

COGNITION AND EPISTEMIC CLOSURE

Radu J. Bogdan

I

JUSTIFICATION and knowledge are thought to be closed under known implication..¹ This widely shared assumption is embodied in the following principles of epistemic closure:

(CJ) If *S* is justified in believing that *P* and *P* implies *Q* and *S* realizes that *P* implies *Q* and accepts *Q* as a result, then *S* is justified in believing that *Q*.

(CK) If *S* knows that *P* and *P* implies *Q* and *S* realizes that *P* implies *Q* and accepts *Q* as a result, then *S* knows that *Q*.

These two principles do a lot of work in the analysis of knowledge, the defense of skepticism and the formulation of various constraints on justification and knowledge. A number of philosophers, very few in fact, have recently criticized (CJ) and (CK).² My aim here is to contribute a new sort of criticism, from a different angle, that of our cognitive functioning. My conclusion will be that, as standardly construed, (CJ) and (CK) are invalid for cognitive reasons that have epistemological repercussions.

My argument will run as follows. The key cognitive concept in our discussion is that of *belief*. The notions of justification and knowledge, on the

other hand, should be construed as *evaluative* in the sense that they ascertain and measure- the epistemic worth of a belief. If this is granted, then (CJ) and (CK) hold only if the following closure principle for belief also holds:

(CB) If S believes that P and P implies Q *and* S realizes that P implies Q and accepts Q as a result, then S believes that Q.

This last principle needs cognitive modulation. The reason is this. We believe under specific mental categories, some (concepts) more optimal than other (ideas). So we can distinguish between sorts of beliefs, according to their *categorical profile*. I will argue that the categorical profile of a belief shapes the role that belief plays in cognition and thus contributes to the (opaque) determination of the cognitive type identity of the belief. The categorical profile of a belief is therefore constitutive of its cognitive type identity. This is bound to affect our understanding of both doxastic (CB) and epistemic closure (CJ and CK).

In a nutshell, the argument is this. Suppose an inference enables one to see when a belief implies another. Suppose the inference does not necessarily transfer and preserve the categorical identity of the beliefs involved. Then, in spite of the implication, the categorical -profile of the initial belief may be different from that of the resulting belief. (CB) fails to reflect this. The categorical profile of a belief will also be shown to affect its epistemic worth. We thus have all the needed ingredients of a recipe for constructing counterexamples to, and so showing the invalidity of, (CJ) and (CK). The recipe is this: Start with a belief that P whose categorical profile is such and such, make sure that in some plausible way the believer realizes the implication from P to Q and accepts its outcome but have him end up with a belief that Q the categorical profile of which is different from and inferior to that of the initial belief that P. This will make the two beliefs not only

cognitively type distinct but also epistemically different. It follows that a known implication does not always preserve the cognitive type identity and the epistemic worth of the beliefs it engages, which is what traditionally the notion of epistemic closure was thought to imply. (CJ) and (CK) must therefore be wrong. So much, in anticipation, for the outline. Now to the details.

II

I need first a few preliminary notions and distinctions. When, in stating the closure principles, I say that *S realizes that* P implies Q, I construe the concept of mental realization functionally in the sense that I intend it to cover a variety of ways in which S may come to see the implication, whether by analogy, intuition, some form of reasoning, an explicit, rule-governed argument, visual imaging, whatever. But I assume that in general the realization of an implication, its mental accessing, takes the form of some thought process or inference. Furthermore, I understand the notion of *acceptance* in a very narrow sense here, as a sort of belief merely generated by and entirely geared to the realization of an implication (or of some other logical relation). I construe the notion of *belief* proper more generally, as a mental attitude generated by and sensitive to factors other than logical or conceptual such as evidence, relevant information, epistemic and practical goals and so on. So, on this distinction, when one realizes an implication, the conceptual factors involved (i.e., what is involved in the relevant semantic network available to the believer) may provide a sufficient reason to *accept* its outcome, although other sorts of facts may be needed to invite *belief*. When an implication, or any other logical relation, is both mentally realized or

accessed and accepted, we say (again in a rather narrow sense) that it is *known*.

The notion of *implication* itself is understood here in terms of (necessary) consequence. Thus P implies Q iff Q is a (necessary) consequence of P . This is to be contrasted with the much narrower, purely formal or syntactic construal, better called *entailment*, on which P entails Q in virtue of logical form alone. Entailment will not concern us here. Unlike entailment, implication is not self-sufficient since it typically requires further data (in the relevant semantic networks), such as meaning postulates, conventions or theoretical assumptions, in order to hold. Thus, for example, whereas " $3 + 2 = 5$ " entails " $2 + 3 = 5$ " by virtue of logical form, "A is a physical object" only implies "A has a spatio-temporal location" when some appropriate meaning postulate or assumption about the concept of physical object is also activated. When, therefore, S realizes and accepts an implication, it should be understood that S does so relative to whatever it takes (meaning postulates, assumptions, guesses, etc.) to represent that implication.

I need next a distinction between *concept* and *idea*. Following Armstrong³, this could be seen as a distinction between a capacity and its exercise. Thus, to be able to speak French is to have a capacity; to utter a French sentence is to exercise that capacity. In this sense a concept would be a capacity to have or form ideas. But one may utter a French sentence or even engage in some limited conversation, and know what one is saying, without being in general able to speak French. This is a possibility I need for my argument. For what I want to show is that one can have and operate with ideas without having the corresponding concepts. To account for this possibility I will distinguish between having ideas as first-order capacities to represent and having concepts as second-order

capacities to represent. Ideas represent *particular configurations of* objects and properties. Concepts represent the *objects and properties* themselves, whatever the particular configurations which instantiate them. A concept, then, is an idea-tic invariant. A concept is a more optimal category than an idea.

Concepts typically involve ideas in their applications to particular cases. This sounds like a Berkeleyian theme. It makes psychological sense. We do appear to have privileged representations, of some concrete sort, which act as *prototypes* for the class delineated by the concept. Ideas *are* these privileged representations. Often ideas can do their job without concepts. This also makes psychological sense. First, there is learning. In the early stages of concept formation, one's first exposures to the instances of a concept are retained as ideas. If, for some reason, the invariance required by the concept is not internalized, one is left only with ideas. Then there is also the maturation of some mental competence, a process which typically goes from ideatic to conceptual grasp. For example, before a certain age, children can correctly judge some volume relations but misjudge others. They have some idea of volume, fitting particular contexts, but not yet the concept. We find a similar story in language learning. In the beginning and for a while, the meanings of many words are merely ideatic. The learner tends to use those words in sentences which mimic those which first introduced him to the meanings in question. One cannot take risks with ideatic meanings. In other words, one cannot vary the linguistic configurations at will and still control a meaning, if that meaning is ideatic ⁴. Quite frequently, ideas are also formed on a more adhoc basis. Suppose that I do not have the concept of quark, which is true, and so I decide to look for an explication in a science magazine for laymen. I take a copy of *EASY SCIENCE FOR HUMANISTS* and look at a few diagrams and simple explanations populated with such entities as electric charges, hadrons, gluons, and the like, some of which I already understand

conceptually, others only ideatically. As a result, I form a specific and rather poor idea of what a quark is. Do I have *some* idea of quarkness? I surely do. Do I have the concept? I don't.

Finally, another notion I need is that of *ideatic redescription*. The elements for it are already in place. What that article about quarks did for me was to ideatically redescribe or recategorize the concept of quark in terms of other concepts and ideas already available to me. This is, generally, how learning goes. It is a particular ideatic redescription provided by that article that, until further developments, is going to constitute my idea of quarks. Many of my scientific ideas have been formed this way and remain, I am afraid, just that, ideas. This should provide ideas with enough psychological reality.

III

We identify beliefs in terms of their *role* in cognition and behavior. The cognitive role of a belief is manifest in the inferences it allows us to make, the thoughts we entertain, the other beliefs we can form. Concepts and ideas are basic ingredients of our beliefs: We believe under some concepts and/or under some ideas. This seems axiomatic. Indeed, most of our beliefs are (sets of) concepts and ideas deployed or applied. It is therefore natural to expect that the distinction between concepts and ideas will be reflected in the *types* of beliefs we form. The roles that concept-driven, or conceptual, beliefs play in cognition will be systematically different from the roles played by idea-driven, or ideatic, beliefs. I have some idea of quarks which, when the occasion demands it, participates in and is formative of many beliefs I have about quarks. The physicist, we may assume, has the concept of quark which also participates in and is formative of his many beliefs about quarks. But we do not want to say that, type-wise, his beliefs about

quarks are like mine. They cannot possibly be, because his beliefs play a vastly and systematically different role in his cognitive life than my beliefs about quarks play in mine. The physicist can draw inferences that I simply cannot, can read the same evidence in ways which are inaccessible to me and can make mental associations which escape me entirely. If beliefs are determined by such mental roles, and mental role is shaped by available categories, then dramatic differences in both role and categorial profile should be captured by an equally dramatic difference in the overall cognitive identity of the beliefs involved.⁵ So the following stipulation seems in order:

(BEL) To believe_c that $P \#$ To believe_i that P

(where $\#$ stands for cognitive type-difference and the subscripts indicate whether one believes under concepts or ideas, respectively).

(BEL) should be taken to indicate that, via mental role, the categorial profile of a belief is constitutive of its cognitive type-identity. A particular consequence of (BEL) is crucial to our argument. Suppose that S believes_c that P . Suppose also that P implies Q , S realizes the implication and accepts Q , yet his mental grasp of Q is only ideatic. So S believes_i that Q . If such a doxastic scenario is logically possible, and in fact psychologically real, the (CB) is untenable since it assumes the categorial and hence cognitive type-identity of the beliefs in the antecedent and consequent clauses. This result in turn affects the principles of epistemic closure.

Indeed, the cognitive type-difference between conceptual belief and ideatic belief *must be* reflected in their epistemic worth. The reason is this. Concepts can represent what ideas cannot, namely, invariant properties. We show mastery of a concept when we pass the invariance test: the configurations instantiating a target

property are varied and yet that property is still identified as such. Ideas cannot pass the invariance test because, as earlier indicated, it is in their nature to be tied to *particular* configurations which are experienced as exemplifying the target property. Our intuitive judgments seem to assume that a belief cannot be regarded as conclusively justified or as knowledge unless it is a conceptual belief. This makes epistemological sense. Suppose that my idea of X identifies X only in terms of a subset of properties which also happen to be shared by Y. X and Y fail to share many other properties but my idea of X cannot register this. So there will be circumstances in which, faced with an Y, I will take it for an X. Even when I am faced with an X and form the belief that it is an X, I cannot be said to be conclusively justified in my belief or to know that it is an X, for if it were an Y I would still take it for an X. Nobody can escape this ideatic predicament. Most of my birds and trees representations are of this sort.'

So we want to codify all this by saying that a conceptual belief cannot share the same epistemic worth with an ideatic belief. In particular, they cannot share the same justification even when they are both supported by the same evidence. Other things being equal, justification of a conceptual belief is stronger than and qualitatively superior to that of an ideatic belief. Furthermore, conclusive justification as well as knowledge require conceptual belief. In short, other things being equal,

- (JUS) (a) Justification of belief_c that P # Justification of belief_i that P.
- (b) Justification of belief_c that P > Justification of belief_i that P.
- (c) A belief is conclusively justified only if conceptual.

(KNOW) Knowledge requires conclusive justification.

We are now in a position to give a general formulation of our main argument. Suppose that S is conclusively justified in believing that P. There is an implication from P to Q which S sees and accepts. But whereas his belief that P is conceptual, his resulting belief that Q is only ideatic. This means, according to (BEL), that the latter belief is cognitively type-distinct from the former. According to (JUS), their respective justifications are qualitatively different, with that of the former much stronger than that of the latter, other things being equal; the resulting belief that Q, if justified at all, cannot possibly be conclusively justified. Conclusive justification is therefore not closed under known implication. (CJ), then, must be invalid. A parallel argument works for knowledge and (CK). The notion of epistemic closure, as standardly construed, is therefore inconsistent.

IV

I will now illustrate the argument with three examples. All are examples of the format conceptual belief-known implication-ideatic belief but in different contexts of cognition, one mathematical, another phenomenological, the third empirical.

A Mathematical Example. Suppose that Sam is a mathematician who not only knows that $6 + 3 = 9$ is true but also knows, expertly, why it is true. He can prove it. Let us conjoin the proposition that $6 + 3 = 9$ with Sam's relevant arithmetical knowledge and call the result P. So Sam knows that P. Suppose now that the following proposition Q is true: $6 + 3 = 9$ is an apriori true proposition. Sam does not know what an apriori true proposition is. When his friend Jim, the philosopher, mentions Q in their discussion Sam does not know what Jim is talking about. Jim tells Sam that Q follows from Sam's knowledge that P. But this does not help Sam since he does not have the concept of apriori. Patiently,

Jim explains to Sam in ideatic terms, appealing to other things Sam is familiar with, that $6 + 3 = 9$ is an instance of an apriori truth. Other instances are also produced. As a result, Sam forms an *idea* of aprioricity, one tied to the particular matrix of Sam-accessible examples and concepts that Jim has rather didactically put together. Under that particular ideatic redescription Sam has come to realize that Q follows from P and to accept Q .

Does Sam *know that Q*? I think not. For suppose that, under the description given by Jim, Sam's idea of aprioricity happens to share some features with the concept of analyticity, which Sam does not have either. Sam would not therefore know the difference between the two and would easily mistake one for another, even when introduced to some idea of analyticity. Also, given his limited ideatic access to the notion of apriori truth, Sam may well fail to identify an instance of an apriori true proposition in non-mathematical contexts. His belief that Q cannot therefore qualify as knowledge. (CK), then, must be invalid because the knowledge that Sam has, that P , is not *sufficient* for the knowledge (CK) requires him to have, namely that Q . Sam also needs the appropriate concept for Q , namely aprioricity. On the other hand, the latter concept turns out not to be *necessary* for Sam's initial knowledge that P . Sam may know a particular mathematical truth without necessarily knowing some further property implied by it. There are islands of knowledge.

A Phenomenological Example. Each of us has phenomenological or introspective knowledge about various conscious experiences such as images, memories, pains, and so on. We form beliefs about them. Let us call them *phenomenological beliefs*. Although inherently accessible under a mode of presentation, our conscious experiences generate phenomenological beliefs about them which are conceptual. Thus we can always recognize an act of remembering or perceiving,

and not confuse them with other kinds of mental acts. We also form justified beliefs about the contents of these acts which we can then redescribe in various ways. Also, the fact that we can publicly discourse and exchange reports about our conscious experiences is further indication that the phenomenological beliefs whose contents are thusly made public are governed by concepts; for if these beliefs were merely ideatic, that is, tied only to particular experiences, the possibility of public description would be drastically reduced to something like an ostensive or demonstrative grunting and gesturing.

This being said and granted, we can make the next move. It consists in telling our subject S that, for any conscious experience she has, there is an underlying physical process which is necessary for that experience to occur. Semantically, this amounts to making her aware of an implicative connection between her representations of conscious experiences and those of some causative physical processes. In particular, S is told that perception implies transduction. S does not have the concept, nor the slightest idea, of what transduction is. She is given a particular redescription of transduction which, together with other data (concepts and ideas S already has), also makes visible the implication from perception to transduction. So S accepts the implication and comes to believe, ideatically, that she undergoes a process of transduction whenever she perceives something. Can her belief about transduction count as knowledge the way her introspective belief about perception does? Or can the former be as justified as the latter? Clearly not. At least that much phenomenalism has taught us.

An Empirical Example. Suppose S1 is more justified than S2 in believing P. Suppose also that we generate from P a large class of implications. Let Q be the name of this class. Suppose further that there is a subset of Q, call it Q*, itself very large, which contains statements for which both S1 and S2 are conceptually

unprepared. We can think of the members of Q^* as describing theoretical developments concerning P (in the distant future) with which SI and S2 are totally unfamiliar. But they are able to reach and grasp some of the Q^* statements ideatically, after realizing and accepting the implications involved, under suitable redescriptions. Now let us imagine that we can regiment all the statements involved in a manageable first order language and that we have a procedure for measuring comparative justification, perhaps one derived from Camap's inductive logic. This procedure is sensitive to conceptual as well as ideatic categorization. It not only rewards one with so many points for each Q-statement grasped conceptually but also subtracts points for any which is grasped only ideatically. Finer details should not matter at this point. Let us call this process the Q^* -test. Our example is set up in such a way that SI and S2 will lose lots of points within Q^* . Since, by assumption, Q^* is extremely large, sooner or later SI and S2 will both lose so many points that their resulting justifications in believing particular Q^* -statements will either converge toward zero or else will become vanishingly small and practically equal.

According to (CI), if SI and S2 start with different justifications, they should end up with equally different justifications; if, as above, they do not, then either both did not have any justification in believing P in the first place or else both had practically the same minute initial justification in believing P. But this contradicts the assumption of our example. (CJ), then, cannot account for such a case. The alternative I am suggesting is to allow for islands of justification. This, on my account, is possible because of the categorial and hence cognitive type-distinctness of the beliefs involved, a feature which justification must reflect but which the standard (CJ) does not. To illustrate our story, a biologist, I take it, would be more justified than a layman in believing today a particular scientific statement about some brain function. Imagine now that both are transported into

the distant future when neuroscience is vastly different and superior. If repeatedly subjected to a Q*-test by those future neuroscientists, both the biologist and the layman may well end up with what our thought experiment predicts: vanishingly small and practically equal justification for their implicatively-arrived-at ideatic beliefs. Yet, initially, the biologist believed something about brain functions with a given justification while the layman believed, the same thing with much less justification. This does not make sense if standard epistemic closure holds.

V

It is now time to review the argument, mark its limits, explore some aspects left implicit, and answer some objections.

First, its *limits*: What the argument does not show is that the closure principles for belief, justification, and knowledge never apply. They do apply quite often, whenever the logical relations and transformations involved connect beliefs which are cognitively type-identical, that is, among other things, have similar conceptual or ideatic profiles. In such cases, other things being equal, epistemic worth is transferred and preserved. But, as shown, there are also cases where beliefs and their justification are not closed under known implication (or other logical relations) because of doxastic type-differences due to categorial differences. This is enough to show the invalidity of the standard closure principles.

The constraint that emerges from our discussion is that epistemic closure requires doxastic closure, the latter requires the cognitive type-identity of the beliefs involved, and that type-identity in turn requires categorial identity. In particular, the closure of conclusive justification and knowledge requires that the beliefs involved be conceptually type-identical. Therefore, an adequate

reformulation of (CJ) and (CK) will have to allow only conceptual beliefs. Since even ideatic beliefs have *some* justification, a more liberal (CI) may also be needed. But these are details that go beyond the critical ambitions of my argument.

It should be noted that, by itself, my argument against epistemic closure does not necessarily invalidate particular analyses of knowledge or justified belief, nor should it be construed as an argument against skepticism. These consequences would follow *only if* such analyses and skeptical positions take epistemic closure as being analytically or intuitively valid. In other words, if the concept of knowledge has epistemic closure built into it, then I take my argument to invalidate that concept of knowledge. Likewise, for justified belief. The overall moral is that analyses and arguments which rely on and exploit epistemic closure would have first to go through the motions of empirically establishing its ingredients (such as categorial profile, then doxastic type-identity and the rest) before claiming any general validity.

There is a critical twist to my argument that must not go unnoticed: It is the CLOSURE of conclusive justification and knowledge that requires the conceptual type-identity of the beliefs involved. The very notions and attributions of conclusive justification and knowledge do NOT require closure. In other words, epistemic worth itself does not require epistemic closure, although its inferential transmission and preservation does. It has been the thrust of my argument that epistemic worth can be attributed locally and autonomously. That attribution is not hostage to the requirement of epistemic closure. One can thus know or justifiably believe something without necessarily knowing or justifiably believing its implications.' This is what I meant when saying that there are islands of justification and knowledge. The concepts of epistemic worth and epistemic

closure have therefore been pulled apart. This is a result I find not only psychologically motivated but philosophically liberating as well.

Many philosophers, I suppose, would not. One very likely objection they are going to raise is that S cannot be said to know or justifiably believe that P if, after realizing and accepting the implication from P to Q , S ends up with an epistemically inferior belief that Q . The objection begs the very question at issue here. It assumes what it has to prove, and what I have tried to disprove, namely, that knowledge and justification cannot be attributed to a belief unless that belief is closed under known implication. In other words, this objection assumes that epistemic closure is built into epistemic worth.

Consider, again, our mathematical example. The objection is that Sam does not really know a particular arithmetical truth, for if he knew it, he would also know it to be an a priori truth, given of course that he is made aware that the latter is necessary for the former. My position, on the other hand, is that the concept of a priori truth belongs to a theory of (mathematical) truth and that Sam does not need to know that theory to know a particular mathematical truth when he sees one. He does not need metamathematical knowledge to have some specific mathematical knowledge, just as most of us need not know the theory of truth to know when a particular sentence is true. There are many other things that Sam may not know *about* mathematical truths, for instance, that they are analytic, innate, objective and so on, features implied by some further reflection, at *different* level of abstraction and categorization, on what he already knows. Lack of such further reflection in no way cancels his initial, more particular knowledge.

The phenomenological example illustrates the same point from a different angle. Our phenomenological beliefs have access to mental phenomena at a certain level of abstraction, as qualitative contents of conscious experience. As

introspective abilities to discriminate and categorize kinds of conscious experiences, our phenomenological concepts tend to gravitate toward the same level. One can give a coherent neuroscientific explanation of all this, an explanation which has to do with the kinds of information processing systems we are and the ways we think and communicate about ourselves and others. Our mental experiences are made possible by various underlying neurochemical processes and states, but one does not *need* to scientifically know the latter in order to- phenomenologically know the former, just as one does not need to know the microphysical composition of a table to know what a table is, what it is good for and so on.

The moral, then, is that concepts and beliefs operate at certain levels of abstraction (or relative to given semantic networks) where they are epistemically evaluated. An implication can form a new belief at another level of abstraction and categorization (or relative to another semantic network), necessary or otherwise logically connected with the first, yet one for which the believer, for whatever reasons, is only ideatically prepared. The epistemic worth of his new ideatic belief will be lower than, but compatible with, the higher epistemic worth of his initial conceptual belief.

In addition, the third example shows why a *local* understanding of beliefs, their categorial profile and epistemic value is needed. We want differences in learning and expertise as well as various stages of scientific progress to be reflected in our epistemic evaluations. Such local variations are often caused by differences in categorial profile. If the standard construal of epistemic closure is not sensitive to categorial differences, it cannot be- allowed to govern our epistemic evaluations.

Another likely objection may be directed at my understanding of implication and its mental realization. It may go like this. If S understands P

under the right concepts and forms the belief that P, then given his realization that P implies Q and acceptance of Q, why shouldn't all this be *sufficient* to provide S with an equal conceptual grasp of Q and hence with the corresponding belief that Q?

This is really an objection from analyticity: It appears to assume that an implication from P to Q deploys only information already contained in P. By hypothesis, all the information in P is accessible conceptually, even if parts of it are often left implicit. What the implication does is to access some implicit part and thus make it visible. If this is so, then how could S fail to also believe that Q? The answer is: This is *not* what implication, as construed here, does. An implication is very often accessed when a given semantic network is expanded, or connected to another, by appeal to further data (meaning postulates, hypotheses, whatever). In such a case, an implication transcends the boundaries assumed by the objection. This is precisely why the realization and acceptance of an implication is often bound to involve ideas. But then neither the realization nor the acceptance can be analytic. One can of course *stipulate* that implication should operate analytically. Assuming that such a stipulation makes sense, it represents a further retreat from the original ambition of having (CJ) and (CK) cover the entire territory of cognition. It has been the systematic purpose of my argument to force such retreats and thus drastically narrow the scope of epistemic closure.

To the likely challenge that my admittedly tentative account of mentally accessing (or realizing) and accepting an implication does not carry enough weight, my answer is pragmatic: I agree, but I want to add that my account is better than nothing and in fact goes in the right direction. To throw the ball in the critics' court, I would note that the epistemological literature concerned with epistemic closure contains almost no significant effort I know of to spell out

what it takes *cognitively* to know or see or be aware of P implying or entailing Q. Yet it is fair to assume that most people who talk of S knowing or seeing an implication take S to possess *the concept* of implication. That would be a mistake. One need not have the concept of belief to believe or to recognize believing, and likewise one does not need the concept of implication to follow one. An idea of implication will do. The argument of this paper does not depend on the strength of S's mental grasp of implicate on, nor does it have to evaluate the epistemic worth of that grasp. One may again want to stipulate that one should have the concept of implication. But this is just one more retreat, for epistemic closure now appears to require expert knowledge of logic.

Finally, one may ask, shouldn't my argument be sensitive to the distinction between a normative and descriptive account of justification and knowledge? Since epistemic closure is a normative constraint on justification and knowledge, why should a descriptive account of epistemic failure like that of ideatic belief be allowed to count against epistemic closure? After all, it is no news that actual cognizers often fail to comply with all sorts of normative constraints, including closure. The very distinction between conceptual and ideatic beliefs indicates how fallible actual cognizers are. End of objection.

My argument is indeed insensitive to the normative/descriptive distinction, whatever that distinction means. There is a good reason for that. The argument only maintains that what we characterize or evaluate in the consequent clauses of (CJ) and (CK) is *not necessary* for what we characterize or evaluate in their antecedent clauses, *no matter how* we construe the clauses themselves, normatively, descriptively, whatever; and conversely, for sufficiency. In other words, the argument deals with *the scope* of our attribution of epistemic worth, not with the nature or character of what is being attributed. The argument is that that scope is *local*: there are islands of justification and knowledge.

Epistemic closure itself works only locally, when specific cognitive assumptions are being made, as our discussion has amply shown. To universalize epistemic closure while subscribing to its cognitive relevance is to invite invalidity. To universalize epistemic closure while preserving its validity is to invite cognitive irrelevance.⁸

Received December 20, 1983

NOTES

1 Almost all papers since 1963 which deal with the Gettier counter-examples make this assumption. For a recent sample. G.Pappas and M. Swain (eds), *Essays on Knowledge and Justification*, (Ithaca: Cornell University Press, 1978).

2 For example, F. Dretske, "Epistemic Operators," *Journal of Philosophy*, vol. 67 (1970), pp. 1007-1023; 1. Thalberg, "In Defense of Justified True Belief," *Journal of Philosophy*, vol. 66 (1969). pp. 798-803-, and R. Nozick, *Philosophical Explanations*. (Cambridge: Harvard University Press, 1981).

3 David Armstrong. *Belief, Truth and Knowledge*, (Cambridge: Cambridge University Press, 1973), P. 52. Armstrong develops a distinction originally made by Geach in *Mental Acts*, (London: Routledge, 1957).

4 This linguistic discussion ties in with what Hilary Putnam has said about meaning. My notion of idea can be related to his notion of stereotype: a minimal level of competence required for having a word. An idea (which may be inaccurate) of what an X looks like or acts like or is. See his *Philosophical Papers*, vol. 11, (Cambridge: Cambridge University Press, 1975), P. 249. This view allows us to see how, linguistically, one may have an idea or stereotype and be right about what it represents without necessarily having the expertise, hence the conceptual categorization required.

5 To type-identify beliefs cognitively, that is, in terms of their role in cognition is not the only way to type-identify beliefs. One can also type-identify beliefs,

transparently, in terms of their propositional objects or their truth-conditions. In this paper I will ignore these other forms of belief type identification. My critique of the currently accepted principles of epistemic closure will exploit only cognitive, indeed categorial, type-identification.

6 It should not be inferred, however, that concepts ensure perfect discrimination and ideas do not. The only point I am making is that as far as capacities to represent go concepts are superior to ideas. Yet, in their exercise, both concepts and ideas can misrepresent. For example, one may have the concept of a barn and still be unable to distinguish from a certain distance a real barn from an imitation of one. The same is true if one has only the idea of a barn that in the latter case one can always mistake a house for a barn in the best visual conditions, if one's privileged representation of a barn happens to be very much like that of the house in question. This latter mistake is one which is unlikely when the concept of barn is available. (it was Fred Dretske who pressed me to clarify this point).

7 I think that this is also what Fred Dretske has persistently and forcefully argued, from a different perspective, in a number of works published in the last 15 years.

8 Versions of this paper have been read at Tulane, Princeton, and the Western APA meeting in Chicago, April 1983. I want to thank my friends and colleagues in all these places as well as two anonymous APA referees and an anonymous APQ referee for their criticisms and suggestions. Special thanks go to Fred Dretske and Gilbert Harman.