

ISSA Proceedings 1998 - On The Fallacy Of Fallacy: Arguing For Methodological Difference: Producing vs Processing



1. Introduction

Fallacies have always been in the centre, or near the centre, of argumentation studies. In fact they lie at their roots in two senses: most approaches to argumentation have sprung from a consideration of what is amiss in human reasoning or thought, and theories of argumentation stand and fall with their capacity for detecting errors. In other words, fallacies are the cornerstone of argumentation theories very much like paradoxes once perceived by Russell as the stumbling block of scientific theories: they constitute the boundary conditions within which human thought and action remain to be *rational*. For a long time fallacies and rationality had been taken to be the two sides of the same coin, until certain evidences appeared to undermine their interdependence. They came basically from two sources: the psychology of decision making and the semantics and pragmatics of inferences in language use. Now it is no longer the exclusive power of argumentation theory that matters but their inclusivity, i.e. how charitable they are with faulty reasoning, error making and unjustified action. If fallacy theory does constitute a major divide, it works rather like a filter through which the beyond normal is let upon the territory of the rational; or at most it is a tradeoff between the rational and the irrational.

In this paper I am not going to take stock of the enormous data corroborating the “legal status” of irrational moves in thought and action; I only elaborate a little on the diagnosis that with the *cognitive turn* in the 70s a new look on the methodological basis of argumentation is needed. Yet I will not adumbrate a methodology here because, as I see it, there is an important, and not clearly noted, distinction underlying most of the insights in cognitive science that should be reckoned with in the first place before any stand on argumentation can be taken. Since there is not enough room here to fully elaborate this distinction, I have to suffice with some important consequences. Thus I am doing a kind of

archeology of knowledge in the Foucaultian sense, which may fall beyond the proper scope of argumentation theory, but if there is anything wrong with the idea of fallacy, as I think there is, it can only be identified in its undepinnings and its undepinnings are in cognition.

It is a most common opinion that the idea of fallacy is theory-laden: no fallacy without a theory. Now I want to oppose that view and try to argue for a rather strong claim that there are – at least some interesting – cases of language use when what appears to be fallacious or misplaced is not the given move itself but rather the *attempt* to judge what has been said or done as acceptable by some pre-set theoretical standards. Fallacies result then from a fallacious methodology; the methodology is fallacious for two reasons, which are however related.

2. *The outline of the argument*

I start then with the first reason why fallacies are originally methodological. It is constituted by what I take to be a major tension between the *descriptive* and *normative* ideal of argumentation theories. It is the basic claim of this paper that conflating the two inevitably leads to apocalypse. Thus it is because of the trafficking between the two ideals that John Woods could once call relevance theory as developed by D. Sperber and D. Wilson *apocalyptic*.

Since most frequently argumentative structures are the result of re-descriptions of utterances, in illustrating the first reason I will draw upon certain tenets from linguistic theory. This does not mean that I am necessarily biased by linguistic theorizing; rather the principles of understanding and producing language like relevance, graduality, similarity or structure mapping etc. should cohere with the more general principles of argumentation. If our understanding of language, i.e. of what is said, is apocalyptic, there is not much chance of constructing a – let alone sound – argument out of it.

Next I present my second reason by outlining a basic distinction that results from the findings of cognitive science. The distinction is between *producing* and *understanding*. My supposition is that even if the structure of our cognitive apparatus might at some future time be found the same in both cases, the terms of its operation, the aims and the procedural conditions, significantly differ.

The distinction has much to do with the debate of the continuity thesis of *similarity* and *rule-governedness* that has recently surfaced in cognitive psychology. (See e.g. the special issue of *Cognition* (65) 1998) Thus in this part I will cite some examples from categorization and topical research in linguistics

and criticize their treatment for not taking heed of the above distinction. The basic idea is that rules are abstract and context-independent, whereas similarity-driven processes are particular and contextual.

Finally I bring together the two distinctions within general rationality in terms of Donald Davidson's principle of charity. I also hint at an evolutionary framework to be developed along the lines proposed by R. Garrett-Millikan. The basic idea is two-tiered: i.) what is fallacious or not depends on the evolving of discourse and thus it cannot and should not be stated *a priori*; ii.) tampering with a rule is acceptable as long as both verbal and non verbal behavior preserve the biologically and culturally vital boundaries. This may be taken as a solution to the paradox of the sorites to which boundaries which are not fallacy-proof can easily give rise to. It is the reason why I consider my approach anti-apocalyptic.

To sum up: cases of rule-governedness, which is descriptive, cannot always mean rule-following, which is normative, and vice versa: cases of not following a rule does not necessarily result in violation *simpliciter*: it may amount to tampering with meaningful content: the domain covered by the the rules in question. One may wish to distinguish between motor activities, which *appear* to be rule-following to the external observer because they respect the evolutionary important boundaries without a proper representation of content, whereas higher cognitive activities *appear* to be rule-following to the internal observer because they are truth-preserving in inferencing and representing content. However, if the continuity thesis is correct, any attempt to separate out the normative element in the two cases is doomed to fail. One should look instead at how much producing speech and action and interpreting incoming stimuli are task-centred.

3. The graduality principle

Producing and interpreting differ in the first place as to their criteria of success. No doubt that in producing some behavior I have to cope with certain constraints or expectation environmentally determined. My behavior is rule-governed precisely in the sense that the constraints are out there: it is always rational to respect them and set the aim of my action accordingly. Yet their observance need not be normative in the full sense: I may be careless or lazy enough, or too roughly – even differently –disposed to come up with an optimal “solution”. What I thus produce, my performance, is rarely ideal or “well-formed”. This does not exclude that I may consciously chose to follow some abstract rule and approximate an ideal as closely as possible. Most (re)actions are however coarse-grained and/or

come off the target, while their aims may be properly defined.

In contrast when I interpret natural signs or other people's behavior, I always do it by relating it to what is given inside my mind, to what I know and believe. But it cannot be said that they are a kind of inner constraints with which strictly speaking I have to cope; rather they form the background for my understanding. Thus it follows naturally that *any way* I interpret what has been said to, or performed before, me IS rational. In other words the descriptive and normative ideals coincide. What I do is *eo ipso* optimal with respect to the *available* alternatives. Most interpretations are fine-grained and relevant to previous knowledge, although they can many times become automatic and similarity-driven. It appears then that, though rules and similarity in principle form a continuum, they are prototypical of two diagonally different activities: producing and processing. And while rule-following is the prototype of producing and shows more flexibility as a result of the working principle of optimality, similarity being the prototype of processing yields more rigidity in structure because of the underlying principle of mapping.

One - if not the only - reason that producing and processing are not mirror-images and rely on different mechanisms is that language use in humans amounts to more than communicating information. The idea is at least as old as the Gricean maxims. Today the clearest formulation of the common core of its "additional" - if not *sui generis* - dimensions is the *graduality principle* (GP). It is a structural principle of human knowledge in that it places the items in long term memory upon a scale or within a hierarchy of levels on the basis of the similarity among them. (Cf. Dubois \ Resche-Rigon 1996: 37) We can identify three important characteristics of GP. First that it allows for a categorical structure based on typicality *à la* Rosch. Second that it is value-laden in that it expresses a point of view and hence it can be utilized for argumentation. And third that it figures in lexical-linguistic structure. (Cf. Raccach 1993) Thus it results that the structure of cognition need not reflect - counter what Rosch claims - the ontological structure of our world, and neither does it follow formal-logical rules; rather it is governed by the orientations expressed in graduality. Language use "involves the application of general principles which we call *topoi* (*pace* Aristotle)." (Anscombe \ Ducrot 1989: 80) The *topoi* constitute an argumentative potential: they are correspondences among a series of gradations which allow for a set of possible inferences and can be exemplified with a comparative (the more/the less..., the less/the more...) structure.

Clearly, the aim here is to discover a common basis for our conceptual and

linguistic apparatus. Accordingly, the commonality is found in the task-centredness of categorization as well as of the manipulation of knowledge: it is always relative to a given task that category judgements are made and decisions are arrived at. And the list is by all means extendable to many kinds of contextual approaches, especially to relevance theory proposed by D. Sperber and D. Wilson where contextual selection is a primitive, an unreducible hallmark of rationality, rather than something awaiting rational explanation. It is the bare fact that the stimulus is "worth the audience's attention. Any utterance addressed to someone automatically conveys a presumption of its own relevance. This fact, we call, the *principle of relevance*. ... it is not something that they (the people) obey or might disobey; it is an exceptionless generalization about human communicative behaviour." (Wilson \ Sperber 1988: 140)

The authors' purpose is to find the rock bottom of communicative activity where a deviation from the norm comes to constitute the norm itself. No wonder that John Woods found this conception apocalyptic. If relevance theory is however aligned with typicality and topical argumentation, its rationale appears to be not so much the wielding of formal-logical structure – although Sperber \ Wilson do make such a claim – but rather the search for non-logical constraints on interpretation. Whether the constraints imposed by what is known include or not the utilization of demonstrative logic is a separate matter. As prototypical categorization represents a move away from taxomical systems, so do relevance theory – and other context selection approaches – make a step toward informal inferencing. That the idea of relevance in question leads to apocalypse in logic may well be true. Sperber\Wilson's real fault does not lie there. It lies rather in occupying two contrastive positions concerning rationality in cognition and in argumentative behavior. On the one hand they set the task to explain how communication even without an explicite code can become successful; that is how things can be inferred instead of being decoded. But if this is so, it appears on the other hand that what people in fact do is not understanding each other but rather conducting a monologue. In order to be otherwise, the speakers should be saddled with the extra burden of optimizing their talk *in such a way* that it facilitate the context selection by the hearers. To do that they should also be ascribed the mutual knowledge the pertinance of which Sperber \ Wilson argue against. Thus, however, we would soon be lead back to the original code model. And indeed, if the speaker were so keen on communicating the same idea, it would be more economic for her to use the latter than sending the hearer into an a-

mazinglabirynth of dubious and intricate – i.e. non-computable – inferences. Moreover, we have seen that, while we are more often than not optimizers as interpreters, we are quite nonchalant in producing proper behavior. So if the apocalypse is there, it is on the side of the speakers, not on that of the hearers. I will even venture to add that the more we are optimizers as producers, the more hard wired the given reaction becomes. In fact, as we will later see, *it is precisely because we ascribe the same optimizing rationality to others that we are prone to be nonchalant in producing behavior*. Sperber \ Wilson cannot have it both ways: retaining the rich inferential potential on the part of the hearers and securing the uptake of the communicative intent of the speakers. That is they cannot account for the fact that we are cognitive satisficers and productive optimizers at the same time. Yet that is what “the exceptionless generalization about human communicative behaviour” would require them to do. Else there is no rational explanation for language to have evolved.

4. The categorization problem

I illustrate the above point with a categorization problem. Thus the second reason for the methodological character of fallacy theories surfaces in cognitive psychology. Subjects are often tested for categorizing with a selection task in which they must perform pairings of figures and/or names, while it is the whole structure of training and testing they have undergone that should explain why they succeeded or not in their task. Yet it is highly dubious that the structure of the experiment correctly mirrors the structure of “inner” processing, i.e. the bridging between stimuli and output. In many cases “subjects are asked to provide a report under conditions where they would ordinarily not see anything meaningful. Knowing that the figure contains a familiar object results in a search for cues.” (Pylyshyn 1998) Still in other cases subjects must judge a statement like “A canary is a bird” either true or false. Such tasks are rather imposed on them and constitute “closed paradigms”. (Cf. Dubois 1991: 43) What psychological experiments are supposed to show is that the same principles that discriminate among the categories are also working within the categories themselves in producing prototypical effects. Thus – as Rosch puts it – there would be no sense in dissociating these principles. But since furthermore prototypicality is only a matter of best example within a category and not to be confused with the question of belonging, in many cases it seems to be enough if only the boundaries between categories (such as human and non-human, friend and enemy, etc.) are represented and the content either simply does not matter,

or if it matters, it matters only to the extent of delineating contrastive categories. Note that in such psychological experiments what goes on in the mind is taken to be mirrored by how the subject reacts to the target problem, that is by producing. Psychological testing reduces inner processes to simulation, that is to outward behavior and thus it commits the methodological fallacy of pulling down the distinction between interpretation and production. Such analyses are open to the criticism that representations are emptied out of content. By content I mean anything from feature-detection to nearest neighbour or averaged vectorial distance among affiliated items in connectionist networks. Representing boundaries may be as congenial (or conducive) to survival as ranking an instance within some category. Representing boundaries, however, implies that behavior relies so heavily on context that it is neither rule-based, nor similarity-based. It is not rule-based because it is an essential feature of rules that they are non-contextual. But it is neither similarity-based because, as e.g. Ellard reports, certain species “respond to *all stimuli* as threatening or to *no stimuli* as threatening depending on their familiarity with the context in which the stimulus is presented” irrespective of the local configuration of the stimulus, since there is an “obvious adaptive advantage... that it pushes the time-consuming and computationally expensive problem of stimulus recognition to a point in time that actually *precedes* stimulus onset.” (Ellard 1995: 681) In other words it does not imply structural mapping, but rather a pre-tuning to current context. I do not see any reason why such behavior could not appear to be significant in man.

A particularly interesting case is the experiment reported by Smith and Sloman who repeated a test by Rips to highlight the difference between the two categorization processes (similarity-based and rule-based). The task was to decide whether the test object with some characteristic attribute(s) belong to one of two target categories, of which one was fixed, while the other was variable with respect to the given attributes. (The attribute was shape falling in between the regular sizes of quarters and pizzas.) When there was only one such attribute, namely size (a round object 3-inch in diameter), most subjects judged that it belong to the category of pizzas rather than to the category of quarters. The explanation went that in case of boundary conditions subjects categorize on the basis of rules and rank the vague object with the variable category, while, and despite, noticing its similarity with the members of the fixed category. Whereas with the test object having more attributes similar to the members of the fixed category (e.g. silver color) subjects tended to judge it not only more similar, “but

also as more likely to belong” to the fixed category. (Smith et alii 1998: 182) This experiment however does not prove -as the authors want it - that categorization is similarity-based, since the attributes in question were necessary and/or perceptually *salient* features, which attest rather the application of rules. Experiments with boundary conditions do not show that people, if made to give all-or-none responses, indeed *represent* the test object as this or that. They rather show to the contrary that subjects are reluctant to tamper with represented boundaries, and so they temper with content: if presented with something conspicuously similar to the target object, they adjust, or temper with, the precise “rule” of what belongs to that category. Note also that such experiments completely disregard the role of context. How would subjects decide if the test object is presented to them within a restaurant or buy-and-sell frame?

Thus we reach the conclusion: the fact that people follow rules in their behavior above - behavior in processing stimuli - is a phenomenon resulting from the contrived character of the situations they are tested or observed. There is nothing like inherently normative here. It is rather that the horizontal organization of categorial structure appears to be far more *relevant* to selective action than the vertical structure. To sum up:

(T1) Human categorization is such that it reflects the evolutionary important boundaries among the objects of environment, but there is no objective mapping between the content of coded categories and external reality. (Cf. Pólya \ Tarnay 1997)

Coded boundaries may naturally shift with evolution, hence there is objective necessity for the semantic transparency of the boundaries themselves. Yet it is crucial that there be observed boundaries, which can be reflected linguistically as well.

5. Normative vs descriptive: rule-governed vs similarity-based

Thus we are confronted with contrastive evidences or conflictive demands: on the one side we have experimental results in developmental psychology, pathology and animal behavior which attest of high contextuality and dispositionality in behavior; hence they point to similarity-based rather than rule-governed behavior. Yet - and this is partly my point here - they *appear* to be rule-following to the external observer since - at most - coding of category boundaries may be inferred in certain cases. Furthermore, it turned out that prototypical categorization prompted by E. Rosch and her followers frequently mirror prior training and

external activity rather than the inner structure of representation; thus typicality should also be ranked here, which accords well with the fact that they are similarity-based.

On the other side we have the topoi or argumentative inference conceived along the lines of J.-C. Anscombe and O. Ducrot. By all means inferential activity implies rule-following, hence it cannot exclude normativity in its entirety. Given the rhetorical nature of language it arises that the scope of inferential activities cannot be wholly captured by a theory of relevance as Sperber \ Wilson want it. Yet it must have also become clear that their theory occupies a middle position in my ranking in that for them context selection is primary and similarity based, while it is only fuel and/or input to the main operation: the producing of contextual effects by means of - demonstrative - rules.

Suppose for the moment that the picture linguists and psychologist with an argumentative bent is close to the truth. Suppose furthermore that it is the best explanation one can offer of what goes on in the hearer's mind. Then we have a blatant inconsistency. When we interpret we are cognitive satisficers, that is we try to extract with the least effort as much content as we can from what has been said. In other words we set our aims too high: we strive to construct a distinctive - fine-grained - picture of the world on the basis of structural and inferential relations between incoming new and retrievable old information. But when it comes down to responding or (re)acting, unless we are *rationalized* experts - we observe only the most "relevant" - coarse-grained - boundaries of our cognitive structure.

Whence such an inconsistency? I have already hinted at one possible answer: evolution driven selectivity. This may well cover low-level - dispositional - action. But I have presented high-level, categorical thinking very much like autonomous, similarity-based action. Can I be justified in making that move? Now here is the source for a second answer, quite orthogonal to the first; it is the principle of charity proposed by Donald Davidson. It says briefly that in evaluating the speakers' behavior we aim at giving the best possible explanation of their - linguistic - behavior. That is we rationalize their activity. At face value, rationality is not an ideal by which we automatically assess their action, but it is rather a result of our interpretative activity. The question is: Can we reconcile the principle of charity with the principle of relevance or argumentative normativity? At first sight it seems yes, since both approaches aim at at a full-blown interpretation of utterances, at exploiting its inferential potential, at resolving conflicts, etc. This latter task most often amounts to supplying missing premises.

But on what basis should such premises be determined? On the argumentative approach, it is a set of agreed upon rules – either semantic or pragmatic – that constrain both interpretation and possible responses (i.e. speech acts). Violations of such rules would then naturally amount to committing a fallacy. But if so, most argumentative-communicative situations are doomed to break down. For what if the best possible explanation of inconsistent or incoherent speech behavior comes from unique or “irregular” sources of the situation in question, from ideosyncratic aspects which are not given or statable once for all. What if the point of a semantic or pragmatic rule consists precisely in tampering with, or manipulating, it? This is a moral to be learnt from oral communication, ordinary and artistic, in primitive and as well as higher cultures. But a moral also rendered by categorization testing in experimental psychology when it is acknowledged that “a change in the activation level of a feature has the effect of changing the criteria of arbitrarily many categories into which that feature could enter, including ones that the investigator may have no interest in or may not have thought to test.” (Pylyshyn 1998)

If we take the principle of charity in an argumentative vein we have our second answer: we are nonchalant in our behavior *just because* we ascribe the same kind of rationality to others. We suppose there is a rock bottom of rule-following, some abstract set of rules upon which agreement must sooner or later be reached. It is an overgeneralization: *an extrapolation of external behavior onto the domain of what goes on in the head*. But it is just this supposition of general rationality that appears to be a fallacy as soon as we take content seriously. Inconsistency may not be right word to apply to what is meant here: tampering with the rules may well be just another metaphor of constantly jostling the boundaries of our inner categorical structure. Redrawing the horizontal structure of our categories cannot be made to follow some pre-set rule, it cannot be normative. There may well be external constraints originating with the changing of our environment, but there is no direct internal response to that change; cognition has its own plasticity but it is essentially constrained by its former structure. If relevance bears any selective advantage, it is in (re-)utilizing “cognitive parts” as building blocks already there rather than starting anew. (Cf. The Gouldian idea of evolution as assembling old parts together – ones adapted to a previous purpose – for a new purpose.)

This may be taken to be a stretched – even a too charitable – interpretation of the principle of charity to cover cases of blatant inconsistency. Yet I think it is not. I agree with Z. Pylyshyn that inference is an activity “where the semantic property

truth is preserved. But we also count various heuristic reasoning and decision-making strategies (e.g. satisficing, approximating, or even guessing) as rational because, however suboptimal they may be by some normative criterion, they do not transform representations in a semantically arbitrary way: they are in some sense at least quasi-logical. This is the essence of what we mean by cognitive penetration: It is an influence that is coherent or quasi-rational when the meaning of the representation is taken into account.” (Pylyshyn 1998) The use of term “rational” is meant to indicate that in characterizing such processes we need to refer to what the beliefs are about – to their semantics. The important point is that such processes can be suboptimal. I think Pylyshyn hits the right note when he asserts that “most psychological processes are cognitively penetrable, which is why (cognitive) behavior is so plastic and why it appears to be so highly stimulus-independent.” Hence cognition is both stimulus-independent and meaning dependent. That could well be the reason of its suboptimality. Suboptimality does not mean however that cognition is not task-centred as most cognitivist conceive of it. But that is not enough reason to treat motor and processing activity on a par. The difference may be analogous to that between “systems that have constraints on interpretation built into them that reflect certain properties of the world” (Pylyshyn 1998) and systems that access and use knowledge. While the first is cognitively impenetrable, the second is not.

To sum up: higher cognition may *appear* rule-governed to the extent that it is stimulus-independent and meaning-preserving in exploiting more or less abstract structural correspondences. Even so, even if it is cognitively penetrable, it cannot become normative, since it always works with used materials.

6. Concluding remarks

Let me conclude with giving vent to a good and a bad consequence. The good one may be that the black box of the mind has not been wholly and adequately opened yet, so there is much work to be done in this field. It is plausible that man is capable of high-level cognitively penetrable activity, of understanding complex relation structure, etc. I take it to be part of the good news that such ability is highly plastic, and even if abstract, it cannot be ranked with rule-following. Quite clearly so because it involves analogical thinking which has much in common with primary similarity-governed processing. But there is the bad news. It starts with the simple observation that if communication (and cognition) is task-centred, then an important part of it must be constituted by the attempt of securing the uptake and the “correct” interpretation of any utterance. Otherwise – and in lack of some

other meaning independent social function – selection should have driven it out. But should it? If what has been said of production is only partially true, we are surrounded with a huge mess of carelessly formulated and misfired talk and misunderstanding. How is it that selection has not already driven out our communicative ability? I can give two brief answers here. The first is simple and a bit cheating. It runs that the evolutionary story of literacy is too short to be a proof of its selective advantage. The second is more complex but I cannot elaborate it here. It starts by seemingly overturning my argument in this paper in that it claims that what we are almost smothered with is not a mess of misfired talk, but rather a “cognitive” technology, factories of ideologies, which not only reproduce the same forms of talk like e.g. the ads, but they self-reproduce as well. That is they do not overturn the communicative function but overexercise it. So far so good. Communication should not be wiped out then. But there is a corollary to this answer: the overarching function of exact communication will result in the wiping out of the meaning-dependent and rule tampering cognitively penetrable higher activity, since any tampering with the rules slows communication down or may even end up blocking it completely. But once again our past is only a drop in the evolutionary ocean. So we are stuck with our morsel of hope.

REFERENCES

- Anscombe, J.-C. \ O. Ducrot (1989). Argumentivity and informativity. In: M. Meyer (Ed.), *From Metaphysics to Rhetoric* (pp. 71-87), Dordrecht/Boston/London: Kluwer.
- Dubois, D. (1991). Categorisation et cognition \ 10 ans apr\ : une évaluation des concepts de Rosch. In: D. Dubois (Ed.), *Sémantique et cognition. Catégories, prototypes, typicalité* (pp. 31-54), Paris: Éd. du CRNS.
- Dubois, D. \ Ph. Resche-Rigon (1996). Champs topiques, catégorisation et expertises. In: P.-Y. Raccach (Ed.), *Topoi et gestion de connaissances* (pp. 27-40), Paris/Milan/Barcelone: Masson.
- Ellard, C.G. (1995). Context and Consciousness. *Behavioral and Brain Sciences* 18: 681-2.
- Millikan, R. G. (1997). A Common Structure for Concepts of Individuals, Stuffs, and Real Kinds: More Mama, More Milk and More Mouse. *MS. a Behavioral and Brain Sciences target article*.
- Millikan, R. G. (1997). Truth Rules, Hoverflies, and the Kripke-Wittgenstein Paradox. *The Philosophical Review* 99: 323-353.
- Pólya T. and L. Tarnay (1997). Is Context Really a Problem? *Proceedings of*

European Conference on Cognitive Science Manchester UP. 199-202.

Pylyshyn, Z. (1998). Is Vision Continuous with Cognition? The Case for Cognitive Impenetrability of Visual perception. MS. a *Behavioral and Brain Sciences target article*.

Racah, P.-Y. (1993). Quelques remarques sur la sémantique linguistique et la construction du sens. *Travaux de Linguistique et de Philologie* (TRALIPHI) 32.

Smith, E.E. et alii (1998). Alternative Strategies of Categorization. *Cognition* 65: 167-196.

Sperber, D. \ D. Wilson (1986). *Relevance: communication and cognition*. Oxford: Basil Blackwell.

Wilson, D. \ D. Sperber (1988). Representation and Relevance. In: R. M. Kempson (Ed.) *Mental representation: the interface between language and reality* (pp. 133-153), Cambridge: Cambridge UP.