J Gen Philos Sci (2017) 48:371–394
DOI 10.1007/s10838-017-9370-3

ARTICLE

Truth and Scientific Change

Gila Sher[1]

Published online: 7 August 2017

Abstract The paper seeks to answer two new questions about truth and scientific change: (a) What lessons does the phenomenon of scientific change teach us about the nature of truth? (b) 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?

Keywords Truth · Scientific change · Pessimistic meta-induction · Dynamic correspondence · Realism · World-oriented holism

In physics the truth is rarely perfectly clear.
Richard Feynman

Two central questions raised in 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?

1 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 of 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

✉ Gila Sher
gsher@ucsd.edu

[1] Department of Philosophy, 0119, UCSD, La Jolla, CA 92093-0119, USA

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. Dramatic examples of scientific change, noted by Quine (1951), include the changes from Ptolemaic to Keplerian astronomy, from Aristotelian to Darwinian biology, and from Newtonian to Einsteinian physics. Other examples, noted by Laudan, include the rejections of aether theories in physics, the caloric theory in thermodynamics, the vital-force theory in organic chemistry, and so on.

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 [probably] **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 2015, 16)

Quite a few philosophers of science (e.g., Kitcher 1993 and Roush 2010, in addition to Psillos 1996, 1999 and Doppelt 2007) 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.*

2 The Nature of Truth[2]

To focus on those aspects of truth that are pertinent to scientific change, I will approach truth from an *epistemic perspective*. What I mean by this, however, is different from what philosophers usually mean when they speak about an epistemic conception of truth. Normally, what they mean is that truth is reduced to some other, distinctly epistemic, norm or concept, e.g., superassertibility (see Wright 1992) or justification (see Rorty 1995) or

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

[2] This section incorporates and, in the case of truth, further develops, several aspects of a theory of knowledge and truth developed in Sher (2004, 2016a).

acceptance at the ideal end of inquiry (see Peirce 1878). But what I mean is different. What I mean is that we approach truth by asking whether a norm or a concept of truth, distinct from any narrowly epistemic norm/concept, is needed either for the pursuit or for the understanding of knowledge. And if the answer to this question is positive, then we can use it as a key for understanding the nature and/or principles of truth. In theorizing about truth from this standpoint, my approach is neither quietist nor deflationist, nor, indeed, platitudinous. My approach is substantivist, inspired by Kant's philosophy more than by the later Wittgenstein's.[3]

The conception of knowledge I am interested in here is a robustly realist conception, the kind of conception that has been challenged by the pessimistic-induction argument, namely, one that encompasses not just observational knowledge but also abstract, theoretical knowledge. But to leave myself room for an open-ended investigation, I prefer not to be overly specific. Rather than beginning with any of the specialized accounts of knowledge developed in the philosophical literature, such as Quine's naturalism, Goldman's reliabilism, Sosa's virtue epistemology, etc., I will signal my initial approach to knowledge with a broad sketch of the so-called "basic human cognitive-epistemic situation".

The rough idea, expressed in general, common-sensical terms, is that we, humans, live in a world of which we are a part. We aspire to obtain objective knowledge of that world, both practical and theoretical. However, obtaining objective theoretical knowledge of the world is not a simple matter. 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. Nevertheless, we, humans, are ambitious creatures, and we aspire to know the world in its full complexity. What makes these aspirations possibly realizable, at least in part, is the circumstance that, notwithstanding our limitations, we have a variety of fairly intricate cognitive resources of various kinds—sensory and intellectual, innate and learned, passively and actively generated and employed. The latter explain our capacity to play an active role in both discovery and justification—designing experiments, developing research programs, making calculations, drawing inferences, gathering evidence, and so on.

Three cornerstones of my approach to knowledge, reflected in this picture, are the views that (1) knowledge is strongly oriented toward the world, (2) theoretical knowledge is both made possible and complicated by our cognitive make-up, and (3) the acquisition of knowledge is an active, dynamic process. These might seem platitudinal principles, but as we shall see from their substantial consequences for realism, philosophical methodology, and the theory of truth, they are not.[4]

A related approach to (1)–(3) starts from rationality and proceeds to the pursuit of knowledge. It is a common-sensical observation that every rational act requires both *friction* and *freedom*: friction, or constraints, set by our goals, environment, and norms of rationality; and freedom to set and pursue these goals in accordance with our norms (Sen

[3] For an overview of substantivism, differences from deflationism, and references, see Sher (2004, 2016a, b). Essentially, a substantivist theory of truth regards truth as a deep, important, complex subject-matter, with applications in, and ramifications for, many fields, and it sets for itself demanding norms of correctness, explanatory power, systematicity, philosophical interest, etc.

[4] I should note that (1)–(3) can be arrived at through different pictures of the basic human cognitive/epistemic situation, and some readers might prefer other pictures. But all I need here is one such picture to be used as a starting point. Still, in the next paragraph I will offer a somewhat different, though related, perspective on (1)–(3).

2002; Sher 2016a). To the extent that the pursuit of theoretical knowledge is a rational act, it, too, requires both freedom and constraint. Our theories are constrained by our goal of knowing the world, and this means that they are constrained both by the world itself (as their target, what they are charged with providing knowledge of) and by the epistemic norms we ourselves impose on them in light of this goal. But the creation and imposition of such norms requires epistemic freedom, as does the pursuit of knowledge. We need freedom to design experiments, critically evaluate our discoveries, justify or refute our theories, and so on.[5]

2.1 Realism

Given humans' aspiration to knowledge of the world, as it is and in its full complexity, theories developed in pursuit of this aspiration have *realism* (a realistic orientation) built into them. But what kind of realism is it? Our considerations above suggest a kind of realism that I will call "basic realism", where "basic" both magnifies and attenuates the traditional sense of "realism". First and foremost, basic realism is the view that, given our epistemic goals, realism is integral to human knowledge as such. As a result, it views realism as *universal*: integral to *all* fields of knowledge, including abstract fields (sub-fields).[6] Second, basic realism construes "realism" as a *thick*, *substantial* notion rather than a thin, deflationist, platitudinal, or quietist notion. Accordingly, basic realism sets thick, substantial demands on all fields of knowledge. But basic realism also recognizes that our theories of the world, being developed by, and intended for, humans, cannot be shaped only by the world, but are inevitably also shaped by our cognitive make-up and activities. In this sense it is an *attenuated* realism. This, however, does not render it a weak realism. In particular, it does not weaken its demand of seeking and making progress toward knowledge of the world *as it in fact is*. What it does do is allow *flexibility* with respect to *how* we go about making progress toward this goal. To take a simple example, creatures without color vision would not be able to go about attaining knowledge of the macroscopic world exactly as we do, but, assuming they would have sufficiently rich cognitive resources (other than color vision), they would be able to obtain knowledge of the world that is as rich and factual as ours (and, assuming their cognitive capacities would not be superior to ours, we are able to obtain knowledge of the world that is as rich and factual as theirs).

Realism is sometimes viewed as assuming that there is just one way to describe the world *as it is*, but basic realism says that there is *no* unique way to describe it *as it is*. Basic realism requires our theories to be correct about the world, but it grants us considerable flexibility with respect to how we access the world, what facets of the world we target at a particular stage in the development of our theories, and what terms we use to describe those facets. Similarly, realism is often described as requiring *complete independence* of

[5] Both my conception of the basic human cognitive situation and my conception of epistemic friction and freedom are partly influenced by Kant. But my view also differs from his on many counts. For example, there is no room in my view for either the duality of thing in itself and appearance or for the idea of fixed categories. Nor do I adhere to Kant's sharp dichotomies of the analytic and the synthetic or the apriori and the aposteriori.

[6] Among realists who do not affirm the universality of realism are radical empiricists and nominalists (such as Hume 1739–40; Goodman and Quine 1947; van Fraassen 1980), who limit their realism to observational knowledge. Among those who affirm realism's universality is Quine, in some (though only in some) of his moods, as when he says that (even) logic is "world-oriented" rather than "language-oriented" (1970/86: 97).

our theories from our mind.[7] But basic realism regards this characterization as unnecessarily extreme, replacing it by the requirement that our theories be *substantially grounded* in the world and that they pursue *strong norms of correctness* with respect to it.

It is important to emphasize that basic realism is *not* a form of *anti-realism* (either scientific anti-realism or any other type of anti-realism). It does not compromise the requirement that our theories describe the world as it in fact is, although it is flexible with respect to the ways they might describe it as it is. We may say that basic realism is a "realistic" version of realism—"realistic" in the sense of recognizing (not being blind to) the variety of factors involved in determining how creatures like us can adequately describe the world.[8] Making realism "realistic" in this sense does not amount to *relativism* either. In measuring distances we must use some units of measure, but while measuring in inches and measuring in centimeters give different results, both results are equally objective. Our cognitive limitations do limit our ability to measure some distances, e.g., very small and very large distances. But while we will never overcome some limitations, we have and will overcome others. Distances that we could not measure in the past (e.g., distances between galaxies and distances between atoms) we can measure today, and it is likely that our distance-measuring capabilities will further improve in the future. Different rational creatures—or we in some counterfactual scenarios—might have developed different measuring methods, but there is nothing relativistic about this in a sense that compromises realism.

Basic realism bears some similarities to Putnam's (1981) *internal realism*, but it also differs from it in significant ways. For one thing, it is associated with a substantial correspondence theory of truth (though not with one that involves a "God's eye view" and/or "thing in itself"), as we shall see below. Basic realism also recognizes that our theories of the world must take *pragmatic* considerations into account. But it is *not* a form of pragmatism. It is more open to pragmatic considerations than more rigid forms of realism, but it gives clear priority to considerations of correctness. In short, basic realism is a flexible yet robust form of realism: flexible about how we go about studying the world, robust in its unflagging commitment to studying the world as it in fact is.

2.2 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, 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 (our) goals. 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 pessimistic-induction argument attacks scientific realism by attacking the no-miracle argument. It says that success by itself does not guarantee, or even strongly supports, truth (correctness, factuality).

But in the literature there is another argument for realism as well, namely, the "realistic aspirations" argument. This argument is expressed by van Fraassen as follows:

[7] See, e.g., the entries on "realism" and "scientific realism" in the Stanford Encyclopedia of Philosophy (Miller 2014; Chakravarty 2017).

[8] For a similar attitude, see Kitcher (1993).

> *Science aims to give us, in its theories, a... true story of what the world is like; and acceptance of a scientific theory involves the belief that it is true* (van Fraassen 1980, 8).[9]

My own considerations (earlier in Part 2) suggest a similar, though not identical, argument: We aspire to know the world as it in fact is, and this aspiration requires a realistic conception of knowledge. In light of this aspiration, a theory that lacks a realistic orientation does not count as a genuine theory of the world, as conveying veritable knowledge of the world.

The no-miracle argument is 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 realistic 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 the no-miracle argument also complements the realistic-aspirations argument. As pointed out by, e.g., Kitcher (1993), in principle, the realistic-aspirations argument might be empty: what we aspire to and what is feasible for us are two different things. The no-miracle argument shows that it is not. It starts from what is the case and argues that, in effect, science (as it actually is) satisfies, or makes significant progress toward satisfying, our realistic aspirations. But while one way to establish the feasibility of our realistic aspirations is to appeal to the no-miracle argument, it has not been shown that this is the only way to establish it.[10] So, when the non-miracle argument is under attack, it is important to keep in mind that it is not the only argument for scientific realism. To refute scientific realism, one has to refute the realistic-aspirations argument as well. The realistic-aspiration argument says that what counts, and should count, as veritable *knowledge* is only knowledge in the realistic or (in my view) basic-realistic sense. To refute scientific realism one has to undermine this claim as well.

2.3 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, the question arises what methodology of discovery and justification will best serve us as seekers of theoretical knowledge. The two dominant epistemic methodologies in the philosophical literature are foundationalism and coherentism. But neither will do. Coherentism will not do because its focus is on the relation between theories rather than the relation between theories and their target, the world.[11] Foundationalism will not do because it is a highly problematic—arguably self-defeating—methodology. Its main problem, in my view, is the strict ordering, or strict hierarchical structure, it imposes on our system of knowledge. An item or unit of knowledge can only be grounded (discovered, justified) using resources generated by lower units in the hierarchy, and the foundationalist methodology requires ultimately basic or lowest units. The problem is that on the one hand, the basic units carry the whole burden of grounding our knowledge, but on the other hand, the basic units

[9] For further formulations and discussion of the axiological (aspirational) argument for realism see Lyons (2005).

[10] I believe that our discussion of truth below offers another route for showing that, how, and in what sense our epistemic aspirations are realizable.

[11] Even coherentists such as BonJour (1985), who recognize the world's significance to knowledge, do not center their methodology on the mind-world relation.

themselves are in principle ungroundable: there are no *more basic* units that could provide resources for grounding these units. As a result, there are no resources for grounding our system of knowledge as a whole or, indeed, any of its units.[12]

My preferred alternative to both foundationalism and coherentism is *foundational holism* (see Sher 2016a). Like foundationalism, foundational holism requires all units of knowledge to be significantly grounded in the world (i.e., be subject to strong norms of discovery and justification centered on the world), but unlike foundationalism it does not impose a (self-defeating) strict-ordering requirement on this grounding. The key idea is that we can and ought to utilize our cognitive resources in whichever way and order best works for us at each stage of the pursuit of knowledge; there is no pre-fixed ordering requirement on either discovery or justification. We need, and can, exercise flexibility both in finding ways to reach those facets of the world that are difficult for us to reach and in devising methods for evaluating the correctness of theories in these fields. Unlike foundationalism, foundational holism emphasizes the *multiplicity* and open-endedness of knowledge-inducing connections between theory and world, and unlike coherentism, it does not compromise the focus on the world. It requires a critical attitude toward all branches of knowledge and a critical examination of each stage in the development of our theories. But it regards this critical examination as an ongoing process. At each stage we have to hold some things fixed (treat them as given), but these very things must be critically examined ("unfixed") at other stages. In this way, no area is left unexamined, so (discoverable) errors in every area are eventually discovered. Progress is ensured, but no (unrealistic) promise of final, ultimate, infallible knowledge is given.[13,14]

One significant advantage of the foundational-holistic methodology is its ability to deal with circularity and infinite regress. Foundationalism bans all circularity and infinite regress, but these are unavoidable in substantive studies of fundamental philosophical subject-matters such as truth or knowledge. Holistic methodologies in general distinguish between destructive and innocent circularity (regress), rejecting the former while affirming the latter. Indeed, foundational holism recognizes the existence of "constructive" circularity. Three examples of such circularity are Gödel's (1931) use of arithmetic to define its own syntax, Glymour's (1980) bootstrap theory, and Rawls' (1971) reflective equilibrium.

Another significant aspect of foundational holism is its recognition of the multifacetedness of humans' cognitive resources and their complementarity. This does not mean that it views everything that has ever been claimed to be a cognitive resource as an actual

---

[12] One might think that we could solve this problem by allowing self-grounding units, but this would render the grounding of knowledge irredeemably mysterious (magical), and in any case, no adequate account of self-grounding knowledge is known to exist.

[13] To prevent misunderstandings, let me note that the holism intended here is not a "total" or "one-unit" holism" (Dummett 1973; Sher 2016a). *Total* or *one-unit holism* is the view that the smallest unit of knowledge is our system of knowledge as a whole. On this view, our entire body of knowledge is a huge atom or blob, lacking inner structure. Our body of knowledge either provides knowledge of the world in its entirety or it does not provide any knowledge at all. This view is attributed to Quine by Dummett (1973, 1973/81) and Glymour (1980), and it is rejected by foundational holism on the ground (originally due to them, as well as to Fodor and Lepore 1992) that if knowledge has no inner structure, a step-by-step acquisition of knowledge is impossible. Foundational holism is, in contrast, *relational* (*structural*) rather than total.

[14] 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 the two emphasize different aspects of discovery and justification. And while Foundherentism is limited to empirical science, foundational holism applies to all branches of knowledge, including logic and mathematics (see Sher 2016a).

resource. Foundational holism is committed to a critical examination of all claims about would-be resources of knowledge, and as such it is unlikely to recognize, say, spiritual revelation as such a resource. But it does recognize that all theoretical knowledge, experimental as well as abstract, requires multiple cognitive resources, and in particular, that even the experimental sections of scientific theories require resources that go beyond sensory perception. For example, all theoretical knowledge requires *intellectual* resources of *some* kind or another.[15] To maximize effectiveness, intellect and sensory perception must (and do) work in tandem in science.

Summing up, I would say that foundational holism combines holism and realism. It is uncompromisingly geared toward knowledge of the world, i.e., realistic knowledge. Yet it has a broad, flexible, open-ended, and "realistic" view of what it takes for humans to make progress in the pursuit of such knowledge.

2.4 Truth

If one of our main epistemic aspirations is to gain genuine knowledge of the world (rather than to tell interesting, aesthetically pleasing, or merely useful stories about it), if what we aspire to is an understanding of the world as it in fact is (rather than as we wish it had been, as we are afraid it might be, or as we imagine it to be), if, as a result, the theories we build are realistic theories, theories that aim at getting the world right, and if getting the world right is a challenge for us (not something that we automatically achieve or that we are guaranteed to achieve), then the question of *correctness* (in the everyday sense of the word) is a serious and central question for us, both as theorists and as critical assessors of theories. But if correctness is important to us, and if we recognize the difficulties involved in achieving it, then we have to put in place a system of checks for our theories, centered on their correctness. And looking at scientific practice, we see that such a system is indeed at work, for example, in scientists' efforts to provide evidence and counter-evidence for their theories and in the intricate measures they take to provide such evidence (designing experiments, building sophisticated instruments, and so on).

So far, we have talked in terms (e.g., "correctness" and "getting the world right") that might be viewed as somewhat vague. Our task as theoretical philosophers is to develop a systematic theory of what is vaguely designated by these terms. Following philosophical tradition, I will use the term "theory of *truth*" for such a theory. In developing a theory of truth I will use the foundational holistic methodology described above. I will go back and forth between analysis of the cognitive conditions required for truth to emerge in our life and the demands set on truth by its role in the pursuit of knowledge. In accordance with this methodology, I will not reject circularity altogether, though I will do my best to avoid vicious and trivializing circularity.

My starting point is the semi-Kantian question: Under what cognitive conditions does the question of truth—the question "Is X true?"—emerge in our epistemic life? My answer is that for the question of truth to arise at all, we need, in the first place, to be able to think

[15] My own paradigm of intellect differs from the traditional paradigms, which range from immediate intuition to pragmatic conventions. My paradigm is "figuring out", that is our ability to figure out things in the ordinary sense of the word. I do not have room to elaborate on this here, but figuring out is something we do at every stage of our life and in most contexts. A baby figures out new things all the time, a technician figures out what caused an instrument to malfunction and how to fix it, a scientist figures out what experiment would be both adequate and effective in testing a given hypothesis, a mathematician (logician) figures out whether a given mathematical system is complete (in the logical sense of the word). And so on. For further discussion see Sher (2016a).

about something external to our thinking, something in the world (broadly understood). I call this mode of thought the "immanent" mode. To think immanently is to think in the way one thinks when one stands within a theory: one holds the world (or something in the world) in one's cognitive gaze and says something about it (attributes some property or relation to some objects in it). The "bearers" of truth, on this view, are *immanent* thoughts.[16]

But immanent thoughts, by themselves, are not sufficient for truth. To raise the question of truth for a given immanent thought we have to go beyond 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 ask whether it is true. I will call the mode of thought required for such a move the "transcendent" mode of thought. It is important to note that there is nothing magical or mysterious about the transcendent standpoint, as understood here. This standpoint is not a Godly standpoint (Putnam 1981). It is a perfectly human standpoint, where one paradigmatic example of such a standpoint is a Tarskian meta-language (Tarski 1933). There is nothing super-human or magical about a meta-language, yet it is quite powerful relative to the "object language" it transcends. It has in view both the object-language and the world the object-language is directed at, whereas the object language has only the latter in view.[17]

Immanence and transcendence, however, are still not sufficient for truth. From a transcendent standpoint we can ask many questions about immanent thoughts, not just questions of truth. (For example, we can ask whether an immanent thought describes the world using long or short sentences, whether it describes it in a direct or in an indirect way, etc.). Truth requires a *normative* mode of thought. It requires a standpoint 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 such questions we need a normative mode of thought.[18]

The view that truth requires immanence, transcendence, and normativity I call the "fundamental principle of truth".

2.5 Normativity

Truth is commonly viewed as a *property* of thoughts. The fundamental principle of truth says that it is a *normative* property of thoughts. But truth, as I see it, is more than just a normative property of thoughts, it is a central *norm* of thoughts. It is a norm that reflects our aspiration to know the world as it in fact is. As such, it guides and constrains our theorizing. It says that our theories should say of a that it has the property P *only if* it in fact has the property P, that they should say something has the property P *only if* in fact at least one thing has the property P (the property P is in fact not empty), and so on. On this understanding, a thought has the *property* of truth iff it satisfies the *norm* of truth.

[16] "Thought", here, is used as a general term that can stand for truth-bearers of various kinds: thoughts proper, beliefs, sentences, propositions, utterances, judgments, cognitions, theories, etc. For the purpose of the present paper, there is no need to single out a particular type of thought as a truth bearer.

[17] I should note, however, 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 the truth predicate as constructed in *stages*.

[18] 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". That is the reason we attribute a truth value to such thoughts as well.

My view that the role of truth in knowledge is first and foremost that of a norm is shared by, e.g., Engel who says (2001, 38) that "we can, and must, say that truth is a **norm** of belief, and that most of our epistemic norms are grounded in this one". Other philosophers who emphasize the normativity of truth are Dummett (1959), Wright (1992), Price (1998), and Lynch (2004). On the other hand, Horwich (2016) rejects this view: "*Is* TRUTH *a* ***normative*** *concept?*... [M]y answer is **'no'**". This is not surprising. The view that truth is a central a norm of knowledge, stands in sharp contrast with Horwich's (1990/8) deflationist view that the only reason we need a concept of truth is a purely technical linguistic reason. According to this view, the truth predicate is a device we can use to make assertions that without it we might have difficulty making. Thus, instead of having to make infinitely many assertions, one can make the single assertion that *every sentence of the form 'P or not P' is true*, and instead of making theoretical assertions on topics one is not in command of, one can make the assertion that a certain theory is true (say: "Einstein's theory of gravitation is true").[19] This purely technical, linguistic, role is very different from the substantial epistemic role that the theory developed here assigns to truth.

But do we really need a norm of *truth*? 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 truth. If, on the other hand, we were incapable of correcting any of our cognitions, then a norm of truth would be of no use to us. But being the cognitive creatures we are—fallible on the one hand, capable of correcting ourselves (in many cases) on the other—a standard (norm) of truth is something we need and can make use of.

One might, however, acknowledge that humans need some norm(s) of correctness for their theories yet deny that such norms must include a norm of *truth*. Could we not make do with, say, norms of *evidence* or *justification*? Some philosophers seem to give a positive answer to this question. Thus Davidson (1999, 461) says: "I do not think it adds anything to say that truth is a goal, of science or anything else. We do not aim at truth, but at honest **justification**". And Rorty (1995, 287) says: "The pattern truth makes is, in fact, indistinguishable from the pattern that **justification**... makes—so... there seems no occasion to look for obedience to an *additional* norm [besides justification, namely] the commandment to seek the truth." In my view, these claims are incorrect. Justification, by itself, is an empty norm. Justification is always relative to a goal or to another norm. A vegetarian can justify his vegetarianism on grounds of taste, or on ground of health, or on moral grounds, or on environmental grounds, and these grounds are very different from each other. Similarly, a scientist can justify her theory based on different grounds, say, based on purely pragmatic grounds, or based on aesthetic grounds, or based on grounds of utility, or based on grounds of truth. She can even justify her theory based on some combination of these grounds, but she cannot justify it in a vacuum. My claim is that truth is one of the central norms that underlie scientific evidence and justification, a norm that reflects one of our central epistemic goals (though not the only one).

[19] It should be noted that deflationists such as Horwich do not claim that the truth predicate is the only device of this kind. For example, we can state the law of excluded middle without resorting to truth by using substitutional quantification. This suggests that truth, on their view, might in principle be dispensable.

2.6 Dynamic Correspondence

If truth is a norm that says our theories should aim at describing the world as it in fact is, then truth is not a coherence norm or a pragmatic norm. Truth is a correspondence norm in a broad sense of "correspondence", the sense of requiring a substantial and systematic connection between would-be units of knowledge and their targets in the world. But since our cognitive connections to the world are inevitably affected by our cognitive make up and the methods we devise for cognitively reaching the world, and since some aspects of the world are too complex for us to reach in a simple, straightforward manner, it is unreasonable to expect, let alone require, that such connections always take the simple form that traditional philosophy associates with correspondence, namely copy, mirror, or isomorphism. Rather, it is reasonable to surmise that the precise form (pattern) correspondence takes depends on, and changes with, the complexity of the targeted facets of the world, their accessibility to human cognition, and our resourcefulness in forging cognitive routes to them. I call this conception of correspondence "dynamic correspondence".

Dynamic correspondence (like basic realism) is both more demanding and more flexible than its traditional counterparts. It is more flexible in allowing us to connect our theories to their targets in the world in any way we can, provided such a way is sufficiently substantial and systematic. That is, it does not compromise the strength of the correspondence requirement. In a sense, it is just its flexibility that enables it to be more demanding than most other correspondence theories. Due to its flexibility, the correspondence requirement can be applied to *all* fields of knowledge, including fields that cannot be connected to the world by such simple and direct connections as copy or isomorphism. As a result, even fields that hitherto have been almost universally exempt from a robust correspondence requirement, such as logic, can be, and are, subjected to a substantial correspondence requirement by our theory.[20] Dynamic correspondence, however, 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 what pattern correspondence must take, but look and see what patterns are suitable for different fields of knowledge and under different circumstances.[21]

Is dynamic correspondence a *relativistic* theory, relativistic in the sense in which the theories advanced by Kuhn (1962/70), Bloor (1976/91), Feyerabend (1978), Latour and Woolgar (1986), and Rorty (1991) are relativistic, i.e., in a sense that conflicts with *robust realism*? No. The dynamic theory of truth discussed in this paper does not imply, let alone say, that there are no objective facts, that there are no facts external to us, that truth is subjective, i.e., that what is true varies from culture to culture or from person to person, that the norm of truth is the norm of abiding by the outlook of one's community, whatever it is, and so on. Saying that humans reach particular facets of the external world in ways that are partly dependent on the complexity of these facets and on our cognitive resources does not conflict with realism.

A simple analogy: In order to measure distance between two points we need to use some measurement system. The fact that "The distance between *a* and *b* is *n*" may be true when

[20] Logic, under this requirement, must aim at judgments of logical validity that are factually correct, i.e., correct in a strong correspondence sense. It is not enough that a logical theory endorse inferences that *appear* to transmit truth from premises to conclusion; our theory requires that it only endorse inferences that *actually* do so.

[21] "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".

measured in inches and false when measured in centimeters does not threaten realism. It does not make reality itself relative to culture. On a different level, many natural parameters are relative to others in the sense of being dependent on them. For example, mass, distance, and simultaneity are relative to frames of reference. But such relativity is factual and objective. Furthermore, such relativity does not affect the *laws* of nature. Whereas mass changes from one frame of reference to another, the laws of nature do not: the laws of nature are the same in all (inertial) reference frames. Spelling out the analogy, we may say: (a) The possibility of multiple patterns or "routes" of correspondence does not, by itself, compromise the objectivity of correspondence. (b) Whereas patterns of correspondence might vary from field to field, the general principle ("law") of correspondence does not. This principle says that truth consists in a substantial connection between a thought (theory, statement) and its target in the world, a connection that explains why, and establishes that, what it says about its target is in fact the case. The permanence of this principle and the objectivity of the routes (patterns) of correspondence safeguard the realistic character of the dynamic theory of truth.

Is the dynamic theory of truth a pluralistic theory? In the current literature, there are two different conceptions of pluralism with respect to truth, the one more radical, the other more moderate. According to the more radical conception, some "laws" (main principles) of truth may differ radically from field to field: truth in some fields may be based on correspondence, in others on coherence. This view is held by, e.g., Wright (1992) and Lynch (2004). According to the more moderate conception, the same principles are at work in all fields; only the patterns they exhibit might vary. It is not the case that in some fields truth based on correspondence and in others on coherence principles. The major principles are the same; only their applications vary. The patterns their applications take are simpler in some fields, more complex in others; more direct in some, more circuitous in others. Two proponents of this view are Horgan (2001) and Sher (2004, 2015, 2016a). The theory of truth presented in this paper is pluralistic in the second, more moderate sense. To see what this amounts to, it would be worthwhile to consider an example. Since this example will be used as a background for our discussion of truth in science, it would be better to select it from outside science.[22]

2.7 Example: Mathematical Truth as Dynamic Correspondence

To see dynamic correspondence at work, let us turn to mathematics. Mathematics is a discipline where correspondence is both attractive and problematic. On the one hand, the "unreasonable effectiveness" of mathematics in science (Wigner 1960) together with practicing mathematicians' prevalent realism (see, e.g., Russell 1919; Gödel 1944) suggest that mathematical truth is due to correspondence with certain facets of the world, namely, formal or mathematical facets. On the other hand, philosophically, this view is widely thought to be highly problematic. In particular, the commitment to mathematical objects which is often regarded as a pre-requisite for (or, alternatively, a result of) mathematical correspondence, leads to Platonism, which is a highly problematic position, challenged by

[22] More specifically: (1) since we are setting the ground for a discussion of science, we don't want to make too many assumptions about science, so it's better to have an example from another field, and (2) the dynamic structure of truth can take several forms, and the form we focus on in speaking about science is different from the one we focus on in relation to, say, mathematics. By bringing the mathematical example, we will highlight this plurality.

many philosophers. Benacerraf (1965, 1973), for example, challenges both the identity of mathematical objects on the Platonistic view and their cognitive accessibility to humans.

The contemporary alternative to mathematical Platonism is radical empiricism or naturalism, in the spirit of Mill (1843), Carnap (1939, 1950), Quine (1969), and Field (1980), which favors psycho-physical, indispensabilist, conventional-pragmatic, and/or fictionalist conceptions of mathematics, most of which sever the connection between mathematical truth and the world. These views, too, are widely considered to be inadequate and/or problematic.[23]

Dynamic correspondence enables us to overcome some of these problems and avoid others. Its treatment of mathematics has roots in Aristotle (4th Century BC, especially as interpreted by Lear 1982), Frege (1884, especially as interpreted by Hodes 1984), Quine (1955), and the mathematical structuralists (Resnik 1997; Shapiro 1997). But it also has significant differences with all the above.[24]

Dynamic correspondence combines a realistic view of mathematical truth (a Fregean theme) with a broad, common-sensical view of the world that is neither Platonistic nor radically empiricist. This it does by focusing on *mathematical features* (*properties*, *relations*) rather than on *mathematical objects* (an Aristotelian approach). What we need for mathematical realism, according to this view, is the reality of mathematical features of objects (in general), not the existence of a special, elusive type of object in our world (or in some other world). By avoiding mathematical objects, dynamic correspondence avoids Platonism; by recognizing the abstract character of mathematical properties and relations (features), it avoids narrow empiricism. Unlike radical empiricism/naturalism, it is able to account for the abstract yet factual nature of mathematics, and unlike Platonism, it does not disconnect mathematics from the world we live in. Dynamic correspondence leaves the question of abstract objects (as distinct from abstract features—properties and relation) an open question, and it is not dependent on a particular answer to this question. For more on this approach see Sher (2004, 2015, 2016a). Here, my goal is to present this approach as a potentially fruitful *example* of dynamic correspondence, one that could have analogs in science.

Dynamic correspondence approaches the question of mathematical truth from a different perspective than most other theories. It is common to approach mathematical truth by starting with language. Both in natural language and in the language of professional mathematics, mathematical objects are denoted by individual terms. For example, numbers are denoted by individual constants, namely, the numerals "0", "1", "2",..., and fall within the range of individual variables—"x", "y", "z" ("x_1", "x_2",...). Now, it is a principle of traditional semantics—which I will call "the syntax-ontology parity principle"—that if an expression of our language is of level n—0 being the level of individual constants and variables, 1 the level of predicates of individuals, 2 the level of predicates of 1st-level predicates, and so on—then the object it denotes in the world is also of level n—0 being the level of individuals in the world, 1 the level of properties (relations) of individuals, 2 the level of properties (relations) of 1st-level properties (relations), and so on. Thus, according to traditional correspondence, numbers, as the denotations of singular terms (individual constants and variables), are individuals.

It is individuals such as numbers that are at the center of many philosophical objections to mathematical correspondence. What evidence do we have of their existence in the

---

[23] For general overviews of contemporary approaches to the philosophy of mathematics and their criticisms, see, e.g., Horsten (2016) and Paseau (2013).

[24] The similarities will become clear as we go along. For differences, see fn. 32 below.

world? And if they do exist, how can we, humans, cognitively access them? What are their identity conditions? Etc. The dynamic conception of truth avoids these problems. The syntax-ontology parity principle is a *static* principle. It says that our language, with its largely fixed and unchanging syntactic structure, determines what objects our theories are committed to. This, with certain refinements, is Quine's principle of ontological commitment. It implies that if we have no ground for believing that the objects the syntax of our mathematical theories commits us to actually exist, then we have no ground for believing these theories are true, at least in the correspondence sense.

The *dynamic-correspondence* theorist rejects this static principle. Language, she recognizes, is one of our most important cognitive assets, 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 so on. Moreover, those aspects of language that play a central role in Quine's principle of ontological commitment—namely, the syntactic status of linguistic expressions—are especially rigid and difficult to adjust. For that reason, the theory of truth should not use language, and in particular syntax, as a touchstone. The connection between correct theories of the world and what they study in the world need not be uniquely determined by the syntactic status of the 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 theories and the world, i.e., correspondence truth in the strong, yet flexible sense of dynamic correspondence.

Partly for this reason, the dynamic theory of truth focuses not on language but on the world. As a dynamic truth theorist, I begin by asking: What facets of the world, if any, do mathematical theories seek to study? Or, better yet: Are there any facets of the world such that to gain knowledge of these facets we need a mathematical theory (or a theory like it)? And I formulate these questions 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 speak to a broad readership, I start with objects that are largely uncontroversial: planets, humans, trees, etc.[25] My initial question is: Do these uncontroversial objects have *formal, or mathematical, features*?[26]

Now, I think 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 a certain number of planets—say, eight—i.e., that the 1st-level physical property "x is a planet in our solar system" has the 2nd-level *formal* property of having cardinality EIGHT (and not ONE MILLION or ZERO).[27] 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 some of its physical properties are obtained by *formal* operations upon other physical properties: the operations of complementation, union,

[25] Concerning the reality of abstract objects, as distinct from abstract properties/relations, I prefer to leave this question open here.

[26] I use "formal" for the kind of features that, intuitively, mathematical theories would study if they studied any features of objects in the world. For a precise characterization of formality, see Sher (2013, 2016a).

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

intersection, 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., cardinality properties by laws of cardinality—and we, as epistemically ambitious beings, seek to know these laws. What form will such knowledge take? It is reasonable to surmise that it will take the form of a mathematical theory, such as arithmetic or set theory. This theory will be correct or incorrect about formal features and their laws, and in this sense its truth will be based on correspondence.

Our intuitive considerations suggest that formal features are for the most part of ontological level 2 rather than 0 (a Fregean theme). Cardinalities, for example, are not individuals but properties of properties of individuals. Yet our current theories of cardinalities use individual terms (numerals and 1st-order variables—level 0) to theorize about them. Does this necessarily render either the current theories or our intuitive observations incorrect? No. Numerals and individual variables can be connected to cardinality properties systematically, yet indirectly. Since (or to the extent that) 2nd-level cardinality properties are real, 1st-order arithmetic and set theory can, in principle, be true about them in a robust, if indirect, correspondence sense. Thus, arithmetical truths, on this view, correspond to laws (regularities) governing the behavior of finite cardinalities (general laws of addition, multiplication, order, etc. and their instances).[28]

But why would humans use 1st-order theories to describe the laws governing 2nd-level properties? 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, properties of properties, their properties, and so on, but are good at figuring out things when we think 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.[29]

This provides a positive answer to both our questions: it is quite clear that there could be a good psychological explanation of why humans might use 1st-order theories to study higher-level properties, and it is even clearer that it is possible to construct a correct 1st-level model of such properties. One way to do this is to use the mental operation of *positing* (a Quinean theme). Exercising epistemic freedom, we can construct a 1st-level structure of *posited* individuals. The posited individuals will be systematically connected both to

[28] More generally, mathematical truths correspond to laws governing formal (mathematical) properties of objects and properties of properties of objects in the world (whether physical or abstract—the view is not dependent on a particular ontological theory and is compatible with both moderate empiricism and moderate non-empiricism). By "laws" I understand regularities with a strong modal force. For further discussion and explanation see Sher (2016a).

[29] It is in this sense that we talk about, say, 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.

individual terms in our language and to higher-level properties 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*.[30] 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* yields the (posited) individual *seven* when applied to the pair of (posited) individuals *two* and *five*. Step 2: The latter is the case iff in the world, the DISJOINT UNION of TWO and FIVE is SEVEN. Finally, exercising transcendence, we explain how our 1st-order mathematical theories are true of *real* cardinalities.

Now, it is easy to see that this account offers a solution to Benacerraf's identity problem. The identity problem is the problem of determining which object is, say, the number *two*. Is *two* the set $\{\{\emptyset\}\}$, as Zermelo's theory says, or is it $\{\emptyset,\{\emptyset\}\}$, as Von Neumann's theory says? How does one tell? On our account the problem does not arise: $\{\{\emptyset\}\}$ and $\{\emptyset,\{\emptyset\}\}$ are different posits representing the cardinality property TWO. They belong to different, yet equally good models of cardinality properties (analogy to mathematical structuralism). TWO is not identical to either of these posits, but it is adequately represented by both.

When combined with our foundational holism, this account also solves Benacerraf's cognitive-access problem. The existence of mathematical individuals (such as numbers) is often thought to require a Platonic reality, separate from the physical reality in which we live. The question of cognitive access is then the question of how we, physical beings, have cognitive access to a non-physical, Platonic, reality. On our account this problem does not arise. Cardinality properties are properties that objects have right here, in our world, and we do not need trans-world cognition in order to find out about them. We find out the cardinality of planets as well as the laws governing these cardinalities (and cardinalities in general) by using our ordinary cognitive resources, both sensory and intellectual, in various combinations.[31,32]

It should be clear that the pattern of mathematical correspondence discussed above is not relative in any sense that conflicts with a robust, yet flexible, realism. To be connected to the *world as it in fact is* by a *complex* or *indirect* correspondence relation is, in principle, to be *as much*, and *as robustly*, connected to *it* as by a simpler, or more direct, correspondence relation.

In science, reference and correspondence are also likely to be dynamic, possibly in connection with abstract entities. Here, however, I would like to explore a different dynamic, one associated with scientific change. Some of my questions are: Does scientific change support our dynamic conception of truth? What does it teach us about the dynamics of truth that we do not know already? Once we have answered these questions, we can proceed to ask whether the dynamic view of truth might help us to address philosophical problems that arise from scientific change, e.g., the pessimistic-induction problem.

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

[31] Here our paradigm of "figuring out" is especially apt.

[32] I have pointed out similarities to Aristotle, Frege, Quine, and the mathematical structuralists. Differences concern many aspects of Aristotelian metaphysics, Frege's insistence that numbers are objects after all, Quine's radical empiricism, and the structuralists emphasis on structures of objects rather than properties. For more details, see Sher (2015, 2016a).

3 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 to the existence of objects[33]—contribute to our propensity for error.

Second, the perceived threat of the pessimistic meta-induction to science testifies to the centrality (or at least perceived centrality) of *truth* for the scientific endeavor. Without truth (strongly understood as correspondence), the pessimistic meta-induction would not be as threatening (or as "pessimistic"), as it is often thought to be. In the twentieth century, much attention was paid to the role of pragmatic considerations, aesthetic advantages, power relations, spheres of influence, and so on, in the development of science (Kuhn 1957, 1962/70; Bloor 1976/91, and others). But the pessimistic induction 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 correctness (or approximate correctness).

Scientific change, with its pessimistic and optimistic ramifications, also supports our *normative* conception of truth. The pessimistic-induction argument emphasizes our fallibility, highlighting the *need* for a norm of truth; the optimistic induction emphasizes humans' ability to correct their errors, hence the *non-futility* of such a norm. In this respect, the pessimistic and optimistic inductions complement each other. Optimists have argued that the likelihood of error is 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 the likelihood of future improvements means that we will have better tools for detecting errors (including errors of kinds that have not been detected so far), so the likelihood of finding errors in today's theories also increases. This, together with discoveries leading to new, and more difficult questions about nature, 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, new things about the *dynamic nature* of truth. The view of truth delineated in this paper so far is already dynamic, but our attention was limited 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? Before answering these questions, let us emphasize once again the centrality of a balance between fixity and changeability, unity and diversity, for the dynamic theory of truth. In discussing disunity in science, Dyson (1988) says that science requires a *fruitful balance* between unity and diversity. I believe his counsel applies to philosophy as

[33] Thus, reasoning to the existence of a certain object or medium sometimes works, sometimes not. (It worked with Neptune but not with the aether.)

well, including the theory of truth. On the one hand, we want to preserve the hard core of truth and make it universal; on the other hand, we want it to take into account the changing circumstances of the human pursuit of truth. Accordingly, our theory distinguishes two parts of truth: a fixed, universal, and stable part, and a diversified and changeable part. The first part includes 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 includes the changing patterns of correspondence. These may change both in context and in time, but the requirement of a solid correspondence is the bedrock beneath these changes. This duality is especially clear in the *norm* of truth: The norm of truth combines a universal requirement of genuine correspondence with flexibility concerning the ways we may pursue and satisfy this requirement.

Variations in patterns of correspondence are manifested in various ways. Most naturally, they are manifested in differences in *truth conditions*. One type of such variations we have already seen: differences in the *structure* of truth conditions in different fields. Truth conditions in mathematics, for example, unlike everyday physics are, or may be, "composite". Such truth conditions proceed in two steps, one leading from language to posits and the other from posits to the world. Truth conditions in other fields may have a different structure.

The phenomenon of scientific change suggests another type of change: *temporal* change in truth conditions, or more precisely, temporal change in our "working" truth conditions. By "our 'working' truth conditions" I understand the conditions we use in practice to determine the truth value of our statements. Here we are interested in statements made in a scientific context, and in this context, our working truth-conditions are often affected by *changes in our scientific understanding of the world.* Take, for example, statements about *water*. Whereas long ago humans could characterize water only in phenomenological and simple observational terms (transparent, tasteless, evaporates when heated, etc.), today scientific developments enable us to characterize water in a more advanced and accurate way, namely, by specifying its molecular structure, H_2O. Water itself did not change in the course of time, and the actual conditions that have to hold in the world for a statement about water to be true did not change either. But our understanding of these conditions, what we take these conditions to be, 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 our working truth-conditions for such statements.

Another way in which working truth-conditions are important in science is by closing the gap between evidence/justification and truth. This gap is especially deep in advanced sciences, due to the prevalence of indirect justification procedures in those sciences. For example, to establish the existence of a certain subatomic particle, scientists devise intricate tests, the results of which are accessible to us through, say, our computer screen. Scientists might say that an experiment's results, as seen by certain pictures on the computer screen, (indirectly) justify 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 results (as visible on the computer's screen) justify (1). Now, for (2) to be true, there must be some specified connection between the conditions under which certain pictures appear on the screen and the conditions under which (1) is true. And to determine whether such a connection exists, we need to formulate working truth-conditions for (1). What we need are not disquotational truth-conditions. These say nothing that would help us determine whether the pictures on the screen are adequately connected to the reality of subatomic particles to justify (serve as evidence for) the claim that a certain subatomic particle exists. The working truth-conditions have to be informative. They have to provide information about the identity conditions of the particle (would-be particle) in question, about some of its distinctive features, interactions with other particles, and so on. In formulating these conditions we rely on our overall scientific knowledge of the world at this stage in the development of science, using what we already know (or believe we know) to determine (as best we can) whether a given attempt to extend our knowledge meets our norm of truth. In practice, the working 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).

Note that the temporal dynamics of truth in science satisfies our Dysonian requirement of a fruitful balance between fixity and changeability. Truth persists in being a highly demanding correspondence norm, but our working truth-conditions change due to advances in specific areas, improvements in our cognitive tools, and gains in our overall understanding of the world.

Note, too, that for truth conditions to be useful in the ways described above, they have to be substantive rather than deflationist. If truth conditions are merely disquotational, as deflationists say they are and should be, there is no significant difference either between truth conditions in different fields or between working truth-conditions in a given field at different times. To detect (and make use of) these differences, we have to go beyond a simple disquotational formulation of truth conditions to more complex and substantive formulations, ones that are capable of discerning variations in relevant parameters.

4 The Impact of Truth on our Understanding of Scientific Change

The study of scientific change in the last hundred years or so brought about several pessimistic arguments concerning truth in science. One of these is the pessimistic meta-induction, which we focus on here.[34] 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 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 underlie its perceived threat to realism and make it seem more consequential than it actually is. Some of these assumptions—e.g., assumptions concerning the goal of science, revolve around the idea of *THE truth*, an idea that distorts our understanding both of scientific realism and of the relevance of the pessimistic meta-induction for it.

Thinking in terms of "*THE truth*" is very natural for humans, given our cognitive make up and the grammar of our language. It is natural for those who use the aspirational

[34] Others include the underdetermination argument (Duhem 1906/14; Quine 1951), the incommensurability argument (Kuhn 1962/70), the unconceived-alternatives argument (Stanford 2006), and so on.

argument to support scientific realism to express their view by saying that *science aims at THE TRUTH,* and it is natural for those who advance the pessimistic-induction argument against scientific realism to say that *it is a mistake to think that contemporary science has reached THE TRUTH* or that *science will ever reach,* or even approximate, *THE TRUTH.* This way of thinking is made natural by our language, but language is sometimes a trap. (See discussion of language in Part 2.7 above.) Language, in Frege's words, "unavoidably" introduces "misconceptions" into our thinking, and it is one of our tasks (as philosophers) 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 in 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 (norm)* 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 (our future *working* truth-conditions).

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.[35] Throughout the development of science, we are responsible for doing our best 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 (our working truth-conditions at the time). But this is not all. It is also our responsibility to critically examine our working 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 today we subject them to what are, as far as we can see, appropriate truth requirements than if we accept or reject them at random or based on gut feelings or a whim? What is the point of the Large Hadron experiments if we do not connect their results with the conditions that, as far as we know today, have to hold in the world for particles such as the Higgs boson to exist, that is, with the best working truth-conditions that we can assign to claims like "Higgs boson particles exist" today? 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

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

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 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 out of 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 (our current working 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, and of its changeable aspects as well.[36]

Acknowledgements This paper is based on an invited talk at the 2nd International Conference of the German Society for Philosophy of Science, Düsseldorf, 2016. I would like to thank Gerhard Schurz, Holger Lyre, and David Hommen for inviting me to present this talk. I am also thankful to participants in the conference—especially Holger Lyre, Michela Massimi, Stathis Psillos, and Gerhard Schurz—for very helpful feedback. Finally, I would like to thank the participants in my 2015 seminar on this topic and two anonymous referees of the paper for very helpful comments.

References

Aristotle (4 BC). *Metaphysics. Basic works.* New York: Random House, 1941.
Benacerraf, P. (1965). What numbers could not be. *Philosophical Review, 74,* 47–73.
Benacerraf, P. (1973). Mathematical truth. *Journal of Philosophy, 70,* 661–679.
Bloor, D. (1976/91). *Knowledge and social imagery.* Chicago: Chicago University Press.
BonJour, L. (1985). *The structure of empirical knowledge.* Cambridge: Harvard.
Boyd, R. (1990). Realism, approximate truth, and philosophical method. In C. W. Savage (Ed.), *Scientific theories* (pp. 350–391). Minneapolis: Minnesota University Press.

[36] 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 wait for another occasion.

Carnap, R. (1939). Foundations of logic and mathematics. In O. Neurath et al. (Eds.), *International encyclopedia of unified science* (Vol. I, No. 3). Chicago: Chicago University Press.

Carnap, R. (1950). Empiricism, semantics, and ontology. In *Meaning and necessity* (pp. 205–221). Chicago: Chicago University Press, 1956.

Chakravarty, A. (2017). Scientific realism. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/entries/scientific-realism/.

Davidson, D. (1999). Reply to Pascal Engel. In I. Hahn (Ed.), *The philosophy of Donald Davidson* (pp. 460–461). La Salle: Open Court.

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: Princeton University Press.

Dummett, M. (1959). Truth. In Dummett 1978 (pp. 1–24).

Dummett, M. (1973). The significance of Quine's indeterminacy thesis. In Dummett 1978 (pp. 375–419).

Dummett, M. (1973/81). *Frege*. New York: Harper & Row.

Dummett, M. (1978). *Truth and other enigmas*. Cambridge: Harvard University Press.

Dyson, F. (1988). *Infinite in all directions*. New York: Harper & Row.

Engel, P. (2001). Is truth a norm? In P. Kotatko & P. Pagin (Eds.), *Interpreting Davidson* (Vol. 3, pp. 37–51). Stanford: CSLI.

Feyerabend, P. (1978). *Against method*. London: New Left Books.

Field, H. (1980). *Science without numbers*. Oxford: Basil Blackwell.

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

Frege, G. (1879). Begriffsschrift. In J. van Heijenoort (Ed.), *From Frege to Gödel* (pp. 5–82). Cambridge: Harvard University Press, 1967.

Frege, G. (1884). *The foundations of arithmetic*. Evanston: Northwestern University Press, 1968.

Glymour, C. (1980). *Theory & evidence*. Princeton: Princeton University Press.

Gödel, K. (1931). On formally undecidable propositions of *principia mathematica* and related systems I. In *Collected works* (Vol. I, pp. 145–195). New York: Oxford University Press, 1986.

Gödel, K. (1944). Russell's mathematical logic. In *Collected works* (Vol. II, pp. 119–141). New York: Oxford University Press, 1990.

Goodman, N., & Quine, W. V. (1947). Steps toward a constructive nominalism. *Journal of Symbolic Logic, 12*, 105–122.

Haack, S. (1993). *Evidence and inquiry*. Oxford: Blackwell.

Hodes, H. (1984). Logicism and the ontological commitments of arithmetic. *Journal of Philosophy, 81*, 123–149.

Horgan, T. (2001). Contextual semantics and metaphysical realism: Truth as indirect correspondence. In M. Lynch (Ed.), *The nature of truth* (pp. 67–95). Cambridge: MIT Press.

Horsten, L. (2016). Philosophy of mathematics. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/win2016/entries/philosophy-mathematics/.

Horwich, P. (1990/8). *Truth*. Oxford: Oxford University Press.

Horwich, P. (2016). Is truth a normative concept? *Synthese*. doi:10.1007/s11229-016-1208-8.

Hume, D. (1739–40). *A treatise of human nature*. Oxford: Oxford University Press, 2000.

Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. Oxford: Oxford University Press.

Kripke, S. (1975). Outline of a theory of truth. *Journal of Philosophy, 72*, 690–716.

Kuhn, T. S. (1957). *The Copernican revolution: Planetary astronomy in the development of western thought*. Cambridge: Harvard University Press.

Kuhn, T. S. (1962/70). *The structure of scientific revolutions*. Chicago: Chicago University Press.

Latour, B., & Woolgar, S. (1986). *Laboratory life: The construction of scientific facts*. Princeton: Princeton University Press.

Laudan, L. (1981). A confutation of convergent realism. *Philosophy of Science, 48*, 19–49.

Laudan, L. (1990). *Science and relativism: Some key controversies in the philosophy of science*. Chicago: Chicago University Press.

Lear, J. (1982). Aristotle's philosophy of mathematics. *Philosophical Review, 91*, 161–192.

Lynch, M. (2004). *True to life*. Cambridge: MIT Press.

Lyons, T. D. (2005). Toward a purely axiological scientific realism. *Erkenntnis, 63*, 167–204.

Mill, J. S. (1843). *A system of logic*. London: Longmans, 1961.

Miller, A. (2014). Realism. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/win2016/entries/realism/.

Niiniluoto, I. (2015). Scientific progress. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/sum2015/entries/scientific-progress/.

Paseau, A. (2013). Naturalism in the philosophy of mathematics. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/win2016/entries/naturalism-mathematics/.

Peirce, C. S. (1878). How to make our ideas clear? *Popular Science Monthly, 12,* 286–302.

Price, H. (1998). Three norms of assertibility, or how the MOA became extinct. *Philosophical Perspectives, 12,* 241–254.

Psillos, S. (1996). Scientific realism and the 'pessimistic induction'. *Philosophy of Science, 63,* S306–S314.

Psillos, S. (1999). *Scientific realism: How science tracks truth.* New York: Routledge.

Putnam, H. (1975). What is mathematical truth? In *Mathematics, matter and method* (2nd ed.). Cambridge: Cambridge University Press, 1979.

Putnam, H. (1981). *Reason, truth and history.* Cambridge: Cambridge University Press.

Quine, W. V. (1951). Two dogmas of empiricism. In *From a logical point of view* (pp. 20–46). Cambridge: Harvard University Press, 1980.

Quine, W. V. (1955). Posits & reality. In *The ways of paradox and other essays* (pp. 246–254). Cambridge: Harvard University Press, 1976.

Quine, W. V. (1969). Epistemology naturalized. In *Ontological relativity and other essays* (pp. 69–90). New York: Columbia University Press.

Quine, W. V. (1970/86). *Philosophy of logic.* Cambridge: Harvard University Press.

Rawls, J. (1971). *A theory of justice.* Cambridge: Harvard University Press.

Resnik, M. D. (1997). *Mathematics as a science of patterns.* Oxford: Oxford University Press.

Rorty, R. (1991). *Objectivity, relativism, and truth: Philosophical papers* (Vol. 1). Cambridge: Cambridge University Press.

Rorty, R. (1995). Is truth a goal of inquiry? Davidson vs. Wright. *Philosophical Quarterly, 45,* 281–300.

Roush, S. (2010). Optimism about the pessimistic induction. In P. D. Magnus & J. Busch (Eds.), *New waves in philosophy of science.* New York: Palgrave Macmillan. http://philpapers.org/rec/MAGNWI.

Russell, B. (1919). *Introduction to mathematical philosophy.* New York: Simon & Schuster, 1971.

Sen, A. (2002). *Rationality and freedom.* Cambridge: Harvard University Press.

Shapiro, S. (1997). *Philosophy of mathematics.* Oxford: Oxford University Press.

Sher, G. (2004). In search of a substantive theory of truth. *The Journal of Philosophy, 101,* 5–36.

Sher, G. (2013). The foundational problem of logic. *The Bulletin of Symbolic Logic, 19,* 145–198.

Sher, G. (2015). Truth as composite correspondence. In T. Achourioti et al. (Eds.), *Unifying the philosophy of truth* (pp. 191–210). Dordrecht: Springer.

Sher, G. (2016a). *Epistemic friction: An essay on knowledge, truth, and logic.* Oxford: Oxford University Press.

Sher, G. (2016b). Substantivism about truth. *Philosophy Compass. 11,* 818–828.

Stanford, P. K. (2006). *Exceeding our grasp: Science, history, and the problem of unconceived alternatives.* Oxford: Oxford University Press.

Tarski, A. (1933). The concept of truth in formalized languages. In *Logic, semantics, metamathematics* (pp. 152–278). Indianapolis: Hackett, 1983.

van Fraassen, B. (1980). *The scientific image.* Oxford: Oxford University Press.

Wigner, E. H. (1960). The unreasonable effectiveness of mathematics in the natural sciences. *Communications on Pure and Applied Mathematics, 13,* 1–14.

Wittgenstein, L. (1921). *Tractatus logico-philosophicus.* London: Routledge & Kegan Paul, 1961.

Wright, C. (1992). *Truth and objectivity.* Cambridge: Harvard University Press.