

Recall the Gettier problem. One response to the apparent impossibility of providing a satisfactory analysis of the concept of knowledge was to argue that the whole project is wrongheaded. We should never have been interested in our concept of knowledge in the first place. The thing we should be interested in is knowledge – the natural phenomenon. And we should study it in the same way in which we study natural phenomena.

Another response is to insist that our concept of knowledge is something that is worth studying. But we should reconsider our methodology in doing so. The problem with Gettierology was not with *what* they were studying but with *how* they were studying it.

Craig argues that the success of an inquiry into our concept of knowledge cannot just be measured by whether it provides a set of necessary and sufficient conditions. It needs to tell us why we have the concept in the first place. What's the point of it? Why is it useful? (Note that the more complicated and convoluted the analysis is, the more difficult it will be for that account to answer these questions.)

The thought that we want to know not just what our concept of knowledge is but also what it is for motivates an alternative approach.

Reverse-Engineering: We begin with a hypothesis about what the concept of knowledge does for us. Then we ask what a concept having that role would be like. What would the conditions governing its application be?

There is a shift of focus here, compared to the traditional analysis. The traditional analysis focussed on the concept. The project of Reverse-Engineering, on the other hand, focusses on our *practice* of attributing the concept.

Note that there's still a role for our everyday intuitions. Note that the project is not to construct a new concept but rather to explicate our everyday one. So, the proposed account still needs to match our intuitions. But it needn't match those perfectly. The goal is now no longer on providing necessary and sufficient conditions.

Reverse-Engineering versus Traditional Analysis

- We can take into account a much wider range of explananda and data than the traditional project of analysis allows: for example, not just intuitions about Gettier cases but also our responses to skeptical scenarios.
- We don't have to assume anything controversial about the underlying philosophy of language or epistemology.
- We don't have to aim at a set of necessary and sufficient conditions as the final goal of our investigation.
- We are not looking for the "smallest common denominator".

Reverse-Engineering versus Naturalized Epistemology

- Why think that something useful would emerge from such an approach? All languages have a word for water. The purpose is simply to refer to the stuff. Isn't it the same with knowledge? According to Craig, the case of knowledge and water are not analogous:
- While the project here is different from the one pursued by Quine et al. it's still in the spirit of naturalism. We investigate our attributions of knowledge as a natural facts, "to be understood as the (broadly speaking causal) outcome of other natural facts. Which concepts we use, what

linguistic practices are common amongst us, these are special cases of input to the more general naturalistic enterprise.” (p. 9)

Knowledge is not a given phenomenon, but something that we delineate by operating with a concept which we create in answer to certain needs, or in pursuit of certain ideals. The concept of water, on the other hand, is determined by the nature of water itself and our experience of it. (p. 3)

Executing the Project

We start with a paradigm case of knowledge attribution. We assume that the concept of knowledge serves some of our basic needs. One such very basic need is to have true beliefs about the environment. And since we cannot always acquire all the true beliefs that we need ourselves, it's helpful to pool information. But for this, we need to know who we can rely on.

Craig's Hypothesis: The concept of knowledge is used to flag approved sources of information. In particular, it serves to flag good sources of testimony.

We want to know whether or not p. What do we want in an informant?

(1) Either p and he believes that p, or not-p and he believes that not-p.

This vindicates the first condition of the traditional analysis: truth.

(2) If he tells us that p, we shall thereupon come to believe that p.

This vindicates the second condition of the traditional analysis: belief.

(3) We need to be able to identify the informant as being someone whose testimony we should rely on.

This is a crucial condition. Craig:

...we need some detectable property – which means detectable to persons to whom it is not yet detectable whether p – which correlates well with being right about p; a property, in other words, such that if the informant possesses it he is (at least) very likely to have a true belief on that matter. (p. 18)

This condition vindicates something in the ballpark of Nozick's tracking account. Recall, according to Nozick knowledge requires true belief that is both sensitive and safe:

Sensitivity: If p were not true, he would not believe that p.

Safety: If p were true (under different circumstances), he would believe that p.

Why should the person looking for an informant be at all interested in possible worlds? Because he does not know which world is the actual one: there are many things that he does not know about which conditions actually obtain.

But here are the crucial difference to Nozick.

- According to Nozick, the relevant possible worlds are all the close ones. But, Craig argues, the inquirer will not be interested in all close possible worlds. He'll be only interested in the close, open possibilities: those possibilities which so far as the inquirer knows might be actual. That's a narrower range of possible worlds than Nozick takes to be relevant. This helps with counterexamples (for example, that of the bank robber whose masks slips.)
- Second, we can't tell who will satisfy the tracking conditions directly. We identify good trackers by some other property that is more directly epistemically accessible to us.