

Brian Cantwell Smith

On the Origin of Objects



6 • Flex & Slop

Enough ground has been cleared that we are almost ready to begin construction.

The strategy in this chapter is temporarily to accept the contributions of physics as ingredients in the overall picture, and to find out what else is needed for a complete picture of intentionality, semantics, and ontology. Given the results of chapter 5, and the commitment to honor the Criterion of Ultimate Concreteness, that translates into the following more specific goal: to understand how a conception of objects can arise on a substrate of infinitely extensive fields of particularity.

Except of course that this is an untenable way to phrase it. To say “a *conception* of objects” makes it sound as if the achievement is the subject’s, by assuming a split between conception and what is conceived of. It also fails by making it sound as if the achievement is cognitive. Nor is anything gained by striking a more traditionally realist stance, and asking “how objects can arise on a substrate of infinitely extensive fields.” That puts the achievement too squarely on the object. Both ways of putting it violate the mandate of avoiding an a priori subject-world split.

In place of these dichotomous formulations, therefore, I will speak, unitarily, of *registering* the world.

1 Registration

By ‘register’ I mean something like *parse, make sense of as, find there to be, structure, take as being a certain way*—even *carve the world into*, to use a familiar if outmoded phrase. Thus, as I write this paragraph and look out the window, I register a bunch of

trees, a stone barn, and a dilapidated child's bicycle thrown down on a moss-covered granite slope.

Although I will differentiate the two notions in a moment (registration is much broader), there are important ways in which registration is similar to perception. In ordinary parlance, and also as I will understand it, perception is an activity on the part of an intentional subject, and furthermore an activity that relates it to, or engages it with, something in the world. Typically this will be something in the *external* world, something outside the subject, something that may even be quite far away. Thus we talk of perceiving a sparrow, a wry look, a distant thundercloud—and not, except on higher-order reflection, of perceiving our perceptions or sensations. Furthermore, perception not only relates the subject to that distal object, but also implies a kind of intentional success. Perception, that is to say, is an *achievement*. To say that a sailor perceives an island on the horizon implies, unless otherwise noted, that there was an island to be perceived, and that the sailor did indeed see it. Moreover, as I will put it, it is the island itself, not an impression of the island, that figures in the *content* of the sailor's intentional state. As one would say in the philosophy of psychology, perception, at least in the unmarked case, is both veridical and “broad.”¹

Depending on how psychology goes, a perceptual claim may indirectly do other work as well: such as affirm the existence of a complex causal relation between the sailor and island, perhaps mediated by light waves and binoculars; or validate the crucial role played by sense data or sense impressions; or even just imply that a certain structure of stimuli or causal disturbances has impinged on the sailor's surface. But none of these things is what perception actually *is*. People do not perceive light waves, or sense data, or impressions, or causal disturbances, at least not ordinarily. It is always possible to define some other word to mean any one of those things—e.g., to signify the relation between an agent and the structure of the impinging stimulus or sensory

field—some word such as ‘narrow perception,’ or even, but unadvisedly, ‘Perception.’ Even if they could be proved coextensive with perception, or nomologically necessary for perception, or to be superior to perception for purposes of a scientifically respectable psychology, none of those facts would make them *be* perception, in part because they would relate the subject to the wrong thing. To say that a sailor perceives the island is to talk about two entities, several miles apart, and to say that, perhaps in virtue of some number of causally mediating processes or other psychological mechanisms, the sailor has successfully established a specific directed semantic relation towards that distal island, an island with which, as I will say, they are *intentionally engaged*.

Both of these facts—that perception reaches out into the world, and that, at least by default, it is veridical—are reflected in grammar. The direct object of the verb ‘perceive’ is (a term for²) the distal entity or situation towards which the subject ends up being intentionally or semantically directed. It is striking, however, that in the case of higher-level intentional activities—thought, conception, cognition, wondering—we have no such analogous verb. Terms like ‘think,’ ‘conceive,’ etc., usu-

²There was a time, I take it, when the *subject* of the sentence ‘Pat ate dinner’ was a person named Pat, and the *object* of the sentence was an edible meal. Nowadays, it is almost universal to strike a syntactic stance, and to say that the subject of the sentence is the word ‘Pat,’ and the object, the word ‘dinner.’ A parallel syntactification has taken place in mathematics. When we were in high school, the velocity of the river was a constant, the angle that the boat left the dock was an independent variable, and the position where it arrived on the far shore was a dependent variable. By the time we got to college, velocities, angles, and positions were no longer variables; variables had become syntactic entities, entities that *designate* velocities, angles, and positions.

I am not fully in favor this development, which has pulled the terms ‘subject’ and ‘object’ away from their original sense, the sense in which they are being examined in this book. Hence my use of parentheses around ‘a term for.’

¹Fodor (1980), McGinn (1982), Dennett (1982), Pettit & McDowell (1986).

ally take a thought or its content as their object (you think a *thought*, not a table). When they are used with an external situation as object, moreover, as for example in the sentence “He conceived a grand party,” there is almost the opposite implication, that the conceived entity does *not* in fact exist (unless, interestingly, that object is itself a mental—intentional—construct, such as when you conceive a scheme or plan). Similarly, the more common construct with a prepositional object, such as “a conception *of* a nuclear-powered aircraft,” also fails to suggest, and may even contraindicate, the existence of whatever it was that was conceived of.

My intent, in contrast, is to have ‘register’ be much broader than ‘perceive,’ and certainly to include full conception, but to retain both of these real-world characteristics of ‘perceive’: of similarly reaching through to the world or the referent, and carrying as well a default implication of an appropriate notion of veridicality or success or other intentional virtue or worth. Unless marked, that is, the phrase ‘she registered a table’ will be taken to imply not only that there was a table to be registered, but that she successfully ended up in an intentional relationship to it. Or to put it more neutrally but more adequately, the phrase will be read as implying that the overall situation of subject-cum-world was such as to sustain the truth or success or worth or virtue—i.e., the *objectivity*—of the intentional act of her taking there to be a table.

Exactly how broad registration is must await a more developed account, but my general intent is to have it include all intentional actions that can legitimately be said to have some form of *content*. How broad that is is partly an empirical question. It will at least include cognition, sensing, having intuitions about, and a number of other things—and will perhaps include willed action, action that assumes that the world is given or structured in some way. To register the world, that is, is to do or be oriented

towards the world *in such a way that it presents or arranges or constitutes itself as world*.³ Note too that although there is something right about speaking of individual subjects as the entities or agents that register, this is not to deny that in all likelihood it will be whole cultures, language communities, communities of practice,⁴ or collectivities of people-and-instruments-and-organizations-and-documents-and-tools-and-other-essential-but-expensive-entities that are the full sustaining locus of this intentional achievement.

The subject matter of registration is more general than that with which the investigation started, which focused more narrowly on the notion of an object. In the special case of objects, ‘register’ can be taken as being related in meaning to ‘individuate’ or ‘identify,’ but a much closer synonym is ‘objectify,’ which shares with ‘register’ the useful property of not committing itself to the locus of responsibility of the resulting individuation criteria. Because objects remain in focus in the rest of the book, I will continue to consider the registration of objects as a central case, but it is important to keep in mind that registering objects is just one way of taking the world to be present or given, even if current theoretical culture takes it to be paradigmatic.⁵ As a reminder once again of how narrow a case it is, however, imagine standing with a friend at the prow of a boat flying across the sea on a summer’s afternoon, drinking in the air and sky and surf and islands, and reveling in the freedom and luxury of the full month’s vacation opening up in front of you. It would be a trav-

³In spite of this sentence, I do not in general wish to speak of registering the *world*. The basic, simple, directed case should always govern. Forget the world; register potatoes.

⁴See Lave & Wenger (1991).

⁵Current *theoretical* culture takes Objects to be paradigmatic, not objects. On the other hand, because of the fact that so much of human practice is in part ideologically self-constituted—i.e., is in part constituted by reflective registrations on what it is to be that kind of thing—the distinction in real life between objects and Objects will not be all that sharp. Only Distinctions have sharp boundaries like that, not distinctions.

esty to reduce the ensuing registration to a process of chopping up the experience into a series of discrete individuals. So I will take registration to include quality, mass substances—even feature-placing, so long as there is reason to believe that the features are placed in a world. Remember when the Count of Monte Cristo, sewn into his friend's death sack, was unceremoniously dropped outside the Chateau d'If. In that first moment of shocked recognition, he registered, but did not individuate, the sea.⁶

The term 'register' has a number of other useful connotations stemming from its use as a term of ordinary English. Informally, for example, we say both that she did not register his anger, and also that his anger failed to register on her—implying a salutary sharing of responsibility for a subject's taking an object to be an object. Responsibility for intentional success falls on both *registrar* and *registered*, in other words, even if not symmetrically.⁷ In addition, to register something implies registering it *as* something, and thus the term, like representation, carries the requisite aforementioned sense of "aspectual shape"⁸ or *intentionality* (with an 's') characteristic of all intentional (with a 't') activity. Third, 'register' also carries a useful implication of alignment with the external situation, as for example in the printer's notion of registration marks (used to allow one document to be fitted exactly into place on top of another). Even without making any special commitments about lining up how the world is and how we take it to be, it is still useful to retain the residual sense of realism. Fourth, 'register' also carries a sense of involvement or engagement with the world: to register a print requires taking physical action on it. Sometimes, moreover, it

suggests not only a process of alignment, but one of active calibration, even with normative overtones.

Finally, and most importantly, the word provides a convenient way to honor the fact that the ways in which we take the world to be are in part a function of the assumptions and culture and machinery, including conceptual machinery, with which we approach it. For many reasons, some already mentioned, and some that will become more evident in succeeding chapters, I do not want to subscribe to the never-very-well-explicated notion of a conceptual scheme (primarily because I do not want to presume that all registration is *conceptual*).⁹ Instead, therefore, I will adopt the nominal cognate of the verb, and simply say that different cultures, and different people within a culture, and one and the same person at different times, and indeed just about all distinct utterances, speak or talk or deal with the world *in different registers*. The idea is to suggest something like seeing the world in a different key, or of speaking in inexorably subscripted vocabularies. For example, part of the problem with specifying the content of a frog's representation of a fly, alluded to in chapter 1, is that most of us theorists, as I would put it, have trouble (to put it mildly) "speaking in a frog's register."

Some of these last connotations will surface in later discussion. For the moment, it is enough to summarize three essential properties of this notion:

1. Registration is the net activity that leads to (what we theorists register as) a conception of, or take on, or intentional attitude towards, the world as given or available—anyway, as *world*.
2. Registration is originally neutral as to the appropriate locus, if any, of two essential subject/world splits: (i) that between registrar (subject) and what is registered (object), and (ii) that between subject and supporting community (people, instruments, practices, documents, culture, etc.).

⁶Dumas (1894). There is of course a sense—to pedantically admit the standard stalking horse—in which he did individuate the feature *sea*.

⁷Subject and object (registrar and registered) are symmetrically *important*; that does not imply that their roles are in any way the same. Perception and registration are manifestly asymmetrical (directed) relations.

⁸See Searle (1992).

⁹See Davidson (1974) for an analysis of its not being well explicated.

3. Registration does not single out objects as a premier ontological category or class—or even, necessarily, require that objects count as a distinct ontological species.

From this characterization, it looks, at least on the surface, as if of the three properties initially identified as characteristic of objects—objectivity, particularity, and individuality—only the first (objectivity) is a necessary condition of registration in general. But since looks can be famously deceiving, it should be considered an open question as to how much, and in what form, the other two properties, of particularity and individuality, are essential to registration. By way of an initial hint, note the following striking duality: what in chapter 4 I called *feature-placing* would appear, at least at first blush, to involve: particularity but not individuality (+PARTICULARITY, –INDIVIDUALITY) in, as it was said, the metaphysical arrangements that made the feature-placing sentence true—e.g., in the actual rain; but individuality and not particularity (–PARTICULARITY, +INDIVIDUALITY), i.e., exactly the converse, in the abstract feature thereby placed.

2 Flex and slop

Here is where we stand. I started out the project, in chapter 4, by splitting the notion of an object into two orthogonal concepts: of individuality and particularity. I then argued that physics supplied a picture of world-extensive cosmic particularity, but was of no help with respect to individuality. Given that result, I temporarily took up the (physicalist-seeming) project of seeing how a notion of individuality could arise on a substrate of fields of particularity. For methodological reasons I forswore a sharp subject/object distinction along the way, but since that meshes with a sympathy, betrayed much earlier, for the constructive intuition that objects emerge collaboratively from subject-object interaction, that will turn out to be not as much of a felt restriction as it might otherwise have been.

Now the topic has been widened, from individuals to regis-

tration more generally. So the question on the table can be broken into two parts: (i) how does *registration* arise on a substrate of infinitely extensive fields of pure particularity; and (ii) how, and in what ways, is the registration of objects a distinct species?

And with that, finally, we are ready to start assembling a positive picture.

How to start? With this observation: that it is essential to the picture developed so far, and also an anchor of common sense, that the multi-various parts of the world do not march in lock-step together. The world is fundamentally characterized by an underlying flex or slop—a kind of slack or “play” that allows some bits to move about or adjust without much influencing, and without being much influenced by, other bits. Thus we can play jazz in Helsinki, as loud as we please, without troubling the Trappists in Montana. Moths can fly into the night with only a minimal expenditure of energy, because they have to rearrange only a tiny fraction of the world’s mass. An idea can erupt in Los Angeles, turn into a project, capture the fancy of hundreds of people, and later subside, never to be heard of again, all without having any impact whatsoever on the goings-on in New York.

The world’s flex and slop is so obvious that it is a little hard to talk about. As a contrast, therefore, imagine a world quite unlike ours, consisting, as suggested in figure 6·1, of nothing but an endless series of interlocked gears. Suppose, to make this precise, that every gear is constructed so as to mesh with one or more immediate neighbors, and that the entire gear universe is interconnected, but in such a way that it is still possible for them all to be turned—i.e., so that it does not lock up. Suppose, too, that the gears are perfect: no friction, no play between the teeth, and shaped so that rotating one at an even speed causes the others to rotate evenly as well, though at a potentially different speed.¹⁰

¹⁰In our world, so-called “constant-velocity” gears must slip against each other, and therefore have more friction than non-slip ones (requiring better oil).

One could even arrange to have an indefinitely large number of gears with relatively-prime numbers of teeth, so as to obtain a geometrically unbounded configuration space.

The gear world would lack slop. Effects would not dissipate. If one gear were to move by even a tiny amount, every other gear in the universe, no matter how far flung, would instantly and proportionally be affected.¹¹ Admittedly, amplified (but




FIGURE 6.1 THE GEAR WORLD

still non-representational) long-distance effects are possible in our world, too. That is why, famously, it is just remotely possible for a butterfly in Tokyo, with a single flap of its wings, to wreak havoc on New York's weather. Yet even if this were to happen, in our world such effects take time. And anyway, that is not how things normally

work, and certainly not in this regimented and predictable way. Even if we were to discover that every macroscopically observable regularity was the product of such amplified long-distance effects of microdisturbances, it would still be true that far and away the majority of microdisturbances quickly die away. In our world, especially to the extent that we find it coherent, effects by and large dissipate. Think $1/r^2$.

Overall, my aim in this book is to show that the world's primordial flex or play does two crucial things: (i) establishes the problem that intentionality solves; and (ii) provides the where-withal for its solution. 'Problem' and 'solve' are prohibitively expensive words, of course; no reader should take them seriously.

¹¹I have phrased this as if there would be time and distance in the gear world, but that is almost surely false. We claim that there is no action at a distance, but that gets it backwards. *Distance is what there is no action at.*

Nor is this the only poetic license I will take. In order to convey an intuitive picture of the relationship between slop and registration, I want to set the irreductionist principle aside, temporarily, and speak in terms of fully constituted, and thus highly expensive, objects, properties, and subjects, too—in spite of having just gone to such work to forswear them. Think of this as a temporary puppet show: as a band of marionettes being held up, not quite invisibly, with strings reaching up into the hidden recesses of the theorist's registration scheme. Though in one sense pure artifice, what is played out on stage can still unleash true imaginings about the genuine article. And anyway the show will not last long. As soon as a sense of the choreography has been conveyed, it will be important to start to untie the strings, and watch as the cast of characters learns to stand up and dance on their own.

Go back, then, to the flex or slop. Here is the problem. Because of the dissipative nature of the playing field, an enduring entity cannot, at any given moment, be affected by things that, at that same moment, are beyond what I will call *effective reach*. Effective reach is not a yes/no affair; rather, this is essentially the gradual falling off or dissipation of influence familiar from physics. Thus imagine standing on the Marin headlands looking at San Francisco, and slowly losing visual contact with the city as it is slowly engulfed in fog. The buildings only gradually fade from view, rather than vanishing instantaneously. Still, there comes a point where the effective illuminative coupling subsides below the level at which it can any longer do any work. Crucially, though, it does not follow that what is no longer within visual or tactile or other kind of effective reach is thereby rendered irrelevant. San Francisco still matters because of past connections, present interests, future possibilities. Effective connection will soon be re-established, too, when one rushes back to meet someone at Zuni's.

Moreover—and this is the essential point—when some-

thing is blocked from effective view, or removed from effective grasp, it is not thereby eliminated from one's thoughts. Intentional directedness is not held in place with physical glue. Long after the physical tie has been broken, or has stretched too thin to hold anything together, or has aged so much that it has become brittle and cracked and yellow, the semantic relation tying a subject to their subject matter will persevere, far less battered by the buffeting of circumstances, far less deteriorated by the ravages of time.

The overall pattern of separation and subsequent re-engagement, with some form of non-effective coordination maintained through the interim, is an extraordinarily important aspect of the sustenance of registration. It happens at a relatively primitive level, as well. Consider a rabbit keeping a watchful eye on a prowling coyote. While the coyote remains in view, the rabbit's tracking system can rely on causal or effective coupling to keep track of it—by implementing what amounts to a simple servo mechanism reaching all the way from the coyote itself through the atmosphere into the rabbit's retina, and from there to the muscles controlling its eyes, head, and even legs.¹² But now suppose the coyote goes behind a rock. Effective coupling is broken, but that is not the end of the story. It is still "useful" (another expensive word) for the rabbit to keep the coyote's presence and approximate position in mind—for example, by continuing to look in the coyote's general direction, so as quickly to be able to re-establish visual contact should the coyote suddenly reappear. It is a simple strategy, but a very general one, and often very useful: *to stay oriented towards what is no longer effectively discriminable.*

More fancifully, imagine that a species of "super-sunflower" develops in California to grow in the presence of large redwoods. Suppose that ordinary sunflowers move heliotropically, as the

myth would have it, but that they stop or even droop when the sun goes behind a tree. Once the sun re-emerges, they can once again be effectively driven by the direction of the incident rays, lifting up their faces, and reorienting to the new position. But this takes time. Super-sunflowers perform the following trick: even when the sun disappears, they continue to rotate at approximately the requisite $\frac{1}{4}^\circ$ per minute, so that the super-sunflowers are more nearly oriented towards the light when the sun reappears. Since in all likelihood the orientation will be imperfect, super-sunflowers, like the ordinary variety, will rely on servoing to re-establish the correct causal locking, as a way of fine-tuning their positions. But because they have a leg up on the competition, in those crucial few moments when the sun is first visible again, they get just that much more benefit from the incident rays. And flourish.

What distinguishes super-sunflowers is not the fact that they track the sun. Ordinary sunflowers do that. What is special is the fact that they track something *to which they are not effectively coupled.* This behavior, which I will call "non-effective tracking," is no less than the forerunner of semantics: a very simple form of effect-transcending coordination in some way essential to the overall existence or well-being of the constituted system.

One does not need biological examples, real or fanciful, to see coordination across separation in effective coupling. It is endemic to computation in even the simplest cases. Thus consider the use of error-correction codes in computer memory systems to guard against random errors due to radiation and other forms of material decay.¹³ It is a straightforward exercise to design a circuit that takes as input 32 bits of data, and outputs a 6-bit code

¹²Note that this closed servo mechanism is easiest to recognize as a *loop* if one does not register the boundary or individuality of either the rabbit or the coyote.

¹³'Bit rot,' a term used to explain why programs that worked at one time no longer work, refers not to this sort of material decay, but to the inexorable effect of ongoing changes and development in the operating system and environment and other aspects of the context in which a program runs. See the discussion of boundaries and identity conditions in §§ 2 & 3 of chapter 11; also Smith (1994).

indicating how many of the original 32 bits are “on.”¹⁴ Suppose this circuit is used as indicated in figure 6-2, so that whenever a 32-bit word is sent to the memory subsystem, 38 bits are actually stored. Then, when a request to retrieve the word is received by the memory system, all 38 bits are read out from the actual store, and a symmetrical checking circuit ensures that the proper number of the original 32 bits remain on. If everything matches, the word is assumed to have survived correctly; if not, an error is signaled.¹⁵

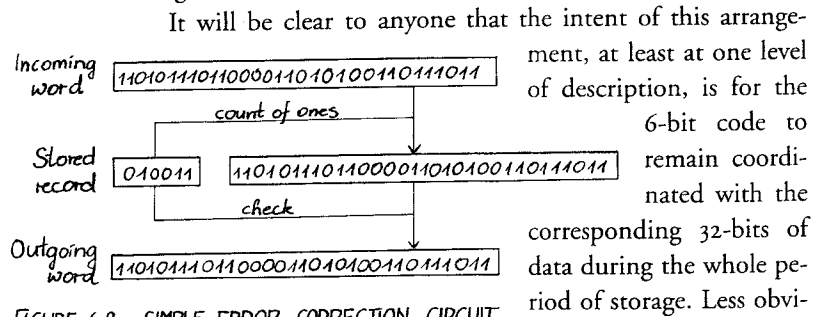


FIGURE 6-2 SIMPLE ERROR CORRECTION CIRCUIT

because it is logically prior—i.e., because it is such a basic assumption underlying the whole situation that we do not tend to

think about it explicitly—is the fact that the 6-bit code should remain *effectively separated* from the 32-bits during storage. It would not do, after all, for the 6 bits to be *invariably* correlated with the 32 bits, locked in nomologically by causal laws, like the gears in the gear world. If that were true, and a cosmic ray accidentally altered one of the data bits, the code would *automatically be updated*, rendering the entire scheme useless.

The situation is similar to that of the super-sunflower. First, code and data are effectively coupled; that is necessary to drive the original circuit. But then the two are crucially pulled apart—sufficiently separated so that effects to one (such as radiation or impurities in the magnetic material) do not affect the other. Separation means flex, in other words; the *point* of the arrangement is to allow the data, at least potentially, to flop around under some external influence without disturbing the code. Then, when the word is read, what was temporarily separated is brought back into contact in order to ensure that the coordination between the two has been maintained.

At an only slightly more complex level, consider file caches. Once again, they are governed by an interlocked pair of normative criteria. The one that normally occupies our attention is the condition that the cached version remain coordinated with the version on disk. Again, however, it is equally important, if easier to implement, for the cached version to be effectively separated from the one on disk, so that causal access to the former does not automatically engender causal access to the latter. If the second criterion were violated, the cache would be correct but useless, because it would not be able to run any faster than the disk version. As it happens, the coordination conditions are more complicated than in the memory system, because of the dynamics. Although perfect correspondence is required over the long haul, temporary misalignment can often be tolerated (e.g., in caches that do not implement what is called “write-through”), so long

¹⁴ Surprisingly, six error bits are needed, not five, because there are 33 different possibilities (0–32), not 32.

¹⁵ Real-world error correction codes are much more sophisticated than this. Not only will they detect multiple-bit errors; they can also correct them. Even more impressively, they can do so symmetrically. Thus on some computers (the DEC-20 was an early example) it is possible to remove any memory module from the machine, while it is running, and the machine will continue on without error—using the “information” encoded in the remaining modules to compensate for the loss. A similar strategy is employed in so-called “RAID” (redundant array of inexpensive disks) storage systems, in which a failure of any individual disk can be compensated for by information retained on the others. Compact disks (CDs) push correction technology in another direction: they are designed to be able to survive the loss of massive errors—hundreds of bits in a row—so that even in the presence of fingernail scratches the band can play on.

as care is expended to ensure that systems using the cache cannot detect the difference.

In all these situations, what starts out as effectively coupled is gradually pulled apart, but separated in a such a way as to honor a non-effective long-distance coordination condition, leading eventually to effective reconnection or reconciliation. There is a great deal more to intentionality than that, and a great deal to say about what constitutes coupling, coordination, and so on, but in various forms these notions of connection, gradual disconnection, maintenance of coordination while disconnected or separated, and ultimate reconnection or reconciliation permeate all kinds of more sophisticated example. There is nothing more basic to intentionality than this pattern of coming together and coming apart, at one moment being fully engaged, at another point being separated, but separated—this is the point—in such a way as to stay coordinated with what at that moment is distal and beyond effective reach.

Throughout, too, it is easy to see how the flex and slop are implicated. To start with, they underwrite the very notion of separation—and hence of connection, disconnection, and the limits of effective reach. If, as in the gear world, there were no flex, there would be no warrant for saying that two parts of the world had come apart. In fact there would probably be no warrant for saying that the world had parts at all, and certainly no warrant for saying that any two parts were far away from each other. Furthermore, exactly because of this slop, it is not automatic that the proximal system will stay synchronized with what is distal. This is why I said at the outset that there is a sense in which slop “establishes the problem” of intentionality.

It is an extraordinarily important consequence of this slop, moreover, that the proximal system can relatively easily adjust its own state (i.e., with modest energy expenditures, like the moth) without having to drag much of the surrounding territory with it. This is how slop provides the wherewithal for intentionality’s

solution (in more familiar language, it is the reason that computation is cheap). The trick, then—again, in some sense a simple strategy, but a very general one—is for the proximal system to exploit this local freedom in order to adjust itself in ways that establish or maintain a coordinating regularity with (at least part of) that from which it is separate.

The same primordial slop implies that this coordination will in general be approximate. It also implies, as I will argue later, that coordination can settle on some but not all aspects of the distal situation.¹⁶ Even if coordination is inexorably approximate, however, it is important to note that the fact that it is possible at all puts conditions back onto the metaphysical nature of the slop itself.

If the flex were too little, as we have already seen, the world would lock up like the gear world, and everything would be correlated with everything else. Such a world would be too rigid, too straight, too stuffy; intentionality would be neither possible nor necessary. If the flex were too great, on the other hand, it would have the opposite problem: things would be too loose, everything would be random, and effect-transcending coordination would be impossible. Imagine, for example, an infinite space randomly occupied by an indefinitely large number of particles, all of which drift aimlessly around, none of which ever interact. Such a world would not have enough structure or regularity for there to be any metaphysical fact of the matter about whether these particles were of the same or different type—or even whether they were particles; it would be insufficiently regular for there to be anything coherent at all. But of course in our world we have neither problem. It is only because the world is *intermediately flexible* that there is anything like life as we know it.

¹⁶The coordinating regularity (what will ultimately become the semantic relation) and the coordinated-with regularity (what will ultimately become the type structure of the referent) will emerge together; they are not as ontologically independent as the phrasing in the text would imply.

The word 'play' is sometimes useful to describe both the world's flex and the intentional behavior that deals with it—i.e., to connote something crucially intermediate between chaos and rigidity. There must be some degree of habit or pattern or at least inchoate regularity in order for it to count as play. Yet play is neither itself, nor does it anywhere require, a straight-laced core of stringent formal rules. It lives—indeed emerges—in the middle, like a spontaneous dance, or like an improvisational session in jazz. It is as fundamental a fact as any about this metaphysics that it is based on an ineliminable notion of “playfulness”—a kind of irreducible, obstreperous, wily refusal ever to be formally captured and written down. Truly a metaphysics for Coyote. And just as a dance is not the property of one or other partner, or jazz the product of a single player, so play is not a feature of subject or object alone. As with so much in this vicinity, it is not even so much shared as it is prior to the distinction between them.¹⁷

That is not to say that there are not ways in which play is quite a misleading metaphor. For even if the world is unbearably light, there is nothing non-serious about these patterns of intentional coordination. Everything, it should be clear, is at stake.

Though massively too brief, this leads to the first characterization:

Intentionality is a way of exploiting local freedom or slop in order to establish coordination with what is beyond effective reach.

3 Disconnection

Because its *raison d'être* is to bridge gaps in effective coupling, it follows that registration, and intentionality more generally, are

not effective phenomena. This will have methodological consequences. If intentionality is not a causal phenomenon, for example, it follows that it will not have a causal explanation.

This is a strong statement, but it is not the strongest statement I will make. In the wings lies another, even stronger: that objects, too, are not causal—and thus that no theory that traffics in objects will ever be entirely causal. In fact it will turn out that *no* theory is entirely causal. In a sense it should already be clear why. If what it is to be an object is in part to be constituted by regularities in which intentional subjects are participants (a thesis that was advertised or at least admitted in chapter 3, and which will presently be claimed and defended), and if intentionality requires distance and non-effective coordination, then it follows that all stories about objects, and indeed all stories in toto (since the point holds of registration, ultimately, not solely of objects) will inherently trade, more or less explicitly, in non-effective regularities.

This is all fine, from my point of view. I see no reason to suppose that anything important is thereby given up (for more details see the “causation and effectiveness” sidebar on page 210). On the other hand, independent of one's reaction to the proposal, the non-effective nature of intentionality should be perfectly familiar, especially on what is taken to be the semantical side of the story.

Thus consider what in contrast to effective reach (that near-and far-engendering bumping and shoving of the physical substrate) I will instead call the *semantic* or *intentional reach* of all registrational phenomena—the fact that, as I have been saying since the beginning, as intentional acts, they are directed, in Brentano's sense, towards something above and beyond themselves. The fact that registration reaches towards something paradigmatically other is an essential component in achieving the transcendent aspect of objectivity. The most impressive thing

¹⁷The notion of play being dealt with here stands in contrast to the much more rigid, as I would spell it, “Play” that underlies formal game theory, the notion that has historical connections with the rise of our current notions of digital computation. See Franchi (forthcoming-b). See also Haraway (1991), pp. 199 & 201, for discussion of Coyote as a symbol of irreducible metaphysical wiliness.

Causation and effectiveness

It is claimed in the text that intentionality is not a causal phenomenon, and will not have a causal explanation, and, even more strongly, that objects are also non-causal, and will similarly not have causal explanations. Strictly speaking, this is only true to the extent that 'causal' means 'effective.' The real claims are that intentionality is not an *effective* phenomenon, and will not have an *effective* explanation, and similarly that objects are non-*effective*, and will not have effective explanations.

In TMD-III I argue that the "theories of computation" mentioned in chapter I—i.e., a rough amalgam of the theory of effective computability, recursion theory, complexity theory, and the like—are not, in fact, theories of computation at all, because they do not deal with computation's inherently intentional character. What they *are* are mathematical theories of effectiveness, simpliciter. If effectiveness could be simply defined in terms of something else, this body of work would not be as important as it is. The way to understand it, I think, is to say that this large body of (mostly twentieth-century) work has recognized, and converged on, effectiveness as a fundamental metaphysical category.

What does that have to do with causation? Who knows, exactly? It is hard to say anything *exact* about causation, for one thing, and I claim no special insight or expertise on such a famously recalcitrant notion. I would guess, however, that the answer will be this: that what in ordinary language we mean by 'cause' (and thus 'causal') will not succumb to intellectual or scientific reconstruction, but that this recent (as it were, computational) formulation of a theory of effectiveness will eventually be regarded as its substantial scientific reformulation. There is another possibility, though: that causation is, in chapter 5's sense, a *material*, rather than a physical *or* effective, notion, and in that form will be retained.

about it is its extraordinary span: the fact that it so strikingly unconstrained by the limits of physical connection. Browning's adage that one's reach should exceed one's grasp¹⁸ may be a difficult injunction to meet, morally or psychologically, but there is

¹⁸Browning's "Andrea del Sarto" (1855), l. 97.

a sense in which it is accomplished by every intentional act of registration. That semantic reach exceeds effective grasp is essentially a theorem of this metaphysical account.

By way of contrast, consider laser beams. They reach a long way, too—to the moon and back, for example, allowing us to measure precisely how far away the moon is. So do radio waves; we are just now picking up signals that were emitted by now-distant quasars only shortly after the beginning of time. But in terms of reach these phenomena do not hold a candle to reference. With a few simple syllables we can reach backwards in time, against the flow of causality, to the Pharaohs of Egypt. Or reach forward, to things that have not yet happened, such as the election of the first female U.S. president. Or to Pluto, without having to wait five or six hours for our reference to succeed. Or to Huckleberry Finn, without even needing our referent to exist. "Reference," Alonzo Church once said, "travels at the speed of logic."¹⁹ And that is just the beginning of its mysterious powers. It also goes where reality fears to tread.

This fundamental *disconnection*, which I am taking to be essential to all registration, explains why being referred to is not physically detectable—e.g., why not even the National Security Agency, with all the money in the world, could build a meter to register when the control room in Cheyenne Mountain was the subject matter of a terrorist's intentional act. Indeed, intentionality's disconnection is so basic that it is impossible to imagine its being false. Error would be eliminated, but so would abstraction. Hypotheticals would be impossible; fantasy lives would be metaphysically banned. You could not even think about continental drift without dragging the tectonic plates along with you.

Although I have forsworn a sharp representation/ontology boundary, these examples of disconnection, in part in order that they can tap into traditional (predecessor) intuitions, have all been taken from the semantical side of the equation. Since, as I

¹⁹Public lecture, Stanford Center for the Study of Language and Information, May 3, 1984.

presume has been implicit throughout, I am going to take registration to be the ultimate subject matter, and therefore will not swear allegiance to a sharp representational/ontological split, it follows that anything true of semantics—i.e., of intentionality—will also and indissolubly be true of objects. Even if it is not as clear why, that is, this essential distance-transcending disconnection is in fact just as much a truth of ontology as it is of semantics.

It would be happier overall, however, to be able to tap accessible intuitions leading to the same conclusion, rather than having to argue for such a strong conclusion on what amount to relatively abstract grounds. And once the point is grasped, I believe, there is no dearth of supporting intuitions. A glimmer of one can be seen, for example, by noting how an ordinary and uncontentious property of individuals betrays their non-exclusively physical character. Part of the “gathering together as a unity” that is constitutive of an object’s being an object is a *temporal* gathering—the fact that objects, by and large, endure over time. Physics, however, is intrinsically ahistorical: there is no action at temporal distance; no dependence at time $t+1$ on how the world was at time $t-1$ unless that effect is propagated through time t . That alone is enough to imply that being an object *cannot be a physically effective property*, since there is no way in which a temporally extended object could have a physical effect distinct from that of a non-temporally-extended instantaneous time-slice of that object. This is yet another reason to support chapter 5’s contention that physics is not ontologically committed to individual objects.

In the end, this point will underwrite the entire story of registration. Abstraction, essential to the notion of an object, is like semantic reach in being physically transcendent.

7 · Registration – I

What, then, do the long-distance patterns of coordination come to—and how do they arise? How do intentionality, registration, and objects take shape on the underlying playing fields?

To answer this question—even to point towards an answer—requires dismantling the puppet show. It is time to get rid of the objects, time to tell the story without the registration of any individuals. But how then is it possible to proceed? We can hardly be asked to set aside all theoretic registration, after all. That would be an impossible standard—especially for a written text (it would be easier if we could have dinner). So the first task is to figure out what options are open to us—what metaphysical and ontological resources we can avail ourselves of.

This is where the rhetorical strategy announced in the last chapter, of adopting a vaguely physicalist or with-reference-to-physics approach, begins to pay off. The overall rhetorical situation, and our position within it, are pictured in figure 7.1. Chapter 4 argued that feature-placing, both intuitively and metaphysically, was a more primitive form of language use than language committed to objects, properties, and relations—and especially that it was simpler than formal languages committed to Objects, Properties, and Relations. That was because feature-placing did not require the individuation of discrete individuals, at least not at the “particular” level. In chapter 5, it was claimed that physics, at least on a field-theoretic interpretation, could (and perhaps should) be understood as an example of such feature-placing. In presenting the physical case, I focused on continuous phenomena, because continuous feature-placing is more obviously unlike objectification (registering in terms of discrete individuals) than discrete feature-placing, and is therefore less liable to be confused with it—though I said at the time,

and still maintain, that this is more a matter of convenience than of substance.

With respect to the current goal of dismantling the puppet show, this means that even if we cannot let go of all registra-

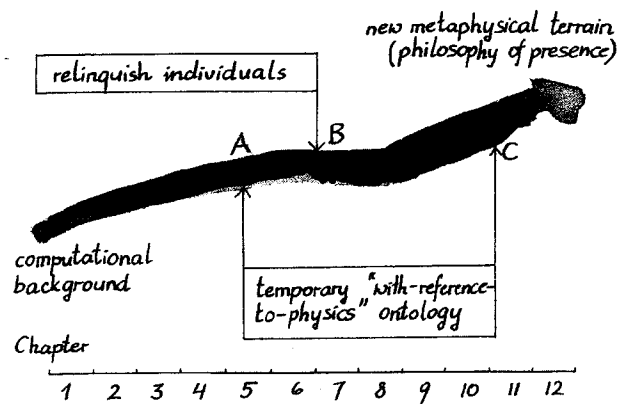


FIGURE 7-1 RHETORICAL STRATEGY

this strategy is more rhetorical than metaphysical, and is accompanied by no commitment to theory-reduction, there is no great need to worry about exactly what kinds of fields are allowed. I will therefore speak very informally about density gradients, simple space-time manifolds, and the like. Whenever it is convenient, moreover, I will let go of continuity assumptions, for example when drawing attention or analogies to turbulence and chaos.

Given this much apparatus, the aim in this and the next two chapters is to convey an intuitively accessible image of how individuation, registration, and the like might arise. Then, in chapter II, with this feature-placing image held as firmly in mind as possible, I will let go of the physicalist props that have been used to support it, and—very quickly, before everything falls down—rearrange the assumptions underneath it, partly in order to set physics back into its proper place, but also, more impor-

tion, a substantial option nevertheless remains open to us. As indicated at point B in the figure, we can let go of individuals, and for the next while register the world in terms of (often, but not necessarily, continuous) particular fields. Since

tantly, in order to rest that temporarily unsupported image on a more long-term, tenable, metaphysical foundation. At that point, finally, the account will meet the requirements (about realism, pluralism, and objectivity) laid down in chapter 3.

1 Tracking

Start, then, first, with a simple case of visual tracking. Tracking, on this account, is not registration, because it is not appropriately disconnected. Or at least nothing is registration merely in virtue of tracking. Sometimes we both track and register, as for example when we consciously follow a bird's flight, or when we register a moving image on a television screen as the runner in a World Series game. By itself, though, simple tracking—by which I essentially mean no more than a causally-driven servoloop—need not be accompanied by, and is by itself insufficient to support, this extra intentional baggage. Still, in spite of these caveats, tracking is a good example on which to gain experience in analyzing a simple situation without benefit of object registration (again, at least at the particular level)—and a far from random one. Even if tracking is not itself intentional, it is an important precursor to some of the simplest forms of behavior that are.

To be specific, therefore, and to develop some concrete intuitions, consider a frog tracking a fly.¹ Except now register the situation without any individuals or boundaries. In the spatio-temporal region of the fly, there is a differential density mass with a complex internal structure. For simplicity, suppose that a

¹ Strictly speaking, frogs do not track—at least not in the sense of moving their eyes (except to compensate for overall motion, as for example when sitting on an unstable lily pad). That does not really matter, here, since the point of the example is purely pedagogical, but anyway it is not eye-tracking that is my concern, but *attention*-tracking—or, even more minimally, the existence of some circuitry to respond to the shadow of potential prey moving across the retina, so as to orient the flicking of the tongue. And that capacity I take it frogs have.

roughly even spatially distributed flux of electromagnetic radiation (constant illumination) flows into the region. In the patch of higher mass density, various reflections and obstructions cause a local disturbance in the radiation pattern. Some of this (reflected) radiation, after making its way through an intermediating region of air, impinges on what we are tempted to register as the frog's retina. But of course there is no frog, yet, nor any retina—merely hills and valleys and spikes and skews in the various relevantly oriented feature fields, corresponding to the fly, the frog, the focusing properties of the eye, and so forth.

If, as theorists, we look in the vicinity of the frog's retina, but let go of any tendency to register objects, we can see (what amounts to a local projection of) a patch of disturbance, extended in both space and time. If we temporarily hold time still and examine a spatial region (i.e., focus on a region on the retina), we see a local distribution of activity—an increased or decreased level of nerve firings in the center, dropping back to a background level in those parts of the retina that are outside the area of disturbance. By the same token, if instead we hold space constant and watch over a period of time (i.e., study the activity at a given retinal position for a moment), we see something structurally homologous: a welling up or welling down of activity while the “fly” passes, subsiding to a level corresponding to the background, as the fly goes on by.²

As said earlier, this localized pattern of activity does not count as even inchoate registration, because it is a phenomenon of pure “connection.” Still, it carries an important moral. A property we ultimately associate with both subject and object—that of a boundary or edge—is in the first instance *a property of their interaction*. What we register as the edge of the fly (strictly speaking, the two-dimensional projection of its three-

dimensional boundary tangential to the direction of illumination) corresponds directly to an “edge” of a very important column: a column of differential energy levels in the incident flood of frog-directed electromagnetic radiation, which corresponds in turn, through all the mediating mechanisms of physical coupling, to an edge of an impulse or differentiable region of energy in the retinal circuits in the frog.

The crucial point to see is that this transient welling-up of activity or difference—this local extended impulse in the manifold of sustaining physical fields—is columnar in shape, cutting at right angles to the frog-fly boundary. It is solid, too: a continuous physical disturbance, made up of a variety of kinds of material reaching all the way from retina to fly. To repeat: it is not yet bipartite. As observers we tend to chop it into two non-adjacent material parts: one out at the fly, another in the frog's eye. That is, we tend to view the mediating air as *conceptually* as well as *optically* transparent—as if it were as evacuated of theoretical significance as it is of light-obstructing material. But the transparency of air is of course not yet a “fact” for the emerging participatory scene. Indeed, the patterns of electromagnetic radiation in the air are just as thick and real as the electrical signals coming off the retina—and not yet obviously distinct from them, especially given that they are fashioned of the same (electrical) material. Indeed, as was said in the introduction, how it is that we and perhaps frogs see flies, not electromagnetic radiation, is the registration problem in a nutshell. So it is vital not to build its achievement into our conception of the initial situation. That would be a classic case of pre-emptive inscription error. Gibson was too mild: it is not just that we do not represent the tree; we are not even distinct from the tree, at least not at first.

And so the initial boundaries or edges that allow the frog to see the fly as fly—or at least allow the frog to see the fly as whatever it sees it as, i.e., in its own register—are what theory will eventually call *relational*. These boundaries matter, too; it is no

²The levels may be inverted, of course: less activity when coupled to the fly, more when coupled to the background. It does not matter to the argument; and anyway it is not clear that we have enough machinery to discriminate the two.

accident that edge detectors are among the first neuronal circuits to get into the registrational act. Except that to call them edge *detectors* is both misleading and expensive. It is expensive because it describes the situation in terms dangerously close to the structure of the solution ('detect' being a fully intentional word). It is misleading because all we have, so far, is a causal loop involving subject and environment. Without further patterns of disconnected coordination, there is no justification for invoking the heavy language of detection—or, for that matter, for saying that the frog *sees* anything at all. To register the circuitry in the frog's eye as an edge detector, without paying its cost, is to over-interpret the situation. It hides the fact that the mechanism simply involves *more accreted boundary*.

Even more specifically, consider the otherwise ordinary-sounding claim that an edge detector in the frog's retina is *fired* or *triggered* by the edge of the image of the fly. For a theorist to register a frog's neuronal activity in that way is to commit to a coincidence of three boundaries: (a) that between quiescence and firing, being triggered or not, in the contended mechanism in the frog (on or off); (b) the corresponding edge of the fly (*being fly*, and *not being fly*, in the mass density region several feet in front of the frog); and (c) the edge of the intermediating pattern of directed electromagnetic radiation. But then, once one realizes that no warrant has yet been provided for separating these three boundaries, one sees that what is really going on is best understood in purely field-theoretic terms: there is a single common edge to a columnar-shaped multimedia disturbance that reaches continuously between (what we register as) frog and (what we register as) fly.

Furthermore, if, in spite of this recommendation, one were to insist on a discrete registration of these retinal mechanisms, it would be more accurate and more modest, at least in the first instance, to call them edge *participators*, rather than edge *detectors*. For they manifest just this simple property: the edges of their

own activity are aligned with, in virtue of being connected to—wired or plugged into or effectively coupled with—the edges of the target entities towards which the systems of which they are part will ultimately be said to be directed. The situation is much closer to that of a sports jacket than to that of recognition. It would be bizarre to say that your jacket “detects” your motion—even though, sure enough, it changes state in a way that is lawfully correlated with your motion. It changes state subject to the constraints of overarching physical law because it is *connected* to you—by coming along with you. So too of retinal activity and flight of fly.

Moreover, what is true of edges (difference) is equally true of stability (sameness). Think of perceptual invariants. In the first instance they too are invariants of neither the ultimate subject nor the ultimate object. They are invariances of *coupled interaction*. Again consider simple tracking. The object moves, the subject moves, a constant relation between the two is maintained. Except now relinquish the registration in terms of objects—i.e., put on your field-theoretic glasses—and look only at the wells of stability in the otherwise seething mass of activity. As before, the important stability lies *across* the boundary that will ultimately be drawn between subject and object.

2 Separation

Simple effective “tracking”—a misnomer, of course³—is an instructive case in which to exercise one's skills at visualizing non-individuated patternings of structure. With respect to the substantive question on the table, however, regarding the rise of registration, there are two problems: (i) ontologically, proto-subject and proto-object are insufficiently distinct; and (ii) physically, they are insufficiently separate. Neither, therefore, yet counts as an object. It will be rhetorically helpful, therefore, to refer to them as *s-region* and *o-region* respectively. Even though

³It is a misnomer because in the simple effective case there is not enough effect-spanning coordination to legitimate the term ‘tracking.’

'region' is a count noun, and so may be said to have singular reference, these phrases are not strongly individualist in connotation (i.e., do not strongly signify boundaries), and so may be of some help in guarding against pre-emptive inscription.

Imagine, then, to consider a simple case of separation, the rabbit, mentioned earlier, at the instant when the coyote it is watching disappears from view. We have already seen how the tracking can, up until that point, be physically driven by the object being tracked. The causal flow—or rather, as it would be better to say, the *flow of effect*—reaches out from rabbit to coyote and back (except for the standard comment that, at the level at which the effective feedback is maintained, there is neither rabbit, coyote, nor boundary—just a feedback loop maintaining a pool of stability in the midst of the flux).

But now, suddenly, the coyote vanishes, and the causal loop is broken. In order to maintain coordination, it is *suddenly incumbent on the rabbit to take over responsibility for keeping "focused" on the coyote*. An internal mechanism has to compensate for what can no longer be relied on to be effectively provided by the environment. In a similar way, the super-sunflower, when it can no longer be driven by the incident radiation, has to employ a different, internal mechanism (some form of chemical drip, or whatever) in order to continue to "track" the sun.

This *retraction* of responsibility into the s-region, as I will call it—this shouldering of effective responsibility by the s-region, to compensate for the break in effective coupling—is no less than the origin of reasoning, representation, and syntax: the effective projection, onto the intentional agent, of the requisite arrangements for maintaining long-distance (semantic) coherence. The break for which it compensates is not a yes/no affair, of course, and so it would be better to use more gradualist language. The point, though, is that, as connection thins out, as it did with the view of San Francisco when the fog thickened, and the o-region recedes from effective reach, responsibility for

Memory and prediction

The examples used to illustrate disconnection in the text are primarily spatial, involving one mechanism (plant, animal, etc.) adjusting its state in order to stay coordinated with a contemporaneous phenomenon that is at the moment out of effective reach—i.e., is too far away to lock onto, causally. But as usual, what is true spatially is also true temporally. Because of physics' ahistoricity, any form of long-distance temporal coordination, either forward in time (prediction) or backwards in time (memory) will similarly involve a degree of disconnection.

Thus consider a simple case of prediction. In order to know now, at time t_1 , what is likely to happen in the future, at time t_2 , it is not possible to be effectively (causally) driven by the relevant t_2 state of affairs. To do that would either require waiting until t_2 , which by hypothesis would be too late, or else depend on some form of backwards-in-time causation, contrary to physical law. Instead, one must *predict*, which involves the same sort of retraction described in the text: a shouldering of effective responsibility by the s-region in order to develop or construct local and current (i.e., *present*) structures standing in coordinated relation to phenomena that are still temporally remote.

The same holds of memory. A memory can be created in virtue of effective coupling, but it cannot be sustained—"held up"—by relying on that coupling, because the driving situation will soon go away. In order to make a memory more durable than a shadow, an s-region or subject must first allow or arrange for an appropriate impression to be formed by the event in question, but must then take responsibility for storing it in such a way as to ensure that it will continue to be effective after the event it records has dissipated. As described in chapter 8, the memory may also need to be updated appropriately, as circumstances relevant to its contextually-dependent content change, in order for it to continue, over the long haul, to signify what is no longer present.

maintaining coordination has to be retracted from the region of interaction onto the body of the s-region. Imagine, to get a feel for this, that the coordinative structure is made of silly putty stretched into some complex shape between two ends, which is then pulled apart, more and more, until it snaps. At the crucial moment, when it “lets go,” a bunch of the putty that had been stretched out in the middle, between the objects, would suddenly spring back and pile up at the end. So it is with the structural responsibility for maintaining long-distance coordination. Separation between s-region and o-region is necessary for registration. (Not having it is like looking at something too closely; you have to back away, have to get enough distance, in order for an image of something as something to snap into place.) But at the same time the separation exacts a cost that the s-region must pay. What will end up being subject must make up for the fact that the now-disconnected o-region can no longer be relied on to hold up its part of the bargain, its part of the dance.

As soon as this happens, moreover, an array of properties characteristic of all registration is exhibited. Even in these elementary cases of non-effective tracking, in other words, with their relatively unsophisticated forms of coordination, features that govern much more complex cases gain an initial toehold.

First, because of the break in effective coupling, and the concomitant retraction of responsibility for maintaining coordination onto the s-region, a distinction takes hold between *what the s-region is doing* (tracking the coyote or incident sunlight) and *how it is doing it*.⁴ The former gets at a non-effective regularity; the latter, at an effective mechanism whose “job” is to implement or sustain it. Among other things, this split provides a toehold for normativity to attach its tentacles.

⁴I have not yet justified the claim that the rabbit really is tracking the coyote, on an intentional reading of ‘tracking.’ The point is only that if it is tracking the coyote, which will depend on the nature and existence of various other patterns of coordination, to be described below, then the break in effective coupling unleashes this familiar explanatory split.

Second, there is suddenly the possibility for error. Because of the underlying slop, there is no reason to guarantee—and every reason to doubt—that the retracted mechanism will be able to mimic perfectly the distal regularity with which it is “striving” to maintain alignment, with which it is striving to remain registered. (Except in rare formal cases, the idea that syntax can perfectly mirror semantics is a nostalgic modernist ideal.) The problem is that the underlying flex eliminates any guarantee of alignment with what is distal. The rabbit has no way to be absolutely sure of what the coyote is doing. The super-sunflower case is a little different: the progression of celestial bodies is astonishingly regular, as it happens, which is why such accurate non-effective solar trackers can be built. But it is unlikely that any sunflower would have access to such precise mechanisms. And anyway the moral is very general: as separation increases, the chances of the s-region and o-region drifting part also increase, perhaps irretrievably. Nor is this a “new” fact, with its own meta-physical price. Being liable to drift apart is what separation is.

Third, the retraction of responsibility onto the s-region is the origin of registration’s asymmetry and directedness. Note, by way of background, that an object’s being an object is a collaborative affair, arising out of the registrational dance of subject and object (s-region and o-region). Logically, therefore, there is a sense in which subject and object are symmetrical in importance; neither can exist without the other. But I have never implied that the dance itself is symmetrical. On the contrary, registration is intrinsically asymmetrical and directed; it is what orients the subject towards the world (hence the use of a transitive verb). And already the origin of this asymmetry can be seen. For although the gap in effective coupling is symmetrical (if x is disconnected from y , then y is disconnected from x), *responsibility for maintaining coordination across the gap is asymmetrical*. It “falls to the s-region,” one might say—except that gets it backwards. Rather: to be an s-region is to assume this re-

sponsibility. In higher cases the responsibility will assume infinitely more complex forms, but even in these aboriginal cases registering—taking the world to be world—requires commitment.

Fourth, the example erroneously suggests that the mechanisms for coordinating with what is beyond effective reach must be inside the s-region. This is obviously false. To maintain orientation towards an invisible sun, for example, one might establish effective coupling with a local clock and rely on it to maintain the long-distance synchronization. S-regions need the capacity to track, in other words, but they need not, and generally will not, do it alone. More generally, the “distance” over which the s-region will be able to maintain coordination will depend on its cleverness—not just its internal cleverness, but the total sophistication of the environmental and cultural resources on which it can lay its hands. Maintaining coordination over vast distances, putting in place mechanisms to “select” what is remote and to “wash out” what is immediate, tracking the subtlest remote coordinations in spite of extraordinary intermediate upheaval—these are no less than the role of instruments, skills, theories, inference, and civilization.

3 Stabilization (or: How can I miss you, if you won't go away?)

I have tried to sketch a way in which an s-region can compensate for separation from the o-region. But there is still a problem. The described non-effective tracking mechanisms still suffice only to maintain perceptual invariants: quietened, somewhat abstract (though still particular), deictic “relations between subject and object.” They *compensate* for the effective separation, and thus succeed in maintaining relational stability. But they do not yet *exploit* the separation—and thus do nothing to stabilize the object as an object (or as anything else, for that matter). Im-

portant as non-effective tracking is, in other words, it is too simple, on its own, to yet count as intentional.⁵

The challenge for the s-region is to exploit the separation in order to *push the relational stability out onto the world*. In order to register an o-region as an object, that is, it is not enough for the s-region to move in synchrony with the o-region, even if that synchrony transcends physical coupling. Pointwise synchronization is just (so-called) perceptual invariance, as described above; to maintain it is just to stabilize the relation between s-region and o-region relation. To register, the s-region must do more: it must play with the relation, even at the cost of a certain amount of destabilization, so as to stabilize the far end. That is, the s-region must “sediment” or “extrude” the o-region as a discrete autonomous individual (at least partially distinct from the s-region) and thereby locate it in the wider world. *It is the o-region itself, in other words, not the relation to the object, that must be “held” in at least partial stasis.*

Doing this requires *detachment*. The s-region must “let go” of the o-region, in a very special sense of letting go: push it away, make it separate from itself. Trying to register an o-region to which you are fully coupled is like trying to see something plastered up against your eyeball. Or like holding onto a jack hammer for dear life. Or like having an insecure sycophant fall in love with you. To be welded to an object like that means being buffeted by its every vibration or variation or whim. Lacking the requisite distance to gather up and “hold” the object as a stable

⁵Though non-effective tracking can and likely will figure as an essential constituent in a larger system that is intentional.

There is something very Gibsonian about this picture: though he would say “represent” instead of “register,” the point is very similar. Looming, for example—the “perceptual” phenomenon whereby an approaching object, such as a wall, casts an expanding image on the retina that is “decoded” by special-purpose circuits that play a direct role in engendering avoidance behavior, and the like—counts neither as symbolic representation for Gibson, nor as registration, for me.

unity, the s-region and o-region are instead (even if only temporarily) liable to fuse, causing both to disappear. Bringing an object into intentional “focus” is thus not only a way of coping with gaps in effective connection; it *requires* gaps in effective connection. The object requires these gaps, in order to settle into being what it is, distinct from the subject, with at least a partial degree of existential independence or autonomy. And the subject requires the gaps, too, to say nothing of a modicum of its own autonomy or independence, *in order to register the object as a unitary whole*. Separation—keeping things at arm’s length—allows the relation between subject and object to relax.

Separation per se is not very interesting, of course. We are effectively disconnected from virtually everything in the world, without thereby registering any of it. Rocks are pretty disconnected, too, and they do not register at all. The issue is what one does with one’s separation—what one makes of one’s “aloneness.” The trick, and it is a considerable trick, is to arrange oneself appropriately, and dynamically, *so that the surrounding world remains present*. By exploiting the local slop, one must arrange one’s own effective structure so as to maintain coordination across the break in coupling. If done just right, one can let go of the world without losing directedness and orientation. More strongly: it is exactly by letting go while retaining appropriate directedness and orientation that the world “comes into presence.”

The key to this coordination, and the importance of the detachment, can be found in an remark made in chapter 4 (§1c): that an essential act of abstraction underwrites the existence of every individual. As was said then, taking an object as an *individual* means gathering up an extended region of the flux and treating it as a unity. This implies that the subject must stand in relation to what is the *same* or *in common* across the constitutive spatial and temporal region, and by the same token must ignore or set aside the multitudes of internal variation attendant to its

Abstraction and error

Note the intimate relation between the processes of abstraction described in the text and the aforementioned possibility of error implicit in the underlying slop. If you are not going to be able to maintain coordination with various details of the separated object, it makes sense to set those details aside. In that way an inability to track them does not eventuate in *failure*; it merely eventuates in loss of detail. It is a good thing that we have a registration of Pompeii that abstracts away from the question of exactly how many people lived there at various times, since we would otherwise be unable to register it at all.

In spite of this convenience, however, it should not be thought that the ability to abstract is a compromise, a good way of accepting limitations. A far more important fact, yet to be dealt with, is that there are regularities—stabilities—that can *only be established at these higher levels of abstraction*. This fundamental fact provides the metaphysical warrant for the so-called “special sciences,” to say nothing of vitiating reductionism and validating pluralism.

parts, or across its life. When I register that passing swoosh as a hawk, I gather up its movements under a general notion of soaring, but insulate myself from the exact number of feathers, or the angle it makes with the sun. When you think of your brother’s upcoming marriage, you think of him as a whole, not as a temporal cross-section, or even as a continuous evolving manifold of space-time points.

Abstraction, in other words, requires separation, in order that the s-region not be buffeted by irrelevant details. Moreover, it also imposes an important symmetric condition on the patterns of coordination. It is not enough for the s-region merely to abstract away from the excruciating particular in order to conceive of the o-region as an individual or unity or whole. In order to register the o-region as an individual, that is, it is not enough

for the *s*-region to take the *object* to be extended (though that alone would be enough to implicate a certain amount of detachment). It is essential that the subject's own registrational state be extended as well. Subject and object must both be stabilized, in other words; both require extension. And not only that; *the two extensions must cross cut*. Each-to-each (piecewise coordination) is not enough; it must be each-to-all.

Consider an example. Today, when I think of you practicing in a loft in Vancouver, I think of a whole you, including the you that will be performing tomorrow. Tomorrow, when the concert is over, and I wonder how it went, I will again think of a whole you—including the you that sweated it out today. It would not be enough, in fact it would not count as thinking, if today I could only relate to the-you-of-today, and tomorrow could only relate to the-you-of-tomorrow. Or rather, more seriously, it would not be enough if at any particular instant, I could only relate to the-you-of-that-instant. Because there is no such thing as you-at-an-instant. You have inherent temporal extent.

Although both subject and object must be stabilized, the two stabilities cannot merely end up pointwise correlated, particular to particular, like a ruler laid down against a run of cloth. As depicted in figure 7-2, this kind of nomic (and instantaneous) pointwise correlation is how ordinary coupling in the underlying physical fields works; at any given instant, a region is impinged upon only by what surrounds it *right then*. Such is the inescapable implication of physics' ahistoricity. Only by making egregious inscription errors—i.e., by pre-emptively assuming object identities, and thereby taking out unpaid debts of temporal extendedness, as in the case of the cell in the cat's brain—could anyone imagine that intentional directedness (semantics) is anything like causal connectedness. "Sufficient unto the day is the evil thereof"⁶ is curiously apt as a description of the physical world. For better or worse, but mostly for better, it cannot stand as a general characterization of registration.

⁶Matthew 6:34.




FIGURE 7-2 POINT-TO-POINT CORRELATION

The cross-cutting nature of the extension of subject and object, in contrast, is reminiscent of Husserl's distinction between the temporality of perception and the perception of temporality, and even of Kant's distinction between the unity of consciousness and the consciousness of unity. That it is implicit in the very notion of an individual can be seen by understanding the topology of the notion of reidentification that is constitutive of being an individual (unless one can reidentify, one cannot identify

—that was why reidentification was built into the initial characterization of what it was to be an individual).

Consider in particular what is involved in reidentifying at time t_2 an object x that was first identified at t_1 . First, at t_1 the subject must be in such a state as to register the *whole* of x ; if at t_1 it were only registering the t_1

slice of x , there would be no chance of *that* showing up again at t_2 .⁷ Second, at t_2 it must again register all of x , so as to be the same extended individual registered at t_1 . As indicated in figure 7-3, both points in the subject extension have to span or cover the whole object extension (except that it is a mistake to think of them as "points"—see the sidebar on page 230). Third, the state of the subject at t_1 and the state at t_2 must themselves be stabilized into an enduring unit. Else there would be no warrant for the claim of reidentification (this is the situation that was described in §4 of chapter 1). Mechanically, the stability of the subject can seem trivial, if one cheats by using unpaid-for registrations. It is the "same cell" that lights up, one says, or the "same name" that is used. Or perhaps the "same record" that was used

⁷The text puts the problem *temporally*, but there is an exactly analogous spatial problem, as well: how it is that one sees a *tree*, not merely a two-dimensional presentation of leaf and bark.

Objects are not points

To treat an object *as* an object is (at least in part) to treat it as a unity, but that does not mean treating it as atomic—i.e., as if it lacked internal structure. Because of a common inscription error we are sometimes misled into thinking of (individual) objects as if they were “points”—so concentrated into unities as to become extensionless with respect to all relevant metrics of variation.

Thus consider figure 7.3, intended to show how states of an *s*-region at times t_1 and t_2 are each coordinated with the *o*-region as a whole. The drawing errs by suggesting that the states at t_1 and t_2 can be represented as points. As better suggested in figure 7.4, the true situation is that the cross-cutting extension mandate requires that *sub-regions* of the *s*-region correspond to the whole *o*-region—sub-regions constituting a single, whole or integral “thought” about, or act of reference to, the object. To put it in somewhat Kantian terms, the sub-regions of the *s*-region effect the (minimal) unity of registration; the *o*-region being corresponded to as a whole reflects the registration of unity. Neither will in general be atomic. Moreover, *there may be intrinsic structure in the finer-grained correspondence*, without challenge to the overarching intuition that a single thought about an object must comprehend that object in its entirety. Thus think of an apple falling from a tree. Sure enough, one can register the fall as a unity. But (at least for me), as long as one registers it as a *fall*, there is an intrinsic temporal extent to the nevertheless unitary thought—an extent that corresponds, topologically, with the temporality of the fall. Few of us, if any, can “think” or “register” a fall from bottom to top.*

Taking an object to be a point (i.e., taking ‘unitary’ to mean ‘atomic’ instead of ‘integral’ or ‘indivisible’) is a piggyback error of the sort described on p. 54 in chapter 1. Probably misled by writing, one starts by assuming that the word or representation or thought is itself atomic, lacking (relevant) internal structure. Since that word or thought must comprehend its object as a unity, it follows that the object, too, must (in all relevant ways) also be atomic, since there is (by assumption) nothing internal to the thought to correspond differentially to different parts of the taken-as-unity object. As usual, the error is in the initial inscription: that thoughts must be atomic. There is no reason why object registrations—

dynamic equivalents of terms—need be any more pointillist or homogeneous than the objects they represent.†

As an example, consider Dennett’s (1991) treatment of consciousness. Dennett argues against the idea that consciousness can be “found” in any particular place in the brain, or identified as any particular entity within the overall mental architecture. Superficially, it can seem as if he thereby undermines the ontological status of consciousness. If it is not *in one place*, then perhaps it does not exist—is unreal, or epiphenomenal. But no such conclusion follows (pace his own instrumental metaphysics). All his argument shows is something we should have expected, anyway: that things conscious—both the seat of consciousness, and its phenomenology—are not unitary points, but distributed and variegated. So too are rivers, lakes, watersheds. To register consciousness *as* consciousness is to treat it as a unity, and to register it as a unity is (at least usually) to abstract away from at least some of that internal (and intrinsic) variability. But registering it as a unity does not require that the thing itself be spatio-temporally or in any other way extensionless, any more than to register the Yukon River as a unity is to compress into a point a river spread out over several thousand miles. We owe as much to acts of registration themselves.

*Given a video tape of the fall, it is famously easy to “see through” the video, registering the fall, not the video itself. If one runs the tape backwards, yielding (at least arguably) a representation whose temporality cross-cuts the temporality of the fall, I can either (i) register the video, or (ii) hypothesize a non-real situation in which the apple flew up into the tree. But I am unable to “see through” this reverse video and thereby register the original fall.

†This raises a problem, of course: given that they are *not* points, what are they? And how can this be answered without committing yet another inscription error? The answer is to take the claim in the prior paragraph and read it backwards. It is not that the “two extensions must correspond,” but rather, that what constitutes a “single, integral term” or a “single act of registration” is *that amount of space-time s-region that corresponds to the whole o-region*. Except how do we pay for the “whole *o*-region”? With *reidentification*. Only by exploiting the redundancy implicit in the multiplicity of extension-extension mappings required by reidentification is there any chance of wresting the appropriate individual boundaries out of the flux, instead of unfairly imposing them.

in the first registrational event is tied—say, with an equality sign—to the record of the second, or to the event itself. In all these descriptions, something is carried in the subject from t_1 to t_2 . But once one lets go of the unwarranted inscription, one realizes that the identity of that internal individual is not God-given, not something on which the subject can unquestioningly rely. It, too, has to be achieved. Except of course it is a conceit to

call it an “individual,” since it is only an individual for the theorist. *Subjects have to be stable* in order to register individuals as stable, but they do not have to be *registered*, by themselves or by anyone else.

I have not yet said how this cross-cutting extension might be achieved. Note, though, that

if it is achieved—if, that is, in virtue of such cross-cutting extensions, the subject succeeds in stabilizing an object with which, by hypothesis, it is at least partially out of effective contact—the subject thereby takes a small but decisive step towards the achievement of objectivity. It is an accomplishment, after all, as Strawson puts it, that “we think of our powers failing rather than the world fading,”⁸ when the slop increases to the point that effective contact with the object is lost. Even the simpler cases of non-effective tracking begin the process of standing in coherent relation to what is distal. But it is really only when this cross-cutting extension takes hold that the subject can be said to begin to have a commitment to the continued existence of something that does not currently impinge upon it—and thus can be said to have a commitment to the *existence* of anything at all.

In this way an essential point comes into view: *the presence of an object inherently involves its absence*. The reason is simply the standard one: in order for a subject to take an object as an object,

there must be separation between them—enough separation to make room for intrinsic acts of abstraction, of detachment, of stabilization. So it is essentially an ontological theorem of this metaphysics that no object, for any given subject, will be wholly *there*, in the sense of being fully effectively accessible. Or to put it more carefully: in order to be present, ontologically—i.e., in order to be *materially present*—an object must also be (at least partially) *absent*, metaphysically, in the sense of being partly out of effective reach.

In high-level cases, this ontological necessity of absence is almost commonsensical. Thus to recognize someone as the person they are is to register something greater than what literally confronts your eyes. “Oh, it is Randy,” you say—thereby reaching out (semantically) to a wealth of relationship and extension not otherwise (effectively) available to you: relationship and extension that are irremediably different and “other.” Perhaps surprisingly, the very same structure—of ontological presence

requiring effective absence—holds at the level of the physical plenum. The spatial version is easier to accept, and is the one I have focused on: in order for an object to be present to a subject, it must be held a certain distance away, where distance includes physical distance. The temporal

version is easier to understand in terms of extension. Because objects are in general temporally (as well as spatially) extended, there is never a moment in time in which the whole object is there, since no instantaneous time-slice can ever “capture” more than a cross-section of a temporally extended whole. It only




FIGURE 7.3 CROSS-CUTTING EXTENSION




FIGURE 7.4 SUBREGION TO REGION CROSS-CUTTING

⁸Strawson (1959), p. 74.

seems perverse because our apperception of time is so different from that of space. Suppose I point at an extensionless point in space, somewhere within your body, and ask “Are you all *right there?*” The question sounds a tad perverse, but the only possible answer is *no*. There is no single point in space where all of you is; you have inherent spatial extent. By the same token, to the extent that it is coherent (and I think it is largely coherent), there is no one *temporal* point where all of you is, either; you equally have inherent temporal extent.

One final comment. I have described this process as one of “stabilizing the object,” but that phrasing may misconstrue the metaphysical picture, as if the object enjoyed an independent existence, and the normative weight in the situation lay on the degree to which it was “stabilized,” whatever that would entail. Although that is false, the opposite is not true either; what happens is a much richer intermediate case of co-construction. As will be described in more detail in the next section, the existence of the object—including the object’s continued existence while out of view—arises in part from the very achievement of registrational stability. This stability can take many forms—ranging all the way from short-term perceptual stability to long-term evolutionary stabilities of the species.

4 Intentional acrobatics

How are these patterns of cross-cutting extension achieved? Once again, I will violate one of my own methodological injunctions, this time by breaking the question into two parts: achievements of the subject, and achievements of the object. Consider the former. From the point of view of the subject, given the inherently deictic or differential structure of the substrate, what is required is a panoply of explicit compensatory mechanisms (or rather, what is required is a range of behavior or mechanism that it is theoretically convenient to register as plural). The s-region must *compensate for its own contribution to the*

deictic relations in which it stands to the o-region in order to stabilize the other end. That is: it must *deconvolve the deixis*.

Think of the problem. Because they (are being assumed to) have arisen out of the underlying physical fields, the relations that have been talked about so far are differential or deictic, having to do with differences in the relative circumstances of s-region and o-region. Both sets of differences, in general, will constantly be changing. What is wanted, through this dynamic flux, is a stable way of relating to the o-region as a gathered and abstracted unity. The separation from the o-region provides the s-region with some buffer from the variation intrinsic to the o-region. What has not yet been eliminated, however, is the inevitable distraction of the s-region’s own (whole) contribution to its relational stand.

Various strategies are available for isolating parts of the s-region from the inner turmoil of other parts of the s-region, or from minor perturbations that are constantly impinging on the s-region as a whole (so that violent thoughts about your advisor do not undermine your ability to think of escaping to Rio, and so that when you are cold you can still remember the friend you met last summer in Amman). For this sort of internal stability *digitality* is an excellent choice. This is why digital computation, at least at the moment, is by far the most popular species. There is a great deal to say about digitality—about its nature, its use, and its achievement, as well as about its role in registration.⁹ Overall, though, Haugeland’s assessment is very much on target: that digitality is “a practical means to cope with the vagaries and vicissitudes, the noise and drift, of earthly existence.”¹⁰

But nothing about the internal structure or architecture of the s-region will enable it to insulate itself from itself as a whole, or from its overall circumstances. Thus think of the part of your registration of the house next door that has been retracted onto the s-region side of the separation, onto which responsibility for

⁹See TMD-V.

¹⁰Haugeland (1982), p. 217.

maintaining coordination has been retracted. And contrast this mental representation with a signpost, planted in the garden. Being unlikely to move, the signpost can simply say “the house next door,” and thereby preserve a stable reference. Like it or not, however, what is in your head will go with you on your next trip to Paris—and so cannot have such a purely deictic structure, at least not if it is to stay referring to the original house.¹¹ Qua *s*-region, if you have reason to remember that specific house, the house next door to where you live, you have to construct something whose directedness will stay fixed on a stable referential point, in spite of, or perhaps in explicit compensation for, your own movements.

That is not to imply that external representations are any more likely than internal representations to be deictic. Deictic signs are often useful—“you are here,” “next bus in 10 minutes,” “campground 200 yards ahead on the right”—but so too are less indexical representations (“4421 Hingston Avenue; 5:13 A.M. Wednesday, November 29, 1950”). That way lies external (derivative) representation, documents, and writing. Skill in using and creating signs involves appropriate combination. In particular, the problem with deictic signs is that the original difficulty recurs: as soon as you drive off and become separated from the sign, it can no longer be of immediate (i.e., effective) help to you in registering the house, campground, whatever. *So in turn you may have to register the signpost.* That sounds redundant (why not just register the house directly?), but in real situations very different costs are associated with what registers what, and in what ways. As a result, complex cascades of registration, more

and less deictic, are often useful, and sometimes essential. Thus you remember where the map is; the map “remembers” the directions to the lake. Or, more fully: you remember where the map is; the map “remembers” the directions to the road that runs by the lake; once there, effectively accessible (encounterable) road signs indicate the turnoff for the campground; at the campground additional (deictic) signs indicate exactly where to pitch your tent. And so on and so forth. The basic problem, however, is inescapable: (i) the circumstances of the *s*-region itself are vulnerable to change; and (ii) the relation between the *s*-region and *o*-region will also change. Given that the underlying fields are deictic, explicit work must be done to compensate for both so as to maintain a stable relation to the *o*-region.

How is this done? By analogy, as depicted in figure 7-5, imagine an acrobat leaping and jumping about on a somewhat darkened stage, putting their body through all kinds of fantastic gyrations, and yet throughout this crazed dance keeping a flashlight pointing absolutely reliably towards some fixed point—a point about four feet off the ground, say, towards the left front center of the stage. What the acrobat would need to do, through a complex series of hand and arm motions—handing the light back and forth from hand to hand, reaching it around behind themselves, and so forth as appropriate—would be to do the opposite with one part of their body (arm and hand, plus flashlight) of what the rest of their body was doing (leaping and dancing), in such a way that the two, when added up together, nullified each other, leaving the focal point of illumination unchanged.

Except of course the word ‘opposite’ is not right. If the acrobat were to leap up four feet, it would not be necessary for them to do the exact opposite—i.e., to drop their arm down, four feet, if that were even possible—in order to hold constant the full six-coordinate position and orientation of the flashlight. In

¹¹A mental analogue of “the house next door” might be retained *opaquely*, as it is said—i.e., without compensation—for example in the general maxim that if, while preparing dinner, you discover that you are missing a crucial ingredient, it often pays to check with “the folks living in the house next door” (a maxim that would lose its practical force if the reference were stabilized—i.e., to use as-yet-unavailable language, if, when learned, it were interpreted extensionally).

fact not a single one of the flashlight's six coordinates needs to remain fixed. It is not the *flashlight* that needs to be stabilized, after all. To freeze the position of the light outright might seem




FIGURE 7-5 INTENTIONAL ACROBAT

to be overkill, and would anyway be impossible, for example if the acrobat were to rush to the other side of the stage. Ironically, moreover, there is a sense in which keeping the light itself locked in position, as if it were epoxied to a particular point in space, would be underkill, *since there would then be no way to be sure where it was pointing*. For a fixed un-

moving flashlight, that is, no single point along its path of illumination is uniquely singled out; all you have is a long gradually dissipating path of light.

Fortunately for the acrobat (and for registration, by analogy) there is a better way: it is only necessary to rotate the wrist in just the appropriate manner. The normative requirement is that there be a fixed point at the “end” of (i.e., at some point along¹²) the line along which the flashlight is pointing—a point that remains stably located in exocentric 3-space. So the dance does not merely compensate for the acrobat’s movements; as a method for stabilizing the object, it is superior. The combined movements

¹²The fact that the stabilized patch is only “somewhere along” the beam of light corresponds to the fact that although we see the car, not just a two-dimensional facade, or even the pattern of incident light-waves, we do not (typically) see *through* the car and register the assembly line, or designer. But the analogy should not be taken to suggest that the phenomenon is simple: it opens up into the whole issue of the “transparency” of signs—the fact that when fluent registrars encounter external representations they typically register not the sign but what the sign represents. Semiotic transparency is almost magical; we just need to remember that part of the magic is that the transparency *stops* somewhere.

of dance and compensation together allow the s-region to triangulate in on the stable point of focus. All told, the focal point is much more stably and redundantly identifiable through the acrobat’s motion than it would have been had the acrobat stayed put. And it is that one point, and only that one point, that needs to be “held” constant.¹³

Something very like this is the function of the vestibular-ocular reflex. In synchrony with rotations of our head, and with displacements in our body, we adjust the angle of the orientation of our eyes in their sockets so that a constant (distal) point of visual focus is maintained. We can even do it when we are not ourselves responsible for the body’s motion. Thus imagine looking out the window of a train. It is possible to hold one’s head and eyes fixed, which turns the visual field into an incomprehensible blur. The natural thing to do, however, guided by visual feedback, is to rotate our eyes and head in such a way as to compensate for the motion of the train—with, again, a net result of maintaining a stable distal point of focus. It is this sedimentation of distal stability that begins to answer the question, raised in the introduction, of why we see trees, not electro-magnetic radiation.

The mechanisms underlying visual perception are just one example of this general strategy of a subject’s *computing its own inverse* (see the sidebar on page 240). When we reach out to grasp a coffee cup, we initially “think” that it is, say, 2 feet away.¹⁴

¹³The power of combined change and compensation is precursor to the common intuition that “multi-modal” perception or engagement is necessary for referential objectivity: a claim that the fact that we can touch, and smell, and use, and burn, and generally interact in all kinds of ways with a table, say, is constitutive of its being a *table* that we touch, smell, burn, use, etc.—rather than there simply being a table-surface that we touch, a table-smell that we smell, etc.

¹⁴Literally, we think no such thing. “Two feet away” is an extraordinarily high-level, sophisticated, and conceptual way of describing distance (see §1 of chapter 9). If distances to coffee cups are registered at all, they are surely registered non-conceptually. In the normal course of events, however, I

The analogies with vision and pointing suggested ways in which the point of focus could be dynamically maintained in the face of the endless ebb and flow of the underlying fields. If a relatively slow observer were to watch the acrobat dance around on stage—say, by half shutting their eyes, and watching the show in just a slight mist—all they would see would be the one single point of constant illumination; the movement and acrobatic compensation would all be washed out in a blur.

So it is with the focus of registration. The speed and complexity with which we compensate for changes we undergo in orientation and access to objects is just as much of a blur—partly because they happen far too quickly for us to be conscious of (as we are reminded by the ethnomethodologists' video tapes), and partly because, not being stable, they do not end up in focus. As a result, the only thing we end up “seeing” (registering) is the single stabilized object-as-entity.

The restless dynamics of the analogy are important, too. To stabilize appropriately—to *bring the world into presence*—very much requires being alive. It is like standing on a boat in heavy seas; you must constantly dance in order to stay upright and focused on the shore. In death, on the contrary, local stability (total quiescence) obliterates the active ability required in order to stabilize the distal world. Coordination ceases, pure deisis takes over, and the only thing that any longer has impact is the press of the immediate physical surround.

There is a crucial respect in which the analogy between registration and vision breaks down, however. Light beams and visual focus are *connected* phenomena. They are all inherently constrained to operate within the pre- and proscriptions of the physical plenum. Registration is different. Not only is its directness *able* to transcend those limits; it *must* transcend them. Go ahead, be expensive: transcending those limits is what registration is *for*.