

DANIEL ANDLER

## THE NORMATIVITY OF CONTEXT

(Accepted December 23 1998)

### INTRODUCTION

This paper attempts to show that context is normative. Perceiving and acting, speaking and understanding, reasoning and evaluating, judging and deciding, doing and not doing, as accomplished by humans, invariably occur within a context. The context dictates, or at least constrains, the proper accomplishment of the act. One may construe this undisputed fact in a naturalistic way: one can think of the context as a positive given, and of the constraints it creates as constituting a natural fact. Whether the act is carried out in conformity with these constraints is then a mere matter of correct functioning of the cognitive system. However, I argue, this is not the only, nor the more plausible way of considering the matter. The context is not a determinate function of situation and task, nor is the outcome of a task a determinate function of a given context: context choice and contextual constraints are irreducibly normative. The norm they obey is *sui generis*, and goes under the (disreputable) name of intelligence.

The paper divides into two parts. First the stage is set: context and related concepts are circumscribed (§1); varieties of doings are distinguished (§2); distinctions are made between ways in which contextual phenomena arise (§3); a methodological predicament is described, having to do with the infinitary character of context, whose solution seems to require some form of normativity (§4); lastly, two forms of normativity are distinguished (§5). My thesis is expounded in the second part: I consider how a given context constrains a given task (§6); how a given situation gives rise to a context (§7), and in both cases I show that the constraints are not of the weakly normative kind, that they irreducibly involve,



*Philosophical Studies* **100**: 273–303, 2000.

© 2000 Kluwer Academic Publishers. Printed in the Netherlands.

in the strong construal, some norm; I give reasons to prefer the old-fashioned term intelligence to the more respectable one of rationality (§8); finally I outline a possible constructivist escape route for naturalism (§9).

This is clearly a book-length project, and this paper cannot purport to be much more than a sketch.

## I. THE SETTING

### 1. *From States of Affairs to Contexts*

Many words are used in the literature in connection with the idea that the outcome of a given task – typically, giving meaning to an utterance – depends not only on the object of the task, officially and narrowly construed, on what the task is *applied* to, or is *about* – typically, the utterance – but on other factors as well, which in some sense surround the accomplishment of the task. Beyond the terminological question lie difficult ontological questions which we shall shun as much as we can, although they have an obvious bearing on our enterprise.

We shall thus pretend that fixing terminology is a simple though necessary task. I assume throughout an agent considered at a given moment of his activity, where being active means not much more than being conscious: for such an agent, immobility, whether corporal or mental, may, if at all conceivable, be activity. The world at that moment, agent included, is, or constitutes, a *state of affairs* which includes every objective entity in existence at that moment of time, and every relation between entities holding then. As nothing crucial in this paper rests on such a construct, we don't need to worry too much about its mode of existence. Strictly included (in some sense) in the state of affairs is the set of entities and relations which may have bearing on the agent and his activity. This I call the *situation*. The situation as I think of it here typically includes features of the environment to which the agent is potentially sensitive, as well as permanent features of his corporal and mental make-up; it typically does not include radiations and fields of forces which have no effect on him, whether corporal or mental, nor entities whose existence cannot be as much as conjectured by him or sensed by some organ or subsystem within him, nor correlations which he is cognitively

incapable of tracking. Admittedly, very hard problems lurk in the vicinity of this concept of a situation, which I shall again ignore at my own risk. Together, states of affairs and situations describe the objective make-up of the world, respectively with no centering and with centering on the agent.

The subjective counterpart is what I mean by context, roughly speaking. Not so long ago, 'context' referred restrictively to 'the structure of a connected text regarded in its bearing upon any particular passage that constitutes it' (*Oxford English Dictionary*). Nowadays its extension is much wider – so much so that a new word had to be coined for the traditional usage ('cotext'). However, to many writers context still calls for text in some suitably generalized sense: they think of an agent as an information-processor, a system with inner states which takes in input and produces output. Input is what they see as the generalized form of text. Two situations then present themselves. First there may be a natural partition of the informational flow, with a central, naturally or canonically distinguished, subset on one hand, and all the rest on the other – as when the process under consideration is the operation of a well-defined perceptual system, or when in addition a natural notion of 'now', a processing window, can be defined: the context is then often taken to be the information not handled by the perceptual system in question, or/and the information which does not belong to the temporal window constituting the 'now' of the agent or of his/its appropriate subsystem. In the other sort of case, there is no clear-cut partition of the information flow, and then one has to rely either on the notion of attentional focus, or on the idea that the agent's goal determines something like an endogenous 'text', a plan, whose execution is sensitive to exogeneous variables, forming the 'context'. Alternatively, the 'text' can be taken to be the total informational flow, and then the 'context' is the set of non-informational factors (this view relies on a robust distinction between what is informational and what is not; whether one such distinction is available or not I shall not presume to discuss).

There may not be anything to choose between: there may be as many special notions of context as there are types of situations, in the sense just outlined. On the other hand, one very general sort of case does not seem to belong to any of those types. That is the

case in which there is no plausible candidate for the function of ‘text’ or ‘input’, however generalized, whether exogeneous or endogeneous. Such is the obsession of contemporary cognitive science – following the steps of utilitarianism in philosophy, and of what I have called elsewhere ‘panproblemism’ (the notion that everyone, and possibly every moving or evolving thing or system, is constantly in the process of solving one problem or another) – with goal-oriented behavior, that the notion of an agent both doing nothing and devoid of a well-defined goal, has disappeared from view. To one, such as the present author, who does not share the panproblematical world view, it is important to leave room for the idle agent.

Together with the hope that something worthwhile may be abstracted from all the goal-directed situations involving context, the notion that there may be non-goal-directed context-involving situations motivates the following ultra-general “definition” of context. We shall take *context* to be whatever the situation becomes in the light of the agent’s expectations. ‘Becomes’ refers to the observer’s move from the objective to the subjective viewpoint, not to a causal or temporal relation. In Gibsonian language, one might think of the context in my sense as something of the nature of the set of affordances a given situation provides to a given agent. Note that although characterized by reference to an agent, the context is not subjective in the strong sense of being essentially private: although it is the situation as it appears to the agent, an objective description of the context is at least theoretically conceivable. In particular, usually a large part of the context is directly inherited from the situation, with very little perspectival contribution (so that in many cases, not much harm is done by using ‘situation’ for ‘context’). However I shall not go as far as saying that the context is a natural unique result (a function as it were) of the situation and the agent’s expectations, first because of the vagueness of the notion of expectation, second because I have already announced a contrary claim: although dependent on situation and expectations, the determination of context is subject to a norm.

One final distinction needs to be made. In the admittedly (both numerically and conceptually) important cases which involve some form of text, whether in the literal or some extended construal, there is a sense in which text and context enter in a codetermining relation.

For instance, in Sperber & Wilson's Gricean relevance-theoretic model of communication, when an agent interprets some speaker's utterance, the text, understood as the linguistically encoded information present in the utterance, brings out a context appropriate for inferential completion, and the context in turn allows the correct completion of the text. Similarly, take an agent pursuing the goal of winning a game of chess; as his opponent completes her move, a context appropriate for the understanding of the situation brought about by the move becomes perceptible to the agent and allows him to interpret the move and react optimally. We are thus led to an apparent contradiction: context seems to be both solely dependent on the agent's expectations, prior to the registering of some text, and brought about by, and dependent on, the text. There is however no contradiction, merely two notions going under the same name. So let us call the context in our first sense, *broad*, and the context in the second sense where it is text-dependent, *narrow*. (I shall omit 'broad' whenever possible; 'narrow' will be the non-default meaning.)

## 2. *Activity, Tasks and Texts: Active vs Reactive Minds*

We have set the stage for a succession of extensions of the ordinary use of the notion of context. The early steps of this process seem rather easy to accept, and have in fact been widely accepted. In the initial situation, 'text' was meant literally as a string of meaningful sentences *and* text and context were homogeneous, so that context really was text, some more or less extended fragment coming before or after the text under consideration. The first extension, performed long ago, consisted in dropping the second condition: context no longer was restricted to the text that came before or after; it could for example include the circumstances surrounding the production of the text, whether written or uttered. The second extension operated on the first condition: phonemes, for example, became acceptable as text, as well as, more generally, any modality-specific perceptual stream. The third step, mentioned earlier, was to forego any reference to a given modality, and to think of text as any kind of incoming information. Finally, context was taken to be anything at all, whether informational or not, which impinged on the processing of the input. At this point cognitive science was ripe for controversies: classical

AI, for example, was taken to task for producing context-insensitive programs, based, according to critics such as Hubert Dreyfus, on a misconstrual of human cognitive capacities; moreover, claimed the rival school of connectionists, a very different sort of computation could exhibit context-sensitivity. The point of this historical reminder is simply to indicate that the generalized notion of context is by no means new in cognitive science (nor in philosophy of mind for that matter).

However I have already indicated my intention to extend the notion even further, and this requires justification. Context without text of *some* sort would seem to be a vacuous concept. So why insist on having it? The answer lies in two simple considerations. First, there exists nowadays a common-sense meaning of ‘context’ which does not rely on prior identification of some “text”; it is often taken as synonymous to ‘situation’. Because it is important to distinguish the objective from the subjective kinds of world, my choice has been to keep ‘situation’ for the former; ‘context’ may then seem a natural choice for the latter. Appeal to loose, everyday parlance is however not sufficient. One needs to argue in favor of some degree of continuity between the notions of context *cum* text, and context *sine* text.

My second point addresses this very issue. A mental life is assuredly occupied in part by episodes of “dealing” with well-defined stimuli: recognizing faces, interpreting utterances, identifying approaching footsteps, reading the X-ray shot of a broken limb, shaking the hand which is coming forward, solving a problem stated on the sheet provided by the teacher, and so forth. But there is more to our mental lives than those episodes, and there is more to those episodes than meets the eye. On the one hand, our minds are active, not merely reactive, and their best efforts are given to discovering, not what to do with or about a stimulus, but what to do *period* – what to say, how to behave, where to go, what to think about politics or our desire to get rich and famous, whether to call up a friend, etc. Our mental lives are driven, in part, by our own initiatives, not solely by stimuli. Now on the other hand, even our handling of a stimulus involves initiative on our part; however automatic it appears, it is no mere release of a one-step routine. We X the stimulus by Y-ing. We turn our attention in this or that

direction. The incoming “text” generates inner activities, which in turn constitute fragments of endogeneous “text” to be “processed” further down the line. Thus, appearances notwithstanding, there is something in common between episodes with text, and episodes without: in all cases, there is taking in, coping, doing something about the situation. In all cases, we bring our expectations to bear on the situation, thus forming what I have proposed to call context. This is not to claim that there are no differences between highly automated, stimulus-driven tasks such as face recognition, and the more reflective and open-ended tasks such as deciding whether to take a walk, change jobs or propose marriage. Rather, there is a continuum stretching from one extreme to the other, and while “text” (suitably generalized) is definitely a relevant variable, “context” is not.

### 3. *Varieties of Context Problems*

Reflection on the issue of context has been heavily influenced by the consideration of some important special cases, and any novel attempt requires a preliminary clarification of the scope of the available approaches. Two classes of cases have given rise to interesting proposals: pattern recognition, and sentence meaning or utterance interpretation. Although context looms large in other areas, such as reasoning, there are no important general lessons to draw from the attempts to account for contextual effects in those areas,<sup>1</sup> and I shall leave them aside.

Very little will be said about pattern recognition (or classification), because the interpretation of the partial successes of neuro-computational approaches in creating context-sensitive models is in no way obvious, and would take us too far afield. But it does invite a distinction which will be helpful in assessing the significance of attempts in other areas. The distinction is familiar enough in a wider context: a functional dependence can be considered either as an ideal, Platonic relation, or as an implementable processual (operational, mechanistic) schema – in other words, a function can be (merely) a mathematical entity, or (also) computable, hence mechanically computable. Connectionist models of context-sensitive processes provide a processual account of the phenomenon: they don’t say very much about the dependence itself, conceptually

speaking, but they show how it can manifest itself, mechanistically, as a causal relation.

When it comes to the area where context-dependence has been studied the most, *viz.* sentence meaning and utterance interpretation, one finds, developed side by side by different people, a Platonic approach and a processual approach. Platonists such as Kaplan, Perry, Recanati,<sup>2</sup> Travis or Searle, are interested in the precise conditions under which context (or rather: certain kinds of contextual clues) contributes to the meaning of sentences (or utterances), in the very idea of such a contribution, in the extent of the phenomenon, and in the consequences which it carries for the notion of meaning itself. For processualists such as Fauconnier or Sperber and Wilson these are not the main questions, and they focus instead on the actual, material way in which an agent can correctly interpret an utterance by using the appropriate contextual clues. What they are after is not meaning, but communication, *i.e.* transfer of information. They want to understand how context contributes to that task. In particular, they naturally feel an obligation to take a stand on what context materially consists in, on what stuff it is made of – not a source of concern for the Platonist, who is content to regard it as a warehouse of meaning.

The two projects would seem to be connected, and at least subject to a mutual compatibility constraint; but in fact they don't sit comfortably together, and there is a reason for this: Platonists are by definition agnostic about naturalization; processualists are by trade naturalists. By and large, philosophy of language is Platonic, philosophy of mind naturalistic. The present author, as will be seen, though more of a philosopher of mind than of language, is wary of premature attempts at naturalization, and in fact suggests that a closer look at the Platonic relation of context-dependence justifies such caution.

Utterance interpretation, although straddling both fields, tips heavily towards philosophy of mind. Let us take a closer look at that paradigm case of contextual dependence, which has been sitting at the center of attention since the development of Relevance Theory by Sperber and Wilson, treading in the steps of Grice. Nothing need be said here about the *solution* they offer. But what is their *problem*? To explain how a hearer exploits contextual information to

retrieve the informational intention which has motivated a speaker to produce the utterance addressed to the hearer. What is of interest to us here is the place which this problem occupies in the space of context problems. It is characterized by four traits: (i) there is a clear-cut division between text and context (and a crisp definition of both: the text is the linguistic content of the utterance, the [broad] context is the pool of “mutually accessible” information); (ii) the text is symbolic, i.e. it is semantically evaluable; (iii) the text is intentionally produced and addressed to the agent; (iv) the agent has a well-defined, predetermined task, which is to understand what has been said.

This quadruple “plus” in the logical space of cognitive processes makes linguistic communication a very, very special case, and the corresponding context problem a very, very special problem. It could well be, then, that the solution offered by Relevance Theory, and the naturalistic moral which some might be tempted to draw from it, have limited bearing on the other kinds of context problems. On the other hand, it could also be that consideration of other cases might bring to light aspects of context in linguistic communication which Relevance Theory doesn’t do justice to, or simply doesn’t address. The property which we are in fact about to discuss may be one such aspect, if I’m right, but the general point should be borne in mind.

#### 4. *The Infinitary Character of Context*

Naturalists don’t believe that we can directly apprehend or represent infinite objects: a natural cognitive system such as ours, because it is finite, cannot accommodate infinite representations. Thus for transcendental reasons those objects which we grasp are finite. Or rather, finitely generated hence finitely describable: a finite system *can* accommodate a finite set of generative rules, and is thus able to “grasp” in some (dispositional) sense, finitely generated infinite objects. On the other hand, objects which admit of no finite description or generation can hardly be finite.

Context, I claim, is infinite in some sense. As I am not willing to commit myself to a precise definition of context, I would be hard put to *prove* that context is infinite in the sense of being composed of infinitely many elements – elements of *what?* But in an intuitive

sense, except possibly for degenerate cases, a cognitive process in the wild involves an infinite set of relational facts, almost all of which are potentially “part” of the context. To be sure, not all of these facts need be represented all at once: they are only *available* for representation, or to use a fashionable phrase, *poised* for inference and other kinds of processes. But they do belong, in a loose sense, to the context nonetheless – the looseness is the reason to speak of the *infinitary*, rather than simply *infinite*, character of context.

The roughly set-theoretic argument which has just been used may not convince every reader. So let us put forward another consideration. Let us ask whether a context is ever *saturated* in the sense that it could not conceivably be included in a proper supercontext, one which would have every feature of the original context, and at least one more. The answer seems to be no: what would prevent a fly from descending on the agent’s nose? And if not a fly, a snowflake? And if not, why couldn’t a fire break out? And if not a fire, the plague? Etc. A positive proof seems out of reach: it is *conceivable* at the present level of generality that some context would actually preclude the existence of any fresh entity, the occurrence of any concrete relation or event not already present. But the burden of the proof, it would seem, rests on the defender of such a possibility: it is he who should give us an idea of what a saturated context might look like. Now if we accept the thesis that every context can be immersed in a richer context, it follows by induction that we can define an infinite sequence of ever richer contexts. The limit of which would then constitute an infinitary context. But now any member of the sequence has among its features the absence of the features which distinguish it from the limit: *no fly, no flake, no fire, no plague*, etc. It follows that *every context* is infinitary.

Although objections will most likely be raised against this argument as well, it is doubtful that many readers will want to seriously disagree with the thesis that in normal, natural circumstances, the context is generally infinitary in some sense, or at least so large as to make it impractical (as well as phenomenologically unrealistic) for every one of its features to be taken into account in the course of the process under consideration. In the terminology of §1, the useable part of the (broad) context is but a subset, within which what

we have called *narrow* context will be ultimately selected: every theorist who has worried about context and related problems (such as the frame problem in artificial intelligence) has had to face the *demarcation* question: how does the agent, or the “system”, cut out the narrow context in the rich fabric of the broad context?

Having probably spent too many words on a commonplace, I come to the main point of the present section. Whenever the student of context wishes to base an argument on the consideration of a *particular* context, he has no other tools, of course, than finite descriptions. Those are, except in special cases, unable to capture the properties of the context, which is infinitary. Thus arises what we shall refer to as the *theorist's predicament*.

The situation may not at first seem at all hopeless. For there to exist a finite description capturing the full set of properties of the context at hand, it suffices that the latter be finitely generated: there could well be features of it (forming a *finite base*) which generate, cause or entail all the other features. It could happen that this is the case for all interesting, ecologically sound types of context. Unfortunately, the argument used above against the existence of finite contexts, to the extent that it shows anything at all, is an argument against the existence of such a finite base in all but degenerate cases. Now the opposition, on whose shoulders the burden of the proof seems to rest anyway, has, or at least had, something to propose: context concepts, known in early AI as scripts, scenarios or frames, were precisely supposed to provide the finite bases of ecologically important contexts. The famous restaurant script invented by Roger Schank in the 1970s, for example, is an informational structure which, when its variables (*slots*) have been fixed on some particular occasion (Fred entering the famous Parisian restaurant *Taillevent* on October 20, 1998 at 8 pm), becomes the finite base of Fred's context as he enters *Taillevent*. Scripts don't work, and the likely reason is that they cannot work, because contexts have no finite base, despite AI's *credo* that they *must* else we humans could not cope with natural contexts.

There remains the hope that natural contexts may enjoy a weaker property, affording the theorist an escape from his predicament. It could happen that a natural context would have a finite base which would generate all of its *non-negligible* features. Such *almost*

finitely generated contexts would be handled by the theorist exactly the way material entities in the natural world are treated daily in the works of scientists and engineers, i.e. by idealizing away immaterial details and keeping the ensuing error under control. However, the sophisticated theorist of context knows how doomed that hope is, even in the ultra-favorable case of utterance interpretation, because he is an expert at coming up with a heretofore unmentioned tell-tale feature which changes the face of the earth, at least locally: this is not to be seen as an *event* which changes the present situation, but as an aspect of a fixed situation which was not included in the purported finite (quasi) base and is nevertheless essential to the agent. In other words, what these examples, which fill volumes, show, is not the non-linearity of certain effects, but the impossibility of even *almost* finitely generating (nondegenerate) contexts.

So how does the sophisticated theorist of context imagine he can escape his predicament? The answer is that he relies on the notion of a *standard* or *normal* context generated by a given finite base. As Recanati aptly remarks, the worlds imagined by Searle, for example, in which sentences true in general turn out to be false, are “characteristically strange: lawnmowers are used to carve birthday pies, there is no gravitation, hamburgers are a mile long, etc.” The theorist could thus hope to rely on a “non-strangeness” clause, and consider that a suitable finite description of a given context does provide a full characterization of it under the assumption that it is a,<sup>3</sup> or the, *normal* context compatible with the description. He could claim that the utterance “Would you pass me the salt?” calls for the handing of the salt whenever the context is normal, e.g. when there is a salt shaker in view yet out of reach of the speaker, *and* the speaker does not continue by saying “if I were your friend again” or “if I weren’t on a low-sodium diet” *and*. . . . He might even be tempted to say that the meaning of that utterance is a request for salt, *ceteris paribus*. The question then would be *what* other things?

As the literature makes abundantly clear, this *is* the way context theorists overcome the predicament, although they usually don’t have much to say about it. And there is no question that it works: readers get the points that are made, theories are constructed, defended and corrected on the basis of countless examples described

with the help of not only finitely many, but usually very few, words. The question then becomes, *why* does it work?

Part of the answer probably lies in the special properties of the phenomenon studied in the most successful part of the literature, i.e. linguistic communication – but I shan't attempt to say exactly which of these properties come into play (cooperativeness is an obvious candidate). I shall restrict my attention to another level of communication: the theorist communicates successfully with the reader, so that his examples are grasped correctly by her on the basis of a few words (the fact that these are examples of successful communication becomes irrelevant; indeed, the topic could be the handling of a terminally ill relative, or of an intoxicated driver, or of a business transaction, . . .). There are two reasons why a few words suffice. One is the principle of relevance, which entails that the reader knows that the absence of any indication in the theorist's presentation that there is something exceptional about the context described is to be interpreted as a guarantee that the context is not exceptional. The other is that it is enough for the hearer to know that the context is non-exceptional for her to retrieve any desired conclusion from the description provided; thus, theorist and hearer rely on a shared sense of what the normal context compatible with the description is. They share a norm which applies to contexts.

### 5. *Normativity*

It is difficult – perhaps impossible – to define normativity independently of one's stand regarding naturalization. The naturalistic philosopher will want to define normativity in such a way as to make an eventual reduction of normativity to natural science constructs at least non-contradictory. The antinaturalistic philosopher favors the opposite strategy. What they may agree on is the methodological necessity of starting off with a rough and ready distinction between two kinds, levels or degrees of normativity: on one hand the *bona fide*, strong, traditional normativity; on the other the pseudo, weak, reduced, thoroughly modern sort.

This is precisely what we need to do for the purposes of the present article, which takes no stand on the general issue of naturalism. I shall say that *strong normativity* rules over a domain of actions or happenings whenever (a) there is an in-principle differ-

ence between the actual unfolding of the action or happening, and an independently defined possible unfolding (the *standard*); (b) it is possible, or at least conceivable, for any action or happening, to be arbitrarily close to the standard; (c) the difference can be grasped, considered, and above all *weighed and debated* by a conscious agent in the full, traditional sense; and (d) an involved agent may desire to make it the case that the difference be as small as possible.<sup>4</sup> *Weak normativity* shall be defined as requiring conditions (a), (b) and (c'): the difference can be grasped or considered by an observer who is in a position to attribute a function to the action or happening.

Weak normativity is in a sense *extrinsic*: the difference between the actual unfolding and the standard is apparent to an observer, not (necessarily) to the system or creature where the happening occurs. Strong normativity by contrast is *intrinsic*: the difference in question is apparent to the agent (insofar as he is an agent), not necessarily to an observer; to try and ensure irreducibility of strong normativity, one can in fact (but I shall not) impose the additional condition that only an observer who shares the norm (and is therefore potentially an agent in the same sense as the agent under consideration) can directly *sense* the difference (a shrewd observer not sharing the norm could be a reliable judge but would lack direct access to the reason for his judgments being what they are). A more straightforward way of putting the difference between weak and strong normativity is to say that in the former case, there is a fact of the matter as to what difference there is, while in the latter there possibly (neutral version) or necessarily (antireductive version) is no fact of the matter.

Typically, human actions are governed by strong normativity, while machine behavior is governed by weak normativity. Homeostatic, or more broadly, cybernetic machines exhibit a form of normativity in which the norm is in some sense self-generated and the deviations can be corrected without outside intervention. As is well known, the development of ever more elaborate machines of that sort has fuelled the hope that strong normativity could be *mimicked* by weak normativity in self-organized machines, and further to the conceptual naturalistic conjecture that there is no objective difference between the two: roughly speaking, according to this line of thought, strong normativity is nothing but a species of weak normativity seen from the first-person perspective.

Although no final empirical or conceptual refutation of that species of naturalism has come forward, as far as I know, another kind of reductive enterprise seems to have taken precedence in contemporary discussions. The neo-Darwinian explication of a norm is based on an idea which can *very roughly* be expressed thus: among the things which an organism can do, some correspond to, or result from, a genetically controlled capacity which contributed, back in the times when the species was evolving, to its reproductive success. What the organism actually does, on any one occasion, can be measured against that independently determined standard. As for conscious access to the norm, it is presumably built in by evolution itself: it might sometimes be adaptively advantageous for the organism to be endowed with an internal representation of the norm, which it can use for self-monitoring purposes. We betas do believe that betas are best, and that belief helps us succeed as betas. Thus strong normativity amounts to innate possession of the norm by the creature, the norm itself having arisen, in weak form, over evolutionary times.

The purpose of this simple, indeed simplistic, discussion is merely to prepare ourselves for the discussion of the thesis of this paper, *viz.* that context is normative. In the weak sense of the term, this would perhaps not merit the reader's attention: *of course* many cognitive processes are context-sensitive, he might say, so isn't it *trivially* the case that context defines a norm over these processes? The point which I shall try to defend is that the norm in question is not merely of the weak kind.

## II. THE THESIS

### 6. *Accomplishing an Act Within a Given Context*

The situation which we will now examine is the following. An agent is about to engage in some activity – typically, performing a task, quite possibly interpreting or otherwise processing some 'text' or input. We are considering a particular occasion, not a class of cases. At that moment, there is an objective situation, which the agent's expectations turn into a full-blooded context. In the present section, we assume this context to be fixed. We want to know what it means for the activity envisaged by the agent to be affected by the context.

One may well reckon that a formal approach would clarify the question. The starting point is the idea that when context varies while everything else remains fixed, the outcome, or the activity itself, varies. So we need to consider the class  $O$  of conceivable occasions which are indistinguishable from a given occasion except for the associated context. For example, a class may comprise all the conceivable occasions on which Mary has just said to Fred: “Would you pass the salt?”. To say that it is context-sensitive just means that as the occasion varies in  $O$ , so does the activity: sometimes Fred passes the salt, and sometimes he says “Yes”, and sometimes he says nothing, etc. It is essential of course to insist that Fred’s intentions and other inner dispositions are fixed: only external circumstances may vary, although the variations are filtered, so to speak, by Fred’s expectations. So for example, Fred’s response in the case of the salt varies not because Fred is being cooperative on the first occasion, mischievous on the second, and cross on the third. The reason might be, rather, that the first situation is “normal”, that the second includes an indication of Mary’s lack of interest in obtaining the salt shaker, and that in the third Mary has been reading aloud from *Relevance Theory*. To relieve the suspicion that these cases work only because they are bizarre, *à la* Searle, let us take Recanati’s favorite example and imagine that Mary has just said to Fred: “Would you pass me John’s book?”. On one occasion, Fred may hand her the book *written* by John, on another the book *owned* by John, on the third the book *read* by John, and so on *ad infinitum*. Nothing out of the most ordinary in such scenarios.

Thus a formal characterization of the claim that a certain process  $P$  is context-sensitive should render the fact that  $P$  receives its input  $x$  from some domain  $D$ , but that there is a possible contribution to the outcome of  $P$  from other sources. In other words,  $P$  is a function of two variables (rather than one):  $P(x) = \phi(x, c)$ , where  $c$  belongs to some set  $C$  of “contextual factors”. Formally this is equivalent to saying that  $P$  takes its values not in the set  $T$  of outcomes, but in the set of functions from  $C$  to  $T$ :  $P(x)(c)$  then is the preferred notation for  $\phi(x, c)$ .

In passing we may note how the idea of a “normal” or “default” context is cashed in: this is the case where  $c$  takes a constant value  $c_0$  across a wide range of naturally-occurring occasions; so that one

may revert to the simplified notation  $P(x)$ , understood as  $\phi(x, c_0)$  or  $P(x)(c_0)$ . Then on those occasions when  $c$  differs from  $c_0$ , and  $P(x)$  takes different values for  $c$  and  $c_0$ , one can say that the context interferes and dictates, or corrects, the outcome. But it is obviously just a manner of speaking: occasion determines context, and context in turn determines, together with input, the result.

Returning to the general case,  $P$  differs only notationally from  $\phi$ , and  $\phi$  from  $\Psi$  defined as the function taking just one argument, the ordered pair  $\langle x, c_0 \rangle$ , such that  $\Psi(\langle x, c_0 \rangle) = \phi(x, c_0)$ . But then, reverting to informal language, this expresses the fact that contextual dependence is no more than the effect of a non-perspicuous characterization of the existence of hidden variables. Our proposed analysis is insufficient by itself for the reason that it rests on a pre-formal distinction between input and contextual factors. What we would need is an independent, intrinsic property of the latter which turns them into permanently hidden variables, not merely unnoticed variables. In the absence of such a property, the theorist has no choice but to admit that his analysis of  $P$  was incomplete, and that a correct analysis yields a characterization of  $P$  as a context-free process ranging over some higher-dimensional domain.

In order to save the very idea of context, and its normativity even in weak form, we must therefore take a step back and ask whether our formalization is on the right track. The answer is a resounding No. First, even if we had a canonical segregation of input and context, the formalization abolishes the asymmetry between them. We can say that the correct response to Mary's question "Would you pass me the salt?" is sensitive to context, not that the correct response to the context is sensitive to Mary's question. Second, some "variables" don't seem to qualify as context at all. Surely we're not tempted to think that the price of the London-Paris Eurostar ticket is sensitive to the context of class, period of the week or season; nor for that matter that the first-class mid-week winter season Eurostar ticket is sensitive to the context of points of departure and arrival! Third, what grounds have we to think that contextual factors can coherently be treated as variables at all? Except perhaps in special cases, there are no candidates for the domain over which the contextual variables would range. Here is a possible exception: Kaplan has suggested that the meaning of

“I” in an utterance depends on the value of the contextual variable “utterer”; this variable ranges presumably over the set of possible utterers. There are difficulties with this idea,<sup>5</sup> but our point is that it does not seem to generalize even to apparently such straightforward cases as “John’s book”. Functions from contexts to extensions are ghostly creatures because they have no definite domains, however abstractly defined. The class of occasions which we denoted *O* above is a chimera. And then, fourth and last objection, what are we to do when we leave the cosy corner of utterances and venture into areas where there is no canonically defined text, or even no input at all?

The contribution of the formal notation is thus entirely negative: it makes salient the limitations of any attempt to cope with context by collecting together the range of values of contextual factors and turning the relevant part of context into an extra variable. From classical AI to formal semantics, attempts at this reductive strategy have not been few.

I suggest that context should not be regarded as part of the problem or task at hand, but rather as what *generates* the problem or task. No context, no problem nor task. On the other hand, once the context is set, the problem is defined, and its solution is determined, and it is dependent, obviously, both on text or input, if any, and on context.

We should note in passing that this way of putting things nicely brings together the two kinds of cases (with and without ‘text’) which appear in the standard information-processing framework as totally separate. Hearing Mary say “Would you pass the salt?” is just one of the reasons Fred might have had to hand her the salt shaker: he could (on another occasion) have thought it appropriate to do so without Mary or anyone else having as much as raised an eyebrow. Intuitively it does seem correct to regard the two cases as related, though clearly different.

On the other hand, the view I am pressing tends to increase the distance between different occasions on which Mary asks Fred the famous question: the standard view brings them together under the general heading “Fred’s interpretation task of Mary’s question ‘Would you pass the salt?’,” and makes the accomplishment of the task context-dependent. We want to reverse the order: context comes

first, and if two contexts arising on two different occasions are sufficiently distant, the interpretation tasks are proportionately distant, and the commonality induced by the fact that it is one and the same question of Mary's which it is Fred's goal to interpret doesn't do much to bridge the distance. The sceptical reader might want to ponder the following example: Mary has been lying comatose for nine days, with her husband Fred at her bedside; now comes Mary's question.

Having thus disposed of the question of how – in what sense – context affects the agent, we must make good on our claim that the way in which it does is “normative”. Weak normativity would mean that while context unambiguously determines a course of action (the “right” interpretation, the “right” gesture, etc.), it is materially possible that the agent takes a different course of action, although he could (in some sense) have made the “right” choice as imposed by context. Thus Fred might have looked around for the salt shaker, although it was clear from context that Mary was requesting a kiss. For reasons of cognitive malfunctioning, Fred failed to correctly apply the contextual clues to the task at hand. Further, the naturalist will insist, it is a *fact* that the context dictates the choice it does – a natural fact, one which follows from natural laws, is computable as a solution to some equations, or results from inferences in some perhaps nonclassical logic. It may take mankind many centuries of hard scientific investigation to discover these laws, equations, inferences, and how to apply them (perhaps at the cost of idealization) to the case at hand, nevertheless *in principle* there is a path from known facts to the correct response which science, or universal reason, can discover.

I want to resist the naturalist's argument. The normativity in virtue of which context dictates a correct course of action is of the strong type: there is no natural law which can *replace* the norm, although there may be natural laws which *explain* it. Whether Fred should have given Mary a kiss is *not* a fact independent of an assessment, a judgment of what the right thing for Fred to do was. It is a matter of debate between human beings who have nothing to rest on but their subjective appraisal of what the context was for Fred. There are no laws to invoke, no calculus to set in motion, and there never will be because the question would remain of whether the laws

or the calculus apply to Fred's context on that particular occasion. And it is not simply the usual predicament of the scientist, who can never *know* whether he has hit upon the truth, nor the logician's problem of needing some meta-rule of inference to confirm that he can apply a rule to a given case: what is to be accounted for is Fred's, and the witnesses', *post facto* deep conviction that the kiss was the right policy, not whatever natural fact might explain the fact of that conviction – if it is a natural fact at all, which we have every reason to doubt: there simply are no convincing cases of universal agreement among humans of what the appropriate course of action is *in a concrete, singular occasion*.<sup>6</sup> In fact, now Fred feels like a fool and regrets not having given Mary a kiss, but who is to say that on the spur of the moment, in the impossible scenario of an exact replay, Fred would not again have looked for the shaker, and the cerebroscope surveying his brain states not confirmed that he felt fully justified in doing so, and would *reject* his own *post facto* criticism? Again, this might look very much like common sceptical arguments, but it is quite different: the arbitration of an intersubjective critical space is not available, and nor is a body of natural laws.

In brief, applying the context in the “right” way to determine the course of action is akin to issuing the “right” sentence once the parties in a trial have been heard, to playing a piece of music in the “right” way, to making the “right” decision in an ethical predicament, to making the “right” choices in one's life, to getting dressed in the “right” way for a social occasion.

Now the norms which govern the latter fields are well known and although there may be endless discussion of how to classify and name them, there are words for them. What about the norm which according to us governs the application of context to the accomplishment of a task? My suggestion is to call it *intelligence*. We shall later return, briefly, to the reasons for this terminological choice. But I envision it, whatever one wants to call it, as a non-reducible, *sui generis* norm.

A more urgent matter is to deal with a couple of related objections which will be raised at this point. Bearing in mind the beginning of this paper, the reader might say this: since, as we are told, everything we do occurs in context, and the manner in which context deter-

mines the proper course of action is normative, it would seem that no human action, be it ideal, can ever be said to be correct in an absolute sense, that it can only be measured against a norm of which we only have an intuitive grasp. Therefore, any rationalistic analysis of the constraints bearing on the accomplishment of any task at all is useless, or at least severely limited. And while this situation might be regarded by some as acceptable in relation to ethical and esthetic norms, it runs against the evidence provided by common sense, by the whole philosophical tradition, as well as by cognitive science. There may be philosophers who are prepared to bite the bullet and acquiesce to this conclusion. Not the present author. To see how we might resist it, let us consider the second objection.

How context bears on the course of action to adopt is, in each specific occasion, a matter for debate. So let us imagine Dr. A and Dr. B discussing, say, how to read the blood test results for some patient. Dr. A bases his therapeutic recommendation on the consideration of features a, b, c, . . . of the patient's case, and features x, y, z, . . . of the epidemiological situation, the available drugs, etc. He may also get part of his motivation, or justification, from a list of similar cases with which he is familiar. Dr. B agrees about a and b, and about x and y, but thinks c' and z', not c and z, significant; he has another recommendation to make, based in part on other similar cases. Eventually, the two physicians will reach an agreement, possibly at the cost of some suspension of belief on one or the other side. Our imagined reader's objection to my analysis of context is this: What doctors A and B are discussing is *objective* medical facts, albeit under conditions of uncertainty.<sup>7</sup> They experience no *context* problem, they are dealing with a *medical* problem. Their predicament is exactly comparable to that of a couple of entomologists, or astronomers, or even mathematicians, writing a paper together. Now, our critic continues, who would dream of regarding entomology, astronomy, or mathematics as areas governed by some irreducible norm, call it intelligence or anything else, rather than bodies of factual knowledge? Aren't we confusing facts and the discovery of facts? And aren't we trying to revive a defunct idea, that of *general* problem-solving techniques (be they intuitive rather than rational), and for that matter the equally discredited notion of *general* intelligence?

This objection, raising as it does several distinct though related points, calls for a lengthy discussion, which we must cut short. One important issue is whether there are general problems involving context on one hand, and domain-specific problems of evidence, proof, etc. on the other. Another is whether the nonfactual character of context-sensitivity has sceptical, antirealistic or antirationalistic consequences on knowledge and other rational enterprises. My view is that while there are in effect general problems involving context, there are no general solutions to those problems, but only domain-specific heuristics which guide the expert in his search. So while no amount of medical, entomological or mathematical knowledge can determine the correct course of action in any specific *singular* episode of scientific research (taken here as a paradigm of rational endeavor), there is (in favorable cases) a know-how which guides the expert towards a proper weighing of the evidence in the light of context. This explains the relative success of the discovery process, when it is indeed relatively successful (i.e. not very often). On the other hand, once the *narrow* context has been determined, with the help of this expert-intuitive heuristic, a purely objective inferential process can develop, leading to a piece of theoretical knowledge. The objectivity of this knowledge is the result (or perhaps the *expression*) of a combination of factors, the intersubjective assessment of the specialised heuristics, the objective and verifiable character of the inferential part of the work, convergence and experimental corroboration of the results, etc. Hence the factual character of at least some collective results of our rational enterprises is preserved, as is the normative character of context application, i.e. once again the fact that on any given particular occasion, the scientist is faced with a singular context problem and can no more rely on some (complete) *theory* of how to conduct his business than the next-door second-hand car dealer, because such a theory cannot exist.

It is of course an approximation to think of context application as an initial stage, followed by a context-free inferential stage. But it is a useful approximation. To put it in a different, terribly metaphorical but possibly less misleading way, any situation raises a problem involving two dimensions. Along one, the contextual constraints are evaluated; along the other, the (partial) results of

this (ongoing) process of evaluation are entered in an inferential module. Rationalistic strategies, as typified in science, consist in relying on intersubjectively assessed expertise to deal with the contextual dimension, and on applying the usual problem-solving, domain-specific techniques to the other.

### 7. *How a Given Situation Gives Rise to a Context*

In the last section, we pretended that the context was fixed – the broad context, that is, not the proper part of it which a given “text” on a given occasion calls forth, not what I called the *narrow* context. Now we want to widen our perspective and ask whether the (broad) context is in fact strictly determined by what we have called the situation and the agent’s expectations. What I have in mind is a familiar albeit ill-defined phenomenon, or perhaps a variety of phenomena which seem related by virtue of involving changes of context: for example, seeing something (someone, some event) “in a [completely] different light”, a happening sometimes known as a *Gestalt* shift; or finding it suddenly possible to solve a problem which seemed intractable before, or to accomplish a gesture which seemed impossible to make, or “rising above the circumstances”, and so forth.

Because the issue is murky and this paper is becoming too long, I shall content myself with a mere outline of an argument. First let us take as our example something less heady than a *Gestalt* shift, and thus perhaps less telling, but easier to grasp and sufficient to provide an initial grip on the issue. Let us imagine an academic selection committee screening applications for a university appointment. Until recently the committee adopted a policy of taking into consideration all the facts that it could learn about the candidates. Following a reform, it now takes into account nothing but the professional background and work of the candidates. Finally, let us imagine a committee member, Dr. X, used to the old policy, screening an application according to the old rules, and who suddenly remembers, or is reminded of, the new policy. Dr. X might experience a change – the sort of change which we have in mind: the situation is constant, the ‘text’ remains the same, and the expectations of the agent are unchanged as well, yet the context changes. The reader might well dispute those claims, but

the discussion would be quite long and not worth our while: there are mostly issues of terminology and underdescription; one real problem concerns what is meant by “expectations”. Let me propose the following temporary resolution: in the case at hand, Dr. X’s expectations are roughly made up of the goal to select a candidate amongst the applicants within the framework of the Committee, together with all the beliefs and desires active in Dr. X’s mind at the moment under consideration. True, one could say that his goal has changed: from “select best applicant on the basis of *all* the available information” to “select best applicant on the basis of the strictly academic information”. But let us stipulate that it is not the case: Dr. X is agnostic on what the proper policy is, he is simply after the best candidate. Yet when he does restrict himself to the strictly academic portion of the information he has on the candidates, he sees them in a different light.

The point of this example is to make plausible the idea that there is leeway in the fixation of context: situation, task, expectations, “text” do not jointly entirely determine context. Thus there is substance to the question of the *proper* context in which to accomplish a task or decide on a course of action. Again, the naturalist will claim that the normativity at work is of the weak kind: some objective optimality criterion can in principle be found, presumably heavily dependent on the domain, on the task and perhaps on prevailing norms, such as the deontological rule which Dr. X’s committee recently adopted: and of course, to the naturalist, in such cases the normativity of context choice is derived. But what are we to do about Dr. X’s *decision* to change his whole perspective and examine the candidates’ dossier in the light of, say, their covering letters? “Decision” is not quite the right word, perhaps: it might have simply occurred to him, but now he cannot but see A’s application in a certain way, B’s application in a certain way, such that in his eyes B is now better off than A while it was the other way round a minute ago, before he turned his attention to the covering letters. What are we to say to Dr. X? We may cite the deontology of the Committee, but suppose there is no rule against taking covering letters into account? We may argue that it does not sit comfortably with the rules and practices of the Committee, which seems pretty close to saying that the letters don’t belong to the proper context! Better

perhaps, we may argue that covering letters are poor predictors of academic excellence. Again, this is missing the point: the argument may carry weight as a general rule, it says nothing final about Dr. X's specific predicament. So we might just end up telling Dr. X that he is being *silly*.

Perhaps the reader, being familiar with the basics of the history of science, will be more willing to consider examples drawn from that area. Lavoisier and Priestley saw combustion in different lights; according to which objective criterion was Lavoisier's choice of context more appropriate than Priestley's? The discussion is entirely parallel, at this point, to the one concluding the last section: of course *post facto* we see reasons drawn from chemistry as we presently understand it; nevertheless the context fixation "problem" Lavoisier faced had no objective solution. There is no fact of the matter about which was the correct context. Here Kuhn's authority can be cited.

My tentative conclusion is that the fixation of context, just like the application of context, is strongly normative, and I propose again to think of the norm which is involved as bearing a relation to what is commonly known as intelligence. I am aware, though, of two problems with this conclusion: first, it relies on an argument made rather fragile by the recourse to ill-defined terms and under-described examples; second, it raises the suspicion that either there is just one basic phenomenon to which context gives rise, rather than two (fixation and application), or that regress threatens – why shouldn't the fixation of context be regarded as the application of some higher-order super-context? We must leave these issues pending for now.

#### 8. *Intelligence vs Rationality*

There probably is a more pressing question on the reader's mind, which concerns my use of the notion of intelligence. Why resurrect a thoroughly discredited concept, one whose use has moreover led to all kinds of social and political perversities, when there exists the philosophers' well studied notion of rationality? Needless to say, I am not about to launch a counterattack on the grounds on which intelligence has been soundly defeated.

And of course, we could make it just a verbal issue, considering how varied and loose the usages of the words in question are. On the other hand, even a superficial inspection of their connotations reveals distinctions which may be relevant to the matter of context.

Rationality, first. Rationality is linked to reasons. In what Bernard Williams calls the 'rationalistic conception of rationality' (1985, p. 18), to be rational is to be able to give reasons. In a more contemporary framework, this can be weakened to the mere *fact* of the proper use, by the agent, of the information at his disposal; although the agent may not be able to *state* his reasons, he does *use*, so to speak, the reasons which a rational agent in the traditional sense would use (and state) in his place. In any case, rationality is strongly associated to the idea of rule-governed behavior. Herbert Simon's notion of limited rationality allows finiteness, or feasibility, constraints to bring the idealization closer to real-life decision-making. Still, the general picture here is that of a computing agent with access to a data base and heuristic rules to limit the inferential search for the solution. A presupposition of this conception, and more generally of naturalistic conceptions of rationality, is the existence of a predefined goal, or more generally of a pregiven problem; the rational agent is the quintessential problem-solver.

Intelligence, on the other hand, in at least one sense in which the word has been used in the Western tradition, has to do with the predicament of a human being needing to figure out what to do; there may be no well-posed problem to solve, no task to accomplish, but instead the opportunity to pose or reformulate a problem, set oneself a task, define a goal. Intelligence is the art not only of reaching a goal, but much more generally of making the best of a given situation, and possibly first and foremost of making sense of a situation. Contrary to rationality, intelligence is not primarily a matter of rules. Next, intelligence, as construed in "folk psychology", seems to come in two degrees: the lower sort is associated with an ability to do well in a given context, the higher sort is linked to creativity, which has something to do with the ability to define a context anew. Creativity is clearly not in the cards for rationality, as it is usually construed: minor mathematicians (administrators, businessmen, etc.) are rational (qua mathematicians, etc.), they are not intelligent in the higher sense. One final difference to

the advantage of intelligence is that while rationality operates within a given framework (indeed, within a given context in our sense of the word), intelligence comes both before and after other norms and constraints have been set. A case in point is Ulysses' sailors and the Sirens, or voluntary limitation of access to information as in refereeing. Rationality (even limited rationality) dictates gathering as much information as feasible, and then perhaps neglecting what need be. Intelligence takes into account the risk of deviation from rationality once the context has been adulterated by the presence of misleading or ethically irrelevant "information".

None of this is meant to be conclusive: I have no definition or theory of intelligence, and am aware of the fact that many writers are happy to use "rationality" where I say "intelligence". Above all I do not want to reify it as some sort of higher-order capacity. What I am urging goes precisely against this trend: intelligence is a *sui generis* norm, and its grammar is that of a norm-word, not a faculty-word. Everyone recognizes that intelligence is 'to itself its own criterion', which is a telling symptom of strong normativity. The name of this norm is of secondary interest, although ignoring the lessons of folk psychology may result in depriving oneself of a precious source of insights.

How tight is the relation between intelligence and context? There are several links which cannot be described here due to lack of space and which support the idea of an intimate connection. But there is no readily available argument to show that intelligence is *exhaustively* characterized as the norm governing the fixation and application of context. This is unsurprising, considering the distinct looseness of the everyday notion of intelligence, and the limitations of folk psychology. It would be instructive nonetheless to find traces in folk or scientific psychology of a natural species of intelligence which would correspond to the norm governing context use.

### 9. *Can Naturalism Survive?*

The reader might well be under the impression that despite cautionary remarks to the contrary, the author takes side against naturalism and conceives of the theory of context proposed in this paper as an antinaturalist weapon. In this very brief concluding section, I would like to set the issue in a different light.

First I should make clear that the focus of this paper is context as a feature of human cognition, not naturalism in general. My imaginary opponent is often called ‘the naturalist’ because of the major assumption he brings to bear to the problems I raise. But what I have in mind is the manner in which he deploys it. What I see my arguments as damaging for is not naturalism *per se*, but a certain influential naturalistic strategy in the philosophy of mind, based on the computational model set within a neo-Darwinian perspective. And even then, I do not regard the difficulties I raise as *proving* that the model cannot be saved. But they do invite us to look into other possibilities.

Alternatives need to be sought in parallel on the side of the phenomenon and on the side of the model. So-called ‘dynamical models’ (those taking dynamical systems rather than computable functions as their fundamental mathematical theory)<sup>8</sup> have a lot going for them, despite being far less understood than traditional computational models; in particular, they feature a kind of autonomy and self-organization which seems more intrinsic and biologically plausible than their rivals. This sits well with a perspective on mind which centers on the constructive abilities of the organism. The basic idea of the constructivist approach is this: as it develops, the creature creates a world of its own out of its biological milieu. Thus whenever it has to defer to the environment in order to know what to do next, it consults or inspects not the entire universe, but only the niche it has created within this universe for itself. This niche is known in a special way by the creature, just as a home, however unordered, is known by its denizen. Clearly the strategy is not fail-safe, but it is a feasible one. Just as clearly, the borderline between home and non-home is not a sharp one, and, importantly, is it constantly reshaped through learning and restructuring.

The knowledge a creature has of its niche is not of a magical nature. What distinguishes it from ordinary, objective, knowledge is its ‘personal’ character, after Polanyi’s apt phrase: it possesses a rich implicit structure which is shaped by the creature’s experience and therefore geared, except for catastrophic novelty, to its future needs.

However suggestive these ideas may be, they clearly require major feats of philosophical and scientific imagination to reach the

stage of a fully autonomous and viable alternative. Context is one of the leads which might profitably be followed.

## NOTES

<sup>1</sup> Except perhaps, if we follow Sperber et al. (1995), for the fact that contextual effects on reasoning may often be reducible to contextual effects on communication.

<sup>2</sup> Recanati's more recent work on pretense tends to set him apart from this group, and brings him closer to the second group, especially Fauconnier; the latter, on the other hand, is a soft processualist, with affinities to the first group: our dichotomy is fuzzy, and would appear even more so if other authors were brought in.

<sup>3</sup> The indefinite article here indicates the possibility that the context described may be a context *type*, with a plurality of tokens distinguishable by space-time coordinates. As a matter of fact, the present author takes the notion of a context-type to be dubious, if not self-contradictory, but this point will not be defended here.

<sup>4</sup> An agent may recognize a norm without thereby wishing to endorse it (I'm indebted to Peter Railton for pointing this out). But *some* agent must conceivably be willing to endorse it.

<sup>5</sup> Take the sentence: "I was torn. Part of me was saying 'I should speak up', and another part was thinking 'I should keep my mouth shut'." Does the part of the utterer referred to by the first 'I' count as a member of the set of possible utterers, as demanded by the interpretation of the second 'I'? This set seems awfully large.

<sup>6</sup> Although space does not permit here to go to the heart of the matter, it should be stressed that the existence of a wide array of psychological universals (as in e.g. Brown, 1991) does not count as an argument against the view defended here, which rests on the consideration of *particular occasions*. Then there is the point that these universals might be the result of a shared predicament, not of psychological laws or even Darwinian selection.

<sup>7</sup> There is a connection with the topics of Bayesianism and of judgement under uncertainty, especially the matter of base-rate neglect, which can be seen as a generic contextual error. We shall not pursue the matter here.

<sup>8</sup> See Amit 1989; Port and van Gelder, 1995.

## REFERENCES

- Almog, J., Perry, J. and Wettstein, H. (eds.) (1989): *Themes from Kaplan*, New York: Oxford University Press.
- Amit, D. (1989): *Modelling Brain Function. The World of Attractor Neural Networks*, Cambridge: Cambridge University Press.

- Andler, D. (1987): 'Problème: une clé universelle?', in I. Stengers (ed.), *D'une science à l'autre: des concepts nomades*, Paris: Seuil, pp. 119–159.
- Andler, D. (1993): 'Is context a problem?', *Proceedings of the Aristotelian Society* LXXXIII, 279–296.
- Andler, D. (in press): 'Context and Background. Dreyfus and Cognitive Science', in M. Wrathall (ed.), *Heidegger, Coping and Cognitive Science*, Cambridge, MA: MIT Press.
- Barkow, J., Cosmides, L. and Tooby, J. (1992): *The Adapted Mind. Evolutionary Psychology and the Generation of Culture*, New York: Oxford University Press.
- Barwise, J. (1986): 'Information and Circumstance', *Notre Dame Journal of Formal Logic* 27, 324–338.
- Brown, D. (1991): *Human Universals*, New York: McGraw Hill.
- Clark, A. (1997): *Being There. Putting Brain, Body, and World Together Again*, Cambridge, MA: MIT Press.
- Cohen, L.J. (1986): *The Dialogue of Reason*, Oxford: Oxford Univ. Press.
- Dreyfus, H.L. (1972): *What Computers Can't Do*, New York: Harper & Row; (1993): *What Computers Still Can't Do*, 3rd edn., Cambridge, MA: MIT Press.
- Elster, J. (1989): *Ulysses and the Sirens*, Cambridge: Cambridge University Press.
- Fauconnier, G. (1985): *Mental Spaces*, Cambridge, MA: MIT Press; (1994): 2nd edn., New York: Cambridge University Press.
- Fodor, J. (1983): *The Modularity of Mind*, Cambridge, MA: MIT Press.
- Hirschfeld, L.A. and Gelman, S.A. (eds.) (1994): *Mapping the Mind. Domain Specificity in Cognition and Culture*, Cambridge: Cambridge University Press.
- Harman, G. (1986): *Change in View*, Cambridge, MA: MIT Press.
- Hutchins, E. (1995): *Cognition in the Wild*, Cambridge, MA: MIT Press.
- Johnson, D.M. and Erneling, C.E. (eds.) (1997): *The Future of the Cognitive Revolution*, New York: Oxford University Press.
- Kaplan, D. (1989): "Demonstratives", in Almog et al. (eds.), pp. 481–563.
- Karmiloff-Smith, A. (1992): *Beyond Modularity, A Developmental Perspective on Cognitive Science*, Cambridge, MA: MIT Press.
- Kuhn, D. (1962): *The Structure of Scientific Revolutions*, Chicago: Chicago University Press; (1970): 2nd edn.
- Perry, J. (1993): *The Problem of the Essential Indexical and Other Essays*, New York: Oxford University Press.
- Polanyi, M. (1958): *Personal Knowledge*, Chicago: Chicago University Press & London: Routledge and Kegan Paul. Revised ed. 1962; rep. New York: Harper Torchbooks, 1964.
- Port, R.F. and van Gelder, T. (eds.) (1995): *Mind as Motion*, Cambridge, MA: MIT Press.
- Récanati, F. (to appear): 'Déstabiliser le sens'.
- Quartz, S.R. and Sejnowski, T.J. (1997): 'The Neural Basis of Cognitive Development: A Constructivist Manifesto', *Behavioral and Brain Sciences* 20, 537–596.
- Schank, R.C. (1982): *Dynamic Memory*, Cambridge: Cambridge University Press.
- Searle, J.R. (1978): 'Literal Meaning', *Erkenntnis* 13, 207–224.

- Searle, J.R. (1995): *The Construction of Social Reality*, New York: The Free Press.
- Sperber, D. and Wilson, D. (1986): *Relevance: Communication and Cognition*, Oxford: Blackwell; (1996): 2nd edn.
- Sperber, D., Cara, F. and Girotto, V. (1995): 'Relevance Theory Explains the Selection Task', *Cognition* 52, 3–39.
- Stich, S.P. (1990): *The Fragmentation of Reason*, Cambridge, MA: MIT Press.
- Travis, C. (1996): 'Meaning's Role in Truth', *Mind* 105, 451–466.
- von Uexküll, J. (1956): *Streifzüge durch die Umwelten von Tieren und Menschen – Bedeutungslehre*, Hamburg: Rowohlt.
- Williams, B. (1985): *Ethics and the Limits of Philosophy*, London: Fontana Press/Collins.

*Université de Paris – Sorbonne and  
CREA, Ecole Polytechnique  
1 rue Descartes  
75005 Paris  
France*

