

A number of manuscripts have been circulated under this title over the last 15 years. This one lacks any sketch of a world view exhibiting the characteristics described—somewhat in response to the first version, which tried to present such a view, without explanation of what was interesting or mattered about it. If it seems worthwhile, I may someday incorporate all the various versions into a single short monograph.

God, Approximately

Brian Cantwell Smith*
University of Toronto

Slowly, but inexorably, science is encroaching on territory traditionally held to be religious. Scientific accounts now regularly deal with such topics as what it is to be a person, the origins and character of altruism, the nature of cooperation, sexual preference, marital and tribal fidelity, social allegiance, subjective experience, first- and second-person perspectives, ecstatic and meditative states, justice, cultural health, and the nature of consciousness. Indeed, the window on subject matters not deemed amenable to scientific exploration is rapidly closing.

If understood as expansionist moves of mechanist science, these developments may seem—especially to a religious person—to reduce hallowed categories to soulless physical arrangements, to desecrate the sacred, to banish humanity once and for all to the void.

I will argue against any such dismal conclusion. But I will do so neither by defending causal scientific explanations, nor by preserving a place for religion as traditionally conceived. Rather, I be-

*Coach House Institute, Faculty of Information, University of Toronto, 90 Wellesley St W, Toronto, Ontario M5S 1C5 Canada.

Comments to: brian.cantwell.smith@utoronto.ca

Last edited October 22, 2009. An earlier version was published in W. Mark Richardson and Philip Clayton, eds., *Science and the Spiritual Quest: New Essays by Leading Scientists*, Routledge, 2002, pp. 207-28.

The paper is dedicated to the memory of my father, Wilfred Cantwell Smith (1916-2000), who not only inspired my interest in these and many other topics, and taught us all what it was to be intellectual, but read and

lieve—and will attempt to argue—that these expansionist developments signal something much more profound taking place in current intellectual history. They presage a profound shift in the foundations of science—as momentous as anything that occurred in the 16th and 17th centuries, when natural philosophy and theology originally sundered.

If we recognize and embrace these impending changes, we will be in a position to see the threat to what matters ease, and the way open up to an increasingly urgent intellectual task: of reconciling our scientific, intellectual, and religious world views.

A note, in advance, 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 questions that have traditionally been asked in theological contexts—about ultimate significance, what it is to be a person, the grounds of justice and humility, what to live for, and for what to die—are too important to leave hostage to pro- and anti-religious debates. I don't care, in the end, who professes religious affiliation, and who forswears it. What matters is whether—and how—we respond to those ultimate questions.

1 Introduction

Consider the rise of the religious right: radically conservative Christians, ultra right-wing Zionism, Hindu and Islamic fundamentalism, militant Buddhism, etc. As well as involving issues of politics, economics, and tribalism, these extremist movements are responding to—and exploiting—a widespread social dissatisfaction: a feeling that reigning secular, scientific, economic, and political world views have failed to provide a sense of what matters: ways to tell right from wrong, guidance to anchor individual lives, the wherewithal to resolve thorny ethical dilemmas.

I find many of the fundamentalists' answers appalling: bigoted, mean-spirited, scary. But what are the rest of us—we on the left, we intellectuals, we scientists, we in the academy, we of economic

commented on a draft just a few months before he passed away. Although our perspectives were superficially very different, one need not dig deep to see the overwhelming debt I owe to him. I promised him, on his deathbed, that I would endeavour, as best I could, to “run the family store”: to sustain his tradition of wrestling with questions that matter.

privilege¹—doing about this heartfelt lack? Do we really think that a mixture of mechanist philosophy and late neo-capitalism, gilded with a dollop of liberal humanism, is a strong enough framework on which to build a just, compassionate, and humane society? If not, and if we do not respond to the yearning—if, willfully or unwittingly, we remain blind to the unsatisfied hunger—then we have no leg to stand on, in criticising others’ responses.

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.

Now you might take this as a call to arms: that scientists and intellectuals should set aside their professional concerns, and take a stand.² You might especially think this if you subscribe, explicitly or implicitly, to what we can call the **classic dialectic**: the roughly Cartesian supposition that we can divide our understanding of the world into two parts: (i) a roughly causal, deterministic, value-free, third-person theory of material objects—a knowable, empirical science of the physical world; and (ii) a more mysterious, phenomenological or experiential, inexorably subjective or first-person, value-laden theory of spirit.

A realm for the **body**, and a realm for the **soul**. It is a convenient dualism, an historically entrenched dualism—and still, unfortunately, a widely-accepted dualism.³

¹This is not to suggest that these five groups are coincident; but neither are they independent.

²Although my primary teaching is now in philosophy, my training was in science, as no doubt shines through.

³As is evident to any contemporary intellectual, numerous analogous dialectics crosscut the academy. Perhaps most visible is the resilient tussle between the constructivist, “embedded in language and culture” views of knowledge most often associated with literary and cultural studies, but also embraced in various forms in feminist epistemology and science studies, and the (to some surprisingly) stubborn mythos about the value- and culture-free objectivity of scientific truth. The shadow of C P Snow’s two cultures also still stalks academic corridors, though (in part for some of the reasons being explored here) it has more trenchancy these days in a contrast between first-person feelings and passions, on the one hand, and

That such a dialectic is still entrenched is betrayed by how many current views can be understood in its terms. Most obvious are two forms of extremist:

1. **Scientific triumphalists:** those who trumpet the idea that science is not only right, but *comprehensive* (ultimately capable of explaining everything that is real)—and that religion is not only wrong, irrational, and useless, but outright pernicious; and
2. **Religious ideologues:** those, including some creationists, who, perhaps hitchhiking a ride on poststructuralism or critical theory, remind us that science is “just a story,” and then, invoking Divine authority, personal revelation, or *a priori* intuition, assert the equal if not greater infallibility of religious myth.

Fundamentalists on either side are not the only ones who pledge allegiance to the classical dialect, however—between sacred and profane, between body and soul, between mechanism and spirit. Just as troubling, in my view—and equally in the thrall of the dialectic—are a group I will dub:

3. **Quietists:** those who take scientific and religious perspectives to be *complementary*, perhaps incomparable or even incommensurable, but potentially equally valid perspectives on a common underlying reality.

Neither view (scientific or religious) *really* challenges the other, the quietists say, since they are talking about different things—so why should they not be jointly and amicably held, in impeccable pluralist fashion?

Compatibilist quietism seems increasingly popular, as evidenced in the increasing literature of prominent scientists who will accept or even embrace the viability of religious affiliation.⁴

third-person studies and analyses, on the other, than between different objects of those affective attitudes.

For present purposes, though, I will focus on what I take to be the original dialectic (and better interpretation of Descartes): how our increasingly successful “mechanist” science has amplified Abrahamic tendencies to distinguish the sacred and profane, riving a realm of mechanism from a realm of spirit.

⁴«Cite Stephen Jay Gould.»

Quietism is also extraordinarily convenient. Do science during the week, go to services on the Sabbath—and laud science as showing the wonder of God’s creation. It is personally reassuring, too: to see one’s own life manifest the very epistemological pluralism and tolerance for which one argues in one’s op-ed pieces.

I believe that none of these stances will do—not only neither extreme version, but not the allegedly accommodating stance of the quietists, either. To put it as bluntly as possible: *the problem is harder than that*. The only tenable alternative, I believe—metaphysically, philosophically, theologically, personally, and politically—is to reject the originating distinction: between “scientific” and “religious” worlds, world-views, accounts, attitudes, etc. Or to put the same point another way: it is time to reject the meta-epistemological categories ‘scientific’ and ‘religious.’ Not only should they not be assumed to be natural or given; I believe it is too late, and the situation in the world too urgent, for them to be acceptable.

If on the one hand we are sensitive to what is left out in contemporary concoctions of mechanism, neo-capitalism, and liberalism, that is, and on the other truly recognize the significance of science’s expansionist moves into the realm of the sacred, then we have no choice but to start over, essentially from scratch.⁵

⁵I hear a cacophony of cries: that “Everything is already and always has been located in its particular moment in history. You cannot start over!” That is certainly true. What I am recommending that we subject to resolute critique, however, are not so much our world-views *per se* as two rather elite *meta*-world views: the idea that (some of) our views or understandings can be fit, unproblematically, into one or either of the two dialectically defined meta-categories ‘scientific’ and ‘religious.’ These abstract meta-notions are far less “always already” than the mundane categories of participatory life.

Some people—including those sympathetic to the epistemological critiques of feminist epistemology and science studies—may rightly feel that they have made progress in dismantling the categorical claim of ‘science.’ (Science and technology studies does study science; but it is not too far off base to understand it as the anthropological, sociological, historical and philosophical study of the reigning vaunted epistemology.) While there are perhaps fewer who would claim to have set aside the grip of the category ‘religious,’ sentiments along those lines certainly pervade discussions of Buddhism, Zen, and the other so-called “wisdom traditions.” What

2 Project

It is not hard to see why a physical/spiritual dualism might need to be replaced. Of many considerations, three are particularly forceful.

Start first, and rather flat-footedly, with something that has been evident to all of us since childhood: that the world, *au fond*, is ultimately *single, one*...or perhaps to put it better, *whole*. To say this is not in any way to deny what is right about cultural differences, to “disappear” the peculiarities of contingent experience, or to deny, as I would put it, that for different purposes, in different cultures, at different times, through different genres, and in different projects, we **register** the world in partially distinct, not certainly not fully intertranslatable, ways.⁶ To affirm the non-plurality of *the-world-au-fond*, that is, is neither to embrace nor legitimate imperialist master narratives, or even to grant pride of place (especially: uncritical place) to any preferred ontology. There is no reason to yearn for a single story, and many reasons to reject any pretenders to such a throne.⁷

But it is wimpy to hide the task of reconciling our ultimate

seems to me still rare, however, is anyone who, *without discarding what ultimately matters most about them*, has been able to demonstrate an easing of the categorical claims of both

⁶*On the Origin of Objects*, Cambridge, MA: 1996, pp. ■■■.

⁷Note that while we are right to reject the possibility of there existing *the* story of all that there is—i.e., *the* story of everything—that in no way implies that we need reject *any* story of all that there is, or *all stories* of everything. Hopi origin myths are not just about the origins of the Hopi people; they are myths about the origin of the world. As such they are legitimate: they are *the Hopi story* of how *everything* got started. What leads to trouble (imperialism, xenophobia, fundamentalism, and all manner of pernicious) is the idea that allegiance to one story in and of itself militates against allegiance to another—i.e., they it is intrinsic to the multiplicity of stories that they inevitably conflict and that (only) one must be right, rather than for example complementing each other in order to allow us to triangulate on what is in fact (even if perhaps only ineffably) the case.

How stories work is an epistemological claim. In §■■■ I will argue that the a number of recent scientific developments are working to bring epistemology within the purview of science, rather than remaining (as it has over the last several centuries) external to it. The claim made in the previous paragraph, therefore, will end up as internal, not external, to science-as-revised.

world views behind a convenient veil of pluralism. In fact it is not even intelligible to be worried about the rise of the religious right unless one recognizes the fundamental metaphysical unity necessary to the possibility of worlds colliding. If suicide bombers live in a different world, how is it that their bombs destroy our families? How could economic privation and the unjust allocation of opportunity be factors influencing the rise of fundamentalism, if economics, politics, and religious belief take place on incommensurably different strata?

The issue of how epistemological and perhaps even ontological pluralism can arise on (or be compatible with) an unsundered metaphysical foundation has bedeviled world-views for as long as they have existed. It is not just the purview of philosophy, science studies, or post-structuralist debate. Even the Christological image of Jesus as God-and-Man can be understood as wrestling with the same question. In the end, as I will argue, the question will be unable to be left outside of science-as-it-is-changing—left as an extra-theoretical topic of scientific interpretation, for late night debate. On the contrary, as I hope to demonstrate, it is rapidly becoming a science-internal issue.⁸ But that is for later. For now, it is enough to note that it is part of neither science, religion, nor commonsense to think, no matter how separate their views of the world, that what their respective views are views *of* can be so thoroughly and metaphysically ripped apart.

A second reason to challenge the physical/spiritual dialectic is the simple fact that all the world's major religious traditions (not just the Abrahamic) are based on natural philosophies hewn long before the rise of anything resembling modern science. As many have noted, the classical views of women, procreation, cosmology, and a myriad other subjects are sadly in need of being brought up to date in a wide variety of religious mythoi. It is nostalgic, at best, to assume that the profound changes in world-view that the human race has undergone in the last five centuries could (or anyway should) leave a religious myth intact.

In fact it is not too far off, I believe, to view the world's major religions as cultural and historical attempts by major civilizations

⁸Though 'scientific' in the sense that that endeavour is changing profoundly, as will be discussed.

to tell stories that rendered intelligible, in the best terms available at the time, and that served the most number of people, that framed and codified or embodied those culture's best understandings of ultimate issues. So why don't we forge a story, *now*, that embodies the highest reaches of our understanding of what matters most, but that frames that understanding in terms of a world view that builds on or out from the very best scientific accounts we now have? At least arguably, such a project would be radically truer to the depth of the religious traditions than to carry over their centuries-old stories as if they were unaffected by any developments in the last thousand years.

A third and not incidental reason to reconcile scientific and religious views stems from stark recognition of the state that the world is in: shrinking distances, rampant secularism, increased population of the global village, globalized cultural and market forces, etc. This is no more to say that we need a single story than it is to say that the world village should settle once and for all on a single kind of music. Appropriate, compassionate pluralism should reign. But the world is rent by strife and tension over allegedly competing conception—not only between and among the major traditions, but between religious and secular world views. If nothing else, we should have the guts to hold religious claims of truth, presumptions of competition, and tenacity of allegiance to the same standards to which we hold claims, presumptions, and allegiance for any world view.⁹

Here, though, I will focus on a fourth reason to reject the classic dualism, already mentioned several times: the fact that science is so rapidly encroaching on territory that has traditionally been understood from a religious viewpoint. From within the confines of any single discipline, scientific or not, or even from within the broader spectrum of the humanities or arts as a whole, it is difficult to appreciate how encompassing is the reach of these relentless and concurrent developments. Neuroscience is ablaze with

⁹This is not to say that there should be a single, univocal notion of truth—and *especially* not to say that all stores should be held to a traditional “scientific” notion of objective truth. But neither is it to relinquish the notion of truth to whim, convenience, or idiosyncrasy.

excitement about elucidating the phenomenology of consciousness; psychology is pushing headlong into the realm of emotions, including jealousy, love, and hate; gender differentiation and sexual orientation are staple subjects in biology; mathematical models are being developed of such basic notions as order, organization, autonomy, and self; evolutionary accounts are being developed of altruism, tribalism, and belief; analyses in analytic philosophy increasingly argue that information, language, belief, representation, normativity, and other (erstwhile?) semantic and/or intentional phenomena should be understood as *biological* categories, taking their place alongside digestion and locomotion; computational models traffic in meaning, representation, and interpretation; first- and second-person points of view are incorporated into sciences as diverse as medicine, physics, and computer science. And so on. Indeed, this encroachment is by now so well established that for young scientific minds it has taken on the status of the obvious. Not since the 19th century has “natural science” been restricted to phenomena that would traditionally have been viewed as “merely physical.”

As mentioned at the outset, some find such developments sacrilegious. And as I do for those who use religious language, I have some sympathy for their reaction, though in this case in fact I disagree with them. Sure enough, I am willing to hazard, these developments would be sacrilegious, if understood in traditional terms.¹⁰ If one unquestioningly understands them in terms of a classical understanding of what natural science is, that is, one is

¹⁰More strongly: current developments may actually *be* sacrilegious (i.e., may be accurately judged by future history *to have been sacrilegious now*) if we fail to respond to the challenges with which we are now faced. If, under the influence of the currently reigning scientific mythos, that is, we go on *thinking* that the understandings being yielded up by these on-going developments are purely mechanistic—i.e., if we go on meta-understanding them in purely causal terms—over time we are likely to *render* them entirely causal, nourishing and validating only their causal content, and eliminating or destroying anything that outreaches it, thereby killing any seeds of normative, metaphysical, and/or intentional insights they contain.

Still, this does not count as agreement with the (often automatic) reaction that the scientific domestication of classically sacred subject matters is intrinsically profaning. It is blind acceptance of that intrinsicity that I challenge.

likely to conclude that these new forays and encroachments are doing exactly what religious people fear: advancing a reductionist, materialistic picture of spirit, values, persons, and other entities traditionally held to be transcendent.

But here the epistemologies of feminism and science studies are right. To understand current intellectual progress traditionally is to put new wine into old wineskins. And it leads to ironic results.

Consider in particular our three species of latter-day dualist: scientific triumphalist, religious ideologue, and quietist. It is not surprising that the traditionalist approach (of understanding or incorporating these expansionist scientific projects within a traditional conception of science, without revising the core conception of what science is) is the explicit or implicit goal of the scientific fundamentalists. In a moment I will argue that it is doomed to failure—that, for purely science-internal reasons, such an approach cannot keep up with the very scientific progress it is designed to celebrate. But this is not the received position. Present-day writers of triumphalist persuasion would certainly embrace a “science as we know it” approach.¹¹

Similarly, and again not surprisingly, quietists also accept (in fact rely on) the traditionalist conception of science—that it is concerned with physical or mechanical phenomena, with causal explanation, etc. It is exactly because they support this traditional view of science that they understand it as not competing with religious doctrine.

Ironically, however—and perhaps surprisingly, but at any rate in a way that is important for us to understand—the traditionalist understanding of these new scientific developments is also, *and necessarily*, the approach of the *religious ideologue* as well. Although they deride scientific expansionism, they are as committed as any triumphalist to keeping (and being able to label) all such forays inside their “scientific” box. After all, the ideologue’s ability to proclaim the legitimacy (and perhaps superiority) of religious over scientific explanation, at least in the case of human

¹¹This is not to say that they would agree that a scientific account of such heretofore sacred notions as justice, consciousness, and normativity are reductionist or desiccating.

and/or sacred realms, depends on their ability to view scientific accounts not only as competitive (e.g., in the case of evolution), but as “less.” And given science’s undeniable success in the mechanical quarters of the pantheon, that “less” is invariably cast in terms of science’s restriction to physical phenomena and mundane causal relations.

In sum, to defend the “cause of science” in a physicalistic or mechanical way:

1. Allows the debate to remain agonistically framed in terms of the classic physical/spiritual dialectic, with very little likelihood of satisfying either side, triumphalist or ideologue;
2. Continues to prop up the quietists, wittingly or unwittingly, since it is the classical dialectic that underwrites the vantage point from which their compatibilism is staked; and
3. Because of what I take to be legitimate dissatisfaction with the adequacy of the mechanist/neo-capitalist/humanist option, *plays straight into the hands of fundamentalists.*

I worry, that is, to put it as starkly as possible, that **retaining** allegiance to a physicalist/mechanical understanding of science will lead to more of something of which we have already seen too much on the world stage: intellectual understanding, progressive politics, and compassionate commonsense falling to the religious right.

What is needed is something radically different. We need to develop new foundations, capable of meeting two simultaneous demands. We must:

1. **Update** our fundamental understanding of science, so as to accommodate new theoretical and technological developments, rather than assuming they will fit into a 19th-century conception of the “merely physical”; and
2. **Reconcile** the classic dialectic, by doing simultaneous and uncompromising justice to (i) what has traditionally been held within the purview of science, and (ii) what is good

and matters most in—i.e., what lies behind the valuable belief structures of—the religious traditions.

In the end I do not believe that these mandates are distinct. Not only can they be simultaneously met; at the deepest level they come to the same thing.

For the time being, though, it will not hurt to treat the two projects asymmetrically. There will be those who have a prior or autonomous interest in the second step, of reconciliation, and would thus be more likely to start there. It follows from the normative discursive standards within which I am currently operating, however, that an argument for the necessity of the second step depends on the result of the first.¹² Here, therefore, I will focus primarily on the first task: of updating our foundational conception of science.

Details will follow, but the overall shape of the reason why transforming the fundamental epistemology of science entails the reconciliation project can already be seen. By the time the updating is done, nothing will remain with which to retain allegiance to a distinction between the scientific and *that which matters most*—which I take (almost definitionally) to be the province of the religious. It will not turn out, in other words, as the triumphalist hoped, that science will remain special, and thereby take over the world. Neither will it turn out, as the ideologue wanted, that a preëminent place for classical religious positions will be secured. And as for the quietists' complacency, that will grow less and less tenable as the investigation proceeds. In the end science-as-revised will be so thoroughly permeated with mattering and hu-

¹²Roughly: any argument that is framed so as to meet normative standards currently governing science, analytic philosophy, and much of the rest of the current intellectual academy will likely need to be phrased in such a way as to address the first point (updating the foundations of science) en route to the second (reconciling science with "what matters" about religion). There are equally powerful reasons, in my view, to start with reconciliation, but until the first project has been completed, such arguments would be unlikely to meet current norms on scientific discourse.

This is why I say, in the text, that the two projects are in a sense the same, but from different perspectives. The more we have not yet accomplished either the tasks (one or both!), the more they seem different. Once they have been achieved, we arrive in a position from which to see that they have always been the same.

mility that it will no longer be special enough to be distinguished from a complete understanding of ultimately ineffable world.

But all of that is for later. For now it is enough to take stock of where we are now, to examine the consequences of a variety of current results for the epistemic project of scientific inquiry, and to make room for several impending new scientific and technological developments. Only when that is accomplished will it be possible, with any real plausibility, to see that adopting appropriately reconfigured foundations for what has heretofore been, and will henceforth be, scientifically understood will put us on the way towards meeting the former goal, of reconciling the intellectual and the religious.

One final clarificatory comment, before we turn to intellectual history.

Although I claim to address issues that have classically been

Christianity without stories

When asked whether he was a Christian, during his 1983 term as President of the American Academy of Religion, my father, to whom this paper is dedicated, replied: “Ask my neighbour.” This answer surely reflects a deeper understanding of Christianity than what is indicted in such anti-religious critiques as Dennett, Dawkins, and Harris[†]—or, for that matter, than what is being explored under the label “neurotheology.” But it is hardly a rare sentiment, having been recognized throughout the ages as what matters most about the ‘religion’—e.g., in such novels as ..., films as Rossellini’s “Open City,” and theologies as diverse as ...

It is that deeper orientation to the world that I am concerned with, not the stories, myths and rituals that have given it expression. Even as a child I rejected the stories, and as an adult the larger sectarian categories that separated the world’s religions, and so for many years have simply said “no” to the same question. But that does not mean that I find the tradition empty of insight. In fact I am sometimes tempted to distill what I take to matter most about the Christian tradition into a single, three-word sentence: that **love trumps justice**. There is a richness and truth to that maxim that, in my view (and without denying that this is the tradition from which I come) should not be lost in any metaphysics we forge for the future.

[†]«refs»

viewed as sacred or religious, in doing so I will not engage with what many readers may take to be virtually constitutive of religion. First, I will not take up any issues of stories, parables, myths, personages (human or divine), historical truth, etc. Nor will I tackle issues of history, politics, allegiance, community, relations between church and state, nationalism, persecution, strife—or even, for that matter, myth, ritual, practice, or identity (let alone meditation, prayer, or a myriad other notions often taken to be constitutive of religious practice). Similarly, even more strongly, I do not intend to deal directly with questions of faith or belief. I will set these things aside out of a conviction that I know will be contrary to that of most readers: I believe that these are all, far more than is usually realized, *surface trappings* of religion, resting on deeper, more fundamental, orienting world views.¹³ Even the category of ‘religion’ itself reeks, in my mind, of a colonialist attempt by a suspicious West to bracket alternative world views.

If, by avoiding these topics, I seem not to deal with religion after all, so be it. Remember: I am not interested in the ‘r’-word.

3 The State of Science

Turn, then—without preconception, or anyway with as little preconception as possible—to the present state of science. As I will attempt to show, the tradition we are inheriting, in the new millennium, is radically unlike the classical image: the conception inscribed in the traditional dialectic. In fact modern science is more unlike that conception, I want to argue, than anyone has yet recognised.

Three developments are particularly important.

First, what gripped the imagination of the mythical person-on-the-street throughout most of the twentieth century were primarily results in physics: quantum mechanics, relativity, and uncertainty, in the first part of the century; quarks, chaos, strings, and fractals, more recently. Extremely odd stuff, especially at the quantum level: ten- or eleven-dimensional curved (maybe even branching) universes of strange forces governed by inscrutable logic. Even on a classical picture, modern physics is alien: a stupefying spray of interpenetrating waves of every conceivable fre-

¹³See the sidebar “Christianity without Stories”

quency—turbulence, attractors, vortices smashing and piling up on top of each other in dizzying disarray. Imagine falling overboard in a perfect storm—and opening your eyes to nothing but salt and spray. *Now subtract you.* It’s a little like that, out there—only a million times worse.

A far cry from tables and chairs, obedient street lights, moribund committees, and 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 will mean our familiar day-to-day realm of medium-sized macroscopic objects: people, cars, continents, lunch, and elections (and possibly even such arguably concrete entities as fame and *détente*¹⁴). It is possible that the material world is, or can be, derived from the physical world by abstraction or synthesis—but “abstraction” and “synthesis” are concepts in epistemology, not in physics, and so themselves stand in need of explanation.

So that’s the first thing about twentieth-century science: the physical world is surpassingly strange, and quite unlike the material world. Many (most?) of us believe that the material world is nevertheless the same world, overall, as the physical world, in at least this sense: that if the physical world were to disappear, the material world would go with it.¹⁵ This is one of the commitments generated by the comment made in the last section: that *au fond* the world is whole. But how it is that the *world-qua-material* rests on (is founded in, emerges out of, whatever) the *world-qua-physical* is a spectacularly complex question—to which we certainly don’t know the answer, yet—and in detail may never (because, among other things, of issues of finitude and complexity).

¹⁴I believe that fame and *détente* are concrete, not abstract, not only because if the physical universe were to go up in smoke, they too would cease to exist, but also because they are essentially temporal. If relativity taught us anything, it is that time and space are not delivered in separate packages.

¹⁵This is an informal characterization of what in philosophy is called “global supervenience.” What matters most about supervenience, as a notion, is that it is *not* reduction. That is: the concept of supervenience is a product of philosophy’s efforts to understand how one thing (the material world, in this case) can “rest on” or “arise out of” or “be founded in” or “emerge from” another (the physical world), without the former being reducible to the latter.

In particular, to take just one example, it is a commonplace that in “packaging up” parcels of the plenum consisting of a maelstrom of entangled, superimposed quantum waves into finite, discrete mesoscale objects we ignore (“abstract away from”) an almost unbelievable amount of complexity. But what details in particular do we abstract away from? What principles do we rely on, to simplify in this way? And why? These are the kinds of things that a theory of the material world would have to explain.

This divergence between the material and the physical has a powerful corollary. It is commonly thought that we understand the world in terms of *objects*—discrete reidentifiable individuals, exemplifying properties, standing in relations, grouped together in sets or collections, etc. To borrow a phrase from philosophy, it is universally assumed that taking the world to consist of objects is part of our *natural ontological attitude*. What is implicit in what was said above, but is worth highlighting, is that in doing this we part company with our classical mechanical sciences. *Science does not treat the world in terms of objects.*

Imagine an old redwood tree, from the base of which numerous younger trees have sprouted, sharing the same root system. Suppose some are large, virtual trees in their own right; others are mere shoots, which could easily be pruned away; and others occupy the full continuum between the two. How many trees are there? One, with off-shoots? Many? And if many, what criterion is used to separate the shoots? *Biology doesn't care.* Or imagine, similarly, a bacterium dividing through a process of mitosis into two. Did the old bacterium die, with two new ones being born? Or is there now just one, physically disjointed? Or if just *one* new one is created, which is the original? The answer might matter if bacteria had developed cultural institutions leaving inheritance to the oldest child. *But once again biology doesn't care.*

Even in physics, where one imagines gravity acting, for example, on the “center of gravity” of an object, in point of ontological fact the gravity acts uniformly¹⁶ across the entire 3D volume occupied by the object—and acts *qua* volume, not *qua* individual. It is an *epistemic convenience* that, for purposes of gravitation, we can treat a solid volume as an object with a center of gravity—but be-

¹⁶“Uniformly” in the sense that the laws are uniform; some parts of the mass will be further from others parts exerting gravitational pull, etc.

fore it can be used, the legitimacy of that convenience has to be proved, with respect to the “real” facts. There are no facts about objects, centers of gravity, etc., in the ontological claims of gravitational theory.¹⁷

The second scientific development that is important to our project that has developed over the last century or so, is also relatively well-known: scientific knowledge has proved epistemologically recalcitrant—far more so than was classically expected. Traditionally, a sharp line was taken to divide knower from known. But that simplistic model has turned out to be unsustainable, in principle as well as practice.

Relativity brought perspective into scientific claim; quantum mechanics shattered the myth that measurement was innocent; experiments proved more invasive than suggested by glib disquisitions on the empirical method; debates rage about whether mathematical entities are legitimate part of the subject matter or better understood as mere theoretical equipment; considerations of (epistemic, effectively, because of being equipmental¹⁸) tractability permeate scientific models.

In human affairs, of course, many of these facts—the perspectival character of knowing, and the violent character of finding out—are almost truisms. What has been sobering, over the last hundred years, is the realization that these issues are deeply implicated in science as well: in the measurement problem, the collapse of the wave function, quantum indeterminacy, complexity bounds. Maybe all knowledge is violent, perspectival, implicate.

One way to understand the the most important scientific,

¹⁷We write the law of gravitation as indicated in part (a) of the figure to the right, in other words. But that is an epistemic idealization. Properly speaking—i.e., to best reflect the ontological claims that the calculation relies on—we should write it as the double triple integral indicated in part (b) (‘d’ is a spatial mass density function).

$$\begin{aligned}
 \text{(a)} \quad F &= G \frac{m_1 m_2}{r^2} \leftarrow \begin{array}{l} \text{epistemic idealization} \\ \text{truer representation of reality} \end{array} \\
 \text{(b)} \quad F &= G \iiint_{\substack{x_1, y_1, z_1 \\ m_1}} \iiint_{\substack{x_2, y_2, z_2 \\ m_2}} \frac{d(x_1, y_1, z_1) d(x_2, y_2, z_2)}{(\langle x_1, y_1, z_1 \rangle - \langle x_2, y_2, z_2 \rangle)^2} .
 \end{aligned}$$

¹⁸I.e., they have to do with the process of modeling, not with the phenomenon thereby modeled.

mathematical, logical, and computational achievements of the twentieth century, in fact, is to see them as a succession of nails driven resoundingly into the coffin of the enlightenment dream. The century may have opened hope still alive that science would deliver a transparent, deterministic, objective, value-free account of at least the physical world. But starting with the Frege/Russell discovery of foundational problems in set theory, eventuating in Gödel's incompleteness results in mathematics, it became clear that, even at the mathematical level used to formulate scientific theories, the idea of a perfected, closed system of mathematics would forever exceed our grasp. And then the concrete empirical results started rolling in: relativity at the large scale, and quantum mechanics at the small, both undermining the idea that physical nature could be completely and effectively grasped. So too for the Turing non-computability results, about which there will be more to say later: even determining simple answers to seemingly simple questions about simple programs proved beyond the capabilities of any finite mechanism. So too for dynamics extended past the simple linear case (initially called 'chaos theory,' then 'non-linear dynamics,' and now simply 'dynamics,' out of a recognition of how little of the world is justice to by linear dynamics), where even the tiniest alteration in initial conditions can have arbitrarily major consequence at a different place or time. This, too, was epistemically sobering, since all measurement ignores some detail—and it became clear that there is no way to guarantee, in any real system, that the approximations and idealizations that we are forced to use, in virtue of our finitude, will not miss major (even cataclysmic) characteristics of those systems.

At the epistemic level, that is, twentieth century science was more than anything else a lesson in *humility*.

The third development to which I want to draw attention (besides ontological and epistemological issues in physics) is the rise of what I will here call the **intentional sciences**: those dealing with *symbols, meaning, reference, interpretation, content, 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 nineteenth 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 twentieth century, unleashing modern logic, meta-mathematics, psychology, linguistics, cognitive science—and the entire computer revolution.

The nature of the intentional sciences, and their consequences for our understanding of the scientific world view, and hence of the prospects for reconciliation of scientific and religious perspectives, are so important to this argument that I will devote the whole next section to them, rather than summarize them here. Suffice it to say that their incorporation of issues of semantics, meaning, and interpretation, rather than just mechanism and cause, will deal the classic mythos of science a blow from which it will never recover. But as predicted in the opening sections, I will argue that it is a hugely positive development, that if properly stewarded offers us the promise of radically opening up our understanding of the world.

So those are the three features of the last 100 years of science that I want to highlight: (i) the ontological weirdness of modern physics, (ii) the epistemological recalcitrance of physics in particular




Figure 1 — Classical Science

(and science in general), and (iii) the rise of what I am calling the “intentional sciences”: logic, psychology, linguistics, computer science, linguistics, and mathematics (underwritten by a dash of philosophy). As we’ll see, the three developments are related. They also have an interesting implication. On the classical image (figure 1), knower was viewed as external to known. Physics has been epistemologically problematic in part because, as I said, physicists (and their practices) are implicated in physical knowledge, but physics is not a science that can explain knowing or practice. So physics has lived with an unresolvable predicament: though it aims for completeness, it cannot explain itself. It takes the intentional sciences to complete the picture, and resolve the

tension (figure 2). Since scientific knowing is one kind of knowing, and knowing is intentional, extending science to include intentionality brings doing science inside science.

4 State of the Intentional Art

The ontological oddity of physics fits largely within the dominant conception of science as a mathematical theory of the physical/causal world. Its epistemic consequences (for the doing of physics) are also substantial, but as indicated in the figures at the end of the last section, they haven't been considered to be physics internal. The other two developments—the recalcitrance of epistemology, and the reconfigurations mandated by the rise of the intentional sciences—have more structural consequences, which it is now time to explore. We have put epistemic and intentional phenomena onto center stage, and redrawn the boundaries of inquiry; but we haven't yet said much about what intentionality is like.




Figure 2 — Science as Revised

In this section I want to dig deeper into this third development. Only with more specific results in hand will we be able to see the full power of what is happening, in intellectual history, its impact on potential reconciliation.

The first task is to highlight two critical facts about intentional phenomena: a way in which they are, and a way in which they are not, “merely physical.”

The first important characteristic of intentionality—the non-physical aspect—is most obvious in the case of reference. Something that is so self-evident in our experience as to defy our imagining its 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 away, in time, space, or 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 U.S. President—without violating forward causality. I can refer to events outside my “light cone,”¹⁹ 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 most easternmost corner of the reception room in Mother Teresa’s clinic in Calcutta, and my reference will “land” precisely there, with unswerving accuracy. Reference is zippy, too—traveling, as Alonzo Church once said, “at the speed of logic”:²⁰ 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 situation we would be in if American presidents were elected by simple popular majority. And so on. Forget angels: reference rushes in where no one can tread.

There is a sense, that is, in which reference is physically transcendent—in the sense that it reaches beyond the bounds of even potential physical interaction.²¹ The amazing thing is that it man-

¹⁹Physicists define the “light cone” of an arbitrary point to be that region of space-time with which communication or transport is possible—i.e., with which interaction can happen at less than or equal to the speed of light. If you imagine your current self to be at a point in four-dimensional space-time, your light cone consists of two infinite cone-shaped regions: one, containing all events that could have affected you to date, converging in towards you, from the past; another, containing all events or points that you could possibly affect, from now on, spreading out away from you, into the future. “Outside your light cone” refers to the infinitely large remainder of the universe not contained in those two conic sections.

²⁰Public lecture, Center for the Study of Language and Information (CSLI), Stanford University, Stanford, California, May 3, 1984.

²¹Someone might object that even if someone is currently outside the light cone of that to which they refer, either they, or someone else, has been or will have been or in physical interaction with that referent. Similarly, someone might also say that we are in daily physical contact with the sun, because of sunlight. But both objections miss the point. First, the enduring (or perduring) nature of individuals, the structure of society, etc., are not physical constructs; and are unlikely themselves to be explicable without reference to intentional capacities and practices, so with respect to the present it merely shifts the place of the argument to bring them in. And as for the sun, how do we know that sunlight comes from the sun? How can we think that thought? That is what is at issue: how does our thought “reach over” the immediately impinging liminal press of the immediate

ages to do these things without requiring spooky metaphysics. That is—and this is the second aspect of intentionality I want to highlight—there is a sense in which reference is also physically immanent. Referring outside your light cone doesn't contradict natural law. Somehow, without violating the inexorable spatial and temporal locality of physical law,²² we are able to direct our thoughts outside the confines of the $1/r^2$ envelope to which physics restricts physical engagement.

Figuring out how we do this, I believe, is as important an issue as any discovery in science. With considerable caveat, it can be imagined as breaking into two parts: (i) an issue of how words or symbols or thoughts or other mental activities “reach out and touch something,” as AT&T would put it; and (ii) how, if the ingredient words or symbols do refer, how a mere mechanism—a physical device or a body made of clay—can operate in such a way as to do sensible things in regards to those referential properties. To put it bluntly but simplistically, and without making any assumption that the answers to the two questions are anything but mutually co-constituted: (i) How, as physically immanent beings, given that we think, do we refer? and (ii) Given that we refer, how do we think?

The former question has by no means been answered. Partial theories and suggestions abound,²³ but none take on the ontological issue mentioned above, of how we refer to *objects*, given that reidentifiable individual objects are not “supplied” in any current science, let alone in physics. (Later I will argue that the two ques-

surround, so as to allow us to register a *world* out there? It is that very ability that, though unimaginably common, is also unimaginably impressive.

²²By “temporal locality” I mean immediacy: the fact that an event, at time t_1 , cannot influence another event, at time t_2 , without passing that influence through all intervening moments t_i between t_1 and t_2 . Temporal “action at a distance” is as forbidden by the laws of physics as spatial action at a distance.

There is this nagging issue: some non-local phenomena appear to be validated by modern quantum mechanics. In spite of some suggestive but extremely inchoate reasons to suppose that quantum non-locality may be essential for intentionality, macroscopic locality (spatial and temporal) remains the dominating “feature” of the world that intentional phenomena surreptitiously devise (without violating physical law) to referentially circumvent.

²³«Refer to causal theories of reference, descriptive theories I suppose, etc.»

tions—of distal reference, and material ontology—are related: that the abstraction constitutive of identifying an object as an object *allows* for that object to be “tracked” or referred to when out of contact—or to say it most starkly: it is not so much that we succeed in referring to objects, as that objects are that to which are successfully able to refer.)

As for the second question, considerable progress has been made over the last century or so. Whereas in Descartes’ time it was essentially inconceivable how a physical mechanism could be capable of rational thought, much of that original mysteriousness has since been eliminated. The achievement 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, their solution was codified within the formal tradition, and that in turn within logic and mathematics, and so their insights are not so much publicly appreciated.²⁴ But they can be extracted from the clutches of the formal tradition, to survive another day.

Here is the basic idea. Physical regularities—causes and effects—are, as I have 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

²⁴“Formality,” it turns out, or anywhere so I argue elsewhere, is a perversely abstract form of digitality: an assumption that useful theoretical categories can be completely, unambiguously, and absolutely divided—a particularly stubborn form of thinking in black and white.

—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 with respect to—develop hypotheses concerning, stand in appropriate relation to—that to which we are not, at the moment, physically 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 (correlations 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, in other words, 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 now is cheap; that’s what brute physicality gives you. Living in the there and then—that takes smarts.

Studying these things—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, how they can successfully take the world in which they live to be ontologically intelligible—this is the work of what I am calling the intentional sciences.

I have traced the origins of these developments to the mid-nineteenth century, but of course there are far earlier precedents: Hobbes’ famous comment that reason “is nothing but ‘reckoning,’”²⁵ to say nothing of calculating devices stemming back to the abacus. But in spite of these forebears, and Turing’s work in the 1930s, it has really only the latter half of the twentieth century that the project has been pursued in earnest.

So we have about fifty years of serious development under our belt. How are we doing? From one point of view, achievements have been extraordinarily impressive; from another, astonishingly modest. There is no doubt that history will look back on us and smile at our innocence and naïveté. Without siding with either judgment, I want to identify three developments that are important for our larger task.

²⁵Hobbes, Thomas ... 1651

4.a From formality to engagement

First, as the new century dawns, we are seeing something of a sea change in our understanding of symbols, meaning, thinking, interpretation, and the like. The models of reasoning and knowing that were dominant in the first decades of last century, and that reigned in cognitive science and ai as late as the 1970s, betrayed what we might call a rationalist view from nowhere.²⁶ 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: neat realms of distinct, well-behaved, durable objects, exemplifying properties and standing in clear structural relations (rather as logicians imagine, when doing 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 increasingly 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 paradigmatic intelligent activities, rather than hypothetico-deductive ratiocination. Driving around, not proving theorems, is taken as emblematic.²⁷ Finally, it is increasingly agreed—especially in practical trenches—that those neat ontological assumptions about the a priori structure of the world 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 scan lines emerging from a robot's television camera quickly dispels any such myth. On the contrary, a growing cadre of researchers believe that that

²⁶Nagel, Thomas, *The View from Nowhere*, Oxford University Press, 1989.

²⁷It 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, was 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.

figuring out how to parse or “carve” the environment into workable pieces or “objects,” in ways appropriate to the task at hand, is the major task that cognitive creatures face.

All of these changes—captured in such slogans as situated ai, embodied cognition, etc.—can be viewed as retreats from the aforementioned formal tradition: in various ways, they share a recognition that many classic dichotomies (inside-outside, symbol-referent, abstract-concrete, continuous-discrete) are partial: negotiated, plastic, problematic. Discovering and stabilising such distinctions, when necessary, in project-specific ways, is an achievement of the subject (robot, animal, person), and therefore cannot be assumed by the theorist as a priori or given. Fitting an appropriate conceptual scheme to the world may lie closer to the heart of intelligence than working within one that is already established.

So that’s the first result, regarding the state of the intentional art: intelligence is concrete, messy, and participatory.

4.b Computability

The second result has to do with limits. Gödel’s incompleteness theorems,²⁸ proved in the 1930s, rocked the mathematical and philosophical community. For the first time, it seemed, intrinsic limits were placed on possible intellectual achievement. But again, as with the “Big Idea” underwriting intentionality, the formulation of these insights in formal terms blocked appreciation, in the wider intellectual arena, of their true significance. One way to put the result is that semantics can never be wholly reduced to syntax. Pragmatically, though, in terms of the story told above, there’s a simpler way to say it: although 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. In all but the most trivial cases, proximal surrogates can never entirely capture what matters about distal subject matters.

This moral was driven home by the “complexity” results of the second-half of the twentieth century, having to do with how hard,

²⁸Gödel proved that no formal (syntactic) axiomatisation of mathematics could entirely capture what we intuitively (but correctly, most mathematicians think) take mathematics to be.

in terms of space and time, it is to achieve things that are theoretically possible. It turns out that these “relative computability” results are much more consequential for embodied cognition than the absolute results discovered earlier. It doesn’t matter that chess is finite, when the number of possible moves in a game is on the order of 10^{120} . Or for another example: the number of possible 16 × 16 bit cursors on your computer screen is approximately 10^{78} , which is 10^{42} times smaller than the number of possible chess games, but still almost a million times greater than the total number of electrons in the universe. Combinatorics can kill.

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 the ontological issue of **determinism** and the epistemic issue of **predictability**.

Classically, it was assumed—and I suspect most people still assume—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, at time $t+1$, and thus iteratively for any future time $t+n$. But it turns out that it is not so—at least not for any finite, concrete, embodied intelligence.

The problem is that the “calculation” of the consequences is likely to be too hard—not just a little bit too hard, in the sense that it would take a lifetime, or that a future supercomputer might be able to do it, but too hard in the sense that the universe is not big enough to contain enough machinery to pull it off—and even if every atomic and molecular vibration were exploited in aid of the issue, it might take vastly longer than the universe has existed or is likely to exist to figure it out. This is the point of the chess example: from any metaphysical point of view worth its salt, even some finite deterministic problems are effectively infinitely difficult.²⁹

²⁹Classically, it has been assumed that there are differences between small and large numbers, but that the real conceptual divide holds between the finite and the infinite—a familiar refrain in religious as well as popular imagination. One consequence of the last fifty year of experience with

The Optimal Universe

In computer science there is a notion of an “**optimum algorithm**” for figuring things out—a way of doing things that is provably faster than any other possible contender (known or unknown).[†] In a few cases, for relatively simple problems, optimal algorithms are known; often they are not. Even when we don’t know the algorithm itself, however, we can often prove properties that such algorithms would have—e.g., how fast they would solve the problem in question.

This raises the intriguing possibility of whether the universe running the optimal algorithm.[‡] What if there were no way, that is—even for God—to figure out the future, except by letting the world continue into the future, and then observing how it all comes out? “You can’t predict it; you just have to wait and see what happens,” in other words—is a sentiment entirely compatible with complete determinism. This puts paid to the idea often associated with determinism: that somehow, in a deterministic setting, there is a fact, now, about what the future will be like. Commonsense might think otherwise, but this would not be the first time that commonsense has been wrong.

[†]There are complexities: some issues of optimality have to do with time; others with “space”—i.e., with the amount of extra resources that the algorithm requires (e.g., how much memory it would use to keep track of things en route). But the complexities don’t affect the point being made here.

[‡]I have no idea whether this is true—and as far as I know no one else does, either. It doesn’t really matter, since the amount of information contained in the universe would swamp any concrete attempt to predict it in detail. The idea is instructive only because it forces a recognition of the profound difference between determinism and predictability.

Now it may be—in fact likely is the case—that physics is not deterministic, because of quantum mechanics’ famous indeterminacy. But my present point is not about the empirical facts; it is that, even if the system were deterministic, nothing follows about its being predictable. If one’s predilections run to stipulating an omniscient, infinitely powerful God, then perhaps he or she or it could know. But as for anyone or anything or any system that is part of this world, waiting for the results to happen may not be just the epistemologically most pragmatic, but metaphysically op-

computing has been an increasing sense that the important conceptual divide is in fact between small and (very) large numbers; that between the large numbers and the finite is of no great consequence.

timal (indeed: perhaps the only) way to know what will happen in detail (see the sidebar on the Optimal Universe).

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, through an astounding number of the major results of twentieth century science. Though not the sort of thing that can be captured in a simple formula or theorem, there may by now be no more thoroughly established result in all of science.

Taken together, these first two features of intentionality (finite, messy, embodied participatory creatures, subject to massively strong computability limits) radically undermine the classical image of “man” as a rational, all-knowing, *übermensch*. In its place we get an ever-deepening sense of a world of paltry, finite, embodied creatures, struggling to make their way around in—struggling to make sense of—the world around them, using intrinsically partial, flawed, perspectival, incomplete, knowledge and skill. These aren’t just practical limitations, either. Embodiment is necessary for reference, but intrinsically (and radically) limits epistemic achievement.

It is a humbling image.

4.c Dynamical norms

I want to argue that the intentional sciences have led to a profound third result—something that as far as I know remains virtually unremarked on, but that I believe is if anything even more consequential.³⁰ It has to do with norms. I haven’t said anything about norms, but to enter the realm of representation—description, language, interpretation, truth, etc.—is to enter a world of phenomena governed by what philosophers call normative predicates: asymmetric, evaluative pairs such *true* vs. *false*, *good* vs. *bad*, *working* vs. *broken*, *beautiful* vs. *ugly*, etc.—where one option is better, more virtuous, more worthy, than the other. Accurate descriptions are better than inaccurate ones; information is better than misinformation; helpful behaviour is better than un-

³⁰In fact I know of no explicit account of it; the formulation presented here is my own.

helpful behaviour...and so on. In fact one good definition of intentional systems 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 what I will call **statical**, in the sense of applying to (passive) sentences or thoughts—i.e., to *states*.³¹ Full-blooded intentional systems, however, are dynamic, and hence governed by what I will instead call **dynamical** norms—norms that govern process.

In formal logic, the seminal intentional science, the only dynamical norm that has received much attention is ontologically derivative, defined in terms of a statical norm. In particular, logic's dynamic processes (reasoning, deduction, inference to the best explanation³²) are mandated to *preserve truth*, where it's assumed that truth and "best explanation" can be defined independently of, and prior to, their preservation or production.

This explanatory strategy—of starting with (presumptively autonomous) statical norms, and then defining dynamical norms in terms of them—has been adopted by other intentional sciences. Economic models of rationality and decision-making, for example, use dynamical norms of utility maximisation³³—where utility is (again) presumed to be static, explanatorily prior, and autonomous. But the strategy doesn't generalize. And no computer scientist believes it. What computational experience teaches us is that things generally work in the opposite direction: the semantic content (meaning) of a symbol or expression or data

³¹By "statical norms" I don't mean norms that don't change, over time. Evaluative metrics on book design, or abstract art, may evolve through the ages, but they would still be counted as statical, on this typology, because what they are evaluative predicates on—books or paintings—are static things. I use the (simpler) phrase 'static norm' to denote norms that don't change—that are temporally fixed. By the same token, a dynamical norm is one that holds of or governs processes; a dynamic norm, one that itself changes, over time. Thus statical norms can be dynamic; dynamical norms, static. (The distinction is parallel to more familiar distinctions between 'historic' and 'historical,' 'strategic' and 'strategical,' etc.)

³²A norm, devised as a simple model of scientific theorizing, in which one aims to infer, from data or some other set of "facts," the best possible explanation of those facts.

³³I.e., a process of maximizing some utility: the worth of an investment, the pleasure of the participants, the good of society, etc.

structure is typically determined by (even: exists in virtue of) how it is used—i.e., by the role it plays in the overall system of which it is a part. Rather than define dynamical norms in terms of static ones, that is, programmers derive static norms from dynamical ones—in a (perhaps unwitting) endorsement of the Wittgensteinian maxim that “meaning is use.”

If we get our static norms derivatively from our dynamical ones, where do we get the original dynamical 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 so baldly, though a variety of alternatives are being explored. But the dynamical norm that is currently receiving by far the most scientific attention—in cognitive science, artificial life, evolutionary epistemology, research on autonomous agents, and of course biology—is **evolutionary survival**.

It’s clear how you get a dynamical 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.

The idea of resting normativity on evolution has proved almost irresistibly seductive to many “naturalistic” philosophers, who see in it a way of resting the normative on the natural—and thereby (so they believe) establishing even more securing an underlying causal or “scientific” picture of the world. The idea 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’s claimed, and not to make a “lub-dub” sound, because hearts were evolutionarily selected for their capacity to pump blood (or creatures with hearts evolutionarily adaptive because of the fact that their hearts pumped blood), not for their sound-making capabilities. Similarly, the function of sperm is to fertilize eggs because that’s 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 dynamical norm to explain all human norms: justice, altruism, authenticity, caring, freedom, and the like? I doubt it. But in a way that's not 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 what the dynamical norms are, in terms of which we understand activity as human (even: as humane). And that, I hope, is obvious: dynamical norms on human activity govern what it is to live—live well, be committed, do good, strive for what is right. That is: **ethics**. And not just ethics, but whatever governs whatever you do: ethics, wonder, curiosity, eroticism, the pursuit of knowledge for its own sake...and so on and so forth, without limit.

In sum, the intentional sciences' taking on of full-fledged dynamical normativity (our third intentional result) is unimaginably consequential. 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, 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, goodness. 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 or "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 statical notions as truth and reference, or limit itself to a hobbled set of dynamical norms (or even an isolated case, such as survival).

Moreover, to up the ante—in case this seems too mild—something else, if anything even more consequential, is implied by these same developments. And with this the pieces of the story all start to fit together.

I said that the classical model of intentionality 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've also said that few working scientists believe the 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 structures 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. That is, as anthropologists, phenomenologists, and poststructuralists generally realise: you can't identify the objects, first, and then tell a story about the lives people live involving them. Rather, the objects themselves—what things exist, what type of thing a given entity is, what differentiates one thing from two—can only be determined with reference to the lives lived in their terms. For high-level social entities this isn't surprising: date-rape didn't exist, I take it, for aboriginal singers of Australian song-lines; baseball's strike zone isn't part of the furniture of the world for earwigs. But the present claim is more radical: it says that what is the case for date-rape and strike zones is the case for food, clothing, rivers, people—perhaps even for the number four.

Ontology is inextricably linked to epistemology, in other words, and epistemology inextricably linked to ethics. And not just ontology and epistemology, in the sense of the study of things and of thinking. Things and thinking themselves are inexorably linked, and linked in turn to life, and the good. Fair enough; these are conclusions we can live 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 recognising that dynamical norms have explanatory priority over statical 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 as we saw in §2; so have norms, values, and mattering. And not just values, but fundamental questions of ontology and metaphysics, too—which of realism, irrealism, formalism, or idealism, is right, for example. Indeed, the answer to that last 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, your insecticides pollute my water supply, my car bomb rains on your parade.³⁴ Ontologically, though, our worlds are many. They are many because objects and properties involve synthetic abstraction—and synthetic abstraction is normatively governed. Our objects are as constitutively different (and, of course, as constitutively the same) as our projects, 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's no pun, in other words, or historical accident, that we use '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 must heal the temporary rift that for 300 years has torn matter and mattering apart.

5 The Age of Significance

In a moment, it will be time to combine these results about intentionality with the general scientific developments rehearsed earlier (§2), in order to take stock of the present and future state of science, and to ask about the prospects for reconciliation. But one issue needs to be addressed first—to deflect misunderstanding.

If you ask most scientists—including modern intentional scientists—whether they *think* they are engaged in partially irrealist metaphysics, probing the ethical structure of the human condition, or constructing scientific models of the good life, their answer, it's safe to say, will be: no. But that is not the question at hand. It doesn't matter, for our purposes, what people (currently) think they are doing. As social critics, philosophers of science, and writers in science studies have repeatedly emphasised, scientists are not usually very reflexively self-critical, or necessarily experts about the nature of their own activities. Rather, what will matter in the long run (irrealist sentiments notwithstanding) is what they—i.e., what we—are actually doing.

That said, I will confess that when considering the situation I am sometimes reminded of Road Runner, who ran off the edges of cliffs and then hung out there, motionless, for a moment, until he looked down—and only then fall. A great many cognitive scientists, computer scientists, evolutionary biologists, linguists, logicians, and the like, in my view, have already run off the edge of the “natural science” cliff, but haven't yet looked down. They have long since abandoned the allegedly “safe” terra-firma of pure, local, causal explanation and pure, unadulterated physical phenomena—mainstays (even bastions) of science as we know it. Locally—in the thick of moment-to-moment interchange—they defend a variety of sensible views, such as that what a data structure means depends on how it is used. What few seem to realise is how profoundly such seemingly innocent adjustments shake the foundations of what for 300 years we have called “natural science.”

³⁴Smith, Brian Cantwell, *On the Origin of Objects*, MIT Press, 1996, p. 100.

To make this concrete, it may help to consider the case of computing—one of the intentional sciences, as already mentioned, and (as it happens) my own area of expertise. For many years I have been engaged in a foundational inquiry into the 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. After decades of work, 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 is nothing there.

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-efficient 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, or have some or or *characteristic* identity—digitality, for example, or formality, or abstractness—that makes them a worthy subject matter for scientific investigation. That computation is special is implicit in the idea that there might be a theory of computation.³⁵

That is what I now claim we will never have. Admittedly, and somewhat distractingly, there does exist a body of work called the “theory of computation”—but I am now prepared to go to court to show that it is wrong, in the following strict sense: in spite of its name, it is not a theory of computing after all. That is not to downgrade it. What the received (so-called) “theory of computation” actually is, I believe, is something of incalculable importance, worth a passel of Nobel prizes: neither more nor less than a *mathematical theory of effective causality*. 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.

The ultimate problem, however, is not with current theory; it is deeper than that. We will never have a theory of computing, I

³⁵It is also implicit in what is called the *computational theory of mind*: the idea that underlies artificial intelligence and much of cognitive science—that we might be computers, too—in a way that is neither tautologous nor obvious.

claim, because there is nothing there to have a theory of. Computers aren't sufficiently special.³⁶ 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.

While that might seem like a dismal result, in fact I believe the opposite: that it is the most positive result that a computational triumphalist could possibly hope for. Moreover, it is a result that




Figure 3 — The Age of Significance

has a direct bearing on our present subject matter. In fact—this is why this digression was warranted—we can lift this “nothing special” result about computing, integrate it with the general results about intentionality surveyed in §3, merge

those into the long-term scientific developments catalogued in §2, and finally get a fix on what exactly is happening to the foundations of intellectual inquiry, as we enter the third millennium.

The situation is depicted in figure 3. Time runs along the bottom axis—from the 14th century up through the present and into the future. 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. To its left is an indication of alchemy: a rag-tag bunch of exploratory practices, conducted in a disheveled and atheoretical way, that for several hundred years after Newton were shunned as “unscientific,” but that are now recognised (i) to have involved far more knowledge

³⁶That is, computers in general: not just present-day computers, but past computers, future computers, computers we will build and computers we will never build—computers in the fully general sense (which is what a theory of computing, if there is one, should be about).

and subtlety than for a long time was realised, and (ii) to have served as a crucial precursor to the very possibility of the emergence of an intellectually satisfying, theoretical mechanistic science.

What about computing—and the intentional sciences more generally? They are represented on the right. Whereas traditional natural science embraced causal explanations, and dealt with mechanism and matter—i.e., with purely physical stuff, in its myriad forms—the full range of intentional sciences, as I’ve said, involve a dialectical interplay of mechanism with issues of meaning, norms, interpretation, semantics, ontology, and the like. That is: they deal with issues of significance. Moreover, what computer science (or computational practice) is, I believe—what history will ultimately recognise it to have been—is an experimental, synthetic precursor to the emergence of this new era: 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.

Computing, that to say, together with biotechnology, with which it is rapidly fusing, is building the laboratories of the intentional sciences. Together, these overlapping endeavours are allowing us to conduct experiments with language, meaning, interpretation, function, normativity, perspective, thinking—experiments of middling complexity, between the frictionless pucks and inclined planes of mechanics and the full-blooded complexity of the human condition. Computers give us raw material, in terms of which to experiment with—and thereby come to understand—the primordially intentional.

What this all means is that we are gradually but inexorably undergoing a social and intellectual transformation every bit as consequential as the Scientific Revolution. We are leaving the age of (purely mechanistic or physicalistic) science, and entering what I have called 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 the intellectual pantheon, alongside

matter, materials, and mechanism.

I am not 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, and something else will be introduced for the intentional variant. But given that ‘scio’ is just Latin for ‘know,’ “natural” presumably means “whatever isn’t supernatural,” and thinking and referring (once one lets go of the hegemony of physicalist science) is about as natural a form of activity as there is—and with an eye, too, to the power and prestige that the scientific enterprise wields in society—I would guess that the term ‘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 first in line 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.

Let me be clear: this is not science as anyone has ever known it in any recent century. It is something profoundly novel—something liable to change forever the reservations people have (perhaps correctly) had about the impossibility of explaining ultimate questions in scientific terms.

6 Reconciliation

Return, finally, to where we started. The call to arms, we can now see, 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. 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. Not just traditionally scientific questions of great social or moral impact (such as whether to develop nuclear weapons), but question of what it is to be weighty, what it is to be serious—even: what should matter. Traditional concerns of the religious traditions, that is, are being subsumed within the scientific juggernaut—independent of preference or

protestation. So the gauntlet—the call to arms—is simply this:

Will our scientific and intellectual answers be strong enough to live by, to sustain the world? Is our moral vision and political fiber tough enough—does our understanding go deep enough—for our answers to be able to counter the influence of the religious right? As scientists, academics, intellectuals—are we up to the task?

I don't yet know to answer these questions. But answer them we must. And if what I have said here even points towards the truth, then maybe—just maybe—answer it we can. The possibility of our answering it positively depends not on responding to the fear with which we started—the prospect of a desiccating science domesticating territories traditionally considered religious—but on recognizing the much deeper fact that science itself is in for an almost total metaphysical overhaul. If we can open our eyes to those transformations, and do a responsible job of stewarding it (and ourselves) through the upcoming changes, then—and only then, I believe—do we have a chance of succeeding at the project indicated at the beginning: of founding a vision strong enough to underwrite the world.

And remember: it is the whole world that is at stake: mattering, as well as matter—significance, as well as silicon.

For just a hint of how this will go, turn, one last time, to intentionality.

I said earlier that one characterisation of intentional phenomena is that they are subject to norms. Another characteristic, widely attributed to Brentano, was implicit in our discussion of reference and non-locality: the fact that intentionality involves orientation or directedness. We have 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 questions with which we started: what it is to be a person, what it is to be us.

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; that much is a consequence even of traditional science. As intentional agents, moreover, we exploit the capacities of our physical embodiment to transcend the proximal, causal limits of that same embodiment. That is: we reach out, commit ourselves, interact, and—in what may be our most impressive achievement—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 in and of the world, too. But unlike us, they’re not directed towards the world. Like us, they are infinite in various ways, though 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 own transcendence. Nothing much matters, for rocks. That’s in part because rocks aren’t even rocks, for rocks. Not being synthetic, abstracting, intentional creatures, they’re incapable of distilling the world’s richness into predicates like “rock,” much less into individual entities like a rock.

We, however, can take rocks to be rocks. Doing so requires commitment. Moreover—and this matters—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 most fundamental sense, is not the “post-intentional” world of thereby-arrayed material ontology.³⁷ Ontology, after all, as we have already seen, is plural; the world is prior, unitary—and more profound. Something like a Tillichean

³⁷I am hesitant to use the term ‘world,’ since it is so unquestioningly assumed to consist not of that which we synthetically abstract material ontology out of, but for the thereby synthesized ontology itself (the “mundane world,” one might say). I can’t exactly use ‘God,’ either—in spite of what is said over the next two pages. Perhaps we should just call it ‘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.

Put it simply. The world that science, without knowing it, is leading us to, or at least the world to which we are being led by whatever science is morphing into—the unutterable world, the ground of being—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 (or anyway, the ultimate realm of the sacred) to be. It is because of that fact, and only because of that fact, that we have a prayer of forging a sustaining intellectual vision.

Moreover, it is because our commitment is to this unutterable world 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, 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 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 one’s unutterable particularity; accepting responsibility for the inevitable violence one does to the world by describing it at all—all these things involve orientation and directedness: orientation up and out of oneself, to the encompass-

ing 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 (and anyway objects are post-intentional, whereas the world is prior). 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 a 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.³⁸

³⁸Thanks, for comments and discussion, to the members of the “Science and the Spiritual Quest I” workshop on information technology: Michael Arbib, Char Davies, Anne Foerst, Kevin Kelly, Mitchell Marcus, Mark Pesce, Henry Thompson, and (the late) Mark Weiser; to Margaret Wertheim and Billy Grassie for their participation—and to Mark Richardson, Philip Clayton, and Bob Russell, of the Center for Theology and the Natural Sciences at the Graduate Theological Union in Berkeley, California, for organising the SSQ conference, which occasioned its writing. For comments on more recent drafts, I am also indebted to John Coleman, Gillian Einstein, Will Oxtoby, Tom Settle, and Arnold Smith.