

3

How the Philosophy of Language Grew Out of Analytic Philosophy

Daniel W. Harris

3.1 Introduction

My task in this chapter is to tell the origin story of the contemporary philosophy of language as a chapter in the history of analytic philosophy. Rather than aiming for completeness, I will trace out several threads that combine to give a useful perspective on how the philosophy of language got to where it is now.

In § 3.2, I focus on the development of semantics, which began as a collection of methodological tools for applying formal languages to philosophical problems. In § 3.3, I trace the origins of contemporary debates about the nature of propositional content. In § 3.4, I give a brief history of philosophical work on speech acts and pragmatics, emphasizing the origins of current debates in conflicting threads of Wittgenstein's writing.

Many of the ideas to be canvassed here began their careers as tools for thinking about philosophical debates but have gradually been repurposed as part of the philosophical and scientific study of natural language. This process has been both an effect and a cause of increased collaboration between philosophers and linguists. Whereas the philosophy of language began the twentieth century as philosophy's methodological R&D department, it increasingly belongs to the theoretical wing of an interdisciplinary scientific research program. This reorientation is an overarching theme of what follows.

Before I begin, a disclaimer: my aim is to outline the major philosophical influences on the philosophy of language as it is now practiced. This must be distinguished from two other possible aims. First, I will not attempt to give an exhaustive summary of twentieth-century work in the philosophy

of language, or to allocate attention to ideas in accordance with their importance in their historical contexts. I will, for example, give short shrift to ideas that were highly influential in their day but that have been mostly cast aside. Second, I will not be concerned exclusively with what I take to be the correct readings of the historical figures whom I discuss. Misreadings sometimes have greater influence on subsequent work, and so I will attend to some of those here.

3.2 The Origins of Natural-Language Semantics

The aim of natural-language semantics is to build a computational model of how the meanings of complex expressions compose as a function of their structure and the meanings of their parts. Semantics is now a thriving interdisciplinary research program. Philosophers have contributed directly to this program since its beginnings, often in collaboration with linguists and sometimes with computer scientists, mathematicians, and psychologists. Philosophers have also tended to take a particular interest in the applications of semantics to debates elsewhere in philosophy. As semantics has developed into an empirical science, it has also become an increasingly authoritative source of empirical constraints on philosophical theorizing about language.

The most important figures in the creation of this research program were the philosophers Donald Davidson (1967b, 1967a), Richard Montague (1970c, 1970a, 1973), and David Lewis (1970), though many others made important early contributions.¹ Natural-language semantics has been done in a variety of methodological frameworks, and philosophers have played an important role in comparing the foundational assumptions of these frameworks. At present, the most influential framework is the one codified in textbooks by Heim and Kratzer (1998) and von Stechow and Heim (2011).

The early natural-language semanticists drew on a methodological tool kit that was almost ninety years in the making – one that analytic philosophers had developed to study formal languages created for various philosophical ends. What was new in the work of Davidson, Montague, and Lewis was the idea that this tool kit could be fruitfully adapted to understand natural language. Earlier philosophers had been explicitly pessimistic about the prospects of such a project. Nonetheless, it was these philosophers' work that made natural-language semantics possible. In the rest of this section I will survey some of the major contributions to this prehistory of natural-language semantics.

¹ In particular, Barbara Partee deserves considerable credit both for her own seminal contributions (collected in Partee, 2004a) and for her role in spreading the influence of Montague's framework. See Partee (2004b) for a history of early work in natural-language semantics and a memoir of Partee's role.

The semantic study of formal languages grew out of work by Gottlob Frege and Bertrand Russell on the foundations of mathematics. Their pursuit was logicism, the attempt to reduce mathematics to purely logical concepts and axioms. This project demanded a more powerful logic than those previously available, and the development of this logic required formal languages with greater expressive power than what had previously been available. The creation and investigation of these formal languages inaugurated many of the discussions that make up contemporary philosophy of language.

Frege contributed several lasting insights to the study of how the meaning of a complex expression is constrained by its structure. He changed the way that we think about sentence structure by replacing the simple subject–predicate sentence structure of earlier formal languages with a more sophisticated analysis. Each atomic sentence is built up from a predicate expression and one or more terms (names or variables), complex sentences may be built up from other sentences using connectives, and quantifiers may take scope over any sentence or subsentence, binding variables within it (Frege, 1879). Frege (1891) argued that expressions' referents combine in the way that functions combine with their arguments: sentences refer to truth-values, names refer to objects, predicates refer to concepts (which Frege identified with functions that map objects to truth-values), quantifiers refer to functions that map concepts to truth-values, connectives refer to truth functions, and so on. As we would now put it, Frege organized the referents of expressions in different grammatical categories into a hierarchy of semantic types.

Of Frege's ideas about expressions' semantic types, perhaps the most enduring and influential has been his view that quantifiers refer to second-order functions – functions that contemporary semanticists categorize as type- $\langle et, t \rangle$ functions.² For example, in asserting *every philosopher is a scholar*, what we are doing is asserting that the second-level function denoted by *every philosopher* is one that maps the first-level function denoted by *is a scholar* to truth. By way of tweaks and refinements at the hands of Mostowski (1957), Montague (1973), and Barwise and Cooper (1981), this idea has become the basis for generalized-quantifier theory, which has been one of the most fruitful areas of natural-language semantics.³

The general picture embodied in Frege's ideas about semantic composition has been so influential that Heim and Kratzer (1998: 1) open their textbook by describing their project as “the Fregean program.” As most semanticists see it, their main goal is to reverse-engineer the function by which the meanings of complex expressions can be computed from their

² This way of typing functions is a variant of the notation used by Alonzo Church (1940, 1941) in formulating his typed lambda calculus. Church's lambda notation has become a crucial part of the metalanguage in which contemporary semantics is done.

³ See Peters and Westerståhl (2006) for a survey of work on quantifiers and Partee (2013) for a history of their “starring role” in natural-language semantics.

syntactic structure and the meanings of their parts. The claim that such a function exists – a claim that can be made precise in various ways – is the principle of semantic compositionality.⁴ Philosophers and linguists at least since Carnap (1947: 121) have credited Frege with formulating this principle, which has sometimes been called “Frege’s principle.” Although there can be little doubt that something like compositionality was at work in Frege’s thinking about his formal language, there has been scholarly debate about which version of the principle, if any, he believed, and how explicitly (see Pelletier, 2001; Neale, 2001; and Dever, 2003: §2). Even if attributing the principle to Frege involves a misreading, however, this misreading would constitute a momentous influence over the history of semantics.

A second major influence on contemporary thinking about the mechanics of semantic composition arose from Russell’s (1905) theory of descriptions. Russell argues that what he calls “denoting phrases” – phrases constructed from a determiner and a noun phrase, as well as words like *everything* and *nothing* that appear to play the same semantic role – do not “have any meaning in isolation, but a meaning is assigned to every proposition in which they occur” (1905: 480). For example, denoting phrases seem, on the surface, to occupy the same grammatical positions as names and pronouns, apparently as the arguments of predicates. But Russell argues that denoting phrases in fact take scope over predicates and bind hidden variables in the argument positions that they appear to occupy. Russell thus posits a layer of grammatical structure, logical form, that underlies sentences’ surface structure and that must be uncovered via analysis in order to understand how their contents are structured.

Russell applies this idea in a more extreme way to definite descriptions, arguing that the logical form of (1) is best represented (using contemporary notation) as (2):

- (1) The king of France is bald.
 (2) $(\exists x)((Kx \ \& \ (\forall y)(Kx \supset y = x)) \ \& \ Bx)$

English gloss: There exists one and only one king of France, and he is bald.

At the level of logical form, then, a definite description is broken apart into a collection of quantifiers, connectives, and bound variables, and sprinkled throughout the sentence. Russell therefore calls descriptions “incomplete symbols” – symbols that appear in the surface structure of sentences but that disappear under analysis and therefore have no unified meanings of their own.

Taking things a step further, Russell (1911: 114) argues that many apparently syntactically unstructured expressions, including most proper names, must be analyzed as definite descriptions (and then further

Erratum: this x should be a y.

⁴ For overviews of issues about compositionality, see Dever (2003, 2006) and Szabó (2017).

analyzed as above). For example, the name *Bismarck* might be analyzed as “the first Chancellor of the German Empire” (Russell 1911: 115). For Russell, then, the logical form of any given sentence may turn out to be radically different than its surface structure would suggest.

Although Frege speculated about how some of his ideas applied to natural language (e.g. in Frege, 1892, 1918), the real purpose of his ideas about the syntax–semantics interface was to increase the expressive power of his formal language. Russell (1905) shared this motivation and also saw his theory of descriptions as a solution to various logical and metaphysical paradoxes. Russell’s theories of descriptions and proper names also came to play a central role in his theory of empirical knowledge, whose aim was to show how we can have knowledge of the external world beyond what is given in sensation. According to Russell (1911, 1918), the components of the contents of our thoughts are all things with which we are directly acquainted – sense data, sensible qualities such as colors and shapes, and perhaps the self. Other entities about which we seem to think and speak are “logical constructions” out of the objects of acquaintance. What this means in practice is that when we find a sentence that contains an expression purporting to refer to something that is not given in perception, we must use Russell’s theory of definite descriptions to analyze it away. This procedure would normally have to be iterated many times, transforming most sentences into extremely complex statements about sense contents.

Russell thus motivated his ideas about logical form by tying them to a foundationalist epistemology, with results that would strike contemporary semanticists as bizarre. Nonetheless, the idea that natural-language sentences possess a hidden level of syntactic analysis that is relevant to their semantic interpretation is alive and well. One vector into contemporary work runs through Davidson (1967a), who, building on some ideas of Reichenbach (1947), argued that sentences describing actions have logical forms that feature event variables bound by covert existential quantifiers. This idea has turned out to have an enormous influence on how semanticists understand verbs, tense, and aspect in natural language.⁵ The idea of Logical Form can also be recognized in generative linguistics, where it is common to distinguish PF (short for “Phonetic Form”) from LF (short for “Logical Form”) as distinct levels of linguistic representation, the former encoding structure relevant to pronunciation and the latter encoding structure relevant to semantic interpretation.⁶ Some of the main claims of contemporary theories of LF are even recognizably Russellian – for example, in that they render the scopes of quantifiers and other scope-taking expressions structurally explicit.

⁵ See Maienborn (2011) and Williams (Chapter 20, this volume) for overviews.

⁶ Chomsky (1976, 1981); May (1977, 1985). Some contemporary linguists reject LF; see, e.g., Barker and Jacobson (2007).

Even Russell's theory of descriptions lives on, though in altered forms and divorced from his epistemology. For example, Neale (1990) defends a version of the theory in which definite descriptions are treated as restricted quantifiers rather than as incomplete symbols in Russell's sense. Meanwhile, several of the main alternative treatments of definite descriptions can be traced to Frege (1892), who argued that definite descriptions, like names, presuppose the existence of their referents and fail to refer if their presuppositions are false. The concept of presupposition, which is often credited to Frege, has now become the subject of a massive literature in semantics and pragmatics.⁷ Many still think of definite descriptions and other definite noun phrases (including names and pronouns) as triggering presuppositions of various kinds (see, e.g., Roberts, 2003).

Even Russell's view that most proper names must be analyzed as descriptions has contemporary successors, although Kripke (1980) is widely seen as having refuted the specifics of Russell's view. Most recently, Fara (2015) has argued that what appear to be syntactically simple occurrences of proper names are actually the nominals of definite descriptions that have unpronounced definite determiners, so that the LF of *Bertrand was right* could be more perspicuously represented by *the Bertrand was right*. Although this is not Russell's view, he does suggest that names should sometimes be analyzed as descriptions of the form "the person called *n*" (Russell, 1911: 119). There are, of course, many alternative theories of names, some of which I will discuss below. In general, names and descriptions have occupied an outsized role in the philosophy of language, and Frege and Russell continue to loom over these debates.

Another early twentieth-century logician whose work contributed some of the basic ingredients of contemporary semantics is Alfred Tarski, whose theories of truth (1935, 1944) and logical consequence (1936) have both had lasting influence.

Tarski's theory of truth is a recipe for assigning truth conditions to sentences in a formal object language with limited expressive power using an axiomatic theory stated in a formal metalanguage with greater expressive power. Much of this procedure is visible in most natural-language semantic theories since Davidson (1967b). In particular, semantic theories begin with axioms assigning meanings to primitive expressions in an object language, generate theorems assigning truth conditions to object-language sentences, and are normally given in a notationally enriched metalanguage. However, the superficial similarity of Tarski's truth definitions to contemporary semantic theories masks deep differences in their respective explanatory goals. Tarski's procedure begins from assumptions about the meanings of object-language expressions in order to offer a stipulative definition of a truth predicate. Modern semantic theories, by contrast,

⁷ For an overview of the literature on presupposition, see Beaver and Geurts (2014).

take the notion of truth as primitive in order to understand the meanings of natural-language expressions. As Burgess (2008: 166) puts it, “We constantly find in the writings of Davidson and disciples mentions of a ‘Tarskian’ theory of truth, where ‘counter-Tarskian’ or ‘anti-Tarskian’ would have been more accurate.”

Aside from lending a general shape to modern semantic theories, Tarski’s theory of truth also provided an influential model of variable binding.⁸ The method is this: a variable’s referent is relativized to an arbitrary sequence, or assignment function. This relativization is inherited by expressions containing the variable, up until it is bound by a variable-binding expression, such as a quantifier, which render assignment functions inert. In recent work, the role of assignment functions has been expanded. Context-sensitive expressions, including unbound pronouns, are often treated as having assignment-relativized contents, and assignments are thought of as formal stand-ins for the utterance context (Heim and Kratzer, 1998: 242–243) or the speaker’s referential intentions (Heim, 2008: 35–36).

Tarski’s theory of logical consequence, unlike the proof-theoretic accounts of earlier logicians, defines consequence model-theoretically, as the preservation of truth under arbitrary reinterpretations of non-logical vocabulary (1936). Repurposing Tarski’s ideas, Montague (1974: 188) argued that “the construction of a theory of truth . . . under an arbitrary interpretation [is] the basic goal of serious syntax and semantics.” In the hands of contemporary semanticists, Tarski’s theory has become an essential tool for empirically assessing semantic theories, since it allows predictions about logical consequence to be generated, which can then be tested against the intuitions of native speakers.

A fourth early analytic philosopher whose work continues to exert a major influence on semantics is Carnap, whose *Introduction to Semantics* (1942) synthesized and disseminated the semantic ideas of Frege, Russell, and Tarski, and whose *Meaning and Necessity* (1947) articulated much of the framework of intensional semantics that, via Montague, most semanticists still work with.⁹

Carnap’s main innovation was the use of possible worlds, which he modeled as state descriptions – maximal consistent sets of atomic sentences.¹⁰ His use of possible worlds allowed Carnap to distinguish between each expression’s intension and extension – a distinction that is inspired by, though distinct from, Frege’s sense–reference distinction (Carnap, 1947: secs. 28–30). In Carnap’s usage, the extension of a singular

⁸ It is noteworthy that the small minority of contemporary semanticists who reject a Tarskian account of variables call their theory “variable-free semantics” (Jacobson, 1999, 2014), suggesting that Tarski’s account has become synonymous with his subject matter.

⁹ Church (1940, 1946) developed his intensional logic in parallel, but Carnap’s formulations have proven to be more influential.

¹⁰ Carnap (1947: 9) credits Leibniz and Wittgenstein (1922) as the inspirations for his use of possible worlds.

term is the entity to which it refers, the extension of a one-place predicate is the set of entities of which it is true, the extension of a sentence is its truth-value, and so on. An expression's intension is a function from each possible world to its extension at that world.¹¹

Carnap's aim in developing his intensional semantics was not to understand natural language. In *Introduction to Semantics*, he had distinguished "descriptive semantics," which is the study of meaning in natural language, from "pure semantics," which is the stipulative "construction and analysis of a semantical system," and made it clear that he was interested only in the latter (Carnap, 1942: §5). But Carnap's ideas were refined over the next two decades, culminating in Montague's application to natural language. Along the way, important contributions were made by many logicians, often independently and in parallel. (For the messy details, see Copeland, 2002.)

One important advance over Carnap's model was the addition of a binary accessibility relation over worlds – an idea that was developed independently by Prior and Meredith (1956), Hintikka (1961), and – most famously – Kripke (1963, 1959). Carnap had treated modal operators as unrestricted quantifiers over all state descriptions: "Necessarily *S*" is true if and only if *S* is true at every possible world. By contrast, Kripke's models treat modals as quantifiers whose domain is restricted to the worlds that are accessible from the world of evaluation. By placing different conditions on the accessibility relation and thereby restricting modals' quantification in different ways, many different modalities can be expressed. The original point of this complication was to devise soundness and completeness results for a range of modal systems, but the same idea has turned out to have applications in natural-language semantics. Most influentially, Kratzer (1977, 1981) showed how to account for the fact that some modals can express different modalities – e.g. *must* can be a deontic or epistemic modal – by arguing that modals are sensitive to contextually supplied "conversational backgrounds" which serve to restrict their domains of quantification in different ways. Most of the enormous contemporary literature on modals takes Kratzer's work as a jumping-off point.¹² A similar story can be told about propositional attitude verbs, which, building on Hintikka's (1962) early work on epistemic logic, are now likewise often understood as modals that quantify over sets of worlds that are epistemically accessible (in various senses).¹³

The treatment of modals as restricted quantifiers over worlds also led to an ongoing explosion of work on conditionals. Two threads in the history of analytic philosophy had provided fuel. First, C. I. Lewis demonstrated

¹¹ The idea that intensions are functions from worlds to extensions is implicit in Carnap (1947), but Montague (1974: 145) reports that Carnap made the idea explicit in conversation, and this is the implementation that Montague and Lewis popularized.

¹² For a survey on the semantics of natural-language modals, see Portner (2009).

¹³ For the standard textbook treatment of these ideas, see von Stechow and Heim (2011: ch. 2).

a range of inadequacies of the truth-functional conditional of classical logic, gave an axiomatized treatment of a “strict conditional” that, he thought, did better, and used this strict conditional to axiomatize several of the systems of modal logic for which Kripke and others would later prove soundness and completeness theorems (Lewis, 1918; Lewis and Langford, 1932). Lewis thus linked conditionals to modal logic and sparked interest in both. A second tradition – exemplified by Goodman (1955) – identified a web of connections between counterfactual conditionals and a collection of weighty philosophical topics such as the metaphysics of dispositions and causation, the nature of scientific laws, and human agency. The explosion was sparked by Robert Stalnaker (1968) and David Lewis (1973), who proposed variations on the following idea: for a conditional “if A then C” to be true at a possible world w is for its consequent, ‘C’, to be true at the world or worlds that are most similar to w and at which its antecedent, ‘A’ is true. This treatment makes the conditional a kind of restricted modal operator that quantifies over a set of worlds determined by a similarity relation on worlds together with the conditional’s antecedent. Although the Stalnaker–Lewis approach is far from the only school of thought on the semantics of conditionals, it remains the default view that others attempt to either refine or challenge.¹⁴

A further important advance over Carnap’s (1947) semantics was the treatment of possible worlds as primitive elements in the model rather than as state descriptions. For Carnap, “necessarily S ” is used to make the claim that S is true in every state description, which is tantamount to saying that there is no way of reinterpreting the non-logical symbols of the language so as to make S false. The notion of necessity involved is thus linguistic, corresponding to logical truth or analyticity. Treating possible worlds as primitive elements in models opens them up to various interpretations, and allows nonlinguistic modalities to be expressed. Most influentially, Kripke (1980) took the worlds in his models to be the *metaphysically* possible worlds and argued that necessity, analyticity, and a priority are distinctively metaphysical, semantic, and epistemic concepts, respectively.¹⁵

Meanwhile, Prior (1957, 1967) devised logics in which tenses are treated as modal operators that quantify over times rather than worlds. Kripke’s and Prior’s ideas were combined and generalized by Montague (1970b, 1970c), who took intensions to be functions from *indices* to extensions, where each index includes a world, a time, and possibly other factors on which expressions’ extensions might depend. This technique would later be further generalized in a wide variety of ways. For example, Lewis

¹⁴ For a summary of work on counterfactual conditionals, see Starr (2019). For a summary of work on indicative conditionals, see Edgington (2014).

¹⁵ On the philosophical significance of this idea, see Soames (2003b: chs. 15–16).

(1979a) posits a form of essentially first-personal, or *de se* content that can be modeled as functions from a world, a time, and *an agent* to a truth-value. More recently, relativists of various stripes have argued that indices also include features that represent the standards relative to which some claims are true or false.¹⁶

Another generalization of intensional semantics has formed the basis of the most influential theories of the context sensitivity of natural-language expressions. The standard presentation of this generalization is due to Kaplan (1989).¹⁷ In addition to an extension and an intension, Kaplan's semantics assigns each expression a *character*, which can be thought of as a rule for determining the expression's content in a given context of utterance. Formally, Kaplan models characters as functions from contexts to intensions, and treats contexts as ordered tuples of entities that will be present in any real-world context of utterance and on whose identity the intension of a context-sensitive expression could depend. One of the coordinates of a context is the speaker, for example, and Kaplan models the character of the word *I* as a function that maps each context to the speaker in that context. Whereas Montague's semantics took each expression's extension to be dependent on a single index, Kaplan's semantics introduced a kind of double indexing, in which one index represents the way in which expressions' contents depend on the contexts in which they are uttered and a second index represents the way in which the circumstances in which a content is evaluated determined its extension. The idea of double indexing – sometimes called “two-dimensional semantics” – has been put to a surprising number of uses in semantics and in philosophy more generally.¹⁸ And, in general, an enormous number of natural-language expressions have been claimed to be context sensitive and given treatments along Kaplanian lines.

Take a course in natural-language semantics and you will likely be taught how to construct an axiomatic truth theory for a fragment of a natural language. Taking as its input a sentence's LF, the theory will assign meanings, modeled as functions of various types, to the sentence's simple parts and will then provide a recipe for deriving the sentence's intension by combining these functions with one another. Although I am leaving out many details and advancements, it should be clear that contemporary semantic theories are covered with the fingerprints of Frege, Russell, Tarski, and Carnap.

At the same time, I hope it's clear that the aim of the game has shifted from the stipulative construction of philosophical tools to the empirical description and explanation of natural language. Although many philosophers of language still attempt to draw philosophical conclusions from semantic theories, empirical adequacy now trumps other ways of

¹⁶ For an overview of recent work on relativism, see Baghramian and Carter (2018: §5).

¹⁷ But see also Kamp (1971) and Vlach (1973). ¹⁸ For an overview, see Schroeter (2017).

evaluating the results. This shift in explanatory goals happened slowly and not always self-consciously, but it has had widespread ramifications for the philosophy of language – most notably, a much greater integration with contemporary linguistics.¹⁹

3.3 Propositional Content

Propositional content is another topic that has been central in the philosophy of language since the late nineteenth century. Beginning with Frege and Russell, most philosophers have taken propositions to play the following roles: (i) they are abstract entities that exist independently of human activities, (ii) they are the contents of (uses of) declarative sentences, (iii) they are the contents of beliefs and other propositional attitudes, (iv) they have truth conditions and are the primary bearers of truth and falsity, and (v) they can be communicated from speakers to hearers. These properties of propositions played an important role in Frege's broader logicist project by bolstering his vehement anti-psychologism about logic. Many earlier philosophers had taken logic to be the study of abstract properties of human cognition. But if logic was the study of the objective relations between mind-independent, truth-value-bearing propositions, as Frege believed, then there was better hope for logicists of explaining the necessity and objectivity of mathematical truths. Like contemporary work on semantics, contemporary work on propositions thus began as a precondition for the logicist program.

Although (i)–(v) articulate what is standardly taken to be the theoretical role of propositions, each of these claims has been rejected in at least some cases.²⁰ And of course, some philosophers – most notably Quine (1960), Davidson (1967b), and one stage of Russell (1912, 1918) himself – have denied the existence of propositions altogether. Among philosophers who accept (i)–(v), there remains considerable disagreement.

One influential view is due to Russell (1903), who thought that the proposition expressed by a sentence is a structured entity whose parts are the referents of the expressions that make up the sentence's LF. For example, a logically proper name – i.e. a name that has not been analyzed away at LF – contributes its referent to the propositions expressed by sentences in which it occurs, and the contribution made by a one-place predicate is a property.

Frege (1892) denied that referential contents can do the work of propositions. The sentence "Hesperus is identical to Hesperus" expresses a trivial truth while the truth of "Hesperus is identical to Phosphorus" is an

¹⁹ For a more detailed discussion of this shift and its consequences, see Harris (2017b).

²⁰ See Soames (Chapter 4, this volume) and Hanks (Chapter 19, this volume) for some of the reasons to reject the idea that propositions are independent of human activities.

empirical discovery. But since *Hesperus is Phosphorus*, Frege concluded that names' contents must be something other than their referents. Likewise, it could transpire that *Beyoncé believes that Jay Z is the greatest rapper* is true while *Beyoncé believes that Hove is the greatest rapper* is false, though Jay Z and Hove are the same person. If the truth conditions of a sentence are determined by the referents of its parts, then this would seem to pose a problem for Russellians, since the two sentences differ only in that apparently co-referring names have been swapped. To account for these and other phenomena, Frege and latter-day Fregeans have argued that each expression – each word, phrase, and sentence – has both a sense and a referent. An expression's sense is a particular mode of presentation of its referent and serves as the expression's referent in "indirect" contexts, such as when embedded under attitude verbs. The senses of sentences, which Frege called "thoughts," are now usually called "Fregean propositions."

Russell (1905) denied the existence of Fregean senses. He handled Frege's data by claiming that *Hesperus* and *Phosphorus* are different definite descriptions in disguise. They therefore contribute different components to the contents of sentences in which they appear and can enter into complex scope relations that he took to explain their puzzling interactions with attitude verbs. This strategy, which depends on Russell's alliance of semantics and foundationalist epistemology, is no longer a popular one. But contemporary Russellians abound, buoyed by Kripke's (1980) arguments that names, unlike descriptions, are rigid designators, Kripke's (1980) and Putnam's (1973, 1975) arguments that natural-kind terms are rigid designators, Kaplan's (1978, 1989) arguments that indexicals and demonstratives are directly referential, and Salmon's (1986) and Soames' (1987) sophisticated defenses of Russellian propositions and directly referential theories of names.²¹

Meanwhile, the emergence of intensional logic in the wake of Carnap (1947) (see §3.2) gave rise to the view that propositions are functions from possible worlds to truth-values – or, equivalently, sets of possible worlds. The idea behind this proposal is that the essential job of a proposition is to embody a truth condition. Sets of worlds are truth conditions incarnate: they do nothing except to specify a truth-value for each way that the world could be.

The idea of possible-worlds propositions has been generalized and modified in many ways. Possible worlds are *maximal* ways that a world could be. Semanticists interested in their non-maximal counterparts – roughly, parts of worlds – tend to build their propositions out of "situations" or "truth-makers" instead.²² As I described in § 3.2, semanticists have also

²¹ For a recent summary of the debate between Russellian and Fregean theories of propositions, see Caplan (2006).

²² Barwise and Perry (1983) were the founders of situation semantics. For an overview, see Kratzer (2020). On truth-maker semantics, see Fine (2017). J. L. Austin (1950) is often credited with the idea that statements are true or false only with respect to situations.

found it useful to generalize intensions so that they become functions from more complex indices to truth-values. And so we find some theorists debating whether propositions are sets of world–time pairs (Montague, 1970c), sets of world–time–agent triples (D. K. Lewis, 1979a), or sets of some other complex indices. We can group these proposals together by saying that they treat propositions as sets of truth-supporting circumstances.^{23,24}

3.4 Wittgenstein, Language Use, and Speech Act Theory

A third major area of current work in the philosophy of language deals with pragmatics – the theory of how we use language to communicate and to perform speech acts of various kinds, and of the mechanisms by which discourses evolve. Contemporary pragmatics developed out of a mid-century movement away from the analysis of formal languages and toward naturalistic attention to the messy details of ordinary speech.

The most important source of this tradition is Wittgenstein’s late work. In the opening passages of the *Philosophical Investigations* (1953/2009), Wittgenstein criticizes a “philosophical notion of meaning” that “is at home in a primitive idea of the way language functions” (2009: §2). Among his targets is the idea that understanding the meanings of natural-language expressions is a matter of understanding what they refer to. He attributes this picture to Augustine, but also to unnamed “logicians” (think Frege and Russell) and “the author of the *Tractatus Logico-Philosophicus*” – i.e. his younger self (2009: §23). In place of this picture, Wittgenstein suggests that we should attend to the multifarious roles that language use plays in what he calls “language games,” a term that he uses both for actual games involving language (§7) and also for all of the human activities, or “forms of life,” in which we use language (§19). To understand an expression is to understand the roles it plays in one or more broader activities. “For a large class of cases of the employment of the word ‘meaning,’” Wittgenstein says, “this word can be explained in this way: the meaning of a word is its use in the language” (2009: §43).

According to an influential reading of Wittgenstein, language use and all other rule-governed activities are essentially social.²⁵ Driven by their foundationalist epistemology, Russell and some of the logical empiricists held that all scientific truths could, at least in principle, be analyzed into essentially private statements about the contents of individual agents’

²³ See Soames (1987) for this terminology, and for an influential argument against treating propositions as sets of truth-supporting circumstances. For philosophical defenses of possible-worlds propositions, Stalnaker (1984) and Lewis (1986: §1.4).

²⁴ For surveys of the contemporary literature on propositions, see McGrath and Frank (2018) and Hanks (Chapter 19, this volume).

²⁵ On the controversy over the accuracy of this reading, see Canfield (1996).

sensations. One important part of the philosophical context for Wittgenstein's philosophy of language was his ambition to demonstrate the incoherence of this project, most famously by objecting to the idea of private languages that it presupposes.²⁶

Wittgenstein's ideas about language played a central role in his anti-theoretical, therapeutic approach to philosophy, whose aim was to diagnose and dissolve self-inflicted philosophical confusion rather than to give theoretical answers to philosophical questions (see, for example, Wittgenstein, 2009: §133). One source of philosophers' confusion, Wittgenstein argued, is their tendency to take a piece of language that has a clear role in a particular language game and attempt to extend its use beyond this natural habitat and into philosophical theorizing. "Philosophical problems arise," he says, "when language goes on holiday" (2009: §38).

One way that Wittgenstein's anti-theoretical posture manifests itself is in his imperviousness to straightforward interpretation. Rather than a theory, what we find is a discussion with many evocative but difficult-to-reconcile thoughts. This is particularly true of Wittgenstein's multifaceted idea of meaning as use. Somewhat ironically, many of the conflicting facets have grown into competing positions in contemporary pragmatics, as philosophers have attempted to build Wittgenstein's insights into theories. Some of these attempts to tame Wittgenstein's ideas have been self-conscious and acknowledged. For example, Brandom introduces his own influential project in the philosophy of language by saying that "one of the overarching methodological commitments that orients [his] project is to explain the *meanings* of linguistic expressions in terms of their *use* – an endorsement of one dimension of Wittgenstein's pragmatism." But in order to "work out the details of a theory of meaning or, for that matter, of use," Brandom says, we must reject Wittgenstein's "theoretical quietism" (1994: xii). Most of the other major figures in contemporary pragmatics have been less explicit about their debts to Wittgenstein, though all can be viewed as attempting to theorize the relationship between meaning and use that he was the first to posit.

A case in point is Grice's intentionalist project, which aims to reduce facts about the semantic properties of linguistic expressions to facts about what speakers mean by them, and in turn to facts about speakers' psychology. On Grice's view, to mean something is to behave in a way that is intended to change an addressee's mind, in part by revealing to them the intention to do so (Grice, 1957, 1969). Speech acts of different kinds are intended to change the addressee's mind in different ways.²⁷ For an expression to have a meaning for a group of speakers is for the group members to have a shared "procedure in their repertoires" to mean

²⁶ Wittgenstein's private-language argument and related remarks about rule following have spawned an enormous literature, much of it centering on an interpretation due to Kripke (1982). For overviews of this literature, see Candlish and Wrisley (2014) and Miller (Chapter 23, this volume).

²⁷ See also Bach and Harnish (1979), Schiffer (1972), Strawson (1964).

something by uttering it (Grice, 1968).²⁸ Grice's view remains highly influential, both in philosophy and in the cognitive sciences, where it has animated theories of the cognitive mechanisms underlying communication, their evolution, and their development in children.²⁹ At least part of Grice's philosophical motivation was to counter the widespread mid-century skepticism of semantic and psychological notions typified by Quine (1951, 1960) and Ryle (1949). Later intentionalists would articulate this motivation more explicitly, arguing that their project was part of a larger attempt to find a place for meaning in the natural world (Schiffer, 1982; Loar, 1981).

Grice (1975) is also responsible for developing the most influential theory of how we communicate in nonliteral and indirect ways. We do this, he argued, by exploiting our interlocutors' tacit cooperativity. By making an utterance that would be uncooperative if literal, we prompt our interlocutors to avoid that conclusion by seeking an additional or alternative hypothesis about what we intended. Grice dubbed this sort of nonliteral or indirect act of meaning something an "implicature." Grice's original philosophical applications of his theory of implicature were to defend a causal theory of perception (1961) and to dissolve the apparent methodological tension that divided approaches to philosophy centered around formal logic and ordinary language, respectively (1989: chs. 1–2). But his theory has now become enormously influential mainly as a contribution to natural-language pragmatics.

A second theory of language use to emerge from postwar Oxford was J. L. Austin's (1962, 1963, 1970) theory of speech acts. Austin took speech acts to be conventional procedures – acts that are performed by conforming to social conventions. As one of his paradigm examples, Austin considered the case of performing a marriage ceremony – a ritual whose nature and conditions of successful performance are bound up with an elaborate social institution.

Austin held that a speech act can be analyzed at several levels of abstraction. A single utterance may constitute a *locutionary* act of saying that the addressee's dog is poorly trained, an *illocutionary* act of insulting them, and a *perlocutionary* act of offending them. Locutionary acts are individuated in terms of their sense and reference (Austin, 1962: 93), illocutionary acts are individuated by their force, which Austin takes to be governed by social conventions, and perlocutionary acts are individuated in terms of their extra-conventional effects.³⁰ Searle (1968, 1969) reworked Austin's

²⁸ Later intentionalists, inspired by Lewis (1969, 1975), replaced Grice's talk of procedures with theories of linguistic convention (Loar, 1976, 1981; Schiffer, 1972, 1982).

²⁹ Scott-Phillips (2014), Sperber and Wilson (1995, 2002), Tomasello (2003, 2008).

³⁰ Austin spent the first half of *How to Do Things with Words* on a precursor view according to which utterances can be divided into constatives (statements and their ilk) and performatives (see also Austin, 1946). Austin spends so much time on this view only to abandon it because one of the negative goals of Austin's lectures is to undermine the distinction between cognitive and noncognitive sentences, which had been central to the philosophical methodology of the logical empiricists.

locution–illocution distinction as the distinction between the propositional content and force of an illocutionary act, and this distinction is still widely assumed, even among non-conventionalists.³¹

A third approach to the connection between meaning and use to arise in Wittgenstein’s wake was due to Wilfrid Sellars (1954, 1969), one of whose defining legacies is his formulation of a functionalist theory of both linguistic and mental content. On Sellars’ view – *contra* (e.g.) Grice – mental states have no explanatory priority over speech acts, and the contentfulness of both is to be explained in terms of the overall functional roles they play in an agent’s perceptions, inferences, and actions. In developing this view, Sellars emphasized the sociality of language, saying that, “As Wittgenstein has stressed, it is the linguistic community as a self-perpetuating whole which is the minimum unit in terms of which conceptual activity can be understood” (1969). For Sellars, this makes language use, and intentionality in general, a form of “norm-conforming behavior” – an activity whose moves are governed by social rules (Sellars, 1954: 204).

Two contemporary approaches to speech acts build on Sellars’ ideas.³² The most influential defender of the first is Robert Brandom (1994, 2000, 2008), who has sought to understand the nature of human thought, language use, and rationality in terms of the roles that these activities play within a language game of giving and asking for reasons (1994: ch.3). Speech acts, on Brandom’s view, are public moves within this language game, and Brandom thinks of them, fundamentally, as undertakings of social commitments.³³ A second strand was first developed by Ruth Garrett Millikan, who understands speech acts in terms of their proper function of producing certain effects in addressees. A speech act’s proper function may be the result of a natural-selection-like process of differential reproduction and needn’t involve intentions on the part of the speaker (Millikan, 1984, 1998). This idea has been further developed by signaling theorists using the tools of evolutionary game theory (Skyrms, 2010; Zollman, 2011).

Another influential mid-century attempt to connect linguistic meaning to language use is due to Michael Dummett, who argued that the aim of a theory of meaning is to say both “*what* the speaker knows, but also how his knowledge is manifested” (1975: 128). Since our knowledge of language is implicit, Dummett argued, to understand how a speaker manifests their knowledge of language is to understand how they use language. Synthesizing ideas that he finds in Frege and Wittgenstein, Dummett argued that the publicity of the sense of a word required it to be “uniquely

³¹ The basic idea behind this distinction goes back to Frege (1879), whose formal language distinguishes thoughts from the act of judging or asserting them, and whose *Foundations of Arithmetic* (1884) was translated into English by Austin.

³² On the relationship between the two, see Millikan (2005a).

³³ For some related views, see Geurts (2019a), Kukla (2014), Kukla and Lance (2009), MacFarlane (2011); Peregrin (2014); Tirrell (2012).

determined by the observable features of [the word's] linguistic employment . . . ; it follows that a grasp of its sense is fully manifested by the manner in which the speaker employs it" (1975: 135). One consequence of this view, according to Dummett, is that the notion of warranted assertibility must play some of the roles that truth is normally taken to play in semantic theorizing. Although most of the details of Dummett's philosophy of language have lost currency, the idea of warranted assertion – and, in particular, the broader question of what epistemic norm(s) governs assertion – has turned out to be enormously influential.³⁴ Most notably, Dummett's student Timothy Williamson (2000, 1996) has argued that what makes a speech act an assertion is that it is governed by the norm that one must assert only what one knows. This claim has given rise to a substantial literature in which theorists assume that assertion can be characterized by an epistemic norm and proceed to debate the nature of this norm.³⁵

Two other traditions of theorizing about language use take inspiration from Wittgenstein's (1960: 67–74) claim that first-person attitude ascriptions should be understood as direct expressions of the states that they purport to report. For example, Wittgenstein says that uttering *I am in pain* is better understood by analogy to moaning in pain than to a description of someone's mental state. This view can now be seen as an early instance of expressivism, which is a loose collection of theories united by the idea that some or all apparently factual claims are actually something else in disguise.³⁶ For example, metaethical expressivism is the view that what appear to be ethical assertions are actually better understood in some other way – as expressions of emotion (Ayer, 1936; Stevenson, 1937), as prescriptions (Hare, 1952), or as expressions of motivational states (Blackburn, 1998; Charlow, 2015; Gibbard, 2003).³⁷ Expressivisms have also been developed to make sense of epistemic vocabulary (Yalcin, 2007, 2011, 2012), ontological claims (Carnap, 1950; Flocke, 2018), and, following Wittgenstein, first-person ascriptions (Austin, 1946; Lawlor, 2013; Wisdom, 1952). In the extreme, global expressivists extend non-factualism to all purportedly factual statements (Price, 2013).

A second tradition that draws on Wittgenstein's ideas about self-expression takes the expression of mental states to be the fundamental concept in a theory of speech acts. To perform a speech act, on this view, is just to express a state of mind, and speech acts of different kinds express different kinds of mental states.³⁸ Expression theorists disagree about how

³⁴ The notion of warranted assertibility is originally due to Dewey (1941).

³⁵ For a summary, see Pagin (2016: §6.2).

³⁶ Another source of expressivism is the view, which can be traced back to Wittgenstein (1922: §§4.003, 6.53–54) via the logical empiricists (e.g. Ayer, 1936; Carnap, 1931, 1950), that many normative and philosophical claims lack cognitive content.

³⁷ For a history of metaethical expressivism, see Schroeder (2010).

³⁸ See, e.g., Bar-On (2004, 2013), Davis (2003), Devitt (2006), Green (2007), Pagin (2011), Rosenthal (1986), Turri (2011). On the connection to Wittgenstein, see Green (2001: 18) and especially Bar-On (2004).

to understand the expression relation, but they agree that it needn't involve Gricean intentions, conventions, or normatively loaded social commitments, and this makes their view a distinctive take on the nature of speech acts.

A final contemporary take on language use thinks of communication as unfolding around shared contexts, understood as bodies of representations that change in response to speech acts and that affect how future speech acts are affected in turn. One source of these ideas is Stalnaker (1978), who models contexts as sets of propositions that participants treat as common ground. Stalnaker's most direct influences were Grice and Schiffer, who had understood speech acts in terms of their intended effects on addressees' minds. Stalnaker, by contrast, understands speech acts in terms of their effects on the information shared by all participants in a conversation. A second source of context-centric models is Lewis, who, with a nod to Wittgenstein, titled his seminal paper on the topic "Scorekeeping in a Language Game" (1979b). Lewis generalizes Stalnaker's model by adding various components to contexts, which he dubs "conversational scoreboards." Just as a baseball game's current score tracks past events in the game and constrains what can happen next, Lewis argues that a conversation's score (i.e. its context) tracks information put there by previous speech acts and determines which ones can be felicitously performed going forward. Models along these lines animate much contemporary work in pragmatics and can be interpreted in ways that are compatible with many of the foregoing foundational views about language use.³⁹

3.6 Loose Ends

With more space, other tales of the recent history of the philosophy of language could be told. Of the topics that I regret having to leave out, some are covered elsewhere in this book: truth (see Raatikainen, Chapter 11), generative linguistics (Ludlow, Chapter 26; Jakielaszek, Chapter 21), internalism and externalism (Collins, Chapter 8), the relationship between semantics and metaphysics (Fox, Chapter 5), and the semantics–pragmatics interface (Del Pinal, Chapter 10; Kijania-Placek, Chapter 14) are all areas with interesting historical relationships to analytic philosophy.

With the space available, I hope to have given a variety of instances of what I take to be a general historical pattern. Analytic philosophy, though preoccupied with its own philosophical concerns, was also an enormously productive workshop for the creation of tools for studying natural language.

³⁹ For an overview, see Harris, Fogal, and Moss (2018: §1.2).