UNKNOWABLE TRUTHS

Paul HUMPHREYS

ABSTRACT: This paper addresses a solution due to Michael Fara to the Church/Fitch paradox of knowability. Fara's solution has significant interest but the paradox can be resurrected within his approach by considering a slightly more complex sentence. The issue of what counts as an epistemological capability for enhanced agents is then discussed with some emphasis on the developmental heritage of agents and their ability to transcend conceptual frameworks.

KEYWORDS: paradox of knowability, epistemic limits, Michael Fara.

1. Introduction

The question of where the limits of knowledge lie has usually been discussed with the tacit understanding that the knowledge concerned is knowledge accessible to humans. That is natural, but for philosophers of science, and especially those who hold that science has access to knowledge that lies beyond the epistemic grasp of humans, a broader perspective on the limits of knowledge is required. By `lying beyond the epistemic grasp of humans,' I do not mean the kind of knowledge that is usually the focus of limit science, the knowledge that will be available once the scientific method has completed its task. For that limit is construed as the limit of science constrained by observations, where observations are observations made by, or accessible to, humans.[1] What I mean here by `lying beyond the epistemic grasp of humans' is for the potential knowledge to be about some aspect of reality that is inaccessible to human perceptual and cognitive capacities, but may be accessible by enhanced agents or agents with different epistemic capacities that practice an identifiable kind of scientific method.

---

[1] The infamous `model-theoretic argument.' Hilary Putnam, *Reason, Truth and History* (Cambridge: Cambridge University Press, 1981) is an exception to this because the `observational' sentences used there are simply a set of constraints that are formulated in some special sub-vocabulary of the language.

2. The Church/Fitch Argument

What has come to be known as the Church/Fitch knowability result has a direct bearing on this discussion of the limits of scientific knowledge, construed broadly. Here is the special form of the argument, special in the sense that the argument is specifically about knowledge in the standard sense, restricted to human knowledge. There are three premises in the argument. Let the sentential operator K represent `some epistemic agent at some time knows that' and ◊ the standard modal operator representing metaphysical possibility, that is, truth at some metaphysically possible world. Then the first premise is:

$$1K.\ \Pi p(\ p \rightarrow \Diamond Kp)$$

which asserts that for any truth p, it is possible for at least one epistemic agent to come to know that p. Premise 1K is often taken to state a verificationist principle, one that denies the realist position that there are parts of reality, representable in some conceptually accessible language, that it is impossible for any agent ever to know. Note that 1K is not stated as a necessary truth. In the original argument, this could be strengthened because verificationist positions generally link truth conditions and verification conditions so tightly that the concept of an unknowable truth is incoherent.[2] In contrast, the realist position is best construed as contingent. It is certainly metaphysically possible for there to be a world in which every contingent truth is not only knowable, but known; it simply happens that the evidence we have strongly suggests that our world is not like that. A solipsistic and idealistic world inhabited by an inquisitive simpleton is an example of the former, because the only contingent truths that need to be known consist in the few thoughts that pass through the sparsely populated mind of the agent. Once mathematical truths are allowed, the existence of a gap between what is known and what is knowable becomes more plausible.

The second premise is:

$$2K.\ \forall p\ \Box(Kp \rightarrow p)$$

I shall take this as the unproblematical assumption that knowledge requires the truth of the sentence that is known.

The third premise is also unproblematical, simply asserting that knowledge of a conjunction entails knowledge of the conjuncts:

$$3K.\ \forall p\Box(K(p\&q) \rightarrow (K(p)\ \&\ K(q)))$$

The Church/Fitch argument is then a straightforward consequence of 1K,

[2] Traditional forms of intuitionism in mathematics take this form.

2K and 3K. Take as an instance of 1K the sentence p & ¬Kp. This asserts that there is a true sentence that is never (actually) known by any agent. It is taken as an empirical fact that there are unknown truths in that sense.[3] Typical examples are historical facts, such as the exact number of bacteria in the body of Joan of Arc when she was burnt at the stake in 1431. Assuming clear countability criteria for bacteria and the impossibility of resurrecting the pre-cremated Joan of Arc, this is a good candidate for a truth that no agent has known or ever will know. Although an extremely strict verificationist might reject this sentence as a candidate for knowledge, it is metaphysically possible that the truth of that sentence is known and so satisfies 1K. If the appeal to unobservables such as bacteria troubles you, take instead the number of hairs on Joan of Arc's head at the appropriate time. Using this instance in 1K we have (p&¬Kp) → ◊K(p&¬Kp) and given the truth of p&¬Kp we have ◊K(p&¬Kp). Using 3K, we have that ◻(Kp & K ¬Kp). Since K¬Kp entails ¬Kp, using 2K, we conclude with ◻(Kp & ¬Kp), which is inconsistent.

Once again assuming that 2K and 3K are uncontroversial (and also that the inferences made are unproblematical) we have three choices for what to do.[4] The first is to retain the verificationist principle 1K, and to maintain that sentences of the form p & ¬Kp do not exist, despite the considerations given above. The verificationist would then be committed to the position that:

$$\mathrm{CV} \quad \forall p(p \rightarrow Kp)$$

that is, that all true sentences are not just knowable but are known by some agent at some time. That is not an inconsistent position, and it perhaps captures what a strict verificationist is committed to, but at the least it drives the verificationist to deny the existence of truths that correspond to states of affairs that occurred before the evolution of cognitively capable agents, as well as many other kinds of contingently inaccessible facts. This is a sufficiently unattractive position that I shall not pursue it.

The second option is to reject assumption 1K on the basis that the verificationist position is false. A third option is to claim that 1K does not represent what verificationism intends and that an alternative premise should be used. Before discussing this third option, let us generalize the argument. Let O be a sentential operator satisfying the following three conditions:

$$1K^*. \quad \forall p(p \rightarrow \Diamond Op) \text{ (Possibility Principle)}$$

---

[3] For arguments to this effect see Nicholas Rescher, *Unknowability* (Lanham, MD: Lexington Books 2009). Further reasons are given in section 2 below.

[4] An alternative approach is to reject one of the inference rules used in the proof on the grounds that one or more of them is unacceptable to a verificationist.

2K*.$\forall p\ \Box(Op \rightarrow p)$ (Veracity Principle)
3K*. $\forall p\Box(O(p\&q) \rightarrow (O(p)\ \&\ O(q)))$ (Conjunction Principle)

From these three premises, if some claim of the form $p\ \&\ \neg Op$ is true, then it is provable that

CK*. $\forall p(p \rightarrow Op)$

Here is a proof that can be found in Rescher, *Unknowability*.

1. $\Box(O(p\&\neg Op) \rightarrow (Op\&O\neg Op))$	by substitution in premise 3
2. $\Box(O\neg Op \rightarrow \neg Op)$	by substitution in premise 2
3. $\Box\neg O(p\&\neg Op)$	From 1, sentential equivalent to 2, modus tollens
4. $\neg\Diamond O(p\&\neg Op)$	From 3 by definition of modal operators
5. $(p\&\neg Op) \rightarrow \Diamond O(p\&\neg Op)$	By substitution in premise 1
6. $\neg(p\&\neg Op)$	From 4,5
7. $p \rightarrow Op$	From 6.
8. $\forall p(p \rightarrow Op)$	By generalization

3. Fara's Argument.

The first part of my discussion will be based on a recent analysis by Michael Fara of the argument.[5] Fara demonstrates that the general form of the Church/Fitch argument can be applied to a number of non-epistemic situations, such as the actions of an omnipotent agent. I shall concentrate here on an application (not discussed by Fara) to scientific agents that have access to knowledge lying outside the domain actually accessible by humans, but constrained by the norms of science so that they are not omniscient. Three features of 1K then require clarification.[6] What counts as the scope of quantification over p is important. I shall take p as representing an arbitrary sentence in some interpreted language, rather than a proposition, so that the concepts occurring in p must be specifiable. p describes some actual and not merely possible fact (state of affairs); the facts can be about individuals or can be generalizations, including laws. The modal operator $\Diamond$ in 1K is usually taken to represent metaphysical possibility. Because the knowledge operator K occurs within the scope of the modal operator, the agents, past, present, or future, must be located in the appropriate possible, usually non-

[5] Michael Fara, "Knowability and the capacity to know," *Synthese* 173 (2010): 53-74

[6] I shall assume here that quantifying over the set of all sentences is unproblematical, setting aside here concerns raised by Patrick Grim,"There is No Set of All Truths," *Analysis* 44 (1984): 206-208.

actual, world. 1K is then a curious mixture of epistemic and metaphysical content. The premise says that for every actual truth, in some metaphysically possible world an agent knows that truth. Fara explicitly represents p as being preceded by an actuality operator A, so that the premise 1K becomes

$$\forall p\,(Ap \rightarrow \Diamond KAp) \qquad (A)$$

According to Dorothy Edgington what is needed to capture the anti-realist position is to interpret (A) as[7]:

> (p is true at the actual world w) → ∃w'(It is known by some agent at some time at w' that p is true at w) (B)

As Fara then points out, when p is not known in the actual world, the consequent of (B) claims that an agent located outside the actual world knows at least one factual truth about the actual world and "this kind of counterfactual knowledge of actuality seems starkly impossible."[8] We can have knowledge of non-actual situations, because it is commonly held, for example, that we can stipulate for a given world that it have the same laws as our world. But because we are concerned with scientific claims having empirical content, a priori modes of investigation or stipulative claims about what the actual world contains are inappropriate and this extra-worldly epistemic access is impossible. This means that 1K does not correctly represent what a verificationist needs.

Modifying Fara's own solution a little, his interpretation of premise 1K is this: actually p → actually ∃x(x has the capability to know that actually p) (C) where `having the capability to know that' means (a) that the agent possesses the appropriate epistemic equipment, whether exercised or not, that would allow the agent in appropriate circumstances to know that p and b) `to know that' is an operator that satisfies both 2K and 3K.[9]

We then have by substituting

$$`p\&\forall y\neg(y \text{ knows that } p)' \qquad (D)$$

for p in (C):

> actually(p&∀y¬(y knows that p)) → actually ∃x(x has the capability to know that actually [p and that no cognitive agent knows that p])

and, as Fara notes, this does not does not lead to a contradiction because the agent

---

[7] Dorothy Edgington, "The paradox of knowability," *Mind* 93 (1985): 557-568.

[8] Fara "Knowability," 64.

[9] Fara uses the term `capacity' rather then `capability' but the former has nuances in the philosophy of science that are best avoided here.

x has the capability to know that (actually) p and also has the capability to know that (actually) no cognitive agent knows that p. In the Joan of Arc example, we can see that, on a reasonable interpretation of `has the capability to know that,' this is true for the sentence about the number of hairs on Joan of Arc's head. A key aspect of this response to the Church/Fitch argument is that we do not run afoul of the analog of 2K. It does not follow from the fact that an agent has the capability of knowing that p that p is in fact true.

4 Another Problem

Fara's solution to the paradox is appealing, not the least because it has a robustly epistemological focus. But the paradox can be revived in what is perhaps a more worrying form. We can begin by asking whether (C) correctly captures what a verificationist needs. According to Joe Salerno the kind of anti-realism involved in the argument asserts that:

> ...barring vagueness and ambiguity, we could in principle know the truth value of any fully understood proposition, given only finite improvements to our epistemic capacities, resources, or environment.[10]

Key placeholders in this claim are `We,' `in principle,' `fully understood,' and `finite improvements.' As with most epistemological discussions, the agents are taken without question to be we humans. Concordantly, the `in principle' refers to idealizations of our existing cognitive capacities, as does the `finite improvements.' Let us see what happens if we relax this restriction. Because Fara's (C) has no modal force beyond the capability feature, the truth of (C) depends on the contingent development of cognitive agents and science in our world. So we have to make some decisions about where to draw the line between what lies within the domain of applicability of an acceptable capability and what does not. Traditional empiricist-based verificationism about factual truths ties verification procedures tightly to the perceptual capabilities that evolution has provided to humans. Even if not universally agreed upon, it is plausible that evolution could have taken a different course and endowed humans with different kinds of perceptual abilities. In addition, the advent of technological supplements to our perceptual and intellectual capabilities are developments that should be welcomed by a scientifically sympathetic verificationist. On the other side, epistemic agents endowed with supernatural powers would be unacceptable to most verificationists. So I shall take as an acceptable cognitive agent any agent that could have developed during the lifetime of our universe, including agents

[10] Joe Salerno, "Introduction to knowability and beyond," *Synthese* 173 (2010):1.

possessing artificial intelligence or hybrids that combine human cognition with the different capabilities of instruments and computational science, the existence of which is compatible with the actual laws of nature, and including those that for contingent reasons did not come into existence.

We now need to ask whether it is reasonable to argue, under the more liberal account of epistemic agents, that (C) is true. The answer, I think, is that the weakening of 1K from a metaphysical possibility to the actual existence of an agent capable of knowing that p, as (C) asserts, renders (C) false. For if we need sophisticated technology or superhuman computational abilities to know that p, as we do in many scientific cases, then there is evidence that no *actual* cognitive agent, whether one that is human and supplemented by artificial devices, or one that is entirely artificial, is capable of knowing that p for a sentence p the content of which involves an extreme degree of unobservability or uncomputability. Here, the capability can even be taken as capability in principle as long as the idealizations involved do not take us out of the domain of what is actually possible given the laws of nature and the finite cognitive resources available.

To deny the truth of (C) is, of course, to deny what a particular kind of verificationist asserts. I am primarily interested in whether Fara's modification of the argument allows the verficationist to completely escape the problems stemming from the original argument. As I shall now show, the modification used in (C) is insufficient to completely escape from trouble. The substitution instance (D) appropriately follows the original line of argument in taking a true but unknown sentence as the target case. But consider a different substitution instance of (C):

(D*): `p&$\forall$y$\neg$(y has the capability of knowing that p)'

where the quantification is over actual agents. Does it follow from the fact that y has the capability to know that p, that p is true? One might dispute this because on some interpretations of capability, it is not true that for every sentence, necessarily if some agent has the capability to know that p, then p is true. Capability is a dispositional property, available for use on a wide variety of applications, if the conditions are right. I have the capability to know that my front door is painted purple were it to be in fact purple, but as a matter of fact it is not. However, Fara restricts himself to actual truths, and any substitution instance of C satisfying the antecedent will therefore involve actual truths. Now consider what the relevant instance of (C) says under substitution of (D*):

actually (p & ¬∃y(y has the capability of knowing that p)) → actually ∃x(x has the capability of knowing that actually [p and ¬∃y(y has the capability of knowing that p)]).

Consider the antecedent. p is an actual truth and in the entire history of the actual world, no agent has the capability of knowing that p. Now take the sentence schema

> `there are exactly N objects of mass 10,000 kg in a closed orbit around star S' (E)

where S is the name of a star, N is a non-negative integer, and S is sufficiently distant from us that the measurement of N would require technological innovations that are never actually carried out, let us say on grounds of the expense needed to make the required instruments. For some N, possibly zero, the sentence schema is true.[11] Call that instance p. Because the necessary instruments were never constructed, no agent possessed the capability of knowing that p. Remember that some cognitive agent in the history of the universe must actually have that capability, not just be in a position to possibly have it. So we can detach the consequent and distribute the capability over the conjunction to obtain:

(This should be in regular font, not displayed)

> actually ∃x(x has the capability of knowing that actually p and x has the capability of knowing that actually ¬∃y(y has the capability of knowing that p). Simplifying the conjunction, we have that ∃x(x has the capability of knowing that actually p). But we also have from (D*) that ¬∃y(y has the capability of knowing that p). And this is a contradiction since p is an actual truth.

One aspect of this argument needs attention.[12] There is a sense of `capability' for which it is true that if an agent has, for example, the technological capability to construct an instrument that would provide evidence for a true sentence p, and the existence of that instrument would endow the agent with the epistemic capability to know that p, then the agent has the capability to know that p. There is another sense of `capability' in which collapsing the successive capability claims is illegitimate. There are two things to mention about this dual sense of `capability.' First, Fara makes a distinction between capacities and abilities in this way:

---

[11] In assuming that some instance of the sentence schema is true, I am assuming that classical logic holds, at least in the form that the truth values of all instances of the schema exist. This assumption is not stronger, I believe, than the factual assumption that there is a true but unknown sentence, an assumption that is used in the original argument.

[12] The need to explicitly address this point was suggested by Majid Razvi and Ken Akiba.

> ... one has the capacity to perform some feat provided one's internal constitution does not rule out the performance of that feat; to have the ability to perform that feat one must, in addition, be disposed to succeed in performing that feat if one tries.[13]

I believe that Fara's account of capacities (called here `capabilities') does not permit us to say that an agent lacking the physical equipment to carry out a cognitive task has the capacity to do so. Just as a blind person does not have the capacity to perceive visual shapes, so an agent lacking the physical equipment to determine the number of objects orbiting a given star, however contingent that lack is, does not have the capability to know that number. Secondly, since both senses of `capability' are legitimate, there is one interpretation of (C) under which the inconsistency exists for that stricter kind of verificationist. Moreover, the more liberal version of verificationism, the one that allows the compression of iterated capabilities, will need to provide some kind of criterion about where the boundary between in principle capabilities and in practice capabilities lies.[14] Too far towards the in principle wing and one loses much of the original motivation for being a verificationist. Too far towards the in practice wing and the risk of permitting inconsistencies returns.

To summarize: Fara's modified representation of verificationism, as contained in (C), itself leads to a contradiction taking plausible auxiliary assumptions as true. So if (D*) is true, (C) must be false and there are actual truths that lie beyond the epistemic capabilities of any actual agents. The truth of (D*) depends upon contingent facts about what extensions of scientific capabilities are actually made and what kinds of capabilities are needed to access all factual truths. The first of those facts will influence the strength of what verification position is reasonable, whereas the second will determine the epistemic gasp between a realist position and that verificationist position. Once one reaches the extreme of an agent that is omniscient about factual truths (rather than possible truths), not only will (D*) be false but the modal collapse of knowability into knowledge will not be objectionable in that situation.[15]

---

[13] Fara, "Knowability," 68.

[14] For a discussion of this distinction with respect to traditional empiricism and contemporary computational science, see Paul Humphreys, *Extending Ourselves* (New York: Oxford University Press, 2004).

[15] Thanks to Gene Mills for raising the issue of where on the spectrum between current humans and omniscient knowers we should draw the appropriate epistemic boundaries.

5 Further Issues.

We can now explore some more general issues involved with unknowable truths. In our discussion above, we left aside Salerno's requirement that the propositions involved in the argument should be fully understood. It is here that some worries might enter. For (D*) to be false, it is enough for a single agent to be capable of knowing that p, but such a situation would violate standard norms for the objectivity of scientific knowledge. Even for cases of non-scientific knowledge, situations in which the justification for p could not be conveyed to others would be grounds for suspicion but we can address that problem at the price of a small amount of imprecision by requiring that a necessary and sufficient condition for (D*) to be false is that a small class of agents is capable of knowing p.

This requires a specification of the conceptual resources that agents have and that they can understand. For cognitive agents with extended abilities, this specification is not easy to give because the quantification in (D*) is over all agents. Conceptually weaker agents do not have the ability to understand what it is that cognitively more powerful agents can know and a given agent cannot understand, let alone fully understand, a proposition formulated using a set of concepts he or she does not have and cannot comprehend. It is often claimed that scientists of bygone eras could not have known true propositions about contemporary theories such as quantum field theory because they lacked the appropriate concepts. This is no doubt correct for historically situated agents, but the more general issue is whether there are concepts that no agent could understand. Because we are dealing with sentences and not propositions, the linguistic relativity of conceptual resources does seem to preclude languages that are incommensurable with all known languages from being understood even by the cognitively most powerful agents. One kind of reason for maintaining a realist position would thus be to maintain that there are sentences that it is impossible for any cognitive agent to ever know because the concepts involved lie beyond understanding. I am not sure that it is possible to provide any clear content to that claim, if only because it requires the idea of an interpreted language that no agent could ever understand.

When discussing science as it might be conducted by aliens, Rescher says that the science of different civilizations is

> inevitably closely tied to the particular pattern of their interaction with nature as funneled through the particular course of their evolutionary adjustment to their specific environment.[16]

and that

[16] Rescher, *Unknowability*, 26.

> ... in science, as in other areas of human endeavor, we are prisoners of the thought-world that our biological and social and intellectual heritage affords us.[17]

Although various degrees of constraint are mentioned, this conceptual framework is said to "determine what is seen as an appropriate question and what is judged as an admissible solution."[18] Let us begin with the social and intellectual heritage.

At least within our own species, one of the most striking features of science has been its ability to transcend in large part, if not completely, the limitations of the social and intellectual environment, at least in the natural sciences. Saying "It's a science thing, you wouldn't understand" is an exhortation to get into the lecture room and learn something, not a claim that because you are from Nepal or Rwanda, the concepts of quantum field theory are blocked by your culture. The extensive debate about the existence of cultural universals such as kin groups, sexual attraction, and concepts of fairness has revealed the existence of at least several hundred solid candidates for transcultural concepts.[19] So claims about the effects of cultural environments must be about a deep kind of environment that perhaps rests on the biological underpinnings, in ways analogous to the ways in which surface grammars reflect cultural accidents and deep grammars reflect a common linguistic apparatus, because the differences in terrestrial environments that allow a common kind of scientific understanding to exist show that surface environments cannot preclude understanding. There is a plausible case that can be made for a certain degree of first-person perspective knowledge that can only be acquired via cultural embedding (*Verstehen*) but at best it will rule out some features of psychological and social knowledge, even if we allow such knowledge to count as scientific.

So it must be the biological heritage that constitutes the barrier to knowledge at least in the sense that the biological equivalent of embedding yourself in a different cultural environment is possible only to a limited extent. There exist virtual reality simulations that allow one to experience, at least to a certain extent, the experiences undergone by an individual suffering from paranoid schizophrenia. Current evidence suggests that schizophrenia is a heritable trait, so the simulation does allow us to enter a world that is in some way biologically different from, for most of us, our own. It is also worth noting that humans who have had their visual abilities made operative by cataract surgery

---

[17] Rescher, *Unknowability*, 32.

[18] Rescher, Unknowability,

[19] Donald Brown, *Human Universals* (Philadelphia: Temple University Press, 1991).

have been able to acquire concepts they previously lacked.[20] We should also keep in mind that `naive physics' or `intuitive physics' is not only very wrong about the physics of our world, but has incorrect conceptualizations about what counts as a natural state and other matters. To overcome this, we use forms of representation that are increasingly remote from our instinctive conceptualizations. So, fractal dimensions, although analogous to spatial dimensions, are not the kinds of things that our ordinary ways of interacting with the world give us, even though fractals are approximately realized in objects such as crystals. This suggests, although it is not presently possible to test it, that a species equipped with different kinds of sensory inputs could arrive at scientific results that we could understand. To take a simple but slightly fanciful example, dogs, which operate in a massively olfactory world that we will probably never properly understand, had they brains of the appropriate level of representational and computational power, could have developed molecular chemistry. This is not necessarily a matter of our results from science converging on results from other ways of doing science; it is simply reflects my scepticism about the degree to which accidental contextual factors can determine what we can and cannot know.

One of the things we have learned from the intensive debates about neural networks is that they possess an enormous degree of conceptual plasticity. Unlike Kantian frameworks in which the fundamental concepts are fixed, inputs leading to a redistribution of the connection weights in the network result in changes in the sub-conceptual representations used by that network. Maybe we could not understand what advanced aliens with connectionist based intelligence know, but they have the capability to understand what we know. Moreover, despite Kuhn's denial, based on a behaviorist approach to concept formation, we can understand what previous scientists did, so this indicates some kind of progress.

I want to finish with a couple of remarks about the notion of limit science. As Nick Jardine noted thirty years ago, the very notion of a limit for scientific activity is not well defined and we have no evidence either that such a limit exists or that science will not perpetually shift from one representational framework to another.[21] These are reasons for not taking seriously discussions of scientific realism and anti-realism, physicalism, and especially what can and cannot be known that are couched in terms of limit science. Although such discussions often take place in the context of the temporal development of science, the issues are

---

[20] John Z. Young, *Doubt and Certainty in Science* (New York: Oxford University Press, 1960).

[21] Nick Jardine, "The possibility of absolutism," in *Science, Belief, and Behaviour: Essays in Honour of R.B. Braithwaite*, ed. D.H. Mellor (Cambridge: Cambridge University Press, 1980), 23-42.

equally relevant to the development of reductive programs. Kuhn focused on the successive replacement of theories about particular domains over time. Yet if incommensurability arguments based on theory-relative semantic content are valid, they should apply to synchronic reduction as well, presenting serious difficulties for any reductive procedure similar to Nagelian inhomogeneous reduction. For example, `gold' has a different meaning at the nano-scale than it does at the human scale and similarly with `water' and `H2O,' as Mark Johnston has argued.[22] This synchronic incommensurability affects the program of fundamental physicalism because the evidence in favor of fundamental physicalism is supposed to come in part from progress in reducing other sciences to physics. It also renders suspect the familiar argument by Putnam that because terms such as `red' can apply to unobservable entities as well as observable entities, the division between the observable and the unobservable cannot consistently be maintained. This kind of argument is fallacious when an inference from the syntactic identity of terms to the sameness of semantic content is made.[23]

H_2O

[22] Mark Johnston, "Manifest Kinds," *Journal of Philosophy* 94 (1997): 564-583.

[23] Earlier versions of this paper were read at the *Limits of the Knowledge Society* conference, Romanian Academy, Iasi, Romania, the University of Pittsburgh workshop on Scientific Progress, and Virginia Commonwealth University. Thanks are due to each of those audiences for comments.