TRUTH AND SCIENTIFIC CHANGE

Gila Sher

Forthcoming: *Journal of General Philosophy of Science*

In physics the truth is rarely perfectly clear.
Richard Feynman

Two central questions of this paper are:

1. What lessons does the phenomenon of scientific change teach us about the nature of truth?

2. What light do recent developments in the theory of truth, incorporating these lessons, throw on problems arising from the prevalence of scientific change, specifically, the problem of pessimistic meta-induction?

I. Scientific Change and the Challenge of the Pessimistic Meta-Induction

The phenomenon of scientific change is naturally associated both with an optimistic and with a pessimistic view on the present and future state of scientific knowledge. Recently, these views have been crystalized in the so-called *optimistic* and *pessimistic* meta-inductions. Optimistically, we conclude that humans will continue to make significant progress in their scientific endeavors. Pessimistically, we conclude that human theories will continue to be found incorrect. The pessimistic challenge is commonly traced to Laudan (1981), who formulates it as a challenge to scientific realism. His claim is that it follows inductively from the fact that most past scientific theories were found to be incorrect that our present scientific theories will be found to be incorrect as well; hence realism with respect to our present theories is unwarranted.

Laudan formulates his argument in terms of *truth (falsehood*):

> [W]e have only to look at the history of science to see that theories eventually get **falsified** ...[, and] we should deem current theories **false** because past ones have been shown to be. [Laudan 1990: 39][1]

Critics, too, describe the pessimistic-induction argument in terms of truth (falsehood):

> The history of science is full of theories which at different times and for long periods had been empirically successful, and yet were shown to be **false** in the deep-structure claims they made about the world. ... Therefore, by a simple (meta-) induction on scientific theories, our current successful theories are likely to be **false** (or, at any rate, are more likely to be **false** than **true**). [Psillos 1999: 101]

> [P]essimistic meta-induction [proceeds] from the many past successful-but-**false** theories to the likelihood that our best current theories are likewise **false**. [Doppelt 2007: 96]

And neutral encyclopedia articles also formulate the argument in those terms:

> [A]s many past theories in science have turned out to be non-referring, there is all reason to expect that even the future theories fail to refer – and thus also fail to be approximately **true** or **truthlike**. [Niiniluoto 2011: 20]

Quite a few philosophers of science (e.g., Kitcher 1993, Psillos 1996, 1999, Doppelt 2007, and Roush 2010) have argued against the pessimistic meta-induction, establishing their arguments on a variety of grounds. Here I will approach the pessimistic meta-induction from a new perspective: *the nature of truth*.

II. The Nature of Truth[2]

To focus on those aspects of truth that are pertinent to science, I will start with a few general, common-sense observations on what may be called *the basic human cognitive/epistemic*

---

[1] I use bold font in citations to indicate my emphases.

[2] This section focuses on, and further develops, several aspects of a general account of truth (not directed toward science) in Sher (2004, 2016).

situation. Briefly expressed, these observations are:

1. We live in a world of which we are a part.

2. We (our species) developed in such a way that one of our significant aspirations is to obtain objective knowledge of our world, and not just practical knowledge, but also theoretical knowledge.

3. Obtaining objective theoretical knowledge of the world, however, is not a simple matter for humans. Our cognitive resources are limited relative to the complexity of the world or, to put it otherwise, the world is complex relative to our cognitive capacities. And while some facets of the world are easy for us to access, others are difficult, and sometimes very difficult, to reach.

4. Nevertheless, we, humans, are ambitious creatures, and we aspire to know the world in its full complexity.

5. What makes these aspirations partly, yet significantly, realizable is that, our limitations notwithstanding, we have a variety of fairly intricate cognitive resources, sensory as well as intellectual, innate as well as learned. Among these is a capacity to play an active role both in discovery and justification: designing experiments, developing research programs, calculating, inferring, exercising our imagination, and so on.

The basic human cognitive/epistemic situation gives rise to two overarching principles of human knowledge: *epistemic friction* and *epistemic freedom*. Knowledge is constrained by the world (requires "friction" with the world). But to attain knowledge, humans must put their freedom to work. It is important to note that friction and freedom come together in knowledge. For example, *norms* in general are products of freedom and sources of constraint. In particular, the epistemic norms of correctness, evidence, and justification are freely generated and imposed on our theories by us, not by the world. But they constrain us vis-a-vis our target, the world. This has direct ramification for one of the issues at the center of the debate on the pessimistic meta-induction: *realism*.

1. *Basic Realism.* Given that knowledge of the world is a central human aspiration, the body of theories developed in pursuit of this aspiration must be directed toward the world. That is, human knowledge has a *realistic* orientation. But what kind of realism is right for creatures like us? The principles of epistemic friction and freedom point to a type of realism that I will call

"basic realism", where "basic" is both a magnifier and a qualifier. *Basic realism* regards the goal of "getting the world right" as critical to *all* disciplines, from experiential science to logic and mathematics, hence requires all disciplines to be subject to significant *veridicality* demands. Logic, for example, must aim at judgments of logical validity that are factually correct. It is not enough that a logical inference *appear* to transmit truth from premises to conclusion with a strong modal force; it should *actually* do that.

But this does not mean that the relation between human knowledge and its targets in the world must take a single form, exhibit a single pattern. Given our cognitive limitations on the one hand and the complexity of the world on the other, we cannot hope to reach all our targets in the world in the same way. And given the diversity of the cognitive resources we do have (e.g., intellect and imagination alongside sensory perception) together with our epistemic freedom, it is possible for us to forge diverse cognitive routes to the world utilizing our diverse resources. Basic realism allows considerable flexibility in connecting our theories to the world, taking full advantage of our freedom, the diversity of our resources, and the incremental growth of these resources.

2. *The No-Miracle Argument and the Realistic Aspirations Argument.* The most widely discussed argument for scientific realism today is the *no-miracle* argument (often traced to Putnam 1975 and 1981). This argument is based not on humans' aspirations or the goals they set for their theories, but on a feature of the theories themselves, regardless of their aims. It says that the best way to explain the applicability of our theories to the world, or their *success* in the world, is to assume that they are *true* about the world. In the case of scientific theories, it is their empirical success which supports a realistic stance with respect to them.

The no-miracle argument is thus independent of the *realistic-aspirations* argument. In a sense, the no-miracle argument is most effective when it is directed at those who deny that science aims at knowledge of the world. It implies that even if one denies the realistic aspirations of science, one should accept scientific realism based on the applicational (including predictive) success of scientific theories. But this independence means that to refute scientific realism it is not enough to refute the no-miracle argument; one has to refute the realistic-aspirations argument as well.[3]

3. *World-Oriented Holism.* Given the complexities of the basic human epistemic situation, the complementarity of friction and freedom, and the basic realist stance of our search for knowledge, it is reasonable to presume that the most promising methodology of discovery and justification open to us is *holistic*. The key idea is that given our aspirations, limitations, and capacities, we have to utilize all our cognitive resources, in whichever way and order best works for us at each stage. We need to exercise flexibility and ingenuity both in finding new ways of reaching those facets of the world that are difficult for us to reach and in developing methods for evaluating the correctness of our theories at each stage. This requires a different methodology from the traditional foundationalist and coherentist methodologies. We need a methodology that, unlike foundationalism, emphasizes the multiplicity and open-endedness of knowledge-inducing connections between theory and world, and unlike coherentism, does not compromise our goal of a correct understanding of the world. The methodology in question will be both holistic and world-oriented – a holism designed to serve, rather than take the place of, a realistic project of

---

[3] The question arises whether a realism based on the aim of science is weaker than a realism based on the no-miracle argument (see Chakravartty 2011, p. 2, and references there). From the present perspective, the aim-of-science argument is the starting point of scientific realism. It provides a fundamental ground for scientific realism, but it does not uniquely determine the extent of such a realism.

knowledge.[4,5]

4. *A Norm of Truth.* Our basic epistemic aspirations, the need for friction with the world, the realism built into our conception of knowledge, and our world-oriented holistic methodology, all point to the centrality of *correctness* for the enterprise of knowledge. But if correctness lies at the heart of our epistemic enterprise, then it stands to reason that norms of correctness are essential for this enterprise. Truth is first and foremost such a norm. It is our principal norm of correctness. It reflects our deepest epistemic aspirations, guides us in constructing theories of the world, and constrains our theories. To the extent that correctness lies at the heart of the human pursuit of knowledge, so is truth.[6]

But do we really need a special norm of correctness? And are we capable of making good use of such a norm? It is clear that if correctness were automatic for humans, if by merely directing our mental gaze at anything in the world we would automatically have correct cognition of everything about it, then we would not need a norm of correctness. If, on the other hand, we were incapable of correcting any of our cognitions, then a norm of correctness would be of no

---

[4] To prevent misunderstandings, let me note that the holism intended here is not a "total" or "one-unit" holism" (Dummett 1973, Sher 2016). *Total* or *one-unit holism* is the view that the smallest unit of knowledge is our system of knowledge in its entirety. On this view, our entire body of knowledge is a huge atom or blob, lacking inner structure and known only in its entirety or not at all. This view is rejected by our holism on the ground that if knowledge has no inner structure, a step-by-step acquisition, criticism, and justification of knowledge is impossible. (See Dummett 1973 and 1973/81, Glymour 1980, Fodor and Lepore 1992, and Sher 2016.) The holism presented here is, in contrast, *relational* or *structural.*

[5] There are significant similarities between my world-oriented holism and Haack's (1993) foundherentism, but also significant differences. Both methodologies affirm some elements of both foundationalism and coherentism and reject others. But Haack's methodology is limited to empirical science while mine applies to all branches of knowledge, and the two emphasize different aspects of discovery and justification. (See Sher 2016.)

[6] We have emphasized correctness in the context of knowledge. Yet correctness might be important in other contexts as well: e.g., the context of morality. For a discussion of truth in relation to *truthfulness*, see, for example, Williams (2002).

use to us. But being the cognitive creatures we are – fallible on the one hand, capable of correcting ourselves on the other – a standard of correctness is something we need and is extremely valuable to us. This, in my view, is the central reason, or one of the central reasons, that we, humans, have a norm of truth and that the associated concept of truth is so important to us.[7]

But while there is little doubt that humans need some norm(s) of correctness for their theories, one might still ask whether they need a norm of *truth*. Could they not make do with, say, *justification*? The answer to this question is negative. Justification, by itself, is an empty norm. Justification is always relative to a goal or to another norm. If I wish to be entertained I will be justified in doing one thing; if I wish to learn about human suffering, I will be justified in doing something else. How we justify our theories depends on what we care about and what we want them to accomplish. If we look for, say, short, or simple, or esthetically pleasing theories, our justification procedures will be guided by pragmatic or aesthetic norms; if we look for correct theories, they will be guided by a norm of truth. And once we have a norm of truth, we can add the norm of (theoretical or factual) justification.[8]

Two points of clarification are in order here: (a) I have been talking about the *norm* of truth; but in the philosophical literature it is more common to find discussions of the *property,* or *concept,* or *predicate* of truth. Is truth a property, concept, predicate, or a norm? My answer is:

---

[7] This view stands in sharp contrast with the deflationist view that the reason we need a concept of truth is a purely technical linguistic reason. According to this view, the word or concept of truth is a device that enables us to say things we cannot *directly* articulate. Thus, instead of having to make infinitely many assertions, we can make the single assertion that, say, *every sentence of the form 'P or not P' is true*, and instead of saying something that we are not in command of, we can make the assertion that, say, *Einstein's theory of gravitation is true*. See, e.g., Horwich (1990/8). (It should be noted that deflationists do not claim that "is true" is the only device of this kind. An alternative device is substitutional quantification. It follows that truth is, on their view, in principle dispensable.)

[8] Clearly, there is room for other norms as well, e.g. practical, pragmatic, and methodological norms.

All. There is a norm of truth, a property of truth, a concept of truth, and a predicate of truth, and they are all closely related. For example, the property of truth can be viewed as the property of satisfying the norm of truth. From the perspective of knowledge, the normative aspect of truth is the dominant aspect. (b) My conception of truth is *epistemic*. But what this means is very different from what philosophers usually mean by talking about an epistemic conception of truth. Normally, by speaking about an epistemic conception of truth, philosophers mean a conception that reduces truth to some other, distinctly epistemic, norm or concept, e.g., justification. But what I mean is that a norm of truth *distinct from justification* is central both to the human pursuit of knowledge and to our understanding of this pursuit. This leaves us with the task of explaining the positive nature of truth from an epistemic perspective.

5. *The Fundamental Principle of Truth*. We have said that truth is a standard of correctness for beings intent on knowing the world yet aware of their cognitive limitations. What does it take to construct a standard of truth under these conditions? What modes of thought do we need in order to have such a standard? My answer is that truth requires three basic modes of human thought: *immanence*, *transcendence*, and *normativity.* I call this "the fundamental principle of truth".

Truth, according to this principle, requires, in the first place, an *immanent* mode of thought. To think immanently is to think about some object external to our thinking, or to think in the way one thinks when one stands within a theory: namely, hold the world (or something in the world) in one's cognitive gaze and say something about it (for example, attribute some property to it). Immanent thought is thought about the world, and without it there is neither room nor need for truth. The "bearers" of truth are immanent thoughts.

In the second place, truth requires a *transcendent* mode of thought. Since we cannot take it as given that our immanent thoughts are correct, we have to evaluate these thoughts, determine whether they are correct or incorrect. But to determine the correctness of an immanent thought we have to *transcend* it to a point of view external to it and see it in relation to those facets of the world it is directed at. Only then can we evaluate its correctness.[9]

Finally, truth requires a *normative* mode of thought. We need to be able to occupy not just any transcendent standpoint, but a distinctly normative standpoint, one from which we can evaluate immanent thoughts with respect to their *correctness*. Is the world as a given immanent thought says it is? Do objects in the world have the properties (relations) a given immanent thought attributes to them? These are normative questions, and to ask and answer them we need a normative mode of thought.[10]

The fundamental principle of truth thus says that *truth is a transcendent norm of correctness for immanent thoughts*. It follows that an immanent thought has the property of truth iff (if and only if) it satisfies this norm.

6. ***Dynamic Correspondence.*** Another positive principle of truth is the principle of **dynamic correspondence**. This principle, too, is closely related to the human cognitive-epistemic situation. It says that (i) truth is a standard for getting the world right, that (ii) this standard requires a substantial and systematic connection between immanent thoughts (theories, statements) and their target in the world, and that (iii) the precise form this connection takes - its pattern and various details - change in accordance with the complexity of targeted facets of the

---

[9] It is important to note that there is nothing magical about such a standpoint. It does not provide a *God's*-eye view (Putnam 1981); it provides a distinctly *human* view. A paradigmatic example of a transcendent standpoint is a Tarskian meta-language (Tarski 1933). There is nothing super-human about a meta-language, but it is quite powerful relative to the "object language" it transcends. It has in view both the object-language and the world it is directed at, whereas the object language has only the latter in view.

I should also note that transcendence does not have to take the form of a move to a separate language. Kripke's (1975) solution to the Liar paradox, for example, achieves transcendence within a single language, namely, by thinking of the extension of truth as constructed in *stages*.

[10] I should add that most transcendent thoughts, including normative transcendent thoughts, are immanent. For example, the thought that "Snow is white" is true is an immanent thought: it attributes the property of being true to an object in the world, namely, the immanent thought "Snow is white".

world, their accessibility to human cognition, and our resourcefulness in forging cognitive routes for reaching them.

Dynamic correspondence is both more demanding and more flexible than other correspondence principles. It is more demanding in not excluding any field of knowledge from the requirement of a substantial and systematic connection with the world. Even logic is not exempt from this requirement. (See Sher 1999, 2016.) And it is more flexible in allowing us to connect our theories to their targets in the world in any way we can, provided it is sufficiently substantial and systematic. Thus, while traditional correspondence limits truth to a single and simple pattern – copy, picture, mirror image, or isomorphism – **dynamic correspondence** leaves the type and complexity of correspondence patterns an open question. It approaches this question in the spirit of "look and see". Don't decide in advance, but look and see what patterns of correspondence are suitable for different fields of knowledge and under different circumstances.[11]

7. *Example: Mathematical Truth as Composite Correspondence.* To see **dynamic correspondence** at work, let us turn to mathematics. What form does correspondence take in this field? To answer this question we have to find out what facets of the world, if any, are targeted by mathematics. Is there anything in the world for mathematical statements and theories to be true of? to be connected to by a significant correspondence relation? Is there any facet of the world such that to gain knowledge of this facet we need a discipline like mathematics? In answering these questions we have to be careful not to fall into a linguistic trap. Both in natural language and in the professional language of mathematics, it is common to use individual terms (individual constants and variables) to talk about mathematical objects, say, numbers. For that reason, questions of the kind we are asking are usually formulated as questions about mathematical individuals: "Are there mathematical individuals in the world?", "Do the individuals zero, one, two, etc. exist in the world?" But this formulation introduces a bias against

---

[11] "Look and see" is inspired by Wittgenstein (1921). But unlike Wittgenstein, I do not contrast "looking" with "thinking" or with giving a rational account of what we "see".

mathematical correspondence, since the existence of mathematical individuals raises thorny problems (e.g., problems of identity and cognitive access as in Benacerraf 1965 and 1973). To avoid such a bias, I will formulate my question in an open-ended way: "Does the world have, or do objects in the world have, *formal features*?", where both "object" and "feature" are non-specific terms ranging over objects and properties of any type. Furthermore, to avoid empiricist objections at the outset, the individuals I will refer to, at least at the outset, will be observable physical individuals. My initial question will then be: Do physical individuals and their physical properties (relations) have *formal features*?[12]

Now, it is quite clear that the answer to this question is positive. Take the solar system. It is part of the reality of the solar system that both the sun and the planets have the *formal* property of self-identity. It is also part of its reality that it has EIGHT planets, i.e., that the 1st-level physical property "x is a planet" has the 2nd-level *formal*, cardinality, property EIGHT (and not A MILLION or ZERO).[13] Furthermore, it is part of its reality that the 1st-level physical relation "planet x is farther from the sun than planet y" has the 2nd-level *formal* property of being ANTI-SYMMETRICAL. And it is part of its reality that it has physical properties obtained by *formal* operations upon other physical properties: the operations of complementation, intersection, union, etc. (e.g., "x has the intersective property of being BOTH a planet AND gaseous", or "x stands in the INTERSECTION of the properties of being a planet and being gaseous").

But if objects in the world have formal features, then these features are likely to be governed by laws – e.g., laws of cardinality – and we, being epistemically ambitious as we are, are likely to want knowledge of these laws. What form will such knowledge take? The form of a mathematical theory, such as arithmetic or set theory. Such a theory will be correct or incorrect

---

[12] If the answer to this question is positive, then, (i) it should be easy to extend it to non-physically-observable objects (if there are any), and (ii) it would provide a (prima facie) basis for the view that mathematical truth is based on correspondence (of some kind).

[13] I use capital letters for 2nd-level properties.

about the formal features and their laws, and in this sense its truth will be based on correspondence. There is much more to say about these things (see Sher 2015, 2016), but what we have said here is sufficient to provide an example of a nontraditional form of correspondence.

Now, one of the most important cognitive resources we have is language. But language is also a source of limitations. Human language was developed at a time when our understanding of the world was far more primitive than it is today, it is constrained by our biology, psychology, and history, it is designed to play multiple roles, including roles other than tracking the correctness of our theories (e.g., social communication roles), it has accidental features and, overall, it is not an ideal tool for formulating advanced theories of the world. Our standard of truth has to take these factors into account, hence, it cannot pose a blanket requirement of always taking language at face value. That is to say, the connection between correct theories of the world and what they study in the world need not be uniquely determined by the syntactic status, or the literal meaning, of their linguistic expressions. It is possible that in some cases, a direct correlation between syntactic and ontological categories is not available. This, however, does not rule out an indirect, yet systematic, connection between our language and the world.

In investigating the patterns exhibited by the correspondence relation in different fields, therefore, we should not be bound by a *syntax-ontology parity principle*. We should not take it as given that the only way to develop a correct theory of something in the world of ontological level n is to use terms of syntactic level n to talk about it. To take this as given is to fall into the "language trap". To say that the only way for 1st-order arithmetic to be true about the world is for the world to have numerical individuals is to fall into this trap.[14]

Our intuitive considerations suggest that formal features are for the most part of ontological level 2 rather than 0. Cardinalities, for example, are not individuals, but properties of

---

[14] This poses a challenge to Quine's principle of ontological commitment. But it is to be expected that the complexity of the basic human cognitive-epistemic situation would challenge this simplistic principle in any case.

properties of individuals. But our current theories of cardinalities use individual terms (level 0) to theorize about them. Does this necessarily render either the current theories or our intuitive observations incorrect? No. Numerals and numerical individual-variables can be connected to cardinality properties indirectly, though systematically. Since 2nd-level cardinality properties are real, 1st-order arithmetic (set theory) can, in principle, describe them (their properties, the laws governing them) correctly.[15]

But why would humans use 1st-order theories to describe 2nd-level properties, properties and relations of these properties, and the laws governing them? Here we have to distinguish two questions: the psychological question, "Why would humans choose a theory of one level in order to describe phenomena (objects, laws) of a different level?", and the methodological question, "Is it in principle possible to use a theory of one level to study phenomena of a different level?". As philosophers we do not need to give a detailed answer to the psychological question; it is sufficient to show that in principle there could be compelling reasons for such a choice. Suppose, for example, that we, humans, do better in figuring our relations between things (of any level) when we think of them as individuals. Suppose we become distracted or confused when we think about structures of properties of properties, but are good at figuring out things about structures of individuals. In that case, to optimize our investigations, we might build a lower-level model of higher-level structures, where "model" is understood in its everyday sense, as in "a model of a skyscraper". To design a skyscraper, the engineer might find it useful to construct a small plastic model of the skyscraper she is designing. The skyscraper and the model are things of different kinds – the former is made of concrete and steel, the latter of plastic; the former is tall, the latter short. Still, there can be a systematic relation between the two that would render the latter a faithful model of the former.

This provides a positive answer to both our questions: it is quite clear that there could be a

[15] For the sake of familiarity, I sometimes use "nth-order" instead of "nth-level".

good psychological explanation of why humans might use 1st-order theories to study higher-level objects, and it is even clearer that it is possible to construct a correct 1st-level model of such objects. One way to do this is to use the mental operation of *positing.* Exercising epistemic freedom, we can construct a 1st-level structure of *posited* individuals. The posited individuals will be systematically connected both to individual terms in our language and to higher-level objects in the world, and as such could mediate between the two. Reference and correspondence will then be *composite*. For example: the numeral "2" would refer to the 2nd-level cardinality property TWO in two steps. First Step: "2" refers directly to the individual posit *two.*[16] Second Step: *two* is systematically connected to the 2nd-level cardinality property TWO. Correspondence truth-conditions will assume a similar pattern. Step 1: The sentence "2+5=7" is true iff on the level of posits, the posited operation of *addition*, applied to the posited pair of individuals *two* and *five* yields the posited individual *seven.* Step 2: The latter is the case iff in the world, the *DISJOINT UNION* of TWO and FIVE is SEVEN.[17] Finally, exercising transcendence, we explain how our 1st-order mathematical theories are true of *real* cardinalities.[18]

Similar methods can be used in science, though they do not have to involve the same patterns. Here, however, we turn to different questions: What can scientific change teach us about truth, in light of the present account, and what can the present account of truth teach us about the pessimistic claims associated with scientific change?

---

[16] I use lower case italics for individuals in the world.

[17] Here we use "*DISJOINT UNION"* for an operation on properties understood extensionally – in the present case, on formal properties, which are inherently extensional.

[18] There are many further questions I will not discuss here. Nor will I explain the way this account solves many outstanding problems in the philosophy of mathematics. For both, see (Sher 2015, 2016).

III. The Impact of Scientific Change on our Understanding of Truth

The phenomenon of scientific change lends support to the general approach to truth advanced in this paper. First, it reinforces our understanding of the basic human epistemic situation by suggesting that theoretical knowledge is both arduously pursued by, and complicated for, humans. The pessimistic induction emphasizes our cognitive fallibility, the optimistic induction – our cognitive resourcefulness. The historical record further highlights the complexity of the human cognitive-epistemic situation. Even our cognitive assets – from our fertile imagination to our ability to reason that certain things (a planet, a substance) must exist – contribute to our propensity for error.

Second, the pessimistic meta-induction testifies to the centrality of *correctness*, hence *truth* (broadly understood as *correspondenc*e), for science. Our goal is to build *correct* theories of the world, and although we have other goals as well (e.g., significant, explanatory, and unified theories), these add to, rather than replace, truth. In the 20th-century, much attention was paid to other aspects of scientific change: pragmatic considerations, aesthetic advantages, power relations among scientists and their communities, spheres of influence, and so on. (Kuhn 1957 and 1962, Bloor 1976/91, and others). But the pessimistic-induction argument redirects our attention to *veridicality.* Originally, the advantages of Copernican astronomy over its Ptolemaic rival may have been extra-veridical. But the Copernican revolution would not have withstood the test of time if later evidence had not supported its basic correctness.

The centrality of correctness to science thus suggests the importance of a *norm of truth* for it, or, more specifically, a *correspondence* norm of truth (broadly construed). The pessimistic induction suggests that, due to humans' fallibility, such a norm is needed; the optimistic induction suggests that such a norm is not futile. In this respect, the pessimistic and optimistic inductions support one another. Optimists have argued that the likelihood of error is constantly decreasing due to improvements in discovery, evidence, and detection of error. And to some extent, they are right. Certain types of error are unlikely to be repeated in the future. But significant improvements are likely to be made in the future as well, and these will increase the likelihood of detecting errors in today's theories. That, together with the continually increasing difficulty of the questions asked by scientists, suggest that a norm of truth will continue to play a significant role in science.

These are a few ways in which the phenomenon of scientific change supports the view of truth described earlier in this paper. However, the phenomenon of scientific change teaches us new things as well, in particular, things about the ***dynamic nature* of truth**. The view of truth delineated in this paper is already **dynamic**, but so far we have limited our attention to one type of **dynamics**, namely, **contextual dynamics**, and in particular the potential variability of correspondence patterns from field to field. Scientific change, however, is largely *temporal*. And this suggests that **our standard of truth is dynamic not just contextually but also temporally**.

What form does the **temporal dynamics of truth** take, and what would be examples of such **dynamics**? As a preliminary to answering these questions, let us briefly consider the question of the *unity* and *diversity* of truth. In discussing the disunity of science, Dyson (1988) says that science requires a fruitful balance between unity and diversity. I believe that this holds for the norm of truth as well. On the one hand, we want to retain the hard core of truth, namely, its connection with correctness, its reflection of our persistent desire to get the world right. But we also want it to take into account the changing circumstances of our pursuit of knowledge. To that end we may distinguish two parts of this norm: a general, and fairly stable, part, and a part consisting of particular, changeable elements. The first part will include the fundamental principle of truth (immanence, transcendence, and normativity), the principle of systematic connection between theory and world underlying correspondence, and other general principles. The second part will include the truth conditions of scientific (and other) theories and their statements. The truth conditions of statements and theories may change both in context and in time.[19]

[19] Note that this division requires a substantivist, rather than a deflationist, approach to truth. If the truth conditions of all sentences are disquotational, as deflationists say they are, then there is no significant difference between truth conditions in different fields or at different times.

To better understand the **temporal dynamics of truth**, we need to better understand the role of truth conditions in knowledge. And to understand their role in knowledge we need to understand their role in justification. We have already discussed the reason justification requires a norm of truth, but this had to do with the fixed part of such a norm. Justification, however, requires truth conditions as well. Their role is to close the gap between justification and truth. This gap is especially deep in advanced theoretical sciences, due to the prevalence of indirect justification in those sciences (or parts thereof). For example, to establish the existence of a certain subatomic particle, scientists devise an intricate test, the result of which is accessible to us through, say, our computer screen. Scientists might say that the experiment's result, as seen by a certain picture, or a set of pictures, on the computer screen, (indirectly) justifies their theory's claim that the particle exists.

Let (1) be the claim that the particle exists, and (2) the claim that the experiment's result (as visible on the computer's screen) justifies (1). Now, for (2) to be true, there must be some specified connection between the conditions under which a certain picture appears on the screen and the conditions under which (1) is true. And to determine whether such a connection exists, we need to know the truth condition of (1). What we need are not the disquotational truth conditions of (1), which tell us nothing informative, but substantive truth conditions, truth conditions that provide information about the identity conditions of the particle in question, about some of its distinctive features, and so on. In practice, the truth conditions of (1) need not be officially or fully formulated; in practice, they might be partial and implicit. But we need a clear idea of what these conditions are in order to determine whether the results on the screen justify (1).

Having a clearer idea of the role played by truth conditions in knowledge, we can look for

---

(The disquotational truth-condition "'2+5=7' is true iff 2+5=7" cannot distinguish whether "2+5=7" is systematically connected to the world in a simple or a composite manner and what facet of the world its truth connects it to – a facet involving numerical individuals, or a facet involving 2nd-level cardinality properties.)

examples of a temporal change in such conditions. It is reasonable to expect that the discovery of, say, significant new details about the identity conditions of some object or substance will lead to such a change. Take, for example, *water.* Whereas long ago humans could characterize water only in phenomenological and simple observational physical terms (transparent, tasteless, evaporates when heated, etc.), today we can characterize it in more precise identifying terms, namely, chemical terms specifying its molecular structure (H_2O). Water itself did not change in the course of time, but our understanding of the conditions that have to hold in the world for a statement about water to be true did change. Once the molecular structure of water was discovered, we could, did, and should have taken advantage of this discovery to improve our ability to distinguish true and false statements about water by updating their truth conditions.

It is important to note that the **temporal dynamics of truth** preserves the desired balance between its fixity and changeability. Truth remains a norm of getting the world right, but the increase in knowledge, the improvements in our cognitive tools, and the progress in our overall understanding of the world, provide us with new and better means of pursuing the goal of building true theories of the world, theories that get the world right.

IV. The Impact of Truth on our Understanding of Scientific Change

The study of scientific change in the last five decades brought about several pessimistic arguments concerning truth in science. One of these is the pessimistic meta-induction, which we focus on here.[20] This simple, yet potentially powerful, argument says that scientific theories embraced by past generations were eventually rejected by later generations; hence, it is likely that scientific theories embraced by our generation will be rejected by future generations. This argument is usually presented as an argument against scientific realism, and this points to a

[20] Others include the underdetermination argument (Duhem 1906/14, Quine 1951), the incommensurability argument (Kuhn 1962), the unconceived-alternatives argument (Stanford 2006), and so on.

significant fact about this argument: what it questions are not the practical or pragmatic credentials of present science but its *veridical* credentials. In particular, it questions our belief that contemporary science is *true*, or even roughly true, in the *correspondence* sense, broadly understood.

Does the conception of truth developed in this paper have anything to contribute to our understanding of the pessimistic-induction argument and its significance? I think the answer is "Yes". Our understanding of the nature of truth helps to put the pessimistic-induction argument in perspective by alerting us to certain unwarranted assumptions that are implicit in it and make it seem more consequential than it actually is. These assumptions revolve around the idea of *THE truth*, and detecting them leads us to question whether what science *aims* at, even at the ideal limit, is *THE truth*, and whether a negative answer to this question undermines either scientific realism or the related view that truth (correspondence) is integral to the scientific endeavor.

Thinking in terms of "*THE truth*" is very natural to humans, given our cognitive make up and the grammar of our language. But language, as we have noted earlier, is sometimes a trap. Language, in Frege's words, "unavoidably" introduces "misconceptions" into our thinking, and it is one of our tasks to "lay... bare" these misconceptions and "free... thought from that with which only the means of expression of ... language, constituted as they are, saddle it" (Frege 1879: 7). The idea of reaching *THE truth*, taken literally, suggests that there is something which is *THE truth*, and scientists' ultimate job is to reach it. Reaching *THE truth*, on this understanding, is like reaching the top of a mountain. Once we have reached the top of the mountain, we are done. Our goal has been achieved. Once the scientific community has reached *THE truth* about nature, its job is done. But this implies *the end of science*. Once we have reached *THE truth* about nature, nature has been conquered. End of science. While this picture is harmless when we think about very limited scientific endeavors – e.g., finding the truth about the molecular composition of water – it is not when it comes to more expansive epistemic endeavors, such as knowing the physical forces acting upon nature, or knowing nature in its entirety.

What is wrong with this picture? Two things that are wrong with it are the existence and uniqueness assumptions implicit in the idea that the goal of science is reaching *THE truth.* According to this picture, there exists one and only one thing that science aims at knowing, and this thing is *THE truth.* Truth is an object we want to know, and there is only one such object. But both implications are wrong: truth is not a thing; and there is more than one thing that scientists want to know. Indeed, there is no end to the things that science seeks to know .

Let us begin with existence. Truth, according to our theory, is a *standard* rather than a thing (something that exists in the world). It is a standard for our theories of the world that reflects our epistemic aspiration of getting the world right. This aspiration, in the case of science, focuses on *nature* – the natural world, or those facets of the world that are open to empirical investigation. What scientists seek to know are natural objects (events, phenomena), their properties (relations), the natural laws governing them, etc. But truth is not one of these things. Scientists seek to know, say, the origins of the universe, but they do not seek to know an object named "truth" (or "the truth") – either in place of seeking to know the origins of the universe, or in addition to it. Truth is not like gravity, which is something scientists seek to know in addition to the origins of the universe, or like the electromagnetic force which they seek to know in addition to the gravitational force. Truth is a standard that we, humans, constructed for our theories, a standard that guides and constrains our epistemic enterprise. To say that truth is the aim (goal, end) of science is just an imprecise way of saying that science aims at knowledge of the world (various facets of the world), and that such knowledge takes the form of theories (or models) that are guided, constrained, challenged, and justified by (or in light of) our standard of truth.

Turning to uniqueness, not only is *THE truth* not a unique object that science seeks to know, there is no uniquely true thing that science aims at. Science strives to know many things (objects, events, phenomena, laws) in and of the world, and it aims at attaining better and better (more and more significant, accurate, explanatory) knowledge of these things. As we discover new things, ask new questions, and develop new tools, the goal of science expands and its direction shifts. This is another reason it is misleading to talk about *THE truth* as the goal of science, to be reached once and for all or even be approximated. By counting more and more numbers we do not reach or even approximate the end of the number series. The idea of reaching our destination – even at the ideal limit – and ending our journey is not applicable to human knowledge. But the idea of advancing our knowledge is. For that reason, the worry that science will never reach this end, will never reach *THE truth*, is a frivolous worry. There is no *THE truth* for science to reach. There is progress to make, but this is an entirely different matter.

What is the relation between scientific progress and truth? What, exactly, are our truth-related responsibilities as seekers of scientific progress? First, let us consider what truth-responsibilities we do *not* have. We are not responsible for anything that is not achievable by humans. We are not responsible for satisfying any standards that require super-human capacities. We are not responsible for anything that requires *absolute* transcendence (or, in Putnamian language, a God's-eye view). In the present context, this means that we are not responsible for measuring the correctness of our scientific theories based on any absolute truth conditions, truth conditions that we, as humans, are incapable of knowing. Moreover, we are not responsible for measuring our *current* theories against truth conditions that we have no idea about at the *present*, even if these will be the truth conditions that we will use in the future.

What truth responsibilities *do* we have? We are responsible for developing a standard of truth that reflects our aspiration for theories that get the world right, for applying this standard to our theories, and for doing our best to satisfy this standard (or have our theories satisfy it). This standard, I have argued, is a correspondence standard – a substantial yet flexible correspondence standard (taking into account the complexities of the human cognitive-epistemic situation). As a standard devised by and for humans, it has, we have seen, two parts: a fixed part, and a part that changes in time.[21] Throughout the development of science, we are responsible for doing our best

---

[21] It is important to note that "fixed", here, does not imply "infallible". All aspects of human thought, including our understanding of truth, are fallible. But a fallible understanding

to satisfy the fixed part by connecting our theories to their target in the world in a substantial and systematic manner. And during each period we are responsible for doing so in accordance with the truth conditions we attribute to our theories during that time. But this is not all. It is also our responsibility to critically examine our standing truth conditions with an eye to sharpening, refining, correcting, and/or improving them, using all the knowledge we have and the best cognitive tools available to us at the time.

But is this too little? Given the likelihood of future changes in science, including changes in our judgments about the truth credentials of our current theories, how can taking responsibility for their credentials today (in light of our current understanding) contribute to future progress? The answer to this question brings us back to holism, basic realism, and the **dynamics of human inquiry**. Ask yourself: Do we have a better chance of improving our theories in the future if at present we subject them to what, as far as we can see today, are appropriate truth requirements, than if we accept or reject them without regard to these requirements, that is, at random, or based on gut feelings or on a whim? What is the point of the Large Hadron experiments if we do not connect their results with the conditions that have to hold in the world for particles such as the Higgs boson to exist, that is, with the (substantive) truth conditions of claims like "Higgs boson particles exist"? These connections do not have to be direct by any means (recall "composite correspondence", though the indirectness here might be of a different kind). But scientists have to establish, explicitly or implicitly, substantial and systematic connections between their experiments and what they take to be the conditions that have to hold in the world for the tested theories (or hypotheses) to be true. That is, they have to do so if they share the basic human aspiration of gaining correct theoretical knowledge of the world, including those facets of the world that are not easy for us to access. We may thus say that we are responsible for satisfying the truth conditions of our theories, as we understand them today, so that tomorrow we are in a better

---

may include two parts, a part designed to be fixed, and a part designed to change along some parameters.

position to discover problems either with their satisfaction or with the conditions themselves. The discovery of such problems will provide us with resources for developing better theories, better truth conditions, and still better theories that satisfy these better truth conditions.

This outlook takes the bite off the pessimistic-induction argument. What follows from the history of science is not that human theories are bound to be false, but that theoretical change is likely to continue. There is nothing hopeless about this conclusion, no ground for giving up either our aspirations for true theories, or our basic realism. On the contrary. Since (as far as we can see today) many of the changes that took place in science in the past turned out to be (by our current standards) improvements, including improvements in correctness (truth, correspondence), it is reasonable to expect that a significant number of future changes will have the same (desirable) feature.

Scientific progress is a holistic process. One of its central aspects is truth. We use our current understanding of truth (truth conditions) to check the correctness of theories, and we use our developing body of theories to check the adequacy of our current standard of truth (truth conditions). This process bears some similarities to Rawls's (1971) reflective equilibrium process. Its **dynamics** reflects both the **dynamics of scientific change** and the **dynamics of truth** (on the present account). The tale of scientific change is both hopeful and cautionary. It cautions us against the existence of undiscovered errors, and it gives us hope of discovering and correcting these errors, as well as of expanding our knowledge to hitherto unknown regions of the world. Both discovery and correction require a heightened awareness of truth: of the fixed aspects of truth, as well as of its changeable aspects.[22]

---

[22] The discussion of truth in this paper has ramifications for other topics of current interest, including the topics of *approximate truth* (Boyd 1990 and many others) and *unconceived alternatives* (Stanford 2006), which are both related to scientific change and the pessimistic induction. These, however, must await another occasion.

References

Benacerraf, P. 1965. “What Numbers Could Not Be”. *Philosophical Review* 74: 47-73.

----. 1973. “Mathematical Truth”. *Journal of Philosophy* 70: 661-79.

Bloor, D. 1976/91. *Knowledge and Social Imagery.* Chicago.

Boyd, R. 1990. “Realism, Approximate Truth, and Philosophical Method”. In *Scientific Theories,* (ed.) C.W. Savage. Minnesota University Press: 350-91.

Chakravarty, A. 2011. “Scientific Realism”. *Stanford Encyclopedia of Philosophy.*

Doppelt, J. 2007. “Reconstructing Scientific Realism to Rebut the Pessimistic Meta-Induction”. *Philosophy of Science* 74: 96-118.

Duhem, P. 1906/14. *The Aim and Structure of Physical Theory.* Princeton, 1954.

Dummett, M. 1973. “The Significance of Quine’s Indeterminacy Thesis”. *Truth and Other Enigmas.* Cambridge: Harvard, 1978. Pp. 375-419.

----. 1973/81. *Frege.* New York: Harper & Row.

Dyson, F. 1988. *Infinite in All Directions.* New York: Harper & Row.

Fodor, J. & E. Lepore. 1992. *Holism.* Oxford: Basil Blackwell.

Frege, G. 1879. "Begriffsschrift". *From Frege to Gödel*, ed. J. van Heijenoort. Cambridge: Harvard, 1967. Pp. 5-82.

Glymour, C. 1980. *Theory & Evidence.* Princeton.

Haack, S. 1993. *Evidence and Inquiry.* Oxford: Blackwell.

Horwich, P. 1990/8. *Truth*. Oxford.

Kitcher, P. 1993. *The Advancement of Science: Science without Legend, Objectivity without Illusions.* Oxford.

Kuhn, T.S. 1957. *The Copernican Revolution: Planetary Astronomy in the Development of Western Thought.* Cambridge: Harvard.

----. 1962. *The Structure of Scientific Revolutions.* Chicago.

Kripke, S. 1975. “Outline of a Theory of Truth”. *Journal of Philosophy* 72: 690-716.

Laudan, L. 1981. “A Confutation of Convergent Realism”. *Philosophy of Science* 48: 19-49.

----. 1990. *Science and Relativism: Some Key Controversies in the Philosophy of Science.* Chicago.

Neurath, O. 1932. “Protocol Statements”. *Philosophical Papers 1913-1946.* Dordrecht: D. Reidel, 1983. Pp. 91-9.

Niiniluoto, I. 2011. “Scientific Progress”. *Stanford Encyclopedia of Philosophy.*

Psillos, S. 1996. “Scientific Realism and the ‘Pessimistic Induction’. *Philosophy of Science* 63: S306-14.

----. 1999. *Scientific Realism: How Science Tracks Truth.* Routledge.

Putnam, H. 1975. “What is Mathematical Truth?”. In *Mathematics, Matter and Method.* Cambridge.

----. 1981. *Reason, Truth and History.* Cambridge.

Quine, V.W. 1951. “Two Dogmas of Empiricism”. *From a Logical Point of View.* Cambridge: Harvard. Pp. 20-46.

----. 1960. *Word and Object*. Cambridge: MIT.

Rawls, J. 1971. *A Theory of Justice.* Cambridge: Harvard.

Roush, S. 2010. “Optimism about the Pessimistic Induction”. In P. D. Magnus & J. Busch (eds.), *New Waves in Philosophy of Science.* Palgrave Macmillan.

Sher, G. 1999. “Is There a Place for Philosophy in Quine's Theory?” *Journal of Philosophy* 96: 491-524.

----. 2004. “In Search of a Substantive Theory of Truth”. *The Journal of Philosophy* 101: 5-36.

----. 2015. “Truth as Composite Correspondence”. *Unifying the Philosophy of Truth.* Eds. T. Achourioti et al. Dordrecht: Springer. Pp. 191-210.

----. 2016. *Epistemic Friction: An Essay on Knowledge, Truth, and Logic.* Oxford.

Sosa, E. 1997. “Reflective Knowledge in the Best Circles”. *Journal of Philosophy* 94: 410-30.

----. 2009. *Reflective Knowledge.* Oxford.

Stanford, P.K. 2006. *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives.* Oxford.

Tarski, A. 1933. "The Concept of Truth in Formalized Languages". In *Logic, Semantics, Metamathematics,* Indianapolis: Hackett, 1983. Pp. 152-278.

Williams, B. 2002. *Truth and Truthfulness.* Princeton.

Wittgenstein, L. 1921. *Tractatus Logico-Philosophicus.* London: Routledge & Kegan Paul, 1961.