APhEx 27, 2023 (ed. Vera Tripodi)
Ricevuto il: 04/04/2023
Accettato il: 22/09/2023
Redattore: Andrea Borghini & Matteo Morganti

APhEx
PORTALE ITALIANO DI FILOSOFIA ANALITICA
GIORNALE DI FILOSOFIA
NETWORK
N°27, 2023

O P E N P R O B L E M S

Can Ignorance Be An Epistemic Virtue?

Marco J. Nathan

Ignorance has traditionally been perceived negatively, as a hurdle. Our knowledge inevitably has gaps and blunders, which are gradually filled in as more is discovered. Over the last few decades, agnotologists and epistemologists of ignorance have challenged this standard story emphasizing that ignorance is not merely the culpable absence of information. Ignorance also has a "virtuous" component, when a specific inquiry ought to be left untouched. Yet, can ignorance become an epistemic virtue in science, a goal in and of itself? This essay shifts the spotlight to a more constructive side of ignorance and its philosophical implications. I begin by distinguishing three kinds of ignorance. Next, I illustrate the claim that ignorance can play a strongly productive role by connecting it to debates over reductionism and the status of black boxes. Finally, I offer some suggestions for incorporating the pursuit of ignorance into our textbooks and other pedagogical tools.

Periodico On-line / ISSN 2036-9972

APhEx
PORTALE ITALIANO DI FILOSOFIA ANALITICA
GIORNALE DI FILOSOFIA NETWORK

1. Introduction

There is a dusty image, mainstream well into the twentieth century, which remains popular, at least in some circles. This is the figurative depiction of science as a slow, painstaking accumulation of knowledge. The chief goal of the scientific enterprise, from this hallowed stance, is to provide a reasonably accurate and complete description of the universe, or some relevant portion of it. In this sense, the pursuit of science is akin to the erection of a wall or the tiling of a mosaic. The building blocks of science are facts, basic truths about our universe. New research adds on to such grandiose endeavor.

Over the last few decades, scientists, philosophers, historians, and scholars from many other fields have vocally denounced the misleading nature of this characterization of science as erecting a wall of facts. To be sure, truth, knowledge, and objectivity remain paramount goals of the scientific enterprise. Still, in the wake of the seminal work of Kuhn (1962), Lakatos (1970), and Feyerabend (1975), it has gradually become clearer that this is merely one side of the story. The remainder involves what we do not know, grasp only partially, or got wrong. In short, what is missing from the brick-by-brick model of scientific progress is the productive role of ignorance. *Pace Pink Floyd*, all in all it is *not* just another brick in the wall.

The inevitable presence of ignorance, within both scientific practice and everyday settings, is neither shocking nor controversial. Generations of philosophers have explored the nature and limits of human knowledge. Yet, with the notable exception of work on skepticism (Smithson, 1989, Rescher, 2009), epistemologists and philosophers of science traditionally focused on the knowledge side. Limits—ignorance—remain strikingly neglected. But what exactly is ignorance? Why has it been slighted for so long? What roles, if any, does it play in science? Let's address these questions in turn.

For starters, what is ignorance? Providing a clear and widely-accepted definition is no trivial endeavor. One can distinguish kinds of ignorance, including *propositional*, *objective*, and *practical* varieties (Nottelmann, 2016). Furthermore, restricting our attention to ignorance of propositional ilk, it remains widely debated whether ignorance is best construed as lack of knowledge or lack of true belief, or whether further conditions must be added (El Kassar 2018, Pritchard 2021, Peels 2023), and how to formalize it (Fano and Graziani 2021, Aldini, Graziani, and Tagliaferri 2023, Bonzio, Fano, Graziani, and Pra Baldi 2023). In what follows, I steer clear of nuances and technicalities, and merely assume that ignorance can be characterized as absence of knowledge or true belief *of some sorts*.

Next, why has epistemology neglected ignorance for so long? Peels (2023) suggests three reasons. First, some philosophers define epistemology simply as the study of knowledge—a narrow conception of the discipline that is becoming increasingly unpopular. Second, ignorance has often been considered a primitive, as opposed to a concept in need of direct elucidation. From this standpoint, an appropriate analysis of knowledge will *ipso facto* shed light on ignorance too. Third, ignorance has traditionally been widely considered negatively, as something to be avoided or as a hurdle to be overcome. As we shall see, the assumptions that ignorance is either indefinable or the flip side of knowledge are now frequently questioned. The crumbling of the "wall" metaphor helped expose these as crass mistakes. There is more to ignorance than the mere absence of understanding. Ignorance is a rich, nuanced topic that deserves to be in the limelight, addressed on its own terms.

Finally, moving on to our third query—the role of ignorance in science—over the last few decades, the landscape has started to shift at last, dragging ignorance under the microscope. This move was largely due to the emergence of *ignorance studies*, an interdisciplinary effort cross-cutting philosophy, psychology, sociology, cognitive science, anthropology, and economics (Gross and McGoey 2023). Two broad areas of ignorance studies have recently thrived (Arfini 2019, 2021). First, *agnotology* refers to the study of culturally-induced ignorance or doubt, especially ignorance created or maintained by the publication of misleading or inaccurate scientific data (Proctor and Schiebinger 2008). Second, *epistemology of ignorance* studies how the generation of knowledge and ignorance is filtered through the production and distribution of empirical data. Two key subfields of the epistemology of ignorance are *feminist philosophy* (Frye 1983, Haraway 1988, Tuana 2004, 2006, Tanesini 2015) and the *philosophy of race* (Sullivan and Tuana 2007).

The present essay aims to contribute to the growing ignorance studies literature on the relationship between ignorance and scientific inquiry by asking whether ignorance can be a bona fide epistemic goal (Pritchard 2016, Peels and Pritchard 2021). It has been widely observed that ignorance cannot be merely dismissed as something negative, as mere absence of knowledge. Ignorance can be blissful, even virtuous (Townley 2011, Kourany 2020). Nevertheless, most examples of virtuous ignorance involve cases where ignorance is desirable as a consequence of adopting certain values, which are not exquisitely epistemic. Science is well-advised to steer clear of issues when they are too costly or dangerous, or because they conflict with other core values, such as privacy. But can ignorance ever be epistemically desirable in and of itself?

Without denying that science is value-laden at multiple levels, can ignorance contribute specifically to our knowledge?

Here is the masterplan for the ensuing pages. §2. kicks off by distinguishing three kinds of ignorance: *disruptive*, *weakly productive*, and *strongly productive*. §3. points the spotlight on the controversial notion of strongly productive ignorance by applying it to a classic debate in the philosophy of science—reductionism—and to a more recent discussion of the construction of black boxes. §4. wraps things up by exploring how the idea of productive ignorance can be incorporated into textbook presentations of science.

2. Three Kinds of Ignorance

A promising place to begin an investigation of the role of ignorance within scientific theory and practice is a book by neuroscientist Stuart Firestein (2012), aptly titled *Ignorance: How It Drives Science*. At the onset, Firestein maintains that "Science (...) produces ignorance, possibly at a faster rate than it produces knowledge" (p. 28). At first blush, this remark may appear like a pessimistic strand of traditional skepticism, denoting the incapacity of science to deliver on its epistemic promises. As the view gradually unfolds, it becomes clear that such initial assessment would be widely off the mark.

> We now have an important insight. It is that the problem of the unknowable, even the really unknowable, may not be a serious obstacle. The unknowable may itself become a fact. It can serve as a portal to deeper understanding. Most important, it certainly has not interfered with the production of ignorance and therefore of the scientific program. Rather, the very notions of incompleteness or uncertainty should be taken as the herald of science (Firestein 2012, p. 44).

But what does it mean to assert that ignorance can become a "portal to deeper understanding"? Firestein's choice of the term "ignorance" is deliberatively provocative. It will be helpful to provide some preliminary clarification. As mentioned at the outset, I shall treat ignorance tout court as the epistemic state of an agent, or community thereof, who lacks access to some piece or information, without drawing further subtle distinctions. With this propaedeutic characterization under our belts, we can clarify Firestein's point by introducing a distinction among three kinds of ignorance.

The first strand of ignorance is what I call *disruptive ignorance* (DI) because it has no constructive role whatsoever to play, neither in science nor elsewhere. DI may manifest itself in several guises. In some instances, it involves a form

APhEx
PORTALE ITALIANO DI FILOSOFIA ANALITICA
GIORNALE DI FILOSOFIA
NETWORK

of willful stupidity, a callous indifference to corroborated facts, a travesty of sound reasoning, a stubborn grasping onto uninformed opinions disregarding any ideas, perspectives, or data that may shake our credal states, leading us to revision. To illustrate, consider conspiracy theories floating around the web, such as the claim that planet Earth is flat, political "fake news," and various other forms of quackery and superstition. These are not merely harmless albeit discredited beliefs, at least when literally construed ("the sun sets in the west"), or strictly false statements that may still be useful ("trees maximize exposure to sunlight"). In centuries past, it may have been excusable, perhaps even rational, to hold on to the belief that the Earth is flat. But it is no longer so. It is an epistemically blameworthy credal state where an agent fails to grasp a piece of information but should know better and has no rationale for not doing so.

Such unawareness, lack of enlightenment and informativeness, is strikingly and scarily widespread, including among citizens who occupy elected offices and other positions of power. Beliefs in ghouls, black magic, and other urban legends are rampant, even among educated people. To give a sense of the scale, recent research suggests that roughly three quarters of Americans hold at least one paranormal belief. Even more worrisome, such credences show no sign of waning, with the younger generations beings just as credulous than their elders—if not more so, as in the case of astronomy. This "pandemic of poppycock," as Pinker (2021) aptly dubs it, calls for explanation. Nevertheless, what is intriguing is its infectious nature, triggering its ubiquitousness. The epistemic state in question is neither especially interesting nor controversial, as we should all agree on its toxicity.

Not all DI takes such crass form. Over the last few years, agnotologists have identified and discussed subtler and more dangerous varieties of DI. Ignorance and doubt mongering, for instance, can be actively constructed and exploited to promote corporate interests, from the fossil fuel industry (Oreskes and Conway 2010), to tobacco companies (Proctor 2011), to the consumption of junk food (Moss 2013). More recently, it has become evident that ignorance may also be a passive construction, a by-product of the way a research program is selected and conducted (Kourany and Carrier 2020). These are all cases of disruptive ignorance, which is is an obstacle to the advancement of knowledge, something that all rational agents ought to shun.

At the same time, these is also a less disparaging sense of ignorance, conceived as a guiltless absence of knowledge, understanding, insight, or clarity. Firestein describes it as "knowledgeable ignorance, perceptive ignorance, insightful ignorance" (2012, pp. 6-7). I call it *productive ignorance* (PI) since, unlike its disruptive counterpart, it contributes to the enterprise of knowledge. Follow-up: what is PI and what is its role within scientific research? Firestein characterizes this epistemic state not as individual lack of information, but as a communal gap in competence. Data may not yet be available, as with the long-term impact of social media on young extensive users, or the evidence at hand may not support any coherent explanation or robust prediction. Its usefulness is that "It leads us to frame better questions, the first step to getting better answers. It is the most important resource we scientists have, and using it correctly is the most important thing a scientist does" (Firestein 2012, 7). There is much to agree with Firestein's separation between—what I labelled—DI vs. PI. At the same time, we must draw a finer distinction between PI of two different kinds. Allow me to elaborate.

On the one hand, there are situations where we cannot offer an adequate solution to a problem partly or wholly because of some epistemic gap. The missing piece of the puzzle may be a fact, a technological gear, insufficient computing power or funds, or something altogether different. While the conundrum cannot presently be addressed, what we don't yet know may provide steppingstones toward acquiring the answers we seek, or point to the direction where to focus, discover, invent, or fundraise. I call this *weakly productive ignorance* (WPI) because, as Firestein notes, when properly harnessed it may become a portal to success, paving the way toward further discoveries. Contrary to DI, WPI is *cognitively blameless*. (For a discussion of the culpability of ignorance, see essays collected in Peels and Blaauw 2016 and Peels 2017). It is not that agents refuse to face the facts or to seek relevant information. The data just isn't there, as of yet. However, such ignorance is *weakly* productive because it is at best a temporary scaffolding to be replaced with knowledge as more epistemic gaps are filled in.

A couple examples should help drive the point home. Newton was acutely aware that his principle of universal gravitation flew right in the face of the mechanistic tradition pioneered, among others, by Galileo and Descartes, and that he had brought to perfection. Nevertheless, in a bout of candor, Newton professed ignorance. *Hypotheses non fingo*, "I feign no hypotheses," he famously proclaimed in the General Scholium appended to the second 1713 edition of *Principia*. Seeking to fill this void has driven both physical and philosophical agendas, as witnessed by the work of Laplace, Kant, and Einstein: giants standing on the shoulders of giants. Once an explanation for universal gravitation became available, there was no longer any reason to feign hypotheses. Ignorance led to groundbreaking knowledge. Like a scaffolding, it had fulfilled its task. The discovery of non-Euclidian geometries, sparked by attempts to show the necessity of Euclid's parallel postulate, provides an intriguing variant of this story. Here the productive role is played not by the absence of knowledge, but by a false assumption.

On the other hand, there is a different strand of ignorance, which I brand as *strongly productive* (SPI). Just like WPI, SPI is epistemically blameless, psychologically harmless, and potentially conducive to various kinds of breakthroughs. What distinguishes SPI from its weakly-productive counterpart is that, in the "strong" case, the state of ignorance is epistemically preferable to the corresponding state of knowledge. WPI is still an instance where ignorance is a hurdle to be overcome on the path to understanding. SPI, in contrast, is an epistemic goal, an objective in and of itself.

Neither DI nor WPI are especially controversial. As noted, conspiracy theories, fake news, and other false beliefs are a well-documented plague in our contemporary society. Agnotology has revealed how ignorance can be socially constructed, actively or passively. At the same time, agnotologists have also stressed that ignorance can be "virtuous," when lack of knowledge is deliberately accepted as a consequence of certain values (Kourany 2020). This may happen for various reasons. In some instances, the pursuit of knowledge may be too risky or dangerous. Governments classify military, political, pathogenic, and other kinds of information that could wreak havoc if it were to fall into the wrong hands. Some of this research should arguably not be conducted at all, as people in power may well be tempted to use nuclear, bacteriological, and other weapons of mass destruction in extreme situations. A different set of cases involves knowledge that could only be acquired via morally dubious, or even straight-up wrong means, such as human or animal torture, human cloning, or inhumane forms of laboratory testing. Some knowledge may be uninteresting, prohibitively pricey in absolute or relative terms, or the cost of acquisition may be greater than the benefit (Caplan 2001, Somin 2015). Research may also have problematic implications. To wit, searching for sex-based or racially-driven differences in "intelligence" or other cognitive features may trigger further harm and marginalization. Patients may wish not to discover their susceptibility to the ravaging effects of incurable diseases, and genetic testing may also conflict with other values, such as our desire for privacy. The location of valuable resources is often best kept secret to avoid looting, poaching, stealing, and other activities that could tempt people with less-than-noble intentions.

In short, there is no lack of research projects unworthy of pursuit. Questions we do not want answered should not be investigated (Kitcher 2001, 2011). I classify these scenarios as cases of WPI for the simple reason that they are all instances where knowledge would improve our epistemic states. There are valid ethical, socio-political, prudential, and other normative concerns, which may well make it best not to obtain the knowledge in question. But, from a purely epistemic perspective, such knowledge would put us in a better cognitive situation than the corresponding state of ignorance, in the sense that we would know more about a specific phenomenon, whether or not such knowledge will do harm or good. SPI feels like a different story, bringing a brand new challenge to the foreground. Set aside the important issue of norms and values. Are there cases where we ought not to replace ignorance with knowledge not for ethical, moral, or cautionary reasons, but on epistemic grounds alone? In other words, are there scenarios where we may be in a better cognitive position when we know less than more? Recent work has revealed that ignorance may be instrumental for pedagogy and other educational purposes (Pritchard 2016, Peels and Pritchard 2021, Peels 2023) or from a cognitive perspective, in abductive and other forms of reasoning (Arfini 2019). But what about actual scientific practice? From an exquisitely epistemic standpoint, is ignorance ever preferable to knowledge?

3. Strongly Productive Ignorance in Focus

The previous section distinguished three types of ignorance. Both *disruptive* (DI) and *weakly productive ignorance* (WPI) call for further scrutiny. The former state is epistemically blameworthy; the latter is blameless. But neither is especially controversial. People sometimes unduly hold on to irrational beliefs and gaps in knowledge often steer us in the right direction. “Virtuous” ignorance involving thorny research better set to the side because of inherent danger or conflicts with other core values, is similarly “weakly productive.” The more quixotic case involves *strongly productive ignorance* (SPI), which turns uncertainty into an epistemic goal in and of itself. Are there instances where a state of oblivion, guiltless as it may be, is actually preferable to the corresponding state of knowledge? The answer is hardly obvious. This section sets out to explore it, summarizing an argument which I originally articulated in greater detail in previous work (Nathan 2021).

The nature, status, and existence of SPI mirrors the *reductionism* debate that has dominated the philosophy of science since the mid-twentieth century. During the heyday of positivism, reduction was conceived as a logically deductive relation between theories characterized as interpreted axiomatic systems. This "classical" model of reduction, epitomized by Nagel's *The Structure of Science* (1961), has been eroded by powerful objections, from the multiple-realizability of higher-level kinds, to the lack of appropriate laws and bridge principles, to the shortcomings of "syntactic" conceptions of theories. In the aftermath of the classic derivational model's demise, the reductionism debate has morphed into an epistemic feud over *explanation*.

Neo-reductionists maintain that lower-level descriptions invariably enrich our understanding of systems, rendering higher-level macro-depictions technically superfluous. Hence, all of science is reducible to its most fundamental level because—barring considerations of computational tractability and convenience—micro-depictions provide deeper, more comprehensive, and more powerful explanations. Antireductionists beg to disagree. While structural details are often crucial for addressing macro-explananda, this downward trajectory neither extends all the way to the foundations nor does it apply across the board. In some cases, abstract and idealized models are not made more powerful by filling in all the details. There is a threshold past which additional information does not enhance higher-level explananda. If anything, unnecessary details muddy the waters. This, in a nutshell, is the *autonomy* of the special species, official trademark of antireductionism.

It should now be evident how the reductionism dispute, thus epistemically construed, can be couched in terms of SPI. Reductionists essentially contend that, practicality aside, uncertainty is never preferable to the corresponding state of knowledge. SPI is a chimera. Antireductionists rejoin that, once we've optimized our "explanatory bang for the buck," additional details become detrimental. Ignorance can be bliss, especially when strongly productive. If this is along the right track, the existence of SPI boils down to the status of epistemic reduction. So, should we be reductionists or antireductionists about science? Unfortunately, over half a century of heated debate later, the dispute remains as open as ever, with no sign of resolution in sight. What's worse, the two sides appear to be entrenched, with modest reductionism and sophisticated antireductionism overlapping, and the disagreement increasingly appearing more semantic and less substantial than initially supposed. In short, the (anti)reductionism research programs have reached a stalemate, becoming regressive. It would not be wise to build our analysis of ignorance on such swampy foundations. We deserve better.

The first step toward a resolution is a diagnosis of where the current feud has turned sour. Reductionists and antireductionists share the presupposition that explanations pitched at different levels provide competing accounts of the same target. This is a mistake; they do not. Explanations—scientific and commonsensical alike—are context-relative. Hence, higher-level explanations can be both reducible to *and* autonomous from lower levels. Reduction is a relational issue, concerning the reframing of a macro-explanandum as a micro-inquiry. Autonomy, in contrast, is an intrinsic feature, pertaining to whether a model provides the resources to adequately address its explananda. Thus framed, these two properties become perfectly compatible.

As a simple illustration, consider Putnam's (1975) famous square-peg-round-hole scenario. A square peg will go through a square hole slightly larger than its cross-section. Yet, that same square peg will not go through a round hole whose diameter is the length of the square's edge or smaller. Are structural physical details relevant to the case at hand? If so, do they increase explanatory power? Unless we clarify what features of the behavior of the system we are trying to explain, such questions are meaningless. If the explanandum is *that* the peg does not pass through the round hole, then the physical details are not just unnecessary, but detrimental and should be black-boxed. In contrast, if the objective is *why* such state occurs, then looking at the physical structure of the system, down to its molecular constituents, becomes relevant. In this latter case, a black box ain't gonna cut it. It has to be "opened" and replaced by a proper mechanistic account.

Followup: what is a black box? Simply put, black-boxing a phenomenon involves isolating some of its core features in a way that can be assumed without further explanation, justification, or detailed structural description. More precisely, a black box can be defined as a placeholder in a causal explanation represented in a model. These placeholders can either stand in for behaviors in need of explanation (*frames*) or for whatever produces the behaviors in question (*difference-makers*). Yet this cannot be the entire story. Not all black boxes are the same. Some work well; others do not. We thus need to distinguish the true black boxes from the red herrings.

So much for square pegs, round holes, and other toy examples. Does this story apply to real scientific illustrations? Yes. Consider Darwin and Mendel's black-boxing of the mechanisms of inheritance, the black-boxing of mental states occurring both in behavioristic psychology and neoclassical economics, as well as a case buried in the history of science: accounts of combustion by phlogiston theories, eventually discarded and replaced by atomic chemistry. In all these instances, we have explanations that are perfectly autonomous, in the sense that they can be—and have been—assessed and verified without shifting down to lower, more fundamental epistemic levels. At the same time, uncovering the underlying mechanisms opens new avenues of research. In short, the reductionism vs. antireductionism dichotomy misses the mark. The history of black boxes is a dual tale of reduction *and* autonomy. The paramount issue, from our current viewpoint, is whether the prominence of black boxes can address the legitimacy of SPI within scientific theory and practice. The answer is not quite clean-cut.

Simply put, there are two ways to conceive of and construct black boxes. From a reductionist standpoint, black boxes can be viewed as disposable scaffoldings, epistemic crutches to be archived when no longer needed. Science, in the long run, aims at the complete elimination of all black boxes. On an alternative antireductionist reading, black boxes can be ascribed a more permanent status: not scaffoldings but frescoes adorning a wall. From this perspective, our best scientific models, our most fundamental physical descriptions of the universe included, involve various forms of abstraction and idealization. This suggests that the role of black boxes can be reconciled with either reductive or antireductionist philosophical inclinations.

In the end, the reconciliation of autonomy and reduction underlying black boxes reveals that bona fide SPI does not undermine the value of additional, lower-level mechanistic details. Ignorance actually has a constructive role to play. It promotes autonomy, as antireductionists point out. At the same time, opening a black box always constitutes progress, as rightly emphasized by reductionists. The key remark is that this operation transforms the old inquiry into a new problem. As long as we view explanation as framework-dependent, cases where ignorance is preferable to the corresponding state of knowledge do not entail that, relative to other epistemic aims, these same details may not turn out to be significant, indeed crucial.

4. Black Boxes & Red Herrings: Closing Remarks

Time to take stock. Our discussion kicked off with the observation that ignorance has traditionally flown under the philosophical radar. This trend has finally begun to shift. Recent research has revealed that ignorance is a rich and nuanced topic that deserves to be addressed on its own terms. Piggybacking on Firestein's insights, we distinguished three kinds of ignorance. First, *disruptive ignorance* (DI) manifests itself as a callous indifference to well-known facts that, as pinpointed by agnotologists, may often be actively or passively constructed and maintained. Second, *weakly productive ignorance* (WPI) is either a way of opening new inquiries, leading to new or better knowledge, or an attempt to set aside questions that, for various reasons, are better off left alone—“virtuous ignorance,” as it is sometimes called. Third, *strongly productive ignorance* (SPI) occurs when ignorance becomes a goal in an of itself, an epistemic stance preferable to the corresponding state of knowledge. But is SPI a legitimate category? Is ignorance ever to be privileged over knowledge? In the epistemology of ignorance, this issue has become a pedagogical one, regarding usefulness of ignorance in education. In traditional philosophy of science, these questions have been extensively discussed in the guise of the debate over reductionism. Influential as it has been, this is a dispute that has reached a stalemate and turned regressive. By focusing on black boxes, we saw that a theory can be *both* autonomous *and* reductive, as long as we are willing to view explanations as relativized to their own epistemic contexts. In the end, there are legitimate cases of bona fide SPI, but this doesn’t undermine the value of our pursuit of knowledge, not merely for teaching, but for actual scientific research.

It should be evident that our current discussion barely scratches the surface of much deeper waters. From the relationship between forms of knowledge and their counterparts, to the uses and abuses of ignorance in science, to a reconciliation of autonomy and reduction, much remains to be done. Open problems are plentiful, making this an exciting area of research. This short essay can be conceived as an intellectual *amuse-bouche* that will hopefully wet your appetite for more. Yet, before wrapping things up, let me briefly bring up two matters that, in my opinion, are especially pressing.

First, as several authors have emphasized, ignorance is a key component of science and scientific education. Then why is its role still not explicitly incorporated into the academic curriculum? Why is ignorance treated at best in the weakly productive sense, as steppingstone toward the acquisition of more knowledge or as confounding variable, as opposed to active ingredient of our research? These preconceptions, I suggest, are a legacy of the old dusty metaphor that kindled our discussion at the outset of this essay. Allow me to clarify.

The classic image of science as a progressive accumulation of facts—now disavowed by most insiders—remains popular among the general public. Journalists, politicians, entrepreneurs, and other non-specialists typically presuppose, more or less consciously, the discredited vision of science as erecting a wall of knowledge or tiling the mosaic of the universe. Furthermore, the source of this misconception is evident: the mismatch comes from the standard packaging of the enterprise. Textbooks and courses introduce science as a bunch of laws, theories, facts, and a handful or other cut-and-dried notions to be internalized and memorized uncritically over time. This resilient, pernicious stereotype trickles down from schools and universities to television shows, newspapers, magazines, and other mainstream channels that perpetuate it among the general educated public. Only the small portion of students who attend graduate school and pursue research careers is exposed to the true face of science. Real practice at the bench is quite different, and way more exciting, than the ossified caricatures presented to outsiders in journals, books, and conferences. But, note, the textbooks by which students learn science are written by those same specialists who, when pressed on the issue, adamantly eschew the dusty brick-by-brick metaphor. Why are experts perpetuating an image they themselves disavow? The answer, as far as I can tell, is that there is no viable alternative model available. The wall of knowledge has overstayed its welcome, but we haven't found any adequate replacement. We need more fruitful ways to expose the new generations to the fascinating world of science. Whether the answer lies in black boxes, discussions of ignorance, failure, and uncertainty, or elsewhere, we must do a better job popularizing our remarkable body of work.

This leads to a second and related remark. I suspect that many people are wary of preaching and teaching ignorance, at least in part, because ignorance can be a double-edged sword. Ignorance has both a productive and a disruptive face. Ignorance may be a pathway to more, even better understanding. But it can likewise be a callous indifference to the facts or a smokescreen. In the wrong hands, black boxes can justify shortcuts purporting to cut corners or avoid the hard work. In a world where science is already facing a crisis of public trust, do we need to bring further attention to its gaps and limitation? My answer is a resounding: Yes! We cannot fight the spread of misinformation by fueling more misconceptions. The history of science is one of staggering success and we have many reasons to be proud. Ignorance, both disruptive and productive, is part of that same story. Hence, a core issue becomes how to incorporate the productive side of ignorance, while leaving its evil twin at the door. We need to separate true black boxes from red herrings. It won't be easy. Are we up to the task?

In conclusion, ignorance can play a key role within scientific theory and practice. Coming to terms with it is paramount for the presentation, teaching, and pursuit of science: its successes as well as its limitations. Paraphrasing Twister Sister, what you don't know sure can hurt you, and what you can't see makes you scream. But it may also be a blessing in disguise.

Acknowledgements

The author is grateful to Bill Anderson, Andrea Borghini, Pierluigi Graziani, as well as anonymous editors and reviewers, for constructive comments on various versions of this essay. An early draft was presented on February 23, 2023 in Denver, during a symposium on *Black Boxes* at the Central Division Meeting of the American Philosophical Association. Audience and fellow panelists provided helpful feedback.

REFERENCES

Aldini, Alessandro, Pierluigi Graziani, and Mirko Tagliaferri (2023). "A Hierarchical Characterization of Ignorance in Epistemic Logic." *Argumenta 8*(2), 373–86.

Arfini, Selene (2019). *Ignorant Cognition: A Philosophical Investigation of the Cognitive Features of Not Knowing*. Cham: Springer.

Arfini, Selene (2021). "Ignoranza." *Aphex 23*.

Bonzio, Stefano, Vincenzo Fano, Pierluigi Graziani, and Michele Pra Baldi (2023). "A Logical Modeling of Severe Ignorance." *Journal of Philosophical Logic 52*, 1053–80.

Caplan, Bryan (2001). "Rational Ignorance versus Rational Irrationality." *Kyklos 54*(1), 3–26.

El Kassar, Nadja (2018). "What Ignorance Really Is: Examining the Foundations of Epistemology of Ignorance." *Social Epistemology 32*(5), 300–10.

Fano, Vincenzo and Pierluigi Graziani (2021). "A Working Hypothesis for the Logic of Radical Ignorance." *Synthese 199*, 601–16.

Feyerabend, Paul K. ([1993] 1975). *Against Method* (3rd ed.). London and New York: Verso.

Firestein, Stuart (2012). *Ignorance. How it Drives Science*. New York: Oxford University Press.

Frye, Marilyn (1983). *The Politics of Reality: Essays in Feminist Theory*. Freedom: Crossing Press.

Gross, Matthias and Linsey McGoey (eds.) (2023). *Routledge International Handbook of Ignorance Studies* (2nd ed.). Abingdon and New York: Routledge.

Haraway, Donna J. (1988). "Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective." *Feminist Studies 14*(3), 575–99.

Kitcher, Philip (2001). *Science, Truth, and Democracy*. New York: Oxford University Press.

Kitcher, Philip (2011). *Science in a Democratic Society*. New York: Prometheus Books.

Kourany, Janet A. (2020). "Might Scientific Ignorance Be Virtuous? The Case of Cognitive Differences Research." In J. A. Kourany and M. Carrier (2020), pp. 123–44.

Kourany, Janet A. and Martin Carrier (Eds.) (2020). *Science and the Production of Ignorance: When the Quest for Knowledge Is Thwarted*. Cambridge and London: MIT Press.

Kuhn, Thomas S. (1962). *The Structure of Scientific Revolutions* (1st ed.). Chicago: University of Chicago Press.

Lakatos, Imre (1970). "Falsification and the Methodology of Scientific Research Programmes." In I. Lakatos and A. Musgrave (Eds.), *Criticism and the Growth of Knowledge*, pp. 91–195. Cambridge: Cambridge University Press.

Moss, Michael (2013). *Salt, Sugar, Fat: How the Food Giants Hooked Us*. New York: Random House.

Nagel, Ernest (1961). *The Structure of Science*. New York: Harcourt Brace.

Nathan, Marco J. (2021). *Black Boxes: How Science Turns Ignorance into Knowledge*. New York: Oxford University Press.

Nottelmann, Nikolaj (2016). "The Varieties of Ignorance." In R. Peels and M. Blaauw (2016), pp. 33–56.

Oreskes, Naomi and Erik M. Conway (2010). *Merchants of Doubt: How a Handful of Scientists Obscured the Truth on Issues from Tobacco Smoke to Global Warming*. New York: Bloomsbury.

Peels, R. (Ed.) (2017). *Perspectives on Ignorance from Moral and Social Philosophy*. New York: Routledge.

Peels, R. (2023). *Ignorance: A Philosophical Study*. New York: Oxford University Press.

Peels, Rik and Martijn Blaauw (Eds.) (2016). *The Epistemic Dimensions of Ignorance*. Cambridge: Cambridge University Press.

Peels, Rik and Duncan Pritchard (2021). "Educating for Ignorance." *Synthese 198*, 7949–63.

Pinker, Steven. (2021). *Rationality: What Is Is, Why Is Seems Scarce, Why It Matters*. New York: Viking.

Pritchard, Duncan. (2016). "Ignorance and Epistemic Value." In R. Peels and M. Blaauw (2016.), pp. 132–43.

Pritchard, Duncan (2021). "Ignorance and Inquiry." *American Philosophical Quarterly 58*, 111–23.

Proctor, Robert N. (2011). *Golden Holocaust: Origins of the Cigarette Catastrophe and the Case for Abolition*. Berkeley: University of California Press.

Proctor, Robert N. and Londa Schiebinger (Eds.) (2008). *Agnotology: The Making and Unmaking of Ignorance*. Stanford: Stanford University Press.

Putnam, Hilary (1975). "Philosophy and Our Mental Life." In *Mind, Language, and Reality*, pp. 291–303. New York: Cambridge University Press.

Rescher, Nicholas (2009). *Ignorance: On the Wider Implications of Deficient Knowledge*. Pittsburgh: University of Pittsburgh Press.

Smithson, Michael (1989). *Ignorance and Uncertainty: Emerging Paradigms*. New York: Springer.

Somin, Ilya. ([2023] 2015). "Rational Ignorance." In M. Gross and L. McGoey (2023), pp. 313–22.

Sullivan, Shannon and Nancy Tuana (Eds.) (2007). *Race and Epistemologies of Ignorance*. New York: SUNY Press.

Tanesini, Alessandra (2015). "Epistemologie e Filosofie Femministe della Scienza." *Aphex: Portale Italiano di Filosofia Analitica 11*.

Townley, Cynthia (2011). *A Defense of Ignorance: Its Value for Knowers and Roles in Feminist and Social Epistemologies*. Lanham: Rowman & Littlefield.

Tuana, Nancy (2004). "Coming to Understand: Orgasm and the Epistemology of Ignorance." *Hypatia 19*(1), 194–232.

Tuana, Nancy (2006). "The Speculum of Ignorance: The Women's Health Movement and Epistemologies of Ignorance." *Hypatia 21*(3), 1–19.