

# What's the Point of "Knowledge" Anyway?

Christoph Kelp\*

## Abstract

In *Knowledge and the State of Nature* Edward Craig defends the thesis that the function of the concept of knowledge is to flag good informants. This paper aims to show that Craig's thesis (CT) is false. In order to establish this, I will point to some data that CT fails to explain in a satisfactory manner. I will then introduce an alternative thesis that is not only able to secure the acclaimed benefits of CT, but also provides a neat explanation of the recalcitrant data.

## Introduction

In *Knowledge and the State of Nature* Edward Craig develops an elaborate and sophisticated genealogical account of the concept of knowledge. His central thesis is that the function of the concept of knowledge is to flag good informants. Recently, a number of epistemologists—notably John Greco (2007; 2008), Duncan Pritchard (2009; 2010), Miranda Fricker (2008) and Martin Kusch (2009)—have appealed to Craig's thesis (CT) in order to support or motivate a variety of further epistemological views. Of course, whether arguments that rely on CT are successful will depend on whether or not the thesis is true. The aim of this paper is to show that CT is false and that, consequently, said arguments fail, at least in their present form. In order to achieve this aim, I will show that there are some data that CT has difficulties explaining and then outline an alternative to it that not only secures the acclaimed benefits of CT, but can also explain the recalcitrant data effortlessly.

---

\*Institute of Philosophy, University of Leuven, Dekenstr. 2, #3220, B-3000 Leuven, Email: christoph.kelp@hiw.kuleuven.be

# 1 Craig's view

## 1.1 Methodology

Craig's project is initially motivated by dissatisfaction with the traditional analytical project in the theory of knowledge, which ventures to identify necessary and sufficient conditions for the (correct) application of the concept of knowledge. Apart from the fact that the traditional project has thus far led to nothing but a series of failures, Craig (1990, 2) points out that even if it were brought to a satisfactory conclusion, some significant questions in the theory of knowledge would remain unanswered. Most importantly, a successful definition of the concept of knowledge would not (or at least not obviously) also deliver an account of its value. On the face of it, however, such an account would be just as important as a correct analysis of the concept of knowledge.<sup>1</sup>

In view of these considerations Craig proposes an alternative approach to the theory of knowledge, one that puts the issue of value centre stage. Craig describes his project as follows:

Instead of beginning with ordinary usage [as the traditional project does], we begin with an ordinary situation. We take some *prima facie* plausible hypothesis about what the concept of knowledge does for us, what its role in our life might be, and then ask what a concept having that role would be like, what conditions would govern its application. (Craig 1990, 2)

It may be worth noting that the ordinary situation he envisages does not involve us, actual inhabitants of planet Earth in the 21<sup>st</sup> century. Rather, it involves a community of agents, our (imaginary) ancestors, who are like us in that they share a range of needs and (cognitive) resources with us, but are also unlike us at least in that they do not yet have the concept of knowledge. Accordingly, Craig's hypothesis concerns, in the first instance at least, the role the concept of knowledge plays not in our life, but in the lives of our ancestors. Or, to be more precise even, it concerns the role an *ancestor* of our concept of knowledge plays in the lives of our ancestors. In order to relate the story about our ancestors' concept to the concept of knowledge we are familiar with today, Craig tells a story about how, in response to further needs in the agent community, our ancestors' concept evolved into our familiar concept of knowledge.

But why think that the concept that the evolutionary process yields really is our concept of knowledge? Craig's answer is that we have reason to think that it is really our concept of knowledge provided, but only provided, that the concept delivered approximates the concept of knowledge as we are familiar with it sufficiently closely. Here is how he puts the point:

Such an investigation would still have an anchorage point in the everyday concept [of knowledge]: should it reach a result quite different from the intuitive intension, or one that yielded an extension quite different from the intuitive extension, then, barring some special and especially plausible explanation of the mismatch, the original hypothesis about the role that the concept plays in our life [that is, CT] would of course be the first casualty. (Craig 1990, 2)

## 1.2 Craig's hypothesis and concept of protoknowledge

With these remarks about the general shape of the project in play, I will now outline how Craig fills in the details. Recall that his project is set in a community of agents, our (imaginary) ancestors, who do not yet have a concept of knowledge. Recall, furthermore, that agents in this community share certain needs and (cognitive) resources with us. More specifically, they need true beliefs about the environment and can get them either through their "on-board sources" or, alternatively, from other agents in the community, i.e. from informants. (Craig 1990, 11)

What, in this situation, could induce our ancestors to introduce (an ancestor of) the concept of knowledge? What needs do our ancestors have which such a concept answers? Craig rightly points out that our ancestors need concepts to evaluate informants. In view of this he states his *prima facie* plausible hypothesis. An ancestor of our concept of knowledge—following Martin Kusch (2009) I will henceforth also call it the concept of "protoknowledge"—was introduced in response to exactly this need: its function is to flag good informants. (Craig 1990, 11)

The next step in Craig's project is to develop the application conditions of the concept of protoknowledge. In order to achieve this, Craig invites us to imagine an ancestor inquirer who does not yet have a belief about *P* but wants to acquire one from a prospective informant. The idea is that we can get at the application conditions

of the concept of protoknowledge by asking what general properties the ancestor inquirer would want his informant to have. He makes the following suggestions in this regard:

**PK-1** The informant tells one the truth on the question.

**PK-2** The informant is as likely to be right about *P* as one's concerns require.

**PK-3** The informant is detectable by one as likely (enough) to be right about *P*.

**PK-4** The channels of communication between oneself and the informant are open.

**PK-5** The informant is accessible to one here and now. (Craig 1990, e.g. 85)

It may be worth pointing out, as Craig (1990, 12-3) also does, that, typically, an informant won't tell one the truth one is after (i.e. he won't satisfy PK-1) unless he also has the corresponding true belief. Cases in which an informant hosts a false belief yet tells one the truth are bound to be exceptions. Similarly for cases in which an informant does not have the relevant belief: after all, typically, an informant who does not believe what he tells one will not speak with sufficient conviction to make one believe what he says. It thus becomes plausible that, typically, a protoknower on the question whether *P* will also have a true belief on whether *P*. In this way, PK-1 approximates the true belief condition we are familiar with from our ordinary concept of knowledge.

Notice, furthermore, that we can also find a condition that may already be recognisable as an ancestor of the equally familiar reliability requirement on the concept of knowledge, to wit, PK-2 above. That said, the reliability requirement on the concept of knowledge differs markedly from PK-2 in that it is not purpose relative in the same way as is PK-2. If, for some inquirer, the benefits of being right outweigh the costs of being wrong and the latter are next to non-existent, a barely reliable informant may be sufficiently reliable for this inquirer. If the informant will also tell this inquirer the truth on the issue and satisfies PK-3 to PK-5, he qualifies as a protoknower. However, he will not count as a knower. Someone who is barely reliable does not qualify as a knower, no matter what anyone's cost-benefit balance of being right comes to.

Finally, our familiar concept of knowledge does not feature anything like conditions PK-3 to PK-5. It becomes clear, then, that quite a bit of work remains to be done in the second part of Craig's project, which, recall, is to explain the evolution of the concept of protoknowledge into our concept of knowledge.

### 1.3 Objectivisation of Craig's concept of protoknowledge

The guiding thought in the second part of Craig's project is that the concept of protoknowledge evolved into the familiar concept of knowledge in response to further needs within the agent community. More specifically, Craig (1990, ch.10) holds that the need for an *objective* conception of good informants explains the evolution of the concept of protoknowledge into our familiar concept of knowledge. To put it in Craig's own words, the concept of protoknowledge undergoes a process of "objectivisation" and our familiar concept of knowledge is what remains after objectivisation. (Craig 1990, e.g. 91)

Let me explain this in more detail. The concept of protoknowledge is tailored, so to speak, to the needs and capacities of the individual inquirer. As a result, the concept has a number of highly subjective features: whether the informant is *as likely to be right as the inquirer's concerns require* depends on the inquirer's concerns; whether he is *detectable by the inquirer* as likely (enough) to be right depends on the inquirer's cognitive capacities; whether he is *accessible* to the inquirer depends on the inquirer's location and his available methods of communication; whether the *channels of communication are open* between the inquirer and the informant depends on what languages the two speak, whether the inquirer is the kind of person to whom the informant would divulge the information at issue and so on. It can easily be seen that whether the informant satisfies these conditions may vary depending on who fills the role of inquirer.

Now, Craig invites us to consider a somewhat more complex situation in which the agents in the community have diverse concerns, different cognitive capacities but can also collaborate in order to achieve their goals. In such a community, argues Craig (1990, 86-91), it will be extremely useful to have a conception of good informants that abstracts from the various subjective features at issue in the concept of protoknowledge. To see this, suppose that an agent in such a community, INQUIRER, is trying to find out whether *P* is true. Since he does not know anyone who would be particularly likely to have the right answer to this question, he asks his friend

POINTER, who is the best educated person he knows. POINTER himself does not know the answer but knows that INFORMANT is an expert in the field to which *P* pertains and so recommends INFORMANT to INQUIRER as an informant on the question. Here it is extremely useful for INQUIRER, first, that POINTER operates a concept of good informant and, second, that the concept of good informant is applicable to INFORMANT. What's more, it may be useful for agents in the community who are not directly involved in the episode—for instance, if INQUIRER is a leading decision maker in the agent community and the piece of information sought crucial for a decision that will affect the community's agents at large. Notice, however, that INFORMANT is not obviously a protoknower for INQUIRER and certainly not for at least some agents in the community who are not directly involved in the episode. After all, these agents may have no way of detecting INFORMANT as likely to be right on the issue, they may never be in contact with INFORMANT and INFORMANT may not divulge the information to them. For a wide range of agents in the community, INFORMANT may fail to satisfy any or all of PK-3 to PK-5 (henceforth also "epistemic accessibility conditions", for obvious reasons). Yet, it is useful (even for them) that INFORMANT be classifiable as a good informant. In this way, the pressure towards objectivisation of the concept of protoknowledge arises.

Notice that the above considerations also suggest a particular way in which objectivisation will proceed. After all, it has become clear that it is useful that INFORMANT be classifiable as a good informant despite there being a notable number of agents relative to whom INFORMANT does not satisfy the epistemic accessibility conditions. If so, there is reason to believe that objectivisation will *relax* these conditions.

But what about the remaining subjective feature of the concept of protoknowledge—that the informant be as likely to be right as the concerns of the inquirer require (PK-2)? There is reason to believe that objectivisation will affect PK-2 as well. To see why this is so, consider a situation in which an agent inquires on behalf of some other agent. To return to the above example, suppose INQUIRER has been employed by BOSS to find the answer to the question whether *P*. Situations like this one call for objectivisation of PK-2 because what matters now are not the concerns of INQUIRER but of BOSS. The need arises to have a conception of good informants that finds application independently of the concerns of the particular agent who happens

to fill the role of inquirer. Notice also that, according to Craig, objectivisation will *tighten* rather than relax the reliability required of good informants. Craig maintains that such cases

...edge us towards the idea of someone who is a good informant as to whether *P* whatever the particular circumstances of the inquirer, whatever rewards and penalties hang over him and whatever his attitude to them. That means someone with a very high degree of reliability, someone who is very likely to be right—for he must be acceptable even to a very demanding inquirer. (Craig 1990, 91).

In other words, in Craig's view, a good informant in the objective sense is someone who is reliable enough to suit anyone's concerns. Objectivisation turns PK-2 into a strong reliability condition.<sup>2</sup>

To take stock, objectivisation relaxes PK-3 to PK-5 and tightens PK-2. At the same time, claims Craig (1990, 90), PK-1, the truth-telling requirement, remains unaffected by objectivisation. Apart from what may remain of PK-3 to PK-5 after objectivisation, we thus get,

An "objectivised protoknower" whether *P*

**OPK-1** tells one the truth on whether *P* and

**OPK-2** is highly likely be right about *P*.

Of course, it remains the case that, typically an informant won't tell one the truth unless he also has the corresponding true belief. In this way, just like its ancestor PK-1, OPK-1 approximates the true belief condition on the familiar concept of knowledge. Crucially, objectivisation has also eliminated the problematic purpose relativity of PK-2 and has transformed it into a strong reliability requirement, which is also familiar from the concept of knowledge as we know it today. Finally, objectivisation dilutes the epistemic accessibility conditions on the concept of protoknowledge, thereby moving it in the direction of our familiar concept of knowledge even further. It becomes plausible, then, that the product of objectivisation approximates our familiar concept of knowledge closely, closely enough to make plausible the suggestion that our familiar concept has indeed evolved from the ancestor concept in the way Craig envisages. This completes Craig's genealogical account of the concept of knowledge.

## 2 Recalcitrant data

In this section, I will adduce a couple of cases that Craig's account of the concept of knowledge has difficulties explaining. I will start with a case of a knower who is not an "objectivised protoknower" and then move on to another case of an "objectivised protoknower" who isn't a knower. It may be worth pointing out that I do not consider these cases to be *counterexamples*, for reasons that will become clear in due course. Rather, I would like to think of them as *recalcitrant data* for Craig's account.

### 2.1 Knowers who aren't "objectivised protoknowers"

As I am about to argue, there remain cases in which knowers aren't "objectivised protoknowers". Before moving on to the actual case, however, I would first like to take another look at PK-3 to PK-5, which Craig argues are relaxed by objectivisation. Does "relaxed" here mean that they simply drop out of the picture? The answer, as Craig (1990, 89-90) also acknowledges, is "no". After all, objectivisation is said to operate on the subjective conception of good informants, i.e. the concept of protoknowledge captured in PK-1 to PK-5. Moreover, it is said to transform this concept in response to practical needs within the agent community. It is hard to see how a practical need could arise to extend the applicability of the concept of protoknowledge to agents who would not be detectable as likely to be right on the question at hand *by anyone*, who would not be accessible *by anyone at any time*, and whose channels of communication would *never* be open to anyone.<sup>3</sup> As a result, it must not be expected that objectivisation will remove PK-3 to PK-5 from the concept of protoknowledge without trace. Rather, the objectivised concept of protoknowledge will feature descendants of PK-3 to PK-5. The thought is, however, that these descendants are weak enough for it to be uncontroversial that they are satisfied by virtually anyone we would intuitively attribute knowledge to: at some time, there will be someone by whom the agent would be detectable as likely to be right, someone to whom the agent would divulge the information, and someone to whom he would be accessible. And if we should be able to construct an apparent counterexample, it will surely be a freakish one and as such does not pose a real threat to Craig's project, which after all requires only an approximation not a perfect match of the objectivised concept of protoknowledge and the concept of



knowledge.

It is this last thought that I find unconvincing. Consider cases of professional secrecy. Prominent examples here include doctors, lawyers, accountants and social workers in certain positions. Perhaps the strongest cases for my purposes, and the ones I will focus on here, are cases of priests under the seal of confession. Here is but one such case:

*Seal of Confession.* Don Camillo is the priest at the local parish. The members of his parish, who are all devout believers, regularly come to Don Camillo to confess their sins. As an ordained priest, Don Camillo is bound by the seal of confession. That is to say, he must not divulge information about his confessors' sins in any way or for any reason and cannot be forced to break this obligation even by the authorities.

Intuitively, during confession Don Camillo comes to know a wide range of facts about his confessors' sins. Yet, arguably, he is not an "objectivised protoknower" on these facts. To see this, recall that some weak descendant of PK-5 will survive objectivisation: at some time, there is someone to whom the agent would divulge the information. The problem for Craig is that Don Camillo does not satisfy even this weak descendant of PK-5. After all, he is under the obligation to keep his knowledge to himself, an obligation that he, a conscientious priest who takes his office very seriously, respects. Don Camillo is committed to taking his knowledge of his confessors' sins into his grave. His channels of communication simply would not be open to anyone at any time. That is to say, Don Camillo does not satisfy the version of PK-5 that remains after objectivisation.<sup>4</sup> At the same time, the case of Don Camillo only illustrates a general phenomenon, *viz.* that of knowledge under the seal of confession. Once one considers how common cases of knowledge under the seal of confession are, not to mention cases of professional secrecy in general, it becomes clear that cases like Seal of Confession are not at all freakish. As a result, Craig is in trouble.<sup>5</sup>

## 2.2 "Objectivised protoknowers" who aren't knowers

Recall that, according to Craig, a protoknower is someone who is sufficiently "likely to be right" on the issue. There are two ways of understanding this expression: it can mean, first, *likely to give the right*

*answer* or, second, *likely to have a true belief* on the issue. It is highly plausible that the first interpretation correctly captures the relevant conceptual truth about good informants. Accordingly, this is how the expression is to be understood in PK-2. Now recall that objectivisation is to operate on the subjective conception of good informants. Thus, when objectivisation tightens the reliability requirement at issue in PK-2, that means, in the first instance, that an “objectivised protoknower” is someone who is highly likely to answer correctly. Now, as Craig points out, typically, people won’t give right answers unless they also have the corresponding true beliefs. Similarly, it is plausible that, typically, people won’t give reliable answers unless they also have the corresponding reliable beliefs. The fact remains, however, that Craig can secure this result *only* for typical cases. The possibility remains that, in atypical cases, someone is an “objectivised protoknower” in virtue (partly) of being highly likely to give the right answer, whilst also being a hopelessly unreliable believer. In what follows, I will adduce a case in which this possibility materialises and explain in what way the existence of such a case constitutes a problem for Craig.

Let me begin with the case:

*Secret Sect.* Dick is a member of a secret sect and for that reason shares the sect’s belief that our planet is gradually warming. However, this belief is held not on the basis of scientific findings but is instead grounded in the sect’s belief (also shared by Dick) that global warming is the result of God’s decision to punish humanity for the fornicatory practices that, in recent times, have become so outrageously widespread among his once beloved sheep. Since the sect is secret, Dick is not allowed to assert its beliefs. For that reason he has adopted a policy of asserting on the relevant issues in accordance with what the experts in the field have to say. Fortunately, Dick is a government spokesman on environmental issues and thus particularly well acquainted with expert views on global warming.

Dick is an “objectivised protoknower” on the issue of global warming. For starters, Dick will tell one the truth about global warming thus satisfying the objectivised version of PK-1. Moreover, given his office and policy, he is also highly likely to give interlocutors the right answer here and so satisfies the objectivised version of PK-2. Finally, given his office, he is accessible to and detectable as likely to be right

by a wide range of people and his channels of communication on the issue will also be open. In other words, he satisfies the objectivised versions of PK-3 to PK-5.

Now, why think that case Secret Sect constitutes a problem for Craig? To begin with, notice that Dick's belief that global warming is happening is highly irrational as it is held for reasons that only the raving mad would conceivably take to support it. In consequence, it does not qualify as knowledge. We thus have a case in which an "objectivised protoknower" isn't a knower, a mismatch in the intuitive extension between the two concepts. By Craig's own lights, one might think, this is cause for concern.

On second thought, the case might not carry a enough weight to seal the case against Craig. After all, recall that Craig's project requires for its success only an approximation not a perfect match of the objectivised concept of protoknowledge and the concept of knowledge. Moreover, we have already seen that cases in which unreliable believers are reliable asserters are bound to be atypical. As a result, one might think that the remaining mismatch does not undermine the success of Craig's project.

I am happy to take this point for the time being and to accept that this case does not seal the case against Craig. In fact, I am even willing to grant that Seal of Confession, which cannot obviously be dealt with in the same way, does not refute CT either individually or in conjunction with Secret Sect. The fact remains, however, that the cases are not easily explicable by Craig's account, they constitute recalcitrant data. In the following section I will introduce an alternative hypothesis that, when plugged into the Craigian framework, delivers a concept approximating the concept of knowledge at least as closely as the one delivered by Craig's hypothesis. What's more, the data that are recalcitrant for Craig pose no difficulty for the concept delivered by the alternative hypothesis. Craig's hypothesis will thus be disconfirmed not by the existence of recalcitrant data but by the existence of such data in conjunction with an alternative that explains them neatly.

### **3 The case against Craig's view**

#### **3.1 Methodology and scope of the argument**

The aim of this section is to implement the Craigian project with a different hypothesis in place of CT and to show that the alternative

hypothesis yields a better result than Craig's. For the purposes of my argument, I will help myself to Craig's methodology. While I think that Craig's methodology (or at least something in the ballpark of this methodology) may turn out to be defensible<sup>6</sup>, I will not be able to concern myself with its defence here. Instead, I will adopt Craig's methodology without further argument. With this point in mind, it ought to be clear what my paper can and cannot hope to achieve. It can and explicitly aims to achieve a refutation of CT. After all, Craig himself relies on his methodological assumptions. In consequence, it is unobjectionable for me to adopt them for the purposes of arguing against him. Moreover, it may succeed in establishing the conditional claim that if Craig's methodology is defensible, then there is at least some reason to think that the alternative hypothesis is defensible also. However, it cannot achieve an unconditional defence of the alternative hypothesis. Accordingly, while I think that the alternative hypothesis is correct (be that because Craig's methodology is in essence correct or for other reasons), I would like to emphasise that the central argumentative aim of this paper does not consist in its unconditional defence.

### 3.2 An alternative hypothesis and concept of protoknowledge

Let's return to the question what sorts of need could have induced our ancestors to introduce an ancestor of the concept of knowledge. Recall that our ancestors are like us in that they need true beliefs about their environment and can get them through their "on-board sources" or from informants. Craig observed that our ancestors need concepts to evaluate informants. He hypothesised that the ancestor of our concept of knowledge was introduced in response to this need and that its function is to flag good informants.

While it is correct that our ancestors need concepts to evaluate informants, this is certainly not their only conceptual need. At the very least, they also need concepts to evaluate various inquiries agents undertake. The alternative hypothesis that I will explore here is that the concept of protoknowledge was introduced in response to this need: its function is to flag when agents *may adequately terminate inquiry* into a given question.

To continue the parallel with Craig, let's reflect on what conditions would govern the application of a concept with this role. In order to do this imagine that an ancestor is interested in the question whether  $P$  and sets out to inquire. What properties would our ances-

tor want himself to have upon terminating inquiry? My suggestions here are as follows:

**PK-A** He has formed a belief on whether  $P$ .

**PK-B** His belief on whether  $P$  is true.

**PK-C** His belief on whether  $P$  stems from a source that is as trustworthy on the question whether  $P$  as his concerns require.

The alternative hypothesis delivers a concept of protoknowledge approximating our familiar concept of knowledge at least as closely as the one delivered by CT. After all, PK-A corresponds to the belief condition, PK-B to the truth condition on knowledge and PK-C states a condition that may, again, be recognisable as an ancestor of the familiar reliability requirement on the concept of knowledge. However, just as in case of Craig's PK-2, the reliability requirement on the concept of knowledge differs markedly from PK-C in that it is not purpose relative in the same way as is PK-C: given a suitable cost-benefit balance of being right, a true belief acquired from a barely trustworthy source can qualify as protoknowledge even though it could not qualify as knowledge. Notice, finally, that the present concept of protoknowledge features no analogues to Craig's PK-3 to PK-5. This is as it should be: Why should it matter to whether one has adequately terminated inquiry that one is detectable as likely to be right on the issue, that one is willing to share one's results and that one is accessible to others?

### **3.3 Objectivisation of the alternative concept of protoknowledge**

Recall that the guiding thought of Craig's evolutionary story was that in sufficiently complex agent communities the need for an objective conception of good informants arises. Again, I will adopt Craig's guiding thought and tailor it to my purposes. In my version of the evolutionary story, it is the need for an objective conception of adequately terminated inquiry that arises in sufficiently complex agent communities. Thus consider a community of agents who can have different concerns at different times. Suppose, furthermore, that these agents are able to store the results of their inquiries in memory but may also forget facts about the sources of stored contents. In such an agent community the pressure towards an objective conception of adequately terminated inquiry arises, one that abstracts away from

the agents' concerns at the time of inquiry. After all, what matters now are not just the agents' present concerns but also the concerns of their future selves, which may be very different and can also be opaque to them at the time of inquiry.

Further pressure towards objectivisation arises from various social facts in sufficiently complex agent communities. Consider, for instance, communities with agents who have differing concerns and capacities, who may form groups which may again have differing concerns from individual agents, including (at least some of) its members. Suppose, furthermore, that agents in this community traffic in information in the way we do: agents can inquire on behalf of other agents or as members of groups and can enter the results of their inquiries into various databases from which they may subsequently be retrieved by themselves or other agents who have access to the databases. In such a community the pressure towards objectivisation of the concept of adequately terminated inquiry only increases. After all, what matters now are not only the concerns of the individual agent at the specific time of inquiry, but also the concerns of other agents and groups of agents, present and future, which may be very different than the ones of the inquiring agent at the time of inquiry. In such a community of agents, the pressure to filter out the purpose relativity at issue in PK-C increases.

Regarding the trajectory of objectivisation, I am, of course, free to adopt Craig's story for my purposes. If I do, in my story, objectivisation will tighten the requirement of source trustworthiness at issue in the concept of protoknowledge in such a way that only sources trustworthy enough to suit anyone's concerns will qualify as sufficiently trustworthy. In this case, just like Craig, we end up with a very strong reliability condition on adequately terminated inquiry objectively construed.<sup>7</sup>

To the extent that it is plausible that PK-1 survives objectivisation in Craig's story, it also plausible that PK-A and PK-B survive it in my story. We thus get:

An "objectivised protoknower" whether *P*

**OPK-A** has formed a belief on whether *P*,

**OPK-B** his belief is true, and

**OPK-C** his belief stems from a highly reliable source.

The result of objectivisation of the concept of protoknowledge as captured in PK-A to PK-C is thus fairly similar to the result Craig obtained. There are, however, a couple of crucial differences. Recall, first, the concept of protoknowledge in my story does not feature analogues to PK-3 to PK-5. As a result, the problem that arose for Craig's story from cases like Seal of Confession can be avoided. Openness of channels of communication does not even figure in the application conditions of a concept with the function of flagging when agents may adequately terminate inquire. Hence, the fact that there are agents who would not at any time publicise certain bits of information is immaterial to the question whether they have adequately terminated inquiry.

The other crucial difference between the results of objectivisation is that, in my story, objectivisation outputs a strong reliability requirement on agents' *belief sources*, while in Craig's story it outputs a strong reliability requirement on agents as *truth tellers*. It may be that, typically, agents won't be reliable truth tellers unless they are also reliable believers so that Craig's account will pass the same verdict as the alternative account in all typical cases. Yet, as Secret Sect nicely illustrates, it turns out that there are atypical cases in which agents are reliable truth tellers despite being unreliable believers. Such cases are exactly the kind in which the two accounts pass incompatible verdicts and are therefore ideally suited as test cases for the two positions. The fact that, intuitively, agents in such cases do not know thus disconfirms Craig's hypothesis vis-à-vis mine.

## Conclusion

I have argued that there is an alternative to CT that, when slotted into Craig's framework, delivers a better result than CT. The alternative hypothesis delivers a concept that not only approximates our concept of knowledge as closely as does CT, but it also explains a set of data that Craig's account struggles with. For that reason the prospects for CT are very dim indeed. For those who have ventured to appeal to CT in order to motivate or support further epistemological views, this has significant implications. At present their arguments are bereft of a solid foundation. The onus is thus on them either to show that their arguments can be run with the alternative hypothesis in place of CT or else to seek different ways of providing the support needed.

## References

- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. In Tomberlin, J., editor, *Philosophical Perspectives*, volume 13. Blackwell, Oxford.
- Craig, E. (1990). *Knowledge and the State of Nature. An Essay in Conceptual Synthesis*. Clarendon Press, Oxford.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52:912–29.
- Fricker, M. (2008). Scepticism and the genealogy of knowledge. *Philosophical Papers*, 37:27–50.
- Greco, J. (2007). The purpose of knowledge and the nature of ability. *Philosophical Issues*, 17:57–69.
- Greco, J. (2008). What's wrong with contextualism? *The Philosophical Quarterly*, 58:416–435.
- Kusch, M. (2009). Testimony and the value of knowledge. In Haddock, A., M. A. and Pritchard, D., editors, *Epistemic Value*. Oxford University Press, Oxford.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press, Cambridge.
- Millikan, R. (1984). *Language, Thought, and Other Biological Categories*. MIT Press, Cambridge/MA.
- Pritchard, D. (2009). *Knowledge*. Palgrave Macmillan, Basingstoke.
- Pritchard, D. (2010). Anti-luck virtue epistemology. Manuscript.
- Riggs, W. (2008). The value turn in epistemology. In Hendricks, V. and Pritchard, D., editors, *New Waves in Epistemology*. Palgrave Macmillan, Basingstoke.



## Notes

<sup>1</sup> The idea that value questions are just as important as definitional questions in epistemology has become more and more widely acknowledged (see e.g. Kvanvig 2003). Some have even talked about a value turn in epistemology (see Riggs 2008).

<sup>2</sup> Recent debate over the semantics of knowledge attributions suggests that people's concerns can sometimes be extraordinarily pressing (as DeRose's (1992) bank case and Cohen's (1999) airport case illustrate nicely). Accordingly, Craig's proposal will arguably lead him to a so-called "high standards" version of classical invariantism. Notice, however, that this result can be avoided. For instance, if one were to hold that it is enough for an "objectivised protoknower" to be sufficiently reliable to satisfy *ordinary* concerns, one will end up with a moderate version of classical invariantism. Thanks to Erik Olsson for pointing this out to me.

<sup>3</sup> The intended sense of "would not" here and in the relevant uses below is: "It would not be that *P*" is true if and only if in no case that has a non-negligible chance of obtaining, *P*. I take it that the intended sense of "would not" is sufficiently strong to make the claim very plausible indeed. After all, it appears to be a matter of conceptual fact that there could be no *practical* need to extend the applicability of some concept to cases that, by definition, have no more than a *negligible* chance of obtaining.

<sup>4</sup> Notice, furthermore, that it seems vital to the survival of the practice of confession that priests would not divulge the information obtained during confession. After all, otherwise, there would be a strong incentive not to go to confession or at any rate not to confess all one's sins without reservation.

<sup>5</sup> Craig (1990, 90) mentions the possibility of supplementing his evolutionary story by "additional explanatory principles" which may allow him to get the cases right. This, I take it, should be no surprise, as we can nearly always explain recalcitrant data by supplementing our theory with additional explanatory principles. Crucially, as I will argue below, the data can be explained without having to appeal to such principles by a theory that is like Craig's in all aspects except that the initial hypothesis is different. As a result, explaining the data by wheeling in additional explanatory principles doesn't help Craig here. The alternative theory is still preferable as it is simpler.

<sup>6</sup> Suppose, for instance, that a theory of the content of mental representations that emphasises the importance of the historical function of the representation (see e.g. Millikan 1984) is correct. Given the plausible assumptions (1) that concepts are mental representations and (2) that mental representations (and hence concepts) are individuated by their contents, a methodology like Craig's becomes more and more appealing.

<sup>7</sup> Again, this result is by no means inevitable. For instance, if objectivisation tightens the requirement of source trustworthiness to comply with ordinary concerns, we will get a less strong reliability condition on the objectivised conception of adequately terminated inquiry of the kind preferred by moderate classical invariantists.