

What use the Realist for a Theory of Truth?

Corey Dethier

[[Pre-print; forthcoming in *Philosophy of Science*.]]

Abstract

Recently, Hasok Chang has argued that we should adopt a pragmatic account of truth on which truth is constituted by practical success. Central to his argument is the claim that truth in this sense can be pursued and that other senses of truth can't be. In this paper, I present a dilemma for Chang's view. Either practical success constitutes truth in a way that implies infallibility—a conclusion that I'll argue is untenable—or the argument from pursuit fails. I end by suggesting that scientific realism alone commits us to very little with respect to truth.

0 Introduction

In *Realism for Realistic People* (Chang 2022), Hasok Chang centers his argument against “standard” scientific realism on the correspondence theory of truth. Which is surprising. As a matter of sociological fact, debates about truth and scientific realism rarely intersect; most work on the latter doesn't commit one way or another on the former (cf. Chakravartty 2018). It's also easy to find work by scientific realists—and not just fellow travelers to Chang like Ronald Giere (2006), but more traditional realists like Jody Azzouni (2000)—who explicitly argue against correspondence theories. Of course, the charge that “standard” scientific realism must in some sense presuppose a correspondence theory is not original: Arthur Fine (1986) alleged the same thing in the 1980s. Even if we accept Fine's argument—and there may be reasons not to (cf. Asay 2017; Musgrave 1989)—“standard” scientific realism is a very different beast today than it was in 1986. The connection seems dubious to me, to say the least.

For now, though, put aside Chang’s arguments against “standard” scientific realism—what about the alternative that he offers? On Chang’s preferred account of truth, successful integration into our practices is *constitutive* of truth: truth *just is* a matter of the right kind of practical coherence. His main argument for this conclusion is that if truth is to be a useful concept, it must be something that can be pursued, and traditional accounts of truth in which there is some sort of gap between it and success according to our criteria for recognizing it make it impossible to ever “verify” (Chang 2022, 200) or confirm that we’ve achieved the truth.

In this paper, I present a dilemma for Chang’s view. Either success according to our criteria is sufficient for truth in a way that makes us infallible—a conclusion that I’ll argue is untenable—or there’s no substantive sense in which Chang’s “truth-by-operational-coherence” can be pursued that doesn’t equally apply to (e.g.) truth-by-correspondence. The problem, in essence, is that the “gap” between truth and success according to our criteria is what buys us fallibility: we can always be wrong because no matter how thoroughly we’ve checked a conclusion, it’s always possible that we’ve made a mistake. If our practices are constitutive of the truth in a way that truly bridges this gap, then there’s an important sense in which we can’t be wrong. But if we maintain fallibility, then truth à la Chang is no more verifiable than truth as it is understood by more traditional accounts.

I begin by clearing some ground regarding theories of truth (§1). I then present the dilemma (§2) and pursue both horns (§§3,4) before considering a couple potential rejoinders that Chang (or any sympathetic fellow-traveler) could make (§5). I end by briefly returning to the connection—or lack thereof—between realism and the correspondence theory of truth (§6).

1 Theories of truth

There are two different things that are termed “theories of truth” in the literature (Soames 1984). On the one hand, there are theories of the semantics of the word “true” and its cognates. These are theories about what the word “true” contributes to the sentence “What John said isn’t true” or “It’s true that it will rain tomorrow.” On the other hand, there are theories of the metaphysics of truths. These are accounts of what kinds of things can be truth-bearers (sentences? propositions?) and what it is that “makes” a sentence (/ proposition) true—what property or properties distinguish a true sentence (/ proposition)

from a false one?¹ My contention in this first section is that insofar as Chang has a disagreement with traditional theories of truth, his disagreement should primarily be understood as a disagreement regarding the metaphysics of truths.

To see why, it will be helpful to have a (brief) outline of what I take to be the cutting edge in the semantics of “true.” The key insight of the various forms of deflationism inspired by Frank Ramsey (1927) and Alfred Tarski (1936, 1944) is that a sentence like “It’s true that it will rain tomorrow” has the same semantic content (and/or meaning and/or truth-conditions) as the sentence “It will rain tomorrow.” That is, from the point of view of the semantics, *this* use of the word “true” is redundant. The same is not true of “true” as it appears in “What John said isn’t true”: sentences like “Not what John said” are strictly-speaking not grammatical. Moreover, while it’s easy enough to negate what John said if we know that it was something like “It will rain tomorrow,” we’re often in a position where we don’t know, or can’t restate, the content of whatever it is we’re appending the truth predicate to. It’s perfectly coherent for me to say—perhaps on the basis of warranted mistrust in John—“I don’t know what John just said, but whatever it is, it isn’t true.” So “true” is not always redundant from the perspective of semantics: while there are other ways that we could achieve the same effect (Azzouni 2006), our assertoric practices require some tool for blindly assenting or dissenting to statements.

To be clear: the account of the semantics of “true” being offered here is not about assertion, at least not directly. The point is that “true” and its cognates have a particular semantic role. Because “true” plays the role that it does, there is no semantic difference between the assertion “It’s true that it will rain tomorrow” and the assertion “It will rain tomorrow.” But this is a consequence of the theory of the semantics of “true,” and not definitional of it. After all, the theory also says that there is no semantic difference between the questions “Is it true that it will rain tomorrow?” and “Will it rain tomorrow?” or between the commands “Make it rain tomorrow!” and “Make it true that it rains tomorrow!”² That is: “true” plays the same role in other speech acts as it

¹There are of course other things that one could mean by a theory of truth. For instance: one might have in mind a theory of the pragmatics of the word “true.” Rightly or wrongly, the literature has not tended to focus on such alternatives.

²This way of putting the point is (to my knowledge) original to this essay. Though they shouldn’t, deflationists, minimalists, and fellow travelers often do write as though there’s nothing more to truth than its role in the speech act of assertion and/or the communication or expression of acceptance. But it would be uncharitable to read even someone like Crispin Wright (1992) as saying that truth is exhausted by its role in *assertion*, as opposed to its role in our *assertoric practices more broadly*—practices that include other speech acts, such as questions, commands, pleadings, etc.

plays in assertion. In each case, it acts as a device of *semantic descent* (Azzouni 2018): by attaching “is true” to a name or description of a sentence, we can use that name or description to assert, or question, or command the sentence that it refers to. Assertion—or, really, our assertoric practices understood broadly—explains why we have a device of semantic descent like “true,” but plays no essential role in defining or characterizing its function.

Notice that we have said *nothing* about the metaphysics of truths. Deflationist theories of truth do take the semantic account just laid out as a jumping off point; broadly speaking, they hold that there is nothing more to be said about truth than the semantic theory. One way of putting this is that if we take the class of sentences that the truth predicate is correctly applied to, and the class of sentences that it isn’t correctly applied to, there is no metaphysical property that distinguishes the former from the latter. There is no property in virtue of which true sentences are true and false sentences false. Those who endorse correspondence or coherence theories disagree. The correspondence theory says that true sentences are true in virtue of corresponding with the world in the right sort of way. The coherence theory that they are true in virtue of cohering in the right sort of way. Pluralism about truths holds that there are multiple distinct properties in virtue of which sentences are true. Notice that this is a distinct position from pluralism about “true,” which holds that there are multiple meanings of “true” with distinct semantics. One can think that there is only one semantics for “true” (at least in a properly regimented language) while simultaneously thinking that the properties that distinguish true sentences from false differ from domain to domain.³

So disagreements over semantics come apart from disagreements over the metaphysics. This isn’t a new point; on the contrary, I take it that the basic outline of the above argument can be found in the literature at least as early as Scott Soames’ “What is a Theory of Truth?” (Soames 1984). I’m belaboring it for two reasons. First: I think Chang goes astray in his criticisms of the realist in large part because he conflates realist views about semantics with metaphysical views about transcendental reality that I’m skeptical most realists actually endorse; I’ll justify this claim more as we proceed.⁴ Second, and more immediately relevant: it’s important we locate the disagreement between Chang and traditional theories of truth in the right place; so far as I

³As Nicolaj Pedersen and Cory Wright note, these two kinds of pluralism are run together “not infrequently” (Pedersen and Wright 2018).

⁴As one reviewer helpfully noted, *precisely* this same complaint has been lodged against prior attempts to connect realism with a positive view of truth. See, e.g., Jamin Asay (2017) or Alan Musgrave (1989). This fact raises a number of interesting sociological questions that unfortunately fall outside the scope of this essay.

can see, Chang’s main concern is the metaphysics, not the semantics.

The motivation for this second conclusion is that the basic semantic theory outlined above is compatible not just with realism and Chang-style Kantianism but with anti-realism, constructive empiricism, and much else besides. Consider Chang’s view in particular. Chang argues for the Kantian view that there is some important sense in which the world that we encounter is (in Chang’s terms) *mind-framed*, meaning roughly that how we divide it up and conceptualize it is dependent on us.⁵ That’s entirely compatible with the semantic point made above: “It’s true that it will rain tomorrow” has the same content as “It will rain tomorrow” but it is not as though one of these is mind-framed and the other is not. On the contrary, both rely on the same way of dividing up and conceptualizing the world. We’ve said nothing here about how true sentences “hook up” to the noumenal world or in “in virtue of what” those sentences are true; we’ve merely made a point about what it means to say that they’re true.⁶ So if mind-framing is going to move us to adopt a new view about truth, it’s going to have to be a view about the metaphysics of truths, not about the semantics of “true.”

Indeed, the semantic point is compatible with Chang’s more skeptical position that *we cannot even justifiably infer* conclusions about what the world is like absent the framing that we provide. I call this a skeptical position not simply because it involves the denial of knowledge of some domain but because it follows from mind-framing only if we grant a paradigmatic skeptical move.⁷ Both the skeptic and the non-skeptic agree that there is some sense in which what Robert Pasnau (2017) terms “epistemic perfection” is impossible in some respect or domain. We certainly cannot describe the world without depending on concepts that are (to some degree at least) up to us. The non-skeptic points to all the ways that we can do better and worse in that domain regardless. In this case, we can study our own psychology and practices to gain insight into what we add and what we don’t. The skeptic—and here David Hume (1740) on induction is the paradigm—argues that better and worse collapse absent perfection: without a foundation of non-mind-framed results, we cannot pos-

⁵Notably, contrary to what Chang (2022, 73) suggests, mainstream scientific realists have distinguished something like this kind of mind-dependence from the mind-dependence that they care to deny for decades now; for citations, see Chakravartty (2017).

⁶Compare Azzouni (2000, 132, 2025, 45), John Collins (2025, 113–14), Musgrave (1989, 392), or Howard Stein (1989), all of whom make essentially the same point in slightly different contexts.

⁷It’s also dialectically important that the move being made here is a skeptical one. Chang presents himself as offering a version of realism suitable for “realistic people.” If the position reduces to *mere* skepticism—if what we get is actually realism for skeptics—then he’s failed in that task.

sibly claim to have warrant, let alone knowledge, for claims about the world beyond our framing. Skepticism in this sense goes beyond mere mind-framing: we need the additional Humean move to get from one to the other.

To my eyes, Chang endorses this Humean move on multiple occasions (e.g., Chang 2022, 75–77, 250–51). I have little sympathy for the move in any context, including this one; I think we have good reasons for rejecting it. I won’t pursue them here, however. My point is simply that the semantic picture offered above doesn’t tell us anything about the metaphysics (or epistemology, for that matter): to say “It’s true that it will rain tomorrow” just is to say that “It will rain tomorrow.” Is science capable of telling us that it will rain tomorrow? What makes it true that it will rain tomorrow, but false that it is raining now? How do concepts like “rain” and “tomorrow” hook up the noumenal world? Perhaps we can never have any evidence for a view on these subjects. Perhaps we cannot even meaningfully ask some of these questions. Regardless, the semantics doesn’t answer them. And insofar as Chang and (e.g.) the correspondence theorist disagree about questions like these, they’re interested in the metaphysics of truths, not the semantics of “true.”

2 The dilemma

So how does Chang depart from traditional theories of truth? The centerpiece of his position is the claim that the “operational coherence” of our practices or activities is constitutive of the truth of propositions that are relied on within those practices:

The operational coherence of activities relying on a true proposition are not *consequences* or *indications* of its truth, from which we may infer the truth. Rather, operational coherence is *constitutive* of truth. (Chang 2022, 170)

Being relied on in coherent activities is *constitutive* of the truth-by-operational-coherence of a proposition, rather than being an effect or evidence of truth. (Chang 2022, 180)

(Note that the emphasis here is Chang’s own.) What is operational coherence? It’s a pragmatic notion: a practice exhibits operational coherence insofar as it “well-designed” for achieving its aim (Chang 2022, 40, 44). Naturally, operational coherence comes in degrees and thus truth does as well (Chang 2022, 171, 186–92).

Now, operational coherence is partly a matter of what *we* do, but it not wholly so. Chang (2022, 189–90) is aware of the worry that operational coher-

ence might be seen as too “up to us” and assures us that it isn’t: “reality offers ‘resistance’” (Chang 2022, 72). No matter how coherent perfectly circular orbits might seem to me, the world does not cooperate and a practice based on this assumption will fail to be coherent except in very narrow applications or to a very small degree. So we can be wrong about operational coherence in the sense that our expectations about it can be violated (Chang 2022, 185).

But Chang repeatedly suggests that we can’t be wrong about it in the sense of thinking that we’ve achieved operational coherence but later learning that we have not:

The only sensible way to shield ourselves from the pessimistic induction from the history of science is to grant that all good theories are true, and remain true, within their own scope. And with my notion of truth-by-operational-coherence, such liberal granting of truth is not only possible but very natural. (Chang 2022, 187)

When we say ‘It may seem as if P were true in these circumstances, but P is actually not true’, what else can we be meaningfully asserting, other than that P will fail to be true in some other circumstances? (Chang 2022, 190)

Elsewhere, he’s insistent that “Truth-by-operational-coherence is something achievable and verifiable in practice” (Chang 2022, 200); it is something that “can actually be tracked” (Chang 2022, 250). If a theory is operationally coherent to some degree, that means that that theory is at least partially true: “Any positive test-outcomes partially *constitute* the theory’s truth” (Chang 2022, 182). So: we might learn that we’re wrong about the *scope* of a theory, but we do not and cannot learn that it is not true within that scope. Nor can we discover that a proposition lacks previously assigned truth-by-operational-coherence by recasting our prior reliance on it: since “reliance [is] a matter of *actual* need *within* a given system of practice” (Chang 2022, 194), the best we can hope to do is argue that a system of practice is no longer useful—not that the propositions that it relies on are not true.⁸

The picture that emerges is pragmatist in a very clear sense. Chang identifies truth with a particular kind of (practical) success—operational coherence—

⁸I will briefly register that I’m not happy with the notion of “reliance” that Chang is working with here (for his elucidation, see Chang 2022, 183–84). Saying that the practical success of our practices relies on a proposition strikes me as uncomfortably close to the claim that “positive test-outcomes are *caused* by the theory’s truth,” a claim that Chang (2022, 182) mocks. In both cases, we appear to be attributing causal powers to our representations that are really held by the worldly affairs that they represent. I’ll put these worries aside in what follows, however.

and while success in this sense comes in degrees, we can tell that we’ve achieved it. (Can we tell *the extent to which* we’ve achieved it? That’s unclear, and will matter later.) To reiterate, though, this is an *identification*: operational coherence is not a sign of truth, or a criterion for identifying when we have good reason for thinking that something is true—it’s genuinely constitutive of the truth.

It’s important that this last feature is not an accidental aspect of Chang’s view, something that is easily amended. On the contrary, the constitutive relationship between truth and operational coherence is central to the argument for his account. According to Chang, we should prefer truth-by-operational-coherence over competitors such as the Peircean ideal of truth as the limit of inquiry or “Truth about transcendent Reality” because these competitors are impossible to pursue. They are impossible to pursue because we do not have “criteria of judgment” (Chang 2022, 217) for distinguishing truth in these senses from their respective falsehoods (Chang 2022, 198, 250); without such criteria, we have no way of aiming at or trying to achieve the goal (Chang 2022, 2). In other words: if we don’t have criteria capable of distinguishing truth from falsity, then truth is not a useful notion for us. If there’s a constitutive relationship between operational coherence and truth, then we have a clear criterion for determining whether a claim is true. If there is not such a constitutive relationship, then truth à la Chang is in no better of a position than alternative accounts according to which the various practices that Chang groups under “operational coherence” are signs of the truth.

We’re now in a position to sketch the dilemma for Chang’s view. As we’ve seen, Chang holds that operational coherence is constitutive of the truth. At face value, this suggests that success according to our practices is sufficient for the truth, meaning that there’s an important sense in which we can’t be wrong—a consequence that Chang explicitly and repeatedly endorses. But if “our practices” in the aggregate can’t be wrong, there’s worryingly little room for the kind of everyday mistakes that bedevil each individual practice. To make the point concrete, Chang seems to be committed to the view that a claim or hypothesis is true to some degree even if it is only “coherent” because we misread the instruments. That’s an undesirable result.

So the first horn involves accepting the apparent implication that we cannot be wrong about truth because it is constituted by operational coherence. The second horn involves rejecting this implication: there is some sense in which we are fallible in our identification of operational coherence and thus truth. To be clear, accepting the second horn requires ignoring quite a few of the passages cited above. But it does allow us to avoid the bad result that everyday errors and mistakes are true. The downside is that the argument for Chang’s

position then falls apart: there’s no sense in which we can “verify” truth-by-operational-coherence that isn’t also a sense in which we can “verify” truth-by-correspondence, and thus no sense in which the former can be pursued but the latter can’t be. In both cases, all we have are signs of the truth.

The dilemma is quite tight: either our criteria for recognizing truth are infallible or they are not. I’ll consider possible routes for rejecting it in section 5. To my mind, however, the more promising line of response lies in showing that one of the horns is not as bad as I have suggested. In the next two sections, therefore, I treat each of them in more depth.

3 The first horn

On Chang’s account, operational coherence is constitutive of truth. He appears to explicitly endorse the implication that our criteria for identifying the truth are infallible; worse, his argument in favor of his account turns on precisely this feature. The first horn of the dilemma involves accepting this implication. While infallibility is bad at face value, perhaps it’s not as bad as it initially appears.

To start, it will be useful to be more precise about what I mean by (in)fallibility. My concern here is with the (in)fallibility of our criteria for recognizing truth. That is: is success according to the criteria for recognizing truth *sufficient* for truth? If our criteria are fallible, then success according to those criteria is not sufficient for truth. If our criteria are infallible, then success according to those criteria is sufficient for truth.

Roughly this same distinction can be found going by a variety of names in the literature. Azzouni (2000) associates this kind of fallibilism with the “criterion transcendence” of truth. Michael Dummett (in, e.g., Dummett 1991) says that realists endorse the “recognition transcendence” of truth, a phrase that has been widely adopted in the literature commenting on his work (see, e.g., Shieh 2018). Hilary Putnam (1977, 485) names a similar commitment by saying that truth is “radically non-epistemic.” Tim Button (2013, 10) attributes to realists the claim that “Even an ideal theory might be radically false.” Tim Maudlin (2015, 491) endorses the view in describing a situation in which a theory is “operationally ideal, but still false.”⁹

⁹Chang insists that Maudlin’s position in particular is “just the opposite of what many scientific realists seek” (Chang 2022, 85). So far as I can tell—it’s unclear who he has in mind—this is simply wrong. Realists do commonly endorse the view that our criteria for recognizing truths are *good indicators* of the truth—that usually or most of the time, those indicators will accurately pick out the truth. But—as Chang (2022, 170, 180) is very careful

Care is needed here, however. For one thing, both Dummett and the literature commenting on him are primarily concerned with the necessity of our criteria for truth. By contrast, my concern in this essay is with sufficiency. Further, while the formulations of Button and Maudlin strike me as broadly in-line with the principle I want to defend, we need to be careful (as always) with the notion of “ideal.” As I’m using these terms, the question is whether *human* criteria are fallible: whether a theory that passed every test that we could possibly subject it to, over the whole course of our species’ existence in the universe, might still end up being wrong. An agent with epistemic powers fundamentally beyond our own would have access to different criteria. For example: an agent who could observe every raven that has ever and will ever exist might be in a position to employ criteria for determining whether all ravens are black that would be sufficient to guarantee the truth of the claim even if our human methods are fallible. It’s possible that such an agent’s criteria would be infallible in the relevant sense—I don’t want to commit one way or the other.

I do want to commit to the fallibility of human criteria, however. The argument is straightforward and does not rely—so far as I can see—on realist assumptions.¹⁰ Our criteria for recognizing truth ultimately bottom out in everyday activities of measurement, experiment, observation, etc. Since success by light of these activities is clearly not sufficient for truth—we can measure inaccurately, an experiment can go wrong for any number of reasons—neither are our criteria. Formally:

- (P1) Our criteria for recognizing truth involve nothing “over and above” everyday activities of measurement, experiment, observation, etc.
 - (P2) Each of these everyday activities is fallible: we can think that we’ve performed them correctly and still get the wrong answer.
- ∴ (C) Our criteria for recognizing the truth are fallible.

to insist with respect to his own view—being an indicator of truth and being constitutive of it are very different things.

¹⁰Arguments in favor of principles in the vicinity of fallibility can be broken into three camps. First, these principles are often taken to be intuitive or commonsensical. Second, there are arguments—particularly by Azzouni (2000)—that it is a descriptive fact that our *word* “true” transcends our criteria for applying it, and (further) that there are good reasons for having a concept like “true” that is criterion transcendent. Finally, particularly within the Dummett-Putnam tradition, debates over “recognition transcendence” are held to turn on deep claims about what it means to grasp or understand a concept. My argument fits into none of these categories cleanly. Though I would not be surprised to discover that it has predecessors within the literature, I have not been able to find any.

Both premises deserve comment.

First, while it is certainly unquestionable that our criteria for identifying truths rely on everyday practices of measurement (etc.), it is less obvious that those criteria in no way exceed or go beyond those everyday practices—should they do so, the argument just given would not succeed. One might think, for example, that while measurement and observation are fallible, there’s something special about measurements and observations that “hang together” in the right sort of way. This might actually have been William Whewell’s view (see Whewell 1840 [2014]), for example, and one can imagine Chang or a sympathetic fellow traveler defending it as well. Of course we can be mistaken in measuring a quantity or recording an observation. But when that quantity or observation productively fits into a whole system of practice (“makes sense,” in Chang’s terms), the only sense in which it could be false is the thin and easily-jettisoned one in which the same quantity or observation does not fit into a different system of practice (cf. Chang 2022, 190).

Ironically, Chang himself provides us with grounds for rejecting this picture:

James warned against the reification of truth as with other ‘words in *th*’ (wealth, health, strength and truth), quoting an epigram from Lessing: ‘How come it’s exactly the richest people in the world who have the most money?’ There is, of course, nothing to explain here, because having lots of money is exactly what being wealthy *means* (James 1907 [1975], 105–6). (Chang 2022, 182)

What lesson should we take from James’ warning?¹¹ Regardless of which of the standard theories of the metaphysics of truths one prefers, a true claim and a false one are indistinguishable when we examine the claims themselves. Truth is not “worn on the sleeve,” as it were, but must be established by examining the subject of the claim—there’s much more to be said here about what “examining” might mean in different contexts (e.g., in the setting of mathematical claims), but this much ought to be uncontroversial. And though I am no James scholar, I take it that this is at least part of what he means in the passage that Chang quotes later on the same page: “truth *happens* to

¹¹Chang wants to draw a lesson about the standard realist’s claim that “positive test-outcomes are *caused* by the theory’s truth.” Certainly, that claim is erroneous, exhibiting a kind of use-mention error (Azzouni forthcoming): what causes the results is what’s going on in the world, not the theory that describes the world (Dethier 2025, cf.). But it’s at best an uncharitable characterization of the realist position, which is that the theory makes accurate predictions because it accurately describes its subject (see Chakravartty 2017). The causal claim is between one world-theory relation and another, not between the theory and the world itself.

an idea. It *becomes* true, is *made* true by events.” There is nothing about an “idea” itself that tells us whether it is true or not—or at least there isn’t the domains commonly studied in the (empirical) sciences.

But it is exactly this kind of reification that we are required to endorse if we are to pry apart everyday activities of measurement and observation from our criteria for recognizing truth. If there is something special and different about our criteria—something that makes them infallible where our everyday activities are not—it must be because those criteria “hook onto” truth in a way that our everyday activities don’t. But it is not as though ordinary activities of measurement and observation involve a certain amount of observing and then we do an additional, final and infallible, observation to push us over the edge. That is, there’s nothing we *do* in recognizing that something is true that goes beyond our fallible everyday activities (cf. Misak 2007, 69). So if our criteria are to “go beyond” these activities in the sense required, it must be because they are sensitive to the truth itself. To be clear: not to features of the world of the sort that we measure and observe, because those are all handled by everyday activities, but to the brute truthiness of the idea. And this brute truthiness that goes beyond the features of the world we measure and observe is precisely like a concept of wealth that goes beyond being rich.

So: rejecting the first premise of the argument requires us to reify the concept of truth, and to suppose that we have some independent and infallible means of hooking on to this concept, in a way that looks epistemically suspect (to say the least). What of the second premise, the claim that our everyday activities are in fact fallible? Presumably, the opponent of fallibility with respect to our criteria could simply bite the bullet here and hold that (e.g.) when you measure twice and get two different results, both are (always!) true or at least have some degree of truth. Indeed, Chang indicates a willingness to go down this path: on his view, operational coherence and truth are always specific to a domain of application; measuring twice to get two different results might yield truth in a “pitifully tiny” domain (Chang 2022, 190), but not in any useful one. So Chang could bite the bullet here—describe both measurements as true—and still say that there’s *something* undesirable about an erroneous measurement.

Of course, the bar is a little higher than simply saying that there’s something wrong with erroneous measurements. Many scientific practices exist precisely because we want to avoid the errors that arise from everyday mistakes. And science would be much worse off without them; these are important practices that should not be jettisoned based solely on a philosophical theory of truth. Which means that anyone who rejects the second premise must offer an explanation of why we have these practices and (moreover) why they are a

good idea—an explanation that cannot appeal to truth, because they hold that both measurements are true to at least some degree. And Chang has resources here as well. For instance: those measurements that you or I would call “true” or “accurate” will typically generalize and project into new domains better than those we would call “false.” So Chang can say that all of our measurements are true (in their respective domains), but sometimes we think that a measurement (or theory or ...) will generalize to larger domains than it does because of a simple error or mistake. Since all else equal we want theories that generalize to larger domains, we need practices for avoiding everyday errors.

I’m optimistic that an account along these lines could be developed in a satisfactory way. The problem I have with it is that it results in a response to the second premise of the argument that seems merely terminological (cf. Stein 1989). For one thing, it makes Chang’s disagreement with “standard” realists merely verbal: in this terminology, he’s talking about truth (as operational coherence), they’re using the word “true” to talk about propositions that are fully generalizable. Since Chang still needs the concept of generalizability to explain and motivate our practices for avoiding errors, he hasn’t offered an account that disagrees with the realist in any substantive way. For another, “true” and associated terms like “false” and “wrong” are useful precisely because on traditional accounts our criteria are insufficient to guarantee truth. It’s useful to have such a concept so that we can capture fallibility; indeed, so that we can capture the fact that even our best-confirmed, most operationally coherent theories might (in Chang’s terminology) turn out to apply to much more restricted domains than we thought because (e.g.) we made an unlucky series of everyday errors. If we make the move to using “fully generalizable” for this role, however, then the account doesn’t really deny the value of the traditional concept of “truth,” it just uses a different word for that concept.

In other words, biting the bullet on the second premise leads to a dilemma of its own. On one horn is giving up on the idea of mistakes and the practices for avoiding them—an option I take it no one endorses. On the other horn is an alternative account of truth that merely changes the subject: neither offers a substantive disagreement with traditional accounts nor fulfils the central role that the concept of truth plays in our current practices. Neither of these options are promising.

Of course, sometimes “changing the subject” is quite valuable; there may also be other roles for the concept of truth besides the one picked out. I’ll consider these possibilities a bit more in section 5. For now, however, I’ll assume that this second horn is unsuccessful, that one constraint on any account of “truth” is that it captures the possibility of everyday errors such as mismeasuring the length of a piece of string. Insofar as that’s a legitimate constraint,

our everyday activities are fallible. As we saw above, our criteria for recognizing truths are nothing over-and-above these everyday activities (or success within them). And so our criteria must be fallible as well: meeting the criteria is insufficient for truth. Which means that we should reject any account—like Chang’s, if the first horn is adopted—on which our practices are constitutive of truth in such a way that said practices cannot be wrong.

Recall the point I made earlier: the argument just given relies only on fairly mundane observations about our criteria for recognizing truth—namely, that these criteria rely on nothing over and above everyday activities of measurement (etc.) that we know go wrong all the time.¹² That’s not to say that there are no bullets to bite—there are, unappetizing though they may be. Rather, the point is that the premises here don’t seem to have anything to do with debates about realism. Indeed, at least on Cheryl Misak’s reading, the central point here is endorsed by no less of a pragmatist than Charles Saunders Peirce (Misak 2022, 279–80), who argued that it was a mistake to try and define truth in terms of success within our activities; the better practice is to treat this success as evidence. Which is to say, to reject Chang’s view that such operational coherence is constitutive of truth.

4 The second horn

Recall: On Chang’s account, operational coherence is constitutive of truth. He appears to explicitly endorse the implication that our criteria for identifying the truth are infallible; worse, his argument in favor of his account turns on precisely this feature. The second horn of the dilemma involves rejecting that commitment. At face value, that undermines the motivation for preferring his view. But perhaps the motivation can be saved—perhaps the argument doesn’t really rely on infallibility.

Chang seems to suggest something along these lines in his response to an argument of Bertrand Russell’s:

it is alleged that ‘assessing the usefulness (etc.) of a belief is no more clear-cut than assessing its truth’, which means that pragmatism fails to turn truth into a concept any more accessible and useful than the correspondence notion. Russell says: ‘it is so often harder to determine whether a belief is useful than whether it is

¹²It is interesting that debates over the parallel necessity claim—success according to our criteria is necessary for truth—often turn on heavy-duty premises about the relationship between epistemology and meaning. The contrast suggests an important asymmetry between the two claims that may be worth pursuing further.

true’ (Russell 1910, p. 138, see also p. 135). This objection draws a false equivalence between the in-principle inaccessibility of the correspondence between idea and world, and the detailed vagueness and multi-dimensionality of very tangible successes. ... The messiness of determining what is useful is only a problem if one has a monist notion of usefulness. (Chang 2022, 203)

On the view suggested here, there’s an important distinction between truth-by-operational-coherence and more traditional alternatives. In both cases, our criteria are fallible: we can never be 100% certain that our current commitments are true. But the former is still (in some sense) accessible in practice while the latter is permanently out of reach in principle. *If* we can draw this distinction, we can maintain the central argument that Chang offers for preferring truth-by-operational coherence without requiring the kind of infallibility that got us into trouble in the last section.

The question is: on what grounds would we draw this distinction? Going off of the arguments that Chang offers against alternative accounts of truth in *Realism for Realistic People*, the idea is a Kantian one: Chang’s opponents hold a view of truth in which it involves matching our words and ideas up with the “noumenal” realm, whereas Chang’s view is one on which we’ve abandoned all talk of the noumenal world to focus only on that which is actually available to us in experience. So we might, on this picture, make mistakes about the phenomenal world, but that world is available to us in a way that the noumenal one is not, and so we can at least offer an argument for truths about the former being pursuable in a way that truths about the latter are not. How might this argument go? Well, suppose that I’m reading too much into Chang’s claim that operational coherence “is something achievable and verifiable in practice” (Chang 2022, 200). Perhaps what Chang means here is just that we can verify operational coherence in the same everyday sense that we can verify that we have two hands. That is, we can verify the truth of various scientific claims—claims like “viruses evolve,” “global warming is primarily anthropogenic,” “electrons carry a negative charge,” etc.—by examining the copious evidence in their favor. While one can imagine that this evidence is misleading because we’re (e.g.) being tricked by a capricious deity, that’s not the sense of verification—or, crucially, the sense of *truth*—that he’s concerned with. Whereas, and by contrast, traditional accounts of truth are concerned with whether we really, not just phenomenally but noumenally, have two hands.

Let’s grant that there’s a coherent and well-motivated picture here. I, at least, am not convinced that this is the right way to view traditional accounts

of truth. That is: it is not obvious to me that “correspondence” should generally be interpreted as meaning “correspondence to the noumenal realm,” and Chang offers no evidence in support of this attribution (cf. Wray and Hoyningen-Huene 2025). Moreover, there are very good reasons for thinking that even if this is what some people mean by correspondence, it’s not what “standard” scientific realists are committed to. After all, if the “standard” realist endorses any view, it’s the view that science provides us with evidence about the way things are; as Anjan Chakravartty opens the SEP article on the subject, “Scientific realism is a positive epistemic attitude toward the content of our best theories and models, recommending belief in both observable and unobservable aspects of the world described by the sciences” (Chakravartty 2017). So the realist thinks that truth—as they understand it—is something that we have in-principle access to. We might not be able to (fully) verify whether we have truth or not, but we have methods that provide us with evidence and reasons for thinking that some things are true and that some things are not. If the realist also endorses the Kantian idea of truth-as-correspondence-to-an-inaccessible-noumenal-realm that Chang presses upon them, then there is a very obvious contradiction in their views. The natural conclusion—at least absent textual evidence to the contrary—is that Chang has misunderstood or misdescribed the “standard” realist position regarding truth. Given that *Realism for Realistic People* only engages with an actual text from a realist in *one* case—and, indeed, does so in service of dismissing them as nonrepresentative without argument (see note 9)—this is a disappointing but not unsurprising conclusion.

But let us suppose that at least some “standard” realists hold something like the Kantian view that Chang attacks. Would the argument just sketched actually save Chang’s argument? I think not. At the very least, Chang would need to show that what matters for avoiding the Kantian trap is the *constitutive* character of his account. So far as I can see, there’s no particular reason to think that’s true. That is, it seems perfectly coherent (to me) to grant Chang most of what he wants with respect to meaning and language, to agree with him that the world described by science is mind-framed, and yet to still think that whether our theories (etc.) are true is a matter of something like correspondence to (mind-framed) reality. On this view, the meaning of words like “force” and “mass” is something we contribute to the world, but once we fix the language, it’s still up to the world (rather than us) whether our sentences are true or not. We might think that we are making sense of the world and rendering it operationally coherent—we might do so for centuries, even millennia, or all of human history—only to discover that we were wrong.

Let me put it another way. If the sense of “verifying” that’s relevant to

Chang’s argument is just the everyday sense involved in verifying that I have two hands, I see no reason why the standard scientific realist cannot claim that truth is just as verifiable for them as it is for Chang. After all, if Chang’s account of truth allows for fallibility, then it must always be possible for us to be wrong. So truth must be “achievable and verifiable in practice” even though we can always be wrong. But then there’s no reason why something like correspondence with the facts cannot be verifiable and thus no reason why it is less pursuable than Chang’s pragmatic alternative. The argument has fallen apart.

All of which is to say that even if we grant Chang (a) that we can properly pursue truth only if we have criteria for verifying our possession of it, (b) that there’s some way of maintaining fallibility even when success within our practices is constitutive of truth, and (c) that opposing views typically endorse a kind of Kantian picture that requires correspondence (or whatever) with the noumena, the arguments he offers for preferring his position still only motivate rejecting said Kantian picture—they don’t actually motivate his account of truth. Given that I don’t think there’s really much of a reason to grant any of these commitments, I’m left thinking that this horn is no more promising than the last.

5 Rejecting the dilemma?

The last three sections have argued that Chang faces a dilemma. On his account, success according to our practices is constitutive of the truth; this would appear to imply that we can’t be wrong. Either he can accept this implication, which leads to the unappealing result that every measurement is accurate (at least with in its own domain), or he can reject it and in so doing undermine the primary argument for his position.

The dilemma is hard to reject—either success according to our practices is sufficient for truth or it isn’t. One response that Chang raises in relation to a similar complaint is that operational coherence is *hard* to achieve (Chang 2022, 189–90). But this won’t help. If my uncle is very wise, thoughtful, and considerate, there is good reason to think that his statements are a good indicator of the truth, but not in the slightest good reason to think that his statements are constitutive of the truth. Similarly, pointing out that a lot of work goes into achieving operational coherence might give us good reason to treat operational coherence as an indicator but does not allay concerns about treating it as constitutive. If truth is not under our control (and if we do not imagine ourselves omniscient), then it should be possible—not necessarily

likely, but possible—for a theory to achieve any degree of success according to our criteria and still be false. No matter how high you set the bar, in other words, the constitutive character of operational coherence is still a problem.

A slightly more promising response is to point out that on Chang’s view, truth is a matter of degree or, really, a qualitative non-binary notion (Chang 2022, 186–87). Now, on its own this is irrelevant, because the same arguments will apply at the level of degrees. Either our practices are constitutive of “how much” truth something has or they aren’t. If they are, then our criteria are infallible; if they aren’t, then the truth-by-operational-coherence has no leg up on its competitors.

But perhaps there is a way out to be found here. Chang might argue that we cannot be wrong about *whether* we have hit on some truth, but that we can be wrong about *the degree to which* something is true. Potentially, this would allow us to have our cake and eat it too: we avoid the worst aspects of infallibility by keeping the grades of truth fallible while giving truth-by-operational-coherence a leg up on competitors by making it “verifiable and achievable in practice.”

An obvious rebuttal is that this “solution” has only the advantages of theft over honest toil: who cares that we can be assured that our current theories have *some* degree of truth if we do not know at least roughly what degree of truth that is? Furthermore, I see no reason why alternative accounts of truth could not themselves adopt a similar graded notion to respond to Chang’s argument. There’s nothing about “correspondence” that requires it to be all or nothing, for example, and nothing about gradations that requires a constitutive relationship between truth and our criteria for recognizing it.

Equally, the response seems to fall prey to the concern I raised in section 3 that we’re merely redefining the word “true.” After all, everyone already agrees that science is shot through with idealizations and approximations that are useful despite being false in some important sense. What have we gained in our understanding of the phenomena by re-working the word “true” so that we need to say of these idealizations that they are true just not to the same degree or in the same contexts as some other theory or model? And, as indicated above, if we did re-define “true” in this way, we would need a new term (e.g., “fully generalizable”) that means just what “true” does now, because it’s extremely useful to have a concept for capturing what our criteria for success don’t guarantee.

Of course, changes in usage are not nothing, particularly when it comes to concepts as central and important as “true.”¹³ Chang or fellow travelers might

¹³In these next two paragraphs—and at the beginning of the next section—I’m granting

be perfectly right to argue that no matter how comfortable philosophers are with our current use of the word, we'd be better off using it in a different way. In this case, however, I see no reason to think that this redefinition of our concept would actually be helpful—that it would serve the project of “activist realism,” as Chang (2022, 209) puts it. For all we can tell, redefining truth in this way would serve only as ammunition for climate denialists, vaccine skeptics, and intelligent designers. This is a tension throughout the whole of *Realism for Realistic People*: arguments for the philosophical view—broadly, a kind of operationalist pluralism—are not the kinds of arguments demanded by the meta-philosophical view, namely that philosophy of science should help science be better. If we take the meta-philosophy seriously, then we need something like empirical studies to determine whether pluralism will actually serve to improve science.

In other words: it may well be a good idea to change the subject and discuss something else using our word “true.” But I don't see a compelling argument for that, yet, and in the meantime the bind presented by the dilemma is very tight. Given how Chang argues for his account of truth, he needs it to be infallibly accessible in the same way that alternative notions of truth are only fallibly accessible. But infallibility has disastrous consequences, in that it implies that we cannot possibly make mistakes. So if both horns are to be avoided, we need to somehow show that the sense of infallibility that is desirable is different from the sense of infallibility that I argued against in section 3. But any way of doing that—at least any way that I can see—is also going to open the door for the defenders of other accounts to make use of the same distinction in defending their preferred views against Chang's argument.

6 Realism as epistemology first and foremost

Where does this leave us? I have offered a dilemma for Chang: either the constitutive relationship between operational coherence and truth implies that our criteria for recognizing truth are infallible or it does not. If it does, the view should be rejected, because there are good reasons endorse fallibility. If it does not, the argument in support of it fails.

So what's gone wrong? In one sense, I think, the problem here is the idea that truth cannot be useful unless we have some sort of definitive criteria for distinguishing truth from falsity. *Definitive* being the key qualifier. Chang

for the sake of argument a number of fairly substantial assumptions concerning conceptual engineering, not the least of which is that we can just *decide* the meaning of our terms, which is (at best) an idealization (cf. Cappelen 2018).

repeatedly writes as though there are no advantages to being open to revising our criteria in the future, to changing practices not just because the old practice has ceased to be useful—or as useful as the new practice—but instead because we think the old practice rested on a mistake (see, e.g., Chang 2022, 194). I must admit that I find this viewpoint almost impossible to fathom. I find it extremely useful to be able to say that our criteria for recognizing truth may be wrong, to allow for future revisions. To say, for example, that Einstein proved that the observed motions of particles in a solution were not accurate proxies for their true motions, and thus that methods and criteria for identifying truths about Brownian motion that treated them as proxies rested on a mistake (see Smith and Seth 2020).

Section 3 can be read as my attempt to show why it’s useful without presupposing any of the realist (or monist) commitments that Chang rejects. In short: it’s extremely useful to be able to categorize something as minor as misreading an instrument as a genuine error or mistake. If I look out my window and accidentally take the Celsius gauge on my thermometer to be Fahrenheit, I haven’t invented a new system of practice; I’ve just done something wrong. Moreover, there’s a robustly epistemic sense in which systems of practices that reliably avoid such mistakes are better than those that don’t. But there’s no way to circumscribe the ramifications of these kinds of minor errors, no level of theory choice that is in principle free from their influence. So we can’t have both everyday errors and an account on which success by our criteria for recognizing truth is sufficient for truth, because it’s possible to succeed according to those criteria because of an everyday error.

All the above presumes that I’m right about the dilemma, of course—that presumption granted, the above is my best attempt to say where the positive view of *Realism for Realistic People* ultimately goes wrong. To be blunt, however, there is a much more immediate source of difficulties for *Realism for Realistic People*: the book simply does not engage with the views of those it critiques. As a consequence, the position that Chang attacks is—at best—twenty years behind what realists actually endorse or argue for. Frankly, I find the “standard” realist position as Chang characterizes it to be basically unrecognizable; Chang’s “standard” realist seems to have more in common with the yard signs that proclaim “In this house, we trust science.” than it does with paradigmatic realist figures like Chakravartty or Stathis Psillos. Unsurprisingly, the positive arguments suffer almost as much as the negative ones do.

I promised to return to the question of the relationship between scientific realism and the correspondence theory of truth. It’s possible that there was once a tight connection between these views; I suspect they’re correlated in

the profession. But it seems to me that given the way the realism debate has evolved since 1975, it's a mistake to treat them as a package. Why? For much of the last 50 years—arguably going as far back as the no-miracles argument (Putnam 1975), but certainly since the publication of *The Scientific Image* (van Fraassen 1980)—the primary doctrine of scientific realists, and the primary point of divergence from various anti-realist alternatives, has been epistemological. Quoting again from Chakravartty's SEP article:

Scientific realism is a positive epistemic attitude toward the content of our best theories and models, recommending belief in both observable and unobservable aspects of the world described by the sciences. (Chakravartty 2017)

For the scientific realist in this modern tradition, the metaphysics (and the semantics) follows from the epistemology: we should believe that there's a "mind-independent" external world (and that our terms really refer) because that's the best explanation of the success of science. So at best the correspondence theory of truth is a consequence of the main realist doctrine—the realist doesn't endorse a correspondence theory as part of their commitment to realism, but because *given the epistemological commitment involved in realism*, the success of science demands that they do.

In fact, though, I think the connection is even more tenuous than this. A (scientific) realist in this epistemological sense is someone who holds that the methods of the sciences provide us with the best means of answering empirical questions—e.g. is the earth warming?—regardless of whether those questions concern "unobservables." As a consequence of this commitment, the realist takes on various commitments regarding metaphysics and semantics: they believe that atoms, viruses, and the like exist and that these terms refer to atoms, viruses, and the like. But if they take on a commitment to the correspondence theory of truth or anything in the vicinity of the idea that the words at the end of the last sentence are not "mind-framed," I can't see how. It is perfectly coherent to hold that the methods of science provide us with the best means of answering the question "do atoms exist?" without taking a stand on whether that question is "really" about some sort of noumenal realm or even whether a true answer to it will be one that corresponds or coheres in the right sort of way.

To be sure, scientific realists might have other reasons for answering those questions in a particular way. The position I'm sketching is not *quietist*; it is simply one on which these questions can and do come apart. And it is important that they come apart, because while the realist rejects the idea that there is an epistemically relevant distinction to be drawn at the level

of observability, they may well hold that the laudatory empirical methods of the sciences are not well-suited to answering certain kinds of questions. The methods of science are simply not the right kind of methods for determining whether the arguments of Kant’s transcendental deduction are correct, for example. So while the realist is committed to saying certain sorts of things about *atoms* in virtue of their epistemological position, they are not committed to saying things about the *noumenal realm*.

Finally, while I will be the first to admit that one can be “more” of a realist than the realist that I have described, by (e.g.) also being a Lewis-style realist about possible worlds, I don’t really see how one can be more of a realist with respect to the epistemological dimension identified by Chakravartty. For though our realist need not believe in every claim or theory in the sciences—the claims of string theory may be worth suspending judgment on, for example—that’s because those claims don’t (yet?) meet the standards of the methods that they endorse. In other words, epistemology-first realism is not a tepid reclamation project. It is “stark, raving” realism (Railton [1986](#)), realism red in tooth and claw. So far as I can see, however, the only thing that such a realist has to say about truth is a commitment to fallibility that everyone else should share as well.

Acknowledgments

I’d like to thank Anjan Chakravartty, three reviewers, and the attendees of the Philosophy of Science: The Next Generation in Miami for helpful comments on an earlier version of this paper. I’d also like to thank the members of the MCPS Spring 2025 Center Discussion Group for reading *Realism for Realistic People* with me.

Funding

None to declare.

Declarations

None to declare.

References

- Asay, Jamin (2017). Realism and Theories of Truth. In: *The Routledge Handbook of Scientific Realism*. Ed. by Juha Saatsi. New York: Routledge: 383–93.
- Azzouni, Jody (2000). *Knowledge and Reference in Empirical Science*. London: Routledge.
- (2006). *Tracking Reason: Proof, Consequence, and Truth*. Oxford: Oxford University Press.
- (2018). Deflationist Truth. In: *Oxford Handbook of Truth*. Ed. by Michael Glanzberg. Oxford: Oxford University Press.
- (2025). *Challenging Knowledge: How We (Sometimes) Don't Know What We Think We Know*. Oxford: Oxford University Press.
- (forthcoming). Scientific Pluralism and the Language-Based Approach to Scientific Theories. In: *From Contradiction to Defectiveness to Pluralism in Science: Philosophical and Formal Analyses*. Ed. by María del Rosario Martínez-Ordaz and Otavio Bueno. Cham: Springer.
- Button, Tim (2013). *The Limits of Realism*. Cambridge: Cambridge University Press.
- Cappelen, Herman (2018). *Fixing Language: An Essay in Conceptual Engineering*. Oxford: Oxford University Press.
- Chakravartty, Anjan (2017). Scientific Realism. In: *Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. URL: <https://plato.stanford.edu/entries/scientific-realism/>.
- (2018). Truth and the Sciences. In: *Oxford Handbook of Truth*. Ed. by Michael Glanzberg. Oxford: Oxford University Press: 602–24.
- Chang, Hasok (2022). *Realism for Realistic People: A New Pragmatist Philosophy of Science*. Oxford: Oxford University Press.
- Collins, John (2025). World and Truth in Conflict. In: *The Question of Linguistic Idealism*. Ed. by Richard Gaskin. Oxford: Oxford University Press: 111–37.
- Dethier, Corey (2025). Use, Mention, and Tangles Around Truth. *Baltic International Yearbook of Cognition, Logic and Communication* 16: 1–12. DOI: [10.4148/1944-3676.1131](https://doi.org/10.4148/1944-3676.1131).
- Dummett, Michael (1991). *The Logical Basis of Metaphysics*. Cambridge: Harvard University Press.
- Fine, Arthur (1986). Unnatural Attitudes: Realist and Antirealist Attachments to Science. *Mind* 95.378: 149–77. DOI: [10.1093/mind/XCV.378.149](https://doi.org/10.1093/mind/XCV.378.149).
- Giere, Ronald N. (2006). *Scientific Perspectivism*. Chicago: University of Chicago Press.

- Hume, David (1740). *A Treatise of Human Nature: Being an Attempt to Introduce the Experimental Method of Reasoning into Moral Subjects*. Ed. by Amyas Merivale and Peter Millican. URL: <https://davidhume.org/texts/t/> (visited on 05/22/2025).
- James, William (1907 [1975]). *Pragmatism*. Cambridge: Harvard University Press.
- Maudlin, Tim (2015). Confessions of a Hardcore, Unsophisticated Metaphysical Realist. In: *The Philosophy of Hilary Putnam*. Ed. by Randall E. Auxier, Douglas R. Anderson, and Lewis Edwin Hahn. La Salle: Open Court: 487–501.
- Misak, Cheryl (2007). Pragmatism and Deflationism. In: *New Pragmatists*. Ed. by Cheryl Misak. Oxford: Oxford University Press: 68–90.
- (2022). A Sensible Pragmatist Conception of Truth. In: vol. 97. 3: 275–94. DOI: [10.1017/s0031819121000310](https://doi.org/10.1017/s0031819121000310).
- Musgrave, Alan (1989). NOA’s Ark—Fine for Realism. In: vol. 39. 157: 383–98. DOI: [10.2307/2219825](https://doi.org/10.2307/2219825).
- Pasnau, Robert (2017). *After Certainty: A History of Our Epistemic Ideals and Illusions*. Oxford: Oxford University Press.
- Pedersen, Nikolaj Jang Lee Linding and Cory Wright (2018). Pluralist Theories of Truth. In: *Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. URL: <https://plato.stanford.edu/entries/truth-pluralist/>.
- Putnam, Hilary (1975). *Mathematics, Matter and Method: Philosophical Papers, Volume I*. Cambridge: Cambridge University Press.
- (1977). Realism and Reason. *Proceedings and Addresses of the American Philosophical Association* 50.6: 483–98. DOI: [10.2307/3129784](https://doi.org/10.2307/3129784).
- Railton, Peter (1986). Moral Realism. *Philosophical Review* 95.2: 163–207. DOI: [10.2307/2185589](https://doi.org/10.2307/2185589).
- Ramsey, Frank P. (1927). Facts and Propositions. *Aristotelian Society Supplementary Volume* 7.1: 153–70. DOI: [10.1093/aristoteliansupp/7.1.153](https://doi.org/10.1093/aristoteliansupp/7.1.153).
- Russell, Bertrand (1910). William James’s Conception of Truth. In: *Philosophical Essays*. London: Longmans, Green, and Co.: 127–49.
- Shieh, Sanford (2018). Truth, Objectivity, and Realism. In: *Oxford Handbook of Truth*. Ed. by Michael Glanzberg. Oxford: Oxford University Press: 432–76.
- Smith, George E. and Raghav Seth (2020). *Brownian Motion and Molecular Reality: A Study in Theory-Mediated Measurement*. Oxford: Oxford University Press.
- Soames, Scott (1984). What is a Theory of Truth? *The Journal of Philosophy* 81.4: 411–29. DOI: [10.2307/2026307](https://doi.org/10.2307/2026307).

- Stein, Howard (1989). Yes, but ...: Some Skeptical Reflections on Realism and Anti-realism. *Dialectica* 43.1: 47–65. DOI: [10.1111/j.1746-8361.1989.tb00930.x](https://doi.org/10.1111/j.1746-8361.1989.tb00930.x).
- Tarski, Alfred (1936). The Concept of Truth in Formalized Languages. In: *Logic, Semantics, Metamathematics (1984)*. Ed. by John Corcoran. Trans. by Joseph H. Woodger. 2nd ed.: 152–278.
- (1944). The Semantic Conception of Truth and the Foundations of Semantics. *Philosophy and Phenomenological Research* 4.3: 341–76. DOI: [10.2307/2102968](https://doi.org/10.2307/2102968).
- van Fraassen, Bas C. (1980). *The Scientific Image*. Oxford: Oxford University Press.
- Whewell, William (1840 [2014]). *The Philosophy of the Inductive Sciences, Founded Upon Their History*. Cambridge: Cambridge University Press.
- Wray, K. Brad and Paul Hoyningen-Huene (2025). Discussion Note on Hasok Chang: Realism for Realistic People. *Journal for General Philosophy of Science* 56.2: 275–86. DOI: [0.1007/s10838-024-09704-6](https://doi.org/0.1007/s10838-024-09704-6).
- Wright, Crispin (1992). *Truth and Objectivity*. Cambridge: Harvard University Press.