the first step, engineer the supercomputational and nanotechnological tools *required to even properly begin* (National Academy of Engineering, 2011).

And then there’s its *relation* to its object. Where the brain, thanks to locomotion, possesses a variable relationship to its external environment, allowing it to *selectively* access information, the RS is quite literally *hardwired* to the greater, nonconscious brain. Its information access is a function of its structural integration, and is therefore fixed to the degree that its structure is fixed. The RS must transform its structure, in other words, to attenuate its access.⁵

These three constraints—evolutionary contingency, frame complexity, and access invariance—actually paint a quite troubling picture. They sketch the portrait of an RS that is developmentally gerrymandered, informatically overmatched, and structurally imprisoned—the portrait of a human brain that likely possesses only the merest glimpse of its inner workings. As preposterous as this might sound to some, it becomes more plausible the more cognitive psychology and neuroscience learns. Not a week

⁴ Since language requires the human brain recursively access and translate its own information for vocal transmission, and since the limits of experience are also the limits of what can be spoken about, it seems unlikely that the development of language is not somehow related to the development of consciousness.

⁵ As seems to be the case with dedicated practitioners of meditation.

passes, it seems, without some new quirk of human cognition finding its way to the research headlines. One might reasonably ask, How many quirks does it take?

Most everyone, however, is inclined to think the potential for deception only goes so far--eliminativists included!⁶ The same evolutionary contingencies that constrain the RS, after all, also suggest the *utility* the information it accesses. We have good reason to suppose that the information that makes it to consciousness is every bit as *strategic* as it is fragmental. We may only 'see' an absurd fraction of what is going on, but we can nevertheless assume that it's the fraction that *matters most*...

Can't we?

The problem lies in the dual, 'open-closed' structure of the RS. As a natural processor, the RS is an *informatic crossroads*, continuously accessing information from and feeding information to its greater neural environment. As a consciousness generator, however, the RS is an *informatic island*: only the information that is integrated finds its way to conscious experience. This means that the *actual* functions subserved by the RS within the greater brain—the way it finds itself 'plugged in'—are no more accessible to consciousness than are the functions of the greater brain. And this suggests that consciousness likely suffers any number of profound and systematic misapprehensions.

This will be explored in far greater detail, but for the moment, it is important to appreciate the truly radical consequences of this, even if only as a possibility. Consider Daniel Wegner's (2002) account of the 'feeling of willing' or volition. Given the information available to conscious cognition, we generally assume the function of volition is to 'control behaviour.' Volition seems to come first. We decide on a course of action, *then* we execute it. Wegner's experiments, however, suggest what Nietzsche (1967) famously argued in the 19th century: that the 'feeling of willing' is *post hoc*. Arguing that volition as it appears is illusory, Wegner proposes that the actual function of volition is to take social *ownership* of behaviour.

As we shall see, this is precisely the kind of intuitive/experimental impasse we might expect given the structure of the RS. Since a closed moment of a far more extensive open circuit is all that underwrites the feeling of willing, we could simply be wrong. What is more, our intuition/assumption of 'volition' may have *no utility whatsoever* and yet 'function' perfectly well, simply because it remains systematically related to what the brain is actually doing.

The information integrated into consciousness (qua open) could be causally efficacious through and through and yet so functionally opaque (qua closed) that we can only ever be deluded by our attendant second-order assumptions. This means the argument for cognitive adequacy from evolutionary utility in no way discounts the problem that information asymmetry poses for consciousness. The question of whether brain makes use of the information it makes use of is trivial. The question is whether self-consciousness is *massively* deceived. A beggar who delivers a million dollars from one inscrutable neural mandarin to another is a beggar all the same.

The RS is at once a neural cog, something that subserves larger functions, *and* an informatic bottleneck, the proximal source of every hunch, every intuition, we have regarding who we are and what we do. Its reliability as a source literally depends on its position as a cog. This suggests that all speculation on the 'human' faces what might be called the *Positioning Problem*, the question of how far

⁶ Even Paul Churchland (1989) eventually acknowledged the 'epistemic merit' of folk psychology—in a manner not so different than Dennett. BBT, as we shall see, charts a quite different course: by considering conscious cognition as something structurally open but reflectively closed to the cognitive activity of the greater brain, it raises the curious prospect (and nothing more) that 'folk psychology' or the 'intentional stance' *as reflectively understood* (as normative, intentional, etc.) is largely an artifact of reflection, and only seems to possess utility because it is reliably paired with inaccessible cognitive processes that are quite effective. It raises the possibility, in other words, that *belief as consciously performed* is quite distinct from *belief as self-consciously described*, which could very well suffer from what might be called 'meta-recursive privation,' a kind of 'peep-hole view on a peep-hole view' effect.


consciousness can be trusted to understand itself.⁷ As we have seen, the radicality of information asymmetry, let alone the problems of evolutionary contingency, frame complexity, and access invariance, suggests that our straits could be quite dire. The Blind Brain Theory of the Appearance of Consciousness simply represents an attempt to think through this question of information and access in a principled way: to speculate on what our ‘conscious brain’ can and cannot see.

Chained to the Magician: Encapsulation

One of the things that make consciousness so difficult to understand is intentionality. Where other phenomena simply ‘resist’ explanation, intentional phenomena seem to be intrinsically *antagonistic* to functional explanation. Like magic tricks, one cannot explain them without apparently explaining them away. As odd as it sounds, BBT proposes that we take this analogy to magic at its word.⁸ It presumes that intentionality and other ‘inexplicables’ of consciousness like presence, unity, and personal identity, are best understood as ‘magic tricks,’ artifacts of the way the RS is a prisoner of the greater, magician brain.

All magic tricks turn on what might be called *information horizons*: the magician literally leverages his illusions by manipulating what information you can and cannot access. The spectator is *encapsulated*, which is to say, stranded with information that *appears sufficient*. This gives us the structure schematically depicted in Fig. 1.

Encapsulation: Magic Act



⁷ You could say the ‘positional reappraisal’ of experience and conscious cognition in the light of psychology and neuroscience is well underway. Keeping with our previous example, something like volition might be called a ‘tangled, truncated, compression heuristic.’ ‘Tangled,’ insofar as its actual function (to own behaviour post hoc) seems to differ from its intuitive function (to control behaviour). ‘Truncated,’ to the extent it turns on a kind of etiological anosognosia. ‘Compressed,’ given how little it provides in the way of phenomenal and/or cognitive information. And ‘heuristic’ insofar as it nevertheless seems to facilitate social cognition (though not in the way we think).

⁸ Dennett (2005) makes extensive use of magic to explain the intuitions of those convinced of the irreducibility of consciousness, but more as a loose metaphor than anything else. As will become apparent, BBT utilizes the analogy in a much more thoroughgoing manner.

Fig. 1 In a magic act, the magician games the spectatorial information horizon to produce seemingly impossible effects, given the spectators' existing expectations. Since each trick relies on convincing spectators they have all the information they need, a prior illusion of 'sufficiency' is required to leverage the subsequent trick.

Apparent sufficiency is all important, since the magician is trying to gull you into forming false expectations. The 'sense of magic' arises from the disjunct between these expectations and what actually happens. Without the facade of informatic sufficiency—that is to say, without encapsulation—the most the trick can do is surprise you. This is the reason why explaining magic tricks amounts to *explaining away* the 'magic': explanations provide the very information that must be sequestered to foil your expectations.

I once knew this magician who genuinely loved ad hoc, living-room performances. I found his coin tricks particularly impressive—he preferred using a coin the size of a soup can lid! Gangs of us would peer at his hands, baffled as the thing vanished and rematerialized. Later in the night, I would eventually have a chance to watch him perform the very tricks that had boggled me earlier *from over his shoulder*. In other words, I was able to watch the identical process from an entirely different etiological perspective. It has struck me as a provocative analogy for consciousness ever since, especially the way the mental seems to 'disappear' when we look over the *brain's* shoulder.

So how strong is the analogy? In both cases you have encapsulation: the RS has no more recursive access to 'behind the scenes' information than *the brain gives it*. In both cases the product—magic, consciousness—seems to vanish as soon as information regarding etiological provenance becomes available. The structure, as Fig.2 suggests, is quite similar.

Encapsulation: Consciousness




Fig. 2 In consciousness we find a similar structure. But where the illusion of sufficiency is something the magician must bring about in a magic act, it simply follows in consciousness, given that it has no access whatsoever to any 'behind the scenes.' In this analogy, intentional phenomena are like magic to the extent that the absence of actual causal histories, 'groundlessness,' seems to be constitutive of the way they appear.

In this case we have multiple magicians, which is to say, any number of occluded etiologies. Now if the analogy holds, intentional phenomena, like magic, are something the brain can only cognize *as such* in the absence of the actual causal histories belonging to each. They require, in other words, the absence of certain kinds of information *to make sense*.

It's worth noting, at this juncture, the way Fig. 2 in particular captures the 'open-closed structure' attributed to the RS. Given that some kind of integration of differentiated information is ultimately responsible for consciousness, this is precisely the 'magical situation' we should expect: a system that is at once open to the functional requirements of the greater brain, and yet closed by recursive availability. If the limits on recursive availability provide the informatic basis of our intentional concepts and intuitions, then the adequacy of those concepts and intuitions would depend on the adequacy of the information provided. Systematic deficits in the one, you can assume, would generate systematic deficits in the other.

So, is it really possible that consciousness is a kind of *coin trick*?

The problem, of course, lies in the extreme eliminativism this seems to entail. But, given that we are genuinely open thinkers, savvy to the dismal findings of cognitive psychology (Mercier and Sperber, 2011) and so painfully aware of the way our prejudices thoroughly warp and skew our every theoretical encounter, we should set this consideration aside. There's the consciousness we *want* to have, and then there's the consciousness we *have*. The trick to finding the latter could very well turn on finding our way past the former.

The analogy warrants an honest look, at the very least. In what follows, I hope to show you a genuinely novel and systematically parsimonious way to interpret the first-person perspective, one that resolves many of its famous enigmas by treating them as a special kind of 'magic': something to be explained away. As it turns out, you are entirely what a roaming, recursive storm of information *should* look like—from the inside.

The Unbounded Stage: Sufficiency

If the neural correlates of consciousness possess *information horizons*, how are they expressed in self-conscious experience?

Is it just a coincidence that the first-person perspective *also* possesses a 'horizontal structures'?

We *already know* consciousness as it appears is an informatic illusion *in some respects*. We also know that consciousness only gets a 'taste' of the boggling complexities that make it possible. When we talk about consciousness and its neural correlates, we are talking about a dynamic *subsystem* that possesses a very specific informatic relationship with a greater system: one that is not simply profoundly asymmetrical, *but asymmetrical in a structured way*.

As informatically localized, the RS has to possess any number of information horizons, 'integration thresholds,' were the information we experience is taken up. To say that the conscious brain possesses 'information horizons' is merely to refer to the way the RS qua consciousness generator constitutes a closed system. When it comes to information, consciousness 'gets what it gets.'

As trivial as this observation is, it is precisely where things become interesting. Why? Because if some form of recursive neural processing *simply is* consciousness, then we can *presume encapsulation*.⁹ If we can presume encapsulation, then we can presume the *apparent sufficiency* of information accessed. Since the insufficiency of accessed information will always be a matter of *more information*, sufficiency will be the perennial default. Not only does consciousness get what it gets, it gets *everything* to be gotten.

⁹ This is not to be confused with 'information encapsulation' as used in Pylyshyn (1999) and debates regarding modularity. Metzinger's account of 'autoepistemic closure' somewhat parallels what is meant by encapsulation here. As he writes, "'autoepistemic closure' is an epistemological and not (at least not primarily) a phenomenological concept. It refers to an 'inbuilt blind spot,' a structurally anchored deficit in the capacity to gain knowledge about oneself" (2003, p. 57). As an intentional concept embedded in a theoretical structure possessing many other intentional concepts, however, it utterly lacks the explanatory resources of encapsulation, which turns on a non-semantic concept of information.

Why is default sufficiency important? For one, it suggests that neural information horizons will express themselves in consciousness in a very peculiar way. Consider your visual field, the way seeing simply vanishes into some kind of asymptotic limit—a *limit with one side*. Somehow, our visual field is literally encircled by a blindness *that we cannot see*, leaving visual attention with the peculiar experience of experience *running out*. Unless we suppose that experience is utterly bricked in with neural correlates (which would commit us to asserting that we possess ‘vision-trailing-away-into-asymptotic-nothingness’ NCs), it seems obvious to suppose that the edge of the visual field is simply where the visual information available for conscious processing comes to an end.

The edge of our visual field is an example of what might be called *asymptotic limits*.

The Limit With One Side




Fig. 3 The edge of our visual field provides a striking example of the way information horizons often find conscious expression as ‘asymptotic limits,’ intramodal boundaries that only possess one side. Given that we have no *visual* information pertaining to the limits of vision, the boundary of our visual field necessarily remains *invisible*. This structure, BBT suggests, is repeated throughout consciousness, and is responsible for a number of the profound structural features that render the first-person perspective so perplexing.

An asymptotic limit, to put it somewhat paradoxically, is a limit that *cannot* be depicted the way it’s depicted above. Fig. 3 represents the limit as *informatically framed*; it provides the very information that asymptotic limits sequester and so dispels default sufficiency.

Limits *with one side* don’t allow graphic representation of the kind attempted in Fig. 3 because of the way these models shoehorn all the information into the visual mode. One might, for instance, object that asymptotic limits confront us all the time without, as in the case of magic, the attendant appearance of sufficiency. We see the horizon knowing full well the world outruns it. We scan the skies knowing full well the universe outruns the visible. Even when it comes to our visual field, we *know* that there’s always ‘more than what meets the eye.’ Nevertheless, seeing all there is to see at a given moment is what renders each of these limits asymptotic. We possess no visual information regarding the limits of our visual information. All this really means is that asymptotic limits and their attendant sufficiencies are mode specific. You could say our ability to informatically frame our visual field within memory, anticipation, and cognition is the only reason we can intuit its limitations *at all*. To paraphrase Glaucon from the epigraph, one has to see more to know there is more to see (Plato, 1987, p.317).

This complication of asymptotic limits and sufficiencies is precisely what we should expect, given the *integrative* function of the RS. Say we parse some of the various information streams expressed in consciousness as depicted in Fig. 4.

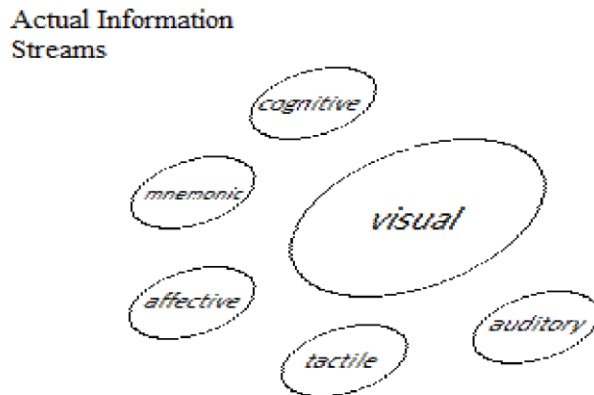


Fig. 4 This graphic, as simple as it is, depicts various informatic modalities in a manner that bears information regarding their distinction. They are clearly bounded and positioned apart from one another. This is precisely the kind of information that, in all probability, would not be available to the RS, given the constraints considered above.

Each of these streams is discrete and disparately sourced prior to recursive integration. From the standpoint of recursive availability, however, we can look at each of these circles as ‘monopolistic spotlights,’ points where the information ‘lights up’ for conscious awareness. Given the unavailability of information pertaining to the spaces between the lights, we can assume they would not even exist for consciousness. Recursive availability, in other words, means these information streams would be ‘sutured,’ bound together as depicted in Fig. 5.¹⁰

¹⁰ As we shall see below, this has important consequences regarding the question of the unity of consciousness.




Fig. 5 Given the asymptotic expression of informatic limits in conscious awareness, we might expect the discrete information streams depicted in Fig. 4 to appear to be ‘continuous’ from the standpoint of consciousness.

The local sufficiencies of each mode simply run into the sufficiencies of other modes forming a kind of ‘global complex’ of sufficiencies with their corresponding asymptotic limits. Once again, the outer boundary as depicted above needs to be considered heuristically: the ‘boundaries of consciousness’ do not possess any ‘far side.’ It’s not so much a matter of the sufficiency of the parts contributing to the sufficiency of the whole as it is a question of availability: absent any information regarding its global relation to its neural environment, that environment does not exist, not even as the ‘absence’ depicted above. Even though it is the integration of modalities that make the local limits of any one mode (such as vision) potentially available, there is a sense in which the global limit *has to always outrun recursive availability*. As strange as it sounds, consciousness is ‘amoebic.’ Whatever is integrated is encapsulated, and encapsulation means asymptotic limits and sufficiency. Given the open-closed structure of the RS, you might say that a kind of ‘asymptotic absolute’ has to afflict the whole, and with it, what might be called ‘persistent global sufficiency.’¹¹

So what we have, then, is a motley of local asymptotic limits and sufficiencies bound within a global asymptotic limit and sufficiency. What we have, in other words, is an outline for something not so unlike consciousness as it appears to us. At this juncture, the important thing to note is the way it seems to simply fall out of the constraints on recursive integration. The suturing of the various information streams is not the *accomplishment* of any specialized neural device over and above the RS. The RS simply lacks information pertaining to their insufficiency. The same can be said of the asymptotic limit of the visual field: Why would we posit ‘neural correlates of vanishing vision’ when the simple absence of visual information is device enough?

So far, on the strength of my magical analogy and our visual field, we have isolated four concepts:

¹¹ Thus the profound monotonicity of consciousness: As an encapsulated *product* of recursive availability, the availability of new information can never ‘switch the lights out’ on existing information.

Information horizons: The boundaries that delimit the recursive neural access that underwrites consciousness.

Encapsulation: The global result of limited recursive neural access, or information horizons.

Sufficiency: The way the lack of *intra*-modal access to information horizons renders a given modality of consciousness ‘sufficient,’ which is to say, at once all-inclusive and unbounded at any given moment.

Asymptotic limits: The way information horizons find phenomenal expression as ‘limits with one side.’

We began by asking how information horizons might find phenomenal expression. What makes these concepts so interesting, I would argue, is the way they provide direct *structural* correlations between certain peculiarities of consciousness and possible facts of brain. They also show us that how what seem to be *positive features* of consciousness can arise *without neural correlates to accomplish them*. Once you accept that consciousness is the result of a special kind of informatically localized neural activity, information horizons and encapsulation directly follow. Sufficiency and asymptotic limits follow in turn, once you ask what information the conscious brain can and cannot access.

Moving on, I hope to show how these four concepts, along the open/closed structure of the RS, can explain some of the most baffling structural features of consciousness. By simply asking the question of what kinds of information the RS likely lacks, *we can reconstruct the first-person*, and show how the very things we find the most confusing about consciousness—and the most difficult to plug into our understanding of the natural world—are actually confusions.

The Always Absent Fakir: Presence

What could make an experience come into existence?

This eminently reasonable and apparently innocuous question lies at the heart of what might be called the ‘accomplishment fallacy.’ For decades now, naturalists like Dennett have argued that consciousness is just another natural phenomenon, something that will be figured out in due time, while naturalists like Chalmers have argued that consciousness is anything but ‘another’ natural phenomenon. Both camps have their apparently inexhaustible wells of apparent reasons. But neither camp has been able to do much more than flounder when it comes to the most perplexing characteristics of consciousness.

What, for instance, could make the experience of *now* come into existence?

The now is a perplexity as ancient as philosophy itself, the conceptual hinge from which both being and becoming seem to swing. As Aristotle puts it, “it is not easy to see whether the moment [*nun*] which appears to divide the past and the future (1) always remains one and the same or (2) is always distinct” (1969, pp. 79). Given the way it has divided philosophers from ancient times to the present day, Aristotle’s observation that time is ‘not easy’ has to count as one of philosophy’s great understatements. As Augustine famously writes: “What then is time? Provided that no one asks me, I know. If I want to explain it to an enquirer, I do not know” (1992, pp. 230). In 1908, John McTaggart used the contradictory nature of the now to argue the unreality of time altogether, engendering the long-standing tensed versus tenseless time debate in Anglo-American philosophy. In 1926, Martin Heidegger (1996) used a novel

reinterpretation of time to anchor both his positive and his critical projects, setting the paradoxical stage for several major strands of Continental philosophy.¹²

The philosophical significance of the now simply reflects its importance more generally. The first-person perspective is paradigmatically ‘now.’ Any theory that fails to account for it fails to explain a central structural feature of consciousness as it is experienced. It certainly speaks to the difficulty of consciousness that it takes one of the most vexing problems in the history of philosophy as a *component*!

So then what is the now?

Time is generally thought to consist of succession of *times*, a series such as,

$$t1 > t2 > t3 > t4 > t5$$

The problem with such conceptions, many philosophers have argued, is that it fails to capture the peculiar *phenomenological* nature of the now.¹³ The above representation, for instance, is only apprehended insofar as it is present, which is to say, *now*. In phenomenological terms, other times only seem to exist insofar as they are framed within some present time,

$$t0 (t1 > t2 > t3 > t4 > t5)$$

where $t0$ represents the ‘now’ of apprehending our original notation. Our first-person *experience* of temporal succession, in other words, seems to be far different than the series suggested above. Rather than experience time like a sequence of beads on a string, each passing moment seems to somehow *encompass* the moment preceding, to become the very frame of reference through which the past and future can be thought. The movement of experienced time is what might be called ‘meta-inclusionary,’ with the present now somehow ‘containing’ the past and the future. From the perspective of phenomenological time then, succession actually looks more like this,

$$t1 > t2 (t1) > t3 (t2 (t1)) > t4 (t3 (t2 (t1))) > t5 (t4 (t3 (t2 (t1))))$$

where the stacking of parentheses represents the movement of inclusion. A subsequent now, $t2$, includes $t1$, only to be included by a subsequent now, $t3$, which includes $t2 (t1)$, and so on, where ‘includes’ means, ‘becomes the temporal frame of reference for...’ Given the limitations of memory, of course, past nows are not conserved in the manner represented here, but dissolve into a kind of ‘generalized past’ instead, allowing us to simplify our notation thus,

$$t1 > t2 (t1) > t3 (t2 (tN)) > t4 (t3 (tN)) > t5 (t4 (tN))$$

where ‘ tN ’ represents the past as conserved in long term memory.

But even this falls short representing time as it is experienced to the extent that it fails to capture the *continuity of the now*, the paradoxical sense in which this now is somehow always the *same*, even though it is manifestly different. What we require, in other words, is a notation that represents the movement of temporal succession within some kind of *temporal stasis*,

$$t0 (t1) > t0 (t2 (tN)) > t0 (t3 (tN)) > t0 (t4 (tN)) > t0 (t5 (tN))$$

¹² Like Husserl (1964), Heidegger sees temporality as the framework through which experience becomes possible, and so makes it the cornerstone of his interpretation of *Dasein*. He also makes it the basis of what he calls the ‘metaphysics of presence,’ the Big Wrong Turn he attributes to Aristotle’s ancient analysis of time, a diagnosis shared by Derrida and an operational assumption of deconstruction.

¹³ Following Heidegger (1996).

where ‘t0’ represents the sameness of the now that frames our experience of temporal difference.

Represented this way, the paradoxical nature of first-person time is clear to see. It’s almost as if time passes within experience, but not *for* experience. The now, as Aristotle (1969) pointed out so very long ago, is both the same and different. So from the accomplishment perspective, the challenge for any theory of consciousness is to explain *what makes this happen*. As Thomas Metzinger puts it in his magisterial *Being No One*:

The phenomenal experience of time in general is constituted by a series of important *achievements*. They consist in the phenomenal representation of temporal identity (experienced simultaneity), of temporal difference (experienced nonsimultaneity), of seriality and unidirectionality (experienced succession of events), of temporal wholeness (the generation of a unified present, the ‘specious’ phenomenal Now), and the representation of temporal permanence (the experience of duration). (2003, p. 127)

The problem Metzinger confesses, is one of fitting these ‘achievements’ together in a phenomenologically satisfying way. He uses the analogy of an island, representing the simultaneity of presence, standing within a river, representing succession. “It is not only the island that is located in the river but in a strange sense the river is flowing *through* the island itself” (2003, p. 127). The now, in other words, somehow frames the experience of what frames it. Phenomenologically, it somehow stands both inside and outside time. ‘Not easy to see,’ indeed.

The neuropsychological investigation of temporal perception has revealed a variety of ‘thresholds’ for the conscious perception of simultaneity, nonsimultaneity, temporal ordering, and presence. In terms of what Wackerman (2007) calls the ‘inner horizons’ of perception, many of the findings merely refine those of William James (2007) and his contemporaries at the end of the 19th century. The findings vary according to a diverse variety of factors, particularly with regards to modality, but also with respect to stimulus type and intensity, training, locomotor, emotional or attentional state, and so on. In each case however, it is possible to track the point at which subjects simply cannot accurately discern simultaneity, nonsimultaneity, or temporal ordering.

In addition to these inner horizons, Wackerman discusses what he calls the ‘outer horizons’ of temporal perception and cognition. Referring to James’ conception of the ‘sensible present,’ he writes, “[c]ontemporary research into time perception and timing behavior has surprisingly little to add to it, except an updated terminology and an extended experimental database” (2007, p. 25). Findings from a variety of sources converge on a ‘window of presence’ some 2 to 3 seconds in length, beyond which ‘perceived unity in time’ disintegrates, and reproductive memory (cognition) takes over.¹⁴ In effect, the now is a kind of *temporal field*, an ‘integration window’ which binds stimuli into singular percepts.

The now, in other words, *possesses its own asymptotic limit*, one analogous to the edge of visual field. Where the limit of the visual field simply marks the point at which conscious access to immediate visual information ends, we could surmise that the limit of the temporal field marks the point at which conscious access to immediate temporal information ends (and where, likewise, we are forced to rely on cognition, which is to say, alternative modalities of information access). Since the conscious brain cannot access information regarding the limits of the temporal information it accesses, the information it receives always appears modally sufficient: as with the visual field, the temporal field becomes something only cognition can ‘situate.’

In other words, the same way modal sufficiency means that we see against an inability to see, we can say that we *time* against an inability to time. Insofar as both modalities are finite, there is a sense in

¹⁴ See also, Pöppel (2009, p. 1891).

which this is simply platitudinal. *Of course* we can't time what we can't time. *Of course* the RS can only process the information it accesses. BBT simply draws our attention to the structural implications of these platitudes, to the way the limitations that we know pertain to the conscious systems of the brain cast experiential shadows—some, like the now, possessing profound implications.

So, to return to our notation above,

$$t_0(t_1) > t_0(t_2(t_N)) > t_0(t_3(t_N)) > t_0(t_4(t_N)) > t_0(t_5(t_N))$$

we can see that 't0,' which was taken to represent the 'sameness of the now,' is actually *the phenomenological expression of a temporal information horizon*.

The point warrants a pause, given that it purports to provide a potential empirical solution to what has to be one of human thought's great conundrums. The apparent sameness of the now, BBT suggests, is simply a 'structural side-effect,' an illusion generated by the modal sufficiency of temporal perception: 't0' is, quite literally, *no time*, the asymptotic limit of our temporal field. Since the absence of information is the absence of *differentiation*, we are stranded with the illusion of some kind of abiding temporal identity—the now.

I appreciate how difficult this is to think, but it directly follows from the identification of consciousness with recursive information integration. Consider Fig. 6.

Temporal 'Blotting'




Fig. 6 Since the information available to the RS at any given moment comprises the totality of temporal consciousness, the first-person passage of time, or ‘flow,’ takes on a peculiar, meta-inclusionary structure, with each now asymptotically ‘blotting’ the now previous. The ‘past’ only exists insofar as it is informatically *embedded* within each successive present. Every ‘this-moment-now’ paradoxically becomes the ‘first moment of the rest of your life.’

Whatever crosses the integration threshold (information horizons) of the RS at any given moment becomes conscious *at that moment*. This information is quite simply *all there is*, at least so far as consciousness is concerned. Each ‘this-moment-now’ is sufficient as a result; the present always ‘fills the temporal screen,’ so to speak. The RS, however, is a finite *mechanism*, one that continually transitions through a dynamically evolving series of information states. Each ‘this-moment-now,’ in other words, arises only to be replaced by some successor. Since the *ongoing* informatic economy of the RS only retains the barest informatic outline of its prior states, preceding ‘this-moment-nows’ find themselves utterly blotted, ‘present’ only in the way they condition ongoing processes. The only way the RS could ‘retain’ all the information previously integrated would be to generate an endless stream of replicas of itself: a physical impossibility. Instead, old information is continually dumped, and the same integrative neuromachinery is continually reused, and each now incessantly blots the previous from existence, becoming the informatic frame for its scant remains.

Each moment of the RS is the *whole* of consciousness, which is why each moment seems the ‘same’ moment, and which is why prior moments only ‘exist for’ consciousness *within this very moment of consciousness*. As I suggested above, consciousness is *amoebic* in a peculiar way:¹⁵ the only differences (information) it can assess are the differences *it assesses*. Any past limit on recursive availability can only become available at the cost of some *new* limit on recursive availability. This is the sense in which I referred to consciousness possessing both persistent global sufficiency and a corresponding asymptotic absolute. Lacking any global first-order temporal information, the RS quite literally cannot distinguish itself *from itself* over time, and so generates no (temporal) consciousness of its transformation. Now is always... *now*. Since it also lacks the means of distinguishing this default from the asymptotic absolute more generally, the now becomes coextensive with consciousness, and ‘persistent global sufficiency’ becomes ‘presence.’

So what could make the paradoxical experience of now come into existence?

BBT answers, *nothing*.

Intermission: Anosognosia

Our experience of *time* is an achievement. Our experience of *nowness*, on the other hand, is a structural side-effect. The same way our visual field is boundless and yet enclosed by an inability to see, our temporal field—this very moment *now*—is boundless and yet enclosed by an inability to time. This is what makes the now so perplexing, so difficult to grasp: it is what might be called an ‘occluded structural property of experience.’ Metzinger, for instance, calls the ‘phenomenal texture of time’ a paradigmatic example of ‘evasiveness,’ something that melts away every time we attempt to grasp it: “Before starting to discuss phenomenological constraints governing the conscious experience of time I herewith fully admit that I am not able to give an even remotely convincing conceptual analysis of what it *is* that actually creates the philosophical puzzle” (2003, p. 153). As a neural accomplishment, Metzinger’s now has to be a special kind of representation, ‘special’ because it isn’t a representation of anything that

¹⁵ These considerations underwrite at least two of Heidegger’s (1996) guiding intuitions in his descriptions of phenomenological time: the notion that temporality ‘temporalizes,’ and the notion that temporality ecstatically ‘stretches,’ that it includes itself in the movement of transcending itself.

actually exists. He needs his illusion to *be something*, a neurally instantiated ‘simulation fiction,’ one capable of generating the experience of, to refer to his earlier metaphor, an island of presence within the river of time that flows through the selfsame island of presence.

According to BBT, however, the now just *is* an illusion. The *assumption* of nowness obviously has some neural correlate, but there is literally *nothing else* to the ‘now.’ As modally sufficient, it seems all-encompassing, which is why the ‘river’ flows through the ‘island.’ As cognitively supplemented by mechanisms such as those involved in reproductive or episodic memory and narrative, it becomes an island ‘in’ the river once again, a ‘window’ onto a greater temporal world. The degree to which attention is directed to its modal sufficiency, the more the past and future seem to be, as Augustine (1992) argued so long ago, artifacts of the present. The degree to which attention is directed to its cognitive supplementation, the more it seems to be an illusory artifact of human finitude.¹⁶

In other words, you could say the now reflects the open-closed structure of the RS *from the inside*. Its continual difference reflects the openness, the continual entry/re-entry of information. Its abiding identity, on the other hand, expresses the closure imposed by recursive availability, the fact of encapsulation. This really is a remarkable parallel, if you ponder it, as is the sense of *structural inevitability* belonging to it. We have the river of openness, the island of closure, and the way the former can only be recursively tracked (cognized) within the latter ‘downstream,’ such that, as Metzinger says, the island always seems to contain the river that contains it.

Once again, the notion that this structural correlation is some kind of coincidence or interpretative chimera seems hard to credit. The subsystematic nature of the RS entails openness, whereas recursive integration entails encapsulation, which is to say, closure. Taken together, we have the general structural template for the now as something that paradoxically contains the time that contains it. Add to this the concepts of sufficiency and asymptotic limits, and presto, we have a *naturalized account of the magical now*.

One need only consider anosognosia in clinical contexts to appreciate the very real experiential impact of *neuropsychological* information horizons. Anosognosia is most commonly seen in stroke victims suffering hemi-neglect, and perhaps most spectacularly in Anton-Babinski syndrome, where patients who are functionally blind continue to insist they can see. Though the etiology behind the specific sensory or locomotor deficits is often well understood, the attendant cognitive deficits remain deeply mysterious,¹⁷ both because of their complexity and the way they “relate to the difficult problem of the brain and mind” (Celesia and Brigell, 2005). Anosognosia not only straddles the vexed relation between phenomenal experience and cognition, it demonstrates the dramatic sensitivity of both modalities to information horizons.

In a sense, BBT argues that the peculiarities of consciousness *as it is experienced*, those things that make it so difficult to understand, are the expression of various *natural* anosognosias. Now is *now* because of constraints on recursive availability. We literally suffer a temporal version of Anton-Babinski syndrome: we are utterly blind to the time of our timing, yet sufficiency convinces us that we see all there is to see. What do we see? *Identity*—or so we think.

And thus we have the paradox upon which entire philosophies have been founded.

The ‘Side-Steal’: Personal Identity

¹⁶ As this instance suggests, I think a tremendous amount of philosophical speculation admits substantive reinterpretation in BBT terms, suggesting that philosophy—or a good portion of it—can be understood as the human brain’s millennial attempt to see itself for what it is.

¹⁷ See Prigatano (2010) for an excellent overview.

Insofar as the ‘now’ seems coextensive with the ‘self,’ we can assume that the same faulty default intuition that renders the now paradoxical also informs our philosophically problematic intuition of ‘person essentialism.’ We are always the ‘same’ simply because we lack the information required to globally self-differentiate.

As we saw above, BBT implies what might be called asymptotic or *marginal identity*. This, one might argue, is what led Hume to declare the problem of personal identity to be merely verbal (1888, p. 262), and Kant to devise his transcendental solution (1929, p. 153).¹⁸ As structural, the identity at issue never appears within experience, and so can only be secured transcendently, or written off altogether when considered in reflection.

The complexities of the ‘self,’ one might presume, are simply built up from marginal identity, with certain elements, perhaps such as willing and agency, evolved and wired in, and others the result of cultural inculcation. The devil is in the details, of course, but it is a testament to the parsimony of BBT that selfhood simply falls out of the explanation of the now. Once the now is revealed as a *structural* expression of encapsulation, which is to say, *nothing in particular*, as opposed to the arcane accomplishment of some arcane neural device that arose for some arcane evolutionary reason, then it becomes something that other apparent features of consciousness can ‘service’—like selfhood.

We know that consciousness is informatically localized, which means we know that the system responsible possesses information horizons—or is encapsulated. It’s safe to assume the limits on recursive availability have to be expressed in consciousness *somehow*. Asymptotic sufficiency, as dramatically evidenced by our visual field, constitutes at least one way these horizons find phenomenal expression. Generalizing from this to the now and the self seems to possess ample abductive warrant, especially when one considers the (apparently insuperable) difficulty of the phenomena otherwise, not to mention the parsimony of the resulting explanations.

The real enigma is what I’ve called asymptotic or marginal identity. Why should the RS (and those cognitive systems mediated by the RS) transform the absence of information into an ‘intuition/assumption of identity’? This question, I suspect, will only be answered when science answers the question of why the *RS results in consciousness at all*. BBT, once again, is a theory about why consciousness appears the way it does, nothing more. Regarding the question of consciousness proper, it promises to clear away many of the explanatory burdens that make the so-called Hard Problem seem so

¹⁸ From the BBT perspective, the ‘transcendental’ is best understood as conscious cognition’s attempt to make sense of the asymptotic structure of its encapsulation—or the open-closed structure of the RS. Much of traditional philosophy, I think, can be systematically interpreted in Blind Brain terms, allowing for a ‘mass naturalization’ of human speculative discourse.

To take Heidegger as an example: his interpretative gestalt clearly turns on a certain interpretation of sufficiency, one where the structural elements of *Dasein*, the asymptotic frame of global sufficiency, are themselves conceptualized in asymptotic terms, such that *Dasein becomes* its various modes of Being, rather than that which employs or contains them. His famous ‘ontological difference,’ the refusal to cognize Being in terms of beings can be parsed in similar terms: to consider things in the register of beings, or ontically, is to interpret them as informatically embedded, which is to say, nonasymptotically structured. To consider things in the register of Being, on the other hand, is to interpret them as not informatically embedded, ‘holistically,’ which is to say, *as asymptotically structured*. His critique of the ‘Metaphysics of Presence’ becomes, accordingly, a critique of the tradition’s failure to recognize the profound structural role played by asymptotic sufficiency.

For his part, Derrida (to consider another apparently ‘incorrigibly intentional’ thinker) made the *illusion* of sufficiency central to his critique of the Metaphysics of Presence. What he calls *differance* can be seen as an interpretative conceptualization of the way the openness of the RS continually undermines the illusion of sufficiency generated by encapsulation, or put differently, the closure imposed by recursive availability. The finite capacity of the RS means that its horizons of access continually reassert themselves, no matter what the incoming information. The ‘complete picture’ is perpetually deferred.

hard. This is perhaps its primary *empirical* contribution: the way it drastically simplifies what neuroscience needs to explain.

One magic trick is far, far easier to solve than many.

The Neural Hand Is Faster Than the Neural Eye: the Unity of Consciousness

Anosognosia in its clinical sense refers to a reduction in neural information access that cannot be consciously perceived. You could say it's the name we give to *pathological* encapsulation. But what if we reversed the situation? What if we could *increase* the information access of the conscious brain? How is 'more information' typically expressed in consciousness?

One might assume more resolution and complexity.

Imagine an injection of trillions of nanobots into your brain, some of them tasked with extending the information horizons that underwrite the asymptotic limits of consciousness, others tasked with providing the resources required to process all the new information that becomes available. As your consciousness transforms, what could you expect?

The simple answer would seem to be, more resolution, more complexity. As nanowires infiltrate every neural nook and cranny of your skull and your recursive processing power expands into the exaflop range, you should experience an embarrassment of informatic riches. Pain would complicate as our access branched out to map its myriad sources. Vision would become less and less 'transparent' as more and more information about the primary visual cortex and elsewhere became available; we would develop entirely new modalities to 'see' all the machinery of perception. Everything that is now murky and etiologically sourceless would become ever more crystalline and etiologically detailed. All the informatic chasms concealed by modal sufficiency would be pried open, and depths of informatic detail that we cannot even imagine would be revealed...

And we would finally see that what we presently call 'consciousness' was in some sense a profoundly deceptive *cartoon*.¹⁹

Traditionally, the question has been, 'What device could turn 100 billion externally-related neurons into a singular, internally related whole?' which is to say, a question of *accomplishment*. And yet, given that discontinuity requires discrimination, and discrimination requires information, the fact that consciousness appears unified, something internally related, immediately speaks to the *lack* of information. Expressed in these terms, the 'problem of unity,' from the accomplishment perspective, is the problem of manufacturing the *lack* of a certain kind of information.²⁰

¹⁹ It's worth noting, however, that *no increase* in access and processing power would allow an enhanced RS to outrun the limits of recursive availability, which means that the persistent global sufficiency and the asymptotic absolute would still obtain in some attenuated sense (an enhanced RS, one might presume, would not be compelled to 'believe' the illusions the way we seem to be). Part of the problem has to do with what might be called 'process asymmetry,' the way information processors always *embody* more information than they can process at a given moment. If, for instance, you think of the human RS as a 'second brain' evolved to track and coordinate the operations of the original, you can see that as the scope of recursive access increases, the overall proportion of recursively available information decreases. More tracking means more processors, means a net increase in untracked processes. The development of human consciousness literally entailed the growth, not the retreat, of the 'unconscious.'

²⁰ Given information asymmetry, the abyssal gap between the brain and consciousness, it is actually *extraordinary* that so few theorists have considered this aspect of the problem. Pöppel (2009) is a rare exception: "Looking at the complexity of the neuronal representation of information, the easy and effortless availability of the basic repertory of conscious phenomena is rather enigmatic. Apparently, the nervous system has developed strategies to overcome inherent problems of complexity" (p. 1889).

As I argued above, the constraints of evolutionary contingency, frame complexity, and access invariance suggest that consciousness has to suffer any number of ‘resolution deficits.’ When one thinks of our fuzzy phenomenal and cognitive grasp on norms or affects, say, this deficit seems fairly clear. But why should a sense of global, internally-related, conscious unity arise out of an inability to self-discriminate?

Encapsulation entails sufficiency: the RS accesses comparatively little second-order information pertaining to the information it accesses: as a result, differentiations are ‘skin deep.’ The various modalities are collapsed into what seems to be an internally related whole.²¹ Perhaps our sense of externally-related multiplicities is environmentally derived, a learned and evolved product of our ability to ‘wade into’ the (externally-related) multiplicities that comprise our proximate world. Consider the *distal* world, the long intellectual labour required to see the stars as externally-related multiplicities. *Access invariance*, along with apparent relational invariance between stars, convinced a good number of our ancestors that the stars were anything but discrete, disconnected objects hanging in empty space. Much the same could be said of the conscious brain. Restricted to invariant ‘channels,’ unable to wade into and interact with the vast informatic cosmos of the greater brain, it quite simply has no access to the information it needs to discern its myriad discontinuities. External-relations are flattened into internal relations; the boggling modular, let alone neural, multiplicity is blurred into the holistic haze of phenomenal experience, and we are stranded with a profound sense of unity, one that corresponds to nothing more than the contingent integration of information in our conscious brain.

Twirling batons blur into wheels. Numerous shouts melt into singular roars. Or as Bacon writes of ignorance: “all colours will agree in the dark” (1985, p.69). Experience is filled with instances of what might be called ‘default unity,’ events drained of complexity for the simple want of resolution—information. You could say reflecting on consciousness is like watching a nation-spanning mob from high-earth orbit: the simple inability to discriminate leaves us ‘assuming,’ ‘feeling,’ a unitary consciousness we quite literally don’t have.

I appreciate how naive or even preposterous this must sound *prima facie*. The thing to recall is that we are talking about the *way consciousness appears*. The question really is quite simple: What information is available to the RS? Information regarding the neural sourcing of available information? Of course not. Information regarding the external-related causal complexities that deliver and process available information? Of course not. Given that the RS is responsible for consciousness, it stands to reason that consciousness will be blind to its neural sourcing and biomechanical complexity. That what it intuits will be a kind of *compression heuristic*, an informatically parochial and impoverished making due. And this brings us to the tantalizing edge of conscious unity: “There is a great difference between mind and body,” Descartes writes, “in that the body, by its nature, is always divisible and that the mind is entirely indivisible” (1968, p. 164). Is it merely a coincidence that this ‘great difference,’ *even as Descartes conceives it*, happens to be informatic?

Once again, the real mystery is why the RS should turn the absence of information into the intuition or assumption of ‘default identity.’ It’s important to realize, however, that the only thing new about this particular mystery is the context BBT provides it. Researchers in psychophysics, for instance, presume ‘default identity’ all the time. Consider the way visual, auditory, and tactile processing ‘blurs’ stimuli together when the information presented exceeds certain thresholds. Below certain intervals, what are in fact multiple stimuli are perceived as singular. One might think of this in terms of filters, beginning with our sensory apparatus and ending at the RS, where the capacities of given neural systems ‘compress’ differences into identities and so strip incoming information of ‘functional redundancies,’ which is to say, information not required for effective action.

²¹ The raises the interesting possibility that the ‘binding problem,’ the question of how the brain coordinates its activities in response to a cohesive world given the variable speeds with which different perceptual information is processed, is altogether separate from the question of the unity of consciousness.

So here's the question: Why should the default identity implied by, say, flicker fusion experiments in psychophysics be any less mysterious than the default identity implied by the above account of conscious unity? Or put differently: Why should the question of conscious, multimodal fusion be detachable from consciousness in a way that the question of conscious unity is not? In the present context, both assume some yet to be determined scientific account of consciousness. Both interpret the resulting identity in terms of constraints on capacity. The former, of course, is amenable to the rigours of psychophysics in a manner that the latter is not, but this only speaks to technical limits on experimentation. The only real difference, I would argue, is one of *scale*. Fusion or coincidence thresholds demonstrate the way informatic capacity constraints find phenomenal expression as default identity. Conscious unity, BBT suggests, is simply a global example of this selfsame effect, 'fusion writ large,' given the limits on recursive availability confronting the RS. Like the now and personal identity, the misapprehension of unity is simply a structural side-effect of being conscious in the first place.

Much of BBT's significance, I think, lies in its ability to parse the question of consciousness from the kinds of questions considered here. By 'explaining away' the perplexing structural features belonging to the first-person, it tremendously simplifies the problem that consciousness poses for cognitive neuroscience. My strategy so far has been to peel away the apparent mysteries, one by one, revealing, at each turn, a way to interpret some perplexing feature *out* of our understanding of the first-person perspective.

What could make the experience of unity come into existence? BBT answers, *nothing*. Conscious unity is every bit as illusory as mistaking an incessant flicker for single abiding light.

'Pinky Break' Semantics: Aboutness and Normativity

So what is 'brain blindness'? Encapsulation, the informatic localization of the RS within the greater brain.

As we have seen, encapsulation means conscious awareness will exhibit what might be called persistent global sufficiency, simply because information pertaining to its informatic poverty and neurological provenance is perpetually unavailable. It's the only frame of reference it's got. This means information that finds its way to conscious awareness *will seem ordinary and complete* no matter how wildly sourced, how parochial or impoverished it is in fact. This doesn't mean we never experience insufficiency—we do all the time—only that we become aware of insufficiency within a frame of implicit sufficiency. To cognize information as incomplete is to possess information to the effect that something is missing. As Dennett writes, "The absence of representation is not the same as the representation of absence" (1991, p. 359). The limits of conscious awareness only come to conscious awareness within conscious awareness.

This is the reason so many people find certain psychological and neuroscientific findings so jarring. Whether it's the distortions intrinsic to human reasoning, the dissociation of suffering and pain in pain asymbolia, or the anticipation of conscious choices using neuro-imaging data, the initial response for most people is outright incredulity.

Once again, the open-closed structure of the RS offers an elegant solution. The above cases (and many others), confront us with information of what lies beyond the pale of recursive availability. Given evolutionary contingency, frame complexity, and access invariance, we can presume that the RS has only parochial, impoverished, and inchoate access to what the brain is *actually* doing. Given encapsulation, however, the RS has no way of cognizing (or mediating the cognition of) itself *as such*. As we have seen, the closure of recursive availability (encapsulation) entails sufficiency, which is to say, the precise opposite of what so many experimental findings are beginning to reveal. The 'back stage' information provided by the myriad examples of dissociation, neglect, bias and so on reveals sufficiency for what it is: a kind of neuropsychological illusion.

Encapsulation, in other words, suggests that consciousness is ‘like a magic trick’ to the extent it exhibits sufficiency. Unlike an environmentally staged magic show, however, the levers of incredulity find themselves reversed. Given that sufficiency is a magic show that we’re born into, it’s the glimpse over the magician’s shoulder that becomes difficult to believe.²²

Asymptotic sufficiency is an illusion. Likewise, presence, unity, and selfhood are informatic magic tricks. The corollary, of course, is that intentionality is *also* a kind of illusion.²³ Consciousness seems intentional *from the standpoint of consciousness*. This just means that intentionality is another informatic artifact of encapsulation, another ‘keyhole glimpse.’ The reason, then, that functional explanations seem to ‘explain away’ intentional phenomena would be because, like magic, they *do*. By looking over the brain’s shoulder, they ‘break the spell of encapsulation,’ introducing etiological and structural information into an informatic economy adapted to the lack of that etiological and structural information.

The Magic of Intentionality

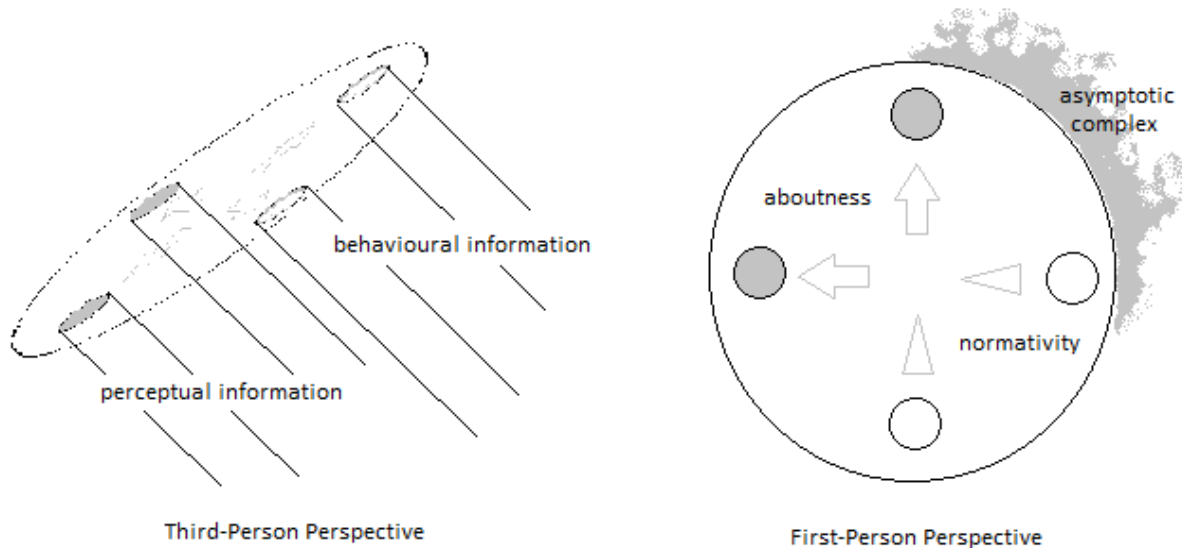


Fig. 6 Given encapsulation, the RS has to ‘make due’ with what little information it has. Both normativity and aboutness, on the BBT account, are signature examples of such ‘making due,’ ways that the brain is forced to cognize informatic relationships it cannot see. Here they are depicted as ‘orthogonal compression heuristics,’ acausal imperatives that apparently constrain acausal relationships to the world—things that appear, thanks to sufficiency and asymptotic limitation, as if ‘from nowhere.’

The third-person perspective we take when we look over the brain’s shoulder reveals the causal complexities that remain ‘back-stage’ from the first-person perspective. Aboutness and normativity, on this account, are artifacts of the way the RS makes due in the absence of this causal information (or worse

²² For my part, I know that a patina of disbelief colours everything I have written so far. I have grown accustomed to the *idea* that the now is an illusory artifact of marginal identity, that the unity of my mind is an artifact of informatic poverty, and that what I call ‘me’ is a kind of social construct raised on these delusional foundations. But do I believe it? Ask me tomorrow.

²³ Note that the issue of the separatism or inseparatism of intentionality and phenomenality is moot in this context, given that it is *informatic access* that is at issue.

yet, of our self-conscious reflection on consciousness). Stranded with fractional information pertaining to the greater brain's behavioural processing, the RS bootstraps an informatically depleted mode of behavioural processing, an asymptotic economy, which we intuit/cognize as normativity. Complex biomechanical processes become ineffable and bottomless 'rules.' Likewise, stranded with fractional information regarding the greater brain's perceptual processing, the RS adopts a depleted mode of environmental processing, another asymptotic economy, which we intuit/cognize as aboutness. Complex causal histories vanish into a hanging sense of *being aimed at*.

Sufficiency means that both modes, despite their murkiness, will seem autonomous and complete, something thoroughly *acausal*. Combined with temporal marginal identity, the illusion of the now, they will seem to dwell entirely outside the natural order. As of course *they do*: aboutness and normativity are little more than second-order chimera, assumptive posits required to cognize what little information conscious reflection has to go on. Thus the endless disputation regarding their nature: as 'compression heuristics' adapted to the requirements of a *parochial* informatic economy,²⁴ one might expect 'processing conflicts,' if and when the brain attempts to harmonize them with compression heuristics adapted to the requirements of *the world at large*, which is to say, its natural environments.

And this would be why they seem to vanish like magic when we look over the brain's shoulder. Doing so inserts them into an informatic economy that quite simply renders them unintelligible. In a sense, they literally are a kind of *inverted* 'coin-trick in the head': like the coin, they seemingly 'come from nowhere.' But where the coin's appearance counts as a cognitive violation, their 'appearance' provides the cognitive baseline. We find the notion of 'seeing through' them absurd or impossible, not only because they are all that we can see, but because we think them embedded in the very architecture of seeing.

From Sleeve to Profonde: The Positioning Problem

The biggest problem with the magic show analogy, however, is the glaring fact that 'magic' *doesn't amount to anything*, whereas many of the intentional phenomena that furnish the first-person perspective are the very engines of cognition—and (apparently) obviously so. This is why the suggestion that intentional phenomena are informatic fictions *tout court* is certain to provoke incredulity—after all, *mathematics and logic* are intentional!

But as I suggested at the beginning of the show, information asymmetry suggests our *conscious* cognitive situation is far more complicated than many have hitherto realized. The bare fact of radical information asymmetry raises the question of whether the information available to consciousness can be deemed 'reliable.' The constraints posed by evolutionary contingency, frame complexity, and access invariance suggest that 'Probably not,' is the most conservative answer. Nor can we simply presume that we can infer, let alone intuit, the evolutionary utility of conscious cognition. We certainly assume that the *cognitive features* of consciousness, no matter how fragmental they turn out to be, must be strategic or operative in some sense. This may be true, but as the case with volition demonstrated, 'strategic or operative in some sense' does not mean strategic or operative in some sense that we can *readily intuit*. Everything considered so far underscores what I earlier called the *Positioning Problem*, the question of how the conscious brain is informatically situated relative to the greater brain.

The RS is a *subsystem* of the brain. As such, the question of its actual functions is one that only neuroscience will decide. There can be little doubt that more than a few surprises await; the question is one of how troubling they will be. What I hope to show is how our sense of conscious cognition, while

²⁴ The question is one of just *how* parochial that economy is. It could be the case that 'aboutness' and 'normativity' are not the phenomenal expressions of the compression heuristics used by the RS, but are instead merely artifacts of *self-conscious reflection*, 'a peep-hole glimpse of a peephole glimpse.'

appearing to be the engine that makes everything work, could turn out to be every bit as distorted as our sense of volition seems to be. What I hope to show, in other words, is how philosophers—many of whom are accustomed to questioning the reality of the now, personhood, and conscious unity—could be also hoodwinked by the ‘manifest image’ of cognition as well.

The RS does not access any substantive information pertaining to its functional role vis a vis the greater brain. How do we know this? Because we seem to have no sense of consciousness’s neurofunctional role *period*. The RS qua open is simply a component of a much larger mechanism, while the RS qua closed is almost utterly blind to the greater mechanism, let alone its details. Encapsulation means sufficiency. Consciousness, you might say, is a cog that congenitally confuses itself for the whole machine.

And this raises the spectre of what might be called ‘metonymicry,’ the way sufficiency, combined with the subsystematic nature of the RS, generates the illusion of *cognitive efficacy*.

Sufficiency, Systematicity, and 'Metonymicry'

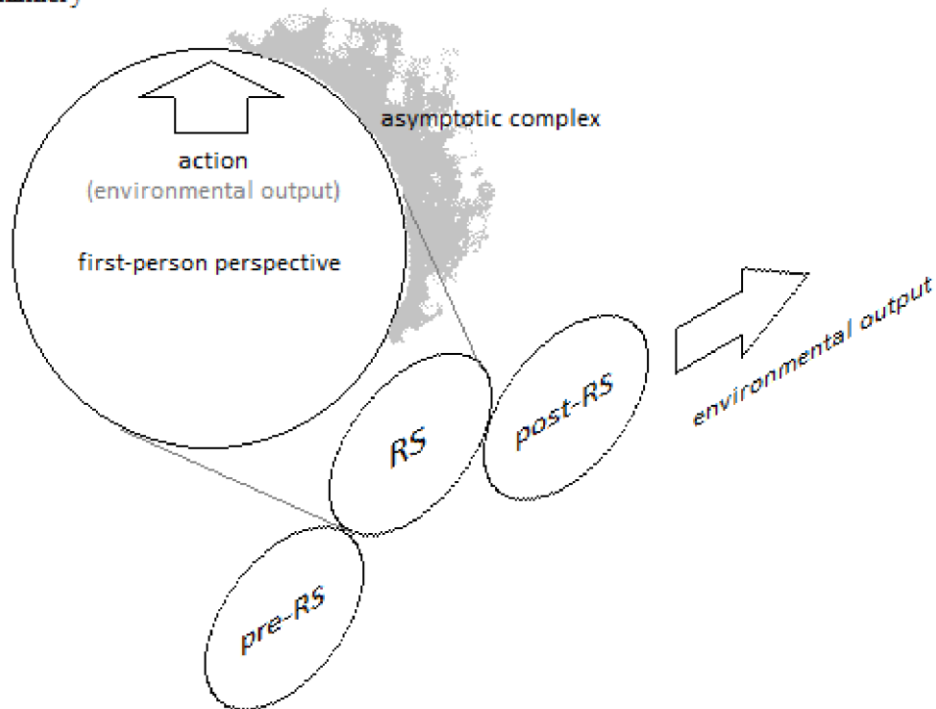


Fig. 7 Given sufficiency and systematic interdependency, the RS is bound to seem like ‘the *only* wheel that turns,’ absent any information to the contrary. This means that consciousness can in fact be massively deceived *about* cognition without substantially compromising the efficacy of cognition.

Of course the RS is nowhere near so integral to the brain’s operations as Fig. 7 implies. The systems it taps for information and the systems it feeds are entangled with each other independently and with the RS, such that each may be ‘pre’ or ‘post’ depending on the neural function. But no matter how complicated the actual system, the upshot remains the same: sufficiency means that *consciousness confuses itself for the whole*. The machinery drops out of the picture, and with it, any comprehensive information pertaining to the actual functions served by the RS. This means that the RS, when brokering cognitive appraisals of its own efficacy, is bound to seem central no matter how marginal its role in fact.

Conscious cognition, in other words, only has cognition *as it appears to consciousness* to correlate with its actions. And since cognition as it appears to consciousness, no matter how fragmentary

in fact, remains systematically related to the efficacious processing economies of the greater brain, *it will seem efficacious*. In fact, given that consciousness has access to *nothing else*, it will seem to be the *very definition* of comprehension, understanding, cognition, or what have you. Thus metonymy: the possibility that we could be thoroughly deceived about cognitive processes which are nevertheless entirely effective. Logic, for instance, could be both a kind of informatic fiction *and* the intuitive baseline of ‘a priori comprehension,’ simply because of its systematic relation to the neural whole.

This is worth a deep breath. Could *logic* be a species of blind brain magic?

A number of independent considerations, I think, suggest that we at least entertain the possibility. For one, our formal semantic intuitions, like so many other intentional phenomena, stubbornly resist naturalization. The thesis proposed here: that intentional phenomena, as the products of a parochial, parasitic informatic economy (the RS), cannot be adapted to the informatic economy of the greater world is as cogent an explanation as any. For another, we have good reason, thanks to Church and Turing, to suspect that computation exhausts calculation. If interaction patterns can capture formal inference structures without any loss *save the semantic*, then a principled means of explaining away formal semantics could very well count as a theoretical gain. Rather than splitting hairs and chasing tails trying to explain mathematical and logical cognition in terms amenable to our *conscious awareness* of mathematical and logical cognition, we could get down to the work of understanding them for what they really are (as perhaps far larger processes merely brokered in various consciously inaccessible ways by our conscious awareness of them). And lastly, there is the way most metaphysicians are prone to characterize the cornerstone of formal semantics, abstraction, *in terms of information loss*.²⁵

In other words, we have a set of concepts and procedures that defy causal explanation. Nevertheless, we generally agree: 1) that causal processes are capable of *doing* everything those concepts and procedures do; 2) that information loss is a signature characteristic of those concepts and procedures; and 3) that the conscious brain views the ‘calculating’ brain through an informatic keyhole.

Is there a real linkage here, or is it simply a coincidence?

We know that our conceptual and intuitive *understanding* of logic as an abstract and inferential economy *has* to suffer informatic deficits of *some* kind. Our shadow likely contains more visual information relative to our appearance than consciousness does relative to our brain! Is it so far-fetched to suggest that what we ‘know’ as calculation is little more than the informatic shadow of the concrete complexities of neural implementation? Could formal semantics be an illusory side effect of a certain encapsulated ‘angle’ on non-semantic neural computation?²⁶

This is where we crash into the wall of the a priori. Suggesting that formal semantics, as we conceive it, is a distorted side-effect of encapsulation is tantamount to psychologism, the apparently question-begging reduction of the logical to the merely empirical. What is true a priori is true by *definition*, not by virtue of information adduced a posteriori. The truth of the conclusion (or theorem) is contained by the truth of the premises (or axioms). The belief—what we might call the ‘semantic totality assumption’—is that everything is already somehow ‘there’ awaiting inferential articulation. Formal

²⁵ Lewis (1986), for instance, describes the ‘way of negation,’ wherein abstract objects are defined in terms of those informatic dimensions (causal, temporal, spatial) they lack.

²⁶ This could, among other things, explain the symbol grounding problem. It’s worth noting that Floridi’s (2011) recent attempt to solve the problem involves the partition of the ‘Artificial Agent’ into *two* processing systems, one directly keyed to the environment, the other directly keyed to this first *in a limited way*. When the first engages in an environmental interaction, the second ‘transduces’ the resulting state at a lower level of ‘granularity.’ Thus it is the *loss of information*, a limit on second-order availability, that allows the symbol attached to the environmental interaction by the second to count as ‘generic’ and so (apparently) emulate the first crucial phase of ‘semanticization.’ Identifying the granularity of the symbol-state with ‘generality’ allows Floridi to attribute ‘meaning’ to the symbol. What he fails to realize, however, is how his story could just as easily be the beginning of an *explaining away* as an explanation of semantics. Otherwise, the lack of any corresponding account of normativity makes it difficult to understand how what he calls ‘meaning’ is actually anything more than ‘granular shape pairing.’

semantic systems, in other words, are autonomous insofar as they supposedly make preexisting information *explicit*.

But this, of course, raises the puzzling question of how formal semantic cognition can be *informative at all*—the so-called ‘scandal of deduction’ (Hintikka, 1973, p. 222). If the conclusions (or theorems) are contained in the premises (or axioms), why do we require inferential articulation? On the BBT account, the RS is a broker, continually mediating far larger informatic transactions, discharging functions that only a mature neuroscience can definitively determine. As such, it constitutes (given the limits of capacity and recursive availability) a neuro-informatic ‘choke point.’ And this, we might speculate, is why we require inferential articulation. In fact, compositionality only makes sense given some kind of ‘cognitive bottleneck assumption.’ Why would we evolve a way to ‘sum truth-values’ if that sum was always available at the outset?

So does this mean the semantic totality of information is already there, only ‘somewhere else,’ requiring (given the bottleneck of conscious cognition) the inferential schemes of mathematics and logic to ‘unpack’ in explicit formulations?

For the sake of argument, imagine that what we call the ‘formal sciences’ of logic and mathematics are simply domains of the *natural sciences* of information. The thing that seems to so profoundly distinguish them from a posteriori sciences like physics is the fact that the brain *is itself an information system*. The human brain is literally its own ‘laboratory’ in this respect. Logic and mathematics are those regions of the natural that we can systematically explore via cognitive *performance*, slowly charting the possibilities of *information interaction patterns* in our brains. And this would be why formal semantics are in fact so tremendously informative: the patterns performatively intuited (observed) are a species of *natural fact*. As an exploration of *how* various interaction patterns organize and select varieties of information, it would seem to have unlimited iterative applicability, to be *prior* to any environmental information.²⁷ As an exploration of the very laws of information processing that constrain the brain, it would seem *procedurally compulsory*, the way information *must* be articulated to maximize environmental applicability.²⁸ It would also appear, given the various consequences of encapsulation we have considered thus far: 1) sufficient, which is to say, asymptotically bounded; 2) nonspatial, atemporal, and acausal, simply because no information regarding the neural processing of that information would be available; 3) internally-related as a result of this loss; and 4) depleted, because of granularity.

In other words, it would accord quite well with the ‘autonomous, preexperiential domain of timeless inferential possibilities’ that generally characterizes our present ‘understanding’ of the semantic a priori...

It seems difficult to believe this is mere coincidence.

When we take the ‘environmental perspective’ to the question of semantics, the difficulty becomes one of understanding how the *apparently* self-evident inferential structure of meaning ‘plugs into’ the kinds of interaction patterns we find in the brain. Strangely enough, the relation between the environment and the brain is far more ‘friendly’ than that between the RS and the greater brain. Where the latter is constrained by evolutionary youth, the former has a pedigree older than the vertebrates. Where the neural environment accessed and tracked by the latter is the most complicated known, the

²⁷ It’s worth noting the way this limitation explains the importance of logical and mathematical notation. It’s the *very inscrutability* of these performances that makes fixing them in perception so instrumental: tracking transformations in visual perception allows the RS to *broker* various forms of training, testing, and experimentation that the limits on recursive availability (encapsulation) would not permit otherwise.

²⁸ The fact that we embody these laws certainly speaks to their fundamental nature. Perhaps we are nature’s most forward fractal edge, a slow evolving informatic replication of a universe that endlessly recapitulates itself in fragments across myriad scales. No norms. No meanings. Only twists that pinch us from our own conditions, and so conjure these proxies (which then become the source of endless second-order controversy).

natural environment accessed and tracked by the greater brain is far more manageable. And where the RS is ‘hardwired’ to its neural environment, the brain can sample its environment from any number of locations. Given these advantages, the environmental perspective allows for a far richer informatic picture than that provided by the first person, one that possesses higher resolution and includes the occluded dimensions of space, time, and causality. Given these *inevitable* privations, the acausal, low resolution, atemporal nature of what appear to be ‘inference patterns’ from the first person perspective become more than a little suspect. The perennial difficulty we have plugging the semantic into the natural appears to be a perspectival conceit not unlike that informing the Ptolemaic conception of the universe, one where the *absence* of information leads to the assumption of ‘default identity.’ Logic and mathematics seem to be the ‘immovable ground’ of cognition, rather than the exploration of the informatic possibilities of a brain too close and too complex to be truly known—short of the environmental perspective.²⁹

²⁹ On the question of physicalism more generally, BBT only postulates systematic differences making systematic differences. So in Jackson’s Mary argument, the whole question hinges on whether ‘seeing red’ counts as *knowledge*, which is to say, a certain kind of *semantic* information. Since it seems clear that Mary ‘knows more’ upon seeing red even though she already ‘knew everything there is to know’ about the brain, we seem forced to conclude there are nonphysical facts about the brain. Once, however, we realize the apparent aboutness of information is a heuristic artifact of the way the brain is blind to the ‘fromness’ of information, the issue entirely dissolves. In terms of information from, Mary first receives information via her third-person environmental perspective, and then via her first person experiential perspective. She receives information that is environmentally embedded, which is to say, *from* the RS as open (or over the magician’s shoulder), and she subsequently receives information that is asymptotic, which is to say, from the RS as closed. Since both perspectives are simply different attentional modes of the *same* perspective, the embedded information *also arrives asymptotically*, which is to say, as blind to its fromness. It becomes information *about* the RS. The asymptotic information (redness), however, is obviously not ‘about’ the RS at all—and so the question becomes one of what ‘fact about the brain’ Mary could possibly be learning.

When cognition ‘flips’ from examining environmentally versus neurally sourced information, it simply crashes against its own asymptotic limits. Since ‘what-is-it-like’ information is asymptotic, it cannot be environmentally ‘embedded,’ and so seems to comprise an entirely unique realm of *nonenvironmental* ‘fact.’ But since Mary’s ‘what-is-it-like’ information is, like her environmental information, *from the brain*, it seems that it should be environmental after all. About is pitted against From, and each side games the ambiguities accordingly. From the BBT perspective, what the Mary Argument really calls into question is the asymptotic structure of what we call *knowledge*. Intentionality is simply our way of coping with the asymptotic consequences of encapsulation. Given the systematic dependency of the RS on the greater brain, it works quite well, *so long we don’t ask questions*. Otherwise we are stranded with square informatic pegs in a world of round holes.

The Informatic Reduction of Neural Interaction to Inference

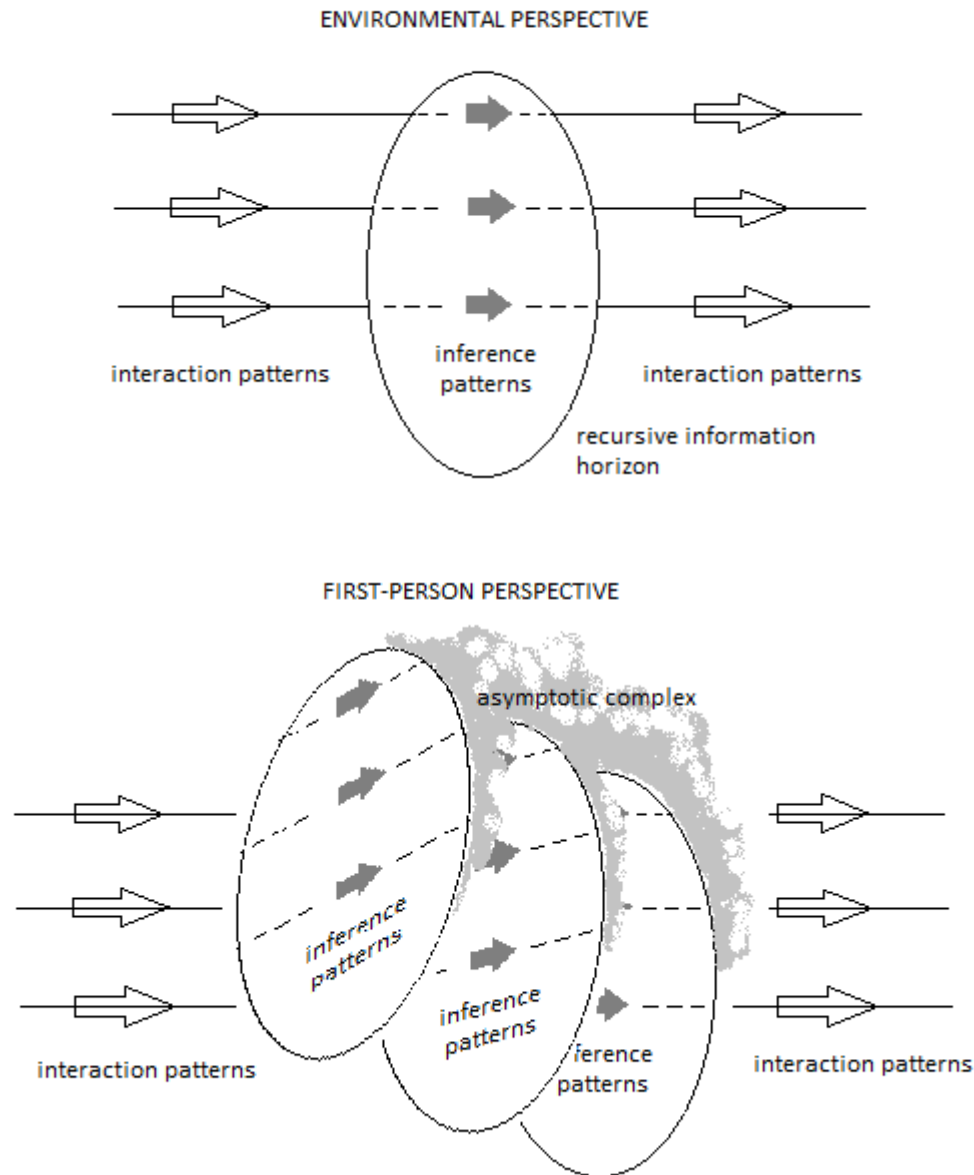


Fig. 8 From the environmental perspective, inference patterns appear anomalous, whereas from the first-person perspective interaction patterns do not appear at all. By showing how the reason for the former lies in the fact of the latter, BBT provides a way of seeing inference patterns as a kind of ‘informatic reduction,’ an artifact of just how little interaction pattern related information is available to the RS.

Absent any definitive demonstration of ‘semantic uniqueness,’ the kind of ‘informatic reduction account’ given here simply has to be taken seriously. As all the zombie and swampman blind alleys seem to attest, the problem isn’t one of accounting for *what* we do, so much as *how we think we do it*.³⁰ We

³⁰ The zombie problem wonderfully illustrates the open-closed structure of the RS and the dilemma this poses for

quite simply do not need meaning to account for our behaviours—only our apparent comprehension of them. BBT offers a *relentlessly* environmental perspective. It paints a portrait where everything is *from*, and absolutely nothing is ‘of’ or ‘about,’ one where information simply bounces around according to natural laws. Since everything is concrete, the functions we assign to ‘abstractions’ are performed by informatic devices possessing regimented applications—vector transformation versions, perhaps, of the computational strategies we already use. We just happen to have information in our brain that is systematically related to information in our environment by virtue of *past* environmental interactions. All ‘cognitive’ relations, BBT suggests, are relentlessly causal, relentlessly diachronic. These patterns are no more ‘true of’ the world than our genome is ‘true of’ our body’s protein matrix. Brains are viruses, bioprograms that hack their organic environments. They either reliably effectuate certain processes possessing homeostatic consequences or they do not.

The fact that it doesn’t ‘seem this way’ simply speaks to the drastic limits placed on recursive availability—limits that we should expect for all the reasons enumerated above. Truth and the inference structures that conserve it, BBT suggests, are artifacts of encapsulation, orthogonal compression heuristics. The brain, among other things, is an environmental information extraction machine. The deep causal histories of extraction—what *actually* drives our relation to the world—are lost, leaving only the information *that* this information is environmental ‘somehow.’

Since the diachronic nature of our environmental relation is only available *within* experience, marginal temporal identity (the now) becomes the apparent frame—our frame. Consciousness as it appears is bootstrapped by identity. Aboutness becomes temporally ambivalent, a halo belonging to all things at all times. Language (what little we experience of it) becomes something apparently possessing parallel *synchronic* relations to everything. ‘Thoughts’ become things *not entirely of this world*, things connected to their environments, not via causal histories, but *logically*, in virtue of their ‘meaning.’

And so a world supposedly purged of ghosts finds itself haunted by other (and in this case, philosopher friendly) means. On this account, what we experience as formal semantics is a tissue of truncations, depletions, and elisions, all of which are rendered invisible by sufficiency, the inevitable fact that the RS lacks any comprehensive information regarding its lack of information. Neuroselective feedback histories are congealed in the near featureless perplexity of ‘truth.’ The diachronic causal intricacies of neural interaction patterns are flattened into bootstrapped inferential cartoons. The RS, qua open, brokers this information in various yet-to-be-determined ways, while qua closed, it brokers the assumption that—far from a broker—it’s the *author*. Its systematic relationship to the actual processes of cognition assures that this authorial conceit will be rewarded by practice. So we find ourselves defending an intuition that is illusory through and through, postulating ‘special ontologies’ or ‘levels of description’ to rationalize the sufficiency of what, from the standpoint of BBT, is obviously an informatic phantom.

From a natural perspective, which claim is more extraordinary? That inferences somehow express something *supernatural*? Or that ‘inferences’ are the truncated, depleted expression of special kinds of neural interaction patterns in conscious? I appreciate how radical—perhaps even crazy!³¹—this must sound. But given the fact of information asymmetry, the likelihood of encapsulation, and the *possibility* of metonymy, how could we know otherwise? Perhaps we should look at mathematics and logic as we

our brains. You could say that BBT draws a distinction between *from* consciousness, the consciousness we actually have, and *about* consciousness, the consciousness this former consciousness *thinks* it has. Given the asymptotic structure of about-consciousness, it resists naturalization, and thus seems to stand ‘outside’ of nature somehow. The zombie provides an ideal vehicle for pressing this intuition, since it has the effect of underscoring the informatic parochialism and poverty that preclude the incorporation of about consciousness into environmental cognition. According to BBT, we are not quite zombies, nor are we quite the *humans* we take ourselves to be. It remains for science to count the degrees of our separation.

³¹ Eric Schwitzgebel (2012) has recently argued that any ‘metaphysics of mind’ must violate ‘common sense’ to the degree the latter is incoherent regarding metaphysical matters more generally.

presently understand them as the last bastions of the soul, a final conceit to be toppled by the natural sciences.

This brings us back to the radical stakes of the Positioning Problem and the question of how far can we trust consciousness to comprehend itself. Even though science has been busily lowering the bar on this trust, very many (myself included) are loathe to let go altogether. Mathematicians *do math*: what could be more obvious? BBT doesn't dispute this so much as it insists on drawing a distinction between what mathematicians in fact do and *what they are conscious of doing*. To say that what mathematicians think they are doing 'exhausts,' or 'captures the essence' (or for that matter 'wildly distorts') what they are actually doing is to answer the Positioning Problem.³² How far we trust consciousness to comprehend itself depends on how the systems responsible for consciousness are *positioned* within the informatic economy of the greater brain.³³

The point can be sharpened into a single (and, I think, quite destructive) question, which I will call the 'Just-So Cognition Query' (JSCQ):

Just-So Cognition Query: Given that conscious cognition, as a matter of empirical fact, only accesses a small fraction of the greater brain's cognitive activity, how do you know it provides the information needed to make accurate second order claims about human cognition?

This question, if you ponder it, is *very* hard to answer. You can't appeal to the sufficiency of experience without begging the question. More corrosively, you can't appeal to the sufficiency of *practice* without begging the question, because your fraction only needs to be systematically related to the brain's cognitive activity to generate the reliable appearance of sufficiency (the problem of metonymy). The Church-Turing thesis, meanwhile, seems to block any claims to calculative uniqueness.

I call it the *Just-So* Cognition Query because of the way it forces noocentric philosophers *to take an explicit empirical position*. Encapsulation, for them to press their second order accounts, has to be 'just-so,' which is to say, provide access to the happy information that renders their accounts, be they phenomenological or transcendental, 'nonmagical.' They are literally making a bet on sufficiency, that the sciences of the brain, despite consistently undermining experiential sufficiency in so many respects, will begin revealing philosopher-friendly exceptions in the near or distant future.

As much as I *want* this to be the winning bet, I fear the odds are long. Once you realize that our intuition of conscious cognitive sufficiency need not indicate anything, the noocentric wager becomes something of a sucker bet, especially when you consider evolutionary contingency ('right enough' is all evolution requires), frame complexity (there is so much to get wrong), and access invariance (what we see is all we get).

Even more troubling is the way BBT seems to explain our perennial inability to arbitrate our second-order speculative disputes regarding intentional phenomena: Encapsulation has stranded us with far less information than we need to adequately theorize what it is we are actually doing. Quite simply, the greater the loss and distortion suffered, the more difficulty the RS should have modeling (or mediating the greater brain's modeling of) its own functions. In other words, the very fact that semantics occasion so much philosophical controversy suggests metonymy...

Or in other words, the Positioning Problem.

³² This is not an argument against functionalism per se: BBT doesn't challenge the multiple realizability of consciousness, only the notion that *what we experience* is anything like a 'program'—something that could be said to 'run' the deeper levels of its implementation. By the same token, BBT doesn't endorse epiphenomenalism: to say conscious doesn't do the work we intuitively think it does is not to say that it is epiphenomenal, only difficult to fathom. What this all means is that the functionality of consciousness is not generally available to consciousness.

³³ In this sense, BBT possesses the time-honoured philosophical virtue of making hitherto invisible epistemic commitments plain.

Why is intentionality so difficult to understand? At what point do we stop blaming our concepts and start blaming the intentional phenomena themselves? One century? Ten? Why does the linkage between the intentional and the natural vex us at every turn? Is it merely a coincidence that causal anosognosia characterizes everything from volition to aboutness to abstraction?

That so much of consciousness possesses the structure of a magic trick?

What we call intentionality *could* be the product of a profound structural incapacity. The integrative crossroads evolution required of our ancestors *may* have stranded us with a profoundly deceptive understanding of ourselves. According to BBT, we are our brains *in such a way* that we cannot know we are our brains. As a consequence, we find ourselves perpetually trapped on the wrong side of the magician, condemned to what might be called the ‘First-person Perspective Show,’ and, like flies born in bottles, doomed to confuse our trap with the way things are. Only the slow and steady accumulation of ‘backstage information’ provided by the sciences has allowed us to conceive, if not countenance, the possibility considered here: that consciousness as we know it is a kind of trick that the brain plays on itself.³⁴

To say that consciousness as a product or bearer of *information* runs out where information runs out is platitudinal. To say that the limitations on information access must be expressed in consciousness *somehow* is also platitudinal. As we have seen, the real questions are, How does it run out? Where does it run out? and, How are these shortfalls expressed in consciousness?

Because they almost certainly are. The question is really only one of severity.

Finale

The Blind Brain Theory of the Appearance of Consciousness could be considered a ‘worst-case scenario.’ There is no such thing as now. There is no such thing as personal identity. There is no such thing as unity of consciousness. Each of these is what might be called a ‘recursive access illusion,’ a kind of magic inflicted upon us by encapsulation.

Recursive availability issues also plague our conscious awareness of cognition, as we should expect, given evolutionary contingency, frame complexity, and access invariance. Information is truncated, depleted, and elided, leaving what can only be called an informatic caricature, one characterized by the same low resolution and causal blindness that underwrite the illusions of nowness and conscious unity. Normativity, aboutness, reference, meaning, universality, abstraction, internal-relationality: all bear the hallmarks of various recursive information integration constraints. Sufficiency assures that conscious cognition will appear to capture ‘enough,’ and systematic dependency upon occluded cognitive processes assures that our practice will seem to confirm our sense of sufficiency. The insufficiencies of the information accessed only arise when we attempt to cognize conscious cognition, which is to say, when we begin to philosophize.³⁵

³⁴ Albeit, one that we *are*. The key here is to always remember that the present discourse, in fact, all conscious discourse, *is encapsulated*, and so necessarily trades in distortions and illusions. This means that ‘distortions’ and ‘illusions,’ as semantic categories, are themselves distortions and/or illusions. One way of viewing these and other intentional phenomena is as kinds of ‘algorithmic fragments,’ apparently effective because of their systematic relationship to inaccessible neuro-algorithmic wholes, but prone to break down when employed as algorithmic wholes in their own right. The language of distortion and illusion, here, is the language of algorithmic functionality. Referring to illusions in this context is no more self-contradictory than referring to ‘design’ in evolutionary contexts. The difficulty is that BBT takes the local problem of intentional terms in evolution, and makes it *global*.

³⁵ Does this explain why consciousness seems to constitute the most puzzling philosophical conundrum of all? Given that consciousness represents the concatenation of all informatic insufficiencies, you would expect it to be the most resistant to cognition.

The first-person perspective is an illusory bubble of identity in a torrential causal stream. Apparently here, apparently now, and apparently *aimed* all for the sake of causal anosognosia. You could say encapsulation is simply perspective naturalized. Recursive information economies are ‘open-closed systems,’ open the way all natural systems are open, and yet closed in terms of recursive availability. Consciousness as it appears is the structural expression of the empirical limits of this availability—the information horizons that constitute encapsulation. Not only does encapsulation allow a symptomatic, naturalistic reading of the philosophical problems pertaining to intentionality, it also explains how a natural information processing system can be transformed into a First-person Perspective Show—the very thing baffling you at this very moment.

If you are anything like me, you find this thesis profoundly repellent. BBT paints a picture that is utterly antithetical to our intuitions, a *cartoon* consciousness, one that appears as deep as deep and as wide as wide for the simple lack of any information otherwise; and yet a picture that we might nonetheless *expect*, given the RS and its myriad structural and developmental infelicities. A chasm of some kind has to lie between consciousness as *possessed* and consciousness as *experienced*. Given the human brain’s mad complexity and human consciousness’s evolutionary youth, it would be nothing short of astounding if it were not profoundly deceptive *somehow*. Sufficiency and systematic dependency, meanwhile, suggest that we would be all but blind to the insufficiencies imposed by encapsulation...

Should we suppose we simply got *lucky*? That conscious cognition is ‘just so’? That all the millennial enigmas and interminable dissent and difficulties with intentionality are simply a matter of ‘getting our concepts right’?

BBT answers, No. We quite literally do not have the *consciousness* we think. The “beautiful many-voiced fugue of the human mind” (Hofstadter, 1979, p. 719) could be the last of the great ancient conceits. Given its explanatory reach,³⁶ I sometimes fear that what we call ‘consciousness’ *does not exist at all*, that we ‘just are’ an integrative informatic process of a certain kind, possessing *none* of the characteristics we intuitively attribute to ourselves. Imagine *all of your life* amounting to nothing more than a series of distortions and illusions attending a recursive twist in some organism’s brain. For more than ten years I have been mulling ‘brain blindness,’ dreading it—even hating it. Threads of it appear in every novel I have written.³⁷ And I *still* can’t quite bring myself to believe it.

I do know that I lack the ability to dismiss it. I used to see it as a kind of informal reductio, a position disqualified by the apparent absurdity of its conclusions. But inquiry is never a ‘rescue operation.’ We’re not here to *save anything*. Given what we know about the myriad biases that afflict reasoning (Sperber and Mercier, 2011, Kahneman, 2011), beginning with an apologetic mindset in matters so fraught with ambiguity and complication is simply asking to be duped. Why should consciousness be any kinder to our preconceptions than, say, quantum physics? Perhaps the nearer we come to understanding ourselves, the more we should expect our ‘commonsense’ to be overturned.

Likewise, the charge of performative contradiction can only beg the question. BBT does not require the sufficiency of conscious cognition to argue its insufficiency. Otherwise, it is an eminently plausible empirical possibility that we could find ourselves in these straits. Not only is it possible that evolution stranded us with a peephole perspective on our cognition, we have good reason to think it probable. In fact, given the informatic constraints BBT proposes, you could argue that *any* sapient

³⁶ The big question unanswered here, of course, is how BBT might bear on the Hard Problem. I think it has simplified the labyrinth simply by explaining away the illusions considered here. The maze becomes simpler still when one applies it to the vexed issue of qualia. BBT, however, is a theory of the *appearance* of consciousness, not a theory of consciousness. It does not explain why any RS should exhibit consciousness at all. What it does do, however, is ‘whittle consciousness down to size.’ Once features such as the now can be explained away, the Problem no longer seems anywhere so Hard. BBT, in other words, could allow researchers to separate the problem from the labyrinth.

³⁷ Only *Neuropath* (2009) deals with the theory in any sustained manner.

biological consciousness would find itself in a similar cognitive quandary—that extraterrestrial aliens, in effect, would possess their own versions of our philosophical problems.

To make matters worse, BBT seems consistent with at least two pessimistic inductions commonly drawn from the history of science. If you think of scientific development in terms of what Sellars called the “gradual depersonalization of the original image” (1962), the slow expunging of intentionality from the natural world, then we should perhaps expect that it would be banished from our little neural corner of the natural world as well. Once again, evolution requires only that our behaviours be effective, not that we be privy to their truth—or anything else for that matter.

And this speaks to a second humbling conclusion that can be drawn from the history of science: that the human answers to the natural *in all ways*. Time and again the scientific ‘view from nowhere’ has disabused us of our first-person perspectival conceits, showing us, most famously, that we were a peripheral speck relative to the universe rather than the substantial centre, and just another twig on the tree of life and not the image of the gardener. The former was relatively easy to relinquish, given that not all ‘centres’ need be spatial. The latter was more difficult to reconcile, but seemed okay, so long as we could keep our gardening *tools*...

BBT argues there are no exemptions, that in a strange sense, we are as inconsequential *with reference to ourselves* as we are to the rest of nature. Given that the RS is informatically localized, why should we view consciousness as a ‘perspective on the world’ as opposed to a perspective *on a perspective* on the world? It recontextualizes us relative to the informatic cosmos of our brain the way astrophysics recontextualized us relative to the universe or the way evolution recontextualized us relative to life. And it does so, moreover, in a way that explains *why* we made these mistakes in the first place.

BBT offers a simple, principled way to puzzle through the riddle of the first-person, if only at the cost of explaining it away.³⁸ Its parsimony alone warrants its serious consideration. BBT turns the question of consciousness around the way Einstein turned around the question of gravity. Where Newton looked at gravity as a discrete *force* acting on objects with inertial dispositions (which is to say, as an *accomplishment*) Einstein transformed gravity into a feature of space-time, into *an expression of structure*. BBT performs an analogous figure-field switch. Where the philosophical tradition looks at unity, say, as an achievement of some discrete neural complex, BBT transforms it into a misleading structural expression of *what consciousness is*, which is just to say, recursive neural information integration. Why is consciousness unified? Because it is too information poor to be otherwise.

Think of each word here as a thimble buzzing with a billion neural interactions. Resist the intuition of sufficiency and the corresponding assumption that conscious perception and thought constitute some kind of functionally autonomous system. Only our cognitive activity *as we are aware of it* requires this happy state of affairs, so set it aside as a possible source of confusion. Dispense with norms. Dispense with representations—all the fractious conceptual children of intentionality. Think of the brain as *just another natural processing system*, one able to arrest and amend its routines via some recursively integrative subsystem. Seen in this light, the kinds of things discussed here are perhaps not so extreme.

Perhaps BBT is on the right track despite its radical revisionism. Perhaps it is time, once again, to acknowledge that we are smaller than we supposed, know less than we hoped, and are more frightened than we care to admit. “Nature,” as Newton famously wrote, “is pleased with simplicity” (2010, p. 398), even if we are horrified.

References

³⁸ This, as should be clear, is what distinguishes BBT from eliminativism more generally: it proposes a mechanism—the open-closed structure of the RS—that allows for a systematic diagnosis of the distortions and illusions belonging to the first person perspective.

- Aristotle. (1969) *Aristotle's Physics*, Translated with Commentaries and Glossary by H. G. Apostle. Grinnell: The Peripatetic Press.
- Augustine. (1992) *Confessions*, Translated with an Introduction and Notes by H. Chadwick. New York: Oxford University Press.
- Bacon, F. (1985) *The Essays*, New York: Penguin.
- Bakker, R. S. (2009) *Neuropath*, New York: Tor.
- Blackmore, S. (2006) *Conversations on Consciousness: What the Best Minds Think about the Brain, Free Will, and What it Means to be Human*, New York: Oxford University Press.
- Burton, R. A. (2008) *On Being Certain: Believing You Are Right Even When You're Not*, New York: St. Martin's Press.
- Celesia, G. G. and Brigell, M. G. (2005) Cortical blindness and visual anosognosia, *Handbook of Clinical Neurophysiology*, **5**, pp. 429-440.
- Chaitin, G. J. (2006) *Meta Math! The Quest for Omega*. New York: Vintage Books.
- Churchland, P.M. (1989). *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. Cambridge: MIT Press.
- Churchland, P. S. and Ramachandran, V. S. (1993) Filling in: why Dennett is wrong, *Dennett and his Critics: Demystifying Mind*, Cambridge: Blackwell.
- Clark, A. Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science, (forthcoming).
- Dennett, D. (1987) *The Intentional Stance*, Cambridge: The MIT Press.
- Dennett, D. (1991) *Consciousness Explained*, Boston: Little, Brown, and Company.
- Dennett, D. (2005) *Sweet Dreams: Philosophical Obstacles to a Science of Consciousness*, Cambridge: The MIT Press.
- Descartes, R. (1968) *Discourse on Method and The Meditations*, Translated by, F. E. Sutcliffe. New York: Penguin.
- Drescher, G. L. (2006) *Good and Real: Demystifying Paradoxes from Ethics to Physics*, Cambridge: The MIT Press.
- Edelman, G. (2005) *Wider than the Sky: The Phenomenal Gift of Consciousness*, New Haven: Yale University Press.
- Floridi, L. (2011) *The Philosophy of Information*, Oxford: Oxford University Press.

- Greenemeier, L. (2009) Computers have a lot to learn from the human brain, engineers say, *scientificamerican.com* News blog, [blog] 10 March. Available at: <http://www.scientificamerican.com/blog/post.cfm?id=computers-have-a-lot-to-learn-from-2009-03-10> [Accessed 30 December 2011].
- Heidegger, M. (1996) *Being and Time: A Translation of Sein und Zeit*, Translated by J. Stambaugh. Albany: State University of New York Press.
- Hofstadter, D. R. (1979) *Godel, Escher, Bach: An Eternal Golden Braid*, New York: Vintage.
- Hintikka, J. (1973) *Logic, Language Games and Information: Kantian Themes in the Philosophy of Logic*, Oxford: Clarendon Press.
- Hume, D. (1888) *A Treatise of Human Nature*, New York: Oxford University Press.
- Husserl, E. (1964) *Phenomenology of Internal Time-Consciousness*, Translated by J. S. Churchill. Evanston: Indiana University Press.
- Jackson, F. (1982) Epiphenomenal qualia, *Philosophical Quarterly*, **32**, pp. 27-36.
- James, W. (2007) *The Principles of Psychology, Vol. 1*, New York: Cosimo.
- Kahneman, D. (2011) *Thinking, Fast and Slow*, Toronto: Doubleday Canada.
- Kant, I. (1929) *Immanuel Kant's Critique of Pure Reason*, Translated by N. K. Smith. London: The MacMillan Press.
- Lewis, D. (1986) *On the Plurality of Worlds*. Oxford: Blackwell.
- Marcus, G. (2008) *Kluge: The Haphazard Construction of the Human Mind*, Boston: Houghton Mifflin Company
- Marois, R. and Ivanoff, J. (2006) Capacity limits in information processing in the brain, *Trends in Cognitive Sciences*, **9** (6), pp. 296-305.
- McLeod, P. (1987) Reaction time and high-speed ball games, *Perception*, **16** (1), pp. 49-59.
- McTaggart, J. M. E. (1908) The unreality of time, *Mind*, **17**, pp. 457-474.
- Mercier, H. and Sperber, D. (2011) Why do humans reason? Arguments for an argumentative theory, *Behavioural and Brain Sciences*, **34**, pp. 57-111.
- Metzinger, T. (2003) *Being No One: The Self-Model Theory of Subjectivity*, Cambridge: The MIT Press.
- National Academy of Engineering, (2011) *Reverse Engineer the Brain*. [online] Available at: <http://www.engineeringchallenges.org/cms/8996/9109.aspx> [Accessed 30 December 2011]
- Newton, I. (2010) *Sir Isaac Newton's Mathematical Principles of Natural Philosophy and his System of the World*, Translated by A. Motte. Whitefish: Kessinger Publishing.

- Nietzsche, F. (1967) *On The Genealogy of Morals and Ecce Homo*, Translated by W. Kaufmann. New York: Vintage Books.
- Norretranders, T. (1999) *The User Illusion: Cutting Consciousness Down to Size*, New York: Penguin.
- Pöppel, E. (2009) Pre-semantically Defined Temporal Windows for Cognitive Processing, *Philosophical Transactions of the Royal Society*, **364**, pp. 1887-1896.
- Prigatano, G. P. ed. (2010) *The Study of Anosognosia*, New York: Oxford University Press.
- Pylyshyn, Z. (1999) Is vision continuous with cognition? The case for cognitive impenetrability of visual perception, *Behavioural and Brain Sciences*, **22**, pp. 341-423.
- Schwitzgebel, E. (2012) The Crazyist Metaphysics of Mind.
<http://www.faculty.ucr.edu/~eschwitz/SchwitzPapers/CrazyMind-120229.pdf> [Accessed 15 March 2012]
- Sellars, W. (1962) Philosophy and the scientific image of man. In: R. Colodny, ed. 1962. *Frontiers of Science and Philosophy*. Pittsburgh: University of Pittsburgh Press, pp. 35-78.
- Tononi, G. (2012) The Integrated Information Theory of Consciousness: An Updated Account,
http://www.architalbiol.org/aib/article/view/1388/pdf_1 [Accessed 21 March 2012]
- Topolinski, S. and Reber, R. (2010) Gaining insight into the ‘aha’ experience, *Current Directions in Psychological Science*, **0**, pp. 1-4.
- Wackerman, J. (2007) Inner and outer horizons of time experience, *The Spanish Journal of Psychology*, **10** (1), pp. 20-32.
- Wegner, D. (2002) *The Illusion of Conscious Will*, Cambridge: The MIT Press.