

How And Why To Draw The Competence/Performance Distinction

Gabe Dupre

University of California, Davis

September 25, 2023

Abstract

It is currently fashionable to assume that, underlying the actual more or less bumbling speech behavior of any human being, there is a subtle and complicated but determinate linguistic “competence”: a sentence-generating device whose design can only be roughly guessed at by any techniques so far available to us. This point of view makes linguistics very hard and very erudite, so that anyone who actually does discover facts about underlying “competence” is entitled to considerable kudos.

Charles Hockett

1 Introduction

The defining progressive step in 20th Century psychology was to reaffirm the seeming tautology that psychology is the science of the mind, against the stifling strictures of the preceding behaviourism.¹ Noam Chomsky's (in-)famous distinction between *competence* and *performance* is, so I shall argue, properly understood nothing more than this insistence that linguistics is a cognitive, rather than a behavioural, science. If this is so, why has this innocuous distinction drawn so much ire over the years, from linguists, philosophers, and psychologists alike? One major reason, I believe, is a confusion, on the part of both allies and antagonists, of exactly what the distinction amounts to. Appeal to 'competence' and 'performance' has been made to draw many, non-equivalent, distinctions. In this chapter, I shall spell out how we *ought* to understand this distinction, distinguishing my interpretation from several prominent alternatives. What I shall elaborate will be, in a sense, trivial: we must distinguish between behaviour and the causal systems underlying and explaining behaviour. But despite this vacuity, proper attention to this distinction has substantial philosophical upshots, some of which I shall canvas in the latter half of the paper.

2 What Is The Distinction Between Competence and Performance?

The distinction between competence and performance (hereafter 'CP') is a cornerstone of the methodology of generative linguistics. It has also been appealed to, explicitly or tacitly, by a wide range of theorists, in areas such as Theory of Mind (Fodor (1992)), mathematical competence (Simms et al. (2013)), human

¹Ryan Nefdt and Kate Stanton provided instructive feedback on earlier versions of this paper, as did an audience at the University of California, Davis. This research was funded by a Leverhulme Early Career Research Fellowship (ref: ECF-424-2020).

rationality (Cohen (1981)), moral reasoning (Knobe (2010)), and semantics and communication (Stojnić (2021)). But what is it?

CPs canonical development is in the opening passages of Chomsky (1965) (although it had already been appealed to in Miller & Chomsky (1963)). As it will be central to what follows, I will here quote at length (pp.3-4):

Linguistic theory is concerned primarily with an ideal speaker-listener, in a completely homogeneous speech-community, who knows its language perfectly and is unaffected by such grammatically irrelevant conditions as memory limitations, distractions, shifts of attention and interest, and errors (random or characteristic) in applying his knowledge of the language in actual performance... To study actual linguistic performance, we must consider the interaction of a variety of factors, of which the underlying competence of the speaker-hearer is only one. In this respect, study of language is no different from empirical investigation of other complex phenomena.

We thus make a fundamental distinction between competence (the speaker-hearer's knowledge of his language) and performance (the actual use of language in concrete situations). Only under the idealization set forth in the preceding paragraph is performance a direct reflection of competence. In actual fact, it obviously could not directly reflect competence. A record of natural speech will show numerous false starts, deviations from rules, changes of plan in mid-course, and so on. The problem for the linguist, as well as for the child learning the language, is to determine from the data of performance the underlying system of rules that has been mastered by the speaker-hearer and that he puts to use in actual performance. Hence, in the technical sense, linguistic theory is mentalistic, since it is concerned with dis-

covering a mental reality underlying actual behavior. Observed use of language or hypothesized dispositions to respond, habits, and so on, may provide evidence as to the nature of this mental reality, but surely cannot constitute the actual subject matter of linguistics, if this is to be a serious discipline.

These famous passages are rich with insight. But, alongside the important methodological, theoretical, and epistemological insight, there is much that is confusing, and indeed confused. To distinguish the former from the latter, we will need to distinguish a *conceptual* distinction between two phenomena, competence and performance, which can, and should, be drawn by any linguistic theory, from *empirical* claims about how this distinction applies within the now-rejected theory of grammar that Chomsky was then developing.

The core distinction is between “an underlying system of rules” and “the actual use of language”. That is, ‘competence’ refers to a specific psychological subsystem, while ‘performance’ refers to linguistic behaviour, including expressed linguistic judgements (“intuitions”). As Chomsky notes, if linguistics is to be a cognitive science, then the former is the target of interest, while the latter serves as evidence about it. Just as vision scientists study not what people actually see, but rather the rules and regularities underlying the capacity to see, linguists aim to describe the principles governing the human capacity to use (and acquire) language, not to describe actual uses of language. All parties should agree that what we actually say is but a fraction of what we *could* say. Further, what we actually say is liable to be biased in various ways. Uttered sentences will tend to be short, and for many of us uttered sentences will tend to be true, and so on, but our strictly linguistic capacities do not similarly display such preferences for brevity and veracity. And so all parties should agree that there is a distinction

between the system of language and particular uses of it.²

I think it is also clear from the passages above that Chomsky views competence as a *causal influence* on performance.³ More specifically, competence is *one cause among many* of performance. What I say is thus sensitive to, and influenced by, but not determined by, the system of rules I, as a competent language user, have internalized. The data for our linguistic theories, i.e. performance, are thus not pure reflections of the target of our theories, i.e. competence, but reflect the confounding influence of a whole suite of cognitive systems. The job of the linguist is then to identify the underlying system on the basis of this noisy and confounded evidence. Thus, the theoretical claim that distinguishes the internal system from its reflection in public behaviour necessitates the epistemological claim that inferences from the latter to the former are empirical and ampliative. That some phenomenon is found in linguistic performance *does not entail* that it is found likewise in linguistic competence, as it may be that this is a reflection of these other, non-target, causal influences. Chomsky mentions “false starts, deviations from rules, [and] changes of plan in mid-course” as examples precisely of this sort, wherein confounding factors distort linguistic performance, undermining the inference from performance to competence.

So, if Chomsky’s distinction is just that between a potentially explanatory unobserved causal subsystem, and the downstream, observable behaviour on which inferences to the properties of this system are to be based, why has this provoked

²A few clarifications. Firstly, even if, like Devitt (2006) Katz (1980), one denies that linguistics ought to be a cognitive science (although see Dupre (In Press) for reasons not to), the question of what it is about the human mind that enables us to use and learn a language remains, as part of psycholinguistics. Similarly, even if one is skeptical about the status of linguistics as psychology, mostly it is agreed that languages are governed by rules. And so the distinction between the system of rules and the utterances they govern remains. Finally, even if one is skeptical of linguistic ‘rules’, as many in the connectionist tradition are, we could restate the CP distinction so as to distinguish uses of language from the processes and regularities governing the linguistic mind.

³Admittedly, Chomsky is often unhappy with the use of causal terminology as applied to cognition and behaviour in general (see e.g. various discussions in Chomsky (2000)). But I think not in ways that undermine the point I am making. The above points could be made, albeit circumlocutously, instead in terms of competence *partially explaining* performance.

so much ire from linguists (e.g. Reich (1969), Ibbotson & Tomasello (2016)), philosophers (e.g. Rosenberg (1988), Hintikka (1999)), psychologists (e.g. Christiansen & Chater (2016)), computer scientists (e.g. Norvig (2017)) and more? I will later come to what I take to be the core concern that theorists have raised against CP, namely that it undermines our ability to empirically test linguistic theory. But before that, it will be important to identify several alternative accounts of what the CP distinction consists in, displaying how the interpretation I have offered above differs from them. Each is, I believe, a perfectly useful distinction, but much of the confusion surrounding CP has come from conflating them.

Firstly, the distinction between competence and performance has often been treated as equivalent to, or as an instance of, Marr (1982)'s famous distinction between levels of description of an information processing system.⁴ Marr thought that explaining the workings of a cognitive system involved answering several different kinds of question: (i) What was the system doing, and why?, (ii) How was the system doing this?, and (iii) what arrangement of physical parts enabled it to do this? Answers to the first kind of question were to be given by identifying the function that the system mapped from inputs to outputs, and explaining what was achieved by computing such a function. A full set of answers of this sort was called a 'computational level' description of the system. Answers to the second kind of question were to identify the algorithm used to compute this function, and thus constituted an 'algorithmic level' description. And finally, answers to the third kind of question specified how a physical system could implement such an algorithm. Marr himself provides a Chomskian theory of grammar as a paradigmatic example of a computational-level theory (p. 28). This fact may well have been responsible for subsequent researchers viewing performance as closely tied

⁴See e.g. Franks (1995). Patterson (1998) identifies and criticizes this conflation in Franks' article.

to the lower-levels of explanation, algorithm and implementation. But this is a mistake. Linguistic competence specifies the computational description of a particular psychological system. This system (sometimes called the ‘language faculty’ or ‘I-language’) will itself have algorithmic and implementational descriptions. But performance is not a matter of the individual workings of this system, but of this system in concert with all the other systems which account for human linguistic behaviour.

Marr’s divisions are what we might call ‘vertical’: they distinguish different, more fine-grained, ways of describing one and the same system. CP, on the other hand, is a ‘horizontal’ distinction, between a cause and its effects. This is essential to the methodological use that CP is put to. When linguists exclude some phenomenon from the scope of their theory on the basis that it is a “mere performance effect”, this move will make sense only on the assumption that the outputs of competence differ from the behaviour constituting performance. But this is precluded by any strictly vertical distinction. A function will, obviously, have the same input-output profile whether it is described at a computational, algorithmic, or implementational level. And so, while it is plausibly very important to ask questions about linguistic competence at all of Marr’s levels, and ensure that answers to these questions cohere, it is a mistake to view CP as an instance of Marr’s distinction between levels of description. Note further that the Marrian distinction will hold also for all the other systems (memory, perception, motor-control, knowledge storage and acquisition, etc.) that are involved in linguistic performance alongside competence. While the distinction between computational, algorithmic, and implementational descriptions apply to one and the same system, it is the essence of the competence/performance distinction that the systems in question are distinct: competence theories describe a psychologically isolable subsystem, while performance is the behaviour of the whole

organism.

Another, possibly more prevalent, misinterpretation of CP has it that competence is an idealized, or cleaned-up, version of performance. Competence describes what we *would* treat as sentences of our language, if only we weren't finite, fallible beings. Just as performance is a personal-level phenomenon, competence is a capacity of an idealized person, unencumbered by the usual sorts of limitations and constraints.⁵ A theory of competence, on this view, stands to actual human speech as the ideal gas law stands to the behaviour of actual gasses. It should be clear what is wrong with this understanding of CP. If competence is a psychological system, while performance is behaviour, there is no sense in which one is an idealized version of the other. No one would suggest that the computational principles of binocular vision are an idealization of the scenes that we actually visually perceive! Much concern over Chomsky's proposal stems from this misunderstanding. If competence is simply idealized performance, one can worry that it is *too* idealized, that focus on mathematical elegance in its statement has overtaken concern for empirical reality (Tomasello (1995)). And while of course it is always possible to let the elegant be the enemy of the true in scientific theorizing, once it is accepted that a description of competence is not even purporting to be a description of performance, and is no more essentially idealized than any other scientific description of a concrete system, the force of such worries should evaporate.

While the quotes above from *Aspects*, and those in numerous subsequent statements, indicate that Chomsky is not blameless in the propagation of this mistaken understanding of CP, it is understandable that these early writings conflated the distinction between performance and competence with the conceptu-

⁵See e.g. Danks (2013): "Roughly speaking, a competence theory aims to characterize what people are capable of doing, while a performance theory aims to describe what they actually do." (p. 2126). As should be clear, on my interpretation of this distinction, 'competence' picks out a structured capacity of a mental subsystem, not the behavioural capacity of a whole organism.

ally quite different distinction between actual speech and some idealized version of this. Early approaches to generative linguistics, especially transformational grammars, did posit competences which produced things that looked quite close to actual utterances. As Chomsky (1957/2002) puts it, “The grammar of L will thus be a device that generates all of the grammatical sequences of L and none of the ungrammatical ones.” (p.13). Such a system thus generated items with the characteristic properties typically assigned to utterances (phonological features, grammatical structures, semantic interpretations, etc.). It was thus reasonable to compare the sets of structures generated by such systems with those actually produced by human speakers, and ask what the relationships were between these two, observing that the latter seem to be a finite, and somewhat distorted, reflection of the former. Competence, on such a view, was quite closely tied to performance, and the deviations were taken to be matters of “false starts”, limitations on memory, and so on.

It is important to note, however, that there is no *a priori* reason why competence and performance must stand in such a close relation. Performance is causally influenced by many, many cognitive systems and processes in addition to competence, and thus there is every reason to think that our linguistic behaviour can be quite unreflective of our competence with our language. And indeed modern generative approaches to language, particularly Minimalism, do posit significant influence of factors other than competence on performance. Minimalist grammars, for example, are typically understood as generating structures that are hierarchically arranged, but which do not determine the surface word order of the sentences they represent. This latter is then determined by the interaction between the grammar and the motor-control system (see discussion in Chomsky et al. (2019), Burton-Roberts & Poole (2006), but see Kayne (2018) for dissent). This exemplifies the way that distance between competence and

performance need not merely consist in errors or imperfections. Much more on this later.

One further source of confusion in this area stems from the relation between Chomsky's distinction between competence and performance and the distinction he draws in Chomsky (1986) between an I-language and an E-language. I-languages, Chomsky tells us, are internal and individual (i.e. are aspects of a speaker's psychology), while E-languages are external (i.e. public).⁶ As Chomsky uses the terms, 'I-language' and 'competence' are equivalent, both referring to the internal psychological system specifically dedicated to allowing the acquisition and use of language ('The language faculty in the narrow sense (FLN)' is another synonym often used here). 'E-language' is not, however, equivalent to 'performance'. Specifically, public languages, like English or Chipewyan, are paradigmatically designated as E-languages, but it is allowed that many of the performances of competent speakers of these languages will not be included within them (false starts, mistakes, deliberate violations of rules, etc.), and that many expressions licensed by these languages will never be found in any speaker's performance. The real trouble, however, stems from Chomsky's claims that I-languages are *intensional*, in that they are individuated not merely by the set of expressions which they generate, but by the procedure which generates them, whereas E-languages are *extensional*, in that there could not be two distinct E-languages which specify the same expressions.⁷ This is fair enough, as far as it goes. But it is liable to suggest some intuitive, but mistaken, infer-

⁶Chomsky also claims I-languages are individuated individualistically, so that which I-language a speaker has does not depend on features of the speaker's environment such as their causal relations or linguistic community. Burge (1989, 2003) presses Chomsky on this point, arguing that we should be anti-individualists, or externalists, about languages, even considered as the object of cognitive scientific approaches to linguistics. This debate is tangential to my concerns in this paper.

⁷This has opened up Chomsky to the notorious 'Martian argument', which runs on the intuition that a non-human alien could speak English (or Chipewyan, or whatever) even though it was psychologically radically unlike us (see Katz (1977, 1980, 1984), Soames (1984), and Devitt & Sterelny (1989)).

ences. Centrally, it suggests that I-languages intensionally specify the extensions characteristic of an E-language. But for all the reasons described above, this is empirically unlikely to be the case.⁸ What I-languages generate are hierarchically structured representations, not public expressions, and there are strong reasons to think that no simple mapping, no matter how idealized, will connect the two. This mistake is further compounded if E-languages are conflated with performance, as of course much of the reason for distinguishing competence (and therefore I-languages) from performance was to deny that all aspects of performance correspond to some structure licensed by competence.

The overall point I have been making, concerning the in principle arbitrarily large gap between competence and performance brings me on to the standard objection to the positing of such a distinction. Namely, that doing so makes empirical linguistics impossible. If, so the argument goes, performance can be dissociated from competence, how can we empirically confirm theories of the latter on the basis of evidence from the former? Performance is, especially in the generative tradition, the point of contact that linguistic theories make with observable reality. While it makes sense to view competence as an unobservable theoretical posit, the standing of such a posit depends on how well it predicts observations, and these observations will, in general, be evidenced in performance. So, by severing the connection between competence and performance, don't we thereby sever the connection between theory and empirical prediction, rendering our theories unscientific? This worry has been pressed by opponents of the generative program for decades, usually in terms of either *immunizing* a grammatical theory from empirical disconfirmation (e.g. Christiansen (1994) pp.18-19) or of rendering the theory *unfalsifiable* (Ibbotson & Tomasello (2016)).

Putting the point in this last way should enable the reader to see the proper

⁸It is plausible that the central argument of Devitt (2006) is premised on this misunderstanding.

response to this worry. It is a commonplace within philosophy of science that empirical observations relate only very indirectly to any particular scientific theory. In general, observations will provide useful information about underlying, unobservable structures and entities only under very specific conditions, specifically when confounding influences are controlled for. As Hacking (1983) forcefully argues, the most compelling work in the sciences does not try simply to account for the mass of our unguided observations, but seeks those rare, and typically artificial, opportunities when our observations might reflect some unique cause of scientific interest. Our theory of grammar needs no more to be “insulated” from the vagaries of everyday speech than our theory of electromagnetism needs to be insulated on account of failing to predict when the filament lightbulbs in my car’s headlights will give out. These are simply not the kinds of observations serious science deals with. Generative linguistics thus perfectly reasonably restricts attention to a fairly small subset of linguistic expressions for which performance (e.g. judgements of acceptability or unacceptability) might be capable of clueing us in to the underlying principles at work. That some performance datum is not as one might predict on the basis of a specific theory of competence *may* show that the theory needs changing. But it may just as well show that the, usually unstated, assumption about the relationship between competence and performance is too simplistic. One can, of course, argue that the posits made by generative linguists are unnecessary for explaining any given data, or that they are incapable of handling some other data. But such assessments are, of necessity, holistic assessments of a battery of assumptions concerning competence and other cognitive systems, never simply of the competence theory alone.

So, appeal to the competence/performance distinction is simply a reflection of the impulse to identify underlying, explanatorily important aspects of the world, and to do so on the basis of such systems’ observable effects, while recognis-

ing that any such inferences will be fallible, and will require complex reasoning concerning causal confounds and other distortions.

While, as noted above, application of this distinction is sometimes viewed as problematic, it is worth stressing the extent to which *all* linguists make, often tacit, use of it. This is for the simple reason that it is obvious that much of what we say does not reflect any underlying competence in, or knowledge of, our native language. Whatever one's views of the psychological status of linguistics, or of the correctness of particular generative approaches, there are whole swathes of our linguistic behaviour that would undermine any linguistic theory that tried to account for them. Uncontroversially, human speech is massively error-prone. We use the wrong words (“...the decision of one man to launch a wholly unjustified and brutal invasion of Iraq ? I mean Ukraine” -George W. Bush, 2022), use non-words (“Despite the constant negative press covfefe” -Donald J. Trump 2017), use constructions with no grammatical basis (“I was a bit hazy what I would find at Peppa Pig world.” - Boris Johnson 2021), and so on. Further, we are perfectly able to deliberately produce sentences we know not to be licensed by our own native tongues. This can happen when we playfully violate the rules of our language (“I wonder what it may be o'clock” - Jerry Fodor), when we want to examine an ungrammatical sentence for linguistic reasons (“The child seems sleeping” -Noam Chomsky), or when we use rote-memorized sentences of languages we have no mastery of (“Dve pivi, prosim” -Gabe Dupre). These cases are uncontroversial because the speakers themselves recognize that their utterances do not conform to their language. Things get more controversial when linguists argue that there are theoretical reasons for excluding some linguistic behaviour from the evidence base of a linguistic theory, when this conflicts with speakers' judgements about what is and is not part of their language. But the underlying reasoning is the same: linguistic behaviour is not a perfect reflection

of the language. So, what is debatable is where the boundaries lie, not whether the distinction between competence and performance is a legitimate one.⁹

While the CP distinction is typically discussed in the context of syntax, it is just as applicable to other branches of linguistics (e.g. phonology, morphology, semantics¹⁰) where we wish to distinguish between the rules governing the formation and/or interpretation of an expression and the features of linguistic behaviour.

To give an example that may appeal specifically to philosophers, consider Stojnić (2021). Stojnić develops an account of the semantic rules governing the interpretation of context-sensitive expressions ('she', 'this', etc.) according to which the content (the contribution to sentential truth-conditions) of these expressions is determined, not, as widely assumed, by domain-general reasoning involving world-knowledge and speaker-intentions, but rather by linguistic rules governing the available topics of discourse. It is a prediction of her view that utterances involving deictic pronouns are "well-formed" only when the referents of these expressions are introduced into the discourse by the means of linguistically-specific conventions, such as pointing gestures, not merely on the basis that the extra-linguistic context makes these referents salient or makes the speaker's intention to refer to them apparent. Stojnić notes, however, an apparent range of counter-examples to her view: "One can use (26) [the sentence: "She is happy"-GD], sometimes, perhaps for a surprising effect, even when such resources [linguistic conventions for licensing the use of deictic pronouns] aren't in place, as

⁹Opponents of the CP distinction often seem to realize this, but fail to draw the correct conclusion. E.g. Reich (1969) (p. 832, fn. 2) in making the charge that Chomsky's distinction between grammaticality (generability by the competence system) and acceptability (a measure of performance) precludes empirical testing, says "I exclude from natural language text sentences dreamed up by linguists, psychologists, English teachers, and poets."

¹⁰Whether we can add pragmatics to this list will depend on our theoretical analysis of what our pragmatic capacities involve. If we view pragmatics as *rule-governed*, then we will likely need to appeal to CP to exclude the kinds of exceptions to the rules (whether in error, deliberately, or unknowingly) that we have highlighted for other cases. If, however, our pragmatic abilities are essentially matters of personal-level, all-things-considered, reasoning, then there may be no pragmatic *competence* in the sense defined in this paper.

for instance, would be the case if (26) were the first sentence of a novel, or the first thing I uttered upon having met you for the first time.” (p. 49). However, she describes such cases as “irrelevant”, noting that the way in which we identify referents in these cases “is markedly different from the seamless interpretation of pronouns on the fly that we see in normal circumstances. The potentially open-ended reasoning about what the speaker wanted to convey kicks in precisely after one is faced with the infelicity of the utterance,” thus treating these examples not as instances of genuine semantic interpretation, but rather of “post-semantic repair” (p.50). This exemplifies the kind of reasoning, in line with CP, that I have been describing. While Stojnić allows that we can, as a matter of performance, use deictic pronouns without the associated reference-determining linguistic conventions, she denies that this fact bears on her theory of semantic competence, on the grounds that it invokes non-target cognitive capacities (“repair”). She further notes that this is not a matter simply of insulating her theory from apparently unaccommodating data, but that there is a testable, empirical signature to drawing the distinction in this way: these anomalous examples are not interpreted “seamlessly” in the way that, for example, pronouns with accompanying gestures are. Given the complexity of human communicative goals, and the fact that human speech is a species of creative yet rational behaviour, I believe that any theory of the rules governing human language will need to make analogous appeal to the CP distinction to account for such apparent exceptions.

3 Competence and Competences

The applicability of the CP distinction to the distinct domains within linguistic science suggests that we ought not speak of linguistic *competence* in the singular, but rather of linguistic *competences* in the plural. There are both empirical and terminological questions here, which are often unhelpfully conflated. In this

section, I shall do a little ground-clearing.

Chomsky (1980) defines ‘grammatical competence’ as “The cognitive state that encompasses all those aspects of form and meaning and their relation, including underlying structures that enter into that relation, which are properly assigned to the specific subsystem of the human mind that relates representations of form and meaning. A bit misleadingly perhaps, I will continue to call this subsystem ‘the language faculty’.” (p.59). One reason this might mislead stems from the use of the definite article. It is, it appears, quite unlikely that there is a “specific system of the human mind” that does all of the work needed to relate all the aspects of human language needed to connect a public sign to a meaning. Despite significant changes to Chomsky’s views concerning the architecture of the language faculty, one abstract commitment has been retained throughout: these approaches are *interpretational*, so that the phonological and semantic properties of an expression are constructed on the basis of the hierarchical structures licensed by the syntactic system.

On current Minimalist assumptions, basic recursive structure-building is accounted for by the operation Merge, combining two linguistic items into one larger structure, and these complex expressions are ‘interpreted’ by the phonological and semantic systems, generating further representations usable by other cognitive systems, specifically, the perceptual or motor systems (in the case of phonological properties) and the conceptual systems (in the case of semantic properties).¹¹ This, however, raises the question of the nature of these ‘interface’ systems, which serve to connect linguistic representations to the rest of our

¹¹I leave morphology out of this discussion, on the grounds that its placement in this computational system is most controversial. Traditional views treated morphology as operating “prior to” syntax (“in the lexicon”), with morphology generating words which could then be grammatically combined. However, there are also a range of apparently morphological operations which seem to operate post-syntactically, enacting a mapping between grammatically complex linguistic expressions and phonology. Some approaches (e.g. Harley & Noyer (2014)) have, for this and related reasons, rejected the distinction between syntax and morphology, arguing that there is no need to view the operations responsible for generating words as distinct from those responsible for generating phrases and sentences.

cognitive systems. Specifically, we can ask: what are the features of the representations that these systems operate on, and what are the natures of the operations sensitive to these features? Answers to these questions will then allow us to ask whether these features and operations are uniquely found in human psychology or are found also in non-human minds, and whether they are unique to linguistic cognition or whether they are utilised also by non-linguistic systems. Related questions arise concerning the nature of the ‘atoms’ of syntactic computation, the lexical items, functional and substantive, that serve as the basic entities to which Merge initially applies. What are the properties or features of such lexical items, and to what extent are they unique to humans or unique to language? These are all, of course, empirical questions.

The most ambitious hypothesis within generative linguistics, the so-called ‘Strong Minimalist Hypothesis’, views all systems beyond the core grammatical system as either not specific to language or not specific to humans, viewing Merge as the *only* human-, and language-, unique cognitive tool (Hauser et al. (2002)). The distinctive characteristics of the phonological and semantic systems of human language users would then have to be accounted for solely by the expanded range of structural representations Merge makes available as inputs to these systems. Weaker claims would allow that these interface systems have antecedents in other species, but perhaps are adapted or re-organized in the human mind. Berent (2013) and Samuels et al. (2016) present cases for such a position with respect to the phonological system. There is less work in these directions with respect to semantics and the lexicon, but it remains a promising avenue. A very different perspective on these issues can be found in Jackendoff’s ‘parallel architecture’ (Jackendoff (2010, 2002); Culicover & Jackendoff (2005); Jackendoff & Audring (2019)), which posits distinct systems for syntax, morphology, semantics, and phonology, which each generate structures in tandem,

in line with constraints ensuring that these representations cohere sufficiently.

The terminological question of what we call ‘competence’ is distinct from, but entwined with, these questions. The above quote suggests two different kinds of answers here. According to the first, we reserve the term ‘competence’ for just those systems that are genuinely specific to language. On Chomsky’s proposal that this is limited to ‘narrow syntax’, or the computational syntactic system applying Merge, this would exclude from competence any of the workings of the post-syntactic phonological and semantic systems. This would have surprising repercussions. For example, as noted above, Chomsky et al. (2019) proposes that Merge itself does not determine surface word order. Thus, the difference between speakers of SVO languages like English and speakers of SOV languages like Punjabi or Armenian would not consist in a difference in their linguistic competence, strictly speaking. If we assume in addition that all features of linguistic behaviour not attributable to competence are thereby performance effects, this has the somewhat counter-intuitive result that word-order is a performance effect.¹² An alternative is allowing ‘competence’ to apply to the whole suite of cognitive systems which enable the mind to relate semantic and phonological representations, adding these interfaces to the narrow syntax. This would allow a wider range of data to be explained by appeal to competence. But it might also suggest some mistaken assumptions concerning the unity of the underlying system. It would not, however, eliminate the distinction between competence and performance entirely. Phonological representations are just that, representations, and thus must be processed by the motor system before a public, perceptible utterance is produced. The difference between these internal representations (products of competence) and these public objects (performance) will remain. Likewise, as noted above, many intuitively linguistic utterances (e.g. deliber-

¹²Note that this conclusion would not follow for someone like Jackendoff who views all these linguistic systems as linguistically-specialized and specific to humans.

ately ungrammatical utterances, or utterances in an unknown language) will introduce into the performance record entities which may bear little resemblance to the representations licensed by the competence system, even under this more inclusive construal. Of course, the terminological decision concerning how to use the word ‘competence’ is ultimately a matter of convenience. What matters is that it is made explicit what sort of system one is appealing to in proposing an account of some linguistic phenomena: does this reflect the workings of a core, language-specific system, of the interpretive systems dedicated to enacting a mapping between grammatical structures and their phonological or semantic interpretations, or is it a product of broader cognitive functions made use of in, but not specific to, language use?

4 Competence, Performance, and Philosophy

Having, I hope, spelled out what the CP distinction is, and is not, and further sketched a picture of human linguistic capacities which situates the different systems involved in the use of language, I now turn to a range of topics within philosophy which seem illuminated by the understanding of this distinction just provided. In the space available, I won’t nearly be able to do full justice to these topics, but hope to at least indicate why philosophers working in a wide range of fields could benefit from getting clear on the ways that competence and performance differ, in the linguistic case and beyond. The core point will be that the CP distinction provides some empirical purchase on the ancient distinction between appearance (performance) and reality (competence).¹³ Failure to ap-

¹³This terminology is not perfect, suggesting as it does that appearance/performance is in some sense ‘unreal’, and carrying as it does some significant historical baggage. As I use the term, I mean only to signal the distinction between those (perfectly real) phenomena that we are able to observe, and the posits of a developed, realist, scientific theory. Sellars’ distinction between manifest and scientific images of the world could play a similar role, but comes with baggage of its own.

preciate the seemingly sizable gap here can undermine various discussions of the relations between language, thought, behaviour, and the outside world.

Most immediately, these points will be relevant to work within the philosophy of language. The gap between competence and performance calls for explicitness in exactly what is meant by ‘language’ in philosophical discussions. Traditionally, philosophers have assumed a notion of language according to which language is essentially an inter-personal, communicative phenomenon.¹⁴ Public symbols, produced by a speaker to be perceived and interpreted by a hearer, are, of course, performance phenomena. As we’ve seen, the linguistic behaviour of producing an utterance is not a straightforward reflection of the underlying representations licensed by the competence system (or systems), and thus philosophers ought to be explicit about which they are talking about. If philosophical theories of meaning, communication etc. are intended to be extensions of scientific theories of the language faculty, it will be necessary to show how the gap between competence and performance can be bridged.

One area in which this distance is easily under-appreciated is in philosophical discussions of the nature or foundations of natural language meaning, i.e. in Metasemantics. One major tradition within this literature stems from Lewis (1975, 1969), and looks to ground the meaning of natural language expressions in conventions, (very) roughly: regularities rationally adopted by a community in order to co-ordinate with one another.¹⁵ On this account, the word “typewriter” means *typewriter* because a community of people (roughly: English speakers) has a need to convey thoughts about typewriters to one another, and by each assuming that others will use the term ‘typewriter’ when this need arises, the

¹⁴This is the view assumed or argued for by many of the formative influences on modern philosophy of language, such as Frege (1956), Wittgenstein (1959/2009), Dummett (1986, 1981, 1991), and Millikan (2003, 1998, 1984). And the view that human language is essentially grounded in social conventions, discussed in the next few paragraphs, is perhaps the dominant modern framework, developed in different ways by, among others, Devitt (2021) and the burgeoning research program stemming from Skyrms (2010).

¹⁵See Rescorla (2019) for fuller discussion.

fulfillment of this need can be facilitated. The core of this picture posits a system of conventions which ensure that different speakers assign the same meanings to the same “strings of types of sounds or of marks” (Lewis (1975) p.3). While there has been much discussion of the nature of the first item of this pairing (are speakers co-ordinating on public referents of expressions, or on private mental states expressed by expressions, etc.? See Ball & Rabern (2018) and Burgess & Sherman (2014)), the second (“strings of types of sounds or marks”) has been largely taken for granted (although see Feinsinger (2021) for a recent exception).

The difficulty that arises with such views is that there is a mismatch between what they take to be basic, and what linguistic theories of competence do. “Sounds or marks” are basic elements out of which Lewisian linguistic conventions are formed. But they are not basic from the perspective of generative theories of language. They are rather performance products of underlying syntactic, morphological, and phonological systems. This fact may be overlooked stemming from an analogy between formal and natural languages often assumed in philosophical discussions. In stating the syntax and semantics for a formal language, such as the predicate calculus, at least within a certain idealized range, the “shape” of the basic symbols is invariant. In a logical deduction, we can identify all the tokens of a *sign*, such as ‘*a*’ or ‘*F*’, just by looking, and we can stipulate that all the signs with this shape have the same interpretation. But this is not the case for natural languages, due to morphological processes which can systematically alter the way a given expression will be produced, for a variety of reasons. The claims made by Lewis and many other philosophers of language to the effect that the basic elements of semantics are mappings between sounds and meanings are thus, at a minimum, significant idealizations.

Consider, for example, inflectional morphology. English speakers who wish to convey a thought about *running* have a range of options open to them: they

can use the words ‘run’, ‘running’, ‘ran’, or ‘runs’. Which option they choose will depend on the sentential context in which they wish to use it. What, on a convention-based account, is the “sound or mark” which is conventionally related to the meaning here? Note that it will not do to posit four *separate* conventions here. For one thing, this would fail to capture the relations between these expressions and the regularities that they exemplify. For another, it is not clear that all of these will count as conventions, on at least some accounts, on the grounds that they are not *arbitrary* solutions to a co-ordination problem, as required by Lewis, nor are their occurrences in language explained with reference to precedence, as on Millikan (1998)’s rival theory (speakers know that ‘running’ is the gerundive form of ‘run’ even if they have not heard others use it). These kinds of phenomena are widespread in even morphologically impoverished languages like English.

Processes like assimilation, wherein speech sounds influence one another’s expression, provide further case studies. For example, “ten pence” will typically be pronounced as “tem pence”, the difficulty of pronouncing an alveolar consonant immediately preceding a bilabial resulting in the speech system “taking a shortcut” and transforming the alveolar \n into a bilabial \m . Again here, questions arise for a convention based approach: if the convention specifies that the *sound* \tɛn conventionally maps onto the meaning TEN, then why don’t we balk at the apparently unconventional use of \tɛm to express this concept (or refer to this abstract object, or whatever)?

I don’t claim that these kinds of phenomena are fatal to mainstream philosophical approaches to language and communication. But they do indicate that focus on aspects of performance, such as the specific sound-type used in speech, can result in over-simplified theories, masking underlying complexity. It is of course reasonable to idealize in theory-construction, and take on one problem at

a time, but if this is what philosophers are doing, they should make this clear and perhaps think more about how these theories could be de-idealized in ways consistent with what is known about linguistic competence and the relations between competence and performance.¹⁶

The significance of the CP distinction is not, however, limited to the study of language. Various theories or approaches within the philosophy of psychology seem likewise to be problematized by proper attention to the distance between competence and performance. This will be particularly true of broadly empiricist approaches, exemplifying the peripatetic axiom that “nothing is in the intellect that was not first in the senses”. If it is correct that linguistic behaviour, which can be perceived, is generally unreflective of linguistic competence, which cannot, then it is doubtful that accounts of this nature will be amenable to integration with work from the linguistic sciences which focus on competence.¹⁷

Currently popular ‘predictive processing’ or ‘predictive coding’ approaches (e.g. Clark (2015)) seem subject to concerns regarding the relationship between language as public performance and as internal competence. On these views, the mind is understood to function, at root, to minimize predictive mistakes concerning its sensory interaction with the environment. As Clark (2013) puts it “Perception, action, and attention, if these views are correct, are all in the same family business: that of reducing sensory prediction error resulting from our exchanges with the environment.” (p. 21). Clark intends for language to be amenable to an account along these sorts of lines (see e.g. Clark (1998, 2005, 2006)). Language acquisition, on such a view, would amount to the construction of a mental model of the linguistic environment which enabled sufficiently accu-

¹⁶Note also that non-conventionalist alternatives to metasemantics, such as a broadly Gricean, intention-based account, seem subject to similar sorts of concerns. See e.g. Harris (2017), who argues that speaker-intentions are unsuited for grounding reference, given that aspects of linguistic competence which play a role in determining the referential properties of our utterances are unconsciously deployed within linguistic competence.

¹⁷See Dupre (2022a) for an extended argument to this effect.

rate and precise prediction of encountered linguistic stimuli. If the above picture, according to which performance is liable to be quite unreflective of competence, is correct, however, there are reasons to be skeptical of such approaches.

As I have stressed, what can be *predicted* must be observable, and thus a predictive approach to language will involve the (minimization of error in) prediction of linguistic *performance*. It is not impossible that the best way to predict performance would be to generate an internal model of linguistic competence, but given the apparently very noisy and messy relation between these two, it is far from obvious that it would. This problem is likely to be exacerbated by the ways that linguistic theories are shaped by observations very unlikely to play any role in shaping the sensory predictions of a given language user. For example, linguistic theories of a particular language are often shaped by phenomena discovered in other languages (see Reiss this volume), or by linguistic phenomena almost never encountered in everyday speech (such as parasitic gaps). These influences seem, *prima facie* to expand the gap between competence and performance, and thus pose problems for views according to which a grammar “earns its keep” in the mind on the basis of its ability to predict linguistic performance. The defender of predictive models of cognition thus appears to face a dilemma. If the picture of natural language competence, and its indirect and partial reflection in linguistic performance, suggested by generative linguistics is on the right track, then an empirical case must be made to show how a predictive mind could acquire a grammar of this sort despite sizable gaps between the grammar itself and the observations it purportedly functions to capture. Clark (2015) (p.19) seems to endorse a view along these lines, saying that “One way to learn a surprising amount about grammar (and lots more too) is to look for the best way to predict the next words in a sentence.” But until a fleshed-out model of this learning procedure is produced, I think it reasonable to remain sceptical of the hypothesis

that the kinds of grammars posited by generative linguists provide optimal ways of predicting linguistic input. The other option, endorsed by Christiansen & Chater (2015, 1999), is to argue from the fact that alternative ways of thinking about grammars provide better models for predicting linguistic stimuli to the rejection of the kinds of grammatical theories endorsed by generative linguists. Of course, such an approach thus commits a theorist to providing alternative explanations for the kinds of linguistic phenomena uncovered and explained by generative linguistics, including developmental and cross-linguistic regularities, but without appeal to the theoretical tools that have been argued to play essential roles in just these explanations. Again this is not taken to be a knockdown argument. Both projects are active areas of investigation. But if we are to endorse such models as general accounts of mind, attention to the CP distinction motivates specific research questions.

Similar concerns arise more practically in the domain of Artificial Intelligence, specifically within the development of Language Models. A Language Model specifies the probability distribution of strings defined over an alphabet, on the basis of a corpus. For a given string, then, a language model can specify how likely it would be to occur in the relevant corpus. And further it can specify how likely subsequent strings would be as continuations of this initial string. So, for example, a functional Language Model trained on a corpus of English utterances could tell you that “Do you have a sister?” is a more likely string than “Have you a sister?”. And that “Yes” is a more likely continuation of the former than “November 9, 1989”. The abstract similarity between the task of a grammar of distinguishing well-formed from ill-formed structures and that of a language model of distinguishing likely from unlikely strings has proven highly suggestive over the years. This has been exacerbated in recent years, with the substantial progress made by very large Language Models, such as GPT-3

or LaMDA, in producing often highly plausible imitations of human linguistic interactions. In a presentation of the first version of LaMDA (Language Models for Dialog Applications, developed by Google), Thoppilan et al. (2022) discuss these substantial strides, and optimistically infer that “A path towards high quality, engaging conversation with artificial systems that may eventually be indistinguishable in some aspects from conversation with a human is now quite likely.” (p. 18).

Of course, taken solely as engineering projects, aimed at developing a piece of kit capable of solving some practical problem, such as translating texts between human languages, replacing human employees in resolving complaints for an online retail company, or producing legal documents, these products may well be quite successful. Machines need not internalize language in the same way that humans do in order to behave in suitably prescribed ways in these sorts of linguistic tasks. The worry arises with attempts to infer from human-level, or better, performance in these tasks to claims that such systems have acquired human-like linguistic competence. There are two ways such an inference could go. The least radical, from the perspective of theoretical linguistics, infers from the fact that these systems produce similar linguistic strings to humans to the claim that they have acquired the kinds of systems and structures posited by linguists in the human case.¹⁸ More radically, one can argue in the other direction: assuming that these systems do not instantiate the kinds of grammars that generative linguists have posited, that such systems are capable of achieving human-level performance on linguistic task shows that these grammars are not necessary for explaining human linguistic capacities (see Piantadosi (2023)). For reasons similar to those discussed above in the case of predictive models of mind, I think there are strong reasons to doubt these proposals.

The analogy between grammars, in the sense of the target of theoretical lin-

¹⁸See Millière, this volume, for arguments pointing in this direction.

guistic inquiry, and Language Models, relies on a blurring of the distinction between competence and performance. A grammar is a description of competence, and when a structure is licensed by a grammar, this means it is generable by the underlying language-specific cognitive system. This, on its own, tells us nothing about how such a structure might relate to an utterance, nor whether an utterance corresponding, in some way, to such a structure has any particular likelihood of being produced. When a Language Model, on the other hand, specifies the probability of a string occurring in a corpus, it is precisely aimed at predicting how likely some piece of performance data will be, on the basis of previously encountered performance data. These models thus, of necessity, confound the influences on performance stemming from linguistic competence with myriad other influences. For example, while “Do you have a sister?” is more likely to occur in a corpus than strings like “Have you a sister?”, which are precluded by some aspect of English speakers’ (broad or narrow) competence, they are also judged more likely than strings like “Do you have more than one sister?” (on the basis that longer sentences are less likely than shorter sentences) and “Do you have a transistor?” (on the basis that ‘sister’ is much more frequently found in the corpus than ‘transistor’). Further, it is likely that many perfectly legitimate expressions, from the point of view of the competence, will be deemed less likely than expressions which ought be excluded from this perspective. For example, very long, or very implausible, or very hard to parse (e.g. center-embeddings) sentences will reasonably be deemed very unlikely, whereas stock constructions (“I kid you not”) or expressions from other languages (“Je ne sais quoi”) which are grammatically illegitimate might be found relatively frequently. Given the aims of the developers of these systems, it is important that these ungrammatical (or, perhaps better, extra-grammatical or para-grammatical) expressions are treated as more likely than many perfectly grammatical ones. A system which

produced the string “I kid you not” is much more likely to be “indistinguishable in some aspects from conversation with a human” than one which used grammatical but communicatively useless sentences like “The candidate the party meeting the pub I drink in hosted elected lost.” In virtue of their central focus on reproducing plausible human linguistic performance, these systems are thus destined to mark distinctions among linguistic expressions which do not correspond to those made by linguistic competence.

Because our grammatical competence is but one determinant among many of our linguistic behaviour, Language Models, and other engineering tasks aimed at reproducing linguistic behaviour, are generally improved by integrating masses of information that are irrelevant from the perspective of theoretical linguistics. It is possible that the best model for predicting linguistic performance will incorporate a true model of linguistic competence. But, as above, this is far from obvious. Thus, any attempt at inference from the behavioural successes of these kinds of models to claims about their underlying competence, in particular to claims that their underlying competences are similar to ours, will need substantial argumentation.¹⁹ Further, for some engineering projects, we may want more than the mere reproduction of human linguistic performance. It is plausible, for example, that the development of general Artificial Intelligence, or for machines with which we genuinely *communicate*, will require something like modular information processing, distinguishing between the various influences on behaviour assumed by generative linguists. Perhaps paradoxically, while lumping together all the distinct sources of influence on behaviour may make these systems better at *simulating* intelligence, it might thereby make them worse at *instantiating* it.²⁰

Finally, and perhaps most generally, this approach to linguistics, stressing the

¹⁹See Dupre (2021) for extended discussion.

²⁰Johnson & Dupre (Manuscript) present an argument along these lines.

difference between the observable behaviour which we use to test and confirm our theories, and the underlying systems we take such theories to describe, provides an interesting case study for debates within the philosophy of science.

Prima facie, the centrality of this distinction in linguistic inquiry poses stark concerns for broadly anti-realist or empiricist philosophies of science. The Chomskian line that linguists are interested *not* in accounting for the observations of linguistic behaviour, but rather in uncovering the unobservable structures partially responsible for some of these observations (namely, those plausibly thought to tell us about this specific underlying cause) seems hard to square with claims that science is ultimately concerned only with making accurate predictions, such as a Quinean instrumentalism, or that the sole standard for evaluating scientific theories is by examining the extent to which they agree with the observations, such as van Fraassen’s Constructive Empiricism. The CP distinction means that many observations, even observations of linguistic behaviour, are simply not relevant to the development of a suitable linguistic theory, and those observations that are confirmationally relevant to theory are so in virtue of an assumed connection to a specific unobservable system. Thus the very methodology of linguistics seems to preclude the observation-centric approaches common in philosophy of science, in favour of a robust, realist approach to linguistic competence.²¹ These sorts of concern are also liable to arise for any approach, such as standard Bayesian Confirmation Theory, which identifies the best confirmed theory with the theory which allows for the best prediction of empirical data. If the above discussions are on the right track, such approaches fail to grapple with the prior question: which data are relevant to the theoretical target of interest?

Even within robustly realist approaches, however, CP seems to open up new avenues. Many realists, especially “entity realists” like Cartwright (1983) and Hacking (1983) center their accounts of the methodology and metaphysics of

²¹See Dupre (2022b) for a full argument to this effect.

science on experimentation. Philosophers in this tradition endorse the general claims made above concerning the messiness of unaided observations. Their discussions then focus on the ways that scientists do not simply observe the world in the manner of Bacon's derided "mere empirics", but rather intervene in it, creating novel situations in which the causal and explanatory structure of the world which underlies this surface disorder can be better identified. It is through experimentation that we are able, often with much difficult labour as detailed by Hacking and Cartwright, to remove causal confounds and isolate the target systems of interest. The ideal of such an approach is a crucial experiment, in which some piece of the observable world is connected to the unobservable target of interest in such a way that the former will behave in a certain way if and only if the latter has a particular theoretically interesting property. A confound is then anything which undermines this biconditional. What is interesting about the CP distinction, and its application within linguistics and other parts of cognitive science, is that it seems to aim at the same thing as traditional scientific experiment—namely, the removal of causal confounds from data taken to be (dis-)confirmatory of a theory— but approaches this goal by different means.

The application of the term 'experiment' within linguistics is somewhat disputed. The standard source of data gathering in generative linguistics involves consulting one's own, and sometimes one's colleagues, intuitions about whether a given sentence or expression is acceptable or not, whether a given sentence has a particular reading, etc. This is sometimes referred to as an 'informal experiment'. On the other hand, the 'Experimental Syntax' movement (Cowart (1997)) (and 'Experimental Linguistics' more generally) explicitly opposes their approach to this methodology, arguing that to be genuinely experimental, linguists must go beyond their own intuitions and adopt more stringent experimental standards, often appealing to large sample sizes, statistical analysis, corpus research, and

so on.²² I don't wish to adjudicate on the terminological aspect of these debates. Whether we call the informal collection of linguistics judgements 'experiments' or reserve this term for more controlled and statistically sophisticated data acquisition is unimportant. What matters are the ways that acquiring linguistic judgements is and is not like experiments as discussed by philosophers of science. On the one hand, acquiring judgements can be experimental in the sense that it may involve creating a novel situation, which would otherwise not occur, and observing the result. Linguists can observe that native English speakers have judgements about sentences that would not ever arise naturally (e.g. "Colorless green ideas sleep furiously" and "Who did Kolo believe the rumour that stole a car?"), in the same way that experimental physicists or chemists can create and observe situations that would never otherwise arise. On the other hand, these linguistic "experiments" do not approximate the ideal of a crucial experiment just indicated, in that the observable behaviour is never *solely* related to the underlying phenomenon of interest. While confounds may be genuinely *removed* from certain kinds of experiment in physics and the natural sciences (e.g. by conducting some experiment in a vacuum to remove potentially confounding air resistance), this is not possible in linguistics. We cannot, as it were, directly observe the operations of the language faculty. All judgement data is performance data, and as such is a causal product of all the extra-linguistic systems used in linguistic behaviour. Reasoning in linguistics thus relies on theoretical hypotheses to estimate which aspects of our linguistic data are reflective of competence as opposed to performance.²³ This alternative approach to experimental reasoning, and the methods and assumptions used to discriminate competence-indicating signal from within the noise of performance, provides a fruitful area of future

²²See Schütze & Sprouse (2013) for an overview of these debates.

²³It is a rough heuristic of recent Minimalist work that as much as possible should be attributed to these broad performance systems, so as to leave the language-specific core relatively sparse.

work for philosophers of (cognitive) science.

5 Conclusion

While widely regarded as a core component of the methodology of generative linguistics, Chomsky's distinction between competence and performance is often taken to apply only in this narrow discipline, and even there it is viewed with some suspicion. I hope in the above to have shown that this suspicion is unwarranted. The distinction is little more than an appreciation of the distinction between the target of our theorizing and the, often noisy and confounded, data used to adjudicate between theories. I hope that my clarifying, in Section 1, exactly what the distinction amounts to, and the minimal assumptions it follows from, will enable us to move on from questions about whether it is reasonable, or scientific, to make appeals to it, and onto the substantial questions concerning how it is drawn, discussed in Section 2. I closed by selectively identifying some of the areas in which a philosophical understanding of this disarmingly simple distinction can bear on much broader questions raised across philosophical inquiry.

References

- Ball, Derek., & Rabern, Brian. (2018). *The Science of Meaning*. OUP.
- Berent, Iris. (2013). *The Phonological Mind*. Cambridge University Press.
- Burge, Tyler. (1989). Wherein is language social? In A. George (Ed.) *Reflections on Chomsky*. Wiley-Blackwell.
- Burge, Tyler. (2003). Psychology and the environment: Reply to chomsky. In

- M. Hahn, & B. Ramberg (Eds.) *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. MIT press.
- Burgess, Alexis., & Sherman, Brett. (2014). *Metasemantics: New essays on the foundations of meaning*. Oxford University Press, USA.
- Burton-Roberts, Noel., & Poole, Geoffrey. (2006). Syntax vs. phonology: a representational approach to stylistic fronting and verb-second in icelandic. *Lingua*, 116(5), 562–600.
- Cartwright, Nancy. (1983). *How the Laws of Physics Lie*. Oxford University Press.
- Chomsky, Noam. (1957/2002). *Syntactic structures*. Walter de Gruyter.
- Chomsky, Noam. (1965). *Aspects of the Theory of Syntax*. The MIT Press.
- Chomsky, Noam. (1980). *Rules and Representations*. Columbia University Press.
- Chomsky, Noam. (1986). *Knowledge of language: Its nature, origin, and use*. Greenwood Publishing Group.
- Chomsky, Noam. (2000). *New Horizons in the Study of Language and Mind*. Cambridge University Press.
- Chomsky, Noam., Gallego, Ángel J., & Ott, Dennis. (2019). Generative grammar and the faculty of language: Insights, questions, and challenges. *Catalan Journal of Linguistics*, (pp. 229–261).
- Christiansen, Morten Hyllekvist. (1994). *Infinite languages, finite minds: Connectionism, learning and linguistic structure*. Ph.D. thesis, University of Edinburgh.

- Christiansen, Morten H., & Chater, Nick. (1999). Toward a connectionist model of recursion in human linguistic performance. *Cognitive Science*, *23*(2), 157–205.
- Christiansen, Morten H., & Chater, Nick. (2015). The language faculty that wasn't: A usage-based account of natural language recursion. *Frontiers in Psychology*, *6*, 1182.
- Christiansen, Morten H., & Chater, Nick. (2016). *Creating language: Integrating evolution, acquisition, and processing*. MIT Press.
- Clark, Andy. (1998). Magic words: How language augments human computation. In P. Carruthers, & J. Boucher (Eds.) *Language and Thought: Interdisciplinary Themes*, (pp. 162–183). Cambridge University Press.
- Clark, Andy. (2005). Word, niche and super-niche: how language makes minds matter more. *Theoria. Revista de Teoría, Historia y Fundamentos de la Ciencia*, *20*(3), 255–268.
- Clark, Andy. (2006). Material symbols. *Philosophical psychology*, *19*(3), 291–307.
- Clark, Andy. (2013). Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, *36*(3), 181–204.
- Clark, Andy. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*. Oxford University Press.
- Cohen, L Jonathan. (1981). Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences*, *4*(3), 317–331.
- Cowart, Wayne. (1997). *Experimental syntax*. Sage.

- Culicover, Peter W., & Jackendoff, Ray. (2005). *Simpler syntax*. Oxford University Press.
- Danks, David. (2013). Moving from levels & reduction to dimensions & constraints. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 35.
- Devitt, Michael. (2006). *Ignorance of language*. Oxford University Press on Demand.
- Devitt, Michael. (2021). *Overlooking conventions*. Springer.
- Devitt, Michael., & Sterelny, Kim. (1989). Linguistics: What’s wrong with “the right view”. *Philosophical perspectives*, 3, 497–531.
- Dummett, Michael. (1981). *Frege: Philosophy of language*. London: Duckworth.
- Dummett, Michael. (1986). A nice derangement of epitaphs: Some comments on davidson and hacking. In E. LePore (Ed.) *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Cambridge: Blackwell.
- Dummett, Michael. (1991). The relative priority of thought and language. In *Frege and other Philosophers*, (pp. 315–24). Oxford University Press.
- Dupre, Gabe. (2021). (what) can deep learning contribute to theoretical linguistics? *Minds and Machines*, 31(4), 617–635.
- Dupre, Gabe. (2022a). Public language, private language, and subsymbolic theories of mind. *Mind & Language*.
- Dupre, Gabe. (2022b). Realism and observation: The view from generative grammar. *Philosophy of Science*.

- Dupre, Gabe. (In Press). Classifying data and interpreting theories: The case of generative grammar. In O. Bueno, & M. Martinez-Ordaz (Eds.) *From Contradiction to Defectiveness to Pluralism in Science: Philosophical and Formal Analyses*. Springer.
- Feinsinger, Ashley. (2021). The variation problem. *Philosophical Studies*, 178(1), 317–338.
- Fodor, Jerry A. (1992). A theory of the child's theory of mind. *Cognition*.
- Franks, Bradley. (1995). On explanation in the cognitive sciences: Competence, idealization, and the failure of the classical cascade. *The British journal for the philosophy of science*, 46(4), 475–502.
- Frege, Gottlob. (1956). The thought: A logical inquiry. *Mind*, 65(259), 289–311.
- Hacking, Ian. (1983). *Representing and intervening*. Cambridge University Press.
- Harley, Heidi., & Noyer, Rolf. (2014). Distributed morphology. In *The second Glot International state-of-the-article book*, (pp. 463–496). De Gruyter Mouton.
- Harris, Daniel W. (2017). Speaker reference and cognitive architecture. *Croatian Journal of Philosophy*, 17(3 (51)), 319–350.
- Hauser, Marc D., Chomsky, Noam., & Fitch, W Tecumseh. (2002). The faculty of language: what is it, who has it, and how did it evolve? *science*, 298(5598), 1569–1579.
- Hintikka, Jaakko. (1999). The emperor's new intuitions. *The Journal of Philosophy*, 96(3), 127–147.
- Ibbotson, Paul., & Tomasello, Michael. (2016). Evidence rebuts Chomsky's theory of language learning. *Scientific American*, 315(5).

- Jackendoff, Ray. (2002). *Foundations of language: Brain, meaning, grammar, evolution*. Oxford University Press.
- Jackendoff, Ray. (2010). *Meaning and the lexicon: the parallel architecture 1975-2010*. Oxford University Press.
- Jackendoff, Ray., & Audring, Jenny. (2019). Relational morphology in the parallel architecture. In J. Audring, & F. Masini (Eds.) *The Oxford Handbook of Morphological Theory*. Oxford University Press.
- Johnson, Gabbrielle., & Dupre, Gabe. (Manuscript). Uncanny performance, divergent competence: A principled barrier to human-machine communication.
- Katz, Jerrold J. (1977). The real status of semantic representations. *Linguistic Inquiry*, 8(3), 559–584.
- Katz, Jerrold J.. (1980). *Language and Other Abstract Objects*. Rowman and Littlefield Publishers.
- Katz, Jerrold J. (1984). An outline of platonist grammar. In T. G. Bever, J. M. Carroll, & L. A. Miller (Eds.) *Talking minds: The study of language in cognitive science*, (pp. 17–48). MIT press.
- Kayne, Richard S. (2018). The place of linear order in the language faculty.
- Knobe, Joshua. (2010). Person as scientist, person as moralist. *Behavioral and brain sciences*, 33(4), 315–329.
- Lewis, David. (1969). *Convention: a philosophical study*. Cambridge, Harvard university press.
- Lewis, David. (1975). Languages and language. In K. Gunderson (Ed.) *Minnesota Studies in the Philosophy of Science*, (pp. 3–35). University of Minnesota Press.

- Marr, David. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. MIT press.
- Miller, George A., & Chomsky, Noam. (1963). Finitary models of language users. In R. D. Luce, R. R. Bush, & E. Galanter (Eds.) *Handbook of Mathematical Psychology*. Wiley and Sons.
- Millière, Raphaël. (MS). Language models as models of language. In G. Dupre, K. Stanton, & R. Nefdt (Eds.) *Oxford Handbook of Philosophy of Linguistics*. Oxford University Press.
- Millikan, Ruth Garrett. (1984). *Language, thought, and other biological categories: New foundations for realism*. MIT press.
- Millikan, Ruth Garrett. (1998). Language conventions made simple. *The Journal of Philosophy*, 95(4), 161–180.
- Millikan, Ruth Garrett. (2003). In defense of public language. In *In LM Antony & N. Hornstein (Eds.), Chomsky and His Critics*. Blackwell.
- Norvig, Peter. (2017). On chomsky and the two cultures of statistical learning. In *Berechenbarkeit der Welt?*, (pp. 61–83). Springer.
- Patterson, Sarah. (1998). Competence and the classical cascade: a reply to franks. *The British Journal for the Philosophy of Science*, 49(4), 625–636.
- Piantadosi, Steven. (2023). Modern language models refute chomsky’s approach to language. *Lingbuzz Preprint*, lingbuzz/007180.
- Reich, Peter A. (1969). The finiteness of natural language. *Language*, (pp. 831–843).

- Reiss, Charles. (MS). Research methods in armchair linguistics. In G. Dupre, K. Stanton, & R. Nefdt (Eds.) *The Oxford Handbook of Philosophy of Linguistics*. Oxford University Press.
- Rescorla, Michael. (2019). Convention. In E. N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2019 ed.
- Rosenberg, Jay F. (1988). About competence and performance. *Philosophical Papers*, 17(1), 33–49.
- Samuels, Bridget D., Hauser, Marc., & Boeckx, Cedric. (2016). Looking for ug in animals: A case study in phonology. In I. Roberts (Ed.) *The Oxford Handbook of Universal Grammar*, (pp. 527–546). Oxford University Press.
- Schütze, Carson T., & Sprouse, Jon. (2013). Judgment data. *Research methods in linguistics*, (pp. 27–50).
- Simms, Victoria., Cragg, Lucy., Gilmore, Camilla., Marlow, Neil., & Johnson, Samantha. (2013). Mathematics difficulties in children born very preterm: current research and future directions. *Archives of Disease in Childhood-Fetal and Neonatal Edition*, 98(5), F457–F463.
- Skyrms, Brian. (2010). *Signals: Evolution, learning, and information*. OUP Oxford.
- Soames, Scott. (1984). Linguistics and psychology. *Linguistics and Philosophy*, 7(2), 155–179.
- Stojnić, Una. (2021). *Context and coherence: The logic and grammar of prominence*. Oxford University Press.

Thoppilan, Romal., De Freitas, Daniel., Hall, Jamie., Shazeer, Noam., Kulshreshtha, Apoorv., Cheng, Heng-Tze., Jin, Alicia., Bos, Taylor., Baker, Leslie., Du, Yu., et al. (2022). Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*.

Tomasello, Michael. (1995). Language is not an instinct. *Cognitive Development*, 10, 131–156.

Wittgenstein, Ludwig. (1959/2009). *Philosophical investigations*. John Wiley & Sons.