Subjectivity & Objectivity

Brian Cantwell Smith[*]
University of Toronto

1 Introduction

Two properties are required, for a creature to be intelligent—for us to extend to it a sense of "we":

1. It must be conscious, "awake"—there must be "someone home" behind its…cameras;
2. There must be a "world out there," for it: a world it knows, moves around in—even: is curious about, respects.

Subjectivity and *objectivity*, in others words—as long as we are generous about the meaning of the words.

Subjectively (as Nagel says), there must be "something it is like," to be it. It must have an "inner" life; a point of view; a perspective on the world that originates from a knowing, feeling, conscious[1] self—a self that it, and it alone, owns, occupies, and knows the world from. Objectively, it must make claims about—have confidence in, rely on—a world beyond its senses, distinguishable from rank hypothesis and flights of fantasy. A world about which to be right…and to be wrong.

This is one way to characterise my project: to understand, enough for us to build it, what is involved in having authentic subjectivity, and authentic objectivity.

---

[*]Coach House Institute, Faculty of Information, University of Toronto
90 Wellesley St W, Toronto, Ontario M5S 1C5 Canada
Presented as the Thomas Langford Inaugural Lecture at Duke University on December 4, 2001, in conjunction with taking up the Kimberly Jenkins University Professorship in Philosophy and New Technology.

Last edited: November 12, 2009

brian.cantwell.smith@utoronto.ca
Draft only (version 0.80)

[1]Prereflexively original

1 Acrobatics

An example, before delving into analysis—an analogy, on which the rest of the discussion will depend.

Imagine a dimly lit stage, in a darkened theatre, blanketed in a fine mist. On stage is a single character: an acrobat, wielding a flashlight (figure 1). Silently, the acrobat leaps and dashes about, twists and turns, flings herself through the air in crazy contortion. All the while, as she weaves about, she adjusts the orientation of the flashlight, so that its beam—a dagger in the mist—passes unerringly through a constant point in space, three feet off the floor, front and center stage. As the beam angles and glances, that single point of illumination fairly glows.

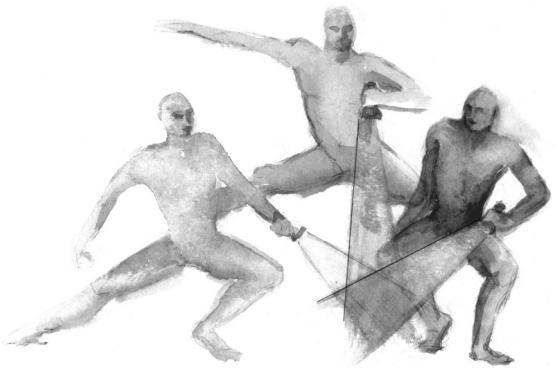

Figure 1 — The Intentional Acrobat

Six observations:

1. Performing this feat would take immense skill, but, as regards *orienting the flashlight*, very little energy.
2. Two motions are relevant: (a) of body (legs and torso); and (b) of wrist. In a sense, wrist motion "compensates" for body motion, with respect to stabilising that single point of illumination. Wrist motion is not opposite of body motion; nor is the wrist held constant (as if epoxied to space). Body and wrist are both lithe—they dance. The only thing epoxied to space is the target glowing fireball.
3. All motion—everything that "happens"—happens in the vicinity of the acrobat. What holds still is the point of illumination: 10, 20, … 100 feet away. From the acrobat's perspective: what is local, is not stable; what is stable, is not local.
4. Though what is stabilised is distal (to acrobat), it is the fluid, dynamic acrobat, not the stable illuminated point, that determines the boundaries of the point of illumination. (This will matter, for constructivism.)

5. Both (a) an explanatory theory of the acrobat, from the outside, and (b) the norms governing the acrobat's performance, from the inside (who was chosen to dance; whether the performance rates a 9.8) would refer to the stability of the distal, illuminated point.
6. If the acrobat stopped, the point would disappear. Nor could a photograph—a single time-slice—capture the phenomenon. Sans activity, nothing determines what position, along the beam, is stilled. Only through activity is the point triangulated upon, and thereby given (durable) identity.

Second example. Turn up the lights; get rid of the mist. This time, instead of a flashlight, the acrobat wields a pointer—a baton. Body movements are just as wild; twists of wrist, as ingenious. There is no illuminated spot, but the audience's attention—*your* attention—just as uncannily, would be riveted to the spot at which the acrobat irrevocably points.

Just two comments, this time:

1. In this version, *really* nothing happens—nothing measurable, at any rate—at the designated point. No energy is transferred, no causal connection exists, between acrobat and point (locus of activity and locus of stability). If the acrobat danced at the back of the stage, we could enclose both acrobat and pointer—the entire behavioural dynamics—in a "black box", across the boundaries of which no energy flowed, which excluded the pointed-at target. That is, we could contain the acrobat, movement, and pointing inside a "closed system." But even if we did this, (a) "what the acrobat is (intuitively) doing; (b) how the acrobat's dance would be judged, and (c) the "invariant" to which an explanatory theory of the acrobat's activity would necessarily advert, would remain external to the energetically bounded region.
2. No amount of local investigation at the point to which the acrobat is pointing—no microscopes, X-ray crystallography, exquisite sensors—could determine, up here at the front of the stage, what was being pointed at. In this (ad-

> mittedly limiting) case, there is a sense in which there is "nothing there," here: no independently-delineated entity or spatial region to serve as the "object" of the pointing. Yet that fact in no way undermines the existence or reality of this spot's being the spot *towards which the acrobat is pointing*.

A third and final example—or rather, suite of them. This time, replace flashlights and pointers with words. Yesterday, I thought to myself: "I have a talk to give tomorrow." Today, I refer to that day as 'today.' Tomorrow, I will think of it as 'yesterday.' Each night, there is a motion of my body—of my self—through time. Each night, compensating for this temporal shift, I adjust my brain: my mentalese. Neither my body, as referenced to the world, nor my mental or brain state, as referenced to my body, remain constant. *Constancy is death!* What is stabilised is once again distal: the day on which I gave the talk, the referent of my thoughts.

Similarly: I say "to my right"; you think "to his left." Last week I referred to "the tallest person in my class." This week, a basketball player signs up. Now, to refer to the original student, I have to say: "the second tallest person in the class." Or suppose I ask: "Would you hand me my glasses, there, on the table behind you?" You turn around. Now where are my glasses? *In front of you.* You unconsciously adjust your egocentric thoughts, to compensate for the movement of your body, in order to stabilise the in-the-world place where my glasses rest. (A good thing, too; without updating, you would go round and round forever.)

2 Logic

These examples illustrate what I call reference-tracking, or the preservation of reference.

Who studies such things? No one!

Actually, that is not quite true.

There is a discipline that studies something similar. Namely: logic—one of the intellectual success stories of the 20th century, and forebear of computer science, cognitive science, and artificial intelligence. Logic does not study reference-tracking per se, for reasons we will get to in a moment; what logic does study is *preservation of truth*, through streams of inference. And while (pace

Frege) I do not consider truth a form of reference, truth and reference have similar properties. Whether a claim is or is not true is not a local, energetic property of it, either. Microscopic investigation of the font will not tell you whether a contentious sentence is right. Whether dinosaurs were warm-blooded is settled—"located," we might even say—a long way away from here. If you submit a term paper containing an especially dubious sentence—unlikely to be true, or perhaps spectacularly shallow—my suspicions will not be relieved by X-raying the paper, to see how deep the claim's ink penetrated the surface of the paper. Truth is not a local property; it (almost always) involves *distal relations*.

Inference, in contrast, and the sentences or mental structures on which it operates, are, in a crucial sense, *local*. They are also *concrete*: physical, subject to causal forces. In order to construct an inference engine, theorem prover (or computer; more on that in a moment)—in order to build a device that can move, dynamically, from one statement in a proof or argument to the next—you must fabricate a local, dynamic mechanism, using small but nevertheless real bits of energy, not unlike the slight rotations of the acrobat's wrist. Crucially, however, (i) the *norms* or *specifications* that govern the transitions, in terms of which we evaluate the theorem prover, and (ii) the scientific theory that explains what the inference system is *doing*, both advert to the non-local, non-causal semantic *interpretations* of the sentences.

Mechanically, that is, because of powerful constraints arising from the physical world, inference is a local, concrete, dynamic activity. Constitutively, in contrast, in terms of meaning, semantics, and *interpretation*, logic is non-local, because the applicable norms—truth, falsity, soundness, completeness, validity, inference to the best explanation, etc. (all technical terms defined in logic)—make essential reference to the (stable) distal interpretation. Put it all together, and we can see what logic is: a *fundamental theory about the dialectical interplay of meaning and mechanism*.

A moment ago, I claimed that the development of logic was one of the great intellectual achievements of the twentieth century. This characterisation shows why: understanding the interplay of meaning and mechanism is one of the most challenging intellectual problems of all time. The problem, however—and it is very grave, especially as regards AI, cognitive science, and the present

interest in subjectivity and objectivity—is that logic dealt with this interplay of meaning and mechanism in an *extremely narrow way*. In particular, it studied the interplay solely within the confines of metamathematics. As a result, it ends up with esoteric and narrowly restricted results, formulated in abstruse technical theorems.

Underneath these theorems, however, and buried in the tacit conceptual frameworks in terms of which they are framed, lies a treasure trove of insights that apply, much more generally, to arbitrary meaningful or (as philosophers say) "**intentional**" phenomena: things like thought, language, representation, etc. Moreover—this is why I have been going on about these things—these buried results are profoundly relevant to our topic of subjectivity and objectivity.

Let me name just three of these buried insights:

1. Meaningful or intentional processes (including thinking) operate under *extraordinarily severe physical limitations*. It is virtually miraculous that evolution, first, and now science, in its footsteps, have figured out ways to exploit an almost vanishingly modest set of resources, so as to enable powerful intellection, so as to give people an astonishingly powerful abstract grasp of the world in which they are embedded.
2. The dialectical interplay of meaning and mechanism is a "relationally emergent" level of coherence or regularity in the world. It involves a fundamental collaboration between the two aspects, and cannot be reduced either to pure meaning, or to pure mechanism. This, too, will have consequences for the prospects of developing scientific accounts of subjectivity and consciousness. The non-reducibility of the intentional is apparent even in logic. Gödel's incompleteness results, Turing's non-decidability theses, theories of computational complexity, and a number of other scientific results (including results from chaos theory and nonlinear dynamics) demonstrate that intentional processes (such as thinking!), can never, except in the most trivial situations,[2] be reduced to proximal physical or mechanical

[22]«Quote Barwise on why completeness proofs are a sign of failure, not success.»

dynamics.

God read Browning. Meaning's reach always exceeds its grasp.

3. The "subject matter" or "content" of a representational structure—what a symbol or thought is about—is not, by and large, the proximal structure itself, nor any of the local activities in which it plays a causal role, but rather the distal situation towards which it (or the agent) is directed. Not *flashlight* or *pointer*, but *what flashlight is pointing at*.

Think about…oh, Sarajevo, ice cream, a long-lost friend. What occupies your mind—what your thoughts are focused on, what you are emotionally directed towards—is a *city*, *food*, or *person*—not a pattern of neural behaviour, not even an image. If you do not believe it, try coming at it from the other direction. Your boss emerges from a meeting with the dean about downsizing the department. You ask whether they discussed you. You would hardly be satisfied by the reply: "Not at all! I only talked about patterns of activity in my neo-cortex."

My brief, here, is not to applaud logic. In fact my very first philosophy paper[3] denounced its limitations, rehearsing any number of ways it is inapplicable to human thought. My only point, here, is to highlight some of its more tacit, but ultimately most powerful, accomplishments—what I will call its core insights. Because what happens next, in intellectual history, is very interesting. Those fundamental insights get lost.

3 AI & Cognitive Science

Logic, as I have said, has two preeminent offspring: computer science, and AI/cognitive science—both fields in which I work. I will describe them separately, because they abandoned logic's insight in separate ways.

In computer science,[4] logic's intentional vocabulary has been redefined to refer to something else: namely, *purely causal relations*—between programs and the processes they engender; and between and among states of machines. There is a long story to be

[3] «Ref 1978»

[4] As we explored in colloquium last week /«ref».

told about how this came about, but the bottom line is that computer science has *projected all intentional vocabulary back onto pure mechanism*. By doing so, computer science can honour science's general methodological predilection for causal accounts—but at a severe cost: genuine (non-causal) intentional directedness is thereby "disappeared." And as we have seen, to do that is to "disappear" the subject matter, since no mechanical reduction can do it justice.

The situation in AI & cognitive science is completely different. Rather than re-appropriating logic's vocabulary, AI and cognitive science have largely overthrown the logicist paradigm on which they were founded—for reasons we have already seen: its being massively too restricted and brittle to account for the sheer ingenuity of on-the-fly human behaviour.[5] What is at stake is often displayed as a list of opposites. Logic is seen as committed to a conception of cognition as: *individual*, *rational*, *abstract*, *disengaged* (from the world), *explicit*, *static*, *discrete*, *generic*, and *context-independent*. The alternative—often called a "situated" approach—rejects all of these assumptions, in favour of a claim that cognition (all human activity) is: *social*, *embodied*, *concrete*, *located*, *engaged*, *dynamic*, *continuous*, *particular*, and *context-dependent*. Something like *improvisational navigation*, rather than *rational intellection*, is taken as paradigmatic human competence. The new view also waxes hugely enthusiastic about the idea that human behaviour is "emergent": not the sharpest tool in the shed, in my view, but a term generally used to index the intricate, seemingly non-reducible patterns of (sometimes self-) organisation, as exemplified in modern non-linear dynamics and complexity theory.

I will label this wholesale (infamous) transformation in cognitive science the **situated sea-change**. I should say that I am completely in favour of it—it is something I have argued for for more than twenty years. If anything, I want something even more radical. For another assumption of logic, less often included in the laundry list, is its commitment to what I will call *"formal ontology"*—an assumption that task domains consist of neatly indi-

[5]Not that it is logic's fault; that field, after all, developed as a theoretical effort to put the foundations of mathematics on rigorous intellectual footing, not to explain how, in finite time, you can work your way across a crowded Tokyo subway station.

viduated, discrete, unambiguous objects, properties, and relations. As will become evident in a moment, formal ontology is not my favourite, either. I want to press for a more constructivist alternative.

In sum: I want to see the situated movement, and raise it one.

Problem is, in discarding logic, cognitive science not only rejected its untenably narrow restrictions, but threw out its core insights, as well—in fact, discarded its entire project: of understanding mind as instantiating a fundamental dialectic of meaning and mechanism. In particular, it threw out what is so distinctive about the acrobat: the centrality of constitutive, non-causal directedness towards a distal object. That is: it threw out what is most important about meaning, semantics, interpretation.

Was this rejection necessary? No. There is nothing about intentional directedness in general, or even about representation in particular, that runs counter to the situated sea-change. Representation—or perhaps intentional directedness, something like representation but far more general—not only can, but must be "rehabilitated" from within a situated perspective. Not only that, such rehabilitation is a necessary precursor to understanding either subjectivity or objectivity.

4 Rehabilitating Representation

Here—very quickly—is how it would go.

Physical regularities—causes and effects—are, as I have said, *local* in essentially all relevant respects, both spatial and temporal. That poses a problem for cognitive creatures. All you get, if you are physically embodied, in terms of *resources*, are two things: (i) the effective arrangements within you, plus (ii) what is pressing in on you, right now, at the surface. You live in a laminar cocoon, with physical coupling limited to the immediate here and now. Moreover, the world is *sloppy* (only weakly correlated), so you cannot 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—those local degrees of freedom—mean that you can rearrange your internal states with remarkable facility (if you are clever), without expending much energy. So what you do—what agents do, *what it is to think*—is to represent the world out there, beyond the periphery, by rearrang-

ing your internal configuration, and adopting appropriate habits and practices, so as to behave appropriately with respect to—so as to develop hypotheses concerning, so as to *stand in appropriate normative relation to*—that to which you are not, at the moment, physically coupled.

That is—to reduce it to one sentence—intentional systems:

1. Exploit what is *proximal* and *effective*, so as
2. To be intentionally oriented towards—i.e., to behave appropriately with respect to, to satisfy governing norms regarding—what is *distal* and *non-effective*.

I call this the **intentional mandate**. It is this mandate for which the acrobat is meant to stand as a metaphor.

As I say, the mandate contravenes none of situated cognition's tenets. Nor does it militate against even radically constructivist metaphysics. Remember the acrobat: what was delineated was established by the acrobat, not by any pre-existing boundaries or identity in the pointed-at spatial region.

5 Subjectivity

So that—glossing a thousand details—is the project: to (i) accept all the revisions of the situated sea change, while at the same time (ii) retaining the intentional mandate.

In what time remains I want to sketch the beginnings of a positive view, by seeing what it takes for subjectivity and objectivity to arise.

The route in is through consciousness. Consciousness has defied understanding for thousands of years because it seems so fundamentally different from things in the material world: sticks and stones, houses, galaxies and quarks. Consciousness has seemed especially inaccessible to scientific explanation, because of an apparently a vast, perhaps even unbridgeable, "explanatory gap"[6] between the two topics we are talking about tonight:

1. The private, **subjective**, inexorably first-person qualitative or phenomenological character of conscious experience;

[6]«Reference Levine»

and

2. The public, **objective**, detached, third-person character of empirical science (from physics and neuroscience to cognitive science to scientific psychology).

It is the first-person character of subjectivity that most people think is the root of the problem. Scientific accounts of third-person perspectives are expected to be (relatively) unproblematic. But when they consider the first-person case, people are driven to say such things as "We haven't the remotest idea of how consciousness could arise, physically."

I believe this intuition is almost exactly backwards. In my view, contrary to received wisdom, for all physically-embodied creatures, it is first-person and second-person perspectives (or at least inchoate versions of them) that are easy. What is hard—what takes skill—is developing a third-person, objective conception. That is what requires acrobatics; that is what leads to the extraordinary intentional dance, to which I alluded at the beginning.

To see why, think back to the example of turning around to retrieve something behind you. The terms used—'in front of', 'behind'—indicate positions from an egocentric or oriented point of view. That is, their reference is defined in terms of a coordinate system established by the concrete circumstances of their utterance. Many words have this property: *here/there*, *you/me*, *now/then*, *come/go*, *bring/take*, … etc. Tense, too, is similarly egocentric. Linguists call such constructions *indexical* or *deictic*, because their interpretation depends on their use.[7,8]

What matters, for us, is a special character of this use-dependence of indexical terms. In particular, what objects such words refer to are not only identified with respect to the use act itself; they are identified differentially, in terms of a "change" or "deviation" from the location and orientation of the utterance. Thus "bring" means to transport something *towards the location of the*

[7]Explain meaning/interpretation Δ.

[8]It is *active* use, not *static inscription*, that matters; Perry has an example of two deaf mutes, so poor they share a single card saying "I'm a deaf mute; can you spare any change?", which they alternately hand to passers-by. What 'I' refers to depends, dynamically, on *who hands out the card*.

speaker; 'you' refers to *the person to whom the speech act is directed*, etc.

The differentially-defined character of indexical terms is important to our story for three reasons:

1. Reaching upwards, it underwrites subjectivity;
2. Reaching downwards, it stems directly from physics.
3. It is what makes achieving objectivity hard.

Indexicality and subjectivity thus serve as something of "bridge" phenomena, connecting higher-level objective understanding with underlying physicality.

5.a Physical deixis

Look downwards, first, towards physics, and consider an example. What does a magnet say to an iron filing? "*You—come here, now! You—come here, now!* …"—on and on and on, forever. Every single linguistic element ingredient in this baldly anthropomorphic projection (all four words, the present tense, the imperative voice, etc.) is indexical. In general: if we give voice to what happens in physics, not from the outside, as in standard physical theories, but—like ventriloquists—*to the physical entities themselves*, so that the content of their communication mimics the operations of the governing physical laws, we are *forced to put indexicals into their mouths*.

This, I argue, is because *physics itself is indexical*. Not physical theory—physics-qua-epistemic-inquiry—but the actual concrete physical world: the force fields, the flux, the underlying plenum. Admittedly, no one talks this way (except me—yet!). But the indexical, or, as I prefer to call it, deictic structure of the physical world is tacitly recognized in our theoretical frameworks. Magnetic and gravitational attraction, bumping and shoving, flashlight beams, etc.—all physical regularities—are: *local* (in space-time), *incremental*, and *differential* (they govern how things change, not how things are).

This deictic, differential structure of the plenum matters *immensely*. It means that as physical signaling pathways well up from, or shade off into, direct physical engagement with the world, they must take egocentric form. Your stomach sends a signal to

your brain saying: "hungry!"—meaning something like "I am hungry, now!"; not "Brian is hungry on Monday night." As you leave the room, your cortex will instruct your motor routines: "turn right, now!", not "turn north at 8:56 p.m.!" Our entire physical existence, in the end, is grounded in such egocentric, deictic signals—symbols trembling on the verge of mere mechanism.

This deictic structure of physical coupling is not subjectivity—yet. But it will underlie it. For consider, taking this deictic physical coupling as the base case, what it is to start on the long and difficult road of beginning, not to stabilise one's relationship to the immediate environment, but to stabilise the things related to—the world out there, beyond the incoming signals, beyond the press of that 1/r2 enclosing causal envelope.

The issue has to do with the relational nature of the egocentric, deictic physical connections. There is no problem if you are being causally driven by something in the nearby environment: then, if your circumstances change, your state will change, too, correspondingly—i.e., you will be "updated by the world." Sunflowers, for example, can *track* the sun by being *driven* by the sun—using a simple servo mechanism. The difficulties arise not when you are engaged with what you care about, but when you are *disconnected*—when you leave. Then, when your circumstances change (position and orientation, say), your internal state, if you allow it to remain constant, will relate you to something new. A displaced magnet will attract new filings; a rigidly held flashlight (rigid with respect to the body of the acrobat) will wildly light up new spots. In order to get a fix on the world, what you want to stabilise is not the *relationship* you bear to it, but *what you are related to*. Consider a fixed entity in the environment, say, like home—if you want to hold that stable, through changes in your circumstances, then you must adjust your egocentric relation to it, to *compensate for what has happened to you*. If you rigidly held onto the idea that it is "four blocks left and nine blocks down," you will find yourself forever referring to a new spot—not what you intend. Think about the acrobat: as she moved, she had to rotate her wrist, so that the light beam could point in a new direction, and thereby arrive at the same point in space—instead of maintaining the "same" egocentric direction, which would have led it to point to a new

point in space.

A less contrived example. Consider: the vestibular-ocular reflex—the fact, wired deep in your brain, that if you rotate your head, with your eyes open, you will invariably find yourself rotating your eyes, in their sockets, by an "equivalent" amount, in the opposite direction, so as to maintain a stable point of visual focus—on a distal chair, tree, whatever. With effort, it is possible to override this, and rotate your head *without* doing the compensating eyeball rotation (i.e., holding your eyeballs fixed in their sockets); in which case—*note this!*—you no longer see the world; everything turns into a blur. But if you allow your eyes to move, just the right amount, the blur suddenly vanishes, and—this is incredibly important—*the world snaps into focus* (not the incident optic array, but the distal world itself).

Objectivity is going to be something like that: *having* (allowing) *the world to snap into focus.*

In some ways, the vestibular ocular reflex is misleading, as an example of intentional acrobatics, because vision can be causally-driven, at least in part, by what you are looking at (like the sunflower, vision can "servo in" on a scene). What is distinctive about long-distance intentional or semantic directedness, in contrast, of the sort objectivity requires, is that it *cannot be "driven" by what it is directed towards*. In order to maintain objectivity—in order for their to be a world, for them—agents (and communities) have to shoulder responsibility for keeping themselves appropriately "pointed."

In spite of these differences, however, two facts about the vestibular ocular reflex are revealing:

1. Like the acrobat, it involves a process that I call **deconvolving the deixis**—compensating for the contribution that changes in your circumstance make to the deictic physical relations in which you stand to the environment. You *wash out your own contribution*—and thereby let the world be stilled.

2. When you "lock onto" a distal object, what is stabilised (via these processes of deictic compensation) is seen or looked out onto *from a point of view*. The very conditions that allow the world to snap into focus entail, as a consequence,

> that the world is "looked out onto," by someone, from some place. Not only that; this "someone" is not anonymous. The "looker" is the very same "you" whose contribution you just had to erase, had to wash away, in order to allow the scene to emerge—like a developing Polaroid picture—before your very eyes.

It is a consequence of physics, in sum, that objectivity is only subjectively achievable.

6 Discussion

Needless to say, there is an extraordinarily long story to tell—a science's worth—about how these deconvolution processes, starting from deictic physical engagement, but gradually "letting go" of what is immediately connected, in order to stabilize what is progressively far away, can ultimately lead all the way to objectivity. In what time remains I will enumerate three final points, to convey a flavour of how that story goes.

6a First, second, & third person

The first set of comments has to do with relations between first, second, and third person perspectives.

You cannot update your *entire* physical state, every instant—it would take too much work. So not all representations can be deictic and/or egocentric. On the other hand, you always have to do some updating, because you need egocentric representations in order to mesh with your basic physical (mechanical) capacities. So in general—if you are going to develop a whole conception of a world—you need an efficient, balanced set of representational strategies, that trades off between local detail, necessary for action, and long-distance stability, necessary in order to "still" the world.

This is where full-scale intentional and representational practices comes in. All sorts of strategies are employed: relatively less perspectival or indexical signs, established landmarks, relying on others, the whole framework of "representational scaffolding" that Clark and others talk about in cognitive science.[9] Think about

[9]Regularities that hold over a full *range* of experience, do not need to be "updated" if one moves around within the range. You need not think about time zone if you stay on the East Coast; or that "to the right of" is a three-place relation, not two, unless you spend a lot of time standing on your

how you get to a national park. Relatively non-perspectival maps get you to the right region; from there, much more indexical signs— "straight ahead," "two miles (from here)," right or left— take you from there to the gate; from there, direct (Gibsonian) engagement with the physical environment guides you the last few feet. Even the road itself—including the markers at the edge of the lane—can be viewed as a final, deictic sign that allow local, connected guidance.

Third-person perspectives do not emerge as so much "from nowhere," on this picture, as "from anywhere"—at least "from anywhere within a broad range" (nothing is entirely deixis-free). Moreover, they are invariably grounded, connected to the world, by being seamlessly connected to progressively more indexical or first-person signs, shading imperceptibly into direct (deictic) physical engagement.

In sum, long-distance conception requires third-person stances, in order to manage the complexity (through abstraction and a degree of perspective independence). Being part of community requires second-person representation, so as to be able to communicate with other deictically-embedded agents. Local navigation requires first-person, deictic directedness, for coupling to the physical plenum.

Two crucial points, about this range of perspectives:

1. Developing an *objective* conception of the world—one that includes the knowing subject—requires not just one of these perspectives, but the full integrated set: skills to move back and forth seamlessly, between and among them (as you did, when turning around to pick up the glasses).
2. All intentionality—directedness—must be anchored to the world via grounded, first-person, subjective skills—skills that ultimately mesh, without a trace, into direct physicality.

6b Ontology

The second set of comments has to do with ontology—with *world-making*.

---

head.

Physical fields are stupefyingly complex: a maelstrom of superimposing and crashing waves and vortices and turbulence—a little like falling overboard at sea, and finding yourself drowning in a buffeting array of turbulence and spray—except without the "you".

The world is so overwhelming in detail, in fact, at the physical level, that if we tried to react to it as genuinely physical (i.e., as field-theoretic), we would be completely swamped. No finite creature could *begin* to deal with all this detail; to say nothing of the fact that huge amounts of it—almost all—is of no especial interest, as regards ascertaining what lies beyond.

So what does the creature do? It *abstracts*—or, as I like to say, "**registers**"—the world, simultaneously (i) failing to attend to a lot of what is there, by ignoring the vast majority of the detail or "information" with which it is presented, and (ii) attending to what is *not* really there, by imposing conceptual structure on what remains, so as to render it modestly intelligible.

A particularly important form of abstraction is what we might call "conceptual":[10] the staggeringly reductive simplification of the world into the familiar cast of characters—objects, properties, and relations. This is the most "third-person," disconnected, non-engaged form of registration. In keeping with its being necessarily disengaged, it is highly non-detailed. It is useful for long-distance coordination. As one moves in to physically engage with the world, the need for—and adequacy of—conceptual abstraction (i.e., the need to parse the world into discrete, reidentifiable individuals) falls away, to be replaced by fine-sensory and motor coupling with inexpressible detail.

Objects, properties, and relations, in other words—products of conceptual registration—are, as I once put it: "the long-distance trucks and interstate highways of normative, intentional life. They are undeniably essential to the overall integration of life's practices—critical, given finite resources, for us to integrate the vast and open-ended terrain of experience into a single, cohesive, objective world. But the cost of packaging up objects for portability and long-distance travel is that they are thereby insulated from the extraordinarily fine-grained richness of particular, *indigenous*

[10]Maybe talk about Evans/McDowell Δ.

life—insulated from the ineffable richness of the very lives they sustain."

There is no reason to suppose—indeed, every reason not to suppose—that, in these conceptualist projects that parse the world into discrete objects, that how the world is registered is "independent of the subject." That does not make the view entirely *relativistic*—i.e., does not imply that ontology (the stabilised world) is wholly dependent on the subject (or the subject's community). Neither limit case is tenable—one hundred percent world, independent of subject; or one hundred percent subject, independent of world. Both are ideological manifestations of exactly the sort of absolutism for which we banished logic. Like gardening, real-world ontology is a *collaboration*—between subject, environment, community, and thereby-stabilised world.

Four notes:

1. This constructivist sentiment is not an extra-theoretical metaphysical assumption; it is demanded by taking field-theory seriously, as an account of the physical basis of existence, and computational complexity arguments seriously, about what finite creatures can do with limited resource. That is: an intermediate level of social constructivism is *naturalistically forced*.
2. This forced view—to which a serious commitment to physics inevitably leads—is one of *ontological pluralism* on top of *metaphysical monism*. I am prepared to argue that the resulting picture does simultaneous justice to what matters, to realists, about realism—and also to what matters, to constructivists, about constructivism.
3. I have said that ontology is not independent of subjects, because it is based on abstraction. The metaphysical world is not independent of subjects, either, but for a different reason: subjects are part of the world—and parts are not independent of the wholes they partially constitute.
4. A note for philosophers: part of what is being said, here is, that ontology needs naturalisation as much as semantics—since (I claim) there are no reidentifiable objects in

science. And not just ontology: but abstraction, as well—its epistemic warrant.

6c Normativity

Third and finally, a remark on normativity.

I have not said much—have not said anything, really—about what establishes the non-effective links that relate subjects (i.e., us) to the entities towards which we are intentionally directed. This directedness, as I have suggested all along, include "truth" and "reference", but that is not all; intentional directedness also involves caring, respect, love and hate, curiosity—and awe.

The bottom line is that the links are normative—as long as we are generous about the meaning of that word. If we had time, I would draw distinctions between (what I call) "statical" norms—norms that govern states, in the way that truth and reference and the values of science and the like have traditionally been conceived—and "dynamical" norms: norms that govern processes. And I would go on to say that dynamical norms, including full ethics, how to live, living in the truth, etc., must be in the driver's seat.

More particularly—to draw out the most important moral—it is the dynamical norms governing people's ongoing projects that warrant the abstractions in terms of which they ontologise the world. Whether there is a single mountain over there, or three mountains—which of the claims "There is one mountain" or "There are three mountains" is *true—depends on what you are doing*, depends on your normative stance towards the normative status of the object of your claim. More generally, there are no norm-free empirical claims. On this the feminists, among myriad others, are right.

Moreover—there is something wonderfully ironic about this—this normative dependence of ontology on life practices is *especially true of conceptual claims*, involving traditional ontology. Direct physical engagement—the local, microdetailed physical engagement with the world, being coupled to the immediate surround, is not nearly so subject-relative. Rather, it is third-person, framed in terms of high theoretical abstractions—including the claims of science—that are the most norm dependent.

The point can be framed terminologically. At least since Des-

cartes, the word 'matter' has been split in two: between a noun, meaning physical stuff, and a verb, meaning something like "is important." On the view I am defending, the two sides rejoin. A "material object," I claim, is *a chunk of reality that matters*. ("To whom?" you ask. That is right: that is the right question.)

Put it this way: my aim is to heal the temporary 300-year schism between matter and mattering.

7 Summary

So that is the picture, in a few thousand words. Radically embodied agents who, in virtue of (i) their singular and collective particularity, (ii) the unutterable contingencies of their existence, and (iii) the enabling constraints that derive from being embodied in physical and social fields, are able to *parlay their material freedom into a commitment to the world*—a world that contains them but transcends their grasp; a commitment that allows them to take it as world, populated with what exists.

Because that, in the end, is what objectivity is: a commitment, on the part of subjects, to take the world to be world, host of objects (and everything else); and objects, to be entities hosted by the world.

Subjectivity and objectivity, in other words, far from being independent or at odds, are inextricably intertwined, for at least three reasons.

1. Objectivity requires subjectivity, because of those underlying physical constraints; there is no other way for reference to be authentic—no other way for a theory to refer to anything at all—except when it is grounded in an agent's subjective experience, and deictic physical engagement.
2. Truly objective knowledge requires recognising one's own subjectivity, in order to take the world to be the world. You cannot be objective, that is, unless you recognise your embodied participation in, responsibility for, and effect upon, and subjective contingency of the conceptual scheme in terms of which you understand, the world about which you care, to which you extend commitment.[11]

[11]Fragment: It also requires admitting that one's attempts to describe or

3. Finally, being conscious, being awake, being aware, in its ultimate sense, in turns requires objectivity. In this the meditative traditions are perhaps right. True consciousness requires commitment to, but openness towards, the self-transcending world.

8 Epilogue

Two projects, I take it, motivate a university:

1. Understanding who we are—our character, our practices, what we are made of, what tests our mettle; and
2. Understanding, without limit, the world that hosts us.

Subjectivity and objectivity, in others words—as long as we are generous about the meanings of the words.

Except now we have expanded beyond the confines of ai, cognitive science, and philosophy, to incorporate history, literature, the humanities and social sciences—including social studies of scientific inquiry as a particular (often privileged) form of human practice. This time, that is, we are concerned, not about what it is to build a creature, what must be true in order for us to extend a sense of "we," but about the human condition, and the world we find ourselves in—how to understand the "we" *that we already are*. This time it is not a case of constructing scientific accounts, from the outside, of subjectivity and objectivity, but of developing stories, from the inside, of what our subjectivity and objectivity are like.

Except that to divide things up this way is untenable, for two reasons—*both of which we reached as results, above*.

1. First-person ("inside") and third-person ("outside") understandings are not separable like that. No third-person account, for starters, *has ever referred to anything at all*, except as grounded in concrete, particular first-person subjectivity. On the classical image, scientific understanding is

---

conceptualise it invariably make reference to one's inescapably embedding normative projects—from which it follows that there is no single story, no master narrative. *That* one is an embodied subject, however, and therefore that there is no single story, no master narrative, neither voids talk of reference, nor commitment of object.

claimed to be third-person; but we have just argued that such a view must be wrong. It is incompatible with physics. Like all understanding, scientific understanding must rest on an integrated suite of first, second, and third-person skills.

Similarly on the other side. No subjective social science or humanities understanding of our place in the world can have the requisite detachment and dispassion to warrant the university's imprimatur, unless it is also objective—unless it assesses our situation fairly, from the outside. Committed detachment—passionate dispassion—require subjective-cum-objective integration, not fracture.

So dividing the university by who studies from the inside, and who studies from the outside, is a mistake.

2. Second, because abstraction is normatively grounded, empirical studies cannot be separated from normative studies. The "takes" on the world that science uses are grounded in the normative projects of scientific inquiry (and perhaps the societies that do it).

All of which has consequences for how we should proceed.

Look at it this way. Science, in the form of logic, ai, and cognitive science—to say nothing, increasingly, of neuroscience and biology—is *theorizing theorizing*—constructing knowledge of what knowledge is, developing accounts of how accounts are developed. Gradually but inexorably, science is growing reflexive—bringing the "doing" of science within science. Though still inchoate, this process of self-incorporation is bound to accelerate. Scientific accounts of consciousness and first-person subjectivity, just now starting to appear, will soon be *de rigueur*.

Increasingly, that is, science will domesticate questions that have classically been viewed as extra-scientific: not just subjectivity and consciousness, which we are already seeing, but also: the relation between models and reality; the character of scientific evidence; the role of the observer; which of formalism, Platonism, relativism, constructivism, or any of a host of other realist and irrealist metaphysical positions is right; the normative status of world-making; what it is to care. Questions like this, which used to be relegated to bars and late night conversations, to conviction

and personal faith, are moving into the heart of this reconstituted science. (They are even starting to be talked about in the middle of the day!)

Problem is, though I find this all intellectually thrilling, I am not politically sanguine about these developments. I do not think we can ask science, present-day science anyway, to shoulder this whole burden. I am still troubled by that erasure of logic's central insight, in computer and cognitive science—by the weight of the methodological pressure to reduce everything to mechanism. Because as we said at the outset, to take on these ultimate questions is to address issues of meaning as well as mechanism—the original dialectic. Symbols, interpretation, sociality, normativity—to say nothing of even deeper issues, such as wonder, creativity, and awe. To do this right, we need collaboration from disciplines that study those things: not just the arts, as well as the sciences, but politics, law, religion. The full gamut. For it is with their insights, about contested stories, pluralism, etc.—as I hope even today's bit of a sketch suggests—that these considerations from science (even: robotics) most naturally mesh.

So maybe this reflexively reconfigured inquiry should not be called science, but science$^{\text{prime}}$. Or perhaps (since we are studying meaning): *reënchantment*. Maybe even: colleges and universities themselves. "*Eruditio* et *religio*", one might even say—if that phrase were not already taken.[12]

How this all works out, we will have to figure out together. All I know for sure is that the intellectual consequences of (i) the rapidly converging trio of computing, biotechnology, and something we have not much mentioned, nanotechnology, in conjunction with (ii) the theoretical disciplines that undergird, nourish, and are in turn buffeted by them…the intellectual consequences of those two things, once we realise they have our being in their sights, are immeasurable. This is why I said, when interviewing for this job, that what matters about STS—about the social impact of science and technology—is not just the spread of the technology itself, its impact on our practices and routines. What matters even more are the ideas on which these technologies are founded: ideas go straight to the heart of our conception of what it is to be us.

[12] '*Eruditio et religio*'—learning and religion—is Duke University's motto.