God, Approximately • 4

Brian Cantwell Smith*

Preamble

Consider the rise of the religious right: Muslim and Hindu fundamentalism, right-wing Zionism, the Christian moral majority. These movements are responding to—and exploiting—a widespread social hunger: a sense that reigning secular, scientific, and capitalist world views don't supply what matters: ways to tell right from wrong, guidance to anchor individual lives and give them meaning, the wherewithal to resolve ethical dilemmas.

I find many of the fundamentalists' answers appalling: bigoted, mean-spirited, scary. But what are those of us on the left—we scientists, we intellectuals, we in the academy—doing about this heartfelt lack? If we don't recognise (and respond to) the yearning—if, willfully or unwittingly, we remain blind to the hunger—then we have no leg to stand on, in criticising others' replies.

What we need are *better* answers: frameworks to stir compassion, give meaning to lives, combat prejudice, secure a modicum of economic well-being, preserve the planet. These frameworks must be global; it is too late for parochial sectarianism. And they must build on the best in science. We need to move forwards, not back.

This talk is offered as a step in that direction.

(A note to those who are allergic to religious language: you have my sympathy. In fact you almost have my company. But I have come to believe that the questions

---

*Cognitive Science, Computer Science, and Philosophy
Indiana University, Bloomington, IN 47405 USE
Comments welcome
smithbc@indiana.edu
Copyright © 1998 Brian Cantwell Smith
Version of November 24, 1998
Note: Several different versions of this paper—or rather, different papers with the title "God, Approximately"—have been presented over the past year. The first was given at a "Science and the Spiritual Quest" conference held at the Center for Theology and the Natural Sciences, at the Graduate Theological Union in Berkeley, California, June 7–10, 1998. The present version (4) was given on November 18, 1998 at the "God and Computers" speaker series at the Massachusetts Institute of Technology (M.I.T.). At some point in the future the various versions will be assembled into a small monograph.
This version is available online at: http://www.ageofsig.org/people/bcsmith/papers/

that have traditionally been asked in theological contexts—about ultimate significance, what it is to be a person, the grounds of justice and humility—are too important to leave hostage to the pro- and anti-religious debate. I don't care, in the end, who is, and who is not, religious. What matters is whether, and how, we answer those questions.)

Now you might take this as a call to arms: that scientists (and I am by and large going to speak of scientists today) should set aside their detailed professional concerns, and take a stand. You might especially think this on what I will call the **classical view**: that we can divide our understanding of the world into two parts: (i) a roughly causal, deterministic, third-person-accessible value-free theory of material objects—a knowable science of a predictable physical world; and (ii) a more mysterious, phenomenological or experiential, inexorably first-person, value-laden theory of spirit. A realm for the body, and a realm for the soul. It's a convenient dualism, an historically-entrenched dualism, and still, I believe, a widely-accepted dualism. But like many dualisms, it is wrong-headed.

But before arguing for change, I want to take stock of where we are now. For the science we have inherited, at the end of the twentieth-century, is radically unlike that classical image. In fact it is more unlike it, I believe, than anyone has yet recognised. In a way, that's the point of this talk. If we truly understand what has happened, intellectually, in the last hundred years, that will overturn our sense not only of where we stand, but of what we need to do to meet the religious challenge.

1. 20th-century science

Three 20th-century scientific developments are particularly important.

First, what's gripped the imagination of the mythical person-on-the-street[1] are results in physics: quantum mechanics and relativity, in the first part of the century—strings, chaos, and fractals, more recently. Extremely odd stuff, especially at the quantum level: ten- or eleven-dimensional curved universes of strange forces governed by inscrutable logic. But even on a classical picture, modern physics is alien: a stupefying spray of interpenetrating waves of every conceivable frequency—turbulence, attractors, vortices smashing and piling up on top of each other in dizzying disarray. Imagine falling

[1]At least the person on Mass. Ave.

overboard in a storm—and opening your eyes to nothing but salt and spray. Now subtract you. It's a little like that out there, only a zillion times worse.

A far cry from tables and chairs, obedient street lights, moribund committees, the PTA. So let's label this distinction. By the **physical world** I will mean the almost incomprehensibly strange, object-less world of modern physics. By the **material world** I'll mean our familiar day-to-day realm of medium-sized macroscopic objects: people, cars, elections. It is possible that the material world is derived from the physical world by abstraction—but "abstraction" is a concept in epistemology, not in physics, and so itself stands in need of explanation.

So that's the first thing about 20th-century science: the physical world is surpassingly strange, and quite unlike the material world. The second result—also relatively well-known—is that scientific knowledge is *epistemologically recalcitrant*—far more so than was classically expected. Traditionally, a sharp line was taken to divide *knower* from *known*. But that simplistic model proved unsustainable. Relativity brought perspective into scientific claim; quantum mechanics shattered the myth that discovery was innocent. In human affairs, of course, both facts—the perspectival aspect of knowing, and the violent character of finding out—are almost truisms. What's novel is their implication deep in science: in the measurement problem, in the collapse of the wave function, in quantum indeterminacy. Maybe all knowledge is violent, perspectival, implicate.

The third development I want to draw attention to (beyond ontological and epistemological issues in physics) is the rise of the **intentional sciences**: sciences dealing with symbols, meaning, reference, interpretation, truth. Logic has been under investigation for millennia, of course, but for most of that time it was viewed as ancillary *equipment for the doing of science*, rather than as itself subject to scientific (especially empirical) investigation. Starting in the 19th century, however, that began to change, with Babbage's engines, Boole's *Laws of Thought*, the ground-breaking philosophies of Frege and Peirce. Their achievements flowered in the 20th century, unleashing modern logic, meta-mathematics, psychology, linguistics, cognitive science—and the entire computer revolution.

I want to highlight two facts about intentional phenomena: a way in which they are, and a way in which they are not, "merely physical".

The first—the non-physical aspect—is most obvious in the case of refer-

ence. Something that is so self-evident, in our experience, that it is impossible to imagine not being true, yet from the point of view of physical science seems almost miraculous, is our ability—with a simple few words, or the merest thought—to refer to things that are far way away, in time, space, possibility. I can refer to the Pharaohs of Egypt, without violating backwards causality; anticipate the great day on which the first woman will be inaugurated as President of this country—without violating forward causality. I can refer to events outside my light-cone, with which physics prohibits any causal interaction at all. Moreover, this "**arrow of directedness**" is exquisitely precise: I can refer to the paint flaking off the ceiling exactly 2 centimeters in from the north-east corner of the examination room in Mother Teresa's clinic in Calcutta, and my reference will "land" precisely there, with unutterable accuracy. Reference is zippy, too—traveling "at the speed of logic," as Alonzo Church once said: we can describe the temperature on the surface of the sun, without our reference taking eight minutes to get there. We can even refer to what doesn't exist: to the different kind of mess we would have been in, if the U.S. had a parliamentary system, and Newt Gingrich had been prime minister, these past four years. And so on and so forth. Forget angels: reference goes where fools merely *imagine* treading.

There is a sense, that is, in which reference is physically **transcendent**. But the amazing thing is that it manages to do these things without requiring spooky metaphysics. That is: there is also a sense in which reference is physically **immanent**. Referring outside your light cone doesn't contradict natural law. Somehow or other, without violating physical locality, we are able to direct our thoughts and ideas outside the confines of the 1/r2 tent or envelope to which physics restricts physical engagement.

How we do this, I believe, is as important a 20th-century discovery as any in science. It is not identified with any single name, though the giants of the early logical tradition—Frege, Russell, Gödel, Turing, Carnap, and others—deserve a lion's share of credit. Unfortunately, the solution was codified within the **formal tradition**, and so is not publicly appreciated. ("Formality," it turns out, is a perversely abstract form of digitality: an assumption that theoretically useful categories can be completely, unambiguously, and absolutely divided—a particularly stubborn form of thinking in black and white.) But the insight can be extracted from the clutches of the formal tradition, and renewed to serve another day—or perhaps another century.

Here is the basic idea. Physical regularities—causes and effects—are, as I've said, **local** in essentially all relevant respects (spatially, in the sense of 3D proximity, and temporally, in the sense of temporal immediacy). That poses a problem for cognitive or intelligent creatures. All you *get*, if you are physically embodied—i.e., in terms of effective resources—is what is pressing in on you, right now, at the surface. You live in a laminar cocoon, that is, with physical coupling limited to the immediate here and now. Moreover, the world is *sloppy* (only weakly correlated), so you can't necessarily *tell*, from what is happening right near you, what is going on elsewhere—behind that rock, or back at home, let alone what went on yesterday, or will go on tomorrow. Fortunately, however, that same slop—the local degrees of freedom—also means that an agent can rearrange its internal states with remarkable facility (if it's clever), without expending much energy. So what we do—what agents do, what it is to think—is to *represent the world out there*, beyond the periphery, by rearranging our internal configuration, and adopting appropriate habits and practices, so as to behave appropriately—so as to be able to think about, develop hypotheses concerning, stand in appropriate relation to—that to which we are not, at the moment, coupled. What we can do, sometimes, is to exploit correlations between the incident, proximal effective array and the (non-effective) distal situations we care about (those correlations are called "information")—though the connection between the two is often intricate, involving lots of internal machination (called "inference").

It's not just *amazing* that semantics outstrips causal bounds, that is, in the referential sense we spoke of a moment ago; that's what semantics is *for*. A "purely" physical entity—a patch of the world that hasn't figured out how to organise itself, locally, so as to be oriented towards what is distal—is existentially limited to the incident press of the immediate physical surround. Living in the here and how is cheap; that's what physics gives you. Living in the there and then—that takes smarts.

Anyway, studying that stuff—how critters and computers can represent and reason and be directed towards what is distal—how they can think and act locally, while honouring what is global—that's been the subject matter of a 100 years or so of intentional science. And the fundamental principle on which it all works—the same principle that underwrites all of computing—truly is a "Big Idea."

So those are the three ideas I want to highlight (the weirdness of physics, physics' epistemological problems, and the rise of the intentional sciences). As we'll see, they are related. They also have an interesting implication. On the classical image, as shown in the figure 1(a), *knower* was viewed as external to *known*. Physics has been epistemologically problematic because, as I said, physicists (and their practices) are implicated in physical knowledge, but physics is not a science that can explain knowing or practices. So physics has lived with an unresolvable predicament. It takes the intentional sciences to close the picture—and resolve the tension. The result is shown in figure 1(b). Since scientific knowing is one kind of knowing, extending science to include intentionality has brought the *doing* science *inside* science.

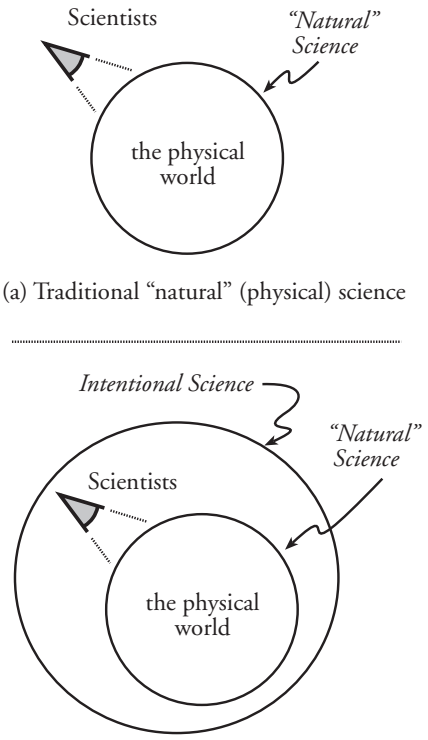

(a) Traditional "natural" (physical) science

(b) The intentional sciences

Figure 1 — Science

2. State of the Intentional Art

How is the last of these three developments—the emerging science of intentionality—getting on, in its self-referential project of understanding understanding?

There are three things to be said.

2.a. From formality to participation

First, as the century closes we are seeing something of a sea change. The models of reasoning and knowing that were dominant in the first decades of the century, and that reigned in cognitive science and AI as late as the 1970s, embodied what we might call a **rationalist view from nowhere**. Intelligence was taken to be a process of (inductive or deductive) deliberation; agents were treated (at the relevant theoretical level) as abstract and unlocated; and the task domains that agents reasoned about were assumed to be ontologically unproblematic: laid out in advance, consisting of distinct, well-behaved, durable objects, exemplifying properties and standing in various pre-established

structural relations (rather in the ways we imagine in model theory).

Recently, however, this whole cluster of assumptions is being set aside. The distinction between agent and world is recognised as problematic, perhaps even illusory. Far from being abstract, agents are seen as crucially active, embodied, participatory—made of the same stuff as the domains in which they act. Perception, action, and behaviour are viewed as more paradigmatic of intelligence than is hypothetico-deductive ratiocination. Think: *driving around*, not *proving theorems*.[2] Finally, it is increasingly well-recognised—especially in the practical trenches—that those neat ontological assumptions don't work. The world simply doesn't come all chopped up into nice neat categories, to be selected among by peripatetic critters—as if objects were potted plants in God's nursery, with the categories conveniently inscribed on white plastic labels. A glance at the input lines coming from a robot's TV camera quickly dispels any such myth. On the contrary, it is now widely accepted that figuring out how to parse or carve the environment in ways appropriate to the task at hand is a (perhaps *the*) major task that cognition must face.

All of these changes—captured in such slogans as *situated cognition*, *embodied AI*, etc.—can be viewed as retreats from the aforementioned **formal tradition**: in various ways, they share a recognition that most binary distinctions (between inside and outside, symbol and referent, abstract and concrete, continuous and discrete) are *partial*: negotiated, plastic, problematic. Discovering and stabilising such distinctions, when necessary, is as much of an achievement, on the part of a robot or animal, as is working within the resulting framework, once stabilised. Establishing an appropriate conceptual scheme, that is, may lie closer to the heart of intelligence than working within one.

2.b. Computability

So that's the first result: intelligence is messy and participatory.

---

[2]It is interesting that Descartes is being turned upside down. Whereas he thought that perception and navigational action were relatively simpler (accomplishable by "mere brutes"), rational high-level reasoning, he thought, were beyond the powers of ordinary physical devices. The recent history of artificial intelligence and cognitive science suggests that high-level symbol-manipulating logical inference (in at least some cases) may be much easier for machines to achieve than in-the-world perception and action.

The second result has to do with limits. Gödel's incompleteness theorems, proved in the 1930s, rocked the intellectual community. For the first time, it seemed, intrinsic limits were being placed on possible intellectual achievement. But again, as with the "Big Idea" underwriting intentionality, the formulation of these insights in formal terms obscured their true significance. Technically, the result is that *semantics* can never be wholly reduced to *syntax*. Pragmatically, however, in terms of the picture given above, the moral is simply this: though local, effective, physical arrangements can do a good job of standing in for remote, non-effective, or abstract situations (especially in highly-constrained circumstances) they are never *perfect*, Proximal surrogates can never entirely capture what matters about distal subject matters.

Moreover, the relative computability (i.e., complexity) results of the second-half of the century drive this moral home. It doesn't matter that chess is finite, when the number of possible moves in a game is on the order of 10^{120}. Remember that the number of possible 16 × 16 bit cursors on your computer screen is almost a million times greater than the number of electrons in the universe. Unchecked combinatorics kill you.

Perhaps the most important consequence of these computability and complexity limits—especially when taken together with the weird ontological claims of physics—has been to drive an **irrevocable wedge between determinism and predictability**. Classically, it was assumed—most people still assume, I suspect—that if you *know* everything about a situation at time t, and that situation is deterministic (in the sense that what happens at moment t+1 is entirely determined by what is true at time t), then an intelligence—at least in principle, if not in practice—could figure out what the situation will be like, at time t+1. *But it is not so*—at least not for any intelligence that is part of this world. In all but vanishingly few cases, in fact, waiting for the results to happen may be the metaphysically optimal—perhaps the only—way to know what will happen in detail (especially on the wonderful suggestion that the universe is running an optimal algorithm).

In sum: the impossibility of accurate (epistemological) prediction is entailed by quantum indeterminacy, by computability limits, by complexity results, by turbulence and chaos and other aspects of non-linear dynamics, by emergent properties and emergent objects … on and on. There is probably no more securely established result in science. But it is still difficult to get our minds around.

Taken together, these first two results (finite, embodied participatory creatures, subject to massively strong computability limits) radically undermine the classical image of a rational, all-knowing, übermensch. In its place we get an increasingly deep sense of a world of finite, embodied creatures, struggling to make their way around in—and struggling to make sense of—the world around them, using partial, flawed, perspectival, incomplete, knowledge and skill. And note that these aren't just practical limitations. Just about every major 20th-century scientific result, both ontological and epistemological, underscores the paucity of potential intentional prowess.

It is a humbling image.

2.c. Dynamic norms

The third technical result of the intentional sciences is less familiar, but if anything more consequential. It has to do with **norms**. I haven't said anything about norms up until now, but to enter the realm of representation—description, language, interpretation, truth, etc.—is to enter a world of phenomena governed by asymmetric (paired) evaluative predicates: true vs. false, good vs. bad, working vs. broken, beautiful vs. ugly— where one option is *better*, or *more worthy*, than the other. Accurate descriptions are better than inaccurate ones; information is better than misinformation, helpful behaviour is better than unhelpful behaviour—and so on. One good definition of intentional systems, in fact, is that they are just those systems that are subject to norms.

Truth is a famous norm—but not particularly general. Scientifically, moreover, it has been treated as a **static norm**, in the sense of applying to (passive) *sentences* or *claims*—i.e., to *states*.[3] Full-blooded intentional systems, however, are *dynamic*, and hence governed by **dynamic norms**—norms that govern *process*. In logic (the seminal intentional science), the only dynamic norm that's received much attention is what I will call a **derivative** dynamic norm, defined in terms of a static norm. In particular, logic's dynamic processes (reasoning, deduction, inference to the best explanation.) are mandated to *preserve* or to *produce* truth, where it is assumed that truth and "best expla-

[3]By static norms I don't mean norms that don't change, over time; evaluative metrics on book design, or on human beauty, may evolve considerably, but would still be counted as static, on my typology, because what they are evaluative predicates *on*—books or motionless bodies) are essentially static things.

nation" can be defined independently of, and prior to, the processes of their preservation or production.

This explanatory strategy—of starting with a (presumptively autonomous) static norm, and then defining dynamic norms in terms of it—has been picked up by other intentional sciences. Economic models of rationality and decision-making, for example, often use the dynamic norm of *utility maximisation*—where utility is (once again) presumed to be static, prior, and autonomous. But the strategy doesn't generalize. And no computer scientist believes it. On the contrary, what practical experience with computing has taught us is that you need to go in the opposite direction: taking the semantic content (meaning) of a symbol or expression or data structure to be determined (even to exist) depending on *how it is used*—i.e., on the role it plays in the overall system of which it is a part. Rather than define dynamic norms in terms of static ones, that is, programmers *define static norms in terms of dynamics ones*—in a (perhaps unwitting) endorsement of the Wittgensteinian maxim that "meaning is use."

If we get our static norms derivatively from our dynamic ones, where do we get the original dynamic norms? What are they like? What governs, what puts value on, what evaluates, the use—i.e., the life and times, the activity—of general intentional processes? This question isn't usually asked quite so baldly, though a variety of alternatives are being explored. But the dynamic norm that is currently receiving by far the most scientific attention—in cognitive science, ALife, evolutionary epistemology, research on autonomous agents, and of course biology—is **survival**.

It's clear how you get a norm out of survival: a process or activity is deemed *good* to the extent that it is *adaptive*—i.e., to the extent that it aids, or leads to, the long-term survival of the creatures that embody or perform it. This idea of resting normativity on evolution has proved seductive. It has been used to define a notion of *proper function*, for example, in terms of which to decide whether a system is *working properly* or is *broken*. Thus the *function* of the heart is to pump blood, it is claimed, and not to make a "lub-dub" sound, because hearts were evolutionarily selected for their capacity to pump blood, not for their sound-making capabilities. Similarly, the function of sperm is to fertilize eggs because that is why sperm have survived (even if only a tiny fraction of them ever serve this function).

Most interesting for our purposes, however, is the use of this same idea to

define semantic content (meaning, reference, representation, truth). The representation in the frog's eye *means* that a fly is passing by, some people claim, because it leads the frog to behave in an adaptive way towards that fly (namely: to stick its tongue out and eat it) in a way that contributes to the frog's (not the fly's) evolutionary success. Similarly, the shadow on the ground conveys information *about* the hawk in the sky to a mouse just in case it plays an evolutionary adaptive role of counterfactually covarying with the presence of hawks in a way that allows mice to escape.

Have we reached the end of the line? Will evolutionary survival be a strong enough dynamic norm to explain all human norms: justice, altruism, authenticity, caring, freedom, and the like? I sincerely doubt it. But in a way that is just the point. For what is at stake is not what will ultimately subserve all the norms we need in order to understand human activity, but to understand *what the dynamic norms are in terms of which human activity is understood*. And that, I hope, is obvious: dynamic norms on human activity govern *what it is to live*—what it is to live well, to do good, to be right. That is: **ethics**. And not just ethics, but whatever governs *whatever* you do: ethics, curiosity, eroticism, the pursuit of knowledge for its own sake … and so on and so forth, without limit.

In sum, taking on full-fledged dynamic normativity is an unimaginably consequential move. It implies that any viable account of intentionality—any transformation of science broad enough to incorporate intentional systems, and thus to treat meaning along with matter and mechanism—will also, thereby, **have to address mattering as well.** Put it this way: in spite of what the logical tradition may have suggested, you can't just bite off *truth* and *reference*, and glue them, piecemeal, onto physical reality, without eventually taking on the full range of other norms: *ethics*, *worth*, *virtue*, *value*, *beauty*. By analogy, think of how computer science once thought it could borrow *time* from the physical world, without having to take on *space* and *energy*. It worked for a bit, but soon people realised what should anyway have been predictable: that time is not ultimately an isolable fragment (not an "independent export") of physics. By the same token, it would be myopic to believe that the study of intentional systems can be restricted to some "safe" subset of the full ethical and aesthetic dimension of the human condition—and especially myopic to believe that it can traffic solely in terms of such static notions as truth and reference, or limit itself to a hobbled set of

dynamic norms (such as survival). To believe that would be to be an ostrich, not a hero.

Moreover, to up the ante (in case this strikes you as too mild), something else, if anything even more expensive, is implied by these same developments. (This is where the pieces of the story start to fit together.) I said that the classical model assumed that the meaning of symbols and representations could be assessed in terms of the objects and properties in the world that they corresponded to, independent of how those symbols and representations were used. But I have also said that no one believes that classical model any more—in part because the physical world doesn't supply the requisite objects (remember, representational contents need *material* objects, which physics doesn't supply). That means that it is incumbent on a theory of representation to explain the objects that figure in the content of a creature's representational states. Objects, that is, are to be explained in terms of the normative structure governing the representations whose contents contain them. And those norms, we've just realised, are ultimately grounded on *dynamic activity*.

The material ontology of the world, in other words—what objects and properties there are, for a given creature (not just what objects and properties the creature *takes* there to be, but what objects and properties there *actually are*, in the world, for that creature)—will be a function of that creature's projects and practices. For high-level social entities this isn't surprising: date-rape didn't exist, I take it, for the aboriginal singers of the Australian songlines; the strike zone (one of my favourite objects) isn't part of the furniture of the world, for earwigs. But the present claim is more radical: it says that what is true for date-rape and strike zones is true for food, clothing, rivers—perhaps even the number four.

Ontology is inextricably linked to epistemology, in other words, and epistemology is inextricably linked to ethics. That's okay; these are conclusions I am happy with; they are also consonant with a thorough-going rejection of formality. What is striking about them in the present context, however, is that we have come to them by making two seemingly innocent moves: (i) by understanding the limited contribution physics makes to material ontology; and (ii) by giving dynamic norms explanatory priority over static ones. That is: we have come to these conclusions not as meta-scientific attitudes—matters of preference or stance—but as *science-internal results*.

Not only has doing science been brought within science, in other words;

so has the nature of ontology. And not just ontology, but fundamental questions of metaphysics—which of realism, irrealism, formalism, or idealism, is right. And not only that; the answer to this question is staring us in the face. Metaphysically, the world is one. That's an anchor of scientific inquiry, to say nothing of common sense: no matter how disparate our cultures, my insecticides pollute your water supply; your car bomb rains on my parade. Ontologically, though, our worlds are many. They are many because objects and properties involve abstraction—and *abstraction is normatively governed*. Our objects are as constitutively different (and, of course, as constitutively the same) as our policies and practices.

We can summarise this conclusion etymologically. A *material* object is something that *matters*. It must matter, in order for the normative commitment to be in place for the objectifying creature to take it *as* an object: to be committed to it as a denizen of the world, to hold it responsible for being stable, obeying natural laws, and so forth—and to box it on the ears, when it gets unruly. It is no pun, in other words, or historical accident, that we use the term 'material' as a term for things that are concrete (made of "matter") and also as a term for things *that* are *important*, as in 'material argument' or 'material concern.' In fact that's one way to see where the intentional sciences are heading: they are going to have to heal the temporary rift that for 300 years has torn matter and mattering apart.

3. The Age of Significance

Now there is one thing I should say right away—to deflect misunderstanding. If you ask most modern scientists—modern intentional scientists, that is—whether they *think* they are engaged in irrealist metaphysics, or probing the ethical structure of the human condition, or constructing scientific models of the good life, their answer, I assure you, will be: *no*. But I don't care. I don't care what people *think* they are doing. I care about what they—which is to say, *we*—are *actually* doing. Irrealist sentiments notwithstanding, I am enough of a pragmatic realist to assume that it is what we are really doing that is going to matter in the long run.

I will admit, though, that I am reminded of that cartoon figure (was it Road Runner?) who runs off the edge of a cliff and then hangs, there, motionless, for a moment, until he looks down—and only then falls. Cognitive scientists, computer scientists, evolutionary biologists, and the like, often

seem to me to have run off the edge of the naturalist's cliff, but not yet to have looked down. The allegedly "safe" terra-firma of pure, local, causal explanation and pure, unadulterated physical phenomena—mainstays (even bastions) of science as we know it—has long since been abandoned. Locally—in the thick of moment-to-moment interchange—people defend a variety of sensible views, such as that what a data structure means depends on how it is used. What no one realises is how much such seemingly innocent adjustments shake the foundations of what for 300 years we have called "natural science."

Which leads me to say something about computing. I have repeatedly mentioned computation as one of the intentional sciences. Some of you will know that for almost 30 years I have been engaged in a foundational inquiry into the basic nature of computing—trying to figure out what it is, where it came from, what its intellectual importance is, what it augers for the future. I have spent 30 years, the project is largely complete … and I have failed. Or rather: I have succeeded, I believe, in coming up with the answer. But the answer is: there's nothing there.[4]

Here's the point: Computers involve an interplay of the two things we have been talking about since the beginning: (i) **mechanism**: in the sense of a materially-embodied causally-efficacious process, wholly grounded in a physicalist metaphysics; and (ii) **meaning**: in the sense of a realm of symbols, information, representation, norms, and the like. It is universally believed, however, that, in addition to these two things, computers are somehow *special*: that they involve some *particular* interplay of these two issues, have some distinctive or characteristic identity—digitality, for example, or formality, or abstractness—that makes them a worthy subject matter for theoretical investigation. That computation is special is implicit in the idea that there might be a *theory* of computing, a *theory* of computation.

That is what I now claim we will never have. Admittedly, there does exist a body of work called the "theory of computation," which is a bit distracting—but I am prepared to go to court to say that it is *wrong*, in the following strict sense: *in spite of its name, it is not a theory of computing after all.* (Which is not to downgrade it. What it is, I believe, is something of incalculable importance, worth a passel of Nobel prizes: a mathematical theory of causality.)

---

[4]There are no computers. Bill Gates has fooled us all again.

But it isn't a theory of computing, because it deals with only one of computing's two constitutive aspects: with mechanism, but not with meaning. But the problem is not just with that theory. We will never have a theory of computing, I claim, because there is nothing there to have a theory of. C*omputers aren't sufficiently special*. They involve an interplay of meaning and mechanism—period. That's all there is to say. They're the *whole thing*, in other words. A computer is anything we can build that exemplifies that dialectical interplay.

Now to spend your entire professional life looking for a theory, only to conclude that there is no theory there to be found, might seem like a dismal result. On the contrary, however, I believe almost exactly the opposite: it is far and away the most positive result that a computational triumphalist could possibly hope for. Moreover, it is a result that has a direct bearing on our present subject matter.

The situation is depicted in figure 2. Along the bottom are the centuries—from the 14th up through the present, into the 21st and 22nd. The vertical axis represents something like "importance" or "weight." On the left is natural science, rising in the 16th and 17th centuries, peaking in the 19th and 20th. The thick part of the curve, on the left, represents alchemy: a rag-tag bunch of exploratory practices, conducted in a disheveled and atheoretical way, that for several hundred years after Newton were shunned, but are now recognised (i) to have involved far greater knowledge and subtlety than was immediately realised, and (ii) to have served as a crucial precursor to the very possibility of the emergence of an intellectually satisfying, theoretical *mechanistic science*.

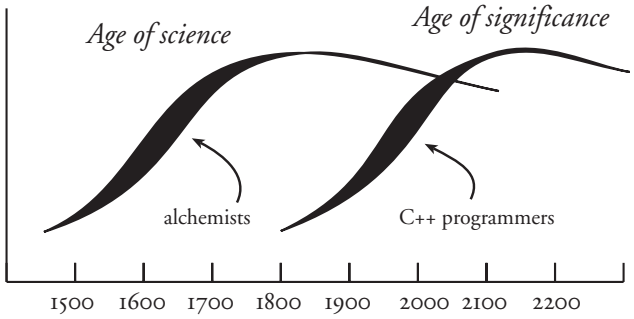

Figure 2 — The Age of Significance

What about computing—and the intentional sciences? They are represented on the right. Whereas traditional natural science embraced causal explanations, and dealt with mechanism and matter—i.e., with physical stuff, in its myriad forms—the intentional sciences, as I've said, involve a dialectical

interplay of mechanism with issues of meaning, interpretation, semantics, and the like. That is: they deal with issues of **significance**. Moreover, what computer science is, I believe—what history will recognise its 20TH-century instantiation to have been—is an experimental, synthetic precursor to the emergence of this new theory: i.e., something I call **semiotic alchemy**. Think of all those C++ and Java hackers, trying to turn web pages into gold! And think, too, of the profusion of rag-tag, untheorised practices, conducted in cottages, basements, and garages, that constitute "computing in the wild": they embody a vast wealth of pragmatic and practical understanding—just not as-yet very well understood. It all makes such good sense. And it is important. Computers are *laboratories of middling complexity*, between the frictionless pucks and inclined planes of mechanics, and the full-blooded complexity of the human condition, in terms of which to experiment with, and come to understand, the primordial intentional dialectic.

In sum, we are leaving the **age of** (purely mechanistic or physicalistic) **science**, and entering an **age of significance**—an age in which mattering, living truly (as much as speaking the truth), importance, etc., will take their rightful place, in our intellectual understanding of the world, alongside matter, materials, and mechanism. I am not quite sure whether to call the new age "science"—i.e., to assume that the term 'science' will broaden to include the new kinds of understanding we are talking about, or whether 'science' will retain its use for the kinds of causal explanation and physical phenomena legitimated in the past 300-year era, with something else introduced for the intentional variant. But since 'scio' is just Latin for 'know,' "natural" presumably means "whatever isn't supernatural," and thinking and referring seem to me about as natural form of activity as there are, my hunch is that 'science' will be extended to incorporate the new sense. However: just as I started with an apology to those allergic to the term 'religion,' I want to be the first to say, to those allergic to the term *'science*,' that the transformation in understanding we are on the verge of—if the diagnosis I am suggesting is right—is of almost unparalleled magnitude. It is liable to overwhelm any reservations people have (perhaps correctly) had about the impossibility of explaining ultimate questions in scientific terms.

4. Conclusion

Return to where we started. The call to arms is not that science should

change its tune, and take on a radically new set of questions. Nor is it that scientists should set their scientific work aside, and devote one day a week to religion. In a way, the call to arms is scarier than that. Whether it realises it or not (whether we *want* it to or not), science is *already aiming in a direction* where it will take on questions of unprecedented weight.

So the gauntlet—the call to arms—is simply this: will our answers be strong enough to sustain us? As scientists, academics, intellectuals—are we worthy to the task?

Now I haven't answered that question today. All I have tried to do is to *ask* it—and make it compelling. Even to suggest an answer is a topic for another occasion.[5] It requires moving past these skeletal considerations, to delve into substantive details. I will say, in passing, that this is where my own technical work begins: as an attempt to answer to these requirements. So what I have been developing is a story of objects and objectivity, self and subjectivity, first-person and third-person and first-person plural, practice and projects, deixis and flux. But I've written about that elsewhere,[6] and anyway we don't have time for it here.

Instead, to draw today's discussion to a close, I just want to hint at a direction in which I think an answer may be found—by drawing out one final theme that has permeated the discussion, but that needs to be acknowledged in its own right.

l said earlier that one characterisation of intentional phenomena is that they are subject to norms. Another, due to Brentano, is that intentionality involves **orientation** or **directedness**. We've already seen this in the case of semantics: to be "about" something is to be *oriented towards* it; to think about something is to cast your mind that way; to mention something involves a directed commitment. But orientation—a profound kind of directedness—is even more powerful than those examples suggest—more important, even, than Brentano may have realised. Rather than talk about it in the abstract, though, I want to illustrate it by considering one of the quest6ions with which we started: what it is to be a person.

We are here, to start with. Not delineated as such, in the physical force fields, but physically instantiated nevertheless. We are *in* and *of* the world.

---

[5]See "God, Approximately: Part II — The Answer" (forthcoming).

[6]*On the Origin of Objects*, MIT Press 1996.

Moreover, as intentional agents, we exploit the capacities of our physical embodiment to transcend the proximal, causal limits of that same embodiment. That is: we reach out to start taking the world *as* world. It is that inchoate directedness—that fact that we're not just *in* and *of*, but also *about* the world—that starts us on the long and difficult road towards humanity.

Rocks are part of the world, too. But, unlike us, they are not directed towards the world. They are infinite, too; and infinitely various—but they don't know it. In a way, there is nothing less transcendently splendid about a rock than about any of the rest of us. In that the Buddhists are right. But rocks are immune to their transcendence. Nothing much *matters*, for rocks. That's in part because rocks aren't even *rocks*, for rocks. Not being abstracting, intentional creatures, they are incapable of distilling the world's richness into predicates like "rock," much less into individual entities like *a* rock.

We *can* take rocks to be rocks. Doing so requires commitment. Moreover—and this will matter—it requires commitment not just to the rock, but also *to the world in which the rock exists*. In order to abstract anything as an object, that is—in order to construct material ontology—we have to be committed to that out of which, and that in which, we objectify. It follows that, in order to say anything at all—in order to refer, in order to stabilise an object as an object—an agent must literally be committed to that which cannot be "said." For the "world," in this sense, is not the "post-intentional" world of thereby-arrayed materiality. It is something prior and more profound. Something like the "ground of being" that, in our faulty, partial, perspectival, self-interested way, we take to *host* the rock. The transcendental grounds for the possibility of objects—if one likes Kantian language. Or the world of "no-thing-ness," if one's preferences run Buddhist (remember: the gauntlet is to develop a global perspective). Perhaps it is a noumenal world, except that to cleave phenomenal appearance from noumenal reality sounds suspiciously like one of those formal distinctions we have been at such pains to eschew. Whatever: it is a world of norms as much as of objects, a world of mattering as well as of matter. Maybe it is a world of enchantment—maybe even a world we can re-enchant. For if we are going to answer the gauntlet, without losing 300 years of insight, it is the only world we've got. It had better be strong enough to sustain us.

Which brings us to the title of the talk. Put it simply. The world that sci-

ence (without knowing it) is leading us to—the unutterable world—is far *less* like the mundane physical world classically set up in contradistinction to God, and far *more* like what the religious traditions, if I understand them, took God to be. It is because of that fact that it has a prayer of being sustaining.

Moreover, it is because it is this unutterable world of significance that our *commitment* is to that *orientation* is such a powerful notion. For notice how much directedness covers. Directedness underlies (and is prior to) purpose or *telos*—directedness in time. Directedness underlies (and is prior to) reference and truth—directedness to *what is the case*. Directedness underlies ethics: loving someone, fighting for justice, treating the world with kindness. And it underwrites curiosity, wonder—even reverence and awe. It is this common application—to truth and compassion and justice and generosity and beauty and perhaps even to grace—that makes Brentano's suggestion so poignant. All sorts of virtues—getting out of bed in the morning and lending a hand; shouldering responsibility for family and friends; taking things seriously; accepting responsibility for the consequences of one's actions; *accepting responsibility for the inevitable violence one does to the world by describing it at all, accepting responsibility for one's unutterable particularity*—all these things involve orientation and directedness: orientation up and out of oneself, to the encompassing world as a whole. Not to the world as *other*, since each of us is inexorably part of that world (the world has no *other*). And not to the world as *object*—for the world is not one, in any sense in which it could have been two. But orientation to the world *simpliciter*: the world entire, the world of which we are a part—a world that so spectacularly defies description that the very notion of "description" is defined over and against it as a way of watering it down. A world of matter and the world of mattering, a world in whose significance our own significance rests, a world unpredictable and risky and hard to master, a world to fight for and preserve, a world to struggle with, play in, defer to.

It may not be God. But it might be enough.

end of file