

Draft version as of Dec. 2018.

The final version has been published in *Science and the Production of Ignorance*, ed. by Janet Kourany and Martin Carrier, Cambridge, Mass.: MIT Press 2020, pp. 195-218.

---

Torsten Wilholt

[torsten.wilholt@philos.uni-hannover.de](mailto:torsten.wilholt@philos.uni-hannover.de)

## On Knowing What One Does Not Know: Ignorance and the Aims of Research

*To the memory of Sylvain Bromberger, 1924-2018*

### 1. Introduction

Most of the chapters in the present volume discuss concrete cases in which ignorance is willfully created, strengthened or deepened by actors seeking to further their own particular interests. My aim in this chapter is at first sight different. I seek to analyze the nature of ignorance and to distinguish between important kinds. My approach in doing so is to focus on the role of ignorance in the search for knowledge. While ignorance can have obvious harmful aspects, the part it plays in research exposes also its constructive potential. The analysis will show that from a philosophical perspective, our ignorance ought not to be regarded as one huge, structureless absence, but rather as a varied realm structured by our varying abilities to articulate and pursue questions. Such an analysis will, I hope, also help us to better understand the conditions under which ignorance plays out its positive or negative influence on our intellectual well-being.

One of the few quite general and relatively uncontroversial things that can be said about the aims of scientific research is that they are usually related to the goal of overcoming ignorance. In fact, rather than saying too much about the aims of research, the problem of this way of characterizing the overarching aim of science is that it says far too little. Our ignorance is, as has often been observed, enormous, and the really tricky problem with analyzing the aims of research is to find out how we identify those bits of ignorance that will be tackled next—or, in the normative version of the question, how we *should* identify them.

In order to be able to select a particular area of ignorance and make it the target of inquiry, one first has to be aware of one's own lack of knowledge in this particular area. Our awareness of our own ignorance therefore plays a role in guiding our efforts to gain new knowledge. My aim in what follows is to explore this connection between ignorance and the aims of research.

I will take for granted not only that science typically *aims* to overcome ignorance, but also that it often *succeeds*. In making this assumption, I may run the risk of being contradicted by a wide range of scholars from philosophy and sociology who have argued that whenever our knowledge expands, our ignorance grows along with it (see e.g. Mittelstraß 1998, 4, Krohn 2001, 8141, Wehling 2006, 83). This claim is sometimes explained by referring to the new questions that arise with every new item of scientific knowledge; sometimes it is also illustrated by the

many uncertainties that are brought into the world with new technologies whose development often goes hand in hand with the creation of knowledge.

However, something about the statement that extending our knowledge always or typically adds to our ignorance seems wrong, even paradoxical. (In fact, sociologists refer to it as a “well-known paradox,” cf. Gross 2007, 743.) The openly paradoxical nature of the claim makes it hard to dispute it, because, as Bernard Williams (2008, 5) once pointed out, in order to criticize an assertion that is articulated as a paradox, you will first of all have to point out that the paradox is paradoxical and hence cannot be true—which was obvious all along and is therefore going to make you look like someone who has missed the point of a joke. However, I think that Anna Leuschner (2012, 100) has managed to avoid creating this impression and to pinpoint the problem very well. She cites the example of our ignorance of the exact contribution of cloud formation to the world climate. This ignorance, Leuschner argues, was not *created* by climate research. Our Stone Age ancestors were no less ignorant about this than we are. What is decisive, she concludes, is that every new piece of knowledge creates *awareness* of a new finite subset of the infinite set of things that one does not know.

I want to bring out this point by identifying *conscious ignorance* as a subset of our *total ignorance*. Total ignorance is just the complement of knowledge. But what we know we don’t know—or conscious ignorance, in short—makes up only a small part of this infinity. The rest, what we don’t know we don’t know, I will call opaque ignorance. Distinctions like these have, of course, been made before: from Jerome Ravetz’ (1993) definition of “ignorance squared,” or “ignorance of ignorance,” to Donald Rumsfeld’s much-debated rhetoric of “known unknowns” and “unknown unknowns.” Paying explicit attention to the distinction can be helpful to agnotology, or so I shall try to demonstrate.

For a start, it helps to avoid paradox. What the many references to our growing ignorance allude to is an increase in our conscious ignorance, not our total ignorance. An increase in our total ignorance is not impossible, but it can only be brought about by destroying knowledge. Whenever we forget something, or erroneously dismiss a thing that we had previously known, our total ignorance grows. But obviously, that is not what those who speak of an increase in our ignorance have in mind. Climate research did not destroy any previous piece of knowledge on the connection between cloud formation and mean temperatures, but it brought the respective lack of knowledge to our awareness, thereby creating new conscious ignorance.

The claim that research is aimed at overcoming ignorance can now also be made more precise and can be differentiated into two claims that operate on different levels. On the level of the individual research project, each such endeavor is targeted at a more or less well-described area of ignorance and aims to convert it into knowledge. This idea—that research is always *directed* at something—automatically means that only *conscious* ignorance can play a role for the goals of research. Perhaps Francis Bacon once thought that the inductive sciences always needed to start from a clean slate, or perhaps this is too simplistic a reading of even Bacon’s radical brand of inductivism. In any case, post-positivistic philosophers, historians and sociologists of science alike have stressed the fact that inquiry is always laden with theoretical preconceptions of the area into which it advances. Science does not blindly forage into the vastness of opaque

ignorance. Occasional spells of serendipity, like Röntgen’s discovery of X-rays, or Ørsted’s of electromagnetism, do nothing to disprove this insight, as two kinds of consideration should make clear. First, the claim is only that research always needs to be *directed* at a piece of conscious ignorance, not that what is eventually found is always identical to what one has been looking for. And secondly, even serendipitous discovery requires a prior awareness of the possibility of a causal nexus, however tentative. As Pasteur famously remarked, “chance favors only the prepared mind” (1939, 131).

Whenever science succeeds in converting a given area of conscious ignorance into knowledge, new conscious ignorance is almost inevitably created. This might at first sight seem to undermine a claim that operates on a more global level: viz., that the central institutional aim of science is to reduce our ignorance. However, if we understand that the global claim refers to a reduction of our *total* ignorance, and not our conscious ignorance, the inconsistency dissolves. An increase in our conscious ignorance can even be regarded as a first and necessary step towards the goal of reducing our total ignorance if, as I have just suggested, it is only conscious ignorance that can be intentionally targeted by inquiry. This means that opaque ignorance must usually first turn into conscious ignorance before it can be converted into knowledge.<sup>1</sup>

## 2. Conscious ignorance

The notion of conscious ignorance thus seems to be of crucial importance. How exactly should it be understood? So far, it has only been defined in a rather vague manner and might even seem to be beset with a few paradoxes of its own. What could it mean to be conscious of a particular piece of one’s own ignorance? What could it mean to target one’s epistemic efforts at something that one does not know? These quandaries may be understood as varieties of the famous question that Meno asks Socrates: “And how will you inquire into a thing, Socrates, when you are wholly ignorant of what it is? What sort of thing among those you don’t know will you set up as the object of your inquiry?” (Plato 1984, 80d) A certain understanding of this question is pertinent to our present topic. If I conceive of a particular piece of ignorance as an item of non-knowledge—that is, a true proposition that I do not yet know—then it seems that I would only be able to direct my epistemic efforts at such an item if I already knew it—and knew it to be a true proposition.

What this shows is that our conscious ignorance in the present sense cannot be understood as a set of true propositions lying out there, waiting to be discovered. Instead, our conscious ignorance is best understood as a set of *questions*. While we ultimately want to *find* correct answers to questions and correct answers are true propositions, we *aim* our efforts at questions.

---

<sup>1</sup> Our total ignorance can only be “reduced” in a very loose sense, because it will always be infinite and the chunks of it that we convert into knowledge are not large enough to diminish its cardinality. The sense in which one might nonetheless speak of a reduction is that, if  $J_t$  and  $J_{t+\Delta}$  signify our total ignorance at different points in time,  $J_{t+\Delta} \subsetneq J_t$ .

Nikolaj Nottelmann (2016, 45-46) has taken issue with this line of reasoning; he finds my focus on questions “excessive”.<sup>2</sup> Nottelmann argues that while it may be impossible to self-ascribe ignorance by means of direct reference to the fact that one is ignorant of (as in “I do not know that my keys are on the dining table”), this only concerns the way in which first person *ascriptions* of ignorance must be made and does not in any way detract from the fact that every instance of ignorance is constituted by a fact (or set of facts) that the subject does not know. For the purposes of the present paper, I might rest content with the claim (acknowledged by Nottelmann) that self-ascriptions of ignorance in the form of questions are of particular importance in the specific context that concerns me here (i.e., conscious ignorance as it informs the aims of inquiry), which is inevitably bound to the first person perspective. But my disagreement with Nottelmann’s objection runs a little deeper. While I concede that every episode of ignorance involves a subject not knowing a fact (or set of facts), I doubt that it is the most helpful analysis to *identify* each episode of conscious ignorance with the subject not knowing a certain fact or set of facts. Suppose that (i) Tom does not know to whom if anyone Sue is married, and that (ii) he also does not know to whom if anyone Sally is married, and that as a matter of fact, Sue and Sally are married to each other. Do the ignorance ascriptions (i) and (ii) really refer to one and the same piece of Tom’s ignorance, as the identification of ignorance with unknown facts would suggest? Nottelmann might insist they do, and that the two ascriptions differ not with regard to the bits of Tom’s ignorance that they pick out but only in the ways in which they pick out one and the same bit. I am somewhat doubtful that there is a fact of the matter grounding the correct answer to the question whether ignorance “really is” constituted by facts or by questions, and I suppose one could adopt either of these ways of speaking. But if ignorance is to be something that can motivate an effort of discovery, that one can direct inquiry at and that one can be aware of in oneself even before one has gained the knowledge that cancels it out, then the latter way of speaking is not merely of particular importance, but is indispensable for coming to a straightforward analysis of ignorance that facilitates this perspective.<sup>3</sup>

In fact, as we shall see, awareness of a question does not even necessarily presuppose awareness of *potential* answers to that question. So what *does* it presuppose? What are the conditions under which a person *P* can be said to be consciously ignorant with regard to question *Q*? Sylvain Bromberger has pointed out that the first and foremost presupposition that is relevant

---

<sup>2</sup> The present volume has been a long time in the making. This puts me in the unusual position to include here a response to a published criticism of (a manuscript version of) this very paper. I have left the passages to which Nottelmann refers unchanged.

<sup>3</sup> One *might* also develop a perspective on conscious ignorance based on a factual conception of ignorance. What I mean by saying that such an analysis would not be straightforward can once more be illustrated with the example of Tom, Sue and Sally. Under a factual conception, to discuss Tom’s awareness of his own ignorance one would have to say that the ignorance that Tom is aware of in case of proposition (i) is the same ignorance as the one that he is aware of in the case of proposition (ii). Possessing awareness in each of these respects, one would have to continue, does not presuppose being aware that one is aware of one and the same item of ignorance in both instances, just as one can be aware of the evening star and of the morning star without being aware that they are one and the same thing. However, this seems to me to posit that there is something that Tom is missing about his own ignorance in both cases (as, in comparison, the underinformed stargazer is surely missing something about the evening star and the morning star). I would submit that Tom’s awareness of his own ignorance in the proposed scenario need not in any way be considered incomplete or lacking. He is missing the correct *answers* to his questions (as he well knows), but there does not seem to be anything that is lacking *in his awareness of his ignorance*.

in this context is what he calls *representation*: *P* must be able to formulate *Q* in a language that she is competent in (1992, 129-31, 147). From the early 1960s on, Bromberger was one of the few philosophers who approached philosophy of science from an ordinary language perspective.<sup>4</sup> He is probably most widely remembered for having devised a famous counterexample to the covering law model of scientific explanation—the story of the flagpole and the shadow. But as I shall shortly return to some of his contributions, I hope to demonstrate that he deserves to be rediscovered as a thought-provoking theorist of ignorance.

### 3. Conditions of conscious ignorance

For now, I will take up Bromberger’s suggestion to take the representation of a question seriously as an important precondition of conscious ignorance:

- (1) *P* is able to articulate or at least understand *Q* and is aware of *Q*.

I take it that *P* need not necessarily be able to find her own articulation of *Q* in order to be consciously ignorant with respect to *Q*, but that understanding some articulation of *Q* is sufficient. In any case, our Stone Age ancestors’ ignorance in Leuschner’s example qualifies as opaque because for all we know they fail to meet even this first precondition of conscious ignorance. As a second condition, *P* will obviously only count as ignorant if she fails to have the relevant knowledge:

- (2) *P* does not know a correct answer to *Q*.

However, *P*’s failure to know a correct answer will only count as an epistemic shortcoming if *Q* admits of a correct answer in the first place. We understand the question “Why does yellow bile cause warm diseases, while phlegm causes cold ones?” We do not know a correct answer to it, and yet we would not count this as an item of ignorance. This observation prompts the addition of a third condition:

- (3) *Q* is sound (i.e., possesses a correct answer).

Conditions (1) through (3) describe what it means to be conscious of a question (condition 1) and ignorant with regard to it (2 & 3). However, consciousness plus ignorance does not always amount to conscious ignorance. Consider the question whether we would describe Christopher Columbus as consciously ignorant with regard to the question: “What is the distance between the Canary Islands and Japan?” To the best of our knowledge, Columbus satisfies conditions (1) through (3): He *did* articulate the question, the question has a correct answer, and Columbus was without a doubt blissfully ignorant of that correct answer, as the distance is almost five times as great as Columbus believed it to be. However, while he was aware of the question and ignorant of its correct answer, he was not aware *of his own* ignorance, for he thought he knew the correct answer to a good degree of approximation. A condition needs to be added:

---

<sup>4</sup> Two others were Stephen Toulmin and Michael Scriven.

(4) *P* believes that she does not know a correct answer to *Q*.

Columbus was not consciously ignorant of the correct distance, because condition (4) was not fulfilled in his case. In addition, conscious ignorance also requires the awareness that the question does have a correct answer. For example, we would not describe Newton as consciously ignorant with regard to the question how light waves are propagated. He was aware of the question, he believed that he knew no correct answer, but he mistakenly thought that this was because there *was* no correct answer. We should therefore add one last condition for conscious ignorance (cf. Bromberger 1992, 131-33):

(5) *P* regards *Q* as sound.

One insight this analysis of conscious ignorance immediately facilitates is that conscious ignorance is quite a demanding affair (as Bromberger [1992, 131, cf. 31] already observed). Ignorance in this sense is not just an absence or a negativity. It is rather a complex ability and can often require quite an amount of knowledge and competence.

The fact that conscious ignorance can be an achievement and an important first step on the way to knowledge is emphasized by Socrates in the *Meno*. Frustrated by Socrates' refutations of his repeated attempts at defining virtue, an angry Meno accuses him of causing only perplexity in others, thus "benumbing" his interlocutors (Plato 1984, 80a). (Plato also has Meno mention Socrates' notoriety for this among his fellow Athenians. He thus calls attention to the fact that pointing out ignorance can be inconvenient and unwelcome and thereby alludes to Socrates' later fate.<sup>5</sup>) But only a little while later, the famous slave boy scene gives Socrates occasion to defend himself. Having just helped the slave boy realize that his initial idea on how to double the area of a square is mistaken, Socrates remarks to Meno: "We have at any rate done something, it seems, to help him discover how things are, for in his present condition of ignorance, he will gladly inquire into the matter [...]." Meno agrees that, in this case, "numbing befitted" the slave boy (84 b-c). Let us call the achievement of conscious ignorance as a first step towards knowledge "Socratic progress."<sup>6</sup>

However, not everyone who dons the mantle of Socrates is interested in enlightenment and the growth of knowledge. The path between knowledge and ignorance can be travelled in both directions. You count as promoting Socratic progress only if you introduce conscious ignorance in order to replace error or other forms of opaque ignorance. If you do it with the aim of replacing knowledge or preventing its generation, you are but an ignorance-monger. While telling whether you are one or the other may therefore often be a difficult and controversial matter (as the question whether the established belief that you are replacing by conscious ignorance constitutes error or knowledge is typically itself disputed), it is by no means impossible. Perhaps most significantly, an advocate of Socratic progress can be expected to be interested in continuing the process further beyond ignorance and in the direction of knowledge, and to not treat ignorance

---

<sup>5</sup> Later in the same dialogue, the dark foreshadowing is continued by the appearance of Anytus, who, as Plato's Athenian readers would have been aware, would become a primary prosecutor in the trial of Socrates only three years after the dialogue's fictitious date.

<sup>6</sup> Socratic progress is, of course, a recurrent topic in Plato's dialogues. Cf. esp. *Apology*, 21a-e.

itself as the goal. (A dead giveaway is therefore when you write up an internal memo which declares that “Doubt is our product,” cf. Michaels 2008, 11, Oreskes & Conway 2010, 34.)

While our conscious ignorance is limited as a result of its dependence on certain cognitive preconditions, it is still vast. (As most of us are aware that an infinite number of questions of the form “Is  $n$  a prime number?” can meaningfully be asked, one might argue that even our *conscious* ignorance is infinite in extension.) Obviously, belonging to the realm of conscious ignorance is only a *necessary* condition for becoming the target of inquiry. The question therefore arises whether one can find further characteristics that help to distinguish between different kinds of conscious ignorance. Are there perhaps kinds of conscious ignorance that are more likely to be targeted by research than others?

To approach the matter of relevant distinctions within the body of our conscious ignorance, let me start by giving you a small sample of the immensity that is my personal ignorance. (As a side effect, this will give me the opportunity to prove my outstanding qualifications for writing on the topic of ignorance.)

#### 4. Torsten Wilholt, philosopher of great ignorance

- (i) I do not know the distance between Calgary and Miami.
- (ii) The roof of Bielefeld University’s main building is covered with pebbles. I do not know their number.
- (iii) I do not know what it is that sometimes makes people sneeze when they look into the sun.
- (iv) I do not know why ice remains slippery at temperatures way below its melting point, even if you (attempt to) stand still on it and thus do not produce any heat by means of friction.<sup>7</sup>
- (v) I do not know whether it is physically possible for there to be a transuranium element with a half-life of a day or more.
- (vi) I do not know whether any person at this moment knows a correct answer to any of the questions involved in (ii)-(v).

Some differences between these six items of ignorance are fairly obvious. Knowing the distance between Calgary and Miami has so far never been of any practical value to me. However, this kind of knowledge is of obvious practical value to some people, which is why I consider it safe to assume that there are people who do not share my ignorance with regard to item (i). In contrast, in case (ii), I cannot easily imagine who might have a practical interest in the correct answer. I do not expect that anyone will ever know it, nor that this piece of ignorance will ever bother anyone. With regard to item (v), whether or not the correct answer to the respective question is of any practical use to anyone depends, among other things, on what the correct answer is. Some people have speculated that stable transuranium elements would have properties that make them

---

<sup>7</sup> The common lore is that a thin film of water between one’s soles (or skates) and the ice accounts for the slipperiness. But how can such a film remain liquid at temperatures of, for example,  $-20\text{ }^{\circ}\text{C}$ ? Water ice can in principle be melted by means of pressure (because the density of water near the melting point is higher than the density of ice), but the pressure that your body weight exerts on the ice will lower its melting point by a few degrees at the very most.

interesting for practical uses, but this would, of course, only confer significant practical value on the correct answer if the answer is an affirmative one.<sup>8</sup> In a similar fashion, the practical value of the knowledge that would cancel out my ignorance in cases (iii) and (iv) depends on what exactly the correct answers to the questions are. The practical benefits of converting ignorance into knowledge thus vary greatly among the questions that make up our conscious ignorance. This is evidently of great importance for the selection of targets for research. How great the practical value will be is, however, often itself an item of ignorance. While our research agenda is shaped by questions, practical value accrues only to answers. In many cases, the value of the answer can only be known after it has been found. This is one of the reasons why the selection of research topics on the grounds of practical considerations typically involves decision-making under ignorance and cannot be regarded as a simple and straightforward matter of maximizing expected value.

However, variations in practical usefulness are not the only differences that stand out when we examine this little sample of my conscious ignorance. I do not expect that answers to the questions in (iii) and (iv) would be of any more practical use *to me* than answers to (i) or (ii), but nevertheless these items of ignorance exert an attraction on me that the others don't. The deficits in my knowledge signified by items (iii) and (iv) seem to be more glaring than the others, they represent a deeper kind of ignorance. It's not just that I lack a piece of information—I'm at an impasse. I'm at a loss for an answer. These descriptions of the relevant difference are, as I am aware, unsatisfying. For help, it is now high time to once more turn to the work of Sylvain Bromberger.

## **5. Sylvain Bromberger, great philosopher of ignorance**

At the very outset of his ordinary language approach to the theory of scientific explanation, Bromberger asks the question: What distinguishes those episodes that we call "explaining" from other information-giving episodes? Is it perhaps the form of the question that is being answered? Or the form of the answer given? Bromberger discards these easy answers. He thinks that there can even be two episodes where the same kind of question is asked and the same kind of answer is given in both cases, but one counts as an explanation and the other doesn't. Consider the following example, given by Bromberger himself in his classic essay "An Approach to Explanation" from 1962.<sup>9</sup> Prisoner Tom escapes from his prison by digging a tunnel. After his escape, the tunnel is discovered, as are the tools and devices that Tom used to dig the tunnel and hide its entrance etc. But one question leaves the prison guards stumped: How did Tom get rid of the dirt that he had to remove from the tunnel? In a second episode, John escapes from a different prison, also by digging a tunnel. Everything is very similar to Tom's case, except that in John's former prison, the guards realize that there are two obvious ways John might have used to dispose of the dirt—by either dropping it into a moat or onto a patch of garden. Both Tom and John are

---

<sup>8</sup> The only practical value of possessing a correct negative answer would be that you could stop wasting time and energy on the search.

<sup>9</sup> Reprinted in Bromberger 1992: 18-51.



later recaptured and interrogated, and both reveal the details of their escapes, including the respective ways by which they actually disposed of the dirt from the tunneling. Now Bromberger observes that, in Tom's case, we would without hesitation say that he *explained* to his jailers how he got rid of the dirt. However, in John's case, who only reveals which one of two fairly obvious opportunities he actually chose, saying that he explained his method of disposing of the dirt does not sound quite right. In Bromberger's own words, it seems "out of place, distorts things, smacks of exaggeration, is at best a near truth" (1992, 25).

In effect, Bromberger's conclusion is that whether or not a question calls for an *explanation* does not so much depend on the question's form as on the profundity of the ignorance that lies behind it. In the case of Tom's jailers, their ignorance is deeper because not only do they fail to know which is the *correct* answer to the question "How did Tom dispose of the dirt from the tunneling?"—they do not even have any plausible *candidate* answers. Bromberger calls this state "p-predicament," formally defined as follows.

*A* is in a p-predicament with regard to *Q* if and only if, in *A*'s views, *Q* admits of a right answer, but *A* can think of no answer to which, in *A*'s views, there are no decisive objections.

I freely admit that I am in a p-predicament with regard to the question why ice is slippery at  $-20^{\circ}\text{C}$ , and similarly with regard to the question what makes you sneeze when you look into the sun. To the extent that some tentative answers spring to mind at all, they all seem to be seriously flawed on closer reflection. I simply have no idea.

Take Bromberger's claim that it is in situations of profound ignorance such as these that questions call for explanations, and combine it with the truism that giving explanations is a central aim of the sciences, and you will raise at least an initial suspicion that, *ceteris paribus*, p-predicaments should offer particularly attractive targets for scientific research.

I think that we can avoid the detour via the problematic notion of explanation altogether and give a more immediate motivation for the basic idea that deep ignorance in the sense of Bromberger's p-predicament makes for a particularly attractive target of research.<sup>10</sup> It has often been noted that in the sciences, *fruitfulness* counts in favor of an idea. A good theory, for example, should, in Thomas Kuhn's words, "disclose new phenomena or previously unnoticed relationships among those already known" (1977, 322). In this spirit, the sciences not only

---

<sup>10</sup> Bromberger's identification of overcoming deep ignorance as a hallmark of *explanation* may lead to some unexpected consequences. Consider the fact that under normal circumstances, someone's p-predicament with regard to any particular question is only overcome once. Once the person has learned at least one plausible candidate answer, the p-predicament is gone and will not return (unless she forgets again what she has learned, or learns of new convincing objections against it). This would mean that you can normally only receive any explanation once. It would also seem to imply that what we would usually call a better, more precise or more complete explanation of something that has previously been explained to us in a more preliminary way, is really an explanation that responds to a new and different question (if it is an explanation at all). For example, Newton's explanation of why the planets are observed in the positions that are recorded in the astronomical tables is not an improved answer to the same question that Kepler's explanation had already addressed; its explanatory import must come from answering a different question ("Why conic sections?"). I will not, however, go deeper into a discussion of the merits of p-predicament as a key to explanation. What matters to me is the recognition of p-predicament as an important special case of ignorance, which is in no way affected by the aforementioned peculiarities.

answer questions, they also raise new questions about every answer they consider. That way, a network of questions and answers is created, bestowing significance on each other along the lines of the answerhood relation. (By describing things in this way, I am borrowing the forceful image of “significance graphs” that Philip Kitcher has created [2001, 63-82].) Even tentative or merely potential answers can be used to raise new questions and establish new relations of answerhood, partial answerhood or potential answerhood. The buck stops only when you arrive at a question to which there are no known plausible candidate answers—a p-predicament. Overcoming such an item of deep ignorance thus holds the promise of opening up a whole range of new possibilities for questions and answers, discovering new connections and establishing new nodes in the network of knowledge and conscious ignorance.

Before we take a step back and revisit the relations between ignorance and the aims of research, one last technical remark is in order. Like with ignorance itself, also its depth in the present sense may sometimes be judged differently by the ignorant person herself and from a third person perspective. Consider the case of Urbain Le Verrier when he was searching for a theretofore undiscovered planet whose gravitational influence would explain the precession of the perihelion of Mercury’s orbit. Presumably, he would have admitted his ignorance of the cause of the precession, but not that it was a deep kind of ignorance. He did, after all, think that he had a good candidate answer at hand. Objectively speaking, however, his ignorance was as deep as anyone’s necessarily was before the theory of general relativity was first articulated. The *correct* answer was not on his radar, and couldn’t have been. The case fits Bromberger’s definition of b-predicament. A person is in a b-predicament with regard to a question if and only if the question is sound, but the correct answer, in Bromberger’s words, “is beyond what the person [...] can conceive, can think of, can imagine, that is, is something that that person cannot remember, cannot excogitate, cannot compose.” (1992, 36) Episodes of information-giving that cancel out someone’s b-predicament also count as explaining episodes in Bromberger’s view. However, with regard to our present topic, the connection between ignorance and the aims of research, our focus should be on the cases where individuals or groups are *aware* that they are at a loss for a potential answer, i.e., cases of p-predicament. It is for these cases that I will reserve the label “deep ignorance.”

If my earlier considerations are not altogether misplaced, then it is this variety, conscious ignorance combined with an awareness that no plausible candidate answers are available, that should hold particular attraction for researchers. Sometimes this attraction finds expression in the writings of scientists. Thus James Clerk Maxwell (1877, 245), commenting on the lack of any plausible explanation for the experimental results for heat capacity ratios (and in particular for why they failed to conform to the predictions of kinetic theory), mused that this was “likely to startle us out of our complacency, and perhaps ultimately to drive us out of all the hypotheses in which we have hitherto found refuge into that state of thorough conscious ignorance which is the prelude to every real advance in knowledge.” (In fact, the anomaly could only be explained in the 20th century and with the aid of quantum mechanics.)<sup>11</sup>

---

<sup>11</sup> Cf. Brush 2002: 121, where Maxwell’s statement is also quoted.

But once more, the reality of the sciences stubbornly refuses to yield to philosophical generalizations. Conscious deep ignorance is not always a guarantee for the sustained attention of scientists. When Alexandre Edmond Becquerel discovered the photovoltaic effect in 1839, it did attract a lot of attention from physicists. But after about two years, interest waned, although the causes of the effect were a matter of deep ignorance at the time. Similarly, Brownian motion, which had been described by the Dutch physician Jan Ingenhousz as early as 1784, for a long time received only scattered bouts of attention before it became a piece of evidence in the case for atomism in the early 20<sup>th</sup> century.<sup>12</sup>

Even questions that represent areas of conscious deep ignorance can thus be passed over when the sciences settle their research agenda. It is easy to think of factors that may contribute to this. Success in converting deep ignorance into knowledge and understanding may be highly regarded in the sciences, but that fact will only provide motivation for tackling an item of deep ignorance if there seems to be at least a reasonable chance of success. In cases where there are no approaches for tackling the question available from the prevalent inventory of methods, the wager of nonetheless confronting the problem is less likely to be undertaken. The research agenda is not only shaped by scientists' conceptions of what seems to be rewarding, but also by what seems to be within reach.

The conclusions of these reflections should therefore be articulated with care. While belonging to the realm of conscious ignorance is a necessary condition for becoming the target of research efforts, being perceived as a problem of practical value and being perceived as a deep problem in the sense that we have identified with Bromberger's aid can be considered conditions that *favor* an item's inclusion in the current research agenda.

## 6. Beyond normal: Kinds of opaque ignorance

While conscious ignorance thus comes in various degrees of profundity, opaque ignorance may upon first reflection seem to consist of one uniform block of negativity. But a closer look reveals important differences between kinds of opaque ignorance. As we have seen, conscious ignorance represents a complex ability that rests on several preconditions. Accordingly, there are different ways in which a subject may *fail* to know what she does not know. In particular, it makes sense to distinguish cases in which she is unable to articulate or comprehend a question from other cases of opaque ignorance.

Consider, for example, the question:

(P) What force makes the planets stay on their orbital paths?

For a pre-modern physicist in the Aristotelian tradition, this question would clearly not have belonged to the realm of conscious ignorance. In the Aristotelian framework, all supralunar bodies by their nature remain in perennial circular motion. To ask for a force that effects this would have made no more sense to an Aristotelian than the question "What force keeps a

---

<sup>12</sup> Both cases have been used as examples for phenomena that remained "meaningless" for contemporaries over an extended period of time by Hacking (1983: 158).

baseball flying after it has left the pitcher's hand?" would make to a modern physicist. Additionally, you could argue that an Aristotelian physicist would not have had the required concept of force at his disposition to articulate or comprehend question (P). Now let us stipulate that a particular hypothetical Aristotelian is also unaware that he does not know the correct answer to the question

(M) How many legs does the mayfly have?

(Perhaps he has been misled into believing that the answer is four by reading the *Historia animalium*.<sup>13</sup>) Assume that his unawareness of his own ignorance with regard to (M) is mainly due to a lack of interest in zoological matters. Then, there seems to be a decisive difference between the two items of opaque ignorance that we have ascribed to him. While his unawareness of (M) is more or less a matter of coincidence and could easily be overcome, discovering his ignorance with regard to question (P) is a matter of near impossibility to him. If he were to manage to do so and transform (P) into an item of conscious ignorance, it would be a tremendous intellectual achievement and possibly a first step towards an episode of revolutionary science in Kuhn's sense.<sup>14</sup> To introduce one last label for convenience, let us use the expression "thoroughly opaque ignorance" for the kind that remains opaque to us because our background beliefs and conceptual repertoire make us either unable to articulate the respective questions or to recognize them as sound.

These last observations may raise a suspicion. Perhaps our reflections on ignorance and the aims of research are bound to remain restricted to what Kuhn (1993, 35-42) calls normal science as long as we concentrate too much on *conscious* ignorance and its importance for the direction of inquiry. Perhaps revolutionary science (science beyond mere "puzzle-solving") can only occur at the interface of conscious and thoroughly opaque ignorance, because it essentially involves the emergence of new conceptual frameworks and thus of answers to questions that could not have been articulated before.

But putting it thus could be misleading. It could be taken to suggest that revolutionary science must of necessity involve "aimless" research—research that is not directed at any question. That, however, would be a mysterious picture of revolutionary science. How can you even engage in any practice of inquiry when there is no question that you are pursuing? But, on the other hand, if the novel questions characteristic of revolutionary science presuppose novel concepts that are not available within the pre-revolutionary paradigm, how can the process ever get started?

The solution to this puzzle is a feature of research that we have encountered already: Inquiry always consists of *pursuing* a question, but what you end up *finding* is not always an answer to that very same question. Besides the possibility that you may chance upon an entirely unexpected empirical discovery (as in the cases of Röntgen and Ørsted, mentioned above), there is also a chance that you may discover that the question is not sound. In your attempt to re-

---

<sup>13</sup> Cf. Aristotle 1910, 490a33-b3 and 552b17-23, where it is claimed that the mayfly (*ephemeron*) walks on four legs.

<sup>14</sup> I am grateful to my colleague Paul Hoyningen-Huene for alerting me to this difference in kind between items of opaque ignorance and to its connection to the distinction between normal and revolutionary science.

articulate, you may initiate conceptual change. Alternatively, you may find a potential answer to the question, but in doing so realize that every potential answer to it stands in stark contradiction to established theoretical belief. This may motivate you to seek your luck in massive theoretical modifications. In either case, it is not necessary that the question you were pursuing in the first place was not articulable in the terms of the old paradigm. Here are some examples of questions that have at some point played a role in periods of massive theoretical and conceptual change: “What would we observe if we chased alongside a light beam at the speed of light?” “Under what conditions are two events that occur at different points in space simultaneous?” “What happens when two freely falling heavy bodies are connected in mid-fall?” “Why does the electron in an H atom not spiral into the core, emitting radiation of greater and greater frequency?”

These questions were, I submit, understandable even before the respective episodes of revolutionary change that they are associated with had occurred. It was in taking them seriously and pursuing them (amongst other questions) that Einstein, Galileo and Bohr encountered deep-seated problems which led them to attempt radical theoretical-conceptual adjustments. These adjustments, in turn, enabled them to pursue other, novel questions, thus opening up whole new areas of conscious ignorance that had been thoroughly opaque before. This mechanism gives a little more flesh to the bones of the idea of working “at the interface of conscious and thoroughly opaque ignorance.” It also once more underlines the importance of *deep* conscious ignorance. A question that persistently appears to point to an item of deep ignorance is an anomaly in Kuhn’s sense; untiringly investigating the phenomena, models and theories surrounding it *may* one day push open the door to a region of ignorance that was hitherto thoroughly opaque.<sup>15</sup>

## 7. Conclusions and epilogue: Ignorance—A construction manual

In summary, examining ignorance under the aspect of the aims and directions of inquiry reveals the following landscape of ignorance. One way to conceptualize ignorance is to think of it as the totality of true propositions that we do not know. While this idea, which I have termed “*total ignorance*,” has certain theoretical uses, I have argued (with Bromberger) that with respect to our *conscious ignorance*—the part of our ignorance that we are aware of—it is best to think of it as a set of questions rather than a set of unknown true propositions.<sup>16</sup> Not only does this help to avoid Meno’s paradox, it also enables us to identify in greater detail the preconditions of conscious ignorance. They are, in short, the ability to articulate a question, recognize it as sound and acknowledge one’s own ignorance of the correct answer.

---

<sup>15</sup> Some of the questions that *appear* to signify an item of deep conscious ignorance will in the process be discovered to never actually have been associated with conscious ignorance as I have defined it. That is because I have opted for a “realistic” understanding of ignorance, under which a question only counts as identifying a piece of conscious ignorance when it does *in fact* possess a correct answer (condition 3).

<sup>16</sup> For the purposes of making sense of figure 1, which represents conscious ignorance as a subset of total ignorance, we may nonetheless think of every question that belongs to our conscious ignorance as represented, in that diagram, by its correct answer.

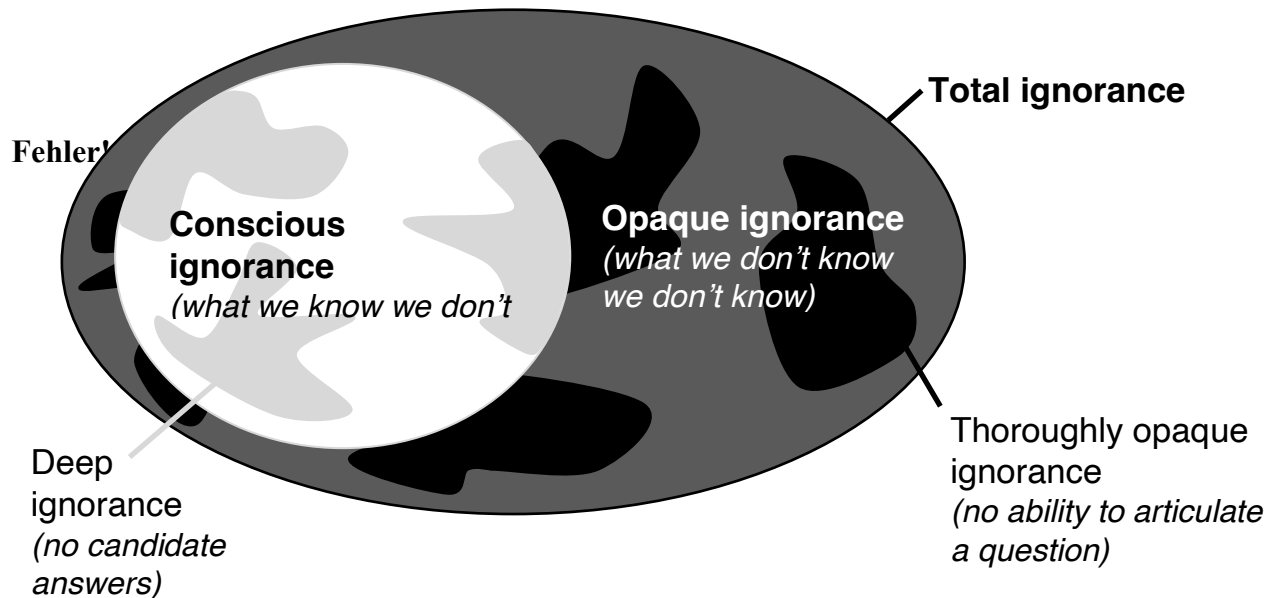


Figure 1: The landscape of ignorance

Conscious ignorance thus always presupposes certain capacities. If and when it replaces error or other forms of opaque ignorance, reaching the state of conscious ignorance constitutes an epistemic achievement—often a quite considerable one—which I have proposed to call “Socratic progress.” Socratic progress is often a crucial step on the way from opaque ignorance to knowledge.

Not all conscious ignorance is of the same kind. Cases where we can identify a range of possible answers and our ignorance concerns only the question which of them is the correct one should be distinguished from cases of *deep ignorance*, where we do not even know any candidate answers yet (or none to which there do not seem to be immediate and decisive objections). These particularly challenging pieces of ignorance promise to also be particularly rewarding once the riddle has been solved, as they may potentially open up new possibilities for questions and answers and unblock new paths of inquiry. (While deep ignorance thus marks a critical subspecies of conscious ignorance, it is important to keep in mind that other differences among items of conscious ignorance, such as differences regarding expected practical usefulness, are equally influential with respect to the aims and directions of research.)

The things we do not know without being aware of it constitute our *opaque ignorance*. Here too, different kinds can be distinguished. In some cases, we *could* easily become aware of our ignorance; it would “only” be a matter of directing our attention to the respective questions. But some items of opaque ignorance, the ones I have termed “*thoroughly opaque*,” are not straightforwardly accessible to us. We lack the conceptual resources and / or the background beliefs to articulate (and recognize as sound) questions that target them. Before thoroughly opaque ignorance can be replaced by knowledge, conceptual-theoretical transformations must precede. Those may come as a result of research that is targeted at conscious deep ignorance.

Is there thus an intimate connection between deep ignorance, on the one hand, and thoroughly opaque ignorance, on the other? Again, one should answer with care. While persistent

deep ignorance *can* be an indicator of the kind of anomaly that will one day lead to a crisis of the existing paradigm, it need not necessarily be so. Our inability to come up with plausible candidate answers may have other causes than the inadequacy of our conceptual-theoretical framework: undiscovered problems with the experimental apparatus, insufficient computational power, or pure lack of imagination, to name just a few. And on the other side, deep ignorance has no monopoly on crisis-initiating potential. Questions to which we *do* have candidate answers can turn out to resist every attempt to settle on the correct one and grow into anomalies. (The most straightforward examples of this are probably dichotomous questions like “Are electrons particles or waves?”)

I opened this essay with the apparent platitude that science aims at overcoming ignorance. Let me close with a few reflections on activities pursuing the opposite objective. Imagine that ignorance with regard to a particular point is what you *want*. What would all this mean to you then?

Obviously, if you want  $X$  to be ignorant about  $Q$ , where  $X$  can be an individual or a collective and  $Q$  is some question, you have to prevent or undermine  $X$ 's knowledge of the correct answer to  $Q$ . The most practical way of doing so is to prevent or undermine  $X$ 's *belief* in the correct answer.<sup>17</sup> But our reflections on different kinds of ignorance suggest that you can do better than that. An item of ignorance that you have created or are trying to protect may still be under threat of becoming the target of someone's inquiry. Better still than just preventing knowledge is preventing conscious ignorance as well.

Of the necessary conditions of conscious ignorance that we have discussed, three seem to present practical avenues to preventing it (while upholding the ignorance, of course). First of all, you can prevent  $X$ 's conscious ignorance by undermining their ability to even articulate  $Q$  (i.e., you can target condition 1). A variety of this strategy can be seen in some cases of strict and thorough secrecy. When the military authorities classified early predictions of global warming in the early 1940s, they in effect prevented the public from even articulating questions about it (cf. Proctor 2008, 19). Alternatively, if you happen to be invested with legislative powers, you might even forbid some questions to be articulated. Arguably, this is what happened to John Phillips, a student of Freeman Dyson's, who in 1977 presented a term paper on the possible construction of an atomic bomb, using only publicly available sources of information. His term paper was confiscated by the FBI on the account of an alleged infringement of the Atomic Energy Act.

If you can't prevent people from articulating a question, perhaps you can prevent them from considering the question sound (condition 5). This is a second avenue for preventing conscious ignorance. To give a fairly simple example: The question “What are the causes of global warming?” can be erased from a group's conscious ignorance if you can get that group to believe that global warming isn't taking place at all. Under that condition, questions regarding its

---

<sup>17</sup> Two other ways of preventing knowledge are (a) undermining the truth of a belief or (b) undermining the belief's justification or warrant. (a) is possible in *some* circumstances, e.g. if someone knows where your diary is hidden and you undermine her knowledge by simply putting it elsewhere. It is, I believe, not an option in *most* real life cases where someone has a vital interest in preventing or undermining knowledge. (b) does not play a role in such cases either, because if someone's knowledge would pose a threat to your interests, then so would her true belief alone. In such circumstances, undermining warrant or justification only makes sense as a means of undermining belief.

causes will not even appear sound to this group, and the group's ignorance in this respect may thus become particularly stable—as long as you can uphold their disbelief in the question's presupposition.

Finally, let us not forget that a necessary precondition for  $X$ 's conscious ignorance with regard to  $Q$  is also that  $X$  believes that she does *not* know the correct answer to  $Q$  (condition 4). A very useful way of undermining knowledge and preventing conscious ignorance at the same time is therefore to introduce error. Many cases of lying and deception in the history of knowledge testify to the effectiveness of this strategy.

It may be the case that none of these options are available to you, and the ignorance you have created is bound to remain glaringly conscious. If this is so, you should at least try to avoid the impression that the ignorance is deep—for the deep kind attracts attention. You should therefore always provide some candidate answers that remain available when you have successfully undermined belief in the true answer. An example for this strategy might be the tobacco industry's efforts to establish the idea of a "sick building syndrome" in order to provide at least a candidate answer to the question of what causes headaches and other health problems among workers in smoky offices (cf. Oreskes & Conway 2010, 140). Similarly, when the manufacturers of vinyl chloride found themselves confronted with a suspiciously high number of brain cancer deaths among their workers, industry-sponsored scientists came up with the concept of "diagnostic sensitivity bias," arguing that perhaps brain tumors are more likely to be diagnosed among workers in the chemical industry than among the general population (cf. Wilholt 2009, 93). What's striking is that such alternative answers already do some work even without any efforts to establish the belief that they are the *correct* answers, or even likely to be the correct answers. The mere presence of a candidate answer that is not yet refuted lifts the pressure of deep ignorance from the question.

Now who said that philosophy of science doesn't have some practical applications. But, alas, the ignorance-mongers seem to know all the good tricks already. In that case, I can only hope that the joint efforts to understand the dynamics between ignorance and the search for knowledge that the editors of the present volume have brought together will make it a little easier to expose their schemes.

## **Acknowledgements**

Early versions of this chapter were presented at the conference *Agnology: The Production, Preservation and Management of Ignorance* at the Centre for Interdisciplinary Research (ZiF) at Bielefeld University, May 29. – June 1, 2011 and at the research colloquium of the Institut für Philosophie at Leibniz Universität Hannover in the summer term 2012. I have benefited from participants' feedback on both occasions, as well as from individual conversations with Nikolaj Nottelman and with Angela Matthies (who contributed the sneezing example). I am grateful for Paul Hoyningen-Huene's very constructive and productive engagement with some of the ideas of this paper. Finally, I consider myself very fortunate to have made the acquaintance of Sylvain



Bromberger towards the end of the year 2015, if only via email. His kind and helpful messages to me displayed a mix of brilliance, wit and intellectual humility, and have left a lasting impression.

## References

- Aristotle. 1910. *Historia animalium*. In *The Works of Aristotle*, vol. 4, ed. J.A. Smith and W.D. Ross. Translated by D'Arcy Wentworth Thompson. Oxford: Clarendon.
- Bromberger, Sylvain. 1992. *On What We Know We Don't Know: Explanation, Theory, Linguistics, and How Questions Shape Them*. Chicago: University of Chicago Press and Stanford: CSLI.
- Brush, Stephen G. 2002. Cautious Revolutionaries: Maxwell, Planck, Hubble. *American Journal of Physics* 70 (2): 119-127.
- Gross, Matthias. 2007. The Unknown in Process: Dynamic Connections of Ignorance, Non-Knowledge and Related Concepts. *Current Sociology* 55: 742-759.
- Hacking, Ian. 1983. *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge: Cambridge University Press.
- Kitcher, Philip. 2001. *Science, Truth, and Democracy*. Oxford: Oxford University Press.
- Krohn, Wolfgang. 2001. Knowledge Societies. In *International Encyclopedia of the Social and Behavioral Sciences*, ed. Neil J. Smelser and Paul B. Baltes, 8139-8143. Oxford: Blackwell.
- Kuhn, Thomas. 1977. *The Essential Tension*. Chicago: University of Chicago Press.
- Kuhn, Thomas. 1993. *The Structure of Scientific Revolutions*. 3rd edition. Chicago: University of Chicago Press.
- Leuschner, Anna. 2012. *Die Glaubwürdigkeit der Wissenschaft: Eine wissenschafts- und erkenntnistheoretische Analyse am Beispiel der Klimaforschung*. Bielefeld: Transcript.
- Maxwell, James Clerk. 1877. Review of H. W. Watson, A Treatise on the Kinetic Theory of Gases. *Nature* 16: 242-246.
- Michaels, David. 2008. *Doubt is Their Product*. Oxford: Oxford University Press.
- Mittelstraß, Jürgen. 1998. Das Undenkbare denken: Über den Umgang mit dem Undenkbaren und Unvorstellbaren in der Wissenschaft. In *Virtualität und Realität: Bild und Wirklichkeit in den Naturwissenschaften*, ed. Gottfried Magerl and Kurt Komarek, 1-23. Wien: Böhlau.
- Nottelmann, Nikolaj. 2016. The Varieties of Ignorance. In *The Epistemic Dimensions of Ignorance*, ed. by Rik Peel and Martijn Blaauw, 33-56. Cambridge, Cambridge University Press.
- Oreskes, Naomi, and Erik M. Conway. 2010. *Merchants of Doubt*. New York: Bloomsbury.
- Pasteur, Louis. 1939. *Oeuvres de Pasteur*, vol. 7, ed. P. Vallery-Radot. Paris: Masson.
- Plato. 1984. *Meno*. In *The Dialogues of Plato*, vol. 1, 151-186. Translated by R.E. Allen. New Haven: Yale University Press.
- Proctor, Robert N. 2008. Agnotology: A Missing Term to Describe the Cultural Production of Ignorance and Its Study. In *Agnotology: The Making and Unmaking of Ignorance*, ed. Robert N. Proctor and Londa Schiebinger, 1-33. Stanford: Stanford University Press.

- Ravetz, Jerome. 1993. The Sin of Science: Ignorance of Ignorance. *Science Communication* 15 (2): 157-165.
- Wehling, Peter. 2006. *Im Schatten des Wissens? Perspektiven der Soziologie des Nichtwissens*. Konstanz: UVK.
- Wilholt, Torsten. 2009. Bias and Values in Scientific Research. *Studies in History and Philosophy of Science* 40 (1): 92-101.
- Williams, Bernard. 2008. Tertullian's Paradox. In: idem, *Philosophy as a Humanistic Discipline*, 3-21. Princeton: Princeton University Press.