

## 6 — Dennett on Smith (and reply)<sup>†\*</sup>

Hugh Clapin  
University of Sydney

*Brian Cantwell Smith holds a rare and valuable intellectual pedigree for a philosopher of mind. As a principal scientist at the Xerox Palo Alto Research Center (PARC) and founder of the Center for the Study of Language and Information at Stanford University (CSLI), he has studied foundational questions in computability and computer programming. Through this work he has come to the conclusion that the representational capacities of artificial systems such as computers raise profound metaphysical and epistemological questions.*

*In his “One Hundred Billion Lines of C++” (1997), Smith illustrates how misleading is the ordinary philosophical conception of computer programming. Standard programming practice is not (as is often assumed) committed to classical cognitive architectures. In particular, the processes implemented by executing programs have nothing like language of thought structure; none the less they make use of representations successfully to negotiate the world. They provide a rich resource of physical representation systems that are effective but don’t fit the ordinary analyses of the philosophy of mind.*

*Smith’s work may be aligned with the situated cognition tradition*

---

<sup>†</sup>Originally published in Hugh Clapin (ed.), *Philosophy of Mental Representation*, Oxford University Press, 2002, pp. ■■–■■.

<sup>\*</sup>Coach House Institute, Faculty of Information, University of Toronto  
90 Wellesley St W, Toronto, Ontario M5S 1C5 Canada

<sup>†</sup>© Brian Cantwell Smith 2009

Last edited: November 20, 2009

Please do not copy or cite

Comments welcome

**Draft only (version 0.80)**

brian.cantwell.smith@utoronto.ca

due to Barwise and Perry (1983). This approach emphasizes the importance of context in determining meaning. The situated semanticist is inclined to begin her theory of meaning with indexicals and other radically context-sensitive representations. Tokens of 'I' have very little meaning independent of how, when, where, and by whom they are used. More generally, the situated approach to cognition places significant emphasis on the contribution of the situation of the organism to that organism's cognitive processes.

Smith argues that as soon as we register the world using a system of representation, we make a set of strong assumptions about the way the world is. His task has been to show the profound consequences of this insight for the study of systems of representation.

Smith makes use of an engaging imaginative strategy to draw attention to the theoretical moves required to explain the occurrence of representation using only the resources of a representation-free physical world. Smith urges us to consider whether we need to think in terms of objects at all. Might an ontology consisting only of Strawson's (1959) 'features' be sufficient? When we declare that 'It's raining' we are drawing attention to a feature (raining) without being committed to any particular object that has that feature. Smith suggests we begin by thinking of the physicist's world as populated not by objects but field-densities. This field-theoretic description can be comprehensive while admitting only of field-densities for a small range of properties (for example, gravitational fields, electromagnetic fields, etc.).

Smith suggests that the common-sense world of middle-sized objects is an achievement of our representational practices. Representation is achieved when one aspect of the mish-mash of fields is able to separate in a certain way from the rest of the mish-mash. This region, the 's-region', is (or is becoming) the subject-something that represents the world. Smith first emphasizes the distance required between the representation and the represented, and secondly the need for coordination between the two. This coordination is likened to the actions of an acrobat who dances around a stage, but keeps a torch beam focused on one spot. The torch must undergo dramatic changes in orientation to maintain its focus at one point. The intentional acrobat is similarly dynamic in keeping its intentional objects stably registered.

Smith builds on this fundamental picture to argue that all representation is partly context-dependent, or deictic. Smith is scrupulous about the reflexive morals which thus apply. Acts of representation

*bring the world's objects and properties into being (as objects and properties), and any attempt to talk about the world will be an act of representation, and thus an act of object-making. This makes likely what Smith calls 'inscription errors' or 'pre-emptive registration'. For example, it is difficult to talk about the world except as containing objects with properties. But if this is due to the subject-predicate structure of language, then it would be an error to infer that the world must be so constructed.*

*In short, Smith says that representation is an immensely complex, powerful, and sophisticated achievement of the physical world. We are so adept at representing that we are apt to neglect this point and think it an easy and simple procedure.*

*The paper 'Rehabilitating Representation' (forthcoming c) amply illustrates what Smith takes to be the more practical implications of his view. Both classical and embodied/embedded approaches to cognition misunderstand representation. The former places too much emphasis on formality and the non-semantic; the latter places too much emphasis on the causal, local interactions between the system and the world, underestimating the importance of disconnection to intentionality.*

*The rehabilitation required involves acknowledging that representation is about causal connectivity to the world, but not a direct, local, or simple connectivity. Representing subjects, by virtue of their representations, participate fully in the world (not just the skin boundary of the world), help constitute the world (by virtue of the entanglement of ontology with representation), but are able to maintain a separateness from the world.*

*Dennett, despite being a self-proclaimed 'reluctant metaphysician', is sympathetic to Smith's metaphysical project (though perhaps is not completely converted). His dispute with Smith concerns the role of evolution in explaining the difficult achievements of representation and objectification. Objectification, says Dennett, is an evolutionary 'Good Trick', which was likely to be stumbled on because it provides significant selective advantage. Dennett also objects to what he takes to be Smith's commitment to the determinacy of mental content.*

## **6a — Brian Cantwell Smith on Evolution, Objectivity, and Intentionality**

Daniel C. Dennett  
Tufts University

### **1 An Original Account of Intentionality and Objects**

Like the rest of us, Smith wants to steer between the Scylla of GOFAI and the Charybdis of Dynamical Non-Representational Systems, and he adds to the feast his own bounty of acute observations and tempting proposals about how such a rehabilitation of mental representations would go. But he and Haugeland, unlike the rest of us, are ontologists who think we need to reach way back and rehabilitate the whole of metaphysics in order to do this job right. Yikes.

What are the less radical alternatives? One might have thought we could safely presuppose the usual catalog of physical objects—ranging in size from sub-atomic particles through tables and mountains to galaxies—and their properties—mass, charge, location, shape, color...and then simply explore the question of which complicated organizations of such objects count as believers, or representations, or symbols...and why. That is the strategy that has worked so triumphantly for magnetism and metabolism, photosynthesis and jet propulsion. Why not for mental representation, too? If we can explain growing an apple, and eating an apple, why not seeing an apple and wanting an apple and reidentifying an apple?

Why not indeed? I have always been a reluctant metaphysician, and Rob Cummins and Andy Clark seem to me to have shared my optimism about the innocence of the standard inventory of what we might call the ontology of everyday life and engineering. We happy sailors on Neurath's ship resist the alarm calls of Smith and Haugeland.<sup>1</sup> Do we have to put on our life-jackets and jump overboard and get all wet doing a lifeboat drill? Maybe, and maybe not. But it can't hurt. A lifeboat drill is a great way to reassure ourselves that we know what we're doing. And actually going through with it—not just imagining going through with it—is

the only way to get this reassurance. If we end up with pretty much the same inventory and explanations we thought we were going to use in the first place, it will be a sounder ship that continues the voyage. And maybe we'll discover something important that has been distorting all our other projects.

For anyone who shares my conviction that traditional or 'pure' metaphysics is a played-out game, a mandarin pursuit so isolated from the rest of human inquiry that it is extremely unlikely to find enough leverage to move us from our comfortable habits, Smith's project is apt to be appealing. Only somebody coming from outside philosophy, somebody whose driving problem is not philosophical but somehow more 'practical' (however abstruse and theoretical relative to farming or building bridges) could hold my attention in a metaphysical exercise, and Smith has been led to his metaphysical vision by decades of struggling with problems that are eminently practical-problems arising not just in the crypto-philosophical arena of AI, but in engineering, for heaven's sake, in the design of hardware and software for all manner of applications. His title "100 Billion Lines of C++" sings to me, then. If disk operating systems, word-processors, and web-browsers confront problems of reference and meaning that can only be alleviated by some revisionary metaphysics, I am all ears. But still, dragging my feet. Constructively, I hope.

Let's start with what Smith calls The Representational Mandate:

#### The Representational Mandate

##### 1. Conditions

- a. A representational system must work, physically, in virtue of its concrete material embodiment (the role of effectiveness).
- b. But it is normatively directed or oriented towards what is non-effective-paradigmatically including what is physically distal.
- c. Being neither oracle nor angel, it has no magic (non-causal, divine) access to those non-effective situations; just caring about them is not enough (physical limitations bite hard!);

2. So what does the system do?
3. It
  - a. Exploits local, effective properties that it can use, but doesn't (intrinsically) care about-i.e. inner states of its body and physical make-up, in interaction with the accessible (effective) physical aspects of its environment.
  - b. To 'stand in for' or 'serve in place of' effective connection with states that it is not (and cannot be) effectively coupled to
  - c. So as to lead it to behave appropriately toward those remote or distal or other non-effective situations that it does care about, but cannot use. (Smith, forthcoming c, hereafter: RR.)

I will be surprised if anybody here has any serious quarrel with Smith's Representational Mandate (though it is easy enough to think of absent theorists who would squirm or rage). But some of us may be taking Smith's Mandate and interpreting it down, understanding it in a less radical way than he would wish. In the hope of giving his vision of it a proper outing, I will first try to give a summary of what strikes me as the dozen or so main points in Smith's work that bear on the issues of mental representation. (Much of the most interesting stuff in his book I'm going to set aside, reluctantly.)

## 2 A Dozen Important Points

1. Why re-tool our ontology? If we don't, if we complacently (or opportunistically) cling to the standard inventory, we will commit what Smith calls inscription errors or pre-emptive registration:<sup>2</sup> a tendency for a theorist or observer, first, to write or project or impose or inscribe a set of ontological assumptions onto a computational system (onto the system itself, onto the task domain, onto the relation between the two, and so forth), and then, second, to read those assumptions or their consequences back off the system, as if that constituted an independent empirical discovery or theoretical result. (Smith 1996: 50, hereafter: OO)

Pre-emptive registration is a sort of metaphysical anachronism, back-projecting onto our vision of ultimate-or at any rate more fundamental-reality a category or assumption that is in fact the

effect or artefact of some later, higher-level, more ‘expensive’ development. [[Discussion point 5.1]]

**2.** The granddaddy case of pre-emptive registration is imagining we can parse the universe primordially into objects, which may or may not be appreciated in their object-hood by any (psychological) subjects in the neighborhood. By objects, Smith means what we (now) mean by objects-things that have spatio-temporal boundaries (at least roughly), that have careers, that can be re-identified, and that can, on occasion, be present to subjects-as objects to be perceived, sought, remembered, thought about, moved, destroyed, gathered, and so forth. As he puts it in RR, the world doesn’t come ‘pre-parsed’ into objects, properties, relations, and other ‘formal’ categories.

**3.** The antidote to this form of pre-emptive registration is hard to swallow, but Smith gives us lots of help with various imagination-aids, temporary ploys, and other delicious candy-coatings. If I understand him right, it is actually strictly impossible to describe the primordial state without committing some sort of pre-emptive registration, since words-any words we can use-already bias us in favor of objectification of just the sorts he wants to describe the birth of. If I understand both Smith and Haugeland (unlikelihood squared) on this matter, they both think one can tiptoe past this problem (of the apparent inevitability of inscription errors in our attempts to do metaphysics). Here is how I put it in my review of Haugeland’s book:

The task facing any ‘Heideggerian/Kantian’ theorist is to do justice to the role of us in constituting the denizens of ‘our’ world without lapsing into awful relativism/subjectivism on one side or caving in to noumena, or a ‘God’s eye view’ on the other. Haugeland’s solution, which grows on me, is to show how and why it is hard to ‘constitute’ a world (that takes care of anything-goes relativism) but not because there is a privileged way that the world-the real world-has always been constituted. His view is a close kin, I think, of my view of the evolution of colors: Before color vision evolved on this planet, sunsets and cliffs and volcanic eruptions had the reflective properties they did, but it makes no sense to ask if those sunsets were, say, red-since that question has

no meaning independently of a reference class of normal observers. We can of course extrapolate back from our current vantage point and fix and answer such questions, using ourselves as the touchstone for colors, but we must recognize that we are doing that. [That is, as it were, acknowledging the pre-emptive registration that you're doing, and just keeping track of the fact that you're doing it. You're keeping yourself and your own categories somehow as a touchstone to talk about something to which they're not really directly appropriate.] Were there dinosaurs before *H. sapiens* came along and invented censoriousness and then ontology so that dinosaurs could be constituted? Of course there were, but don't make the mistake of thinking that this acknowledges a fact that is independent of *H. sapiens*. [[Dennett 1999: 433–4]]

I don't see that Smith's view of this is different, and that's fine, since I think this is a good and defensible view. Discussion point 5.2

**4.** With that apologia (or is it a caveat?) in place, I can now (pretend to) describe the primordial basis, the out-of-which that objects find their origin in. It is (very roughly) a Heraclitean world of flux, dynamically flowing and concentrating and dissolving. What is it composed of? Well, you really shouldn't ask, barefaced, since any answer will involve registration that is to some degree pre-emptive; but since we must advance the discussion, let's just speak of features. Don't worry; this is just a temporary stopgap: 'That the distinction between features and properties and objects is not sharp, on the other hand-that logic is messy, not just finger paints-will not ultimately be a problem, at least not for us' (OO: 127).

Features, I take it, are ways one region can be different from the neighboring region. Here 'neighboring' means, constitutively, in effective interaction with. One of my favorite dicta in a work filled with arresting phrases: 'Distance is what there is no action at' (OO: 200 n. 11).

**5.** This idea of locality underlies Smith's account of another kind of pre-emptive registration, highlighted in RR: the family of errors that occur when we persist in casting what really ought to be a

theory of effective processes (or just effectiveness) as a theory of effective computation. The idea is that the truly important phenomenon of effectiveness is not a particularly computational phenomenon—it is a sort of historical accident that our first intellectual grip on effectiveness came via the work of Turing, Church, and their kind. Smith suggests that all kinds of mechanisms are effective without being computational in the ways that foster spurious connotations (of semantics, of proof in formal systems, etc.). This then seduces us into further pre-emptive registration and more inscription errors, taking ‘logician’ baggage along for trips where it proves worse than useless. The central idea of effectiveness, Smith claims, has to do with local, non-distal causation.

**6.** The importance of ‘flex and slop’: Interactive effects dissipate, diminish with distance and time, due to what Hume once called ‘a certain looseness’ in the world (Hume 1739/1978: 408 (ii. iii. 2)). If the whole universe were like a gigantic interlocked gear-world, in which nothing could move without propagating effects *ad infinitum*, nothing could be out of touch with, or inaccessible to, anything else; nothing could be alone, or individual. Nothing could ‘keep its distance’ without flex and slop, which is a heretofore unremarked precondition for intentionality, because it creates the distance that then creates the problems that the varieties of reference-negotiation solve. [[Discussion point 5.3]]

**7.** Particularity is not individuality. The primordial physics world is everywhere particular, but contains no individuals (OO: 124–5). [[Discussion point 5.4]]

As I said at the outset, what appeals to me about Smith’s project is that he’s coming to this from a career in computer science, not from a career of teaching metaphysics. I’m trying to reconstruct the head-scratchings in computer science that make this seem so attractive, and it seems to me that they are something like the ultimate Y2K problem. The Y2K problem was not having enough bits for the year-settling for 2 when you should have 4, or, if you want to take a longer view, 5, or if you want to take Smith’s view, many many many more. That is, when you start representing the world, if you’re using any sort of data structure, you stop short with *n* fields, and *n* fields is in a certain sense never enough

for a concrete thing, even something as simple as a cup.

The reason we make something a cup is that we have to create our little  $\Upsilon$ 2K problem. There are only so many fields that we can carry along in our representation of the cup. We realize that if we want to keep track of that cup, there are all sorts of futures that we're going to have real trouble tracking if the cup gets smashed and then reconstituted, if it gets sold, if it gets repainted... There are so many different things that can happen to that cup. If we want to have a data structure that refers to that cup—come what may—it's going to have to have too many fields. We just can't do it. This is Haugeland's point, I think: a description of a person can't go into everything that's determinate about that person. [[See Discussion point 5.4—ed.]] It simply leaves out a lot of fields. There's a lot of bits that just aren't fixed and there's no room to fix them. [[Discussion point 5.5]]

**8.** Chiming a point also made rather differently by Cummins (1996a), Smith offers several arguments to show that reference, and semantic relations in general, cannot be effective or causal relations. We can refer effortlessly to things outside our light cone, for instance, and the whole point of having something local by which you keep track of something distal is to overcome (without guarantee) the non-effectiveness of all such distal relations (OO: 157, 210–11, 228; see also the Representational Mandate 1c, above). And, like Cummins, Smith sees this as providing the elbow room for error (OO: 223).

**9.** The sort of 'non-effective tracking' exhibited by Smith's imaginary supersunflower is the forerunner of semantics, the basis of intentionality.<sup>3</sup> It is not what Smith calls registration, but it is the competence out of which registration can ultimately grow.

In all these situations, what starts out as effectively coupled is gradually pulled apart, but separated in such a way as to honor a non-effective long-distance coordination condition, leading eventually to effective reconnection or reconciliation. There is a great deal more to intentionality than that...but in various forms these notions of connection, gradual disconnection, maintenance of coordination while disconnected or separated and ultimate reconnection or reconciliation permeate all kinds of more sophisticated example. (OO: 206)

**10.** ‘The retraction of responsibility onto the *s*-region [forerunner of the subject] is the origin of registration’s asymmetry and directedness’ (OO: 223). (I’ll have more to say about this later, mostly critical, but reluctantly so, since I love the pedagogical uses to which he puts this mythic image of the *s*-region parting from its partner.) Smith’s ulterior aim in this imaginative theme is to highlight the importance of the perspective shift he advocates in the next point.

**11.** It is the emergence of dynamically coordinated variation-systems (illustrated winningly by the intentional acrobat’s flashlight, and the ‘columnar’-shaped ‘sustaining physical field’ that unites the frog to the fly, OO: 217) that explain ‘why we see trees, not electromagnetic radiation’. I think this point is strongly related to some of Ruth Millikan’s (1984: ch. 15; 2000b: §§7.1-2) observations on identifying the reference/function of something by finding what holds constant across occasions, when we ‘turn the knobs’. It is not just co-variance but systematic co-variance—which won’t be perfect since systems are costly and may have weaknesses—that underlies our identification of objects of experience. [[Discussion point 5.6]]

**12.** There are a variety of instances in which philosophers have traditionally dealt with dichotomies and Smith shows us how to see these as extremal points along some axis of variation. Thus the philosopher’s ideal of a purely non-deictic registration is a myth (OO: 249). We have cases that are halfway between implicit and explicit, halfway between ‘pure’ reference and intension (e.g. OO: 251), and so forth. These middle-ground cases are very important in Smith’s larger scheme of things, since as he eloquently says (OO: 254–5), the main lesson to be learned is ‘not to be seduced by limit cases’. (See also the end of ch. 8 (and of the book): Life-what matters-happens in the middle ground.)

These strike me as the main things I have learned from Smith’s book and RR. Let me add what I take to be the main point of ‘One Hundred Billion Lines of C++’ (1997), to make a baker’s dozen:

**13.** The productivity, the compositionality of programming languages (such as C++) should not lead us to suppose that in general

the processes such programming languages permit us to design and implement are similarly compositional. The fact that the programmer can create indefinitely many identifiers (and indeed can create nonce-systems of compositionality on the fly as he goes) does not at all imply that the identifiers thereby created can be treated as manipulable, composition-friendly items by the program itself. The compositionality is in the syntax and semantics of the source code but not in the structures that then get built and then actually get implemented and then run.

This insight, restored to philosophy, shows not that Fodor's language of thought is not the way we work, but does show that there is nothing remotely like a plausible inference to the conclusion that there is a language of thought from the premise that the brain engages in computational-like processes whereby it extracts apt behavior from the information it extracts from the world.

Now I take the upshot of all this to be a multi-path attack on the 'classical' ideal of mental representations as well modeled by 'propositional' symbol systems that obtain their intentionality by composing something like Fregean Thoughts out of Terms with Extensions and Intensions. Every tractable theory has lots of idealizations and simplifications, but the idealizations of that family of theories are trouble-makers, not helpers—largely because of preemptive registration: they create the illusion of sharp distinctions where in reality there is something of a spectrum, from 'non-effective tracking' to the most intellectual of opinions (e.g. my opinion that the shortest spy is a spy). What Smith calls registrations—occupants of the right-hand region of this spectrum, you might say—only work in contexts of 'coordination conditions', in adjustment or compensation (what Smith calls 'intentional dynamics', OO: 262), processes that philosophers have tended to overlook or underestimate the importance of. Smith puts to good use one of my own images to skewer the false view: the classical system of uninterpreted symbols is seen as wearing a thin 'overcoat' of transducers and effectors as the interface between symbol and world. [[Discussion point 5.7]]

Now I want to offer what I take to be a friendly amendment, but I expect Smith will view it askance. If my expectation is mistaken, hurrah; if it's right, there is no question I am more than eager to

explore than why he resists this (to me) obvious improvement.

### 3 The Origin of Objects?

For me, the ghost at Smith's banquet is—surprise, surprise—Charles Darwin. Evolution is hardly mentioned in his book, whose very title trumpets its likely affinity to Darwin's great vision. How can we have a story—a Just So story, in fact, eloquently brandishing its own unavoidable metaphors and anachronisms—of the origin of objects, of their emerging onto the contemporary landscape from some primordial scene in which they were absent—and not have it rely on the fundamental Darwinian principle of natural selection? What alternative shaping forces could do the work that needs to be done? Smith does a wonderful job of showing us the 'expensiveness' of objects and subjects; something has to pay for all this R&D!

Let me draw your attention to a few crucial points in his account where I, Darwinian Fundamentalist that I am, feel an irresistible urge to insert evolutionary considerations. Look again at the Representational Mandate: a representational system is 'normatively directed'(1b); it 'exploits local, effective properties' (3a) 'so as to lead it to behave appropriately' (3c). Smith's examples—the supersunflower, the frog, and (most important) the unnamed simpler organisms who pioneer the passage from proximal irritation to distal 'non-effective tracking'—all bespeak his interest in evolution, in simple minds and their successors, but he strangely eschews the evolutionary perspective. Why? Because, I think, he wants to avoid what he takes to be the pre-emption error of what we might call premature teleology, or premature function. But he overdoes it, methinks. He wants to introduce normativity in his way, not riding on the coat-tails of evolutionary normativity. But I think this is a mistake, too. All normativity does ride on Darwin's coat-tails. In trying not to be 'expensive' Smith goes too far here. Consider, for instance, his excellent summary (OO: 241):

"The underlying spatio-temporal extended fields of particularity throw tufts of effective activity up against each other, and let them fall apart, fuse them and splinter them and push them through each other, and generally bash them around, in ways governed by the pervasive underlying

(physical) laws of deictic coupling. [So far, no hint of teleology; this is all just Heraclitean flux, signifying nothing.] For a subject to begin to register an object as an individual is, first, for a region of the fields (the s-region) not to be connected to another region (the o-region), but in the appropriate way *[[my italics]]* to let go of it....The coordination requires establishing appropriately *[my italics]* stable (extended in the s-region) and abstract (extended in the o-region) focus on the o-region, while remaining separate. The separation helps in maintaining the s-region from being buffeted by every nuance and vibration suffered by the o-region."

Notice how we end with pure engineering: protection of the s-region from buffeting, in order to maintain a 'focus'-on an appropriate object. The fact is that s-regions that happened to begin to register inappropriate o-regions (don't-cares) or to register suitable o-regions inappropriately (inefficiently, counter-productively, etc.) would not last long in the buffeting flux, not long enough to out-reproduce the competition in any case. Once we add this evolutionary point, we can emend Smith's account, adding what strikes me as its most important theorem: the world doesn't come 'pre-parsed' into objects and properties (just as Smith says) but objectification is what I call an evolutionary Good Trick (Dennett 1995a, hereafter: DDI), an elegant solution to the problem of staying alive in the world of flux, flex, and slop, a solution we would expect to find, for instance, in other galaxies in which life had evolved.

Will Smith want to go that far with me? I hope so. His pluralism is sane and temperate. By taking pluralism (and postmodernism more generally) seriously (and not just pre-emptively dismissing it with a sneer, as it is extremely tempting to do) he allows it to tame itself. Yes, there are real problems of pluralism, and yes, there is no guarantee at all of a single, pre-given ontology to which we can anchor all reference, but reference-preservation, or reference-negotiation, is a problem that we can solve, and routinely do solve. (Don't patronize the Others. You can be brought to understand their ontology and they can understand yours, with a little effort.) There is a Good Trick (maybe two or three,

but we know of one for sure that works well) that has been discovered again and again by evolution, and Smith has a deeply insightful account of how it works to generate our ontology.

I think this evolutionary perspective on the birth of intentionality is preferable to the charming myth that Smith puts in its place: ‘In all these situations, what starts out as effectively coupled is gradually pulled apart, but separated in such a way as to honor a non-effective long-distance coordination condition, leading eventually to effective reconnection or reconciliation’ (OO: 206), which ignores the fundamental evolutionary facts: we only ‘want’ to be coordinated to the things that matter to us, and these are not necessarily things we used to be attached to. The food I hope to coordinate with has never been within hailing distance of me till now, but I pounce on it just the same. I love Smith’s imagery—especially his Country and Western song sound bite: ‘How can I miss you if you won’t go away?’—but it reminds me, I fear, of another cool idea (Freud had a lot of fun with it) that we evolutionists have shown how to replace: what we might call the Siamese-twin theory of sexuality, which imagines a primordial time when male and female were happily united, later cruelly sundered, and spending the rest of eternity as ‘halves’ trying to reunite. The evolution of sexuality is a deep and fascinating problem, since it, too, is expensive and needs to be paid for, but we don’t solve the problem by imagining that an m-region and an f-region gradually got pulled apart and are striving to reunite.

Smith says at one point: ‘Third, the retraction of responsibility onto the s-region is the origin of registration’s asymmetry and directedness’ (OO: 223). This serves to balance his various claims about the shared roles of subject and object. As he says, the dance has two partners but is not symmetrical. By leaving out evolution, however, he leaves out what I take to be the deeper reason for the asymmetry. The sun doesn’t give a damn about the sunflower, but the sunflower needs the sun. You need something more like predator-prey (or mate) asymmetries to make sense of the asymmetry of registration.

I think Smith ought to accept all of this, and in some passages he sounds just the right notes. For instance, he notes that ‘a distinction takes hold between what the s-region is doing (tracking the coyote or incident sunlight) and how it is doing it. The for-

mer gets at a non-effective regularity; the latter, at an effective mechanism whose “job” is to implement or sustain it. Among other things, this split provides a toehold for normativity to attach its tentacles’ (OO: 222). Exactly: An evolutionary toehold for normativity.<sup>4</sup>

What might be fueling his resistance, then? In his account of what he calls ‘intentional dynamics’, his name for the theoretical basis of situated cognition, he tells us he wants to keep the normative at bay (OO: 262): he doesn’t want to build the normative condition into the name (by calling intentional dynamics something like ‘rationality’ or ‘reason’). Fair enough; we need to understand the underlying physics, if you like, that any representational scheme, good, bad, or indifferent, must cope with, so we must be careful to describe not just the (presumably) optimal mechanisms, but also the junk that might be lying around interfering. Bad engineering and good engineering live in the same world, and that world should be clearly described without the bias of preemptive registration, if possible. I also think he wants to avoid what might be called ‘premature agency’ a sort of inscription error in which one breaks the world up into things doing things to things, as if this were the primordial catalog. See, for instance, his nice image of getting rid of the potter, OO: 270. But in the end, I gather—mainly from the strong claims in RR (p. 29) about a distinction between static and dynamic norms—that Smith’s reasons for resisting an evolutionary treatment of representation come from...Pittsburgh. The ‘dynamic norms’ claims ring a Haugelandian, Brandomian, McDowellian bell for me, but I don’t buy it. Not yet. I think I’ll stand firm and ask to be shown what’s wrong with my Darwinian fundamentalism, whose motto is All normativity is grounded in evolution and emerges from the cascade of Darwinian algorithms.

#### 4 Coda: Three Reservations

1. *Indeterminacy of Content.* I see a tension between ‘There may not be any compelling reason to believe there is even a metaphysical fact of the matter’ (OO: 55) on the one hand and, ‘We may not know what it is, but that does not mean God leaves the content indeterminate’ (OO: 62) and, on the same topic: ‘Somehow or other—and this I take to be the most important and difficult task

facing the cognitive sciences—it must be possible to have determinate representational content, that is, for there to be a fact of the matter as to how the world is represented' (OO: 68). I ask, 'Why?' Smith says 'it will have to be an answer that does not depend on how anyone registers or individuates those mechanisms—again, for the simple reason that it happens in people, for example, without anyone doing that.' I don't see that as a good reason. This is like Cummins' similarly staunch line on determinacy of content, and I am not yet persuaded. Why can't God leave content indeterminate?

In this tug of war, I tug on the former side, of course. It helps us escape what might be called Cartesian (or 'from the inside') ontology, the view Quine calls the 'museum myth of meaning'. We must not assume that there will be an 'inner' perspective from which semantic facts of the matter can be mined. (See also Ruth Millikan's (1984) critique of 'meaning rationalism'.)

I don't see what's wrong with (my) perspectivalism about this. After all, it is flat true of some computer applications that they can be adopted wholesale for use in another domain (the old chestnut of the chess machine that can play war games, or whatever). See 'The Abilities of Men and Machines' (Dennett 1978b) for an ur-example. Why should it be different when we then look at animals, say? What if the fly-detector machinery is reused (exapted) intact in some later beast? I think Smith is right (and it's a good point) that the semantic/syntactic distinction is not the external/ internal distinction, but I don't think that this further point follows. [[Discussion point 5.8]]

**2. *What about Animals?*** In spite of all the good discussion about frogs (OO: 197, 216–18, and other places) and coyotes, we are left wondering: do clams register? do amoebae? do they objectify? (see OO: 149, 193, 232). Smith (OO: 195) says that larger corporations and communities may be implicated in intentional achievements, but he downplays the role of proper parts of organisms. Why? Smith's bias in favor of human beings is largely uncharted (see my 1999 review of Haugeland on the same topic). Yes, only whole human beings living in whole societies, with slathers of normativity laid on, ever really refer to anything, but then there is lots of quasi-reference. And Smith is the master of pointing to just these

facts. I wish he'd said more about whether dogs reidentify individuals, for instance (a question I take up, and don't answer properly, in *Kinds of Minds* (1996: 113–16), where I explore the case of Ulysses' dog Argos, who seems to recognize him when he returns. Does he? Really? [[See also Discussion point 1.4—ed.]]. In Smith's brief remarks on ethics (RR: 31), there is a clear link to my concern with Smith's silence on evolution. What if there were no people, only animals? There would be no ethics, I gather, but wouldn't there be lots of mattering? There would be lots of survival and extinction for cause, lots of biological norms.<sup>5</sup>

**3.** *C++ and Searle on Programs.* First, I give Smith's essay an A++, and express my main objection: he should have written it twenty years ago and saved us all from a series of dubious battles that have gained precious few insights as by-products. But I also want to add to his concluding point 3 on Searle, about which a bit of clarification is in order: 'Searle's analogy of the mind to a program is misleading. What is analogous to mind, if anything, is process.' Smith adds: 'it is unimaginable that evolution constructed us by writing a program, a syntactic, static entity, which specifies, out of a vast combinatoric realm of possibilities, the one particular architecture that the mind in fact instantiates'.

In his uncharacteristically ill-considered *Daedalus* article on AI, Hilary Putnam (1988) speaks of the Master Program—which is perhaps the closest anybody has ever got to imputing this view to AI or to anybody. (See my critique in the same issue of *Daedalus*, reprinted in *Brainchildren*, 1998.) Smith is right in what he says, but let's see what this leaves available: the mind is, as Smith says, process (or a bunch of processes conspiring together), and while there is—need be—no programming language that specifies that family of processes in nature, that plays the causal role played by the source code in the genesis of new processes inside computers, those processes may nevertheless be usefully specified as if they were implemented programs. That is, to take the Searle case very much to heart, Searle has claimed that whatever consciousness is, it is not like a program in this sense: take a brain that is unconscious, and make it conscious by installing/implementing that program on it.

Now I continue to believe with all my heart and soul that this

is exactly what consciousness is! Consciousness is a set of behavioral competences that depend not so much on the organicity of the brain's neurons as on their global behavioral roles, so that you could in principle have live, healthy neurons by the billions subserving no consciousness at all—a comatose or otherwise utterly demented person—and you could turn that brain into the brain of a conscious person by 'simply' revising the behavioral microdispositions of those neurons, turning them into organelles and tissues that accomplished various 'computational' and 'communicative' tasks. In fact, when people recover from strokes, the resumption of various parts of normal conscious competence is very much a matter of the reutilization of healthy neurons to play new computational roles.

Moreover, of course, I've argued (and here is where my view is most radical, most embattled) that there is something that plays a causal role similar to that of source code in the genesis of much of this behavioral microcompetence: there are virtual machines that are installed by cultural imposition, learning, imitation, and memetic infestation, and whatever it is that hops from brain to brain is, in some no doubt hugely indirect way, a specification of a set of habits of thought. A bit like Java applets. Thus, you encounter Tetris, and find yourself executing shadow Tetris-moves involuntarily for some minutes or hours. Or you learn bridge, and find yourself putting yourself to sleep doing shadow-finesses, or you learn about agreement of adjective and noun in Italian and execute hundreds of agreement-checks...until it becomes second nature. The culture has driven a little rule into implementation in your head, and it is the same rule that all Italian-speakers have somehow or other implemented in their heads. Perhaps, to continue the analogy up to if not beyond the breaking point: native Italian speakers have the rule compiled in their heads, a much 'sleeker, more efficient machine' (as Smith says) than the interpreted version that still occasionally rises to the level of consciousness in my own operating system.

My point here is that nothing Smith has said about the non-compositionality of most executable programs casts doubt on the utility of such treatments. On the contrary, it helps mightily to clarify them, and to ward off likely misinterpretations.

## 6b — Reply to Dennett

Brian Cantwell Smith  
University of Toronto

Let me start by thanking Dennett for two things.

First, I'm grateful for the effort he has put into understanding this project—a project, I admit, that can seem a little like a fire hydrant: the content comes out in lots of different sprays. I learned from his comments, and that's great.

Second, I want to thank him for mentioning the issue he identified as number 12: domestication of the 'middle ground' opened up by all sorts of traditionally dichotomous theoretical distinctions. That focus on the textured intermediate territory, rather than on limit cases, is very important to me. I think of it as the philosopher's analog of *in vivo* rather than *in vitro* analysis. In my experience, people who don't appreciate the importance of this kind of middle-ground stance find it hard to hang on to, especially at first. It is a well-entrenched intellectual habit (especially in analytic philosophy) to think that theoretical rigor demands 'clear and distinct' ideas, even clear and distinct cases. But just as there are dangers of drowning in complexity and detail, so too there are dangers of excessive (especially formal) abstraction, particularly for subject matters—of which I think epistemology and ontology are instances—whose stuff and substance only emerges in these often messy middle regions. Doing such phenomena justice requires a distinctive theoretical style. Although hard to get, initially, this middle-ground approach is also hard to lose once you've got it—particularly when you see its not being appreciated all over the place. So I thank Dennett for noting that right up front.

Needless to say, I can't respond to everything he has brought up. Instead of giving a point-by-point response, I want to make six general remarks bearing on the issues he has raised. In conjunction with his comments, I hope these will clarify what is going on.

### 1 Naturalizing Ontology

The first remark has to do with the project of naturalizing ontology. ‘Why bother?’ asks Dennett. The main reason, of course, is because I believe the subject matter demands it. What ends up as a methodological commitment is grounded in an empirical claim: that the theory of ontology and the theory of representation and intentionality are about intrinsically interconnected phenomena. To study one without studying the other would be like studying time without studying space. Time is not space, of course; no one thinks they are identical. But you would not get an adequate account of either space or time by studying it on its own. So too, I believe, for representation and ontology. How things are and how we take them to be, though by no means identical, are co-constituted in intricate ways.

I might say that I haven’t always believed this.<sup>6</sup> During the 1980s I spent a long time trying to develop a theory of representation independent of ontology.<sup>7</sup> I was particularly interested in taxonomies of representational types (symbols, icons, descriptions, models, simulations, etc.)—a theory, I might say, in which isomorphisms figured.<sup>8</sup> Now I didn’t have the smarts to invent targets to do the work that representations couldn’t do. But my fundamental problem was that I couldn’t hold the ontology fixed—couldn’t stabilize it adequately—in order to develop satisfying accounts of the plethora of correspondences that held between them. I was unable to determine (except by fiat, which didn’t satisfy me) which items were objects or basic elements, which were properties of those elements,<sup>9</sup> and which were relations among them. Small variations in how I registered the basic domains wreaked havoc with how I ended up classifying the representations defined over those domains. In the end I was forced to admit that the (ontological) question of whether something was an object could not be answered except with reference to the (epistemological) question of whether it was being objectified by a representing or cognizing subject. That is: my independence assumption did not work. So there is a lot of failure behind this claim that representation and ontology are parts of the same subject matter. That really is the bottom line.

So I started over, to reconstruct ontology and representation together. It is not just an exercise, at the end of which you end up

with the same recognizable parts. The theory that comes out—the benefits it gives you—are different.

Perhaps the simplest benefit is that it gives you more resources to describe intermediate cases. The notion of feature placing, for example, turns out to be extremely broad and useful—and relevant, I think, to the issue Dennett raised about animals. The basic idea of a feature, which I take from Strawson (1959), is of something logically simpler than a property. Like properties, (concrete<sup>10</sup>) features are spatio-temporally instantiated, but, unlike properties, they do not involve a commitment to a discrete, individual, re-identifiable object, complete with unity or identity or individuation criteria, to serve as the exemplar or ‘holder’ of the property or feature or abstract type. Paradigmatic commonplace features are fog and other meteorological phenomena. The truth of an utterance of ‘it is raining’ requires only that there be raining going on ‘around here, about now’, as is sometimes said. There is no object to which the term ‘it’ refers.

Take another example. Suppose Dennett visits my California house, and on the second day remarks that the fog’s come back. ‘You are a philosopher,’ I ask, ‘has the same individual fog returned, or is it new fog, of the same type?’ I don’t know what Dennett’s answer would be; but I know what it should be: ‘Go away!’ Similarly, suppose you and I go camping,<sup>11</sup> and you, getting up early and looking around, stick your head back in the tent and say ‘It’s amazing; we’re camped right next to a whole ridge of mountains!’ Again, suppose I pedantically inquire, ‘You’ve used the plural “mountains”; just how many mountains are there?’ There is no reason to expect that an answer is possible. The problem is not epistemological: that you don’t know, that you can’t count; that you can’t see. Suppose the air is crisp, the view clear, and that we have all the time in the world. It doesn’t help. The point is that there is no metaphysical warrant, at least no metaphysical warrant up there on the ridge, for one answer over another. Criteria for mountain individuation simply don’t apply to such situations with anything like exact enough grip.<sup>12</sup> Similarly for a host of other examples. The point is simple: pre-theoretic philosophical intuition notwithstanding, much commonplace registration of the world does not require parsing it into discrete individuals.

Philosophers are a rarefied class; many of us, at least since our first course in model theory, have been persuaded that we do take the world in terms of discrete, reidentifiable objects exemplifying properties and standing in relations. Or anyway that that is the right idealization under which to pursue philosophical topics. I myself suffered under this misconception for many years. But I no longer believe it. (This is another of those things that are hard won. At first it is difficult to credit, but then, once you come to see that it is true, it is hard to imagine how you ever believed the traditional story: that it is a precondition for finding the world intelligible that you first parse it into discrete individuals.)

Feature placing is just a stepping-stone, of course. Adopting a richer ontological framework doesn't require the stronger thesis, that ontological facts are in part intentionally constituted. But examples of feature placing are useful because they suggest why that stronger claim is true. If pressed to supply answers to individuation questions in such cases (for example, to decide how many mountains there are), you will notice that the only way to do it is to make recourse not simply to the structure of the world (the details of the shape of the ridge), but also to the demands and contingencies of the projects you are engaged in. If we were committed to climbing all the mountains on the ridge, for example, that might affect our answers as to what distinguishes 'one mountain' from 'two.' If we were geologists, our answer might be different. Likewise, airplane pilots might arrive at judgment different again.

One common way to handle such variation in individuation practice is to claim that the word 'mountain' is ambiguous; that climbers use one sense, geologists another, pilots a third, and so on. But this strategy doesn't work. Senses multiply too profusely-varying per speaker, per occasion, per project. Eventually one is forced to admit that sense is indefinitely variable, and subject to factors anchored in the intentional projects of speakers. But this is an expensive admission: it reduces the 'multiple sense' proposal to no more than a relabeling of the original problem.<sup>13</sup>

In the long run, I believe, there is no credible alternative except to recognize that intentionality is implicated in individuation. Let me put it as succinctly as I can:

*The identification and reidentification of objects involves an epistemic process of abstraction over the infinitely rich (and often surpassingly messy) ur-structure of the world. Among other things, the normative character of the intentional projects that agents are engaged in, when they commit these acts of abstraction, figures in the resulting ‘clumping’ of the world’s effectively infinite detail. To be an object is to be a region or patch of the world that is successfully abstracted—where the issue of ‘success’ is tied into the normative conditions governing the dynamic project of which the act of abstraction is a constitutive part.<sup>14</sup> The fundamental character of (what it is to be) an object is thus intrinsically hooked into the intentional life practices of the objectifying subject.*

One more point on this topic. As a way to muster support for simply availing ourselves of ‘common-sense ontology’, Dennett says ‘Look, why not just assume sub-atomic particles and tables and mountains and galaxies, in the way that science does?’ This leads me to mention a radical thesis that I hold, although I can’t give it much defense here: namely, that science may not be committed to objects at all. Consider: an amoeba splits. Biology doesn’t care about the individuals in the situation: whether one amoeba died and two new ones were born; or whether we now have a spatial distribution of unitary amoeba-ness; or whether one of the two emerging amoebae is the original one, and the other one is new; or any other possibility. Another example: in California I own an ancient redwood tree that has clumps of very substantial shoots (some as much as 50 feet high) sprouting around its base. How many redwood trees are there? Science doesn’t know, and science doesn’t care. Similar conclusions hold for fog, for the units of selection, for a myriad other examples. What this leads me to believe is that scientific laws (like animals) may in fact deal only in features;<sup>15</sup> and that the objects we think of as constitutive of science may merely be simplifying epistemic devices that allow humans to calculate.<sup>16</sup> Objects in science, that is, are in my view properly understood as part of the epistemic apparatus involved in the conduct of science as an intellectual activity (on a par with mathematical models); they are not ontological commitments of the theory as a whole. [[Discussion point

5.9]]

Put it this way: ontology and abstraction need naturalizing as much as meaning, semantics, and content. Assuming a ‘standard ontological inventory’ for purposes of giving a naturalistic account of intentionality, as Dennett suggests, is thus a doomed project: it is viciously circular. Think about how appalled we would be (or anyway naturalistically unsatisfied) if someone were to propose a theory of representation that dined out on intentional notions, as if they were freely available. The naturalistic challenge is to explain intentionality without viciously presuming intentionality. A similar moral holds for ontology, in my view. Because ontological categories are in part intentionally constituted, attempting to explain representation while dining out on ontology is, for analogous reasons, fatally circular.<sup>17</sup>

## 2 From E&M to M&E

Second, I wanted to make a remark about the role of physicality in the metaphysical project.

I remember talking to Fodor once,<sup>18</sup> trying to convey my amazement that reference could point outside a speaker’s light cone. His response stunned me. ‘Look,’ he said; ‘it doesn’t matter what physics is like. Physics could be arbitrarily different, and it wouldn’t have a shred of impact on the theory of intentionality.’<sup>19</sup> It is hard for me to say how deeply I disagree with this sentiment. There is a sense in which I am something of a physicalist.<sup>20</sup> Not, mind you, a reductive physicalist-but someone who takes the character of the physical world to be essential in determining what intentionality is like. As a result, I take the consequences for a theory of intentionality of the structure of the concrete, material world to be enormous (as, I might add, must anyone who takes material embodiment seriously). The trick is to spell this out in a non-reductive way. Note that the issue is not merely one of engineering: that intentional subjects be physically implementable. The connection is much stronger than that. As I tried to show in OO, the structure of the physical world actually establishes the problem that intentionality solves (as well, fortunately, as supplying the wherewithal for its solution). [[Discussion point 5.10]]

For various pedagogical reasons, I take field-theoretic interpretations of physics especially seriously. I’m a complete amateur at

physics (as my readers will know), but for purposes of understanding intentionality, field-theoretic interpretations have a decisive advantage. They make it evident that physics does not involve a metaphysical commitment to discrete fundamental individuals. [[Discussion point 5.11]]

To see this, assume a field-theoretic interpretation of classical, high-school (Newtonian-Maxwellian) physics: spatio-temporally extensive fields of force, mass, charge, etc., subject to various dynamical regularities. And consider what is involved when we talk about individual bodies, as for example we might if we were to ask about the gravitational force exerted by this cup on this pen. In calculating the answer, we might be tempted to characterize the problem as a mass of 200 grams and a mass of 30 grams standing one meter apart. My point is simply that, as everyone knows, this characterization involves some simplifying idealizations. It makes two acts of abstraction over the raw fields: one to collect up the region of space-time we call ‘the cup’ into a dimensionless unity; the other to collect up another region, which we call ‘the pen’, into a similar dimensionless unity. That is, we objectify both cup and pen: treat them as discrete, individual, infinitely dense space-time points.

Why do we do this? For a very good reason: the simplifications are necessary in order to yield a problem that is epistemically tractable. Staying true to the field-theoretic interpretation would require treating the cup as a full three-dimensional mass density manifold, the pen as another three-dimensional mass density manifold, and formulating the question as one about the gravitational attraction between two solid regions. Setting up the problem in this way (that is, without any abstracting simplifications) would require an infinite amount of information. And solving the resulting problem (a double triple integral of point-wise gravitational attraction between two regions) would require an infinite amount of work. Neither, in general, will be feasible.

In sum: working with solid 3d regions, which is all that physics is really ontologically committed to, yields epistemically intractable problems. So we simplify, for purposes of calculation. That is where individual objects enter.

Once the distraction of individuals has been set aside, one can see

that the features of the physical world that most affect the nature of intentionality have to do with distance, coupling, and the locality of physical force (this has already come up in discussion, and I will say more about it in a moment). In particular, the 'point' of intentionality and reference, on my view, is to allow agents to be directed toward (ultimately, to care about) the world as a whole, beyond the (causal) limitations of that envelope with which, at any given moment, they are causally engaged. There are additional detailed connections as well. I mentioned an important one yesterday:<sup>21</sup> the differential character of physical regularities<sup>22</sup> engenders a kind of in-the-world deixis, which engenders an ineliminable indexicality in all representation and reference, which in turn underwrites the first-person qualitative character of phenomenological experience.

These are just a few examples of how I mean to take the physical world seriously.

### 3 Effectiveness

The third remark I want to make has to do with the relation of semantics to these issues of causation and local effectiveness. I subsume this under what I view as the problem of physical or material embodiment:

How can small patches of the physical world (for example, us) exploit a small fraction of the sum total of ways of being that the world supports (namely, that fraction that is causally potent or causally effective), so as to allow them to register the whole world (not just the part they are in or constituted of) as exemplifying an almost limitless variety of properties?

How, in other words, do we exploit a small fraction of the properties of a small part of the world to gain access to all properties of the whole world? This, I would argue, is the problem to which reference and semantics are the solution.

I hope this formulation clarifies my disagreement with Fodor. For if my characterization is right, then understanding the character of what I am calling the 'locally effective' (that is, those properties of local situations that can do causal, effective work) is essential to the project of understanding intentionality. And this for two rea-

sons. First, the effective properties are what an agent gets to use; they are the 'material', as it were, from which an agent can construct its intentional solution. Understanding them is thus necessary in order to understand how intentionality 'works'.<sup>23</sup> Second, we need to understand what these properties are not—that is, the vastly larger fraction (99%) of the world's features and properties that aren't effective, or don't hold of the local situation, and hence that the agent can't be coupled to 'directly', by physical coupling—since that is what constitutes the 'rest of the world' toward which the agent is intentionally directed.

One question that inevitably comes up, when I put things this way, has to do with the relevance of quantum mechanics. If intentionality is intrinsically related to physicality, then is the character of the intentional affected by the fact that the physical world is not ultimately classical? The answer may be 'yes', though I confess I am not prepared to say very much about this yet. To date, I have constrained my study of the locally effective to phenomena that, as far as I can see, could supervene on a classical base.<sup>24</sup> I've done this in part because I have yet to see any compelling argument that the human brain reaches further.<sup>25</sup> Even though I take my subject matter to be intentionality full bore—that is, intentionality in any possible material manifestation, not just its human projection—nevertheless, the sheer magnitude of human accomplishment convinces me that basing such an account on classical physics is not too severe a constraint.<sup>26</sup>

In the long run, though, I admit that the study should probably expand to include quantum efficacy. But the nature of quantum influence may be quite subtle. For example, one place where quantum mechanics may bear on the nature of human experience, at least indirectly, is in issues of long-distance coordination, of the sort that violate traditional locality constraints (for example, as characterized in Bell's theorem). Note that the fact that we can register the world, see things, think, remember the location of Dennett's house in Blue Hill, and so forth, is because there is a tremendous amount of long-distance relatedness in the world. Maintaining a (moderately stable) conception of the (moderately stable) world depends on this (moderately stable) relational regularity. Is quantum non-locality a necessary precondition for such long-distance regularity? In informal conversations, some physi-

cists have suggested that the answer may be ‘yes’. If that is so, then that is surely one way in which the human condition may be crucially non-classical.

Other than speculative questions of this abstract sort, however, I doubt that quantum mechanics has much to say about our middle-scale intentional lives. So I lack sympathy for writers (such as Penrose) who feel that in order to penetrate the mysteries of consciousness we need to understand mind in quantum-mechanical terms. In ‘Who’s on Third?’ (forthcoming a) I argue to almost exactly the opposite conclusion: that (again) using no more than a field-theoretic interpretation of high-school classical physics, one can see how the first-person, subjective, qualitative character of phenomenal consciousness must arise in any physically embodied agent that achieves an objective conception of the world around it. [[Discussion point 5.12]]

But return to the issue of simple effectiveness: how an agent can exploit what is effectively available to stand in for, care about, and otherwise direct it toward, that which is unavailable. This, in my view, is the best way to frame the question of intentionality. As you will predict, I ultimately locate the syntax/semantics distinction as a special case of this more general issue. But for pedagogical purposes, logic is not always the most illuminating place to start, in order to avoid being distracted with inessential aspects of the formalism.

In my undergraduate teaching, I start by studying clocks. I choose clocks for several reasons: (1) because they are familiar, (2) because they have a clear mechanism (clockworks), (3) because clock faces raise issues of interpretation and content, and (4) because clocks are so manifestly dynamic. Suppose we want tea at 4 o’clock. I assume that ‘4 o’clock’ is a non-effective property exemplified by passing metaphysical moments (one every 12 or 24 hours). If ‘4 o’clock’ were effective, it would be simple to build a tea-making device: you would construct a detector to respond to a moment’s exemplification of that property, and connect it to a switch. When 4 o’clock arrived, the detector would respond, the switch would flip, the kettle would boil, and out would come tea (or whatever). But of course—to make a point so obvious that we typically don’t realize how crucial it is—you cannot get a meta-

*physical moment's exemplification of the property of being 4 o'clock to turn a switch.* You can't do that because 'being 4 o'clock', as I keep saying, is a non-effective property. So what do you do instead? You construct a mechanism that uses properties that are effective, out of stuff you don't otherwise care about, and arrange it to be coordinated with the property that isn't effective that you do care about (a moment's being 4 o'clock). If the coordination is established properly, the former effective mechanism can stand in for the latter non-effective goal.

As Cummins just said [[in Discussion point 5.12]], one metaphorical way to understand this is to realize that intentional creatures have just a tiny keyhole through which to access the world. How far does their effectiveness reach? Because of proscriptions of locality, it reaches only to the surface of their skin. Strictly speaking, what is absolutely proximal—what impinges on your surface—is all you have to interact with. This is true of any conceivable physical agent: the infamous locality of physics restricts all engagement with the world to coupling with what is immediately present. Here we are, at this very moment, sitting in the living-room of an inn; our coupling to Dennett's farmhouse, even though it is only a few miles away, is at the moment very weak. So weak as to be 'undetected'. When we want to go to Dennett's place for lunch, we cannot be driven by effective coupling to it (as a Gibsonian might imagine we are directed to an opening in a wall by effective coupling). So what do we do instead? We arrange the situation so that we can be driven by things that are effective, in the here and now (such as maps), that will enable us to get us to his house, there and then. The dance that this strategy engenders-of exploiting what is local and effective in order to be directed toward what is non-local or non-effective-this is the phenomenon I am talking about. [[Discussion point 5.13]]

#### **4 Computation**

I trust that it is clear how this effective/non-effective dance relates to issues of mind, reasoning, and logic. My fourth remark has to do with its relation to computing. To explain this, we need to look at the history of computer science.

For almost a century people have been developing a so-called 'mathematical theory of effective computability'—or, as it is often

simply called, the ‘theory of computation’. In spite of its name, however, I do not think it is a theory of computing, because it doesn’t deal with computing’s essentially intentional character. Nevertheless, I still consider it an amazing achievement. What it is, I believe, is a mathematical theory of causality—that is, a theory of exactly what I have been talking about: physical effectiveness. This theory will not capture everything that matters about our pre-theoretic intuitions about causality, such as how you can cause things to happen after you’ve died. But that’s fine; scientific theories never exhaust the pre-theoretic intuitions on which they are founded.<sup>27</sup> What this theory does capture, in the long run, I believe, is what we will end up taking physical effectiveness to be.

So the ‘theory of computing’ supplies half the intentional story: the effective half—the part about what you can use, what you can do, what works, how hard it is to change one physical arrangement into another.

A brief historical caricature may help explain why things developed this way. At the end of the nineteenth and beginning of the twentieth century, following the impressive achievements of the Industrial Revolution, there was a tremendous sense of the power of machines. Some of these machines were bluntly physical (steam engines). Some were targeted at very specific material concerns (smelting iron ore). Clearly, however, some very useful mechanisms, such as clockworks, weren’t so concretely specified. Although it was crucial that they be physically constructed, it didn’t matter what specific materials they were made of.

People realized, from examples of this sort, that if you want to know ‘what can be done’, you can abstract from purely physical considerations—how big the mechanism is, how much energy it uses, etc. Paradigmatically, such mechanisms are used for detection, tracking, and other (at least inchoately) intentional tasks. Suppose you want to know when some particular train passes a spot on a railroad. You might put a sign or indicator on the train, and install a detector next to the track, to signal when the train passes. Sign and detector will obviously have to match, in physical characteristics, so that the latter can respond to the presence of the former. But beyond this, there are no requirements on what they should be made of, how big they have to be, etc. Because, of course, all we really need is to detect one bit of information: ‘the

train is here’.

Many such mechanisms can be imagined, of an essentially physical yet ‘multiply realizable’ sort, from simple detectors, to clocks, to what has seemed like the most powerful mechanism of all: one that could calculate, reason, do mathematics or logic. What people realized (to continue this glib story) is that, to get a theory of such devices, you have to let go of specifically concrete concerns. So what did they do? They went to the other extreme, and considered devices as completely abstract. Since what could be done (for example, by Turing machines) seemed not to have to do with specifics about particular materials, the theory took the opposite pole, and assumed that what could be done had nothing to do with materiality at all. This is why the theory of effective computability is framed as if computability were a purely abstract notion.

We are still living in the shadow of this history. The idea that the fundamental results of computability theory might be anything other than completely abstract is far from universally acknowledged. Many academics treat theoretical computer science as a branch of logic or mathematics.<sup>28</sup> And challenging this assumption bends some people out of shape.<sup>29</sup> Nevertheless, I believe, helped by people studying the powers and limitations of quantum computing, it is going to become increasingly apparent that computability limits are fundamentally material.

In terms of long-range intellectual trends, in other words, we moved from the completely concrete steam engines of the nineteenth century to the completely abstract inaccessible ordinals of the early twentieth century. Now, on the verge of the twenty-first century, we are settling somewhere in the middle. But through it all, the study has been a study of mechanism-of what can be done by concrete, material processes. As I say, I still don’t think it is a theory of computing; real-world computation involves relationship, semantics, non-efficacy. But even if it doesn’t explain our main subject matter, a theory of pure efficacy is a phenomenally important intellectual project, for which I have the greatest respect.

What is distracting—the reason this is not all universally realized—is that, because of its history, the theory of computability is still framed in semantical terms (computing functions, coming up

with answers, representing numbers, etc.). Thus consider the standard practice of taking marks on Turing machine tapes to denote numbers. Though historically comprehensible, this practice, I argue, is actually wrong. The marks don't denote numbers, in spite of what everyone thinks. Rather, the numbers denote the marks. Computability and complexity theory, in my view, are mathematical models of complex configurations of marks. All the regularities captured in the theorems have to do with these marks and their arrangements, not with the numbers we associate with them. Why this is true, why you have to understand it this way—that is a story I can tell you over drinks.<sup>30</sup> The bottom line is that the role of the mathematics, in computability theory, is just like the role of mathematics (and of objects!) in physics: it's a classificatory, epistemic device, employed by theorists. Like all theories, the theory is semantic, but it is not semantical;<sup>31</sup> it is not about semantics. What the theory does is to use semantics (not just terms, equations, variables, etc., but also mathematical modeling relationships, like physics) to classify concrete, in-the-world, non-semantic regularities.

Thus consider the results of computability theory, such as the unsolvability of the halting problem, or the difficulty of factoring products of large primes. Both problems are framed semantically: that you can't decide whether an arbitrary machine will halt, on an arbitrary input, that you can't figure out what numbers are prime factors. As any good theoretician knows, however, if framed in purely non-representational terms, as issues of yes/no decisions, or of pure numbers, these problems can be solved, trivially, if you employ what are called 'non-standard encodings'. For example, if you represent numbers as lists of their prime factors, then factoring them takes no work at all! Given this vulnerability, which applies to all complexity results, my argument has three steps: (1) the only way to bar such non-standard encodings is by bringing into explicit view constraints on the representations (constraints on the marks), not just on what the marks denote; (2) once you bring in the minimal constraints on marks needed in order to preserve the theorems, you have brought in everything you need; there is no more work for the 'denoting numbers' aspects of marks to do; and (3) what happens, in the traditional practice, is that these entirely concrete constraints are implicitly

modeled by numbers, the numbers that the marks are taken to denote.

But enough technicalities; this is not the place for details. The point is merely that what is today called the ‘theory of effective computability’, in spite of the way it is framed, has nothing to do with semantics. It is a mathematical theory of physical effectiveness, pure and simple, of exactly the sort that we need for half the intentional story. It is a mathematical theory of Cummins’ keyhole. [[Discussion point 5.14]]

Before concluding this topic, I should admit one thing: how much work is opened up by the reconstruction I am proposing. If I am right that complexity theory is really about the capabilities of pure mechanisms, independent of semantic interpretation, I am committed to reformulating its results in non-intentional terms. The theorems cannot be framed in terms of decisions, or numbers, but as statements about how certain configurations of the world (that is, certain machines), if started off in given effective arrangements, will or won’t or can’t get into other effective arrangements, or about how, if you give a machine two different input marks, sufficiently complicated, these inputs will essentially ‘drown’ the machine, so that it won’t be able to produce one kind of output mark from one, and another kind of output mark from the other. In other words, I am committed to reformulating all the theorems as claims about effective arrangements, simpliciter, without regard to anything those effective arrangements mean. [[Discussion point 5.15]]

## 5 Objects

Fifth, I want to say a few words about objects—the subject matter of the book (*On the Origin of Objects*, 1996).

If you take logic, or introductory philosophy, you might think that there are two fundamental kinds: (1) concrete, individual, particular objects (called tokens, if they are linguistic or semantically interpretable); and <sup>(2)</sup> abstract, perhaps Platonic, types, which the individuals instantiate. In the general case, a type will have multiple instances: there will be a one-to-many relationship between types and their instances. In this sense, the abstract type acts to ‘bind together’ what is similar across the (extensional) set of objects of a given kind.<sup>32</sup>

One of the things that a career in computing has given me an extraordinary appreciation of is the sheer complexity of real, in-the-world, material objects. Among other things, this has in turn led me to appreciate the profound inadequacy, as an account of reality, of this simple picture of types and their instances. There is nothing magic about computing, in this regard: librarians know it too, in their efforts to catalog copies, editions, translations, reproductions, templates, and so forth. But I came to the lesson through computing.

Here is just one example of the kinds of practical issue that drove me to the story I report there. I normally download my e-mail to the hard disk on my home computer, where I have several hundred megabytes of files, folders, pointers from files to folders, and so on. When I travel, I copy the whole mess onto my laptop. Suppose I come here to Maine, dial in to a local ISP, and get a message from Dan Dennett. Intending to file it, I follow a link I have set up to the folder reserved for messages from him. This would have worked fine, at home. But when I try it here, on my laptop, a system message is displayed, asking me to mount the hard disk I left in California. Why does it do this? Because I copied the pointer from my home machine onto my laptop, and on my home machine this pointer pointed to the copy of the Dennett folder stored on that disk. Of course my laptop has a copy not only of this pointer, but also of the folder in question. The problem is that the system wasn't smart enough to know that the pointer should be adjusted to point to the copy of the Dan Dennett folder that now resides on the laptop.

What I hope this tiny example shows, or at least evokes, is the messiness of real-world issues of concreteness, abstractness, and multiple versions of 'the same thing'. Pointers are normally taken to point to individual files, but my intent, for this pointer, was that it point to something slightly more abstract: the Dan Dennett folder, of which I have multiple copies. You could say that this abstract Dan Dennett folder is a type, of which the individual copies are instances; but other than dressing the situation up in formal guise, that move doesn't much help. The problem is that even an ordinary desktop contains an astounding proliferation of highly related objects, of various sorts, many of which stand to each other in analogous one-to-many or many-to-one relations.

Copies, virtual copies, pointers, caches, back-ups, editions, versions, replications, and so on—seemingly without limit.

Similar issues arise inside programs. Suppose you call a subprocedure with a matrix as an argument. And suppose the subprocedure changes the matrix. Was the original matrix changed, or did the subprocedure modify a copy? It depends on whether you passed it, as they say, 'by value' or 'by name'. Some other examples: one variable, multiple values; one IP address, multiple CPUs; one procedure, multiple call sites; one program, multiple copies, each of which can be run multiple times; one web page, multiple servers; one web page, multiple translations into different languages. And so it goes. Templates generate multiple copies, generators spawn new instances every time they are called, etc. And when the proliferating objects are interpretable, the situation gets even more complex. For example, there are problems of context-dependence: something that means one thing, in one context, can turn up in another context, or a copy of or pointer to it can turn up in another context, and mean something different (the Y2K problem is one especially famous example of context-dependence gone awry).<sup>33</sup>

Needless to say, if you work with these systems, you have to keep things straight. Some properties (the number of messages I have received from Dennett, say) hold of the abstract 'one' of which there are multiple instances or versions or copies. Some properties differ across each member of the group, in systematic ways (as we will see, this has to do with indexicality). Other properties (such as the location of a file on disk) may differ across the 'many' in no systematic way at all.

You might think that the way to avoid confusion would be to be extremely, even aggressively, clear—always knowing exactly which object type you are referring to. You might even want to have different names (for example, to distinguish the program, considered as an abstract object, from concrete copies of the program, from temporal runs of the copy, and so on.) Let me simply report that all attempts I have made at being extremely clear in this way have failed miserably; they drown in inscrutable complexity.

Humans apparently handle such situations in a very different way. They seem to have a feel for the sort of things different

properties can hold of, and to infer the appropriate instance or entity or individual for any given property in question.<sup>34</sup> In a way, you can tell that the term ‘program’ refers somewhat differently, in different cases, because a kind of zeugmatic infelicity arises from combining different types of reference under a single conjunction: ‘Is that program recursive and corrupted?’ sounds ‘off’.

For many years, behind the scenes, I have been trying to develop a calculus in terms of which to understand this kind of proliferating objectification practice. I call it a ‘fan calculus’—a calculus of the ‘one’ and the ‘many’—because so many of these situations involve one thing (what I think of as the point or root of the fan) that devolves or engenders or creates or spreads out into or is exemplified by or holds of multiple copies or versions or instances or tokens.

The classical type-instance distinction is a single fan, on this generalized scheme: what we call the ‘type’ is at the point of the fan; the (extensional) instances constitute its fringe. Even in the case of simple language, however, it is clear that a more complex classification is needed. Yesterday, in a discussion about indexical utterances, I made a three-way division, among type, token, and use. [[See Discussion point 4.2—ed.]] Very roughly, you can think of one fan connecting the type to each different token, and another fan connecting each token to the set of all its uses (if it is used more than once). To see the utility of this double-fan characterization, note that the mentalese word ‘T’ is indexical on the first fan, but not on the second. Your mentalese inscription of ‘T’ (if there is such a thing) and my mentalese inscription differ, systematically, in their referents. But unless I am deranged, all my different uses of my (single) mentalese inscription of ‘T’ refer to one and the same enduring individual.

In the case of types, tokens, and uses, we typically think of the types as abstract, the tokens as enduring and concrete, and the uses as concrete events. What is interesting about the proliferation of computational examples I cited, including files and copies and versions and editions and templates and copies of templates and generators and so on and so forth, is that much more complicated hierarchies seem to exist in which all the entities are apparently concrete. (Whether that is actually true, however, is not so easy to say. As the case of the Dan Dennett mailbox indicated,

there may be slightly abstract unities in terms of which some regularities hold—whatever it is to be ‘slightly abstract’.)

I mention all this only to say that considerations of this sort, involving complex relationships between ‘one’ and ‘many’, have influenced the account of objects (or perhaps I should say of objectification) presented in OO. One question that is of considerable interest, as I have already indicated, is how issues of concreteness and abstractness play out across these hierarchies. And as usual my answer is the predictable one: the most interesting cases, I believe, are somewhere in the middle.

If one is strict about laying out the hierarchies, moreover, intriguing patterns emerge. Even the simple case of a proper name, standardly described as ‘one name, one object’, involves two fans, as indicated in Fig. 5.1: one spreads out from the name qua type to its various utterances or uses; another fans back in from these different uses to a single person. Similarly, the fan structure of the mentalese indexical described above is given in Fig. 5.2: one fan from type to tokens, individual fans from each token to its uses; and finally another set of fans back in from those uses to the holder of the token.

With respect to the overarching project of naturalizing ontology, perhaps the most important observation is the following: there are similar patterns of complex fan-ins and fan-outs underneath or ‘within’ the notion of an object (as above objects, having to do with particulars and classes or types). Suppose we lay out any given concrete object in 4-space, as an extensive space-time worm. If I touch you, in a certain sense my hand will touch one part of you, say, your shoulder, at just one point in time. In saying that I have touched you, not just your shoulder, and not just now, I am saying that the touch, as it were, ‘fans out’ across space, to all of your body, and also across time, to make contact with you as an enduring individual. Suppose I touch you again, ten years from now, and for some reason ask ‘Have I ever touched you before?’ The positive answer that this question warrants can only be defended by noticing that the two spatial and temporal fan-outs end up being coincident, on one and the same enduring object envelope—an envelope, as I hope is obvious, that bears some resemblance to the result of the act of abstraction that we talked about earlier, in the case of the cup or the pen.

What I find intriguing is that so many cases, from the epistemic structures of reference to the ontological structures of individuals, involve various kinds of abstraction: ‘gathering up’ of a bunch of that which is in some ways different, and taking the result as a unity-as that which is one. Getting to the heart of this practice is an essential part of the story I want to tell. Moreover—and in a sense this is the heart of the meta physical story—my ultimate claim is that there is no technical way to deal with the stunning complexities of these interrelated fans except by a single, integrated account that makes simultaneous reference to the aboriginal structure of the world and to the normatively governed intentional projects of the objectifying agent. One can only make sense of these structures, that is, via a ‘blended’ epistemological-cum-ontological account.

Finally, let me say a little about the subject Clark has brought up (in conversation): the role of indexicality and deixis in all this, and how that ties into issues of effectiveness (as you know, I have a claim that the intrinsic indexicality of reference stems from the fundamental character of physical law). Consider a single person’s multiple utterances of the word ‘now’. And assume, again for simplicity, that each utterance is used to refer to instantaneous moments, so that a sequence of utterances (‘now! now! now!’) would be used to refer to a corresponding sequence of (very short) passing moments.

What is evident in this case is that the referential pattern involves something I call ‘point-to-point correspondence’. One concrete use refers to one concrete moment; another concrete use refers to another concrete moment; and so on. Moreover, the temporal sequences line up. As regards the link between deixis and physical effectiveness, the fundamental insight is very simple: physical interactions have exactly the same point-to-point correspondence structure. What takes place now affects what is happening now; what took place then affects what was happening then; what will take place next time will affect what is happening next time. In effect, this point-to-point correspondence (both spatial and temporal) is intrinsic to the structure of physical law.<sup>35</sup>

What happens when we objectify is that we gather a region or patch of the world into a unity. To do that requires extending

these patterns of correspondence from simple point-to-point relationships, of the sort that underwrite physical regularity and simple indexicals, to much more intricate and hierarchical fan-ins and fan-outs. Feature placing, of the sort I described earlier, involves more complex forms of correspondence than simple point-to-point, but simpler than what is characteristic of the exemplification of properties and relations by full-fledged objects. To get to these more sophisticated cases, one needs to start getting involved with types, instances, and so forth, which involve complex, cross-cutting fans.

For example, suppose that last week I thought about Clark, and made a mental note to ask him, when I saw him this week, how his wife Pepa is doing. As indicated in Fig. 5.3, a complex set of fan-ins and fan-outs governs this successful ability to refer to Clark as an enduring unity. Both Clark and my mental token are temporally extensive, for starters. However, unlike the case of 'now! now! now!', and (similarly) unlike physical connection, the two temporal sequences don't line up, point by point. On the contrary, it is essential to the logic of the situation that when I thought about Clark, last week, at time  $t_1$ , I didn't think only about the temporal slice that was him-at-time- $t_1$ —that is, about him then. Rather, I thought about him, which is to say, I thought about him as a temporally extended person (the whole lower space-time worm). By the same token, when I met him here at the workshop, at time  $t_2$ , the person that I greeted and talked to was again not the temporal slice him-at-time- $t_2$ , but the same complete temporally extended individual. So my mentales token fans out into individual uses; each of which fans out to cover the whole singleton person. In order for the reidentification to work, they must be coincident in that temporal extendedness.

Reidentification, in sum, requires this kind of cross-cutting gathering up and spreading out. This is in sharp distinction to the vastly simpler point-to-point correspondence that is true of all physical interaction, and that is true of at least limit-case indexicals.

Needless to say, this is just a whiff of a picture. What I am really doing, I suppose, is diagramming the field-theoretic structure of simple reference to concrete individuals. My point is only that if we are serious about our naturalism, something like field-

theory<sup>36</sup> is all we have to start with. Somehow or other, we objectifying creatures are able to do a sophisticated enough dance to parlay our simple, effective, local, point-to-point field-theoretic coupling with our immediate physical surround into these complex patterns of cross-cutting fan-in and fan-out that characterize objective reference to the world. I don't claim to understand more than one per cent of how it goes. But it is something that I want to figure out.

Enough about objects. I have just one more general remark to make. [[Discussion point 5.16]]

## 6 Evolution

Sixth and finally, let me say something about evolution. In brief: I don't want to accept evolution as a rock on which to build my church. But I will accept it as a flying buttress-as something that supports my church from the outside.

I agree with lots of things that Dennett says. For example, I think he is right to say that nothing but evolution could have gotten us here, could have done all this work. If evolutionary biology is right, which I presume it is, then sure enough, evolutionary adaptivity must be the means by which we learned to register the world. The causal history (of the emergence) of our registrational capacities, the causal history of our opening up into normativity, the causal history of how and why we take the world to be significant-all these causal histories undoubtedly unfolded along evolutionary paths, especially originally (for the last 10,000 years, social and cultural and political histories have presumably carried more of the developmental and explanatory weight).

Dennett is also right to suggest certain corrections of emphasis. Of course it is true that what we want to be intentionally directed toward is what matters to us, not what we used to be connected to. I certainly don't want to be sentimental about primordial or aboriginal union.<sup>37</sup> I also agree that processes of connection and disconnection, processes of registration of that to which one isn't (and maybe never was) connected, and so on and so forth-these constitute an Extremely Good Trick, which evolution discovered and exploited.

Finally, let me say that I am completely open to being instructed, as regards the details of evolutionary affairs. I am no ex-

pert, and look forward to knowing more. Moreover, I feel ready for the handshake. I believe that the constructive tenor of my account ('constructive' in an engineering sense), starting from very simple patches of the world and progressing up through mechanisms of simple non-effective tracking, to featural registration, to full-blown conceptual registration in terms of objects and properties, to a form of situated objectivity—this story of 'building up' registrational capacities seems to me very compatible with evolutionary history. And I look forward to understanding better how such capacities evolved.

If I agree with all these things, then why don't I talk about evolution?

The answer is essentially this: while I will admit that evolution is the means by which we learned to register, *I don't believe that registering is itself an intrinsically evolutionary process*. The explanation of how we came to do it may be evolutionary, that is, but the explanation of what it is we do is not.

To go back to Cummins' metaphor, I think of registration as a way of exploiting the 'keyhole' of what is effective, so as to end up being oriented toward what matters, including (largely) what is not effective—in order to be oriented to the world, to take the world as mattering. I see no reason to believe that this is an intrinsically evolutionary thing to do; that non-evolved creatures could not do the same thing. Moreover, I do not believe—and I suspect this may be a point on which Dennett and I substantively disagree—that the fundamental normativity on which registration rests is intrinsically evolutionary, either. Again, that we take the world to matter may have evolved. But that the world matters is not by itself an evolutionary claim. If creatures were magically placed here, or emerged via a different means (say, on another galaxy), the world could matter to them—just as much. [[Discussion point 5.17]]

So that's the claim: that this intricate, sly, surreptitious strategy for exploiting what's effective in order to stand in relation to what's not effective is a Phenomenally Good Trick. It is a trick on which evolution stumbled, and made out like a bat out of hell. It is even clear why it is so evolutionarily useful. So there must be an evolutionary story about how it was discovered, what kinds of

registration it evolved first, how it worked, what advantages it conveyed, what was tried and failed, and so on and so forth. In order to make good on the claim that it is an Evolutionary Good Trick, however, the character of the trick cannot itself be defined in evolutionary terms. In order of explanatory priority, registration must be understood antecedently to evolution, if we are to understand how it is a trick—a trick that the universe supports, that evolution could pick up and make out with. [[Discussion point 5.18]]

To make this concrete, let me talk a little more about norms, because I think norms are the place where the issues become most urgent.

In RR, as Dennett pointed out, I made a distinction between ‘statical’ and ‘dynamical’ norms.<sup>38</sup> Think about normative notions in formal systems—the norms we apply to processes of inference in logic, for example, and analogs in such related fields as economics and game theory.<sup>39</sup> What we traditionally encounter are norms such as truth preservation, inference to the best explanation, utility maximization, and so on. What is interesting about these norms is that, while they apply to processes (and hence, in my terminology, are dynamical), they are defined in terms of explanatorily prior norms, such as truth, reference, explanation, utility and so on, that are defined on states (and hence are what I call statical).

My point in RR was that an extraordinarily important intellectual shift is taking place, across the intentional sciences: the explanatory order is changing. For a variety of reasons, ranging from evolutionary considerations, to the sheer difficulty of characterizing such age-old statical norms as truth and reference, to concrete experience building and maintaining and using computer systems, people have come to realize that the only way to define norms such as truth and reference in a useful and non-question-begging way is to base them on how things are used. In a sense, we are all becoming closet Wittgensteinians. It is not so much that statical norms are being discarded (truth, function, utility, etc.), as that they are being understood as derived from dynamical norms.

The reason why this shift in explanatory priority from statical

to dynamical is so important is that it puts pressure on a source of dynamical normativity. If you are going to define your statical norms in terms of your dynamical norms, then you cannot define your dynamical ones in terms of the statical ones, on pain of circularity. You need something else as a source of dynamical normativity. What is that going to be? Especially for a naturalist, this is a very urgent question: what naturalistically palatable source of dynamical normativity is available, on top of which to construct the entire normative edifice necessary for full-blooded intentional characterization?

This is the role that evolution is playing, I believe, in many of the intentional sciences. Evolutionary advantage is an extraordinarily convenient candidate on which to rest an intentional story. Ruth Millikan and the general project of teleosemantics can be seen as one example of this trend. And I take it that Dennett is proposing something similar, when he says all normativity rides on Darwin's coat-tails. He is basically asserting that evolutionary advantage is the 'mother of all norms', the ur-dynamical norm.

Unfortunately, I don't believe evolution is a strong enough base on which to rest all of human normativity. For think about how much the general issue of dynamical norms includes. Ethics, for starters; and even more generally, how to live. I cannot muster arguments here, but I simply don't see how one could milk evolutionary survival for this full range of normativity, for all that has inspired individuals and cultures, led people to distinguish good from bad, and so on and so forth, over the ages. Remember: I am not denying that human normativity may have emerged evolutionarily, at least at first (that is, to the extent that there was substantive normativity prior to the development of culture and civilization-something on which I have my doubts). All I am denying, to repeat my standard refrain, is the claim that normativity is an intrinsically evolutionary notion. [[Discussion point 5.19]]

So I have said that evolution is not strong enough to be a basis for all substantive norms. Clark and Dennett ask [[Discussion point 5.19]]: 'Is it strong enough to give rise to them?' The answer depends on what you mean by 'give rise to'.

If by 'give rise to' you mean causally, historically, how did these norms emerge? What engine could have done all this work to get us here? Isn't this the only mechanism through which normativity

could have emerged? My answer to that question is ‘yes’. Evolution may have been the train on which norms arrived at our present station. And as I said before, you can see why it would have arisen, evolutionarily: the ability to care, to register their world, to take things as significant, confers a huge evolutionary advantage.

But if by ‘give rise to’ you mean (as I do) something more metaphysical, something more explanatorily substantive—and something more like the word ‘origin’ in the title of my book—then my answer is ‘no’. Even if it is an a posteriori necessity that, given basic material facts about the universe, about the origins of life, etc., evolution is the only means through which normativity could have arisen in registering creatures, nevertheless, it is not constitutively intrinsic to our normativity, I claim, that we have an evolutionary history. Remember: I want to be able to explain what normativity and registration are, such that we can say of evolution that it stumbled on them. But the fact that evolution stumbled on them merely makes evolution the implementing mechanism whereby we came to be normative.<sup>40</sup>

In detail, I should admit, there is undoubtedly tons to be learned from our evolutionary history—about what kinds of normativity there are, what kinds we have evolved to be attuned to, etc. And (perhaps even more so) there is tons to be learned about what our registrational capacities are, what constraints they have evolved to satisfy, and so on. All those things are interesting and useful. I just don’t want to allow the interest and urgency of asking those questions to seduce us into what I think of as a kind of non-reductive causal foundationalism: wherein we confuse what things are with the causal history of how they came to be. [[Discussion point 5.20]]

It may help, in understanding this, to think of registration first, and normativity second. Remember what I want to say about registration: that it is a mechanism whereby you arrange things so that you can track long distances, put together machinery and external signs and external scaffolding and so on and so forth, so as to stand in relationship to more and more and more-leading, ultimately (this is where consciousness and objectivity merge) to a profound and care-full orientation to the whole world. Do we need evolution to understand that? In a constitutive sense, I think not. In an historical sense, I think we do. It

doesn't seem to me that evolution plays a constitutive story in understanding how the strategies of exploiting local effective structure can stand you in relationship toward that with which you are not physically coupled (though I admit: the hardest issue—the nub of the matter, as it were, and maybe for that reason an appropriate subject for drinks, later-is what 'being oriented to' really means, in all its respectful, concerned richness and depth). For now, let me just say that I believe the same thing about the norm-side: that an evolutionary account may tell us how we got here, and may tell us in detail what the costs and trade-offs are, in regards to honouring norms, but it won't give us a constitutive account of the structure of normativity itself.

It is time to stop; but I hope that this makes it at least a little bit clearer why I want to accept evolution as the train, but not confuse it for the goods that were transported on that train; to accept it, as I said at the beginning, as a flying buttress, but not the rock on which our registration of the world is founded.

## 6c — Dennett & Smith · Discussion

**CUMMINS:** This is what anti-realists do on purpose.

**DENNETT:** Commit this inscription error?

**CUMMINS:** Just, lots of them.

**SMITH:** Here's a simple case, to see what's going on. Suppose you want to know if your cat recognizes you as an individual as opposed to as just more 'Hughness'. And suppose people say, 'Yeah, look, because here's its neurophysiology and this cell lights up every time—same cell, so it must be the same person.' But how do you know it's the same cell? Maybe it's just reinstantiating some type in there. How do we decide that's a second use of a single token, as opposed to new instantiation of a single type? See, when you say 'one cell, so one individual', you're making one assumption about object identity in the head of the cat, and then assuming that you can use that identity to warrant a claim about object identity in the content of the cat, in what the cat's identifying. But I can redescribe the situation in the brain, and then get the presumptively competing suggestion about the content. All you're really doing, in other words, is piggybacking your analysis of content identity off brain identity. We should worry if our analysis of whether the cat recognizes Hugh as a type or an object depends on empirically equivalent ways of us theorists individuating its brain, especially since the cat itself doesn't individuate its brain at all.

I think these things happen very subtly, even in modest cases. Earlier today we were talking about whether we objectify things on a map. And it depended on how we registered the map in the first place. Take a line: do we call that a relationship between two points, where the points are the objects? Or are the lines the objects and the points just relationships between two lines-where they cross? Problem is, in one case you end up saying we're objectifying the line, and the other case you don't. It's that kind of thing. How we as theorists register the problem domain affects our analyses in ways that are stunningly consequential.

**CLARK:** So it's a worry about the baggage that comes along with the

labels. Like all those worries people have about how you label the nodes of your semantic network.

**DENNETT:** So, good. You can go way, way back to Drew McDermott's old paper 'Artificial Intelligence Meets Natural Stupidity' [1976], and that was a sort of ur-anxiety about inscription errors which Brian has generalized.

**SMITH:** In a way, I think what I might say is: Look, inscription like this, it's something you have to do. It has enormous consequences; so you want to be tremendously modest and humble and cautious. It's not black and white; it's not as if you can say 'Here's an inscription error, here's not.' You always have to inscribe. The issue's just this: don't let the fact that you have to inscribe license you to project all sorts of ontological assumptions all over the subject matter without taking responsibility for them.

**CLAPIN:** OK, so that's Rob's point. That's what anti-realists maybe are doing, is that they are allowing a license for them. They say, well we've got to do this anyway, so there are no constraints.

**CUMMINS:** Yeah, other than internal coherence.

**CLAPIN:** So the way you're describing color, Brian's kind of saying is true for...

**DENNETT:** Objects.

**SMITH:** Everything. To think that taking human relativity seriously implies irrealism is only true on a Cartesian view that we're not part of the world. But if someone were to write a book called *Being There*, or something like that, and actually realize that we are here, then from that (correct!) point of view, human relativity shouldn't be metaphysically scary. It is profoundly consequential, but it's not skepticism or irrealism.

**DENNETT:** I think it's like this. I think we can talk about the colors of things on distant galaxies by helping ourselves, and knowing we're doing that, to human color vision, and using that as our standard. And Brian's saying we can talk about primordial objects by using human object vision and recognizing that that's what we're doing, as long as we keep track of the fact that we're using human object and property vision as our standard, as our perspective, we can sort of discount—sort of like discounting the illuminants. That's the idea, and I'm snowed by it, whether I should be or not. I sort of like it, so, we'll see.

**HAUGELAND:** Is this also an example of the same point—I can't remember what I've thought about this in the past—that, before the human race evolved, the moon was 240 thousand miles away, even though there weren't any miles then?

**CLAPIN:** In that kind of an example, is the thought to go sort of de re-because miles didn't exist, *de dicto* how many miles is just not a sensible way of talking?

**HAUGELAND:** Well the trouble with that way of putting it is that it supposes that de re just is as it always was.

**DENNETT:** But thank you for raising it, Hugh, because that helps with one of the reasons why I'm attracted to this: because it helps me fend off that awful de re/de dicto stuff.

**SMITH:** Also, it will have more bite when one realizes that taking the world in terms of objects and properties is underestimating the world. That's when the approach really starts to fight back-not in a way that Rob will swallow, perhaps, but in a way that is at least akin to something Rob would swallow.

**HAUGELAND:** If everything were rigidly blocked in the universe, there were no flex and slop and slippage, nothing could be out of touch with anything else'-that seems to me to be completely wrong. What is rather the case is that there couldn't be a distinction between being in touch with and not being in touch with. You couldn't be in touch with anything else specifically if moving this moved everything.

**SMITH:** Well I couldn't be some thing. I don't disagree. I think the very language we speak so presumes that there is a certain looseness in the world that our attempts to describe what it would be like if there weren't a certain looseness will all fail. But they don't fail completely. We can have this conversation which we're having, and it makes some sense. I actually think it's impressive that we can agree, 'Yeah, all these things fail, but, we actually do get a sense of what we mean.' We can actually reach a kind of consensus which, if we all said, 'OK, now we've got it; let's try to say it,' we wouldn't do any better than we just did.

**CLARK:** So it's not essential for there being objects that the universe not be connected like that?

**SMITH:** Yeah, it is, because in order for it to be an object it's got to have distance. And it's got to have shear.

**HAUGELAND:** It's got to have distance and it's got to have an internal life that is different from what's around it.

- CUMMINS:** This was Cartesian physics. Everything was locked together. It's just that minds were allowed to slop around a good deal-but the physics was all locked.
- SMITH:** Well dissipations of forces are tricky. The problem with Cartesian physics is that it isn't theoretically precluded that you can tell everything from just how one little thing is vibrating.
- DENNETT:** That's just what Newton saw. This is right. You're sort of recapitulating a Newtonian revolution here with your point about flex and slop.
- SMITH:** One thing that might help explain it is that I don't think space-time points are echt individuals. Other than particularity, I actually think they lack all of the characteristics that individual objects actually have to have to be objects. Think about the field theoretic interpretation of classical physics which is actually doing some work in this story. Imagine a rubbery manifold with forces going up and down, and all that kind of stuff. You can imagine that if everything were just spatio-temporally infinite manifolds, there could be lots of space-time points as it were, but there would be no clumping of them together into reidentifiable individuals that have heft, size, or separateness, and so on. I think that field theory, this rubber manifold stuff, is probably the best imaginative route in. By particularity I just mean roughly spatio-temporally concrete occurrence.
- CLARK:** So it's like saying particularity is not objecthood. They're not objects, are they, space-time points?
- CUMMINS:** Cartesian points don't move around, although they do have properties.
- SMITH:** That's right. They don't move around, because they don't do anything-that's the problem with them.
- DENNETT:** A nice way of thinking about it might be to think about the individual cells in Conway's life world-they are particular but not individuals. But a glider can be individual.<sup>41</sup>
- SMITH:** I think if I were to write the book over again I might have said either 'occurrent' or 'concrete' instead of 'particular'.
- HAUGELAND:** Well, I think what you want is concreteness. I'm not sure what you mean by occurrent, or do you just mean actual?
- DENNETT:** Concreteness is not individuality.
- SMITH:** Right; that is close to what I mean. Note, for example, that in the book I embrace a 'criterion of ultimate concreteness', so using 'concreteness' as a word would be relatively

straightforward. I do believe that everything is concrete. But as a word, 'concrete' is misleading, too...

**CLARK:** And what are space-time points here, they are...?

**SMITH:** They are concrete, but not individuals. You wouldn't treat them as individuals.

**CLARK:** That seems kind of funny if space-time points again come out as concrete. They seem like paradigm cases of something that's not...

**HAUGELAND:** Well, look, what do you understand by concreteness? This is what I understand by concreteness. That in every respect in which it can have a feature, say, in some degree, that degree is fixed. Nothing is left free.

**CLARK:** OK, so it's well-definedness or something.

**HAUGELAND:** Well, no, it's more than that, it is in a way, the difference between kinds and particulars. It's a metaphysical thesis and I can't make up my mind whether it's analytic or not, that particulars are concrete. That is, you can have the picture of the man that doesn't indicate whether or not his fly is open, or whether or not he's got a bald spot on the back of his head, this is left open. A sentence likewise leaves things open. But the man is...everything is settled.

**DENNETT:** That's what I find appealing.

**SMITH:** Another thing that I want to say, which seems to me a simple point, but is hard to phrase using traditional terminology, is that objects-individuals, essentially-are also abstractions. And by being abstractions I mean that some of their concreteness has been...

**CUMMINS:** Lost.

**SMITH:** Well, in a way, but it's tricky. Consider a cup. On the one hand, the cup is fully concrete. Taking it as a cup, however-gathering and clumping this chunk or region of the concrete flux, and treating it as an individual unity-saying, 'Okay, this is a cup'-to do that is to ignore some of its concreteness.

**DENNETT:** That's the price you pay.

**HAUGELAND:** That's to say that the kind isn't concrete.

**CLARK:** So this is just the price of my kind of data compression.

**SMITH:** It's tricky. It's not just the kind that is abstract. I want to agree that the kind is abstract.<sup>42</sup> But the cup is not abstract in the way that the kind is abstract. In taking the particular cup to be a cup, to be an individual; that act of objectification is an act that ignores some of the concreteness.

**HAUGELAND:** That is there in that vicinity.

**SMITH:** Right, it's in that vicinity. Taking it as a cup ignores, it packages the thing together, takes this distributed part of the flux as a unity, a whole lot of things like that. Come at it epistemologically. It's really that objectifying is an act of abstraction. At least at first blush, it's not that the thing that's objectified is abstract, really, because it actually is as fully concrete as you think.

**HAUGELAND:** Right, the definition of concreteness is that everything that could be determinate about it, is fully determinate.

**CLARK:** I'm having trouble with the 'it' here.

**DENNETT:** That's where we get the inscription error you can't get out of.

**SMITH:** You are right that 'it' is the problem. You have to realize that there is more to the cup than figured in you're taking the cup as a cup.

**HAUGELAND:** Well, that's really true.

**SMITH:** Perhaps, for now, I should take that platitude that you, John, would agree with, about everything being determinate, and then just locate the individuality of the cup more in the act of taking it as an individual, and less in the concrete patch of the world there might be-

**HAUGELAND:** Here's a motto I would think you would be sympathetic with (even though it isn't actually using the words quite the way you want to), which is to say that the achievement of objectification is achieving an 'it' such that concreteness makes sense. So, to be objective just is to be, and, I think, a thing, an object. Actually, not just a thing, an object. You have to get it into a space of possible determinacies such that for talking about 'this one', full concreteness makes sense.

**SMITH:** Yep, but we still differ, for several reasons. One is that I want things to be concrete that aren't objects-

**HAUGELAND:** That's cool.

**CUMMINS:** That's all right.

**SMITH:** But I want the determinacy-no I'm not sure it's cool with what you just said, is it?

**CUMMINS:** Yeah, sure, that's OK.

**HAUGELAND:** Yeah, I said objectification, is-the mark of success is-that you've gotten the sense of the possible determinacy in place such that it can be fully concrete. But that doesn't limit

where else might be concrete.

**CUMMINS:** OK and then so from the point of view of the structure, the data structure you do create, a lot of futures don't count as cup futures. So you get this distinction between qualitative and substantial change just built into the finitude of your representation.

**CLARK:** So here's what's getting balanced. In thinking about these things, you have to recognize them as objects. To recognize it as an object is precisely to think that it has all these features and that they are fixed. And on the other hand, what you really want to do is have a sense of what differences don't make a difference. So there can be all sorts of things that can change but you still ought to recognize it as the same object; the cup can get chipped, you still want it to be the cup. Hence even the Y2K thing, you might think that the problem here is that there's a difference that does make a difference that we never thought about.

**SMITH:** It also has to do with projects. The commitments that underwrite identity arise in part from one's commitments. It follows that the identity of an object doesn't inhere in the object itself—that is a very important theorem of this view. So that 'being an object is not an intrinsic property' would be a way to say it.

**DENNETT:** And that's why when the hyper-intelligent extra-terrestrials arrive and find that we're still stuck with objects and properties, this is the ultimate Y2K problem for us. Everything stops.

**DENNETT:** What holds constant when you turn the knobs-sort of tuning for the null, as we say, in radio direction finding.

**CUMMINS:** Yeah, or my idea of which invariant you need to track in order to understand the variance in the error signals.

**DENNETT:** Yeah, we're reaching convergence on this way of thinking of the idea.

**SMITH:** This idea is massively more applicable than just in the case of objects.

**CLAPIN:** The non-modularity of the mind as somebody said earlier in the week.

**DENNETT:** Well, the main thing is that Brian is saying, the overcoat is really thick, and that's where all the action is, or a great deal of the action is.

**SMITH:** I don't want us to infer wrongly from the fact that we theorists don't quite know what's going on, that there isn't

something quite precise going on—even if what’s going on is something that neither we nor anyone else can actually say. The lack of being able to say it doesn’t mean that there isn’t a fact of the matter.

**DENNETT:** Let me go back to my ur-example of indeterminacy of content. In ‘The Ability of Men and Machines’ [1978b] I describe a case where the engineers find this device on the beach and they study it and they agree completely about its physical constitution, and every atom of its being. They agree on exactly what trajectories it will follow under all circumstances. They disagree about what it is, what it’s for. And it’s only when we get to their content level, where they treat certain things as malfunctions—one of them treats certain events as malfunctions, the other one has a different gloss and says those aren’t malfunctions, that’s signal not noise—and I claim that it is not the case that there must be a fact of the matter about which is the right content gloss on this object.

**CUMMINS:** There’s still a scope ambiguity, because the way you put it leaves out the possibility that there is a fact of the matter but they’re both right. I get this all the time, because they say ‘which isomorphism?’ You know? All of them. They don’t like that. Somehow something couldn’t have two structures at once.

**DENNETT:** Right, but of course you can have them.

**CUMMINS:** And really, there’s a real fact of the matter that they have all of them.

**DENNETT:** But there isn’t a real fact of the matter about which one is privileged.

**CUMMINS:** That’s right.

**DENNETT:** And that’s the one point that I’ve always wanted to insist on.

**CUMMINS:** Privilege is always observer-relative as it were.

**DENNETT:** Thank you.

**CUMMINS:** Right, but it is a bad argument from the observer-relativity of privilege, and some premise that somehow builds in that it isn’t there if it isn’t unique to just rampant conventional-ity of all this. And you write sometimes in a way that suggests to me that you think that there’s-

**DENNETT:** I issue a tentative mea culpa. I think I probably do write as if, I think you may have me there-

**CLAPIN:** But with the multiple isomorphisms in the same thing-

usually the case is that only one of those structures is actually doing the effective work.

**DENNETT:** That's perspectival too.

**CUMMINS:** Yeah, that's perspectival too. There's any number. It's simple and clean to think of these couplings one at a time. But the fact of the matter is, this is just another one of these things. In Dan's case, you've got one engineer, as it were, coupled into one structure, and another engineer coupled into another structure. And since those two are different there's a temptation to infer that there's no fact of the matter of which structure is there. They're both there because after all, if they weren't, the two engineers couldn't be coupled to them.

**DENNETT:** I have a Quinian crossword puzzle. It's very simple as a crossword puzzle, but there's two solutions to it. I hand it out to my students and I say this is a simple little crossword puzzle, see if you can solve it. And they come up with the two solutions.

**CLARK:** If someone found a third solution, what do you say about that?

**DENNETT:** That's fine, too.

**CUMMINS:** They're all fine.

**SMITH:** Let me just say a bit about these two isomorphisms—two structures in one thing, right? I think what I would want to say is, there is one thing, of which two abstractions hold.

**DENNETT:** Yeah, that's fine.

**SMITH:** But it's not exactly as if there's indeterminacy in the concrete.

**DENNETT:** No, of course not. We agree.

**CUMMINS:** I think there's a metaphor that may be misleading you here, Brian. You tend to think of structures like shape. You say 'Well out in the world there's the cookie, and I've got a whole drawerful of cookie cutters,' right? And, in some sense, if all the cookie cutters are different shapes, then it just seems to follow that they couldn't all fit the cookie. But that's just because you've got a very limited notion of structure.

**SMITH:** No, look, that's not the point. I have no trouble thinking that seventeen different cookie cutters could fit the cookie. All equally well. What's indeterminate in that case is which type this token is an instance of.

**CUMMINS:** Why isn't it an instance of all of them?

**SMITH:** It is an instance of all of them. That's fine. What I'm saying is that one question one might ask is, 'Look, I'm not sure which of these cutters applies,' and the answer is 'All of them.' That's the question you're talking about; it has to do with the cutters. That's not what I am talking about. What I'm saying is determinate is the cookie, not the cutters.

**CUMMINS:** But you don't want to understand the determinacy of the cookie as somehow a matter of how many cutters fit.

**SMITH:** You absolutely don't. I agree with that. So I want to say, 'That's right: in the actual concrete thing there's no indeterminacy.'

**CUMMINS:** Why isn't the determinacy just the determinacy of fit? It determinately fits this cutter, it determinately fits that cutter, and it determinately fits that other cutter, and that's all there is to it. You exhaust those facts you're just done.

**DENNETT:** Let me read the passage I stubbed my toe on, all right? On page 68 of OO, Brian says. 'Somehow or other—and this I take to be the most important and difficult task facing the cognitive sciences—it must be possible to have determinate representational content, i.e., for there to be a fact of the matter as to how the world is represented.' Brian goes on to say, 'it will have to be an answer that does not depend on how anyone registers or individuates those mechanisms—again, for the simple reason that it happens in people, for example, without anyone doing that.' Right. It doesn't depend on how any observer registers or individuates the mechanisms, but there may still be many different ways of interpreting those mechanisms. And no one of those ways is privileged.

**CUMMINS:** And, moreover, all of them might be either a little or a lot wrong.

**SMITH:** Sure, but the point is, if I'm looking at the rug, registering the rug, and you're a theorist and you're looking at my mechanisms, you've got all kinds of ways of doing it. That's fine. But no amount of slippage or indeterminacy or multiple categorization or simultaneous truth or anything in your interpretation of my content has any consequence as to how I take the rug.

**CLARK:** It doesn't follow from that that you take the rug just one way though, does it? The fact that how you take it isn't determined by how someone, as it were, takes you to be taking

it, doesn't imply that you take it just one way.

**DENNETT:** This really is Brian's problem of the indeterminacy of radical translation.

**SMITH:** It may be that the way I take the rug is in fact to take it simultaneously as instantiating three different types or something. I'm not saying that I can't multiply categorize the rug. It's not the plurality there that is worrying me. Nor, given my love of smeariness, do I have any problem with saying that I take it to be an indeterminate category.

**CLARK:** But you do think that there is something about your state that absolutely fixes, as it were, whether it's a plurality of three or four or...

**SMITH:** Think about this fact: that while there may be questions about, as it were, the classification of the cookie, it doesn't make sense that the cookie is indeterminate. I guess what I'm saying is that when I take the rug to be a certain way, there's a concrete situation here and a concrete situation there and a concrete relationship, about which there are then questions of how to categorize. And something that's going on here is that I may be categorizing the rug in some way, or not, or something like that.

**CLAPIN:** Don't you have to be, to be taken as an object?

**SMITH:** That's right, but there are lots of ways to have a take on it, but not take it as an object. I guess I'm saying that this situation here that is happening is itself fully concrete.

**DENNETT:** Of course it is. But now look-

**HAUGELAND:** Fully concrete under some descriptions.

**SMITH:** No, no, that's where we go back about an hour.

**DENNETT:** It's fully concrete.

**SMITH:** Concreteness is not a property of characterizations of things.

**HAUGELAND:** I'm not sure that you can just say that.

**SMITH:** Well, I certainly can't just say it in a peremptory sense!

**HAUGELAND:** But the thesis of the indeterminacy of translation, carefully named, is that there are some characterizations of things which cannot be fully determinate. There are just systematically different ways to characterize them which are all equal, there's no choosing among them except sort of randomly or on convenience or something.

**DENNETT:** They all tie for first.

**HAUGELAND:** They all tie for first. And there isn't a right answer, there

just isn't a right answer.

**SMITH:** The problem is that I object to something in that thesis.

**DENNETT:** I think you do too. That's where there's a problem.

**CUMMINS:** Now I don't think you need to object to that. I think you can allow that-and still hold on to your concreteness.

**CLAPIN:** I thought concreteness was open to all possibilities-you know, it's kind of, before it's been conceptualized, before it's been categorized, so it's open to all of those categorizations, all of those interpretations. There's a sense in which the concrete things are precisely indeterminate, that all categorizations tie for first because concreteness is before categorization.

**SMITH:** That's right, that concreteness comes before categorization. That's why it is a metaphysical position, not an argument. That the world, as it were, comes completely concrete. Subsequent to that, there are issues of categorizing.

**CUMMINS:** Yeah, sure. I think I'm with you on this, because I think that when, as it were, the world gets targeted, when it becomes the case that the thing that is the target is fully concrete, nothing in the structure of my intenders will allow me to read off all that concreteness or anything. The intender is a pretty sloppy instrument, but fortunately, in this case, the world is there to saturate, to fill in all the holes in the cheese, in a way. And of course equally my representations will just sort of skim the surface in various kinds of ways. But the-whatever it is-

**SMITH:** The stuff.

**CUMMINS:** Yeah—it is fully concrete. I don't have any problem with that. But I think you [Smith] ought to have a problem with that.

**CLARK:** It seems like what you're appealing to is just the fact that there's something absolutely determinate going on. But of course there is. It's like, OK, so something absolutely determinate is going on when the thing crosses the electric eye, and that event triggers something else. But that doesn't make it, as it were, determinate whether or not the electric eye is taking it as a bee-bee rather than a fly, or whatever.

**SMITH:** I appreciate that. I appreciate that if I'm looking at the rug, there's something absolutely determinate about the rug. And there's something absolutely determinate about my state. And there's something absolutely determinate about the relation I bear to the state I am in. So there's something

absolutely determinate going on around here (in the vicinity of my head). And the question is, one of the things that's going on over here is a taking, right?

**DENNETT:** Not one. That's the point.

**SMITH:** OK, there's some taking going on around here. (In the end, of course, what one has to say is that part of what is absolutely determinately going on around here can be taken as a taking-or taken as some taking.) The question, though, with respect to that absolutely determinate taking, is whether it has an absolutely determinate content?

**DENNETT:** Yes, that's the question.

**CLARK:** OK, so 'taking' there is irrespective of content. I'm having trouble keeping them apart.

**DENNETT:** Here's the question as I understand it. There's Brian looking at the rug, and three neuroprousts are scoping out the situation. As neuroprousts, they know everything about what's going on in his brain and about the light impinging and so forth. So they agree on the absolute determinate situation vis-à-vis Brian and the rug, right down to the finest details. So that's the absolutely determinate thing, and they all go to write their neuroproustian accounts of Brian's taking. And they come out different because it's like my Quinian crossword puzzle. It just turns out-extraordinarily implausibly, but you want to make this just an actuarial point. That they come out with three different contents for Brian's taking, and now the question is whether one of those is privileged.

**SMITH:** OK, so, go on, suppose they come out with three.

**DENNETT:** Now, there is this Cartesian intuition—it has been Quine's job and my job and a number of other people's jobs to beat it up at every opportunity—which is the museum myth of meanings, which insists that at least two of those neuroprousts are wrong, and you, from the inside, know what the truth is. And that is the fundamental intuition that Quine is setting out to destroy, and I think he's right.

**SMITH:** Right, but notice something. My claim that there is an absolutely determinate content—that my content is absolutely determinate-doesn't imply that there is any way to settle the question of which of these three others is right. It doesn't imply that there is any way, in this world, to grant one of them priority.

**HAUGELAND:** Does it imply that at most one of them is right?

**SMITH:** No. Nor does it imply that the subject of the taking knows, as it were, what is and what isn't right. The reason it implies none of those solutions is because all of these things that the neuroprousts are doing, and what the-what does Descartes call it?-infallible introspection is presumed to do, the transparency of the...

**HAUGELAND:** The natural light.

**DENNETT:** The light of reason, yes.

**SMITH:** All of those things are more registrations of the taking. They are registrations of my original registration.<sup>43</sup> Both the neuroproust's registration of my original registration, and my own meta-registrations of my original registration will approximate and categorize and lose detail and so on and so forth. That is part of my picture: that absolutely every story massively misses what it registers. So part of what I'm saying is that this picture of registration and the location of ontology is in fact a kind of negotiation between the epistemic act and that to which it's directed. Such a picture makes room for all the Quinian kinds of points, and the sorts of points you're making. Just as you, Rob, were saying a moment ago, it is because our ways of getting at the world are approximate and sloppy; there are all kinds of room for error. But the picture also preserves, I think, a durable intuition, which I think is right: not only are there absolutely determinate phenomena, but also that there may be absolutely determinate content, even if that content doesn't totally tie down the part of the world it registers.<sup>44</sup> Note that the fact that the content is absolutely determinate in my story doesn't mean that content captures all of what's absolutely determinate about the rug-as usual, it massively misses; that's why this is a story of loss, as I keep saying in the book. My only point, here, is that I think that in all of the cases that you are bringing up to show that my position is wrong, there is an extra layer of reference or registration or description, between what I am claiming is determinate and what you are claiming is not determinate. And I'm saying that in one sense you are right, but that that is why-

**DENNETT:** All right, I see where that's going. As long as you clarify all these things that don't follow from your position.

**CUMMINS:** Can I just ask one really quick question? Inscription isn't always error?

**SMITH:** No—or rather: yes, you're right. I actually meant to say that.

**CUMMINS:** So it's open for Dan to say, OK, when I tell my evolutionary story I'm inscribing like mad, but I'm getting it right.

**SMITH:** In a way I'm trying to be really clear about this in the irreduction chapter. I say look, what I require is not that you don't inscribe. What I want you to do is to take responsibility for the fact that you are inscribing (as, of course, you must).

**CUMMINS:** Well if inscriptions are errors there must be some mismatch between them and something.

**DENNETT:** They must miss their targets.

**SMITH:** It's a better or worse kind of story.

**HAUGELAND:** Dan, why didn't you reply to Brian just now, when he made the response to you, that he's pointing out that any registration of any phenomenon is bound to fall short of the full determinacy of the phenomenon—and so there's obviously various ways in which different registrations can fall short—that the issue isn't the certainly undeniable fact that any registration must fall short, but rather that, in some cases, any registration which is sufficient to capture a certain kind of richness in what's there, must inevitably overcapture it, and so there's more than one way to do that? And there's no choice between those.

**DENNETT:** I like that because it nicely conveys a point which people have been making in different ways, and that is, there's a real benefit in carving the world one way or another. The cost is—it's presumptive, and you always get some leverage that won't work as it will turn out, that goes beyond what you've been given.

**CUMMINS:** There's no free lunch. Getting A right inevitably means you compromise B.

**SMITH:** By saying 'overcapture' do we mean not only that it doesn't, as it were, represent things that are the case, but does represent things that aren't the case?

**HAUGELAND:** No, it will render some things in a determinate way. It cannot but render them in a determinate way, to capture as much as it does capture, when there would be other ways of rendering it in a determinate way, distinct determinate ways, which are equally good—I mean, indeed capture exactly the same part of what they capture from the original structure.

**SMITH:** But anyway my answer to that is going to be the same: from the fact that no registration is preferable, maybe even in

principle (which is sort of being assumed here) the indeterminacy of the thing registered doesn't follow for me.

**CLARK:** But nor does its determinacy. Isn't that kind of the point?

**SMITH:** No, that is right. That is why the stuff about determinacy is only on page 52, instead of page 252. It is because it is a metaphysical kind of determinacy, not a-

**CLARK:** But what makes you think that's a determinacy of content? The fact that there's more to the content than any story will capture doesn't imply that the content's determinate.

**SMITH:** No it certainly doesn't imply that. However, neither does it imply the falsehood. The fact that there's more to something than a story will capture doesn't imply that the story's content is either determinate or indeterminate.

**CLARK:** No, that's right.

**CUMMINS:** When you're taking scientific laws here, you're thinking of dynamical laws? You're not thinking of laws that, for example, just tell me what's in the kit? I mean, I think it's a scientific law that there are electrons.

**SMITH:** Right.

**CUMMINS:** Rather than tell me what's going to happen next, or anything like that. It's not dynamical. It doesn't-

**SMITH:** Right. I'm saying that it's not obvious to me that physics is committed to there being electrons as opposed to the (weaker) claim that the electron feature is spatio-temporally instantiated in various ways. That, whether it's one electron or seventeen, and so on and so forth—physics doesn't care: it makes no commitment to reidentifiable individuals.

**CUMMINS:** Well, in my view of things, a lot of science is about mechanisms and how things are built and put together out of stuff, and-it all sounds less plausible as a story about those things than it does-

**SMITH:** It may be less plausible about engineering-

**CUMMINS:** As a story about dynamical physics.

**CUMMINS:** And you mean this in some very strong sense, I mean Dan believes this too, because if the laws of physics were arbitrarily different then nothing would ever replicate and so, as it were-

**SMITH:** Right. I mean it very strongly.

**CLARK:** You don't think of them as inscription errors, field theories?

**SMITH:** Like all registration, they are somewhat pre-emptive. But since they don't register in terms of objects, they give us a

leg up on what it is for subjects to register the world in terms of objects.

**CUMMINS:** I want to make sure I understand the project. The image I get is sort of, OK, because, for example the way we're built, there's going to be this sort of keyhole effect, that you can't, your actual coupling with the world is pretty limited. So the question is, how can you see so much through such a little hole?

**SMITH:** That's right.

**HAUGELAND:** Unlike homing in on a magnet if you're an iron filing?

**SMITH:** Yes, exactly: unlike homing in on a magnet. There (if you are made of iron) you can be driven by the magnetic field. The problem is, it is hard for me to be driven directly by most of the things I care about—such as by Andy's philosophical views.

**DENNETT:** So, in the past I've talked about making something that can detect whether something has once been on my desk. It's extraordinarily hard, unless, of course, my desk was made of uranium or something and it imparted some Geiger-countable property—then, you could use that as a proxy. But in fact we are able to detect all sorts of properties for which there are no natural cheat proxies. How the hell do we do it? We have this elaborate technology for tracking things so that we can, with really very little effort, register, think about, all these weird properties.

**CLARK:** There does seem to be a sense in which it's our practices of timing things that brings 4 o'clock into being. It's not exactly as if it's sort of out there and we just have trouble tracking it.

**DENNETT:** You're not accusing him [Smith] of making an inscription error are you!?

**SMITH:** Andy, you are right; o'clock properties are pure human constructs. You might think that they are so stunningly non-effective in part because they were created, but that can't be quite right. What really matters about the o'clock properties, for the point of the example, is not that they are constructed, but that they are purely formal, in a certain (not so simple) sense.

**CLAPIN:** So it's kind of like adding to physics. Physics has numbers, has maths, it's just adding a bit of sort of logic and what we now think of as implementation theories as a bit of extra formal apparatus for physics—there's a bit more maths.

**DENNETT:** It's a different maths.

**CLAPIN:** It's a slightly different maths that's being used.

**SMITH:** Yep, it is new math. But what is important about (so-called) computability theory is not the math per se, but that it makes new claims-claims that get at concrete regularities that seem to hold, in the world, that involve issues of stabilization and digitization and so forth, claims that seem to be level-independent, that appear to hold across different substrates and at different scales.

**CLAPIN:** Like numbers do. Like cardinality does.

**SMITH:** Well I'm a little reluctant to cosy up too close to the numbers, because many people think that they are a genuinely abstract phenomenon, whereas what I am talking about here are concrete phenomena, at different levels of abstraction.

**DENNETT:** So it's sort of dependencies that are scale-independent and substrate-independent.

**SMITH:** Right! Or rather, originally it was exactly that-genuinely independent. Increasingly, though, it is morphing into a theory of dependencies understood in relationship to (arbitrary) scale, in relation to different kinds of substrate.

**HAUGELAND:** Well, substrate-independent and scale-independent don't mean substrate-less and scale-less; rather that you can have different and in some sense perhaps arbitrarily different scales and substrates and see the same phenomenon.

**SMITH:** Right. Exactly. All I am saying is that I think the theory of how they relate to different substrates may actually end up being part of the new theory, the direction the theory is taking.

**DENNETT:** Oh, well, certain features, like resistance to decay independently of the process you're considering, or constancies...

**SMITH:** And compensations for stability, you know, and for tunneling, cosmic rays, things like that-what kind of circuits will be stable in a nanometer scale, what kinds of stability will in fact hold over periods of weeks, years, etc.

**DENNETT:** Insulation properties, in fact.

**HAUGELAND:** Well, you can't make a Turing machine tape with frogs either.

**SMITH:** Right! There was a person, I remember, when [Digital Equipment Corporation's] DEC20 was made, whose responsibility it was to track the radius of curvature of the lines etched into silicon, because the bits tended to fly off

etched into silicon, because the bits tended to fly off the tracks, if they went around sharp corners too fast—like errant Ferraris. The pulses would just radiate, you know; not make it around the corner. I don't know exactly how medium-independent or non-independent that is.

**HAUGELAND:** Well, medium-independence never meant that you could do it in any medium whatever. It never meant that the medium is irrelevant. It just meant that you could have the very same thing in quite different media.

**CLAPIN:** Two media would be enough.

**HAUGELAND:** No, the thing has to be sort of open-ended.

**SMITH:** Suppose I say, 'Look-this table leg is medium-independent because I can take away the wood and put in aluminum...'

**DENNETT:** Try putting in water.

**SMITH:** Right, you can't make a leg out of water.

**HAUGELAND:** You can't make your computer out of water, either.

**SMITH:** Right. It is all a little gray. Absolute medium-independence won't work, as if it didn't make any difference what you build it out of. We all agree with that. And pure medium-dependence doesn't work either, as if it had to be that of this specific set of electrical components, or that specific piece of protein. What we need is an appropriate 'middling' level of dependence and independence.

This gets to a point you've made, John, about the importance of engineering. I think our engineering practices have very refined intuitions about what kinds of properties materials need to have, in order for what we are building to work. Anything that has those properties will serve.<sup>45</sup>

If you want to calculate  $p$ , then a wide variety of materials will work. If you want something that will run at gigahertz, the range is smaller. If you want something to hold up a table, the range is (perhaps) more constrained yet. My prediction is that this whole terrain will eventually be mapped. And as it is mapped—this is really my claim—the maps will tie together physics as we know it today, which is ultimately concrete, and the (alleged) 'theory of computability' as we know it today, which looks very abstract.

**CLAPIN:** Why isn't that logic? Why isn't that properly logic?

**SMITH:** Because they are merely constraints on physics.

**CLAPIN:** But one way to think of logic is precisely as a description of how to just set up the physics the right way. This was the insight of computation: to turn syntax into something

physical. So it seems to me that there is this match between logic and physics and the kind of representational redescription you're talking about, when you use a new representational code...

**SMITH:** What matters about logic is that there are physically-realizable syntactic configurations that you can interpret, in such a way that the effective transitions end up being semantics-preserving. It was that honoring of the semantic that I take to be the fundamental insight of logic. What happened, historically, I believe, is that computer science borrowed all the theoretic apparatus of logic-including terminology that had been developed in order to talk about semantics-honoring transitions, but then, in a deep way, forgot about the issue of honoring semantics. They took vocabulary that comes from a tradition that was interested in things like proof. But they used that vocabulary to study issues that are really about pure (uninterpreted) mechanism. For example, think of what is called 'denotational semantics' in computer science. Obviously, the word 'semantics' occurs in that label; you might think it would have to do with meaning. But what I believe this phrase really refers to, in computer science, is the relationship, given some machine, between effective arrangements that can be given to that machine as an input (called the 'program'), and a mathematical model of the behavior that results, when the machine is started up on that input. That's not semantics! Note, in particular, that you can construct such a mathematical model for any piece of machinery whatsoever. I can construct a denotational semantics of a can-opener, for example, or for arbitrary mechanisms built out of Meccano.

**DENNETT:** So it's logic gates...

**SMITH:** Yes, except the 'logic' part is gone! See, I think the theory of effective computability has a wonderful first name. It is a theory of the effective. *C'est tout!* It is the second part of the name that is problematic: the computing stuff, which I think has to do with issues of interpretation. Real world computing, I firmly believe, is drenched in genuinely semantic issues. But the theory of computability, the body of work we teach in computer science departments, fails to deal with these semantic issues-issues that actually constitute the practice.

**HAUGELAND:** You say that in the history of computer science and the portion that has to do with the 'computability issues' and so on,

took over the apparatus that went with the problem of the formalization of logic and issues of proof and so on. But in fact that wasn't really part of their topic, right?

**CLAPIN:** Part of whose topic?

**HAUGELAND:** The computer scientists'. That, actually, semantic issues weren't the real issues that were being solved with the computability theory. You're not dismayed by the fact that the semantics fell away. If you're dismayed by something it's that they didn't realize that the semantics fell away and kept using the words.

**SMITH:** That's right. Because it's hellishly hard to tell them that semantics matters, when they are already using the word for something else! This all gets back to the points made in the '100 billion lines of C++' paper. What computer science uses the word 'semantics' for is the relationship between a static program and the dynamic process it engenders.<sup>46</sup>

**HAUGELAND:** As in Scott semantics?

**SMITH:** Yes, as in denotational semantics, more generally, of which Scott semantics is a type.<sup>47</sup> In fact denotational semantics and operational semantics are two characterizations of the same relation (one abstract, one more concrete—the relation between a program and the behavior it engenders).

**DENNETT:** This is Allan Newell and what I call my Julie Christie problem.<sup>48</sup>

**SMITH:** Yes, absolutely. So you've got to find a use for it, and all this apparatus and stuff. Why are they interested in Martin-Löf [1984] and the intuitionists and so on and so forth? The fact is computer scientists aren't radical hyper-intuitionists like Yessenin-Volpin [1970]. All that's going on is that they are studying an intrinsically effective subject matter. The relation between a program and the behavior it produces must be effective. Of course, if you want to study that relationship, and use mathematics to do so, you will be interested in mathematics that concentrates on what is effective. With just a bit of detachment, in other words, and a bit of historical perspective, you can understand why the mathematics went in the direction it did.<sup>49</sup>

**CUMMINS:** A footnote about pointers. Pointers are, in a way, the kind of parade case of embedded intenders, or nested intenders. Because the example you had, in a sense, the low-level intender got fooled because, as it were, the world switched targets on it. The problem was that there's a higher level in-

tender in which it was nested to get something else.

**SMITH:** It's more complicated than that because it was a copy of the pointer that got fooled.

**CUMMINS:** Oh yeah, it is more complicated, but do you see what I mean about the nesting?

**SMITH:** That's right. The interaction between that and the object identity...and I'm just saying we're assuming object identity. It blows fuses in your brain. It's amazing stuff.

**HAUGELAND:** The world, or taking the world as mattering, is ultimately what matters?

**SMITH:** The former: the world is what matters.

**HAUGELAND:** Either could be an intelligible claim.

**SMITH:** That's right, but I think the latter is derivative from the former. It is the world as a whole that matters. Understanding that the world matters—taking the world to matter—that stance matters, too. But it is a subsidiary normative condition, a condition on what it is to be human (or perhaps 'humane'). So taking the world to matter is what ultimately matters about you. That's the sort of humility built into this Brentano-esque form of being oriented. The thing that matters most about you is that you recognize that the world matters more than you do.

**CUMMINS:** One of the disputes between myself and Millikan on functions all along has been that I wanted to say she can't understand how evolution works unless she can first identify the functions, and she thinks it's the other way around. I think this is a similar kind of thing.

**SMITH:** Right; I think it's a similar point.

**CLARK:** Strong enough to be the norms, or to give rise to them?

**DENNETT:** Give rise to them, yes.

**SMITH:** Actually I am just about to talk about that.

**CLARK:** But in between is the question, what makes sense of them—what makes sense of the norms. And there you really want to say that evolution doesn't make sense of the norms. On the other hand it's part of what it takes to make sense of them.

**HAUGELAND:** 'Make sense of them' means 'that in terms of which we can understand them'?

**CLARK:** Yeah. It's part of that in terms of which we can understand them.

**CUMMINS:** I'm with Dan on this, in that I think that just like species,

what evolution stumbled on, the 'it' that it stumbled on is something that is constituted by the fact that it's the endpoint of that branch on the tree, and nothing else. And so you said we want to understand registration such that we can figure out why it was a good thing to stumble on, and I think actually that's a really kind of misleading way to put it. Evolution didn't stumble on registration. It built it out of smaller things, and there wasn't any such thing until it got it, because to be a registration is just to be the endpoint of that branch, developmental branch. I don't really believe this, but I've got to the point where I know how to say it in a way that sounds so plausible that I have a hard time resisting it.

**HAUGELAND:** I'm with Brian.

**DENNETT:** When you described this wonderful growth of registration and all those wonderful things, you presupposed a 'we' that wanted to do this. The we for whom this was the obvious product. And one wants to know, OK, help yourself, there's an agent. There's an agent with goals, there's an agent with purposes, there's an agent that is trying to find out more about the world. Where'd that come from? Now-

**SMITH:** No, look. Let me try to explain. You may be right that I often speak as if I were presupposing the existence of an agent, but you're misinterpreting me. And, remember, I do think that evolution is what did the work-I wish that I had that statement chiseled right here on this table, so that you're not tempted to think that I don't believe it. But given that, here is what I want to say. Suppose your 12-year-old grandchild and I are flying along in our spacecraft, and we notice a planet, which has plants on it. I say, 'Wow! Evolution is happening, all over again!' He says, 'Hey, can we speed it up?' I say, 'Yes, you know, there is something-a Really Neat Trick-that evolution on this planet doesn't seem to have been stumbled on yet. This trick involves organisms playing games with their internal structure so that they can track stuff that they are not physically coupled to. It might take another 100 million years for that to happen. Why don't we just drop down and do a little genetic engineering, give evolution a shove?'

In principle, this is something that could be done. Registration is a way of being that works. I don't think it is just the endpoint of some evolutionary branch. The world is such that registration is a way of being that is powerful.

And it would be powerful, whether evolution found it or not.

**DENNETT:** What's fascinating to me about that is that now you're playing the card which I intend to play, usually, and get hammered on by people like Dick Lewontin who says don't, don't, don't think of evolution as these problems that are posed and then are solved by evolution—these sort of Platonic problems that are solved. And I'm very happy to be this sort of minimal Platonist. You know, there really are these problems independent of history that are posed, or could be posed. We can conceive of them being posed again and again and again across the cosmos, and solve them in the same way. That's what a Good Trick is. And I think that's fine and I think that's quite consistent with evolutionary theory. It is of course an idealization. There's danger of inscription errors. But it's a deep way to think about evolution. And when people like Gould and Lewontin chastise themselves and their fellow theorists for doing it, it is at the cost—and this is sometimes glaring and even to the point of being comic—that they can't even talk about convergent evolution. That's why convergent evolution is a sort of non-topic for Gould—it's bizarre—and for Lewontin. The reason it's a non-topic for them is because they can't let themselves talk about the same solution to the same problem they reinvented. But I think it's a deep part of evolutionary theory to be able to separate—just as you say, to separate the discovery from the process.

**CLARK:** But I think that one thing that you don't want to do—and this kind of fits in with something Joan [Wellman] was saying last night—is to get forced into a discussion about where normativity comes from when all that really matters, for most of John's projects, is what it's like when it gets there. Yet understanding where you are actually isn't, I suggest, independent of understanding how you got there. When you want to understand where you are, you want to understand where you're likely to go. And the kind of processes that got us there are still, one way or another, operative.

**SMITH:** Remember, I didn't banish Dan to the other side of the town; I took him as a flying buttress; that's right. There's something extremely important about understanding how we got here.

**WELLMAN:** Brian, that's the second thing you said in your list of three

last night to me. You said it matters—the way that something is implemented matters to the way it can be.

**SMITH:** Yes, I think that's right. I'm all in favor of understanding history and implementation. At the risk of agreeing with you all so much that I sound wimpy, there is even a reading of 'rides on the coat-tails' that I can agree with: that evolution is the train it took—a train made of coat-tails. Our normativity, our registrational prowess; they all came via that route.<sup>50</sup>

One way to understand this is in terms of the design space. That's one thing that being a computer scientist teaches you: to be interested in the entire space, the whole fitness landscape, not just in the structure of a particular solution. Maybe on the other side of some great canyon in the fitness landscape there lies an enormously powerful and possible way of solving some of the problems you face—but evolution will never find it, because it is too bloody expensive to go in that direction. No creature could survive an attempt to cross that canyon.

So just as I think how you build something is really important, but not necessarily constitutive, I also think that understanding what's constitutive, and what the space of possibilities is and so on, has got to be helpful in terms of how the evolutionary story went. So I'm all in favor of this handshake with evolutionary theorists. I just don't think that norms and mattering rest constitutively on evolution.

**DENNETT:** Then I think we agree.

**CLARK:** I don't see any reason to deny that.

**CUMMINS:** I'm not sure I do.

**SMITH:** I don't believe there's complete agreement, but I do think there's a kind of...

**DENNETT:** I think there's still some tension here...

**HAUGELAND:** There's a question which I've been sitting here trying to formulate, and I'm not sure how it bears, but I have this inkling that it does. And that is, we were talking about Good Tricks and having a phrase like, 'it's a Good Trick in one's kit'. But there's some question as to how these are individuated, what makes them a trick and what makes them good. And then how they could be brought about; whether evolution is the only way. For instance, why isn't it a Good Trick to produce planets composed of heavy elements; to get a whole lot of hydrogen to attract itself to the point

where it then fuses into helium and then collapses into a dwarf where the helium fuses-then it explodes and those pull together and then you've got a planet. That's a pretty tricky process, right?

**DENNETT:** Read Lee Smolin [1997] and you'll see somebody who says, add it right in there, another Good Trick. That's what evolutionary cosmology is about.

**HAUGELAND:** There's no selection there.

**DENNETT:** Oh, there is for Smolin. I'm agnostic about Smolin's cosmology, but I think it is a not provably incoherent cosmology that simply embraces what you're trying to do as a reductio, and says no, look, we actually do have an evolution with selection of whole universes, that is, whole ways of having the basic so-called constants of physics. And some of them produce...

**HAUGELAND:** Yeah, but you've changed the subject. I'm talking about this universe with our constants of physics-

**DENNETT:** It is a Good Trick and can be seen to be a Good Trick, but, as usual, you never see it against the background of the failed universes that don't happen.

**SMITH:** What about saying 'Look, why don't I just inhale some dioxin and decompose...'

**CLAPIN:** Why is that a bad trick?

**SMITH:** Yeah, what's the 'good'-

**CUMMINS:** Well, you have to have replication in the picture. Replication and selection wasn't in the picture that John told. It is in the story that Smolin told.

**DENNETT:** It is in Smolin's story.

**CUMMINS:** I don't believe it for a minute. I don't disbelieve it, either. I just don't get it.

**DENNETT:** But my point is, you've asked the right question, but don't presume that there isn't an answer to it.

**HAUGELAND:** Well, the question is not to Smolin, whom I've never heard of, but rather to you, who don't believe in the alternative universes, let alone them procreating and competing. Just to you, believing, like we Weinbergians or whatever, that it started back there in big bang and it's been buzzing along ever since. Well there it's been evolution in the sense of the trajectories but not evolution in the sense of selective pressures. And yet there are things which it's not obvious why you wouldn't call them Good Tricks. Unless you build se-

lectiveness into the notion of good thing, and that's a suspicious move.

**DENNETT:** Well, I've got a long answer to it, but I don't have a short answer to it, and it's time for lunch. But you can read the long answer because it's in my book *Darwin's Dangerous Idea*.

— *Were this page blank, that would have been unintentional* —