INFORMATION AND SEMANTIC COGNITION: An Ontological Account

RADU J. BOGDAN

I Introduction

Information is the fuel of cognition. At its most basic level, information is a matter of structures interacting under laws. The notion of information thus reflects the (relational) fact that a structure is created by the impact of another structure. The impacted structure is an encoding, in some concrete form, of the interaction with the impacting structure. Information is, essentially, the structural trace in some system of an interaction with another system; it is also, as a consequence, the structural fuel which drives the impacted system's subsequent processes and behavior. Information takes various forms because the world has many levels of compositional and functional complexity, under different constraints. The key constraints that matter in the understanding of information are natural patterns of organization, or types, and systematic correlations among types, or laws. These level-sensitive constraints, in the form of types and laws, shape the very form in which information is tokened in some structure, that is, the very form in which it is encoded. As a result, the information-producing interactions bring about different sorts of structures, with various sorts of causal effects and functions, whence so many ways in which information is coded and utilized.

I will consider here two forms of information, material and semantic, needed to understand the semantics of cognition. This is a metaphysical, not a natural-scientific classification. It defines a strategy of analysis which focuses on very general properties of things, events and properties, with their arrangements and interactions, and ignores (by abstraction) their concrete forms of worldly incarnation. Such a metaphysical analysis aligns the notion of information to those of event, thing, property, structure and causation, among many others of the same sort -- notions equally bare, thin, abstract and without worldly flesh, hence metaphysical, yet descriptive of what exists and what does or could happen in the real, thick world.

There is a good reason for telling a metaphysical story of information in general ontological terms. The reason, old as metaphysics itself, is to avail ourselves of the possibility of stepping back from the concrete configurations of the world and examine their underlying and pervasive arrangements, features and relations, in the form of things, events, properties, laws, structures and causation. Information also belongs to this ontological class. It too is present and operative at each concrete level of worldly complexity (whether physical, chemical or biological), and yet is not exclusively or ultimately characterized at any such level. This entails, first, that no particular natural science is the principled science of information, just as none is the principled science of causation; and second, that if we want to understand the invariant nature of information, across its various concrete embodiments, we must run the required analyses and arguments on an idealized conceptual plateau. This is what motivates an ontological account of information.

Given this motivation, the objective of this paper is, first, to sketch a rudimentary ontology of information, needed as a foundation for a naturalist story of the cognitive mind, and, second, to concentrate on the implications of this ontological account for understanding the semantics of cognition. The long term project, of which this is just a part, is a thorough naturalization of the semantics of cognition in informational terms. If the mind is exclusively understood in terms of semantic facts, then this naturalistic account has the implication that a door opener or a crab has a mind. If, in general, the mind is understood as (nothing more than) semantic cognition, as it seems to be understood by a vast majority of philosophers and theorists of cognition, then we ought to be able to give a full scientific account of the mind. I doubt that the cognitive mind is just semantic but the doubt is suppressed in what follows and will only be briefly reactivated in the final section.

II Material Information

Information is an indication, expression and measure of an occurrent and particular material interaction between items of the world organized according to some patterns (or types) and behaving according to some laws. So understood, the notion of information specifies the fact that, the form in which, and the extent to which, something is structurally impacted by interacting with something else. The impact is material in nature, always with structural and often with functional and behavioral results. The notion of information is meant to characterize and measure the structural side or the form of the impact and its subsequent causal repercussions along structural lines, that is, the way the effect of the impact is organized or encoded as well as operative at the receiving end. Cognition is a complicated way of being impacted by and reacting to the environment, a way of coding and exploiting the structural consequences of such an interaction. The structural consequences of an interaction, instantiated in a receiving system, constitute the information from the environment that the interaction makes available. Cognitive information has its roots and key properties in this very elementary fact. To establish this point, we only need a rudimentary, narrow and minimal ontology of information on which to build, step by step, the forms of information implicated in cognition.

The Ontology. The starting point in our metaphysical story of information is a useful abstraction, an analytic device which I call a 'thin ontology'. A thin ontology should be imagined in the form of a world of bare material tokens in arrangements and interactions, under types and laws, whose concrete (or thick) nature is simply left unspecified. For the purposes of our analysis, the only feature that the denizens of this imagined world have is being material -- which just means stuff in space and time. A thin ontology, so construed, is a sort of schematism in need of a domain of instantiation, a sort of system of place holders or slots awaiting actual, thick occupants. It is the latter which constitute the real furniture of our world, its thick ontology.

The real world is ontologically thick, at different levels. A thick ontological level is specified both by the typology of configurations, properties and relations of its key items (such as elementary particles or molecules or cells), and by the laws governing the structural composition of, and the causal interactions among, such items. This means that a thick level is (at least) implicitly defined by the basic theoretical terms and the law statements of the appropriate science. The adjectives characterizing the thick levels as physical, chemical and so on, are science-dependent. The physical, the chemical, the biological, and so on, are scientific ways of characterizing the thickness of the world.

Ontological Filters. At any thick level, the specific types of phenomena and their laws may be said (in terms of a realistically intended metaphor) to 'filter in' only some properties and relations (precisely, those subsumable under the types and laws in question), and to 'filter out' the rest. Only the properties and relations filtered in matter structurally and causally at that particular level. At each such level, the types define the basic sorts of components and their structural arrangements or organizations which the causal laws, operating at that level, range over. At the chemical level, for example, only atoms and molecules (but not elementary particles) matter structurally, as appropriate types of components and arrangements of such components. As a result, causal relations under laws at the chemical level are sensitive to, and so range over, only atomic and molecular interactions and not those involving (say) the subatomic elementary particles. The latter are the business of physical types and laws. At the next level up, it is the turn of the chemical components, arrangements and their interactions to fail to matter structurally and causally, in the eyes of the types and laws posited by biology. This is why a biological type (say, a species) is indifferent to the chemical and a fortiori physical composition of its tokens (specific organisms), as long as biological constraints (say, reproducing or digesting food) and laws (say, passing on the genetic make-up of the species or adapting) are respected.

To push my realistic metaphor still further, the point I am belaboring here is that nature has 'abstractive' powers. This, I think, is an essential clue we must chase if we are to understand the nature of information, both in the world and in the mind. If the notion of information characterizes the structural and causal effects of interactions under laws, then the selectivity that nature displays at each of its thick levels, through the appropriate organizing types and their laws, must have a role in shaping the very forms in which information is tokened and becomes operative at different levels. This means that if a form of information is instantiated at a given thick level, that instantiation can involve only tokens of properties, relations and interactions filtered in by the level-appropriate types and laws.

Here is a schematic example at the biochemical level. An informational transaction at that level could be one in which a synaptic connection conveys some neurally encoded information from one formation of neurocells to another. This is to say that the second neural formation is structurally reorganized following the interaction with the first. This structural reorganization can be said to encode the information transmitted. The information is biochemical in that the aspects which have shaped it (i.e. the first neural formation, the synaptic interaction itself and the relevant factors and constraints involved in the organization of second neural formation) are all biochemical. This means that out of the many sorts of properties that the first neural formation has (e.g., it is made of elementary particles, of atoms too, it has a certain weight, and so on), the interaction and its outcome, the structural form of the second neural formation, retain or filter in only tokens of biochemical properties. This is what the information tokened is sensitive to, this is what makes it biochemical information. In other words, the information encoded in the second neural formation can be said to range only over biochemical properties because the very form or organization of that neural formation (which constitutes the encoding) and its causal reasons (the first formation and the synaptic interaction) have a biochemical nature. When, subsequently, a further neural process taps the information encoded in the second formation to do its own functional job, it too <u>treats</u> the information <u>as</u> biochemical.

All I have said so far is fully compatible with the observation that at each juncture in our biochemical examples there are simultaneous informational encodings and transactions which occur at other, lower as well as higher, ontological levels. A neural or cell formation is, among other things, a subatomic structure. When a neural or cell formation interacts with another, there must occur in one a subatomic encoding of the subatomic impact originating in the subatomic structure of the other. To that inevitable extent, then, there is subatomic, hence physical information encoded in the second formation. In the physical case, the informational transactions and encodings range only over physical (not biochemical) properties and laws; they exploit only the physical (not the biochemical) nature of what is involved. To the extent to which the thick ontological levels of the real world do not match each other, because their types and laws do not, there is no good reason to think that the various types of informational encodings and processes co-occurring at those levels in any natural event match each other either. This remains true even though, of necessity, an informational encoding or process at the physical level is necessary for an encoding or process at the chemical or biological level.

Perhaps the best evidence in support of ontological filters and a discriminating, multilevel ontology, comes from scientific experiment and technological application. Guided by appropriate scientific theories, experimentalists and engineers have for years managed to insert man-made devices in the great chain of nature in order to either observe nature or experiment with it or replicate some of its key functions. This is, for example, a familiar story in current medical practice. The success of the human insertion into biological nature (say, in the human body in the form of artificial hearts or various tubes and tissues) is measured by the nature's inability to detect the inserted device as foreign. How do we explain this phenomenon? When the body fails to recognize and reject an implanted device as foreign it cannot be because it fails to recognize the physical nature of the device; after all, both the natural organ and the artificial replacement, just like everything else in this world, are made of elementary particles. Likewise, when the body rejects the artificial implant, it cannot be just for physical reasons. It is not elementary particle physicists who can explain either the success or the failure of artificial implantation. Biological nature and the human body in particular react selectively, at appropriate thick levels, to whatever physical stuff displays a certain form of organization and plays a preassigned function compatible with a number of biological and chemical constraints. Fooling nature by clever artificial interpolation means taking advantage of its ontological filtering.

The Key Elements . Our task now is to locate the objective features of material information, its thin ontological nature. The account will be simple and entirely qualitative. I will not be concerned with more specific and quantitative properties of information such as quantity or amount, which would require more parameters than our metaphysical story can provide at this stage. I do assume, however, that the basic qualitative aspects are the more fundamental. Given these limitations, it should be clear that no precise and complete definition of material information is intended. The illustrations will be, of necessity, ontologically thick.

Of any given thing which has a property or is involved in a relation we can say that it is <u>in a state</u> of some sort. In this sense, at any given moment, the world is a vast configuration of states. A state comes into being or becomes actual whenever a thing (or several) instantiates one property (or several) or enters into one or more relations with other things. A thing occupies a state because it has <u>interacted</u> with another thing, itself occupying some state.

In order to encode information, a state caused by interaction must be more than a mere effect. The state created by interaction must also encode or structure the impact of the interaction in specific patterns; it must structurally acknowledge the interaction in its very organization. This <u>structural acknowledgment</u> brings us to the very heart of the notion of information. Why do we need the notion of information in an universe of interactions between states of different things? If all that matters is that a state brings about another state as a result of some interaction, then the notion of causation can tell the whole story. Why give it an informational reading? What does an informational reading of an interaction between two things reveal that a causal reading of the same interaction doesn't? The answer is that the notion of causation is calculated to be opaque, while the notion of information is calculated to be sensitive, to the structural form of the interaction relation and its relata. The notion of information is intended to capture the structural aspects and consequences of causal interactions.

Take the environment as a source and the light as a receiver of information. To encode information from the environment, the light needs a specific type of (physical) structure. If the light is homogeneous or blank, thus lacking the required structure, it encodes no information from the environment. In general, a state of a receiver must in its very structure show the impact or trace of the structure of the state at the source which caused it to be. This is a structural requirement on the receiver, if it is to encode information. There is also a structural requirement on the source, if information is to originate there. If the source does not have the right type of information-encoding structure, it may interact with a receiver and yet fail to generate information. Take light again, now as a source, and the eye as a receiver. Mere radiation by the unstructured (blank) light may stimulate the eye's retina and yet fail to provide any visual information (although some physical information is thereby exchanged). Interaction without the right type of encoding structures at both ends yields no information.

When the material structure of the receiver is affected by changes in the material structure at the source, we can say of the receiver that it responds to or is stimulated by the interaction. Any material thing that responds to an interaction displays a pattern of stimulation which is the very state brought about by the interaction. That pattern of stimulation (say, a rearrangement of molecules) encodes information from some state at the source. The notion of stimulation, then, is just another way of characterizing what is happening at the receiving end of an interaction. Nothing fancy, nothing biological, certainly nothing cognitive so far. The receiver gets stimulated by an interaction not only because there is an interaction to stimulate it but also because the receiver can be stimulated by that interaction -- because, in other words, it is (in some thick sense) sensitive to that interaction. Again, there is nothing biological about this notion of sensitivity. It tells us which structural impact gets through and reacted to as information. Information systems, whatever their thick ontological level, have their own dominant formats of encoding the input; these formats in turn specify which kinds of properties and magnitudes from a source get encoded as stimulus information. [1]

The structural aspects of information help introducing a very important condition on information, its essential actuality. This condition construes information as an actual token structure, as essentially an occurrent and complex particular. Only a particular and actual material token can encode the information transmitted from another token because only particular things have the thick structures (say, of the physical or chemical sort) needed to realize the encoding, and because only such things are connected by actual interactions needed to transmit and token the information. If there is no actual token structure, there is no material stuff to encode the information, so no information is encoded, therefore there is no information to talk about. The actuality condition is another way of stating the obvious fact that information takes the form of an encoding and that encodings are actual structures.

There is no way around it. Unless, of course, people mean something else by information, which they very often do.

Whence the critical punch of the actuality condition: Don't treat information as contained in and carried by general or abstract or ideal entities and relations, for it won't work. There is nothing in those entities and relations to do the containing or the carrying. (Remember, if that helps, that Plato and Frege were not in the dirty business of explaining information in the world, nor did they have naturalists among their friends.) Nor can information be contained in either general dispositions, regular correlations or laws, for, again, they do not have the required resources for tokening. Where there is no actual tokening, there is no information.

This is, for the time being, as much thin ontology as we need to get a rough but useful conceptual hold of the notion of material information. The notion speaks of information in a receiver from a source as an interaction between a source and a receiver which is such that, under external constraints, a state of the source has an impact upon the receiver and produces in the latter a stimulation state whose structural organization reflects, under internal constraints, the fact, the nature and the extent of the interaction. The structural organization of the receiver's state can be said to encode information from the source.

From and About. The reader may have noticed that so far I have almost succeeded in avoiding talk of 'information about'. I have not said, for example, that, as a result of an interaction, a state of the receiver contains information about a state at the source. Instead, whenever linguistically and conceptually feasible, I have talked of 'information from', in the form of information-in-the-receiver-from-the-source. The distinction between information-from and information-about is meant to be a principled distinction between material and semantic information. Material information is information understood in terms of material types of facts and subject only to material constraints and laws, at some thick level. What matters in understanding material information is some source-receiver interaction with structural results and structurally sensitive causal effects in the receiver. That is what the notion of (material) information-from is all about.

Semantic information turns out to be a different notion, one which understands information as material information shaped and constrained by <u>further</u> aboutness-ensuring sorts of facts and constraints. The constructions <u>information-from</u> and <u>information-about</u>, with or without hyphen, are (from now on) meant in a stipulational sense to mark and recall the distinction between material and semantic information. I am not trying to capture the ordinary meaning of 'from' or 'about', or of 'information', for that matter. I just need a useful and salient way of talking about various forms of information.

We get the notion of material information when we look at the world as nothing more than arrangements of things, properties and events in some causal interaction which have structural effects of some thick sort. If the world were as just described, then the notion of material information would be <u>all</u> we need to characterize and explain it. Most of the universe is probably of this sort. It only tokens material information. But is material information truly <u>information</u>? Is there really information in a world whose ontology is merely material? Or, perhaps better said, is there any thing in particular about a material ontology which would require a notion of information to conceptualize and explain it? We all seem to have an 'intuition' that there is more to information than the notion of material information allows. I am going to call this the 'semantic or aboutness intuition'. The next section will try to make reasonable sense of it. The intuition, it turns out, is not about the existence of material information but rather about the conditions in which material information tokened in a receiver becomes informative (or informs) about properties of a source. These are the conditions in which material information turns semantic.

III Semantic Information

Semantic information is material information with a functional business determined by teleology. Function and telos are not types of items in a material ontology, nor do their regularities count as laws in such an ontology. This is why the nature of semantic information cannot be understood in material terms. A system must also do something with/to information, for some reason, and those doings must eventually relate or apply to aspects of the world in ways which explain the system's behavior, if the system is to encode semantic information about the world.

<u>In-formation</u>. Aristotle did not talk much about semantic intuitions, yet there is nothing like Aristotle to clean up a good old dusty philosophical intuition, any intuition, any time. If told about our ontology of material information, Aristotle might have said that a state of a receiver R is in-formed (i.e. acquires a certain form or structure) by interacting with a state of a source S. The notion of <u>in-formation</u> indicates that a state of the receiver is structured as it is, or has a form, because a state of the source has interacted with it and brought that structure or form about. In-formation, then, is nothing but interaction with structural (or form-inducing) results. In-formation is all there is to material information. It is all a matter of hyphen. (There is nothing like Aristotle's humour.) Thus, whenever R is in-formed <u>by</u> S, R has information <u>from</u> S. Material information, we said, is information-from. 'From' is the causal 'by' read informationally.

When a receiver R has information from a source S, does R also have information <u>about</u> S? About what at S? About some specific state or property of S? This is the question. How do we get from <u>by</u> (or rather <u>from</u>) to <u>about</u>? From in-formation-by, hence information-from, to information-about? From causation to semantics? How indeed?

<u>How About?</u> <u>About</u> is specific in ways in which <u>by</u> and <u>from</u> are not. <u>About</u> tells us that the information must have some well delineated target; <u>by</u> and <u>from</u> only tell us that the information must have countless sources, indeed as many as a causal chain may sustain. When a state of a source interacts with and produces a state in a receiver, there may well be millions of material (say, physical) microstates and processes along an open-ended causal chain which are involved in the interaction. Many of these are in the source itself, others (to enlarge the context) in more distant causes of the source. All these millions of items may be causally contributing, in various degrees, under various laws, to bringing the receiver in a given state. The latter state, in other words, may be in-formed <u>by</u> all those millions of causes, hence encode information <u>from</u> all of them. Can that state really be <u>about</u> them all, hence about none <u>in</u> <u>particular</u> ? Does <u>this</u> idea of aboutness make any

sense?

Is information-about a sensible notion when fully explicated by the notion of in-formation-by? To the defender of the semantic intuition, this is a reductio ad absurdum: If information (about) is just in-formation, or information-from, then (given the 'about all hence about none in particular' form of argument) there can be no information-about. But there is information-about. It follows that information-about cannot be explicated only by information-from. Semantics is not just matter and causation. Something else is needed, something to take us beyond the pale of bare materiality, in some thick form of ware or another.

Let us not beat around the dialectical bush, by showing the reader (and the professional colleague) how we honestly cast around for various solutions, which are then carefully shown not to work, until we are struck by the illuminating insight. There is no need or time for that. There is only one plausible solution, it doesn't require any special insight, which is why everybody knows it, although very many don't like it. It is called teleology. Information-from also needs some teleology, implemented by various sorts of goals and functions, to become information-about. It takes some appropriate design to make a teleology effective. The point is familiar, has been made by many people (e.g., Dennett, 1969; Bennett, 1976; a self-repudiated Fodor, MS; and others), and it still makes a lot of sense. Let us see why.

The semantic intuition finds that the problem with material information is that it does not inform. Nothing in the receiver is being informed about something at the source when the receiver, upon interacting with the source, comes to encode information from the source. What sort of information could that be which does not inform? Information must be information to or for X in order to be about Y. Aboutness requires some destination for the material information tokened in the receiver. In the austere (material) ontology assembled so far, something is being materially done to produce and encode information, but there is nothing to, or for, whom/what that is information. The receiver is not yet designed to accommodate the additional relation of 'information to/for something in the receiver about something at the source'. If we can specify the 'to' or 'for' (a destination), we get the tools to specify what the information is about.

The notion of material information was calculated to characterize the <u>stimulus</u> information in the receiver, that is, the information which succeeds in stimulating, hence in-forming, the receiver at the appropriate thick level, in the appropriate encoding format. The question <u>now</u> is, What are the conditions in which the stimulus information itself, appropriately tokened and encoded, informs <u>about</u> some particular aspects at the source? How does the receiver reach back, as it were, to select at the source what the stimulus information is all about?

The notion of destination, needed to answer such questions, is liberal. It covers both the immediate roles of an information structure and the more distant objectives such roles help accomplish. I distinguish between the internal role of information, on the one hand, and the more general goals of a system, which such roles serve, on the other hand. Since, in the last analysis, goals shape internal roles and internal roles shape the semantic information which helps realizing the goals, we have to look at a system's goals as the <u>outer</u> constraints on semantic information and at the internal roles as the <u>inner</u> constraints. It is these sorts of constraints which are fundamental in typing the tokened material information <u>as</u> semantic. Thus, the conceptual move in which we indulge now is from destination, in the form of goals and roles of information, to the semantic character of that information. Teleology is what rationalizes this move.

<u>From Teleology To Semantics, I: The Motivation</u>. Why should an organism be so built as to re-encode a stimulus information structure, tokened internally, in the form of information about another, typically distal structure in the environment? Because this re-encoding enables the organism to react internally to, and perhaps engage externally in actions directed at, the things displaying the distal structure. Organisms do this to satisfy their needs and goals, their teloi. Semantic information provides the internal structural means by which goals and needs can reach their environmental targets. [2]

The reality of goals, needs and the functions serving them is the natural reason for there being semantic information in the world. If needs and goals are real, and if it is an objective fact that needs and goals direct behavior toward specific sorts of targets, then there must be a form of information which makes this directing possible by being appropriately constrained and shaped. This is how I suggest we think of semantic information. Or, to put it differently, if the semantic form of information is a fact about the way the world is, then teleology is the best account of that fact. Goals, needs and functions, as teleological types, are there at the right ontological juncture where new forms of material organization (underlying the semantic form of information) reflect new patterns of interactions in the world: from organisms to their environment, in the form of goal-directed behavior; and from environment to behavior, in the form of perception and cognition, specifically, in the form of functional connections and operations which align input to behavior.

The teleological kinds and laws should be thought of as a new sort of ontological shapers and filters of the material structures involved in the new forms of interactions. A goal, for example, is an ontological filter in that it maps an organism's sensory input onto an appropriate behavior only with respect to certain sorts of conditions at the source which can satisfy the goal, while filtering out many other sorts of conditions. Such ontological filtering allows us to explain, for example, how an input structure, tokened in an organism through perception, is aligned to a behavioral output with respect to certain properties of the encoding of the input, and <u>not</u> with respect to its other concommitant properties. Thus, a goal can explain why a lion perceives and goes after a prey (it needs meat, badly, now), by filtering in the relevant (meat-eating) dimensions of the distal target, while abstracting from, by filtering out, various other properties of the target (such as color, smell, shape, etc.). The lion's goal can explain this because it is a constraint on the very shaping of the information which enables the lion both to perceive the target as a meaty prey and behave toward it accordingly.

If we want teleology to do clear and honest work in our account of semantic information, we must delineate and explicate the notions which will do the work. My line is going to be economical. To simplify matters, I reduce needs and other objectives an organism may have to goals. As a working approximation, I will treat a <u>goal</u> as a manner of identifying and rationalizing a class of behaviors which show systematic sensitivity to selected aspects of the distal environment. Concretely, a goal can be construed as specifying a type of condition of an organism which a class of behaviors can bring about. The condition itself need not be specified in semantic terms. Goals are not semantic constructs. (This preempts the likely charge that the teleological account of semantic information is circular.) The goal-satisfying condition may be the condition of landing on a flat surface, or eating a potato, or something like this. (Nothing semantic here.) A behavior aims at bringing about the goal-satisfying condition by (among other things) causally interacting with relevant types of targets in the environment. It is this very fact that explains the presence and operation of semantic information, thus giving us the first grip on the semantics of cognition.

Not all the goals of an information system are of the same caliber. We can distinguish between <u>vital</u> goals, such as eating, reproducing, avoiding danger, things of that magnitude, on the one hand, and <u>active</u>, instrumental, specialized goals, appropriate to various types of environmental circumstances in which the cognitive and behavioral functions of the organism must be exercised, on the other hand. The active goals serve the vital ones. The active goal of a system's vision can be that of seeing where to land, or recognizing the tiger in the bush, or other such targets of immediate behavioral concern. It is the active goals that we need in our pursuit of semantic information. Vital goals are not served by information, so they cannot be expected to shape information. They are too general for that. How could the vital goal of staying alive or that of enjoying oneself shape any sort of information? Vital goals are satisfied only when active, specific goals are.

Goals are implemented internally by <u>functions</u>. Organisms are made of parts, devices and organs which contribute to bringing about the goal satisfaction conditions. The latter can be actualized only if the contributors properly execute their assigned functions. The execution of such functions is in turn activated and guided by the relevant information structures. In order to activate the executors of the goal-serving functions, and thus play its role, the information must be encoded in ways which the executors can recognize and respond to. This is another way of saying that (like goals, at a further remove) the functions to be executed internally shape the form of the information that enables their execution. It is then in this sense that goals and functions, as teleological types, type the material information (which will serve them) as semantic.

Having made some initial sense of the basic teleological notions expected to deliver the semantics of cognition, it is fair to emphasize the reassuring minimality of my teleological approach. All the reader is required to consider is the reality of the facts the teleological notions point to, that is, the reality of goals, of goal-directed behaviors, and of the internal functions whose execution satisfies the goals by appropriately directing those behaviors. Although I call the notions in question and the characterizations and explanations that employ them 'teleological', I make no claim here as to how the notions are to be in turn analyzed and explained. Nor does this matter. They can be analyzed and explained at an autonomous and emergent level or, on the contrary, reductively, say, in terms of evolution by natural selection or in terms of the cybernetics of self-organizing systems or otherwise. No matter how analysed and explained, the teleological types I need must reflect the distinct reality of an organism's selective behavior toward well delineated aspects of the environment, a behavior which originates in and is guided by, appropriate internal states and processes. It is the presence, not the nature, of these types of phenomena that is important for our story of semantic information. Teleological types are simply new sorts of ontological filters shaping new sorts of systems whose behavior displays systematic patterns which are unaccountable in terms of preteleological types and laws.

Although, in the large cosmic picture, appeal to teleology explains the very existence of and the rationale for semantic information, it does not explain everything about semantic information. Teleology does not explain how semantic information is encoded, computed and fed into behavior. Nor does teleology explain the specific aboutness of a particular semantic structure. To have access to these further explanations we need to go from the generic teleological constraints on information systems to the concrete internal conditions, both structural and functional, in which such systems encode, process and utilize information. Only then will we have a portrait of the forms of encoding and processing semantic information, indeed, of the intentional program of a semantic system. That route will be outlined next.

From Teleology to Semantics, II: The Steps. I propose to chart the route from teleology to semantics in the form of several methodological steps. These are steps that a designer of semantic information systems might want to follow in his work, or steps that the theorist of such systems might take in trying to understand what makes the systems semantic. In one form or another, this is a useful expository strategy. Here are the first two steps.

(1) **active goals -----> targets for behavior**

(2) **targets for behavior ---> aspects of semantic significance: semantic types**

(1) tells us that it is the active goals of a system that we must consider if we want to determine which types of features at a source (with which the system is interacting) are of behavioral interest because likely to lead to the satisfaction of vital goals. (2) says that those types of features at the source, (type) targeted by behavior and (token) signalled to the system by the stimulus information from the source, are semantic aspects of teleological interest to the organism. I call them <u>semantic</u> <u>types</u>. If, for example, a bird has the active goal of landing safely, then a suitable surface is going to be a favorite type of target for its landing behavior. A few properties of landing surfaces, such as texture, flatness and size, and a few relations to surfaces, such angle of descent and distance, are going to be particularly instrumental in guiding behavior, hence likely to be candidates for semantic types.

Teleological types identify the semantic types for an organism, for they specify the types of environmental aspects which matter to the organism. But knowing the semantic types for an organism is not yet knowing how the organism is going to make the types available to its behavior. We are faced with the problem of establishing how the organism is organized internally so as to implement its semantic policy, which in turn means establishing how the organism handles internally the information about the relevant semantic types. This is now a matter of <u>design</u>. The semantic types must themselves be captured (encoded) and put to work internally, in formats suitable to guiding behavior toward the sorts of targets that the semantic types identify. The next step, then, is

(3) **semantic types ----> their internal encoding, processing and functions: design problems**

This step deals with the internal implementation of a type semantics for an information system -- specifically, the problems the design of the system must solve to encode, process and utilize the semantic types required by its teleological well being. One problem is that of finding the right form of encoding the semantic types. This is the encoding problem. (In some systems, a syntactic code or even a natural language will be solutions to this problem.) Another problem is how to operate on the encodings to get new encodings and eventually action. This is the processing problem. (In some systems, computation is a solution to the processing problem.) The third and most important problem is that of a semantically sensitive input-output alignment. This is the intentionality problem. We need functions which align the input to the system's behavior in a semantic way, thus enabling the system to act on those types of targets whose selection is teleological and whose current token presence is signalled by the input. We are talking of <u>intentional</u> functions, such as concepts, recognition patterns and the like, which discriminate the input for evidence of the semantic types that the behavior is sensitive to. The set of these discriminative and aligning functions can be called the <u>intentional program</u> of the system.

The reason why the input-output alignment problem is the most important is that the intentional functions which execute the said alignment are directly responsible for implementing the semantic policy of the organism, for they are most instrumental in establishing what the input is all about<u>, as far as the behavior is concerned</u>. The solutions to the other problems, the form of encoding and the form of processing, can be viewed as ways of executing the intentional alignment: the input is encoded and processed the way it is in order to be brought under the right concepts which in turn will activate the right behaviors.

For an example, we can regard the design of a bird species as having been confronted with the (encoding) problem of how to structure the information about, say, the size and texture of a surface or that about the distance to a surface under certain angles of flight; also with the (processing) problem of how to coordinate the fragmentary information about these aspects in real time; and finally with the (intentional) problem of how to recognize the semantic significance of the information (this is a landing surface) and feed the result into the appropriate (now land) behavior.

The fourth step goes from the design problems to their solutions:

(4) **design problems ---> solutions: architectural orrepresentational**

The solutions to the encoding, processing and intentional problems are of two major sorts, architectural and representational. The architectural solutions, primitive and inflexible, are built into the ways the system is assembled and operates; the representational solutions, while inevitably relying on the right architectural opportunities and strictures, are flexibly and explicitly manufacturable encodings of semantic information. These types of solutions are arrived at, by natural evolution or artificial design, under various sorts of pressures and constraints, with an eye to maximal benefits against minimal costs. We will occasionally consider these conditions surrounding a type of solution or another. But the main interest is in the semantic character of the solutions, not their particular origin or rationale.

IV. Semantic Information Systems

Any information system to which the methodology sketched by steps (1) to (4) applies qualifies as semantic. The task now is to examine, illustrate and motivate the key conditions in which, step by step, an information system comes to be semantic. We will meet obstacles and clever traps, which will require some important critical detours. It is going to be a long section.

THE CONDITIONS. We now get acquainted with my paradigmatic example, DOP, the door opener. By gradual comparison with other similar devices and natural organisms, we will build DOP in stages until it becomes a genuine semantic information system. In the process we will introduce and examine key conditions on naturalized semantic information. Given the fact that our favorite system is very primitive, its semantics is bound to be architectural.

Encoding Sensitivity. In the beginning DOP is just a sensor, in the form of a photocell which registers certain light patterns. It is not different from simple measuring instruments, such as altimeters and speedometers. They all display material sensitivity in that they are stimulated only by certain types of input properties, which is why they have their own dominant formats of encoding the input. These formats, we recall, specify the kinds of properties and magnitudes from a source which can be encoded as stimulus information. An altimeter conveys information about altitude by means of material encodings of pressure. Our photocell conveys information about moving objects by encoding information from appropriate light patterns. Is the information about altitude or about moving objects also encoded *semantically* by the altimeter or the photocell, respectively? The answer must be negative. All the altimeter encodes (in its dominant format) is pressure magnitudes, and the photocell just light patterns. This is what they can materially react to or be stimulated by. It just so happens that, given some laws of nature, pressure is correlated with altitude, and light patterns with moving objects, so that the information from pressure magnitudes is also information about altitude, just as information from light patterns is also information about moving objects. But, in both cases, the information about is not *for* the sensors which make it available, because they do not token it in a semantic form. This is because the sensors are not constrained in the required way. And they are not constrained the right way because they do not have to. Nothing follows functionally and behaviorally from their having information about altitude or moving objects. The information those sensors deliver is semantic only for us, because we are constrained in the right (aboutness-ensuring) way and we also know about the correlations between pressure and altitude, or between light and moving objects. It is because their registration and encoding capabilties do not reflect further functional and behavioral constraints that the sensors encode only material, not semantic information.

Behavior. It seems, then, that a form of encoding stimulus information must result in some behavior if the information encoded is to qualify as semantic. Will a thermostat qualify? Imagine one which activates the air conditioner when it measures 70° F. A perceptual value is connected to a behavioral output. Or imagine a speedometer combined with a buzzer which tells the driver when the speed is 55mph, or (better) a speedometer which cuts the gas at 55 mph, screams at the driver and then pours champagne on him? Or indeed our upgraded DOP which now consists of a photocell <u>and</u> a motor part which opens doors when the photocell sends the right information? Of course, not any sort of behavior will do. The information systems we are now considering behave in ways determined by what their sensors sense and measure. There is in each of them a recoding of the stimulus information which goes into specific behavioral outputs.

Is this going to do the trick? Not yet. We may decide, cleverly, to simply correlate selected values measured by the sensors with certain types of behavior. We first observe what a sensor does, find a pattern of registration, make it selective within preassigned limits, and then exploit the pattern by correlating it with forms of behavior we are interested in. (This, I take it, is how thermostats are designed.) Such an adhoc perception-behavior connection cannot token semantic information, for there are no internally reflected constraints on such tokening. What we want <u>inside</u> the system are patterns of encoding the information under constraints which reflect in a systematic way what the system is supposed to act on or toward.

<u>Internal Selectivity.</u> So what our DOP also needs, to token semantic information, is some internal selectivity which would convert the input, as materially encoded, into information about specific distal aspects at the source. (Imagine, for example, that the complete DOP is designed to open doors only for people who look intelligent, think syllogistically, are funny, and carry a copy of the ECONOMIST. Not much work in that case.) The reason such an internal selectivity is a major condition on semantic information is that a tokened information structure counts as semantic only if its shape and function in a system can be explained, under approriate types and regularities, relative to some distal properties. The information structure must therefore be shaped <u>inside</u> the system, by its architecture and modus operandi, in ways which can be explained only by appeal to semantic considerations.

The question now is, How are we to construe this internal selectivity? It is critical, for my approach, to clearly distinguish it from other lower-grade forms of selectivity which I have liberally bestowed on nature. We must rememeber to distinguish it from encoding selectivity. The latter concerns what is encodable at all, in some dominant format. What is encodable is what ends up tokened as material information, as a stimulation pattern. Nothing semantic so far. The difference we are after is exemplified in that between which sounds get acoustically encoded as words (encoding selectivity) and what the sounds encoded as words mean or carry information about (internal, intentional selectivity).

Appearances notwithstanding, digitalization must also be excluded as a condition on semantic information. There are three reasons for that, which I will have to discuss rather briefly and dogmatically. First, any measuring device (altimeters, speedometers, watches, and so forth) can be built to react to and measure what they have to in either analog or digital ways. Since such devices are not semantic to begin with, digitalization cannot ensure semanticity. Digitalization is just a manner of encoding the input. Second, digitalization does not provide an internal selection geared to aspects of semantic significance. The digital format provides the means, not the subject matter, of encoding. The digital means may well be more conducive to implementing genuine semantic discriminations such as those provided by concepts. But then the analog means may be more suitable to implement the semantic recognition required by a template. Third, digitalization fails because its principles of selecting what is encoded may have nothing to do with the system's intentionality and behavior. This failure gives us a clue as to what we need next.

Intentional Alignment. To have semantic information we want internalized links between sensing and behaving. We want such links to constrain internally the very ways in which the system encodes and acts on information. For that to happen, in such infomorons as we consider here, we must provide them with functions which align input to output in ways which make explanatory sense only by appeal to semantic facts. This is the origin of intentionality. The aligning functions, such as concepts, templates, recognition patterns, various coordinating schemes, etc., are executing an intentional program. The alignment of the stimulus information from a source to the appropriate behaviors has internal reality because it is now built to execute an input-output (or intentional) function, which is why it constrains the encoding of information the way it does. What the system does must constrain in a principled fashion what types of aspects it registers, if the system is to token semantic information. The constraining goes from types of behavior to types of aspects registered, not from some individual behavior to some individual object of registration.

Prior to its intentional alignment, DOP failed to be semantic. Although it opened doors for moving objects, and although its actions were triggered by its photocell registering certain types of light patterns, systematically correlated with shapes in movement, we cannot plausibly say that the DOP had semantic information about moving objects, because its door-opening behavior did not yet have the appropriate relation to those moving objects. Neither the presence of the appropriate behavior nor its directedness toward distal targets sufficed to token semantic information. The behavior as well as its distal directedness can be separately faked, or accidentally realized. This shows that the information which drives the behavior and ensures its directedness is still not internally constrained in the right way, the intentional way. Since information is structure under the right constraints, the system in question does not token what it is supposed to. Only when there is a natural, systematic and fully internalized form in which the system's behavior is directed toward certain sorts of distal aspects revealed by the input, and only then, can we begin to say that the information which activates the behavior is about those sorts of distal aspects.

Teleological Behavior. So are we there, after all? Not quite. We have perception and behavior, we have a perception-behavior link up, we have such a link up motivated intentionally in terms of input-output functions, so (to put it roughly) we have behavior constraining what is perceived. What we do not yet have is a complete teleological motivation of the behavior whose targets identify the types of aspects at the source which matter vitally and actively to the system under consideration. Why do we need such a complete teleological motivation to talk of an information system as semantic? Because we want teleology to firmly fix the distal semantic types, so we can go top down and see, step by step, how the teleology is implemented (architecturally or representationally) in the appropriate intentional types, encoding formats, strictures on processing and functional deployments of interim outputs into the final forms of internal reaction or external behavior -- all these aspects being made explanatorily intelligible by the teleological determination of the distal sources.

Another way of looking at the significance of the complete teleological motivation of the distal semantics is the following. On our analysis prior to this teleological motivation, if (counterfactually) there were any semantic information tokened in the systems considered, it could, it just could, be at most about the input or some external but very proximal structure, such as the light in vision. Without a clear teleology hold on distal targets, and a clarification of what this means, we might only get proximal semantics, and we do not want that. For if proximal semantics makes sense, then my entire approach to semantic information doesn't. Hence the urgent need for modus tollens.

I am going to use the notion of proximal semantics rhetorically (because I do not believe in it) and a bit deviantly, as the complement of distal semantics. Distal semantics involves only the source, proximal semantics whatever is between the source and the intentional program or control center or whatever in the receiver that takes the cognitive input and feeds it into its behavior. In this sense, to take vision as an example, both the light and the retinal patterns are proximal, only the object reflecting the light is distal. The prejudices favoring a proximal semantics are many, deep and made clever by a long philosophical tradition. I do not have the space to handle them in a leisurely way, so here is a quick aperitif.

Proximal semantics makes neither logical, nor functional, nor teleological sense. Logically speaking, an end need not reflect its means, just as a result need not reproduce its enabling conditions. An image I acquire visually cannot plausibly be about the visual means by which I acquire it. Why should it? What's the point? Likewise, the very thought I am now articulating is not about its logical and grammatical forms. Yet these forms are among the means by which I articulate the thought. If an image could be about its proximal means (causes), it could be about a lot of things which counts as its means, say, interim neural formations in the brain, retinal patterns, light patterns and so forth. How do we decide? One recognizes here a version of the earlier slippery slope generated by the very weak causal account of aboutness. For good reason, because the notion of proximal semantics is really causation redescribed. No wonder, then, that the logical point against causation explaining semantics can also be made against proximal semantics.

Proximal semantics does not make much functional sense, either. We cannot understand how a perception-action link up works in an information system if we assume that its functions are determined by a proximal semantics. We have to assume semantic distality to understand, for example, how the perceptual representations in humans are assembled and work. The argument is to the effect that the rules and mechanisms of perceptual representation are functionally motivated only by the assumption that they structure information from distal sources (see Marr, 1982; Fodor&Pylyshyn, 1981; also Dretske, 1981, on constancy mechanisms). The argument can be generalized to other aspects of cognition.

Finally, proximal semantics makes no teleological sense because it explains neither the rationale for, nor the very operation of, the link up between perception and behavior in an information system. As mentioned, the explanatory generalizations about a perception-behavior link up are lost under the assumption of a proximal semantics. We do not eat light and air, even though the information from food is made available to the eye by light and to the nose by air. The behavior that matters teleologically is, in this case, eating food, not aligning the arms to the visual information coming from the food, even though that alignment is necessary to our getting the food. The internal functions which execute the perception-action link up do indeed connect the input (in this case, visual data from the food) with the appropriate behavior (grabbing the food). It does not follow, however, that we can understand what information we have in such a case, or rather what shapes the information we have, if we consider perceptual inputs and behavioral outputs as the only important types of constraints on that information. If we take the information tokened in a system to be proximally semantic, we are going to miss the factors (types and laws) which filter, organize and use that information.

Such, then, is the regular connection, secured by teleology, through internal encodings and functions, between distal properties or distal behavioral targets outside a system and the semantic information about them tokened inside the system. It is either distal semantics or no semantics at all.

System-Wide Semantics. The semantic upgrading of DOP now amounts to the following. Taken separately, the photocell itself (DOP's perceptual organ) only detects proximal light patterns and tokens the stimulus information from them. Yet the whole system (perceptual and motor) is designed to react to certain types of distal sources. This means that internal function of the photocell is to register only those input features which, for the motor device, have a precise distal significance: solid objects in motion for which the door must be opened. By design, the perceptual competence of the system is aligned to, because constrained by, the motor competence. In other words, the architecture of the motor part embodies a behavioral capability, whose exercise is to move the steel arm linked to the door, whenever the photocell says so; whereas the architecture of the perceptual part embodies a perceptual capability whose output is fed into the exercise of the behavioral capability. It is this cooperation of the two capabilities linking perception to action, under intentional (alignment) connections, which implements the assigned goal of the door opener. Constrained in this way by its goal and the perceptual and motor functions implementing it, the stimulus information tokened in DOP has semantic significance. DOP can now be confidently said to have semantic information about the distal moving objects for which it opens the door.

What confers semantic quality on the tokening of stimulus information is the structural and functional projection, in the systems's very architecture and/or mode of operation, by way of evolution or design, of the types of distal targets toward which the system's behavior is directed, for teleological reasons. It is that projection, in the form of an intentional program, which internally constrains and types the organization and function of information as semantic. It is the very fact of the projection, not its locus or geographical tokening, that matters in understanding the nature of semantic information. Whereas the conditions in which information is typed as semantic are system-wide (because they reflect la raison d'ê tre of the system), as attested by the alignment of the perceptual to the behavioral part, the tokening of semantic information need not be so. It can be and often is local. It should not follow from our analysis that the semantic information is tokened throughout the system or that it isn't. In most advanced systems it isn't. In such systems the semantic information is tokened in specialized subsystems, such as vision or memory. We can say then that the system has perceptual or memory information about moving objects or about something else. But it does in no way follow that the typing of the token pertceptual or memory information as semantic has only perceptual or memory roots and reasons. Location of tokening is no indication of the source of, and the reasons for, the typing of information. The typing is systemic. [3]

SOME REAL LIFE SEMANTICS. Bacteria are real life semantic infomorons. They are naturally designed to detect nearby types of events which they functionally exploit to meet some simple needs (such as staying alive or multiplying). Bacteria can be said to encode information about whatever there is in the nearby distal events which satisfies their needs. The reason the semantic information attribution is legitimate is that bacterial systems have an internal metabolic behavior which responds selectively and systematically to certain distal properties; their response is generated by their internal make up whose rationale is to make the response possible. In other words, the internal make up was designed or has evolved to constrain the incoming information in ways which make its response systematically sensitive to the distal aspects of vital interest. The fact that bacteria do not always engage in actions directed at the distal properties is no reason for them not to token semantic information. The internal response is sufficient. For what matters is the systematic internalized connection between the relevant distal properties and the nature of the organism's response to those properties, not the reach of the response. Our own mathematical thinking can be semantically about numbers, under the right rules and constraints, without necessarily involving some behavior brutally directed at numbers.

The elevation of bacteria to semantic status is not a very prudent move, for it is likely to reenergize the resistance against my approach to the semantics of cognition. What is it in the bacterium that is about some distal property, or about anything? And if the bacteria has it, why not the photocell? First point first. Most of us are conditioned by the prejudice that only a fully explicit representation (an image, a sentence, something like this) tokened in a system can be about something at all. This is precisely the prejudice challenged here. If, as I maintain, there are forms of encoding information which are semantic (about something) but not representational, it follows that those forms of encoding can be about something without representing it, that is, without being a picture of it or a symbolic expression of it or anything of that sort. If this is true, then we should not expect to find in a semantic system some sort of entity that we can compare with its distal cause.

Neither the bacterium nor the photocell represents anything. They both encode information by tokening architecturally constrained properties. The tokened properties are in both cases material, at some thick level; they bear no representational relation to the distal sources causally responsible for their being tokened; and the shape of their tokening is determined by their design, by how they were built. This is where the similarities end, for one encodes information semantically, while the other doesn't. Where do we look for the difference? In the architectural design, for it embodies the types and constraints which organize the tokened information in ways which reflect what distal aspects the system takes structural and functional account of. In the case of the bacterium, the tokening of the information is made architecturally sensitive to distal aspects of teleological value, while in the case of the photocell there is nothing of the sort. [4]

If this is hard or couterintuitive to swallow, consider the following point which anticipates our story of representation. An explicit representation in your head (say, the very English sentence you are now reading) is nothing but a material structure, at your favorite thick level: clusters of atoms or molecules or cells bouncing around in your brain. This is what's happending up there. How, then, as a theorist, do you know that it represents, not what it represents, just the very fact that it represents? You look at some properties of the material structure and its casual behavior and conclude that they have been typed according to rules and constraints (of grammar, logic, and so on) which do not make explanatory sense at the thick level at which the representation is just a material structure. You conclude that your brain has been built according to a new typescript. This is why it represents. The typing is not visible, either in the bacterium or the photocell or your brain. What is visible in all three are just thick material structures getting tokened and transformed into others. Just looking at them will not tell you whether they are semantic or not, whether they represent or not, let alone what they are semantically about, or what they represent. So, again, how do we decide? And, again, I say, we look at what types and constrains their architecture. If the types and constraints make explanatory sense only materially, at some thick level, then the system encodes just material information. This is the photocell. If the types and constraints make sense only by appeal to a teleology geared to distal aspects, then the system encodes semantic information. If the typing and constraining is merely by way of architecture, then we have architectural semantics. This is the bacterium, and DOP as well. If, however, the typing and constraining of semantic information requires explicit displays and formally sensitive operations on such displays, then we have representational semantics. That is your brain.

To continue our sampling of real life semantics, consider now a more complicated semantic information system than the bacterium or DOP. The housefly, we are told by David Marr (1982, pp. 32-34), has a visual system designed to control its flight and track various types of distal targets: landing surfaces, (eatable) objects of a certain angular dimension, other objects to be avoided, and so on. The fly can track these types of targets only if they are made available by the visual information it has from its environment. Knowing the active goals as well as the behavioral and perceptual functions which implement the goals amounts to knowing that, from the wealth of the stimuli it gets from the environment, the fly encodes and acts only on the information which specifies the distal targets the fly has a vital interest in. This in turn tells us which aspects of the stimulus information are worth encoding (i.e. those which systematically correlate with teleologically significant distal targets) and hence have semantic significance for the fly. Teleology thus explains both the principled need for semantic information and, in particular, the conditions in which the stimulus information from some distal source turns semantic.

Let us now ascend to the human level. Developmental psychology reports that for the first few months the child lacks the concept of a permanent object, and many other concepts besides. As Piaget observed many years ago, if a toy with which the child plays is covered with a blanket, the toy ceases to exist for the child. She makes no effort to retrieve it, nor does she show any toy-related attitude. After the child reaches a certain age, however, the toy acquires permanent existence, even when hidden, for the child is then capable of engaging in search actions which are toy-directed. There must be internal information structures (recognition patterns, concepts, beliefs, etc.) and internal functions (expectations, inferences) which make such actions possible. The actions, then, are the rationale for the responsible types of information structures and functions, not only for their mere existence and operation but for their semantic range as well. The child's actions delineate, by means of the relevant intentional functions which guide them, the types of distal facts her information can possibly be about.

V. Words And Objections

Time for a very brief critical interlude. There are a few terminological and substantive matters that had better be cleared up if our story of information is to continue credibly. This will also allow us to summarize, firm up and sharpen some of the points made so far. [5]

WORDS. A number of misunderstandings and objections may originate in the very terminology, at times unorthodox, that I have adopted in this work. I can anticipate the following.

<u>Why Semantic?</u> One may not like to apply the <u>words</u> 'semantic' and 'about' to most of the infomorons we have considered so far. That is all right. It is the facts we must be after, not their labels. One could call the facts 'Xerxes' or 'Artaban' or whatever. What matters is the ontologically motivated distinction between the material and the semantic forms of information. Even if one is reluctant to call what I call 'semantic', one is still under the intellectual obligation to explain, first, what is it informationally that allows target-directed systems to systematically behave the way they do, and, second, why is this different (because it <u>is</u>) from the behavior of simpler, nonteleological systems. The explanation need not invoke semantic information but it

must handle the facts. Appeal to information, in its various forms, is just a manner of reading the facts.

The Epistemological Angle. Then there is the way in which I have argued about the individuation of information, particularly its semantic version. I have been often saying things in an epistemological tone of voice. Thus, for example, I was talking about how the target of a system's behavior provides evidence that the functionally responsible information tokened in the system is semantic; or about how other constraints allow us to attribute semantic information. This was just to get the thesis across. The thesis itself is not epistemological. It has nothing to do with what we know about sources, receivers, their teleology, their interaction, the information thusly produced, acted upon, and hence semantically significant. All along, the thesis has been ontological, even though phrased epistemologically. It concerns objective facts about the sources and receivers themselves, the worlds in which they operate, their teleological raison d'ê tre and (therefore) the very information they encode and act on.

Behaviorism? Nor should my position be construed, from bits and pieces of terminology, as behavioristic. I have said on a number of occasions that the organism's current behavior may be evidence for, because the effect of, semantic information being tokened and efficacious in the organism. But I do not think that the current token behavior is constitutive of the semantic information. Nor do I think that behavior in general is an antecedent condition on semantic information. (So I cannot be a behaviorist.) My claim rather is that both behavior and semantic information are conjointly made possible by teleology and implemented by design. One can even imagine a behaviorless system (not the computer, however) so constrained as to token semantic information which causes no behavior; the system's possible behavior is envisaged or simulated during design but never given actuality.

OBJECTIONS. The above have been answers to possible objections to how I said what I said. We turn now to objections to what I said, no matter how said.

Circularity. Is a teleological account of semantic information circular? Am I not saying that the individuation of semantic information in a system requires the individuation of the targets of the system's behavior in terms of its goals? I am. Isn't it then the case that to individuate behavioral targets and their goals, we have to already know their semantic range, what they are about, in which case one aboutness (of the information) is being defined in terms of another aboutness (of goals and behavior)?

I have already anticipated and almost answered this charge. There is no circularity in the teleological account of semantic information because no semantic notions back up the teleological ones. Recall that aboutness emerges only at a stage in our analysis when we go from (already defined) goals, behaviors and their targets to the internal structures and functions which enable behaviors to reach their targets and thus satisfy the goals. It is only the internal structures and functions, not the goals, behaviors and their targets, which token semantic information and hence are about something. The point of the entire analysis is that, as an abstract designer, one has first to know what goals and targets of behavior a system has before one knows what semantic information the system tokens -- or, ontologically phrased, the system must have goals, functions and targets of behavior in order to token relevant semantic information. To reflect this fundamental feature of the present analysis, it had better be the case that goals, functions and behavioral targets are conceptually independent of semantic considerations. I think they are.

A Cluster of Objections. There are also a few other objections which, in a somewhat related manner, presume as necessary for semantic information certain conditions that my analysis appears to either ignore, rule out or simply contradict. Let me compress them in one long Teutonic formulation. Consider again our paradigmatic DOP. Neither the fact that its photocell cannot tell the difference between light and distal layout, nor the implied fact that it (and the whole system) can be fooled in various ways, nor the fact that the identity of what is semantically specified is as coarse and indeterminate as to be exploded by familiar slippery slope arguments, nor the fact that there is no representation code for encoding its semantic information, nor the apparent lack of any underlying and intrinsic intentionality, should be construed as undermining the case for assigning semantic information to DOP. Let us take them up in order.

Distal and Semantic. The first two objections sensibly assume that semantic aboutness must be distal and that, therefore, if distal aboutness cannot be secured, the system (our DOP) is not semantic. (Bishop Berkeley was no semantic system.) I agree and indeed have spent some time arguing this very point. The objections must prove the antecedent (no distal aboutness) to show that DOP is not semantic. The proof exploits an obvious fact: DOP's photocell cannot tell the difference between a distal object and a proximal input, and it can be easily fooled about the former by faking the latter.

Sensors are built with proximal scope, on the independent assumption (enshrined in their design) that the input features they register connect in some regular fashion with relevant distal objects. It is not the task of the photocell, or of perception in general, to ascertain and check on the truth of that external connection. So there is nothing special about the photocell's failure to reach beyond the input (see Fodor&Pylyshyn, 1981; Dretske, 1986).

This is a good answer but it concedes too much. The photocell is not alone, is part of DOP, and is built the way it is because it is part of DOP. It is DOP, as a perception-response link up, where the semantic information is (type) constituted or constrained (although not necessarily where it is tokened). We have seen that the factors responsible for typing the information structures (tokened in DOP) as semantic are systemic in that they reflect properties and functions of the entire DOP, not just of its photocellular (or perceptual) part. DOP, after all, opens doors for moving objects, not for light inputs, so the information it tokens had better be constrained so as to accomplish its proper function. The photocell may fail to see the difference between proximal input and distal source but that is not true of DOP itself. If the photocell keeps activating the motor center when the input is all right but moving objects fail to materialize, DOP may have to go for a thorough redesign. Such is often the fate of biosystems which have to evolve or perish when the environment changes in ways they may fail to perceive. This goes to show that the semantic information they token has system-wide significance.

<u>Fooling the System</u>. The second objection is inadvertently friendly. To fool an information system is to fool a semantic system. This can be shown in several ways. Any light manipulation which evidences, but does not originate in, the distal objects to which DOP is designed to react will trigger the door-opening behavior. Good news for DOP as a semantic system, for it shows that what is being fooled is not the photocell but DOP itself. The photocell just registers, faithfully, what it ought to register, a proximal input. The photocell does not have, nor does it care, to reach beyond that input. But DOP has and does. We have shown that semantic information is a system-wide, not a local (photocellular) phenomenon, so we have the connection between deception and the systemic significance of semantic information. The same diagnosis is true of any cognitive system, no matter how complex. Play with my brain and I will see, feel and think whatever you please to input me with. If I am semantic, then so is DOP. I am, so it is.

<u>Coarse and Indeterminate</u>. Another objection concerns the coarseness and indeterminateness of what the semantic information in DOP (or in any very simple semantic system) is about. This opens the way to a slippery slope where the information can be about practically anything, hence about nothing, hence nonsemantic. True, DOP's information about moving objects lets in quite a bunch of candidates, from bugs to cats to deans, chairs, fleeting shapes and continental philosophers. But this is all right. This is why the system <u>is</u> simple. The simpler a system is, in its design and behavior, hence in its internal functions, the coarser the semantic information which guides its behavior. Knowing the environment, the DOP's designer counted on people doing most of the moving in that neighborhood, which is why he did not bother to built in further discriminators. It is all a matter of objectives and costs. If bugs were to multiply unreasonably, we can count on the designer to fit DOP with appropriate antibug discriminators. That will refine DOP's semantic information. Natural organisms have to do this, from time to time, on their own.

<u>Representation and Intentionality</u>. The last two objections need not

detain us long, as they are going to be answered in what follows. No semantic information without representation? The surprise betrays a confusion between a way of encoding semantic information and the very nature of the information in question. This entire paper is an attempt to separate the notion of semantic information from that of representation.

Does semantic information requires an 'intrinsic intentionality' or 'mental directedness' which specify what cognitive states are about? It all depends. If by these nice words we mean teleological constraints and specific solutions (architectural and representational) to the design problems raised by those constraints, then we have began to answer these questions positively and naturalistically. If these nice words mean something else, some ghost in the machine, some irreducible subjective experiencing or some ur-biochemical juices, then we shouldn't bother.

VI Representation

The door opener and other simple systems have been said to have semantic information. Nothing so far has been said about representation. This was deliberate. DOP has semantic information about moving objects but does not represent them. The notion of semantic information must be disengaged from that of representation. There are systems which encode and act on semantic information but do not necessarily represent it. This line of analysis makes naturalist sense because (at least as a working hypothesis) it allows for continuities in the way nature handles and uses semantic information at various levels of structural and functional complexity. Semantic information is the key ontological phenomenon. Representation is one type of solution to the problem of encoding and utilizing semantic information. Architectural encoding and execution of functions is another type of solution to the same problem.

The tone in this section is going to be mostly programmatic, the pace quick, the details few. This is because, given the space, my objective is not so much to offer a theory of cognitive representation but rather to show that, and why, representation is just a manner, itself versatile, of encoding semantic information.

Why Representation? To improve semantic aboutness, one may answer, that's why. But why improve semantic aboutness? And why improve it through representation? Let us begin intuitively.

It is all a matter of what you want to do in life. Imagine you are DOP itself. You are static, simple, fond and capable of opening doors when stimulated by moving shapes (no other action, no other stimuli). What a life! And yet, one sunny day, feeling underutilized and slightly bored, you decide to aim higher. You want more action. So you ask to be moved downtown. You ask the chief designer to do a few things for you. It is called cognitive upgrading. It comes in several stages. Stage one. You ask to have some mobility along certain areas of the wall you are now going to inhabit. Your employer likes the idea because mobility gives you more semantic precision, hence a better performance for the money. Your task remains that of opening the door upon detecting certain light patterns which (for your designer) are indicative of objects heading toward the door. You are also endowed with better sensors attuned to the human shape and with a visual template with which the input is going to be compared. You even acquire a minimal form of control which allows you to defer or cancel behavior if the input does not reasonably match the template.

With such extra fun, new problems. One is this. Whereas before the change, an instant detection of the right light pattern was activating your behavioral function directly, like a reflex, now, because you are able to move, take more input in, you end up with several detections. Each detection is likely to be partial and fragmentary, so none is likely to contain the semantic information which will match the built-in template and open the door. The detections may only allow you to establish things like 'high intensity on the left, therefore a bit of a shape moving rightwards, oooops, now nothing, aha, again that bit of a shape still further rightwards', and the like. From such partial inputs, when compared to the template, you must be able to conclude, more or less, that 'That's it, composite shape has human dimensions, moves rightwards, better open that door, quick'. (As a DOP you do not get to talk like this. This is our talk as we attempt to make sense of your fancy cognition.) It is this 'conclusion', which integrates the semantic information needed to activate the behavior, as well as the way it is arrived at, that hold the key to the notion of representation.

Aspiring even higher, you escalate to stage two and desire to improve the range and resolution of your semantic information. You now ask for full three-dimensional vision. This requires even more integration of information. You are cognitively upgraded by being wired to some sophisticated devices such as edge detectors, intensity graders, and the like. Your architecture is instructed that it is along the edges of shapes and other surface discontinuities that the intensity of the reflected light changes, thus revealing three-dimensional objects. Such gadgets and others allow you to organize and integrate the stimulus information from edges, discontinuities and other such external clues into the semantic information about three-dimensional (moving and humanly shaped) objects, not just mere shapes and surfaces, as before. Your internal template is upgraded to three-dimensional human volumes. And you can do the following as well: you can now build complete encodings of information about people in motion at successive time-space points and thus follow their progress before opening the door.

Stage three. To make sure that you open the door only for people, perhaps some bigger animals and walking robots, but nothing else, you ask the designer to fit you with a second sense, hearing. You want to be able to recognize by sound how people step on the pavement while heading toward the door. Again, a sound template is built into your architecture, so you can compare it with what you hear. And you also ask that the sound track be integrated with the vision track in order to improve your semantic accuracy (the visual template may disqualify big animals whose walk may fool your ear) and have a back up in case one track breaks down.

There is now more integration and commerce with information, but there is also something qualitatively new. Vision and hearing have to communicate their outputs, to each other if they are to check on each other, back each other up, or even contribute different but complementary parts of the final semantic information. Yet their inputs are very different in nature and so are the formats of encoding their stimulus information. Their cooperation require that your design ignore the differences and find a common format of encoding, a common code.

Stage four is a bit insolent, but why not? You want the supreme social reward which you cleverly sell to the chief designer as a dramatic but useful cognitive upgrading. Just a split second before opening the door, you want to be able to shake the person's hand. Now the behavioral component becomes itself a sensor, your third type. One improvement. Coming after you have formed the complete encodings of information about people in motion, the hand-shaking acts as an interim feedback: it tells the whole process whether its earlier informational output was any good, and even allows for correction. The need for inter-sensors comparisons and hence a common code is reinforced. There is another very serious improvement. You have finally accomplished something which so far has been a bit of a convenient fiction. Your behavior is now in directly and physically related to the types of objects which are the sources of your information <u>and</u> toward which you behave as you do. This is the fulfilment of the semantic ideal. Not that before you were <u>less</u> semantic: everything that mattered in the typing of your semantic information was build into your design. But only now, with this upgrading, what you <u>are</u> is fully matched by what you <u>do</u>. You are ontologically fulfilled, your essence is your being serving your doing, Herr Doktor DOPPEL.

<u>A Working Characterization.</u> To put all this into a perspective which motivates the need for representation, let us remember where we came from. The original DOP lacked the internal resources to build up, integrate, compare, adjust to and share the semantic information it had. There was a good reason for that. Like all the other semantic infomorons we have been considering, DOP was only a perception-behavior system. Its job was to directly convert the stimulus into prearranged forms of behavior. Its architecture did the job. That the semantic policy was implemented in an architectural manner, directly and without any signficant internal mediation, makes all the difference between old DOP and new DOPPEL. The latter is quite a different semantic machine. All its aspirations, just surveyed, raise new problems as to how to encode and compute the stimuli and harness them to feed the new forms of behavior. Whence the upgrading. The end of the upgrading remains semantic, which is to allow DOPPEL to encode information in forms which render its behavior responsive to preassigned targets. But the <u>means</u> are now different. This is where representation comes into the picture, as a nonarchitectural solution to the same

semantic problem. So let us first summarize the difference before examining it in more detail.

There is semantics by architecture and semantics by representation. Both architecture and representation token semantic information, that is, both token a function from stimulus information to targets of behavior. Schematically:

semantic information: stimulus information ---> targets of behavior

The notion of representation captures a local function, constitutive of the semantic information function. Whereas semantic information is a form of encoding and using material information, representation is a form of encoding and computing semantic information. A representation is a function from stimulus information to its semantic encoding. So is the architectural form of semantic information. This means that the semantic information function can be realized under either architectural or representational conditions:

Architecture/Representation: stimulus information ---> semantic information

The first difference between the two semantics, hence between DOP and DOPPEL, concerns the format of the intentional types (concepts, recognition patterns and the like) which execute the system's semantic policy. For its semantic purposes, DOP's architecture exploits a number of natural and lawful connections, one between source and input, another between internal states tokened by the input and output, and so on. DOP's intentionality is architectural. DOPPEL's intentionality works quite differently. Instead of being executed architecturally, DOPPEL's intentional types are explicitly encoded in specially designed structures. We opted for iconic concepts or templates. The template must be activated by the input if DOPPEL is to token the semantic information which guides its behavior. For this to happen, the input must provide what the template expects. But it doesn't, at least not <u>in the form</u> that the template needs. This creates the second major difference between architectural and representational intentionality. In our story, the input is just patterns of electrical activity, whereas the template is a structure encoding a geometrical form.

How is the gap between input and template to be bridged? Through analysis or computation. This is a process which extracts from the input the information in a form which the template can recognize and use to do its job. In DOPPEL, computation is a formal or form-sensitive process. It deals with the input along its formal dimensions by analyzing and recoding its formal (geometrical) aspects commensurate with the form of the template. The result is a representational reencoding of the input in an explicit format along formal dimensions. It is as representation that the input can be compared with and matched against the template.

The representation of information is assembled in DOPPEL out of bits and pieces which indicate proximal aspects (such as light intensities and discontinuities) which in turn can be analyzed to reveal distal aspects (such as edges, corners, textures, etc.) at the source. It is the analysis of distal aspects that the resulting representation submits to the template's recognition because (in our story) the template is designed to be a replica of those distal aspects. This process of analysis requires some sort of data space where successive encodings of the stimuli can be tokened and held for a while so that the formal operations can recognize, engage and transform them into further encodings. This means that the tokens in the data space are explicitly encoded. Since to encode is to give structure or form, to treat information tokens as explicit encodings is to treat and operate on them in virtue of their structure or form. It is the formal aspect of the encoding of information which is made explicit in the data space.

What we have so far is this. Processes which operate on semantic information according to its explicit form are computations. A representation is a form computed. A representation, we can now say, is a data structure. The notion of data structure is meant to remind us of three things: first, that the business of representing information is to make explicit the form of the semantic information; second, that data structures are objects of computations in virtue of their form; and third, that, as data, such structures are relied on ('premised') in the computation of further representations and in the use of representations in more complex cognitive affairs (belief, decision, intention) as well as in action.

This last feature, of representations as data, is a very important condition on cognitive representation. It tells us that an explicit encoding of information becomes itself the object of computations whose output is another encoding, and so on. There are two senses in which representations are made out of representations. One is the compositional sense: a representation (say, a verb phrase) can be the component of a more complex representation (say, a sentence). Another is the transformational or generative sense: a representation is generated under rules from another (the way, say, an active sentence is generated from a passive one). We are now concerned with the second sense. It tells us that in a transformational process any successive representation has access only to the preceding representation, as far as the semantic information encoded up to that point in the system is concerned.

To spell this out a bit, imagine a system S which at time t computes a representation R7 out of a prior representation R6. Imagine that we freeze S at t in this R6 to R7 transition frame, and that nothing else matters. The semantic information that S has at t is encoded in R7. As far as S is concerned at t, this information originates in R6. R6 is the only source of information. For all that S knows at t, the world is as R6 says it is -- or indeed, the world may well be R6. It is R6 that is the object of S's computations, the source of its semantic knowledge and the basis of its behavior. It is thus the essence of representational cognition to treat, at each computational stage, the world as represented at that stage along dimensions teleologically worth representing and functionally worth computing.

This bring us to a final and crucial condition on representation, already amply anticipated. It is the fact that representations are analysed and computed according to their formal properties. This concerns both the formation of a single representation and the transformation of one representation into another, since both processes are governed by formal rules. To explain this, we need a terminological agreement. Recall that, at any thick level, an encoding of information is structural (or materially organized) under some thick laws. So far we have used the notions of structure and form, and their adjectives, as equivalent. No more. From now on the notions of form and formal get a strict, technical sense meant to indicate that, when it comes to representations, the structural aspects of their formation and transformation are characterized in a proprietary manner by formal scientific theories such as topology, geometry, grammar, logic, and so on. This is an extension of the earlier stated policy of letting science define how a domain of reality is shaped and regularized under relevant types and laws. In our approach to material information, we thus let various empirical sciences define the thick level (physical, chemical, etc.) at which material structures were typed and legislated.

The theoretical angle of teleology, with its types and laws, was needed then to define what counts as semantic. At that point we had a bifurcation into semantics by architecture and semantics by representation. When the teleology is implemented by architectural intentionality, attuned to exploit various distal-proximal and proximal-input-output connections, we have one sort of semantic structure, of the DOP sort. When the teleology is implemented by way of formal constraints on the formation and transformation of semantic structures, made intelligible by the appropriate formal sciences, we are talking of representation.

It is important to see that construing the formal as distinct from the architectural does not imply that the formal program of semantic representation is not executed by the system's architecture and, by implication, its hardware. The distinction exploits a familiar philosophical gambit. What counts as formal is a manner of typing the information, that is, a manner of selecting the properties which really matter for its semantic job. Those properties (or types) are now identified from the standpoint of the formal sciences because the former display features and patterns which only the latter can make intelligible and explain. The task of the architecture is to execute the formal typing, whereas the task of the hardware is to token the architectural execution.

Form is now construed as the structure of semantic information under formal constraints (types and laws) of formation and transformation. When a semantic information state has a form, it represents. To represent, then, is to encode semantic information in a formal manner. One can reasonably speculate that the reasons for the formal strictures on representation may well have to do with the very nature of the manufacture of representations, relative to quickly changing targets of distal interest, as well as with the very optimization of the execution of the intentional program of a semantic system, given what the system's teleology requires of it. We should look at it with an open evolutionary eye: since architectural intentionality can do only so much for an information system, as our DOP's limitations have shown, representational intentionality must be a step forward in several directions. Formal computation is one of them.

Semantics by Representation. The job of representation is to provide an analysis of the input by (successive) computational assignments of form which in the end can engage the intentional types of the system, be they templates, recognition patterns, concepts, rules, whatever. To put it in Kantian terms, the job of representation is to bring inputs under intentional types along formal dimensions. A representation encodes semantic information in ways which makes such computational subsumption possible and functionally effective.

It does not matter for the argument made here whether the formal program whose execution enables a system to represent is itself explicitly encoded (as a further representation) and consulted, or else is simply built-in architecturally and just followed or causally complied with. The key conditions on representation can be met in either case. In particular, neither the explicitness of the datal representations themselves (as objects of formal computation) nor the formal character of their formation and transformation depends on the formal program being itself explicitly represented or not. Our notion of how grammatical competence is formally exercised in language production and understanding is no way affected by the decision that grammars are themselves represented or not, although such a decision is bound to alter our overall theory of cognition.

Now a risky move. Representation can do its cognitive job in several formal formats, which therefore identify (with respect to form) several types of representation. My aim here is not to engage in an exhaustive survey, analysis or motivation of these types but only to establish that some important candidates are no less and no more than representations, as formal encodings of semantic information. In so doing, I want to push a bit further our understanding of what is semantic and representational in cognition. (I also have to protect DOPPEL as a genuine representer.) As these are very hot waters, I will only make a quick dive and splash and get out.

Syntactic and Topographic. We have seen earlier that digitalization does not even come close to distinguishing semantic from nonsemantic information. So, a fortiori, it cannot be a condition on representation, which is a subspecies of semantic encoding. It is true that cognitive representations can be digital, most are, but not in order to be representations but rather in order to represent in specific ways. If we want representations to be made out of discrete units (say, symbolic atoms) aggregated and operated on in syntactic ways (like in grammar or logical reasoning), then digitalization is a formal implementation requirement. The digital form of encoding will satisfy certain types of formal constraints but not others. This means that digitalization is needed to comply with specific formal constraints on computation. By contrast, the analog form of encoding may be required when the computations operate nonsyntactically, on continuous yet formally constituted patterns of encoding information.

The digital and analog properties do not by themselves type representations. They only type formats of material, presemantic encoding of information. If representations differ in respect of formal properties, the difference cannot be that between their analog and digital encoding but rather that between what the analog and digital form of encoding make possible, namely, distinct sorts of formal computations on (appropriately) distinct objects of computation. What genuinely types a representation must be the formal constraints on its formation and transformation. In this respect, we can talk of syntactic versus topographic representations. The notion of syntactic representation is very familiar from grammatical and logical theory, in particular, from the way those disciplines characterize the formation and transformation rules of grammatical and logical expressions. We can let those rules implictly define syntacticity.

The notion of topographic representation is meant to suggest continuous patterns of variation, deformation, change, and so forth, which (a) obey distinctly formal constraints defined in respectable formal sciences, such as geometry, topology, tensor theory, etc., and (b) are essentially sensitive to aspects of space and location. Although a topographic representation is likely to be analogically encoded (just as a syntactic representation is likely to be digitally encoded), the fact that the representation is of the formally continuous type does not originate in its analog encoding but rather in its typing, for example, by geometric or topological constraints.

Topographic Representation. The fact that topographic representation is formally sensitive to aspects of space and location shows semantic promise. If we find cognitive systems which token and act on semantic information in ways which are made intelligible and explained by the notion of topographic representation, then we can say that such systems are topographic representers. Are topographic representers possible? Are some also real? First, a gentle apriori push, then some empirical speculation.

Apriori, there is no reason why a representation cannot be formed and transformed in continuous patterns across space and time, under topographic constraints of some sort or another. We can imagine systems whose telos, behavioral objectives and intentional resources make topographic representation the most optimal frame of implementing their semantic policy. This would be, for example, the sort of system which must coordinate its continuous movement with the continuous scanning of distal targets in motion. The representation process could of course be digitalized and then mimicked by some syntactic encoding, so that the scanning yields only discretely structured samples the scene, but that is not the point. If there is no further reason for digitalization and syntacticization, or no resources to accomplish them, wouldn't a representation formed and computed topographically make more sense? Could this be real sense? Let us turn now a bit empirical, as empirical as philosophers can be.

To begin with, current and sharp disagreements about how vision and mental imagery represent give our line of speculation some elbow room. The syntactic view of vision has been systematically developed by Marr and his colleagues, the alternative nonsyntactic view by Kosslyn and his colleagues, among many others. No matter which view is right, it is important to note that they share the notion that vision (and the rest of cognition) is representational and computational. The same observation can be made about the competing views on mental imagery. The debate there is (generally) not about mental images being or not computed representations (they are), but rather about the manner in which they are computed and in which they represent. It is clear that both visual and mental images have semantic force and that, in addition, they represent what they are semantically about. The important conditions on representation are satisfied: both visual and mental images are formed and transformed computationally along formal lines, relative to definite intentional types, such that subsequent representations feed on, because have formal access only to, antecedent representations. The disagreement is (or ought to be) about the formal story. It is here that we run out of intuitions, a good sign that expert speculation and science ought to take over. (This should not suprise us. Common sense had almost no inkling that physical reality is mathematical. Why would it be better off when it comes to the formal character of cognition?)

A more decisive empirical step now. For years neuroscientists have studied many sorts of topographic representations or maps in the brain. [6] The basic idea is that the brain displays a good number of layers which are structured and operative as topographic maps of various areas of the sensory and motor peripheries. The topographic representation appears to have the key function of projecting spatial and particularly neighborhood relations among cells from the peripheral surfaces to the specialized formations inside the brain. The surface of the visual cortex, for example, tokens topographic representations of stimulations patterns on the retina, just as the surface of the auditory cortex tokens topographic representations (called 'tonotopic') of the neighborhood relations in the frequency spectrum in the cochlea. There are plenty of other such topographic representations for others sorts of inputs. There are clear teleological and functional constraints on these topographical maps, constraints which are responsible for what aspects of the stimulus information are being selected and explicitly encoded as semantically relevant. Moreover, the constraining is formal, so the mapping appears to be representational.

Let us see how this works by feeding the story of topographic representation into our analytic framework for semantic representation. I will use, for this purpose, Paul Churchland's (1986) example of a topographic representer, a schematic crab. (The crab is not so different in intent from DOPPEL, although it is more complex and versatile.) Its mission in life is to be able to grasp what it sees (which is why its sensory maps are projected onto motor maps), yet antecedently what the crab is going to see is (teleologically and functionally constrained by) what it can grasp and eat. The topographic maps are representations with a teleological motivation, for they are in the business of securing sensory-motor coordination. This is the evolutionary problem these maps are a solution to. (Marr would say the same about his syntactic visual sketches and representations.) Furthermore, the topograhic maps are explicitly encoded in an internal data space as data structures (representing state-space positions), and are operated on by specific computations (of coordinate transformations of points in one neural space into points in another), leading to further representational maps before turning on the behavioral engine. An important point now. Any post-input map in this circuit is build exclusively from the information explicitly encoded in a preceding map by means of formally (in this case, topologically) constrained computations. The type identity and work of these computations is revealed in a principled manner by (applied) topology, a science as formal as logic or grammar.

Is <u>this</u> really representation? Programmatically, I have argued for a positive answer. The reasons are clear. Key conditions on semantic representation seem to be met. The typing and constraining is formal in a principled manner, in terms of the appropriate formal sciences. The theoretical advantages, for a naturalist, are also evident. No abhorrent gaps are allowed. Syntactic representation does not pop up out of the blue of nature but appears, rather, as a particular solution to the more general problem of representation, just as representation in general (whether syntactic or topographic) appears as a solution to the more general problem of encoding semantic information.

Whence the uneasiness, then? Is it because whatever is formally continuous and materially encoded in an analog fashion has long been assimilated with what is physical, with hardware, whereas the syntactic, discrete and digital with what is computational, with software? Or, even more decisively, is it because showing something to obey formal constraints is not yet showing it to be anything more than material, at some thick level? After all, physical reality is mathematical because in some sense it tokens formal types and constraints. The behavior of planets obeys all sorts of fancy equations. Are we now going to say that the planets compute the equations describing their behavior? Or are we making one step beyond this lunatic idea and say that planets also represent something?

Excessive worry is premature. Even before we come to representation, planets and the physical nature in general are banished from the semantic game altogether, just as photocells and altimeters were, earlier on. This is the tactical beauty of making representation a particular case of semantic information. The worry, as I vaguely see it, lies elsewhere. Topographic representation seems to be a specific sort of solution to the semantic problem of <u>sensory-motor coordination</u> relative to not very distal targets of behavioral interest. (The bacterium is semantic but not topographically representational because, unlike the crab or even DOPPEL, it has no distally targeted behavior.) Syntactic representation, however, seems to aim further or higher than that. What we represent syntactically is not obviously or not primarily or no longer related (hence it is not motivated and constrained just by) sensory-motor aspects. The <u>what</u> determines the <u>how</u>. Syntactic representation is likely to be a specific sort of solution to the semantic problem of representing these further or higher aspects, whatever they turn out to be.

My diagnosis, therefore, is this. If we have problems admitting nonsyntactic forms of representation, it is not because we have problems admitting other forms of representation but because we have problems admitting <u>objects</u> of representation <u>other than</u> those of syntactic representation. We have been terrorized for too long by natural language and higher-level cognition in our understanding of semantics and representation. As a result, we have dissociated the latter from behavior and sensory-motor coordination. The dissociation (pushed by every major philosopher from Plato to Descartes to Frege and Husserl) has in turn shaped our philosophical intuitions. We may of course <u>decide</u> that only syntactic representation is representation because we want representation to delineate behavior-free and deeply distal aboutness. That could be a legitimate decision. It may well be the decision made by such cognitive syntacticists as Fodor and Pylyshyn. But this decision will not absolve us from explaining the structural and functional difference between the semantic policy of the bacterium and that of the crab, or between those of DOP and DOPPEL, or in general between the semantic policies of systems which react very simply and directly to stimuli and systems which take the trouble of going through intermediate stages, by analysis and coordination, along well motivated and semantically respectful formal routes, before behaving. Why the trouble? And why in this form?

VII Looking Back and Ahead

There is a feature of my account which sounds a cautionary note. One can tell a plausible naturalist story of semantic information and representation, hence of the semantics of cognition, <u>without</u> saying anything important about our real life cognition -- that is, about thoughts, beliefs, intentions, decisions and inferences. That should warn us that the semantic cognition is not yet real life cognition.

The story of the semantics of cognition, no matter how told, is ultimately a story about a designed architecture, which embodies some competence in handling information, and about how that competence types and constrains the information. The story was told here from an ontological angle. Yet at various thick and/or formal levels, this <u>is</u> the story told by the sciences of cognition. These sciences are in the business of figuring out cognitive design and competence and how they type and constrain semantic information. We did invoke behavior and teleology to explain the rationale for cognitive design and competence and hence for the very typing of semantic information. But that did not mean that the <u>tokening</u> of semantic information in real life human cognition is purely semantic, or that the tokening is operative under semantic constraints and values <u>only</u> (as in grammatical processing or logical reasoning), as the sciences of cognition may suggest and as the current philosophical consensus assumes. Our real life cognition does not token pure semantic information.

Our story of naturalized information ends with representation. This is its semantic end. Most philosophers and theorists of cognition think that the semantic end is <u>the</u> end of the informational story of cognition. It isn't. They think that if we can figure out the semantics of cognition, we can figure out cognition <u>tout court</u>. One figuring out does not follow from the other, as I have argued on a number of occasions (Bogdan, 1985, 1986, 1987, 1988, and forthcoming). The semantic typing of information, up to and including representation, does not reflect (because it is not yet subject to) the sorts of facts and constraints which are responsible for the type of information handled in believing, planning, inferring or acting. There is a third form of information, mental information, whose typing reflects factors operative in genuine central cognition. Mental information is semantic information doing real cognitive and behavioral work. The laws and constraints on this work, hence the typing of the information doing it, can be shownn to be essentially different from, and irreducible to, those responsible for the typing of semantic and material information.

The fact that semantic cognition is not yet full cognition should also warn us that the prevalent scientific and philosophical models of various aspects of cognition, such as thought, belief, intention or inference, have something artificial and seriously incomplete about them. Some of us have felt this all along. Clever experimental methods in psychology and superimaginative and science-fictional feats of philosophical abstraction are needed to extricate the semantic competence and the information it types from the natural flow of information in real life cognition. My objective here was to retrace, identify and make some ontological sense of what exactly is it that is thus extricated. I have concluded that it is just the semantic prelude to genuine cognition. [7]

Department of Philosophy
Tulane University
New Orleans, LA 70118

FOOTNOTES

1 The notion of dominant format of encoding is modeled after Dretske's notion of proper dimension (Dretske, 1981, p. 161).

2 I am making, for the moment, the simplifying assumption that the source of material information is also the reachable target of a system's behavior because it helps the exposition. This need not be so and very

often is not. There are semantic systems which just react internally without behaving externally. More on them later.

3 The idea that appeal to behavior is needed to warrant the attribution of distal semantic significance to cognitive states has been forcefully argued by Daniel Dennett in his important (1969). Also implied in Dennett's conception is the system-wide ('afferent and efferent') character of the attribution of semantic information. I doubt, however, that Dennett is committed, as I am, to the reality of semantic information in the guise of a type of structure in the world, prior to and indepenedent of any attribution. Moreover, Dennett's position may at times be epistemological and vaguely behaviorist in a way in which I hope mine is not. See section V below. There are other differences as well but this is not the place to consider them.

4 Fred Dretske (1986) has convinced me that bacteria deserve semantic citizenship. His distinction betwee natural and functional meaning is partly reflected in the present distinction between material and semantic information.

5 Although many people have objected to various claims made in this work, three in particular have forced me to allocate enough space to answer them. This is what happens if you see them every Friday. With his doux sourire, Francois Recanati has made the circularity charge, Dan Sperber has shrewdly defended proximal semantics (which I dealt with earlier), while Pierre Jacob has queried me about the naturalist notion of fooling a semantic system and about the slippery slope.

6 The term, but not the entire meaning, of 'topographic representation' is borrowed here from neuroscience. A detailed survey of the neuroscientific work on topographic representation can be found in Patricia Churchland,1986, chs. 3 and 10. Paul Churchland (1986) provides an imaginative model of a topological representer, also summarized in Patricia Churchland's book.

7 Work on this project has been financially supported in part by Tulane University through a summer research grant and then a blessed but finite sabbatical leave. I register here my thanks and appreciation. I

have started this work in New Orleans, where some early ideas have been read and even taught, and continued it during the leave in France, with presentations in Paris and Cerisy-la-Salle, in morose and charming Normandy. Many friends, colleagues and students have asked tough questions and made important criticisms and suggestions, for which I am very grateful. I am thinking of the members of my Tulane graduate seminar on mental representation and of my philosophical colleagues in New Orleans, particularly, Harvey Green and Carolyn Morillo, and of my good hexagonal friends Daniel Andler, Dick Carter, Pascal Engel, Pierre Jacob, Francois Recanati and Dan Sperber in Paris. At a further remove, the work of Dan Dennett, Fred Dretske, Jerry Fodor and Zenon Pylyshyn have been a constant frame of reference and learning. During most of my work on this paper, Pierre Jacob has been such a wonderful copain. I dedicate the result to him. Salut Pierre.

Replies to Commentators

RADU J. BOGDAN

As the commentators have noticed, my essay shows the symptoms of being the programmatic roof of a building under construction. Work goes on, things are added, removed or improved upon, but the top is still representative of the original blueprint. The replies reflect a bit of all this.

I. TO FRED DRETSKE

Fred Dretske is the information man in contemporary philosophy. I have learned a lot from him. As this occasion shows, I still do. Dretske's focus is on my move from teleology to the semantics of cognition. He finds two main problems with it, both equally damaging. One is that the move is not naturalist, the other, that it is circular. (Deep breath.)

I begin with the naturalization problem. Dretske and I appear to take different routes to naturalization. He takes naturalization to entail both type reduction and genetic explanation. I take naturalization to establish the nature of something in terms of objective types and laws. It is an empirical (not a logical) question whether those types and laws can in turn be reduced to, or genetically explained in terms of, some other types and laws. It is thus an empirical question whether, with respect to types and laws, the world is organized in levels or not. Either result is consistent with naturalism. One need not go reductive or genetic to prove one's naturalist credentials. One has only to figure out the nature of things.

Dretske may be construing naturalization with a critical eye to circularity and mysticism. He is strategically right. There is no better defense against both mental diseases than explaining things reductively and/or genetically. The question is, Do I need this defense? Is my account caught in the circle of semantics? And is it clouded in the mysticism of teleology talk? Obviously, what I said in the paper (that it isn't and it isn't) has not helped much, so this is another try.

The circularity problem now. The move from teleology to semantics rests on the notion of goal. Is this notion semantic? Dretske thinks it is. So do many others. I do not. The semantics of cognition is ordinarily construed in terms of concepts and attitudes (beliefs and the like). If the semantics of concepts and attitudes were shown to rest on the semantics of goals (which is not my project), we would have not circularity but explanatory progress -- not unlike the progress made by the Gricean explanation of the semantics of language in terms of that of cognition.

But this may be only local explanatory progress. Dretske's real objection seems to be that, by going from the semantics of something to the semantics of something else, my overall project is circular. I do not think it is, but the deeper argument, barely sketched in the paper, takes us into the foundations of biology. All I can do here is to outline its direction.

My notion of goal should not be understood, with Dretske, as the satisfaction condition of a desire. A goal, for me, is not an end in view, or the object of a desire. Nor is it the desire itself. Indeed, a goal is not an occurrent entity or state of any sort, internal or external to a system. A goal is a type of constraint on an organism's organization and behavior. A vital goal can be thought of as an instruction encoded in the genetic (and preintentional) program directing the organism to bring about some condition. Goals can bring about biodevelopments (growing wings) or specific conditions (warmth). The story then is one of an instruction (the goal), the execution of the instruction (via cognition and behavior), and the outcome of the execution (the condition satisfying the goal).

The goal satisfaction relation need not be semantic. In most organisms it is not. An instruction need not be about the outcome of its execution. The instruction to reach Fred's office in that dreamy land of book writing (right on Stanford, then right again on Junipero, left, up the hill, etc.) need not contain anything that is about the office itself. In general, there need be no sign or representation, in the genetic program, of the condition the organism is instructed to arrive at.

I haven't yet said anything about goal execution. It is at this stage that the semantics of cognition comes into the picture, not earlier. An organism's behavior must be guided in ways which ensure the satisfaction of its goals. So there must be constraints on how information does the guiding. This is where I locate the roots of intentionality. Concepts and attitudes execute goal instructions by typing the information in the semantic forms needed to steer behavior toward the goal satisfying aspects of the environment. It is as intentional executors of goal instructions that concepts and attitudes are about targets of behavioral interest. Desires and plans, in particular, may project states of affairs in the form of ends in view. They seem to be Dretske's candidates for goals. They are not mine. Desires are semantic all right but are not genuine goals qua basic instructions. As a matter of intentional execution, desires typically constitute interim objectives ("derived" goals) for the actions required to satisfy the basic goals.

Dretske's next objection invites a distinction between bioteleology and teleology in general. His query is whether bioteleology is logically necessary for information being semantic? I think not, since I see no logical relation between bioteleology and semantics. Other forms of teleology (say, robotic) could constrain a semantics. Is then bioteleology in some natural sense necessary for semantic information? I am moved to deny even that by an extended functionalist argument. A narrow functionalist argument disengages the essence of software from the accidentality of the hardware. I want to push the disengagement even further, in two stages.

Stage one separates the history of an intentional program from its nature. Even when history shapes nature, as it does in biosystems, the explanation is not necessarily transitive. Explaining from evolution is not explaining from bioprograms, even though evolution shapes bioprograms. Why and how a biosystem comes to be semantic (through evolution) must be distinguished from the aspects (say, genetic program) in virtue of which it is semantic. I agree with Dretske that an evolutionary story of biointentional programs is essential to a naturalist account of cognition. Still I insist that it is the nature of the program itself that is responsible for the semantic typing of the information. Since a biointentional program serves goals, the first stage of my argument assumes biological goals but brackets out their genesis.

In the second stage, it is the very nature of the goals that is left out. Information can be typed semantically whether the responsible teleology is biological or not, whether, in particular, its goals are biological or not, primitive or acquired, natural or simulated, intrinsic or artificially implanted. All that matters is that the (telic) constraints on semantic information be operative in directing the system to type the information in ways which are systematically responsive to distal properties. A system remains semantic even if, after being designed with goals and behaviors in mind, it is somehow deprived of both and left only with its intentional program. When activated, the program still types information semantically. It is the fact that goals and behaviors initially constrained the design of the intentional program, not their current presence, that makes a system semantic.

Dretske is right to doubt that I substitute explanatory role for real ontological beef and take the path of instrumentalist perdition. I have been saying all along that semantic information is typed and legislated into real existence and causal efficacy by real types and laws, whatever explanatory use we make of them. But he is wrong to think that I think he is wrong, so that he must think I am wrong. (Got that?) As editor, I read Dretske's paper on 'Misrepresentation' [1] more than four times. Each time I was happily reinforced in the view that Dretske sees a bioteleology of goals, needs and functions as being responsible for the presence and operation of semantic information in an organism. So I thought we agree on some key moves from teleology to semantics. Now Dretske says we don't. So I am myself unsure as to what his view is.

When unsure, deconstruct. So here is a tentative deconstruction of our diagreement. When Dretske says we disagree on the 'derivation of meaning (or mental content) from material information', he may well be construing (i) the notion of derivation in a strict type reductionist sense (see the first part of my reply) and (ii) the teleological constraints on that derivation with the most solemn respect for their actual biohistory and bionature (see the second part of my reply). We seem to agree on the naturalist need to derive semantic from material information under teleological guidance but disagree on the nature of both the derivation and the guidance.

II. TO DAVID ISRAEL

David Israel's thoughtful and detailed comments touch on most of the key issues raised in my paper. Two in particular stand out and need further discussion: my structural account of material information and the formal angle on representation. And then there is the thermostat, of course.

Structure, encoding, form. My notion of material information is meant to be very minimal. Its structure is an organization of material units (atoms, cells, etc.) under appropriate constraints. Like information, the notions of encoding and form were deliberately downgraded and used as synonyms for structure to emphasize a common origin and nature. The notion of semantic information reflects the fact that material structures and their causal effects are constrained in ways which are additionally sensitive to some distal or abstract equivalence classes of material properties and relations (such as eggs-to-eat or critics-to-reply-to).

This being said, when Israel notes the oddness of saying 'that the information about the viral infection is encoded in the spots', I note the same. About, for me, marks the presence of semantic information, which spots cannot possibly carry, for they are not constrained to do so; they only carry material information from the infection. Yet when

Israel says that the tree's rings carry semantic information, in propositional form, about the tree's age, I am puzzled. It must be concrete growth features (not age) that cause the rings. But the information relation here is no different from that between the infection and the spots: it is only material information. [2]

To ease us now into the vexing matter of structural acknowledgment and the Aristotelian sin of structure transference, remember that even the material information in a tree's rings or in anything else is subject to level and type specific constraints. Since the world displays a variety of levels and types, there must be some level and type underdetermination of information. A physical interaction may fail to have biological expression (level underdetermination), just as some inimical bioprocess may fail to trigger the right immune response (type underdetermination).

Does all this presuppose, as Israel fears, some vodoo Aristotelian notion of transfer of structures from sources to receivers under some formal (isomorphic or homomorhic) constraints? Does information transmission generally require structure preservation? That would be too simple and beautiful to be true. So it is not true. All I am committed to is the idea that a structure at a source is causally responsible for generating a structure in a receiver under receiver specific constraints, which are level and type specific. There is no need for the source and receiver structures to be alike in any respect other than being material. One must cause the other. That is all. The constraints on the receiver decide, as it were, whether the information to be tokened conforms to their level and type, or not. This in turn affects the form in which the information is structured.

Representation. I am still dialectical about representation, so let me begin with an innocent certainty. Representation must be a superior format of encoding and processing semantic information. But in what sense superior? A basic insight, which I think I share with Israel and others, is that we must find the rationale for representation in a policy of semantic encoding and processing which is not simply the result (as architectural semantics is) of directly embodying patterns of sensory discrimination aligned to fine tuned behavioral responses. What makes representation possible are new architectural capabilities (such as short and long term memory, control) and new constraints on the combinatorial formation and transformation of the resulting structures.

For a number of reasons, this insight has so far favored (what I call) the syntactic format of representation at the expense of the pictorial or geometrical format. A priori at least, this need not be so. We can imagine systems forming and transforming geometrical representations in a compositional way from basic symbols (points, lines), under appropriate rules, by helping themselves to items in long term memory and explicily tokening the results step by derivational step in some short term memory.

The problem that gives me dialectical fits is that the syntactic bias reflects a principled decision to identify representations with semantic structures fixed and utilized centrally and penetrably (as Fodor and Pylyshyn believe and phrase it), qua thoughts or beliefs. If representations were such structures, they may indeed have to be syntactic. Their modular generation must then be counted as an exclusively architectural, not representational affair. Israel seems to buy this idea. I might buy it too, at least on the even days of the month, were it not for one major implication which I find troublesome. The animal implication. Since animals do not seem to engage in centrally penetrable cognition, it would follow that they do not represent. But that doesn't follow. We too form representations which activate behavior rigidly and unreflectively, without attitudinal mediation. If this is good for us, why not for them? One reason it is good for us (to directly and very quickly align action to input, relative to distal aspects of major interest, in semantic ways which only representations can capture) must be a reason it is good for them. We are animals.

So I suggest we make room for representation between its (modular) formation and its utilization and THEN let the manner of utilization decide the format of representation. When the utilization is centrally attitudinal, the representation is bound to be syntactic, if it is to be general and abstract (as in this paper), creative (as in natural language), projective (as in desires or plans), and negatable (as in 'this can't be true'). When the utilization is peripheral and direct but requires explicit semantic structures, deployed in formal ways and rigidly fed into appropriate behaviors, then nonsyntactic formats are possible and often likely. [3]

Why, asks Israel, do I characterize the representation as formal?

What is to be gained? A new level of abstraction and explanation, I answer, comparable to that acquired by ascending from the behaviorally motivated functional architecture (of the thermostat) to the teleologically motivated semantic architecture (of simple organisms). [4] To serve representation, the semantic architecture must be constrained in ways which can be called (more explicitly) formal^{sem} , i.e. formal with respect to semantic distinctions.

The intuition is not simple, its explication is even less. My initial move was from explication to intuition. The intuition was that representation is semantic information encoded under formal^{sem} constraints. Then the explication took over by letting particular theories of such constraints implicitly define what is formal^{sem}, for various semantic domains such as visual space, linguistic meaning and so on. [5] Let me now briefly try the other direction and meet explication half way.

An architecturally semantic system is designed to exploit causal interactions by treating its internal states (effects) as natural signs of some proximal states (causes) which in turn are natural signs of distal states (still prior causes). These sign-signified connections are so general as to make the semantic aboutness quite coarse. In a representational system the sign-signified relations are preserved, for its semantics must be implemented by architecture and hardware, but they get layered, organized and read quite differently.

We are told that at the most peripheral level of vision the internal signs states (retinal patterns) signify precise types of input states (intensity changes). At the next level, the sign states are edge segments, boundaries and the like, and their signifieds (at the inferior level) are now discontinuities of intensity in the retinal image. And so on. This is not just a more complex and powerful architectural semantics in which signs pile up on signs in some functional sequencing which exploits a natural semiotics. How all this is done is equally important because semantically new and significant.

Some signs are primitive (symbols), others are derived (expressions). The latter are generated under rules in combinatorially constructive ways from simpler ones. The theory of the semantic domain in question can explicatively map these structure generating patterns onto formulae and their derivations under semantically respectful rules of formation and transformation, relative to primitive terms and axioms. This is the gist of what I take to be formal^{sem} about representation, whether in vision, language or thinking. This is still the last word from Delphi (I mean Boston).

Department of Philosophy
Tulane University
New Orleans, LA 70118

Notes

1 See Radu J. Bogdan (ed), Belief, Oxford University Press, 1986.

2 Which brings me to the unavoidable thermostat, easily the most nightmarish candidate to semantic competence ever invented. Israel says the thermostat is semantic, I say it isn't. He says it is semantic (in my sense) because the constraints on its semantic relation to environmental targets of interest (ambient temperature within limits) are reflected internally in how the thermostat is built and calibrated: in particular, the metals comprising the bimetalic strip are said to meet the semantic constraints. The reason I think they don't is that the internal embodiment of the semantic constraints is not at the right juncture. (In a different form, the point was first made by Fred Dretske in his (1981)). It comes too late in the process, at the motor stage. To make the thermostat semantic, the internal embodiment of the semantic constraints should have been present in the very forms in which the thermostat registers the incoming information. A selective behavior simply associated, via control, with some registration values does not make a system semantic. Not in my book. To be semantic, a system must have the guides to its selective behavior built into how it registers and hence organizes the incoming information.

3 I chose 'topographic' over 'topological' on pure Greek grounds, to emphasize the lack of any (syntactic) logos and the presence of mere graphe in that format of representation. A number of workers in neuroscience seem happy with topographical talk of representation. But of course no terminological decision in this business is without cost. If pressed hard, I might opt for 'topoactic' to emphasize the space-action

connection with an unprecedented, hence uncontroversial word.

4 Although based on Pylyshyn's (1984) now classic analysis, my notion of functional architecture is much more liberal than his. Pylyshyn's notion refers to an architecture designed to run a representational semantics in syntactic format. Mine also allows architectures designed to run a nonrepresentational semantics as well as a representational but nonsyntactic semantics.

5 Israel says that topology, logic or grammar have nothing substantive in common. This may be so, generally. But my point was not general at all. It was specifically confined to how such theories (or appropriate fragments thereof) realistically reconstruct the encoding and processing of semantic information. Nor was the point pertinent to the theory of formal systems, methodological rigor or mathematical definability.