ARTICLE

Realism and Observation: The View from Generative Grammar

Gabe Dupre

School of Social, Political, and Global Studies, Keele University, Newcastle, UK
Email: g.g.dupre@keele.ac.uk

Abstract

Standard proposals of scientific anti-realism assume that the methodology of a scientific research program can be endorsed without accepting its metaphysical commitments. I argue that the distinction between competence, the rules governing one's language faculty, and performance, or linguistic behavior, precludes this. Linguistic theories aim to describe competence, not performance, and so must be able to distinguish observations reflective of the former from those reflective of the latter. This classification of data makes sense only against the background of a psychologically realistic view of linguistic theory. So the very methodology of the science commits one to its realistic interpretation.

> If all you want is to be able to predict your experiences, the rational strategy is clear: Don't revise your theories, just *arrange to have fewer experiences*; close your eyes, put your fingers in your ears, and don't move.
>
> Fodor (1991, 202)

1. Introduction

Once perhaps the central topic in the "general" philosophy of science, the realism/anti-realism debate has fallen out of favor in the last few decades. One reason for this is skepticism about whether there is much profitable to say about science *in general.* Questions about the reliability of scientific methods in generating knowledge about the unobservable world have largely been superseded by narrower questions about the methodologies, epistemologies, and ontologies of specific sciences. The attitude toward the realism/anti-realism debate within the philosophy of particular sciences (physics, biology, cognitive science, economics, etc.) has become somewhat indifferent: figure out the successes and failures within a scientific field, make as much sense

as one can of the fundamental concepts and approaches adopted therein, and let the high-level debates about the reliability of science generally play out as they may.[1]

In this article, I attempt to connect these traditional concerns with this recent "particularist" stance, providing an argument for scientific realism based on the methodology of a particular science: generative linguistics. I shall argue that at least this research tradition is committed, by its very methodology, to a realistic interpretation. I shall do this by identifying a crucial assumption made by mainstream scientific anti-realists: detachability. Anti-realists are committed to viewing the metaphysical commitments of a science as *detachable* from its methodology, in that scientists could, in principle, deny the former while retaining the latter. I shall argue that one of the central methodological tools of generative linguistics, the competence/performance distinction (hereafter, CP), precludes detachability in this domain. In the remainder of the article, I shall detail how this lesson generalizes beyond generative linguistics.

2. Detachability: The methodology of anti-realism

One difficulty with the realist/anti-realist debate is that these labels are applied to a wide range of positions, and therefore arguments for/against one position may not apply to others grouped under the same banner. For this reason, it will be more profitable to simply stipulate the theses I shall be discussing rather than going through the elaborate and subtle distinctions in the literature between different strands of realism and anti-realism. I believe that what I say will apply widely to the views discussed in the literature because it targets claims that broadly characterize either side of the issue.[2] *Realism*, as I will use the term, involves two commitments:

Axiological realism: Scientific theorizing *aims at* producing true and justified theories of the world, including its unobservable features.

Epistemological realism: Science has been reasonably successful in meeting these aims. Scientific theories are often justified and even true.

Of course, much in these statements is vague or ambiguous. Although the statement of anti-realism presupposes some, perhaps flexible, boundary between the observable and unobservable, exactly where this boundary is found is debated. Likewise, truth is a notoriously contentious notion, but again, the anti-realist seems committed to some relatively robust notion of truth by their denial that we can reasonably view certain scientific claims as true. And gradable expressions like "reasonably successful" prompt arguments about where to draw a line. I don't wish to enter into these debates in this article. Further, the two realist theses are logically independent. One could think that science aimed at truth about unobservables without endorsing its prospects for achieving this goal (as argued by Lyons [2005]) or that science managed to accurately describe unobservables but did so strictly in the service of some other goal.

Despite these complexities in the uses of the term, hopefully, these statements indicate clearly enough what I mean by *realism*. The scientific realist is optimistic

---

[1] Schindler (2018) represents a notable exception to this trend.

[2] I take Psillos (2005) to be a canonical example of the realist position. I am ignoring the differences between the realism defended by Psillos and, for example, semirealism (Chakravartty 1998), critical realism (Bhaskar 2013), measured realism (Trout 1998), real realism (Kitcher 2001), and so forth.

about science as an epistemic enterprise: science is in the business of describing the world and is a good tool for this purpose. Of course, the realist accepts that evaluating the truth of specific scientific theories must be done on a case-by-case basis, and any given current theory could turn out to be false. What is denied is that there is an in-principle barrier to such theories being true and known to be true.

With this characterization of realism, there are two broad ways of being an anti-realist:

> *Instrumentalism:* Scientific theorizing does not aim at producing true and justified theories of the unobservable world.
>
> *Empiricism:* Scientific theories of the unobservable world are not justified.

As with *realism*, the terms *instrumentalism* and, even more so, *empiricism* are used to identify various different philosophical positions, but I will again use them stipulatively. Instrumentalism is the rejection of axiological realism, the denial that scientific theorizing is successful to the extent that it generates true and justified theories of the unobservable world. Empiricism is the rejection of epistemological realism, the denial that the methods of the sciences can confer justification on theories about the unobservable world.

Empiricism is broadly motivated by high-level epistemological worries. Whatever justification scientific theories have must come from empirical, observational contact with the world. Given this, the empiricist asks, how could such confirmation accrue to theories about entities/processes that we cannot observe? Speculation about what we cannot empirically engage with seems, from this perspective, to go beyond natural science and into the troublesome realms of metaphysics. These concerns are then bolstered by appeals to the spotty record of such attempts to go "beyond the data" in the history of science.[3]

These kinds of issues also motivate instrumentalism. Philosophers of science, even anti-realists, generally don't want to diminish the successes of science too much. But if the empiricist is right, whatever successes science has had don't consist of producing justified theories of the unobservable world. So if science isn't to look like a large-scale failure, it would be good if science wasn't aimed at producing such theories. The instrumentalist thus proposes exactly this: scientists shouldn't be embarrassed to have failed at the task of truly and justifiably describing the unobservable world because this was never their goal. In lieu of this, it is typically claimed that science aims at something like predictive accuracy. If a scientific theory can adequately describe what we can observe, including predicting what we *will* observe, then that is sufficient.[4] So although the two strands of anti-realism are, in principle, independent, they are typically put forward as a package: science can't provide justification for our beliefs in the unobservable, but it can justify claims about observables, even those we haven't yet observed, and that is all it aims at in the first place. This package view is most famously defended by van Fraassen (1980).

---

[3] See, for example, Laudan (1981) and Stanford (2006).
[4] One could, of course, propose various other goals—reducing human suffering, developing useful technology, and so forth. I will focus on predictive success because it has been most frequently suggested in the literature and seems to underwrite any other plausible goal.

This desire to avoid downplaying the successes of science, in addition to providing a central motivation for instrumentalism, seems to commit the anti-realist, at least tacitly, to the assumption of what I will call *detachability*:

Detachability: One can endorse the practices and methods of a given science without accepting the metaphysical commitments of the scientific theory appealed to.

Anti-realists want to differentiate themselves from those who simply dismiss the scientific project. Anti-vaxxers and climate-change deniers accept the empiricist claims, in their respective areas, that scientific theories about the unobservable world are unjustified. To differentiate themselves from this bad company, anti-realists accept that science has had successes, but they restrict the range within which such successes are to be found. In particular, they endorse scientific methods as suitable for conferring justification on *some* claims, provided that these claims are themselves empirically testable. To the extent that, say, viruses are unobservable, the empiricist and the anti-vaxxer agree that scientific claims about viruses are not justified. However, many statements about the effects of vaccination made by physicians and epidemiologists will be testable against our observations: we can *see* whether vaccinated people cough, sneeze, die, and so forth less frequently.

Anti-realism can thus avoid lapsing into outright denialism by insisting that science is highly successful, just within a limited sphere. The empiricist allows that the methods of science have developed so as to enable scientists to make highly accurate and reliable predictions. What's more, the instrumentalist claims this is all that science aims to do, and so there need be no shame in restricting our claims in these ways. It is only when we attempt to go beyond this, and make claims about the unobservable causes of the observable phenomena, that our justifications fail us and we leave behind the respectable scientific project. In this way, the scientific practices of using a theory, say, to make predictions, develop technology, and so forth can be endorsed as worthwhile and success-prone without accepting these theories' ontological commitments (in roughly the sense of Quine [1948]).

This rather tidy anti-realist picture thus seems to assume a reasonably neat boundary between science and ontology. The anti-realist portrays the realist as advancing a "two-step" procedure for justifying claims about unobservables. The first step involves accumulating observations and figuring out which theory is most compatible with them. Once this is done, the realist argues that we should believe whatever the most empirically successful theory says. The anti-realist declines to take this second step. In this picture, the real scientific work is done in the first step, and the metaphysical commitment made by the realist is extraneous and risky. This picture embodies detachability: the assumption that the metaphysics of science is some sort of "optional extra" to be accepted or not once the empirical work is done.

Something along these lines must be assumed by the anti-realist in order to distinguish that part of science that they wish to endorse from that they wish to reject or remain agnostic about. However, realists have often, at least tacitly, endorsed it. They argue, on very general epistemological grounds, that this extra step is justified. For example, it is argued that only these metaphysical commitments to unobservables can *explain* why we are so good at making the predictions we can make about the observables (Smart 2014; Putnam 1975) or that the inference we make in taking this

further step is no different in kind from those we make in predicting the observables (Kitcher 2001).

I will argue that allowing the anti-realist to frame the debate in this way is to grant too much from the start. My response to the anti-realist draws not on general epistemological grounds but rather on the very practices of the sciences themselves. By examining the construction and testing of scientific theories, we see that detachability fails. The practices of scientists are inexplicable except in light of the metaphysical commitments of the theories they adopt. In particular, I shall argue that in contrast to detachability, generative linguists[5] do not begin by seeing which of their theories are consistent with all the known observations and then infer to the truth of one which is. Rather, the assumed truth of some general picture of the target of inquiry provides guidance on which observations are relevant to confirmation and which are not.

As applied to generative linguistics, the realist position is that linguistic theories aim to provide, and are reasonably viewed as, true descriptions of the psychological competence of human language users.[6] The anti-realist, on the other hand, denies this. The posits of linguistic theories (linguistic expression types, such as noun phrase [NP] or verb phrase [VP], and combinatorial operations, such as Merge) are, at best, theorists' tools for the prediction of observable linguistic behavior and need not correspond in any way with any of the causal machinery involved in the generation of such behavior. My claim is that such a position is inconsistent with the methodology adopted by mainstream generative practice. In the next section, I shall sketch the methodology of generative linguistics, focusing on the application of the (in)famous CP distinction, before turning to how such an approach precludes detachability and thus anti-realism.

3. Competence and performance

At the most general level, linguistic theory aims to describe the properties of human language. Given that human languages are productive, in that there are infinitely many well-formed and meaningful expressions in every natural human language, such a description must consist of identifying the *rules* that govern the construction of such expressions. For example, the rules for the formation of relative clauses provide a way of indefinitely extending noun phrases. A finite number of symbols can thus be used to generate indefinitely many expressions: "the woman," "the woman whom the man loves," "the woman who wishes the man didn't love her," and so on. In addition to enabling the generation of complex constructions from simple constituents, these rules *preclude* certain expressions and certain interpretations of expressions. For example, *"the woman who devours" is not a noun phrase in standard English, and although "the woman whom the sandwich devoured" is, it cannot be interpreted to refer to a sandwich.

This goal of describing the rules governing human language suggests an obvious methodology: propose rules, and compare the expressions generable by such rules

---

[5] I shall sometimes refer simply to *linguistics*, but it should be kept in mind that I am speaking only about the broadly generative tradition.

[6] I am here rejecting "noncognitivist" accounts of linguistics, such as those of Katz (1980) and Devitt (2006). This is because I am drawing the CP distinction on *causal* grounds, and neither abstract objects nor public symbols cause linguistic behavior in a suitable way.

with competent (e.g., native) speakers' judgments. To the extent that the rules generate only expressions that competent speakers deem to be acceptable on the predicted interpretation, the theory is confirmed. To the extent that the speakers judge these expressions to be ill-formed or well-formed but not interpretable as predicted, the theory is disconfirmed.[7]

To get a better sense of this methodology, let us consider a very quick and simple piece of linguistic theorizing. One of the driving questions of linguistic theory is this: In what ways do/can natural languages differ, or what are the rules that differentiate languages? This question can be answered by looking at similar expressions in different languages and seeing whether they can be captured by the same rules or whether distinct rules are needed. For example, the distribution of subject pronouns in spoken Spanish and English appears to differ. In English, indicative sentences (as opposed to imperatives) must have explicitly pronounced subjects: a sentence like "I am hungry" is uncontroversially well formed, but when we leave out the subject "I," the sentence becomes totally degraded: *"am hungry." This is not so in Spanish, in which subject pronouns may be absent from produced utterances. The expressions corresponding to the previous English examples, "Yo tengo hambre" (literally: *I have hunger*) and "Tengo hambre" (*have hunger*), are both acceptable. This shows that the mandatory presence of main subjects in spoken English indicatives is a contingent feature of the language, and so characterizations of the rules of English ought to include something to this effect.[8]

We now have some data ("I am hungry" is good, *"am hungry" is bad) and a rule (English indicative sentences must have an explicit subject). The rule is compatible with the data in that the good sentence is consistent with the rule, and the bad sentence is not. So far, so good. Even better, this rule allows us to explain phenomena that might seem unrelated. For example, there are various phenomena in English wherein the grammatical subject of a sentence is mismatched with the sentence's semantic properties. Sentences that semantically require no agent are found with dummy ("expletive") pronouns ("It's raining"). This puzzling phenomenon doesn't arise in Spanish: "Está lloviendo" (*is raining*) or simply "llueve" (*rains*) are the natural ways of expressing this thought. Our initial rules (English requires explicit subjects; Spanish doesn't) seemed to simply describe our observations but can now be used to *explain* other, novel phenomena. Although the weather is not semantically associated with an agent, the rules of English grammar require an explicit subject, so the dummy pronoun "it" is inserted. This requirement is absent in Spanish, precluding similar phenomena there.

So far, linguistic methodology seems to be roughly what one might call "naive inductivism": gather data, propose generalizations consistent with and explanatory of such data, repeat. However, what happens when a well-confirmed generalization

---

[7] For completeness, I am here including reference to the semantic properties of these test expressions: speakers do not merely judge that a sentence is well formed but also that it is well formed *in a certain interpretation*. This will require a semantic theory in addition to a syntactic one. I will mostly be concerned with syntactic phenomena and so will have little to say about this semantic component.

[8] For ease of exposition, I am describing such rules in a fairly loose manner. Careful linguistic analysis would distinguish explicitly between the proposal that this is a rule governing the construction of linguistic structure (syntax) and that this rule governs the "externalization" of such constructions (phonology).

confronts data that seem to be inconsistent with it? The flatfooted answer seems clear: reject the generalization, and propose a different rule capable of covering the novel data. However, this is not always the approach in generative linguistics.

Consider commonly uttered expressions like "found it" and "love you" or discourses such as A: "How did you like the new Michael Bay movie?" B: "Hated it." In a wide range of contexts, such utterances sound fine. But these are English sentences without pronounced subject pronouns.[9] Prima facie, these utterances are counterexamples to the proposed rule. Linguists following the naive inductivist strategy would thus reject this rule and seek to find an alternative rule capable of capturing the original data that made the proposed rule seem plausible as well as these new observations. But this is not the typical approach. Generative linguists mostly continue to accept the aforementioned rule as a feature of English, despite these apparent counterexamples. What is going on here?

The crucial move involves distinguishing between *competence*, the rules and constraints governing the operation of speakers' specifically linguistic psychological capacities, and *performance*, observable linguistic behavior that is partially reflective of competence but is also influenced by many other nonlinguistic factors. Generative linguistics, since at least Chomsky (1965), has been dedicated to the study of the former. Performance, in this conception, serves as *evidence* for linguistic theorizing but is not its target. This means that apparent counterexamples to linguistic generalizations can sometimes be dismissed as features of performance rather than competence. In the previously described case, it can thus be argued that it is part of English speakers' competence that sentences require explicit subjects. However, because speech *production* is a rational activity, general considerations of efficient communication lead speakers to produce utterances that deviate in certain ways from the forms licensed by their competence. The factors that make such elision acceptable seem to be of this pragmatic sort (e.g., one can elide a subject only when the communicative context makes it obvious who/what the subject would be). These sorts of speech "shortcuts" should then not be taken as indicative of the underlying linguistic rules.[10]

The clearest examples of the need for such a methodology come from everyday linguistic slips that can be recognized as such by their speakers. For example, it is quite common for speakers to incorrectly mark verbs so as to correspond to the number marking of the nouns nearest to them in the sentence rather than to their arguments (Eberhard 1997). That is, speakers produce sentences analogous to "The author of the books are being awarded a prize." But speakers can easily recognize that such utterances are ill-formed. Thus, such sentences ought to be excluded from the data used to (dis-)confirm theories of grammar (although they may be very illuminating in

---

[9] Note also that such utterances do seem to be genuine *sentences* rather than sub-sentential phrases. Stainton (2006) argues compellingly that sub-sentential expressions can be used to make full-fledged assertions, and it may be tempting to assimilate these data to this phenomenon. However, the cases described previously are marked for tense ("found it" is acceptable as a declarative, whereas the bare VP "find it" is not) and are thus genuinely sentential.

[10] This approach to language and linguistics is controversial in a couple of ways. Firstly, it assumes a distinction between properly linguistic rules and other sorts of psychological activities. That is, it assumes the existence of a *language faculty* in the sense of Hauser, Chomsky, and Fitch (2002). Secondly, it assumes that the central aim of linguistic theory is to describe this faculty. That is, it assumes a psychologistic interpretation of linguistics (as opposed to that developed by Devitt [2006]).

other areas, such as psycholinguistics). This uncontroversial point seems to force the CP distinction onto anyone theorizing about natural language. More complex and controversial cases arise when the disparity between competence and performance is not apparent to speakers, as in the previously described cases involving absent subjects and famous examples such as center-embeddings. But I see no reason, in principle, to view these cases as different in kind: the rules of the language are one thing; linguistic behavior is another.

When faced with a counterexample to a generalization, then, there are two options. One is to reject the generalization. The other is, in a sense, to reject the data. This latter is an option frequently (although not, of course, exclusively) taken in generative linguistics. The reason is that the generalizations proposed within this tradition are not claims about utterances but about the operation of a psychological system centrally involved in their production and interpretation. The behavioral phenomena (performance) serve as a source of evidence rather than the target of inquiry.

One worry that could be raised at this point is that such a methodology is viciously circular or offers a "get-out-of-jail-free card" (Goldberg and Gonzálvez-Garcia 2008, 351) to linguists. If which observations are treated as performance effects and thus irrelevant to (dis-)confirmation of a theory is determined by the theory itself, it will always be possible for the theorist to treat any apparently disconfirmatory evidence as ipso facto irrelevant, insulating the theory from counterexemplification. I believe this worry is overstated. Classifying an observation as a performance effect is to deny that it is primarily explained with reference to competence. It is not to deny that it is explicable in general. Thus, drawing the CP boundary makes predictions about which observations are explained by what theory. If a purported performance datum seems stubbornly resistant to explanations involving extra-linguistic cognitive systems (e.g., memory, perceptual systems, personal-level beliefs and intentions), this places an empirical strain on the claim that it is indeed a performance datum. That is to say that the CP distinction shifts the explanatory burden, and the empirical sensitivity, from linguistics to some other field but does not eliminate it entirely. In this way, although when taken individually, linguistic theory might be insulated from specific data, in the broader context of multiple noncompeting theories of human behavior, all the data remain relevant.[11]

Briefly, it is worth noting how the CP methodology differs from various well-known approaches to apparently counterexemplifying anomalies in the literature, such as those of Kuhn (1962/2012), Cartwright (1983), and Jerry A. Fodor (1974). These approaches all assume that anomalies are genuinely counterexamples to the generalizations proposed within the sciences. They differ in their accounts of how anomalies should be understood and responded to by scientists. Kuhn allowed that all theories face anomalies, but he argued that scientists did, and should, ignore them, at least for the time being while the theories were developed. But over time, scientific crises and, ultimately, theory replacement result from the accumulation of anomalous data. Cartwright similarly accepted the inevitability of anomalies but argued that

---

[11] There is more to say here in responding to this worry. But because my goal is to show that accepting the methodology of generative linguistics commits one to realism about linguistic theory, not to defend this methodology against its detractors, I will leave it with this brief comment. For a fuller account along these lines, see Dupre (2019).

we should take this fact at face value and infer that all purported scientific laws are therefore false, motivating a pragmatic and localist approach to science. Fodor, and the large literature on ceteris paribus laws that followed, accepted that anomalies did in fact falsify universal interpretations of scientific laws and argued instead for a weakened view of scientific laws, according to which they could be true, "all things being equal," despite counterexamples.

Lakatos ([1976](#)) provides a slightly different account, according to which the "hard core" of a science is insulated from apparent counterexemplification by modification to the "auxiliary hypotheses." This view is closer to mine, but it still retains the idea that such avoidance of apparent counterexamples is a bad thing. In Lakatos's terms, a research program that relies too frequently on such tinkering with auxiliaries is liable to be disvalued as a "degenerating" program. In my view, excluding an observation as "mere performance" has, in general, no such negative results.

Despite their significant differences, all these proposals view apparently counter-exemplifying data as relevant for theory confirmation, differing in how they mitigate this. This is the key difference between these approaches and my understanding of CP methodology. When an observation is excluded from (dis-)confirming a theory on the grounds that it is a "mere performance datum," it is being claimed to be strictly irrelevant to the truth of the theory/generalization. It does not counterexemplify the rule, even as a case where not all things are equal. It simply reflects confounding causal influence, which is not the target of the theory and so can be ignored.

4. The impossibility of anti-realism in generative linguistics

But what does this talk of linguistic methodology have to do with the realism/anti-realism debate? As indicated earlier, the difficulty for the anti-realist stems from the commitment to detachability: the assumption that the ontology of a science is an optional extra that can be accepted or not without changing the character of the science itself. CP shows this to be impossible, at least in generative linguistics. The argument for this can be stated simply:

1. Scientists distinguish between those observations that are pertinent to confirmation and those that are not.
2. This distinction is drawn by appeal to the unobservable causes of our observations.
3. But the anti-realist cannot appeal to the unobservable causes of our observations.

   Conclusion 1. So the anti-realist cannot draw this distinction.
   Conclusion 2. So the anti-realist cannot make sense of scientific practice.

This section elaborates on and defends this argument, applying it to both instrumentalist and empiricist strands of anti-realism. Note that, as argued in section [2](#), these positions are most plausible when taken together, and so an argument against either serves as an argument against both. However, I believe a stronger case can be made: both theories taken individually succumb to this argument.

First, the instrumentalist. Remember that on this view, science does not aim at correct description of the unobservable world but instead at accurate prediction of the observable world. Scientists are free to appeal to whatever unobservables will

help with this goal, but in doing so, they act *as if* such things exist, and their success is judged independently of the truth of such posits. The difficulty that CP raises for instrumentalism is that generative linguists often appear to *eschew* prediction of the observable world. In the case described previously, the linguistic theory (containing the rule that English sentences must have explicit subjects) fails to predict various phenomena. But this is not necessarily seen as a shortcoming of the theory to be improved upon. Rather, these observations are simply dismissed as irrelevant to the theory in question. They are, in the terminology of generative linguistics, rejected as "mere performance."

Sober (2015, 1999, 2002) has argued that instrumentalism can make better sense of certain episodes in scientific theorizing on the grounds that scientists tend to favor theories that are known to be false (e.g., null hypotheses, which, taken literally, are wildly implausible) on the grounds that they are likely to enable better prediction. He is explicit that the alternative case would provide an argument for realism: "Perhaps there are situations in which the choice is between [truth and predictive adequacy] and where scientists prefer [the former]. Realists need to produce such examples" (1999, 27). I take the example of the CP distinction to meet this request.

Note that this is not merely a case of less-than-ideal science being confronted with obstinate anomalies, which is, of course, compatible with instrumentalism. It is in many cases possible, and often fairly easy, to construct a theory capable of predicting these observations. In the case of pronoun dropping in English, one could retreat from the universal claim to a weaker probabilistic claim (as suggested by Norvig [2017]), or revise the rule so as to apply only to some subset of English sentences (sentences without first-person subjects or for which the subject is not contextually salient, etc.), or modify our theories of verbal subcategorization (just as some verbs, like *eat*, may or may not have object arguments, some verbs, like *love*, may or may not have subjects), or some combination of these. All such strategies are, however, deeply problematic theoretically. Beyond being clear cases of overfitting, they require significant complications to the grammatical theory. Because language must be acquirable despite fairly minimal and varied linguistic experience, there are strong motivations for keeping this grammar highly simple. For this reason, we are much better off excluding such anomalous data from the confirmation base of our theory than adapting our theory to generate it. And this is thus what we see in linguistics. In this case, simplicity is taken to be an indicator that the theory is true, *even though* the simpler theory is less likely to allow for better predictions, in conflict with Sober's claims that simpler theories will typically be favored on instrumentalist grounds.

One can press this worry with the following question: Which predictions do we want our theory to make? Of course, we want it to make some (correct) predictions. As we saw earlier, our simple grammatical hypothesis seemed more compelling because it allowed us to account for phenomena for which it was not specifically designed (expletive subjects). This minimal empirical sensitivity is required of any science. However, what this discussion of CP highlights is that not all predictions are desirable. Some predictions seem like significant marks in favor of a theory, whereas others seem, at best, irrelevant and, at worst, actively unwanted (e.g., a grammatical theory *predicting* that no English sentences contain 13,956 words or more would be predictively accurate in this respect but would be *thereby* less plausible).

CP provides a principled distinction between those predictions we want our theories to make and those we don't. Those observations that primarily reflect linguistic competence are relevant to the confirmation of our grammatical theories. Those that are too distorted by extra-linguistic factors, such as communicative efficiency, memory constraints, parsing heuristics, and so forth, are not.[12] This reasoning is, however, expressly realist. Linguists determine the relevance of observational phenomena, performance, on the basis of their relations to their unobservable causes (competence vs. extra-linguistic cognition). It is hard to see how an instrumentalist could even draw this distinction, let alone motivate it. If what matters is prediction, then observable performance, rather than unobservable competence, should be the central target of our linguistic theories.

The point here is not merely that a theory aiming to describe competence is ipso facto a realist theory. This is true, but the anti-realist can simply say that this aim is misguided. The point is that without accepting some truths about the unobservable causes of linguistic performance, we cannot make sense of the practice of generative linguistics. It is precisely because utterances of "Tengo hambre" by native Spanish speakers and utterances of "See you" by native English speakers are causally explained by different psychological systems that a grammar for Spanish had better capture the former fact, whereas a grammar for English need not, and should not, capture the latter. Thus the ontology of the theory, the psychological systems it posits in the causation of behavior, is inextricably intertwined with its methodology and epistemology. Instrumentalism, with its focus on mere prediction, is inherently unable to account for this.

Empiricism suffers from a closely related problem. Empiricism is an anti-realist *epistemology*: science justifies claims only about observable objects and phenomena, not about unobservable ones. This allows empiricists to differentiate themselves from the skeptic by allowing that science grants genuine epistemic authority on certain claims, even about the as-yet unobserved. However, this authority is restricted so as to not apply to theses about viruses, bosons, mental states, and the like. The problem for this view is, again, that it seems unable to draw a distinction between the observations that matter, and are taken to matter, for linguistic theorizing and those that do not.

The empiricist relies on an epistemic distinction between claims about the observables and claims about the unobservables. The former is said to confer no justification on the latter. Claims about observables confer justification only on one another. Going beyond such testable, observational statements to claims about the unobservable world requires leaving the reliable methods of natural science and doing philosophy.

Such a view is unable to account for the differential significance attached to different kinds of observation in linguistics. Whereas empiricism is a thesis about what *can be empirically confirmed*—namely, that only observationally testable statements can be confirmed—CP points to a distinction about *what can confirm*. As we have seen, some observations can (the acceptability of "Tengo hambre" confirms a grammar for Spanish that allows dropped subject pronouns), and some cannot (the acceptability of

---

[12] The distinction here is principled in a metaphysical sense: it captures a real difference in the etiology of our observations. This is not to say that there is a method or procedure that could be used, independent of our theorizing, to classify observations as reflective of competence or not.

"Missed me" does not disconfirm a grammar for English that disallows dropped subjects). This distinction is drawn by linguists on the grounds that the former is causally explained by one kind of psychological state (linguistic competence), whereas the latter is explained by a different kind of psychological state. That is, this distinction is drawn on strictly realist grounds. The empiricist simply can't make sense of this distinction. Without reference to unobservable causes, the set of observations relevant to confirming a linguistic theory will look heterogeneous and arbitrary. Thus again, these anti-realist approaches can't make sense of the actual practice of linguists.

It is not that the differential significance of different observations is itself inconsistent with empiricist epistemologies. All parties agree that some observations are more epistemically relevant than others. The difficulty comes in providing an explanation of why it is this set of observations that is particularly relevant rather than some other set. The empiricist is forced to answer with reference to the differential relevance these observations have to claims about other observables. And this may sometimes be sufficient. Some observations provide a better basis for projecting about future observations than others. Observations in novel contexts provide better inductive support to generalizations than observations in familiar contexts. However, for the empiricist to fully capture the practice of a science, evidential relevance to observables must line up with the importance placed on observations by practicing scientists, and this need not, and will not, always be the case. There is, in general, no reason to think that observations taken to be reflective of competence are more telling about future observable behavior than are mere performance data. Many "performance effects" are highly reliable, such as those stemming from constraints on memory load, or parsing heuristics, or ungrammatical idioms. And many aspects of competence, although reliable, are infrequently found in normal linguistic behavior or are commonly "overruled" by competing extra-linguistic influences. "Eliise and myself request your attendance" reliably sounds more acceptable than "After the boss ran the company went bankrupt," despite the latter being syntactically perfectly well formed and the former being in violation of binding-theoretic rules on the distribution of pronouns.[13] If we simply wanted to justify claims about observables, there is no particular reason to focus on those data that linguists do. However, the realist has a straightforward explanation of linguists' behavior: linguists focus on those observations that are causally explained specifically by linguistic competence. Without the realist commitment to the existence of this unobservable system, this approach doesn't make sense.

Having argued that generative linguists must be realists, it is worth briefly answering the question: Realistic about *what*? That is, just which aspects of a given proposed grammar are linguists committed to endorsing? On the one hand, linguists can't

---

[13] Binding theory concerns rules governing the acceptable distribution and interpretation of pronouns. The aforementioned example is claimed to violate the binding-theoretic principle that reflexive anaphora must be co-referential with a nearby NP. Exact statements of such principles are controversial, but the reality of the phenomenon is near-universally accepted and is appealed to in explaining why "Eliise made herself a drink" is acceptable, whereas *"Eliise asked Sandro to make herself a drink" is not.

Although recent "minimalist" theorizing in linguistics rejects the assumption that binding theory constitutes its own independent linguistic "module," it retains these binding phenomena as explananda. See Hornstein, Nunes, and Grohmann (2005, 13–14) for a discussion of the relation between minimalism and traditional generative approaches.

identify relevant data on the basis of a *fully articulated* linguistic theory. If they took such a theory to be true, there would be nothing for them to test. And further, letting consistency with such a theory be a guide to confirmational relevance would be viciously circular. On the other hand, it cannot be merely the belief that *some linguistic theory or other* truly describes our competence. Such a vague "commitment" falls short of realism entirely and wouldn't provide any guidance in seeking relevant observations. The realism required falls between these two extremes, consisting of some general assumptions about the nature of the linguistic system in question and constraints on possible/plausible linguistic explanations.

   What is assumed about the linguistic system varies over time with theoretical fashion. But some such assumptions must be made in order to constrain the set of explanatory options available to the linguist in identifying the crucial phenomena. It is widely assumed, for example, that "syntax can't count"—that is, that no linguistic rule can apply only some finite number of times successively. If a linguistic phenomenon would require "counting" rules, this may be evidence that it should be explained with reference to performance, not competence. Another example is the assumption that linguistic rules must be binary branching. This is assumed by most minimalists, based on the empirical arguments of Kayne (1994), and has been central to much generative work since X-bar theory became prominent. It is, however, denied by other linguists (e.g., Culicover and Jackendoff 2005). Which such assumptions are made will determine which observations are plausibly explained as reflecting competence and which are more plausibly viewed as performance effects. And indeed, Jackendoff and Culicover are explicit that they view the fact that their assumed architecture allows them to treat a greater range of phenomena as competence, not performance, reflecting as a mark in favor of their view. Crucially, these assumptions themselves will be confirmed, or disconfirmed, by the successes of the fleshed-out grammars they license. In general, linguists will assume the truth of well-confirmed and theoretically fruitful claims about linguistic competence and leverage these assumptions into novel analyses of further observations, as when the widely assumed constraint that English subjects be pronounced is appealed to in ruling such data out as performance effects. What is assumed in linguists' decision making is thus an empirical and changeable matter, utilizing largely established claims about the architecture of the language faculty in determining which observations are pertinent for determining the rest of its structure.[14]

5. Beyond generative linguistics: From retail to wholesale

Although I think it is bad practice to let philosophical theories dictate which scientific theories are legitimate, it is, I suppose, an option for the anti-realist to simply reject the generative linguistic theorizing I have based my argument on. There is significant controversy about what the correct approach to linguistics is, and various traditions and authors have rejected generativism, particularly highlighting discomfort with CP. Many such approaches aim instead to stick closely to the observational data, proposing theories aimed at capturing more "surface-level" linguistic phenomena

---

[14] In this way, my approach bears similarities to the "investigative scaffolding" discussed by Currie (2018), in that judgments about the relevance of an observation to theory confirmation depend on current theoretical knowledge.

(e.g., Chater, Clark, Goldsmith, and Perfors 2015). I think there are strong empirical reasons to retain the generative approach I have assumed in this article, but I will not be defending that here. Instead, I will respond to this objection by briefly arguing that the considerations raised by CP are liable to arise in almost any scientific discipline. Detachability is unlikely to be viable in science in general, and so the commitment to realism underlies scientific methodology across the board.

Magnus and Callender (2004) distinguish between two different kinds of arguments for realism: *wholesale* and *retail.* Wholesale arguments aim to show that science *in general* should be understood in realist terms as successfully describing the unobservable world. Such arguments abstract away from individual scientific achievements, relying instead on general patterns and trajectories across the history of science. Retail arguments, on the other hand, have narrower scopes: they appeal to particular scientific results and argue that these cannot be accounted for without substantial ontological commitments. Famous examples include Salmon (1984, 213–26) and Hacking (1983, 21–31), who provide arguments from the history of physics for the reality of sub-atomic and microscopic particles.

Retail arguments are compelling in that they turn on concrete work in the sciences rather than general epistemological commitments, such as the reliability of inference to the best explanation. Relatedly, to speak of the successes of science *in general* often obscures more than it clarifies. Although some sciences, such as physics and chemistry, have been highly successful, others, especially the human sciences, are widely agreed to have more mixed records. Given this, the retail approach has often been plausibly touted as better suited to the naturalistic impulse in the contemporary philosophy of science. I take the foregoing discussion to provide just such a retail argument for realism in generative linguistics, albeit one that focuses more closely on methodology than is typical. However, retail arguments are only as compelling as the science they draw on. In the cases discussed by Salmon and Hacking, this science is particle physics, uncontroversially one of the best-developed sciences of the unobservable world. Generative linguistics has, even according to its advocates, not developed to the same level of depth or confirmation, and so the argument for realism is correspondingly less compelling. However, I believe that the aforementioned issues will occur across the sciences, and thus this retail argument can be turned into a wholesale argument: any science that relies on its theoretical posits to distinguish relevant from irrelevant observations will be methodologically committed to realism.

The obvious place to begin extrapolating this point from linguistics to science more generally is in other parts of psychology. As in linguistics, psychologists aim to uncover generalizations about the behavior and development of the mind. Such generalizations are not mere statements of observed regularities. Instead, they are descriptions of underlying and unobservable cognitive systems. And indeed, as in linguistics, generalizations that seem to conflict with observations can be retained if the apparent counterevidence is attributable to confounding influences other than the system under investigation in ways that seem difficult to make sense of from an anti-realist perspective. I will illustrate this with a case study from developmental psychology that has been of significant interest to philosophers: developmental accounts of Theory of Mind (ToM).

ToM refers to whatever psychological system enables humans to attribute mental states to other creatures. A core component of this capacity is the ability to view

others as having *beliefs* and, in particular, beliefs different from those of the attributer. Adult human beings can do this, whereas blastocysts cannot. The question for the developmentalist, then, is this: How do we get from there to here?

An intermediate question is, *When* do we get from there to here? That is, at what stage in the human life cycle do we start to view other agents as believers? The tool most commonly and notoriously used for such a purpose is the "false-belief test" (Wimmer and Perner 1983; Baron-Cohen, Leslie, and Frith 1985). This test involves exposing a subject to a scene in which some information is evidently available to a protagonist. However, when the protagonist leaves the scene, thus losing informational access to it, this information is changed, rendering the protagonist's beliefs about the scene inaccurate. The subject is then assessed with respect to their expectations about the protagonist's behavior. A subject capable of attributing false beliefs should expect the protagonist to behave on the basis of their original, now-false, information. In contrast, a subject who lacks this capacity might expect the protagonist to act on the basis of the evident-to-the-subject facts about the scene. In a paradigmatic example, the subject watches as a doll, Sally, places a toy in a box before leaving the scene. While Sally is absent, another doll, Ann, enters and moves the toy from the box to another location, say, a basket. When Sally reenters, the subject is asked where she will look for the toy. Subjects capable of viewing other agents as possessors of, possibly false, beliefs will predict that Sally will look where she left the toy, that is, in the box. Subjects who don't view other agents in this way, who don't distinguish between the world as it really is and as it is mentally represented by others, might predict that Sally will look where the toy actually is, that is, in the basket.

Interestingly, children reliably (see Wellman, Cross, and Watson 2001) transition from "failing" the test (guessing that Sally will look in the basket) to passing the test somewhere between their third and fifth birthday. This discontinuity in observed behavior suggests a discontinuous theory: (some aspect of) ToM is absent in infants but then develops, as a result of learning or maturation, around the end of toddlerhood. For example, infants may lack the concept of a belief and thus may not distinguish between their "take" on the world, the world itself, and other agents' perspectives. But this concept is acquired around age four, at which point children are able to pass the false-belief test.

However, despite this easy correspondence between these observations and theory, the discontinuous theory has not been universally accepted. Various theorists have argued that children's discontinuous *performance* on false-belief tests is not sufficient grounds for proposing a discontinuity in their development of a ToM competence (see Jerry A. Fodor 1992; Surian and Leslie 1999; Bloom and German 2000). Just as in the linguistic case, it is argued that the same observable behavior is compatible with multiple distinct underlying causal stories. Changes in behavior can be attributed to changes in the system of interest (ToM) or to changes in the other systems relied on in producing behavior. For example, classical false-belief tests are *verbal* tests: the subjects are *asked* what Sally will do upon return. This means that their behavior is dependent not solely on whether they can attribute (false) beliefs to Sally but also on how they interpret the experimenter's question. Thus, if their linguistic capacities change in this time period, that change, rather than a change in ToM, could account for their differential behavior at three and five years of age. Alternatively, this change could be attributed to changes in short-term memory

capacity, attention span, or numerous other cognitive traits. Thus, it is compatible with these results that children have just the same conceptual resources as adults for attributing beliefs, but the other mechanisms relied on in false-belief tasks differ in ways that explain their failures on this test.[15]

As we saw in linguistics, psychologists are not merely seeking hypotheses compatible with the observations. When some observations (discontinuous behavior) seem to support one theory (discontinuous conceptual development) over another, this doesn't settle the issue. Rather, this raises the question of whether these observations reflect the target system or some other causally relevant influence. If one can make a compelling case for the latter, then these observations cease to be relevant for the confirmation of the theory. That is, if the difference between three- and five-year-olds is really reflective of mere performance factors rather than of developments in ToM, then this developmental discontinuity doesn't reveal anything against continuous theories of ToM development. And so a theory that seems to cover fewer of the observations is preferred over a theory that more closely captures the data. But this choice is made essentially with reference to the unobservable causes of the observations and thus is outside the scope of anti-realist theories of science.

Of course, cognitive psychology is close to generative linguistics in both its goals and its methods, and so it is perhaps not all that surprising to see that arguments applied to the latter generalize to cover the former. Indeed, the methods of cognitive psychology are, in many cases, explicitly influenced by those of generative linguistics. However, we can extract from these cases a general set of conditions that, when met in any science, will motivate such realist reasoning:

1. The aim is to describe a specific system ("S1").
2. S1 is a component of a larger, complex system ("S2").
3. The behavior of S1 influences but does not determine the behavior of S2.
4. We cannot extract S1 from S2 and identify its properties in isolation.
5. We cannot remove or hold fixed the contributions to the behavior of S2 made by nontarget subsystems other than S1.

In scientific contexts meeting these conditions, scientists make judgments about which observations are relevant to S1 and which are not. And these judgments will be essentially theory based. The entire strategy of eliminating confounds, avoiding "experimental artifacts," distinguishing reliability from validity, and so forth involves realistic reasoning of this sort.

Because my expertise is within the cognitive sciences, I will leave a detailed discussion of cases in which these conditions are met in the natural and social sciences to those who know more about them. I believe that we will see very similar patterns of reasoning there, but I remain officially neutral. Sciences like community ecology seem to be prime candidates. If we want to identify the specific influence of fluctuations in populations of one organism on those of another, it may be impossible to examine this other than in the full complexity of their natural ecosystem, and so the potential for confounding

---

[15] Such a hypothesis is further motivated by analogues of false-belief tests in which much younger children appear to display sensitivity to other agents' false beliefs. See, for example, Baillargeon, Scott, and He (2010) and Luo (2011).

influences is great. Diamond (1986) discusses examples of this sort. On the other hand, largely experimental sciences, such as particle physics, may be less liable to produce cases meeting these conditions, items 4 and 5 in particular, although see Franklin and Preovic (1998, sec. 2.3) and Schindler (2018, chap. 6) for plausible examples.

What is distinctive about linguistics, and cognitive science more generally, is thus not that a realistic attitude toward its ontology is required in order to determine which observations are pertinent to theory confirmation (those causally explained by such theoretical posits) and those that are not (those causally explained by non-target systems). Rather, linguistics is distinctive in the strategies it uses to draw such a distinction. But these differences stem not from deep epistemic or metaphysical disagreements between practitioners of these fields but from practical constraints on evidence-gathering strategies.

It is widely accepted that it is difficult and epistemically risky to infer properties of the underlying system(s) of interest from mere observation of the natural world. Surface appearances are typically products of numerous interacting causes, so there is no clear path from such observations to the properties of these component systems. For this reason, the epistemically best strategy is often one of experimentation, in which an artificial system is constructed so as to, as much as possible, exclude all nontarget sources of causal influence.[16] If it can be brought about that the target system is the *only* causal determinant of observed behavior, then inference from properties of the latter to those of the former is unproblematic. Another good strategy, when this experimental elimination of confounding causes is unavailable, is to compare situations in which confounding forces are present but constant, whereas the influence of the target system varies (as in a randomized controlled trial). Although this won't show us the influence of our target system *neat*, it will show the difference that such an influence can account for. By eliminating or factoring out the causal confounds, these approaches enable theorists to gain insight into the hidden causes of the observables. These two very general strategies together account for a large proportion of theorizing in the natural and social sciences. However, for obvious reasons, they are typically not available to the linguist.

Firstly, we, of course, cannot remove the human language faculty from the rest of the system causally responsible for linguistic performance, in the way that we can sometimes remove organisms from their normal environments and watch their development absent various confounds. Secondly, even if we could observe the language faculty working on its own (e.g., with complex neural-imaging equipment), this would be unlikely to tell us anything much about the level of description we are interested in. As Poeppel and Embick (2005) argue compellingly, there is a "granularity mismatch problem" between neurobiological and psychological descriptions: we have no idea how to map descriptions of brain states onto fine-grained linguistic descriptions. Qua linguist, one is interested in linguistic competence as a *psychological* property, and our descriptive vocabulary must thus be couched in psychological terms. Only by viewing the outputs of a linguistic system as causal antecedents of linguistic performance are we able to characterize this system in this way. So experimental investigation of the target, competence, operating independently of performance systems seems

---

[16] Hacking (1983) provides a detailed discussion of this strategy, as do Nancy Cartwright and Allan Franklin in numerous places.

out of the question.[17] Likewise, given the impossibility of keeping confounding factors constant across trials, a "control and compare" strategy seems inapplicable. Linguistic performance is sensitive to numerous features of prior linguistic experience, and ethical constraints preclude raising children in environments controlled enough to ensure that such confounds are genuinely evenly distributed.

For these reasons, linguistics is stuck dealing with a genuinely causally confounded system and is unable even to attempt to hold confounding influences constant to see the distinctive contribution of specifically linguistic competence. The CP approach is an attempt to de-confound these data without experimental intervention by making educated guesses about the underlying system best suited for explaining such observational properties. But this difference in approach merely reflects the practical difficulties of studying part of the human mind. The epistemological and metaphysical assumptions about the reality of unobservable causes, and about the strictly evidential role of prediction, are the same in linguistics as in these other sciences.

6. Conclusion

In this article, I have presented a novel argument for scientific realism.[18] What I think is notable about this argument is not its conclusion, as I take some weak form of scientific realism to be basically the default position, but its strategy. The epistemologies appealed to in most arguments for scientific realism are very general; indeed, it is typically taken to be a virtue of such arguments that the reasoning utilized in scientific ontology is continuous with everyday inference. Whatever the merits of such approaches, I believe they have distracted from the distinctive and subtle methods used in theory construction and evaluation in specific sciences. I hope this article indicates the value of attending to these science-internal epistemologies and methodologies in addressing perennial philosophical questions. That the ontological commitments of a theory are not detachable from the methods and practices of the scientists investigating it is an important fact about scientific practice, both in generative linguistics and beyond.

This discussion points to a very general lesson. Psychology was able to make significant progress in the twentieth century precisely because it moved away from surface appearances and started theorizing about underlying systems. A productive cognitivist program thus superseded a sterile behaviorist one. But this "retreat from the surface" was not merely a metaphysical move, allowing into our ontology unobservable mental states and processes, but an epistemological one. The significance of this latter aspect of the cognitivist revolution has not been fully appreciated. Once we

---

[17] This is not to say that experimental work cannot be useful in distinguishing competence from performance. Of course it can, and it has done so throughout the history of psycholinguistics (see, e.g., Sprouse and Hornstein [2013] for a range of work aimed at just this). What I deny is that such a distinction can be drawn independently of, or prior to, evaluation of our theory of linguistic competence.

[18] However, see Schindler (2011, 2013, 2018) for a realist argument that scientists often justifiably view data that seem inconsistent with a well-confirmed or otherwise virtuous theory as unreliable. This is akin to my cases of excluding observations on the grounds that they reflect mere performance. Although Schindler's and my approaches are similar and mutually reinforcing, they differ at least in emphasis. Schindler discusses cases in which the data are assumed to be erroneous (e.g., as a result of faulty experimentation), whereas in my cases, the data are perfectly good on their own terms; they just don't reflect the target of interest.

give up on the idea that our scientific theories must be *about* the observable world, we ought to likewise give up the idea that the central test for our theory is how closely and comprehensively it *predicts* such observable phenomena. Prediction is, of course, important. But it should not be seen to trump other aims. There is no reason in principle, nor, as I have argued earlier, in practice, that the theory with better empirical coverage is thereby the better theory. The CP distinction provides one entryway to this point and reorients the discussion concerning the predictive successes of our theories. The question is not, Which theory captures the most observable behavior? Rather, it is, Which observable behavior provides the most insight into the underlying target of interest?

References

Baillargeon, Renée, Rose M. Scott, and Zijing He. 2010. "False-Belief Understanding in Infants." *Trends in Cognitive Sciences* 14 (3):110–18.

Baron-Cohen, Simon, Alan M. Leslie, and Uta Frith. 1985. "Does the Autistic Child Have a 'Theory of Mind'?" *Cognition* 21 (1):37–46.

Bhaskar, Roy. 2013. *A Realist Theory of Science*. New York: Routledge.

Bloom, Paul, and Tim P. German. 2000. "Two Reasons to Abandon the False Belief Task as a Test of Theory of Mind." *Cognition* 77 (1):B25–B31.

Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. Oxford: Oxford University Press.

Chakravartty, Anjan. 1998. "Semirealism." *Studies in History and Philosophy of Science Part A* 29 (3):391–408.

Chater, Nick, Alexander Clark, John A. Goldsmith, and Andrew Perfors. 2015. *Empiricism and Language Learnability*. Oxford: Oxford University Press.

Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.

Culicover, Peter W., and Ray Jackendoff. 2005. *Simpler Syntax*. Oxford: Oxford University Press.

Currie, Adrian. 2018. *Rock, Bone, and Ruin: An Optimist's Guide to the Historical Sciences.* Cambridge, MA: MIT Press.

Devitt, Michael. 2006. *Ignorance of Language*. Oxford University Press on Demand.

Diamond, Jared. 1986. "Overview: Laboratory Experiments, Field Experiments, and Natural Experiments." In *Community Ecology*, edited by J. Diamond and T. J. Case, 3–22. London: Harper and Row.

Dupre, Gabe. 2019. "Linguistics and the Explanatory Economy." *Synthese* 199 (1):177–219.

Eberhard, Kathleen M. 1997. "The Marked Effect of Number on Subject–Verb Agreement." *Journal of Memory and Language* 36 (2):147–64.

Fodor, Jerry A. 1974. "Special Sciences." *Synthese* 28 (2):97–115.

Fodor, Jerry A. 1991. "You Can Fool Some of the People All of The time, Everything Else Being Equal; Hedged Laws and Psychological Explanations." *Mind* 100 (397):19–34.

Fodor, Jerry A. 1992. "A theory of the child's theory of mind." *Cognition* 44 (3):283–96.

Franklin, Allan, and Slobodan Perovic. 1998. "Experiment in Physics." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. Stanford: Stanford University Press. https://plato.stanford.edu/entries/physics-experiment/.

Goldberg, Adele E., and Francisco Gonzálvez-García. 2008. "Cognitive Construction Grammar Works: An Interview with Adele Goldberg." *Annual Review of Cognitive Linguistics* 6:345–60.

Hacking, Ian. 1983. *Representing and Intervening*. Cambridge: Cambridge University Press.

Hauser, Marc D., Noam Chomsky, and W. Tecumseh Fitch. 2002. "The Faculty of Language: What Is it, Who Has It, and How Did It Evolve?" *Science* 298 (5598):1569–79.

Hornstein, Norbert, Jairo Nunes, and Kleanthes K. Grohmann. 2005. *Understanding Minimalism*. Cambridge: Cambridge University Press.

Katz, Jerrold J. 1980. *Language and Other Abstract Objects*. Lanham, MD: Rowman and Littlefield.

Kayne, Richard S. 1994. *The Antisymmetry of Syntax*. Cambridge, MA: MIT Press.

Kitcher, Philip. 2001. "Real Realism: The Galilean Strategy." *Philosophical Review* 110 (2):151–97.

Kuhn, Thomas S. 1962/2012. *The Structure of Scientific Revolutions: 50th Anniversary Edition*. Reprint, Chicago: University of Chicago Press.

Lakatos, Imre. 1976. "Falsification and the Methodology of Scientific Research Programmes." In *Can Theories Be Refuted? Essays on the Duhem-Quine Thesis*, edited by S. Harding, 205–59. New York: Springer.

Laudan, Larry. 1981. "A Confutation of Convergent Realism." *Philosophy of Science* 48 (1):19–49.

Luo, Yuyan. 2011. "Do 10-Month-Old Infants Understand Others' False Beliefs?" *Cognition* 121 (3):289–98.

Lyons, Timothy D. 2005. "Toward a Purely Axiological Scientific Realism." *Erkenntnis* 63 (2):167–204.

Magnus, Paul D., and Craig Callender. 2004. "Realist Ennui and the Base Rate Fallacy." *Philosophy of Science* 71 (3):320–38.

Norvig, Peter. 2017. "On Chomsky and the Two Cultures of Statistical Learning." In *Berechenbarkeit der Welt?*, edited by Wolfgang Pietsch, Jorg Wernecke, and Maximilian Ott, 61–83. New York: Springer.

Poeppel, David, and David Embick. 2005. "Defining the Relation between Linguistics and Neuroscience." In *Twenty-First Century Psycholinguistics: Four Cornerstones*, edited by Anne Cutler, 103–18. Abingdon: Routledge.

Psillos, Stathis. 2005. *Scientific Realism: How Science Tracks Truth*. New York: Routledge.

Putnam, Hilary, ed. 1975. "What Is Mathematical Truth?" In *Philosophical Papers Volume 1: Mathematics, Matter and Method*, 60–78. Cambridge: Cambridge University Press.

Quine, Willard V. 1948. "On What There Is." *Review of Metaphysics* 2 (5):21–38.

Salmon, Wesley C. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.

Schindler, Samuel. 2011. "Bogen and Woodward's Data-Phenomena Distinction, Forms of Theory-Ladenness, and the Reliability of Data." *Synthese* 182 (1):39–55.

Schindler, Samuel. 2013. "Theory-Laden Experimentation." *Studies in History and Philosophy of Science Part A* 44 (1):89–101.

Schindler, Samuel. 2018. *Theoretical Virtues in Science: Uncovering Reality through Theory*. Cambridge: Cambridge University Press.

Smart, John Jamieson Carswell. 2014. *Philosophy and Scientific Realism*. New York: Routledge.

Sober, Elliott. 1999. "Instrumentalism Revisited." *Crtica: Revista Hispanoamericana De Filosofa* 31 (91):3–39.

Sober, Elliott. 2002. "Instrumentalism, Parsimony, and the Akaike Framework." *Philosophy of Science* 69 (S3):S112–23.

Sober, Elliott. 2015. *Ockham's Razors: A User's Manual*. Cambridge: Cambridge University Press.

Sprouse, Jon, and Norbert Hornstein. 2013. *Experimental Syntax and Island Effects*. Cambridge: Cambridge University Press.

Stainton, Robert. 2006. *Words and Thoughts: Subsentences, Ellipsis, and the Philosophy of Language*. Oxford: Oxford University Press.

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

Surian, Luca, and Alan M. Leslie. 1999. "Competence and Performance in False Belief Understanding: A Comparison of Autistic and Normal 3-Year-Old Children." *British Journal of Developmental Psychology* 17 (1):141–55.

Trout, John D. 1998. *Measuring the Intentional World: Realism, Naturalism, and Quantitative Methods in the Behavioral Sciences*. Oxford: Oxford University Press.

Van Fraassen, Bas C. 1980. *The Scientific Image*. Oxford: Oxford University Press.

Wellman, Henry M., David Cross, and Julanne Watson. 2001. "Meta-Analysis of Theory-of-Mind Development: The Truth about False Belief." *Child Development* 72 (3):655–84.

Wimmer, Heinz, and Josef Perner. 1983. "Beliefs about Beliefs: Representation and Constraining Function of Wrong Beliefs in Young Children's Understanding of Deception." *Cognition* 13 (1):103–28.