

## **Extruding Intentionality from the Metaphysical Flux**

Josefa Toribio

*Journal of Experimental and Theoretical AI*, **11**, 1999: 501-518.

Department of Philosophy  
Washington University in St. Louis  
Campus Box 1073  
One Brookings Drive  
St. Louis, Mo. 63130, USA

e-mail: [pepa@twinearth.wustl.edu](mailto:pepa@twinearth.wustl.edu)  
Phone (314) 935 6140  
Fax (314) 935 7349

Brian Cantwell Smith  
*On the Origin of Objects*  
 A Bradford Book  
 The MIT Press  
 Cambridge, Massachusetts  
 London, England  
 ISBN 0-262-19363-9  
 Cloth

## **Abstract**

*On the Origin of Objects* is, at heart, an extended search for a non-circular and non-reductive characterization of two key notions: intentionality (the content or "aboutness" distinctive of mental states) and computation (the familiar but elusive tool of much cognitive scientific explanation). Only a non-circular and non-reductive account of these key notions can, Smith believes, provide a secure platform for a proper understanding of the mind. The project has both a negative and a positive aspect. Negatively, Smith rejects views that attempt to identify the key notions with lower-level physical properties, arguing instead for a more abstract and systemic understanding. This negative effort occupies Part I of the book (Analysis). In Part II (Construction), we encounter the positive side of Smith's proposal: an attempt to develop a non-reductive analysis of computation and meaning able to meet the (rather severe) requirements laid out in Part I.

One purpose of this critical review is to lay out this project in fairly simple terms. This is necessary since Smith's own treatment and prose sometimes obscures the flow of the argument. I suggest that, properly understood, Smith's proposal bears a clear affinity to ideas emerging from the dynamical systems movement within Cognitive Science, and that this tie-in can help put flesh on several of the more metaphorical characterizations in the book. My main criticism is that the book ultimately fails to provide an account able to meet Smith's own requirements for a truly non-reductive account of intentionality. This is especially the case regarding Smith's commitment to licensing a partition of the world based on no a priori assumptions whatsoever. The exercise is a valuable one, however, since it forces us to look harder at some foundational assumptions and at least hints at a new and refreshing perspective: one in which the key explanatory relations are grounded in facts about human practice and pragmatically established social norms.

## 0. Introduction.

Brian Cantwell Smith's book *On the Origin of Objects* (henceforth *OOO*) is a metaphysical tale. The scope is ambitious, the goal almost pre-Socratic, the aspiration: to extrude ontology from the metaphysical flux. The basic topic is the search for a non-circular and non-reductive explanation of intentionality. To this end Smith outlines a set of basic constraints that any such explanation should meet and attempts to show why current efforts to solve the problem are doomed to failure. The structure of the book is simple; the content less so. *OOO* is beautifully written but it is so full of framework twists, reversals, promises, metaphors and caveats that fleshing out the argument is a slow and delicate task.

The project has a negative side (Part I: Analysis) and a positive side (Part II: Construction). In Part I, Smith challenges the standard view according to which intentional properties will be revealed as real in virtue of their identity with, or supervenience on, some set of lower-level physical properties. It is here that the reader finds a radical shift in the way intentional properties are to be accounted for. It is not a question of finding a set of narrow or individualistic constituting properties. Intentional phenomena are instead regarded as context-dependent and institutionally created. This shift has consequences for many of the central notions invoked in the standard view (such as the notion of object, subject, physics, computation, etc). In particular, the notions of subject and object are especially transformed. It is one of the main desiderata of Smith's project that we need to *earn* the subject / object distinction, because otherwise our explanatory scheme would be seriously incomplete and any metaphysics built upon it rendered shallow or invalid. My view, as will become clear, is that Smith fails to meet his own (strong) requirement here. But it is a valiant effort nonetheless.

In Part II Smith develops what he considers to be an adequate explanation of intentionality. Here his own metaphysical view is revealed as one story (although not *the* only one) able to meet all the requirements. To what extent Smith's view marks a real departure from mainstream philosophical positions on this topic is, however, difficult to evaluate, as we are quickly immersed in a superficially new terrain in which many

central concepts have been recharacterized and re-defined. The *spirit* of the proposal, however, bears a clear resemblance to that of a Dynamical Systems perspective in Cognitive Science (more of which below). Moreover this dynamical approach, I will suggest, provides many of the tools needed to clarify some of Smith's more metaphorical characterizations. The position defended in the book (unfortunately) shares the Dynamical Systems approach's inefficiency in reconstructing higher-level, mostly linguistic notions.

Despite all this, there is something overwhelmingly right about the ideas Smith puts forward. If we want to understand how intentionality works, a good explanatory strategy is indeed to confront the complex interactive processes that bind subject and object into a whole, and to try to then account for the subject / object distinction itself. By contrast, the general model of explanation that informs most recent attempts to naturalize intentionality looks to be governed by two (often unremarked) physicalist assumptions: faith in a causal, reductive explanatory strategy, and faith in a (metaphysical) notion of individualistic constitution (again, see below). The multiple and complex interactive relations that characterize the cognizer's dealings with the world are often underrated by such models. Subjects and objects are treated as though they are pre-given entities, with well-defined boundaries that figure in causal explanations of behavior. By challenging this presupposition, Smith helps us to re-think our assumptions and hints at a new and refreshing perspective: one in which the key explanatory relations are grounded in holistically characterized, socio-pragmatic notions.

## **1. The Problem.**

The target problem, as I said, is intentionality. It is, to put it roughly, how some noises, or marks on paper, or computer programs, or neural events reach out to some object or property in the world. The need is for an account of how and why this property of 'being about something else' arises<sup>1</sup>. What we need, to use one of Smith's numerous examples is: 'an account of why, when we look out the window, we see a tree —i.e., have a (potentially) conscious experience of or about a tree, not of a two-

dimensional leafy and barked surface, let alone of a pattern of incident electromagnetic radiation ...' (OOO, p. 13).

Although one of the multiple ways of reading this book is to see it as a discussion about reference, Smith seems less worried about how particular words and representations come to refer than about the putative phenomenon of reference in general. It is the parsing of the world into states that can be referred to and states that can refer that he finds puzzling. His question is what makes the world parsable into states that refer and states that are referred to or, in other words, how it is that the mind/world distinction emerges. What is wanted is an account of intentionality that doesn't presuppose either the subject or the object and that can therefore provide an explanation of how the world *is* as it is taken to be by a subject. This is a foundationalist project. Yet, the challenge is to build this metaphysical foundation without relying on any a priori assumptions.

The problem, thus characterized, has a Kantian resonance also to be found in work such as McDowell (1994) and Adrian Cussins (1992). Seen through this Kantian lens, Smith's message may be summarized as the thesis that representational experiences and practices are actively acquired, and that the faculty by which one exercises one's conceptual and non-conceptual capacities in thinking and judging is *dynamically constrained* both by spontaneity (the subject's cognitive capacities) and receptivity (the realm where empirical content is to be found). As he tells us at the very beginning of the book, his 'metaphysics —a philosophy of presence— ... aims to steer a path between the Scylla of naive realism and the Charybdis of pure constructivism' (OOO, p. 3). In order to steer that path, we are invited to join a long trip (full of detours) that begins in Part I with the analysis of computation, its foundation and its role in Cognitive Science.

## **2. Computation.**

Computers are intentional artifacts. Theories of computation are chosen as *an* example of how various issues involved in the explanation of intentionality are standardly approached, and of how they might

alternatively be treated. A Theory of Computation, it is argued, has to meet two essential criteria: an *empirical* criterion, or the requirement that the theory do justice to computational *practices*, and a *conceptual* criterion according to which the theory should be a foundation for the Computational Theory of Mind. In the justification of these two criteria we can already see two of the main components of Smith's view. The first component is a very strong anti-reductionism. The other is a dynamical approach towards intentionality.

In order to meet the empirical criterion, the theory has to illuminate and explain the differences between intuitively different forms of representation, program and machine. But the notion of a representation, program or machine as being of such-and-such a type is, Smith argues, not explicable solely in term of properties internal to the computational system. What it is to be this particular program or this particular representation can only be individuated by appealing to properties external to the system itself. A theory of computation that takes a reductionist approach towards such complex intentional artifacts will not be able to meet the empirical constraint.

Smith believes the roots of this reductionist mistake lie in what he calls the 'binary model of semantics'. The binary model takes computation to involve only two types of entities: the program itself, and the domain or subject matter. But computation, it is argued, involves *at least* an extra factor: the process via which the program is being executed. The conflation of the notions of program and process leads to what Smith sees as the mistaken view that computation is purely syntactic; it also encourages a static view of the world and the entities that populate it, and has important consequences in Cognitive Science and Philosophy of Mind. According to Smith, an analysis of computation based on a multifactor model of semantics (instead of a binary one) would better fit our computational practices and would invite a more dynamic approach both to what it is to be an object and to intentionality itself. It would shift attention away from '... what it is to *be an object* ... [and] instead focus on what it is to *behave or act or participate or be treated as an object*' (OOO, p. 36). Similarly, 'intentionality will be reconstructed not so much in terms of a static notion of meaning or significance, but instead in active terms of *being meaningful or being significant*' (*Ibid.*).

The conceptual criterion, even though it concerns not so much the contents of the theory but its form or its *use*, moves things in the same anti-reductionist direction. That a theory of computation ought to provide a foundation for the Computational Theory of Mind means that in order to understand how the various practical representational practices collected under the label 'computation' (Smith's *computation in the wild*) belong to the natural world, we should not look *downwards* in the direction of the internal, constitutive properties of computers and programs. Since computers are already intentional devices, it wouldn't be methodologically or epistemologically appropriate to rely on computation as a way of explaining intentionality. If intentionality itself lies at the core of computational practices, there is really no such thing as an independent theory of computation (cf. *OOO*, pp. 73-75) —no abstract theory about the programs and representational powers of these socially constructed intentional artifacts. Providing a foundation as a result involves looking away from the syntactic properties of computational devices and towards the institutional frameworks that make them count as intentional machines. Inevitably, current theories of computation fail to meet these highly demanding criteria<sup>2</sup>.

The traditional use of computational ideas as a means of explaining intentionality thus involves committing what Smith calls an *inscription error*: 'a tendency for a theorist or observer, first to write or project or impose or inscribe a set of ontological assumptions onto a computational system ... and then, second, to read those assumptions or their consequences back off the system, *as if that constituted an independent empirical discovery or theoretical result*' (*OOO*, p. 50). The notion of inscription error is not entirely new, and has a familiar correlate in AI where it surfaces as the handcoding problem. A good example of this kind of error is the AM program, a program designed to prove mathematical theorems. The program was extremely successful but, as Lenat, the programmer, later conceded, the main reason for the success lay in the very representation notation used to write the program, i.e., in the amount of mathematical knowledge implicit in LISP (see Lenat, 1983 and Ritchie and Hanna, 1984).

Smith's treatment of computation ultimately implies a very radical position, completely different to mainstream foundational projects in

Cognitive Science. Such projects inherit important reductive assumptions from the explanatory model of physics. Cantwell Smith's claim, however, is that intentionality is not reducible to computation or physics, and that a theory of computation itself needs a metaphysical underpinning in order to meet both the empirical and the conceptual criteria (cf. *OOO*, pp. 69-76). The position defended here is quite close to that of another radical anti-reductionist, Robert Brandom (1994). Both Smith and Brandom propose a different kind of explanatory scheme for intentionality, in which the *explanans* includes higher level features such as skilled know-how and social frames of action. The rupture of the higher-to-lower explanatory strategy calls into question one of the core metaphysical notions in the framework of scientific explanation, the metaphysical notion of *constitution*. Although Smith doesn't frame his views in terms that involve a criticism of the notion of constitution, it is not difficult to see how this might work.

The notion of constitution I have in mind is the one at work in the thesis that what makes a thing into a thing of a certain kind is the relation between that thing and the stuff it is made of. Although there are other, more contextually oriented, ways of accounting for the exact nature of a particular thing (see below), *constitution*, understood in this narrow and *individualistic* sense, seems to be at the root of the model of explanation that informs most contemporary naturalistic models of intentionality. This notion of constitution is also answerable for the demand of cashing out the connections between the internal properties of an entity and its surroundings in *causal* terms. Roughly, this is how these two ideas fit together. A particular piece of metal is constituted as e.g. a piece of gold if it is made of the right stuff, if its composition has the atomic number 79. What it is for a piece of metal to be gold is then fully *explained* by appealing to the properties of its internal structure. But it is also its internal structure that is important in accounting for whatever causal powers that piece of gold might have<sup>3</sup>.

Most naturalistic approaches to intentionality aim to discover some set of lower level, non-semantic, properties in terms of which semantic properties are metaphysically *constituted*, e.g. those properties that make a given representation into the representation it is and that will explain why it is that representation (and hence why it has the semantic

properties it has) rather than some other. The search is for those non-semantic properties that play a (metaphysical) *constitutive* role in the instantiation of semantic properties. Underlying the model of explanation at work in contemporary naturalistic proposals there is thus not just the general physicalist bias mentioned earlier, but also the additional idea that what makes a thing into a thing of a certain kind is the relation between that thing and the stuff of which it is made.

But there is also a more contextually oriented alternative to this way of accounting for the exact nature of a particular thing. John Haugeland and Tim van Gelder have recently explored this possibility (Cf. Haugeland 1993 and van Gelder, 1993), referring to it as the 'holistic' alternative:

Roughly, holists see the fundamental nature of things as depending only upon some larger whole to which they belong ... metaphysical holists see constitution as solely a matter of context; an individual entity is what it is in virtue of some larger structure or pattern into which it fits ... If an entity is constituted as an A holistically, then to understand that entity as an A you have to understand the relevant larger whole and how As fit into it. To render an A intelligible as an A is to articulate its place in the larger whole.

(van Gelder 1993, p. 67)

To avoid confusion between the individualistic and the holistic notions of constitution, I propose to label the latter notion 'institution'. Institution is thus a context-oriented way of understanding what makes a particular entity into the entity it is. To say of something that it is instituted by such-and-such properties is to claim that its fundamental nature is best accounted for in terms of the properties of some specific context in which the entity is manifest, and thus that to understand the entity *as* an entity of that particular kind involves an appreciation of the properties of that context. John Haugeland gives the familiar example of the ontology of chess pieces (Haugeland 1993, 4). Suppose we ask what makes something a queen in chess? Obviously queens are made of some material or another but this individualistic notion of constitution will not help us understand the fundamental nature of a queen. What makes

something a queen can only be understood by focusing on the role that the piece plays in the larger game of chess and hence by focusing on contextual properties such as the rules of chess<sup>4</sup>.

Constitution and institution are not, of course, exclusive ontological relations. In most cases, that something is an entity of a particular kind results partially from the stuff it is made of and partially from the (environmental, social, legal, political, etc.) context in which the entity is located. However, there are countless entities for which the only way of getting at what they really *are*, and the only way of explaining why they behave how they do, is to look at the contexts that *institute* them, not the stuff they are made of. This is obviously true in the case of conventions and social institutions. Their nature, character and function are completely dependent on the socio-political context and cultural practices of a country or a community. Even though these are obviously instantiated by organizations of matter of some kind or another, the stuff they are made of doesn't make them the entities they are, and it certainly doesn't help to explain why they behave in one way rather than another.

One way of understanding Smith's notion of computation and, as we shall see, of intentionality itself, is through this institutional lens. That doesn't mean that the constitutive properties of computational and intentional phenomena (in the narrow and individualistic sense of constitution just explained) don't play any role. It rather means that what it is to be a program or a particular representation is always something partly context-dependent. It is the result of a delicate balance between what is out there —constitutively speaking— and the structures within which it is institutionally framed. Keeping this idea in mind will help us to better understand the middle path between Realism and Constructivism that Smith is about to steer.

### **3. Irreduction.**

Smith quickly introduces a general methodological requirement that plays an essential role both in the computational context analyzed above and in the metaphysical story to be developed later. This methodological criterion is dubbed a *Principle of Irreduction* and 'it mandates that no

theoretical assumption ... be given a priori pride of place' (OOO, p. 77). The ideal situation would be one in which the foundations of a computational, intentional or metaphysical theory are not grounded in any assumption  $\mu$  for any  $\mu$ . If this turns out to be unfeasible, then for any category  $\mu$  in which our theory is grounded, one should be prepared to say 'where one bought it; how much one paid; how one got it from there to here' (cf. OOO, p. 78).

One deep assumption underlying theories of computation, of intentionality, and metaphysics is the notion of a subject / object split. The role of this particular  $\mu$  begins to be unraveled in Chapter 3, where we encounter a new set of requirements, this time directly related to the metaphysical enterprise and not just to the theory of computation. Think of e.g. traditional Realism (upper-case 'R'), a position associated with natural sciences and roughly characterized by the defense of two theses: that the world is external to the subject and that it exists independently of the subject's cognitive abilities. By taking the world to be *external* to and *independent* of the subject, traditional Realists are committing (according to Smith) an inscription error of the kind mentioned earlier. They are grounding their metaphysics in the (as-yet-unjustified)  $\mu$  that the subject and the world are two independent entities, and also the (as-yet-unjustified)  $\mu$  that there is a univocal perspective on reality (cf. OOO, p. 89).

Against such a position, Smith argues for a unitary, commonsense realist account (lower-case 'r') that incorporates the intuition that there is a world out there, but that denies the idea that this world is *external* to or *independent* of the subject, i.e., it denies the idea that there are, as it were, two different entities: a world—full of clear-cut, univocal objects and properties—and a subject who takes no part in the makings of that world. By denying that duplex, independent structure of world and subject, by claiming that 'the world, our world, is *one*' (OOO, p. 103), we are supposed to leave theoretical space for an explanation of how both sides of the same metaphysical coin simultaneously emerge. Or, at the very least, we leave space that can be filled by a justification of our particular assumptions (our particular  $\mu$ s) concerning the existence and nature of both subjects and objects.

The need to *earn* the subject / object distinction is just one of several constraints offered as desiderata for a successful metaphysics. One might think that Smith's insistence on questioning this distinction is a clear sign that he is conflating the elaboration of a theory of intentionality with the project of constructing a basic metaphysics<sup>5</sup>. However, it is not so much a matter of conflating two different enterprises as it is of establishing the dependence of the former on the latter. Part of what makes Smith's proposal a genuine alternative is precisely this reversal in the order of priorities: in order to get a theory of intentionality, we have first to develop a metaphysical theory. Or, at least, we have to re-think the one we have in such a way that some important conditions are met. I have already mentioned one such condition: to give an account of (not take for granted) the subject / object distinction. The others, some of which we are already familiar with, may be summarized as follows: (i) a successful metaphysics has to be anti-reductionist: ' every distinction must be wrested from the same metaphysical flux ... everything, including physical objects, human societies, truth, beauty, reason and mathematics .. are to be drunk from this one same cup' (OOO, p. 1016); (ii) a successful metaphysics has to acknowledge the perspectival character of knowledge, i.e., it must not be Objective (upper-case 'O') but objective (lower-case 'o'); (iii) a successful metaphysics should find a way of doing justice to pluralism without succumbing to rampant relativist positions; (iv) a successful metaphysics has to be open to the fact that all intentional processes are located, embodied, perspectival and pluralist in a dynamic and potentially unstable sense (cf. OOO, pp. 109-110); (v) a successful metaphysical story has to account for the fact that what makes something a particular object is the result both of the cognizer's work to establish reference despite the troublesome flux and the flux fighting back against the subject's cognitive abilities. And finally, (vi) a successful metaphysics has to provide an explanation of this lower-case realism. Again, in the spirit of the principle of irreduction, the theorist should be able to say 'how, why, and what it is for an intentional agent to have a world-directed commitment' (OOO, pp. 110-111).

One question that immediately arises is just how different the traditional and the commonsense versions of realism really are. At first blush, it seems perfectly possible to be a (more or less traditional) Realist

(upper-case 'R') without flouting —at least most of— the requirements Smith is advocating. Take, the first Realist thesis, the thesis that the world is *external* to the subject. The key issue seems to be how much *structure* we allow that world out there to have. It is typical of a traditional Realist position to take the world to be a certain, *fixed* way. The *duty* of the intelligent cognizer is then to find out about that ready-made world (perhaps through the development of scientific theories). The main difference between this and Smith's position is that for Smith the world is certainly out there but it is not ready-made in any way that renders the cognizer's or the scientists' views uniquely correct. The independent world is a flexible and unstable metaphysical flux: it is not a fully structured realm, already sliced up into objects and properties.

Similar comments apply to the second main thesis of Realism: that the world exists independently of the experiencing subject. This, too, could be accommodated within Smith's view if that *independence* is located at the appropriate level of *unstructured* reality. As we shall see, one of Smith's main contentions is that the metaphysical flux itself (that undefined, unstable putty) only becomes populated by objects and properties —a step up in organizational structure— once one part of that flux (the part that will eventually become the subject) achieves a certain kind of disconnection from the rest. It is only then that other chunks of the metaphysical flux can be seen as objects, as things out there. Smith talks, in this vein, of an s-region and an o-region of the flux: the precursors of the subject and object respectively —new categories designed to avoid the inscription error of talking too soon of subject and object. They are like the proto-subject and the proto-object of standard metaphysics, still too embedded in some non-individuated pattern of structure to count as subject and object proper. The key issue for the realism / Realism distinction thus seems to concern the amount of structure with which we are ready to credit the metaphysical flux. The more structure it has, the closer you get to a traditional Realist view; the less structure the closer you get to the so-called 'symmetrical realism' defended by Smith. The problem with traditional realism, Smith thinks, is that it is unable to account for the emergence of the subject / object division itself. Yet, Smith's own position, for all his intended monism, still keeps falling back into the same old dichotomy, now expressed via

the notions of *s-region* and *o-region*. We are told that this is just a trick, not to be taken seriously; a prop to help us until all the details of the final picture are in place. But the reader is kept waiting rather too long (until the very last chapter of the book). And when she gets there it is not at all clear whether an innocent account of subject / object emergence has indeed been provided. It would be unfair, however, to pursue such radical criticisms before at least sketching the rest of the positive argument. To do so, we must first acquire a few more tools.

#### 4. Physics

Smith introduces a dizzying number of distinctions and re-definitions which I will swiftly rehearse before showing how they fit into his general argument. The main ones to keep in mind are:

(i) The distinction between particularity and individuality. Although Smith acknowledges that the ideas of particularity and individuality overlap in our everyday conception of medium-size physical objects, they, he argues, respond to different metaphysical motivations. On the one hand, 'particular' means 'something like "occurrent": something that is located or that happens, something that is embodied' (OOO, p. 117), as when we talk about that (particular) bottle of wine we had last night or yesterday's cup final. The notion of individuality, on the other hand, is meant to capture 'whatever it is about an entity that supports the notion of individuation criteria ... *Individuality is what allows one to say of one object that it is one; or of two, that they are two*' (OOO, p. 119) (Smith counts not just things like chairs, tables, and people as individuals, but also types, properties and relations). Individuality, unlike particularity, is linked to the idea of discreteness, the idea of integrity within some boundaries.

(ii) The distinction between property and feature: 'the term 'property' is reserved, at least informally, for qualities or ways of being or types that do require objects for their exemplification; the term 'feature', for things that do not' (OOO, p. 125), as e.g. in 'it is raining'.

(iii) The distinction between individuality, identity (lower-case 'i') and Identity (upper-case 'I'). The (traditional) notion of Identity (upper-case 'I', also called identity to avoid confusion) is extensional and it is a clear-

cut notion; either two things are Identical or they are not. The notion of identity (lower-case 'i') that Smith introduces is not related to this property. Instead 'it has to do with what makes something be the thing that it is' (OOO, p. 132). It is an intensional notion. Diffusions (such as fog), abstractions (such as melancholy), and collectives (such as groups of people) can have identities without necessarily being individuals, or supporting individuation criteria.

(iv) The distinction between physics (lower-case 'p') and Physics (upper-case 'P'). The upper-case notion of Physics is characterized as a value-free, Objective, True, Rational empirical inquiry, while the lower-case version is 'physics-as-it-actually-practiced, complete with political battles, elaborate and often messy laboratory procedures, complex instruments and documents, power struggles, aesthetic judgments and other paraphernalia of a socially engrained and highly professionalized discipline' (OOO, p. 139)

(v) The distinction between material and physical. 'Material' is used to refer to 'our ordinary commonsense notion of physicality' (OOO, p. 141), as when we refer to e.g. table and chairs. 'Physical' is whatever belongs to the lower-case realm of 'physics'.

(vi) The distinction between first, second and third realms. The first realm is the realm of complete and total particularity, i.e., forces, fields, spatio-temporal positions. It is constituted by our common understanding of the material world. The second realm includes people, their experimental apparatus, their books, papers, conferences. Tracing this definition back, we can see that this realm is (lower-case) physical since it includes all the social, aesthetic and common practices of the study of physics. The third realm is *abstract and universal*. It contains objects such as types, numbers, sets and physical laws. Physical laws don't apply to themselves though. They apply to events in the first realm.

In order to reconstruct Smith's position, we not only have to get a grip on these distinctions, we have also to be aware of a disconcerting maneuver that will feature throughout the whole book: it has the flavor of a *reductio* and it involves making a particular position the center of the discussion, as though it were the *right* view, only to later show that the assumptions underlying that view are mistaken and that the position itself should be rejected. This is how the traditional view that favors

Physics as a source of ontological grounding for a theory of intentionality is handled: 'The question on the table is not whether physicalism is right ... The aim is ... to see *what contribution it can make towards metaphysical understanding* ... The aim is to lean on physics—or at least on some of its ontological commitments—only in order to make some points in terms of which ... to describe the metaphysical view to which in fact I am ultimately committed' (OOO, pp. 152-153).

Why is physicalism sometimes thus favored? Smith rehearses some reasons for its appeal. One is an interest in naturalism, taken as the methodological correlate of Realism: the idea, roughly, that if at the end of the day everything that exists can be completely characterized at the physical level, then a special interest in physicalism is justified. Another is a matter of ontological hygiene: physicalism is said to provide a general guarantee against *bizarre* ways of partitioning reality. A final incentive spins around the idea of implementation, the idea that building something is a good way to understand what it *is* to be that something.

With all these (apparently) good reasons to embrace physicalism, shouldn't we just go for it? Smith's argument attempts to show that Physics (upper-case 'P') *can* sustain a basic notion of particularity—of located and embodied, continuous and deictic fields—but not a notion of individuality. Physics does not have the resources, Smith claims, to provide the individuation criteria needed to turn this metaphysical flux of particularity into discrete individuals: 'neither the third realm itself, nor the relation it bears to the first realm, is part of the subject matter of physics itself ... *nothing non-particular is itself governed by physical law*' (OOO, p. 156). It cannot provide the interpretative parameters to turn those deictic fields into e.g. a table, the property 'red' or the relation 'to be the mother of', all of which are wholly discrete individuals. The claim 'there are no physical objects' (OOO, p. 178) doesn't sound as radical as it otherwise might when seen in this context. It doesn't mean, as one might think, that there is nothing out there to support our common-sense notion of an object. Instead the claim is that the items Smith has characterized as individuals (a category that includes our common-sense notion of object) are not ontologically basic: 'instead they are part of the epistemic structure of physics-the-discipline. Not only that; they are absolutely crucial, epistemically. They are necessary for us, as epistemic

agents, to *calculate* anything, to *figure anything* out, to allow physics to be *useful*. In sum: they are *constructs of the second-realm*' (OOO, p. 178).

The methodological correlate of this metaphysical conclusion is what Smith calls the Criterion of Ultimate Concreteness, according to which '*no naturalistically palatable theory of intentionality —of mind, computation, semantics, ontology, objectivity— can presume the identity or existence of any individual object whatsoever*' (OOO, p. 184). This means, among other things, that no naturalistically palatable theory of intentionality can presume the subject/object distinction. Otherwise the theory will not be able to account for how intentional properties arise. Overall, the main conclusion regarding the issue of physicalism is that Physics cannot provide any foundational insight as the base for a theory of intentionality. Since objects are ultimately '*inexorably cultural, biological, political, psychological, social, evolutionary, historical, economic, and so on*' (OOO, p. 188), '*physicists will have to look to a theory of intentionality for an account of the notion of an individual, not the other way around*' (OOO, p. 180).

For a reader with anti-reductionist inclinations, these proposals are clearly tempting. The account has the additional virtue of bringing to light some of the problems that arise when we invoke the micro-properties onto which our representations are supposed to supervene as the properties that must account for the causal efficacy of those representations. It is not that there is nothing out there that makes e.g. a chair a chair (a chair is something material made of some chunk of physical stuff). But what narrowly and individualistically *constitutes* something as a chair —as a specific chunk of stuff about which Physics can give us details— is a partitioning of the metaphysical flux that is not provided by Physics itself<sup>6</sup>.

Such a line of thought is a welcome change from the dominant reductionist picture of cognition. However, given the re-definitions, re-partitions and re-characterizations involved in this analysis, the reader (or, at least, this reader) cannot help but find herself wondering how much of this makes real contact with those current positions (with their own definitions, partitions and characterizations). In other words, once the *meanings* imposed upon the key notions in the debate have been altered so much with respect to our standard notions of e.g. 'physics', 'property',

'individual', 'particular' and so forth, how are we supposed to evaluate the real distance between Smith's proposal and the traditional view? The usual way to criticize a view is to keep constant many features of the normal backdrop and to focus attention on just those places that will make a difference to the final outcome (on pain of begging the question or committing an equivocation fallacy). Here, by contrast, we are invited to alter both the background and the key features of the resultant position. Under these conditions, it is difficult to say exactly how the two views relate.

## 5. Registration.

We turn now to the Second Part of the book, in which Smith aims to meet all the requirements and avoid all the problems displayed in Part One. The key positive notion here is the idea of *registration*, and the goal is to understand how it arises out of the metaphysical flux. The notion of registration is introduced to avoid the inscription error of taking for granted the split between a conception and that which it is a conception of. In this sense, registration is quite close to perception. Unlike perception however, registration includes all aspects of intentionality, not only those relating to what is present for the subject at a given moment: '[b]y "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 ...*' (OOO, p. 191).

Chapters 6, 7 and 8 present the essential features of registration. Firstly registration involves an *asymmetrical* division of responsibility. The s-region takes the responsibility while the o-region works against it courtesy of the instability of the metaphysical flux. Secondly, to register something *as* something is an intensional act. Registration is an intentional act and carries both an 'implication of alignment with the external situation ... and a sense of involvement or engagement with the world' (OOO, p. 196). Finally, the subject of registration is not restricted to the biological subject. Entire language communities and bodies of practice can also share the responsibility for taking of an object to be such and

such. They can do this either directly or via the cultural baggage accrued by the individual subjects themselves<sup>7</sup>.

We said earlier that registration is akin to perception. The main difference, though, is that whereas perception can only occur in the presence of what is perceived, registration actually requires separation and distance between (what Smith calls) the *s-region* and the *o-region*. The simple tracking of one region by the other—a connection without *distance*—doesn't count as registration. In simple tracking the *s-region* and the *o-region* are insufficiently distinct and insufficiently separate (cf. *OOO*, p. 219). What we need is to break this causally-driven servo-loop; to break effective coupling and yet maintain a certain coordination. A nice example is that of the super-sunflower. A super-sunflower is a (fiction) species of sunflower whose members have the following property: they not only track the sun (in the way ordinary sunflowers do) but are also able to continue to do so even when the sun disappears. In that way, when the sun reappears again they are better oriented toward the light and thereby enjoy a selective advantage over ordinary sunflowers. The example also helps clarify the point about asymmetrical responsibility: the *s-region*, the super-sunflower, takes the extra responsibility of keeping track of the sun even when it disappears. The sun just moves around in the way it does. It is up to the super-sunflower to 'keep an eye' on it.

The requirement of separation between the *s-region* and the *o-region* yields a set of essential properties of registration. It is because *s-* and *o-*regions are separated that the extra task of maintaining coordination arises, and the *s-region* must assume extra responsibility. It is also this break in effective coupling that ushers in the possibility for error, because 'there is no reason to guarantee ... that the retracted mechanism will be able to mimic perfectly the distal regularity with which it is "striving" to maintain alignment' (*OOO*, p. 223). Smith calls this view a *philosophy of presence* but it should be clear that the presence of the *o-region* essentially involves its absence. It is this occasional absence of the *o-region* that enables the *s-region* to impose a certain kind of *stabilization* in which the metaphysical flux of the *o-region* is treated as a unity and is taken as an *object*. Given the dynamical relations between *s-* and *o-*regions, the *s-region* has to compensate for the various changes in order to *bring the world to presence*, in order to fix some aspect of the flux

into what we think of as an object. In doing that in a sort of acrobatic fashion —full of twists, leaps and unstable movements— the s-region makes itself a subject: ' ... the first emergence of the "subject" as an individual would be as a long-term integral or aggregate of *that which it must compensate for, in order to stabilize the rest of the world*' (OOO, p. 240). The driving vision here is of an active process of stabilization through which 'reference is *achieved*' (OOO, p. 263). This achievement, however, is never completely stable and therefore the idea of a completely context-independent representation emerges as a kind of theoretical fiction.

The notions of s-region and o-region are thus designed to avoid an inscription error regarding the distinction between subject and object. Once again Smith hopes to levitate himself above any ontological commitments in order to provide a non-grounded foundational metaphysics. However, it is difficult to see how the notions of s- and o-region can help him to achieve this goal. We may grant that some of the properties associated with these notions are indeed different to the ones generally associated with 'subject' and 'object'. However, in terms of the metaphysical picture, the notions still play the role of an *ontological commitment*. One difference, however, is that Smith aims to obey the rules of his Principle of Irreduction, i.e., he tries to tell us where he bought those assumptions, how much he paid for them and how he got them from one place to another. So, even if Smith's metaphysical ambition is somewhat overblown, some key elements of his epistemology may well survive.

## **6. Conceptual Registration.**

Smith aims, we saw, to cut a path between Realism and Constructivism. This metaphysical ambition also has a correlate within the theory itself. Chapter 9 explores and characterizes this as the idea of a *middle distance*. The reader is already familiar with the idea behind that expression, the idea of partial interdependence and partial co-constitution between subject and object. This time, however, the aim is to apply that notion in a more sophisticated sphere: the logical structure of property exemplification. What is at issue is how to accommodate the notions of

object and property within the compositional structure of cognitive representation. Unfortunately, given Smith's position, it is difficult to see how objects and properties can be *stabilized* sufficiently to enter into the formal processes of inference that constitute that compositional structure. In moving from primitive sense-based concepts to the concepts involved in linguistic behavior, we encounter typical problems such as those of context sensitivity and collateral information. To solve these problems we may need to posit a symbolic structure over which to define direct operations in the form of context-free rules. Smith's account doesn't seem to have the resources to deal with this problem. To give just one concrete example, one way of approaching the problem of designing connectionist models in which the behavior to be modeled involves representations of higher order than the ones based on observable properties is precisely to reduce the *level of distribution* of the representations. This can be done by building the symbolic structures directly into the hardware (the architecture) of the network or by several other methods. One way or another, the requirement seems to be for some sort of hybrid models in which the lowest level procedures are broadly connectionist but the higher level redescriptions are broadly symbolic (Cf. Hendler, 1989).

What Smith's proposal is in need of, then, is some way of grouping the partitions that give rise to primitive objects and properties into more coarse-grained constituents onto which to project proper semantic distinctions. To be able to do that, we need some previous semantic *criteria* that will let us discriminate between which properties belong to which concepts and which do not. Such semantic criteria are not easily recovered using Smith's own theoretical resources. For, given his perspectival and dynamical view, it doesn't seem possible to posit a principled distinction that lets us discriminate without circularity between those inferences that will count as part of the meaning of a concept and those that won't. What is lacking from Smith's philosophical perspective is thus a proper account of the relations between constituent properties and concepts. What he offers instead is just a development of the earlier ideas pursued in metaphorical terms. The moral of his story, regarding *conceptual registration* is simple, but disappointing: '*we individuate properties and relations in the same way as we individuate ordinary (particular) individuals*

—through analogous processes of separation, stabilization, long-distance coordination across breaks in effective coupling, and the like' (OOO, p. 317). There is also, however, an appeal to the social structures that sustain the phenomena of participation and commitment involved in inferential processes. So perhaps one way of understanding the processes of inference that characterize higher levels of cognition is to take these processes to be, à la Brandom, '*instituted* by social-practical activity' (Brandom 1994, xiii). But while Brandom turns his whole project into an explanation of the explicit normativity of inference in terms of such implicit normative practices, Smith deals with the issue only in rather vague and metaphorical terms.

## 7. The Metaphysical Picture.

It is only at the very end of the book that an overall vision of Smith's project comes to the surface, combining all the twists, distinctions and theses defended so far. Yet, in a way there are no surprises. The metaphysical story is one that we know already. It is the physical story (lower-case 'p') applied to the whole world (cf. OOO, p. 323). It is grounded in what he calls *immanent induction*: 'a way of being grounded that is ultimately unregistered, inexorably participatory (partially *connected*), and to a degree ineliminably first-person' (OOO, p. 321). The claim is not that this is *the* only plausible metaphysical story. It is that it is one plausible metaphysical story but, arguably, one that meets the basic criteria of plausibility described throughout the book. Accordingly, the world is not a clear-cut kingdom of categories. It is rather a kingdom of indefiniteness at all levels. The world is depicted

as one of cosmic and ultimately ineffable particularity: a critically rich and all-enveloping deictic flux. Neither formally rigid nor nihilistically sloppy, the flux sustains complex processes of *registration*: a form of interaction, subsuming both representation and ontology, in which "s-regions" or subjects stabilize patches of the flux, in part through processes of intervention, adjusting and building them and beating them into shape, and also through

patters of disconnection and long-distance coordination, necessary in order to take the patch to be an object, or more generally, to be something *in and of the world* .

(OOO, p. 247)

The aim is not to deny that there exist clear-cut categories. It is rather to claim that in order to create those categories, we have to fight and wrestle with what is, in principle, just a pure metaphysical flux. Smith's *successor* metaphysics is thus meant to occupy a middle ground from which it is possible to (i) do justice to what is right about constructivism and realism; (ii) sustain a notion of pluralism that avoids the risks of radical relativism and incommensurability problems; (iii) be irreduccionist (iv) re-think some of the central notions of modern philosophy (notions such as object, truth, physical, etc.), and (v) be foundational without being grounded in any special category (cf. OOO, pp. 345-346).

What happens then to the popular metaphor of 'nature's joints'? It is here where Smith's position seems to be most vulnerable. He certainly does not deny the existence of 'joints'. Smith is not advocating that how the world turns out to be is arbitrary or subject-dependent. Particular things are supposed to emerge from the interaction between the s- and o-regions. Yet, very little is said about how this interaction takes place, or evolves as to result in a particular structured world. Rather than a settled metaphysical picture, what we get is a recipe indicating the right methodological principles we should keep in mind when constructing our metaphysics. In accounting for how we should go about this, Smith focuses on a reductionist assumption; it is this reductionism he wants to avoid. The cure for reductionism seems to be to whole-heartedly embrace a holistic alternative strategy, one which appeals to the larger context of the conventions and constraining practices of a society of intelligent beings. But even the appeal to this larger context seems to require precisely the kind of a priori assumptions (about entities and boundaries) that Smith is explicitly disallowing. Thus, Smith looks to be violating his own principle of irreduction. This is hardly surprising since a project governed by such a principle is surely likely to become, if not self-refuting,

at least highly unstable. We will constantly have to revise the categorical partitions generated by our metaphysical theory since the central principle of our metaphysics is not to settle on any specific metaphysical principle.

This instability, this on-going revision of our own categorization as the result of a constant interaction between us and the world, makes me think of Smith's position as —to a certain extent— similar to a Dynamical Systems perspective in Cognitive Science. Of course there are clear and important dissimilarities. For example, the prominence assigned to the notion of representation (the general idea, within the Dynamical Systems approach is that internal representations are an unnecessary and unproductive posit for the analysis of real-time, environmentally-embedded behavior) and even the nature of the inquiry (metaphysical in Smith's case and mostly empirical for the Dynamical Systems' theorist). But there are some surprising similarities which we can use as a way of clarifying some of the more metaphorical claims in *OOO* and of making contact with further themes of mainstream Cognitive Science. For instance, both parties agree that cognitive/intentional systems are often characterized in terms of continuous circular exchanges between the physical / biological features of an organism and those of the environment in which the organism is embedded and functioning. In both cases, where content-involving descriptions are suitable for the explanation of such a cognizer's behavior, they involve a notion of content that is characterized by the abilities of the systems to interact in specific ways with the world in which they are situated. The relevant systemic properties are thus not individualistic properties in the classical, standard sense, i.e., they are not intrinsic, context independent properties of the system whose behavior we trying to explain. Instead, the properties that count for the characterization of intentional content are, in both approaches, external, context- and environment-dependent properties of the system.

As a result, the notion of internal content *carriers* or vehicles has vanished from both the Dynamical Systems and from Smith's approach. We know already what we find instead in Smith: complex processes of stabilization of an original and unstable metaphysical flux. What we find in the Dynamical Systems case is a general formalism based on the idea of an abstract geometric state space whose dimensionality depends on the number of relevant system parameters. The system's behavior is then

explained in terms of location and motion within that space. Constructs are defined that capture properties of certain points or regions in the space. The mathematics typically specifies a dynamical law which determines how the values of a set of state variables evolve through time (see e.g. Abraham and Shaw 1992 for more details).

This picture bears a clear resemblance to the dynamical and perspectival view of metaphysics defended by Smith. Nothing in either picture preserves the idea of a particular computational state as the vehicle of the content that explains a given behavior of the system. Without this idea of a 'vehicle', we lose (in both scenarios) the idea of content as constituted by intrinsic properties of the system<sup>8</sup>. Faced with the question of what it is for a system to have a belief with such-and-such content, we find that in both cases it is the system's *abilities to act* in a larger environment that counts. Those abilities are not what they are merely in virtue of some underlying intrinsic physical structure. Instead, what matters are the *superstructural* facts concerning the space of environmental interactions into which the system enters:

One of the key metaphysical shifts in the AAA [Autonomous, Adaptive, Anticipative Dynamical System] approach lies in a shift from understanding complex system identities in static structural terms to understanding them in active process terms ... Intelligent systems ... are adaptable, for unless they can constantly adapt to mitigate or compensate for disturbing signals, they will be disrupted and, losing their cohesion, lose their identity as that sort of system ... This kind of organized responsiveness, especially when based on self-organization, leads to a complex, developing system identity, an individuality of internal conditions distinctive of intelligent, learning systems.

(Hooker & Christensen, 1998, p. 106)

As I said, Smith's proposal is based on the idea of searching for a set of criteria that determine particularity, without making any a priori commitments to the metaphysics which allegedly underlies those criteria. The search is thus constant and, especially at the beginning, unstable. For what will eventually turn out to be such and such object or subject is

the result of an active process in which both the s- and the o- region have to constantly adapt to each other. They both have to compensate for 'disturbing signals' from the environment in which they are embedded. That is exactly how the Dynamical System approach understands the idea of a system's identity.

Furthermore, both in Smith's and in the Dynamical System's approach, the system's skills —as opposed to any supervenient properties of its internal structure— are precisely what gives the system its identity. The system's skills, its 'organized responsiveness' to the environment and to itself are the only reliable signals able to drive some kind of self-regulating mechanism which, in turn, makes the system the system it is.

The parallelism between Smith and Dynamical Systems also helps illustrate what has already been mentioned as a glaring objection to Smith's project. Smith wants to disallow any assumptions about what is relevant to determining particularity. Particulars are supposed to *emerge* from the interaction between the s- and the o- regions within the metaphysical flux. They are the equivalent of attractors in a dynamical system. But, the emergence of an attractor is only possible once the parameters that define the geometric state space have been determined. We can make the dimensions of that semantic space correspond to observable (sensory) properties or we can make the dimensions correspond to social/aesthetic properties. But one way or another, in order to have certain patterns of activity in a dynamical hyperspace, the dimensions have to be fixed. Dynamical Systems approaches, unburdened by Smith's metaphysical agenda, can afford to make the necessary assumption. Smith, alas, cannot.

It must be admitted, however, that the Dynamical Systems approach has thus far been most successfully pursued in cases where the parameters to which the system is responding are straightforwardly physical, and are available to the system as ambient, proximal stimuli. This is the case, for example, in dynamical explanations of reaching and grasping in infants (see e.g. Thelen 1995). Such models are not successful, as yet, in providing explanations of more complex, often language-related phenomena. The same, we saw, seems to be true of Smith's proposal. Smith claims, in his defense, that he is not interested in meaning but in meaningfulness; something more dynamic and unstable

than whatever is relevant for the application of language-related notions. But one wonders whether a theory that has as its main aspiration the provision of a general theory of intentionality can *really* afford not to be more explicit about the fate of linguistic concepts.

Overall, then, although Smith's holistic and anti-reductionist views are forcefully presented, the question still remains as to whether what *OOO* provides is a substantive theory of intentionality or just a set of constraints upon a theory of how intentionality emerges. The metaphysical project was to steer a path between naive realism and pure constructivism. Yet, most of the time the reader seems to be left oscillating between these two poles without a stable resting point. Nonetheless, Smith succeeds in bringing to light, and sometimes challenging, some of the deepest assumptions buried in contemporary attempts to explain intentional phenomena. The book hints at a truly non-causal and non-physicalistic characterization of intentional properties. Unfortunately it delivers only part of what it promises, and many of the arguments will be persuasive only to someone who is already committed to Smith's own desiderata. Such doubts aside, Smith does us a great service in pursuing this genuinely alternative vision. It is still a view from the trenches, and much remains muddy and unclear. But it is also a valiant and thought-provoking meditation that rewards serious study and should be of great value to anyone interested in metaphysics and philosophy of mind.

## **Acknowledgments**

This critical notice grew out of a series of discussions held in the School of Cognitive and Computing Science, University of Sussex, UK. Special thanks to Ron Chrisley, Jo Faith, and all the members of *On The Origin of Objects* reading group. I would also like to thank the two anonymous referees of JETAI for their insights and suggestions.



## References

- Abraham, R. H. and Shaw, C. D. 1992: *Dynamics. The Geometry of Behavior*. 2nd ed. Redwood City, CA.: Addison-Wesley.
- Bickhard, Mark H. 1993: 'Representational Content in Humans and Machines'. *Journal of Experimental and Theoretical Artificial Intelligence*, 5, 285-333.
- Bickhard, Mark H. and Richie, D. M. 1983: *On the Nature of Representation: A Case Study of James J. Gibson's Theory of Perception*. New York: Praeger.
- Brandom, Robert 1994: *Making it Explicit. Reasoning, Representing, and Discursive Commitment*. Cambridge, MA.: Harvard University Press.
- Cummins, Robert 1983: *The Nature of Psychological Explanation*. Cambridge, MA: MIT Press.
- Cussins, Adrian 1992: 'Content, Embodiment and Objectivity: The Theory of Cognitive Trails'. *Mind* 101 (404), 651-688.
- Dennett, Daniel 1975: 'Why the Law of Effect Will Not Go Away', *Journal of the Theory of Social Behavior*, 169-176.
- Lloyd, Dan 1989: *Simple Minds*. Cambridge, MA: MIT Press.
- McDowell, John 1994: *Mind and World*. Cambridge, MA.: Harvard University Press.
- Haugeland, John 1993: 'Mind Embodied and Embedded'. In Y. Houn and J. Ho (eds.), *Mind and Cognition. Collected papers from 1993 International Symposium on Mind and Cognition*. Taipei, Taiwan: Institute of European and American Studies. Academia Sinica, 3-37.

- Hendler, James A. (ed.) 1989: *Special Issue: Hybrid Systems (Symbolic / Connectionist)*. *Connection Science*, 1, (3).
- Hooker, C. A. and Christensen, W. D. 1998: 'Towards a New Science of the Mind: Wide Content and the Metaphysics of Organizational Properties in Non-Linear Dynamical Models', *Mind and Language*, 13 (1), 98-109.
- Lenat, Douglas 1983: 'Theory formation by heuristic search', *Artificial Intelligence*, 21, 31-59.
- Ritchie, G., and Hanna, F. 1984: 'AM: A Case Study in AI Methodology', *Artificial Intelligence*, 23, 249-268.
- Thelen, Esther 1995: 'Time-Scale Dynamics and the Development of an Embodied Cognition'. In R. F. Port and T. van Gelder (eds.), *Mind as Motion. Explorations in the Dynamics of Cognition*. Cambridge, MA.: MIT Press, 69-100.
- van Gelder, Tim 1993: 'The Distinction Between Mind and Cognition'. In Y. Hounig and J. Ho (eds.), *Mind and Cognition. Collected papers from 1993 International Symposium on Mind and Cognition*. Taipei, Taiwan: Institute of European and American Studies. Academia Sinica, 57-82.

## Notes

1 As we shall see, even this way of phrasing the problem contains some of the assumptions that Smith challenges. It presupposes e.g. that there is a sharp distinction between representation and what represents, between something that is out there in the world and something that is not: 'In philosophy, it is traditional to view representation as independent of ontology ... [T]he rest of the book [can be taken] as a proposal for the consequences of its [this view] being false' (*OOO*, p. 42, footnote 18).

2 This is probably why, interestingly, Smith only mentions Turing once in the whole book (*OOO*, p. 79), in the context of the standard —abstract, and mathematical— view about computation that Smith is criticizing.

3 This notion of constitution ought to be distinguished from another kind of relation, also sometimes called a constitutive one, but which is more logical than metaphysical. The *logical* notion of constitution amounts to something like a set of necessary and sufficient conditions that something has to meet for being a thing of a certain kind. In current debates, the phrase 'being constitutive of' —as opposed to 'constituting'— tends to refer to this logical cousin of the metaphysical notion. Thus even though it is constitutive of the number 2 to be the successor of the number 1 (in the logical sense of constitution), 2 is not constituted (in the metaphysical sense) by being a successor of 1.

4 Of course, similar considerations have been invoked in support of functionalism, a paradigmatically internalist position in Artificial Intelligence. The difference is that, for the functionalist, the role involves only permissible moves in a inner economy, whereas, for the institutionalist, the roles are defined in terms of wider, usually socio-pragmatic, settings.

5 Such was the impression of one of the referees of this paper, whom I thank.

6 It is worth mentioning that Smith's recognition of this problem is not entirely new. The circularity in trying to explain intentional properties by invoking supervenient elements each of which is already intentional has been noted in various forms by e.g. Bickhard (1993. See also Bickhard & Richie, 1983), Dennett (1975), Cummins (1983) and Lloyd (1989, pp. 8-9).

7 Smith does not go as far as e.g. Brandom regarding this issue. Brandom explicitly endorse the thought that 'only communities, not individuals, can be interpreted as having original intentionality' (Brandom, 1994, p. 61)

8. The connection between the idea of a 'vehicle' of content and the internalist thesis of content as constituted by intrinsic properties is not presented here as one of necessity. Mark Bickhard's work on interactive representation (see e.g. Bickhard, 1993) is a prime illustration of the contingency of that connection. In Bickhard's proposal, content is explained in terms of internal-to-the-system functional links between current system states and expected internal outcome if a certain action is carried out. This account of content is certainly intrinsic to the system and does not involve a 'vehicle' or content-carrying token.