John Hawthorne
and David Manley

The Reference Book

OXFORD

The Reference Book

This page intentionally left blank

The Reference Book

John Hawthorne and David Manley

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Contents

Acknowledgments

This work is in large part a synthesis and defense of a cluster of ideas, many of which have already been suggested individually by contemporary philosophers or linguists. We are therefore indebted to many thinkers, as we will try to make clear throughout. But it is worth emphasizing those theorists upon whose work we have most heavily drawn. In the first half, Robin Jeshion, David Kaplan, Steven Schiffer, Ernest Sosa, Jonathan Sutton, and Howard Wettstein; and in the second half, Paul Elbourne, Kai von Fintel, Irene Heim, Jeffrey C. King, Stephen Neale, Roger Schwarzschild, Jason Stanley, and Zoltán Szabó.

Many thanks are also due to all of those who have discussed the ideas in this book with us, or provided us with comments on previous drafts of the manuscript. We will inevitably fail to mention everyone who ought to be thanked, but we are very grateful for help from Barbara Abbott, Emma Borg, Herman Cappelen, Nate Charlow, Billy Dunaway, Rachel Goodman, Aidan Gray, David Kaplan, Gail Leckie, Ofra Magidor, Sarah Moss, Craige Roberts, Daniel Rothschild, Barry Schein, Anders Schoubye, Scott Soames, Jason Stanley, Eric Swanson, and Juhani Yli-Vakkuri. In addition, several people went so far as to provide extensive comments on the entire book (at some stage or other of its development), and for this we owe special thanks to Cian Dorr, Paul Elbourne, Jeff King, Ted Sider, Timothy Williamson, and Elia Zardini. Thanks also to audiences at: a symposium at the Eastern APA, a symposium at the Pacific APA, a colloquium at the University of Notre Dame, a workshop at Arché in St Andrews, a Logos conference in Barcelona, a seminar at Princeton, and a seminar at USC. Finally, we would like to thank the Center for the Study of Mind in Nature for its support of this project.

This page intentionally left blank

PART I

Against acquaintance

This page intentionally left blank

1

Introduction: reference and singular thought

Our first aim in this book is to critically examine some widespread views about the semantic phenomenon of reference and the cognitive phenomenon of object-directed thought. In this chapter we provide an initial case in favor of *liberalism*—the thesis that neither of these phenomena is tied to a special relation of causal or epistemic acquaintance. And in the two chapters that follow, we consider and reject various arguments against liberalism.

In the second part of the book (Chapters 4–6) we turn to an investigation of various noun phrases in natural language. In particular, we challenge the alleged semantic rift between definite and indefinite descriptions on the one hand, and names and demonstratives on the other—a division that is often motivated by appeals to acquaintance. Drawing on recent work in semantics, we explore a more unified account of all four types of expression. We conclude by drawing out some implications of this account for the traditional categories of reference and singular thought.

1.1 Preliminaries

The discovery of the twin categories of *reference* and *singular thought* is widely felt to be one of the landmark achievements of twentieth-century analytic philosophy. On the one hand, there is the distinction between a genuinely referential expression of natural language and one that is *about* an object only in some looser sense. On the other hand, there is a corresponding distinction between a thought that is loosely about an object, and one whose bond with an object is robust enough for it to count as genuinely 'singular' or 'de re'. We hope to shed light on these two ideas in this chapter.

Two procedural points should be stressed.

First, we are, like most of our readers, philosophers who have become acclimatized to habits and trends in philosophy that make use of the terms 'reference' and 'singular thought'. Our project is not to survey how these and related terms are used in the tradition, as though we were anthropologists studying the practices of a community from the outside. Neither would it serve our purposes to insist on our own quasi-stipulative definitions. Our project is one of critical examination from within. The working assumption of much contemporary philosophy is that the relevant terms carve

meaning and psychological reality at the joints—that is, they are intended to express a semantic natural kind and a cognitive one, respectively. As such, even a critical inquiry would do best to approach them, not by nominal definition, but by carefully investigating candidate kinds and clusters of kinds, to identify what (if anything) best answers to talk of reference and singular thought.

Second, our inquiry will take place primarily within the philosophy of language. But why focus on language when one of our topics is singular *thought*? To begin with, the category of referential terms affords us with a preliminary grip on a certain kind of content: *singular* contents are those that are expressed by sentences containing referential terms. And the special category of singular *thought* is usually taken to involve bearing a cognitive attitude towards a singular *content*. Moreover, the contents of our thoughts can be studied to some extent by attending to the language with which we ascribe those thoughts to each other. If there are object-directed thoughts of an importantly distinctive type, our reporting practices would seem to offer *prima facie* good guides to their presence. From this perspective, referential terms give us a window into a certain kind of content, and attitude reports that use them reveal a distinctive kind of mental state. By the end of our inquiry we will have qualified the preceding ideas; but for now they will serve as our springboard.

1.2 Themes from Russell

By way of introducing the relevant tradition, let us begin with two themes from Bertrand Russell.

(i) *Logically proper names*. We all have a rough and ready notion of aboutness according to which the phrases 'a certain president', 'the president of the U.S.', and 'Barack Obama' can all be used to make claims about a particular individual.[1] And aboutness seems to connect straightforwardly with truth, at least for simple subject-predicate sentences.[2] In the simplest cases, the result of concatenating an expression that is about an object o with a monadic predicate will yield truth iff the predicate is true of o. All three expressions just mentioned pass this test; no real gap is yet opened up between a special category of reference and the ordinary notion of being about an object.[3]

But Russell held that some such expressions can figure in true and false sentences even if there happens to be no object for them to be about.[4] For the moment, let us call

---

[1] Here we suppress Russell's denial (in 'The Philosophy of Logical Atomism') that there is, in reality, a particular corresponding to 'Socrates': 'the names that we commonly use, like "Socrates", are really abbreviations for descriptions; not only that, but what they describe are not particulars but complicated systems of classes or series' (Russell, 'The Philosophy of Logical Atomism', pp. 200–1).

[2] See Evans, *The Varieties of Reference*, p. 49.

[3] At least reading the 'iff' as a material biconditional.

[4] E.g. Russell, 'On Denoting', pp. 479–84.

an expression *meaningful* if it can combine with other expressions to generate variously true and false claims. The object that 'the president of the U.S.' happens to be about does not enter into its meaning in any constitutive way; indeed according to Russell it can be meaningful without being about any object at all, as in a case where the U.S. has no president.[5] The crucial mistake of the Meinongian, from a Russellian perspective, is a failure to grasp the complex relationship between *meaningfulness* and *being about an object*.

At the same time, Russell held that there is a category of expressions—that of logically proper names—whose meanings do crucially involve the objects they are about.[6] Such an expression's contribution to determining claims as true and false must be fixed by the object it is about. In fact, Russell held that the meaning of a logically proper name *just is* the particular that it is about. He also argued that the way in which a name means a particular object is different from the way in which a predicate means a quality.[7] This encourages us to think that there is a special binary relation—call it *reference*—that relates logically proper names to particulars. This relation differs from the ordinary notion of 'being about' and also from the relation at issue when we say '"red" expresses redness'. Thus, for example, while 'The famous ship that sank on April 15, 1912' is an expression about the *Titanic* and so 'refers' to it in some loose sense, it is not a logically proper name, and so the fundamental relation of *referring* does not hold between its occurrences and that ship.[8] In short, as Gareth Evans put it, Russell 'challenged the unity of the intuitive category of referring expressions'.[9]

(ii) *Knowledge and discrimination*. A further Russellian thesis about logically proper names is that thinkers who employ them must stand in an intimate epistemic relationship to the objects these names are about.[10] Russell used the term 'acquaintance' for the relationship involved here, but this is a term of art. The ordinary use of 'acquaintance' connotes no more than a shadow of the relation Russell had in mind, which must involve unmediated presentation. As a result, our Russellian acquaintances are few: they include universals and sense data—and possibly ourselves—but certainly not ordinary physical objects.[11] This thesis about logically proper names is an instance of a more general constraint: nothing can be a basic building block of thought unless one

---

[5] See 'Knowledge by Acquaintance and Knowledge by Description', p. 122. See also 'On Denoting', where he writes that 'The King of France is bald' is 'plainly false' (p. 484). We will return to this contentious claim in Chapter 5.

[6] See 'The Philosophy of Logical Atomism', p. 201; *My Philosophical Development*, p. 168.

[7] Ibid., p. 46.

[8] Rather than think of reference as a binary relation between occurrences of expressions and objects, one might instead think of it as a ternary relation between an abstract expression, a context, and an object. The latter would be truer to David Kaplan's style of presentation.

[9] Evans, *The Varieties of Reference*, p. 42.

[10] For example, he writes: 'A name, in the narrow sense of a word whose meaning is a particular, can only be applied to a particular with which the speaker is acquainted'. Russell, 'The Philosophy of Logical Atomism', p. 201. See also 'On the Nature of Acquaintance', pp. 130–1, 167–8.

[11] 'Knowledge by Acquaintance and Knowledge by Description', pp. 109–12.

is acquainted with it. A language reflecting the structure of thought would be one whose simple terms in a given lexical category correspond one-to-one with objects of acquaintance in a corresponding ontological category.[12] And every claim would be built up out of a basic lexicon of such simple terms.

One way to understand Russell's motivation for this strict constraint is to notice the apparent connection between linguistic understanding and knowledge. Russell held that 'Socrates' is not a genuine and logically proper name, because understanding the word 'Socrates' involves grasping some kind of description about the man; that is, knowing something *about* him. But a logically proper name cannot be understood by any kind of definition: it is understood by *knowing the particular itself*.[13] This in turn requires direct awareness of the particular: all other knowledge concerns the 'inferred world' rather than the world of 'data', which makes it descriptive and indirect.[14] The result, he hoped, is that there is no way to be *wrong* about which particular is at issue in thoughts involving proper names, because if one understands the name, one cannot be in error about what it is a name of.[15] But if understanding a logically proper name requires us to know which particular 'before our mind' we are naming, then it would seem that understanding two distinct proper names rules out the possibility that they refer to the same object (or different objects) without our knowing that they do.[16] As a result, if two genuine names 'a' and 'b' refer to the same particular, 'a is b' has no epistemic significance.[17] (And as an added benefit, Russell's requirement on under-standing rules out the epistemic possibility that someone understands a logically proper name and yet it fails to refer.)[18]

[12] See 'The Philosophy of Logical Atomism', pp. 194–9; 'Knowledge by Acquaintance and Knowledge by Description', p. 117–8, and also *My Philosophical Development*, p. 169, where this principle is qualified with the remark: 'It is perhaps necessary to place some limitation upon this principle as regards logical words—e.g. *or, not, some, all*.'

[13] 'A proper name, if it is to fulfil its function completely, should not need to be defined in terms of other words: it should denote something of which we are immediately aware' (*My Philosophical Development*, p. 166).

[14] See for example ibid., pp. 22–7.

[15] As Russell writes: 'At any given moment, there are certain things of which a man is 'aware' . . . If I describe these objects, I may of course describe them wrongly; hence I cannot with certainty communicate to another what are the things of which I am aware. But if I speak to myself, and denote them by what may be called 'proper names', rather than by descriptive words, I cannot be in error' ('On the Nature of Acquaintance', p. 130).

[16] However, notice that there are important asymmetries of epistemic possibility here (where *p* is epistemically possible iff it is compatible with what one knows). Suppose we grant that, given strict acquaintance, one cannot understand two logically proper names that co-refer if it is epistemically possible that they fail to co-refer. Even so, this does not rule out wrongly *taking oneself* to understand two logically proper names, and taking them to be co-referential when in fact they are not. Without some supplementary principles, strict acquaintance does not guarantee the impossibility of unnoticed coreference—or unnoticed lack of coreference.

[17] Russell, 'The Philosophy of Logical Atomism', pp. 244–6.

[18] Though the sort of epistemic accessibility issue mentioned in fn. 16 applies here as well. The acquaintance requirement does not by itself rule out that one wrongly takes a given term to be a logically proper name that one understands.

However, the requirement that co-reference must always be discriminable in this way puts extraordinarily demanding constraints on what sorts of objects can count as the semantic values of logically proper names. One might have two encounters with an ordinary physical object and on each occasion perform a baptism, without knowing whether the resulting terms co-refer. In general, it seems that logically proper names cannot refer to objects for which one can have epistemically diverse perspectives: if, say, an object is the sort of thing that one can look at from the front as well as from the back, that object will be rendered unsuitable as the referent of a logically proper name.[19]

Given Russell's assumptions about understanding, then, his acquaintance constraint seems inevitable. But—as we will see in Chapter 3—some sort of epistemic acquaintance requirement can be motivated with less contentious premises about understanding. For example, it would seem that to understand an expression requires knowing what it means. So if the meaning of a proper name is its referent, understanding a proper name would seem to require at least being in a position to know *that it refers*. Now suppose, for reductio, that one has a proper name whose meaning is an epistemically distant object like the black hole Cygnus X-1. Knowing what 'Cygnus X-1' means will require knowing that it refers to Cygnus X-1, and hence that Cygnus X-1 exists. But insofar as one does not *know* whether Cygnus X-1 exists, it appears to follow that one does not really understand the name.[20] (We will return to such arguments in Chapter 3.)[21]

Since most philosophers have abandoned the radical Russellian form of epistemic acquaintance as a constraint on understanding, our concern in the following three chapters will be to address some of its far less stringent successors in vogue today. However, a small current of contemporary philosophy shares some of Russell's original motivations and—as a result—some have attempted to craft a notion of semantic value that satisfies his extreme acquaintance desiderata.[22] They posit a kind of meaning ('primary meaning') such that it is always available to *a priori* investigation whether two things share a meaning of that type. In short, equivalence of primary meaning must always be open to view. It is therefore no surprise that facts about primary meaning are supposed to be fixed by facts wholly internal to the mind.[23]

---

[19] Thanks to David Kaplan for this way of expressing it.

[20] Russell would not himself put things in exactly this way, as he denies that where 'N' is a logically proper name, 'N exists' can be a meaningful sentence. (See 'The Philosophy of Logical Atomism', p. 252–3.)

[21] Note that the epistemic requirement that one needs to know that x exists in order to understand an expression that refers to x does not obviously deliver the epistemic requirement that one must be in a position to know that two referential expressions co-refer when they do. Even if there is some x such that someone knows that 'Hesperus' refers to x and that 'Phosphorus' refers to x, it may not follow that she is in a position to know that 'Hesperus' and 'Phosphorus' both refer to x.

[22] See, for example, Chalmers, 'The Foundations of Two-Dimensional Semantics'; Jackson, *From Metaphysics to Ethics: A Defence of Conceptual Analysis*.

[23] However, even this would not suffice for the kind of transparency required; for example, the 'subpersonal-level whirrings and grindings' of the devices that underlie our recognitional capacities, while

Along with the majority, we hold out little hope for such a project. Not only are primary meanings often disconnected from ordinary ways of thinking and talking about meaning, but there also seems to be a problem *in principle* of crafting primary meanings so that the epistemic desiderata that moved Russell are met. For one thing, we adhere to the widely held view that the indiscriminability of phenomenal episodes is intransitive.[24] (Phenomenal episodes are the contemporary successors to 'sense-data'). This in turn suggests that two phenomenal episodes of different types may nevertheless be indiscriminable. Given this, even predicates whose meanings are phenomenal state types need not satisfy strict Russellian acquaintance conditions. Moreover, consider illusions of consciousness—beliefs that a phenomenal episode is present where it is not. If these are possible, one might fail to refer with a putatively logically proper name and not know it; and one might also refer without knowing that one has done so (there being a real danger of illusion). We postpone a proper engagement with reactionary Russellianism to another time.[25]

Russellian acquaintance, then, has been largely jettisoned, and the category of 'logically proper names' has been vastly expanded. But acquaintance-theoretic ideology has lived on in various degenerate forms. Some retain the idea that there is something deep and important about Russell's requirement that to understand an expression whose meaning is an object, one needs to know *which* object it is. Others adhere to the idea that there must be a special causal connection between an object and a mental representation if the former is to serve as a *bona fide* referent of the latter. Our own view is that attempts to constrain reference or object-directed thought using some version of acquaintance have invariably been misguided. We aim to vindicate this pessimism in the next few chapters.

1.3 Reference after Russell

Let us now set out some semantic ideas connected with Russell's logically proper names that have been widely brought to bear on their successors.

(i) *Object-dependence*. First, Russell held that logically proper names are 'meaningless unless there is an object which they designate.'[26] That is, intuitively, a referring expression's very meaningfulness—again, in the sense of its capacity to successfully contribute to the generation of truth conditions—is somehow bound up with its successfully being about a particular object. However, we cannot simply identify an expression of the

---

internal to the mind, will not fit the bill. See Martin Davies, 'Reference, Contingency, and the Two-Dimensional Framework'.

[24] For exploration of the relevant issues see Williamson, *Identity and Discrimination*.

[25] Though for some critical discussion, see Hawthorne, 'Direct Reference and Dancing Qualia'; Schroeter, 'The Rationalist Foundations of Chalmers' Two-Dimensional Semantics'; and Yablo, 'Coulda, Woulda, Shoulda'.

[26] *My Philosophical Development*, p. 168.

relevant kind as one that is meaningless unless there is an object that it is about. For we may wish to allow—on some ways of individuating expressions—that one occurrence of an expression can be crucially about a particular object, while another occurrence of that expression is meaningful even though it is not about any particular object. (This is a common view about the pronoun 'he': to be meaningful when it is used 'deictically', there must be an object that it is about. But it is not being used in this way when someone says, 'Every boy wishes he did not have to go to school'.) To cordon off this issue, we can put the idea this way. An expression is *object-dependent for meaningfulness* on an occasion of use iff it is meaningful on that occasion by virtue of there being an object that it is about on that occasion.

(Can we state this idea without using a notion like *by virtue of*?[27] It will not quite do to state the idea counterfactually, as in: 'An expression is object-dependent on an occasion just in case: had there been no object that it was about, it would have lacked meaning'. This works for many cases. But it does not get to the heart of the matter. Suppose an expression is actually used in a fashion that is paradigmatically referential, but the closest possibilities where it is not about an object are possibilities where it is meaningfully put to a different use—one can easily construct such cases using 'he'. Then, by the counterfactual test, it would not be object-dependent for meaningfulness.[28] The situation is a familiar one: counterfactual tests are fairly well understood but can rarely play the role of real definitions—at best they are reasonable facsimiles.)

A related idea is that of an expression's having an *object-dependent meaning* (on some occasion of use). Let us say that an expression's meaning is object-dependent on o just in case no expression could have had that very meaning unless it were about o. Given the plausible assumption that no expression can be about an object unless the object exists,[29] this implies that when an expression's meaning is object-dependent on o, then no expression could have had that very meaning unless o existed.

---

[27] Some avenues are opened up if we follow Kit Fine in distinguishing facts that are merely semantic *as to topic*—they pertain to semantic features of things—and facts that are also semantic *as to status*—they are true 'wholly consequential upon the meaning of the expressions' which they concern (Fine, *Semantic Relationism*, pp. 43–5). We can then introduce a sentential operator 'It is a semantic-as-to-status fact that' (p. 135, fn. 7). Equipped with such an operator, we could (as Fine is aware) cash out the idea of object-dependence in terms of it: An expression e is object-dependent for its meaningfulness on an occasion iff it is a semantic-as-to-status fact that a sentence containing e is true on that occasion only if there is some x such that the expression e is about x. And a term has an object-dependent meaning on an occasion iff there is some object x such that it is a semantic-as-to-status fact that e is about x.

[28] One might try to improve things by blending 'occasion of use' ideology into the counterfactual in a suitable way. This is hard to do, however. For example, suppose, as is fairly standard, occasions of use are treated as world-bound. Then counterfactuals beginning 'Had there been no object it was about on that occasion of use, then...', directed at a term that actually succeeds in referring, will have impossible antecedents.

[29] The consequence flows from the natural assumption that, in general, for a binary relation to be instantiated, its *relata* must exist.

Roughly, then, an expression is object-dependent for meaningfulness if *having any meaning at all* requires it to be about an object. And an expression has an object-dependent meaning if *having that meaning* requires being about the object that expression is actually about. There are expressions whose occurrences happen to be about objects and yet they are not object-dependent in either way. For example, 'The first president of the U.S.' is about George Washington; but it is widely thought not to be object-dependent in either sense. In contrast, Russell's logically proper names are object-dependent in both ways. Their meaningfulness, as well as the particular meaning they have, consist in their referring to particular private objects.

As noted earlier, the purported class of referential vehicles has expanded well beyond Russellian restrictions; ordinary proper names, demonstratives, and indexicals are now typically held to be referential vehicles. And it is still common to connect the idea of the reference relation with both kinds of object-dependence. But this connection is not inevitable. Some theorists, such as Gareth Evans, have defended views of reference on which referring expressions can fail to be object-dependent in one or even both ways. We will have occasion to examine such views in later chapters.

(ii) *Exhaustiveness and rigidity*. Another idea that has come to be associated with referential terms, especially since the work of David Kaplan and Saul Kripke, concerns their behavior in modal contexts. We can begin with the picture according to which a sentence can be evaluated for truth relative to ways the world might have been (or 'possible worlds'). Among other things, this procedure determines the truth-value of sentences that result from embedding the original sentence within a modal operator. For example, 'possibly *S*' is standardly treated as true just in case *S* is true evaluated at some (accessible) possible world.

Within this framework, the notion of sentences being true relative to worlds will give rise to the notion of expressions being about objects relative to worlds. Recall the contrast between 'The first president of the U.S.' and 'George Washington'. The former picks out (it is, loosely speaking, about) different presidents at different worlds, while the latter picks out the same individual at every world, regardless of whether he is president—or even called 'George Washington'—at that world. As a result, the truth of 'George Washington played chess', evaluated at a given world, will always turn on how things stand with respect to that individual at that world. But the truth of 'The first president of the U.S. played chess' evaluated at a world will turn on how things stand with whomever is the first president at that world.[30] And what this shows is that

---

[30] Not everyone will prefer a metalanguage involving worlds and truth at worlds as the preferred semantic framework for describing modal language. Those who treat modal operators as quantifiers that can bind variables may prefer to treat the semantic contribution of an atomic sentence when unembedded as true simpliciter or false simpliciter, and the semantic contribution of those words within a modal operator as open-sentence-like. (For more on this see Chapters 5 and 6.) Those who prefer a metalanguage of modal operators will not use the apparatus of worlds to illuminate the semantic workings of a modal object language. But all of these approaches will have something like the non-rigid/rigid distinction—after all, they will likely all agree

the contribution of 'The first president of the U.S.' to determining the truth or falsity of claims is not fixed by the object it is about. And so, in vindication of Russell, it will not do to think of the meaning of that expression as the man, George Washington.

A *rigid* use of an expression is dedicated to being about a certain object—it is not about any other object relative to any world. The *obstinate* kind of rigidity requires a term to be about that object at every world, regardless of whether that object exists at that world.[31] For example, consider the claim 'George Washington fails to exist'. This claim is standardly taken to be about Washington relative to every world of evaluation, so that relative to a world where Washington does not exist, the sentence is true. (This is why 'Possibly, George Washington fails to exist' is true.) But there are weaker forms of rigidity. An expression (on an occasion of use) will be rigid but not obstinately rigid iff there is some object o such that it is about o relative to some worlds, and relative to every world, it is either about o or about nothing at all.

There are at least two reasons why an expression might—in principle—be rigid without being obstinately rigid. First, it might pick out o at every world where o exists, but fail to pick out anything at worlds where o does not exist. Call this *moderate rigidity*. For example, take the expression 'the man identical to Socrates.' Now consider a world where Socrates fails to exist. On the standard Russellian treatment of definite descriptions, the proposition expressed by 'The man identical to Socrates fails to exist' is false evaluated at that world—since it claims existence with the determiner 'the' and denies it with the predicate.[32] It is thus natural for the Russellian about definite descriptions to conclude that 'the man identical to Socrates' is only moderately rigid. A second way a term might be rigid but not obstinately rigid is by failing to pick out o even at certain worlds where o exists. For example, one way of treating an expression like 'Professor Hawthorne' is to think of it as picking out John only at worlds where he is a professor, and nothing otherwise. Call this *weak rigidity*.

A further distinction is worth noting.[33] Insofar as an expression rigidly picks out o, it may be that it does so by virtue of the kind of expression it is; or instead thanks partly to extra-semantic facts. Thus 'the successor of 2' obstinately rigidly picks out the number 3 thanks to the fact that numbers exist necessarily, but not simply because it is a definite description that picks out the number 3. By contrast, the story goes, it is a semantic rule for '3'—at least when it is used as a name—that it is obstinately rigid. In the first case, the rigidity has a metaphysical explanation (this is '*de facto*' rigidity); in the second case it has a semantic one (this is '*de jure*' rigidity).

---

that '∃x (x is not the president of the U.S. but could have been)' is true but '∃x (x is not Nixon but could have been)' is false—though they will articulate that distinction in rather different ways.

[31] See Kripke, *Naming and Necessity*, p. 48ff, Kaplan, 'Demonstratives', pp. 492–3, and fn. 16; Kaplan, 'Afterthoughts', p. 569ff.

[32] We do not, of course, intend for this sentence to be read as though 'fails' somehow becomes an operator and takes wide scope over the definite description.

[33] See the Forward to the 1980 edition of *Naming and Necessity*, p. 21, fn. 21.

How does referentiality relate to rigidity? We can identify two ideas that, working together, would render all referring expressions obstinately rigid.

The first idea is that, regardless of linguistic environment, the meaning of an expression exhausts what is relevant to its contribution to the truth of a sentence. Call this *compositionality*. The second idea is a Russellian one: the meaning of a referring expression is exhausted by the object that it is about.[34] Call this *exhaustivity*. (Following the usual practice, when *o* exhausts the meaning of *e*, we shall say that *o* is the semantic value of *e*.) If we put these ideas together, we get the following picture. Even when an occurrence of a given referring expression *e* finds itself within the scope of a modal operator, the fact that *e* picks out *o* will be its only contribution to what is considered at worlds of evaluation when the resulting sentence is assessed.[35]

In less picturesque terms, the idea is that a compositional semantics for *e* will proceed simply by assigning to any occurrence of *e* the object that it is about. No further information need be encoded. Thus, the fact that the occurrence of *e* is about *o* will exhaust *e*'s contribution to compositional semantics. In particular, the interaction of a modal operator with a sentence containing an occurrence of *e* will be sensitive only to the fact that it is about *o*. This in turn will generate intuitions of obstinate rigidity for referring expressions. (This is a special instance of a more general phenomenon. Let a *rigid predicator* be a predicate that expresses the same property relative to all possible worlds. Suppose one thought that the semantic life of a predicate consisted entirely in its expressing a property. Then it will turn out that all predicates are rigid predicators. Similarly for the definite article 'the': if it has a certain relation between properties as semantic value, then there will be a perfectly good sense in which it is rigid with respect to that relation.)

Of course, insofar as the semantic life of an expression varies from context to context, the idea of exhaustivity will require refinement. The phenomenon of context-dependence raises the question whether, in principle, a referring expression might be about one object when it occurs in an atomic sentence, but about another object when it occurs within the scope of a certain operator. This is consistent with a contextualized version of exhaustivity—that the semantic contribution of *a given occurrence of e* is exhausted by its reference.

(Consider an example. Suppose we introduce an operator 'plossibly', which behaves like 'possibly' except that in the scope of 'plossibly', any name refers to Plato. The truth-conditions of 'Plossibly Socrates is identical to Plato' could be generated in a way that is perfectly compositional, and it would turn out true even though the

---

[34] If we think of predicates as being about the properties they express, and as having those properties as their meanings, this will certainly not be a sufficient condition for being a referential expression. This point is reminiscent of Russell's insistence that the way a predicate means a quality is different to the way that a logically proper name means a particular.

[35] Of course, there are other packages that deliver obstinate rigidity. Even those who reject exhaustivity can hold that modal operators are blind to trans-referential aspects of meaning.

semantic contribution of one of its parts is a little out of the ordinary.[36] But on standard assumptions about the semantics of natural language, 'plossibly' is highly anomalous, since it disallows 'Socrates' from having the referent that it normally enjoys. Insofar as one holds that ordinary names always pick out whatever they refer to in unembedded uses, one must insist that natural language modal operators are not like 'plossibly': there can be no special principles governing the reference of referring expressions within their scope.[37])

(iii) *Variables*. Philosophers have sometimes characterized reference by pointing to the semantic behavior of variables and treating it as paradigmatically referential. Of course, it can hardly be said that 'x' in an occurrence of '∃x(x is happy)' simply stands in a binary reference relation to a particular object. (Which object would that be?) Instead, the point is that *relative to an assignment function*, a variable behaves much like a referential vehicle. To begin with, there is the following structural similarity. An occurrence of an atomic sentence formed by concatenating a referential term with a monadic predicate is true iff the predicative occurrence is true of the object referred to by the referential occurrence. And an occurrence of an open sentence formed by concatenating a variable with a monadic predicate is true relative to an assignment iff the predicative occurrence is true of the object assigned to the variable by that assignment.

The similarities may run deeper. Consider a sentence containing a referring term that refers to *o*, and then the result of substituting that term for a variable. Those who see a deep affinity between variables and referential vehicles may think that every claim about the semantic life of the referring term is also true of the variable, relative to an assignment of *o* to that variable. Hence, insofar the original sentence has a semantic content, the open sentence will be considered to have that very content relative to that assignment.[38]

These reflections provide another standard route to the notion of singular *contents*: they are contents of the sort expressed by open sentences relative to assignments of objects to their variables. Relatedly, suppose one holds that any belief attribution with a referential term in its complement clause attributes a belief in a singular content. The corresponding variable-theoretic idea is this: an open sentence of the form 'John believes that *x* is F' will attribute to John a belief in a singular content about Socrates relative to an assignment that assigns Socrates to '*x*'. Coordinately, if a predicate of the form 'believes that *x* is F' is true of John relative to an assignment of Socrates to '*x*', then John will be deemed to believe a singular content about Socrates. Given the interaction

---

[36] Note that this would not make 'plossibly' a monster in David Kaplan's sense. A monstrous operator is one that operates on the Kaplanian character of an expression (Kaplan, 'Demonstratives', p. 511).

[37] This further idea precludes, inter alia, Frege's view that in the context of attitude verbs, names refer to that which is their sense in ordinary contexts.

[38] Here we are thinking of 'Relative to assignment F' as functioning roughly as a sentential operator.

of existential quantifiers with variables, it follows within this framework that '$\exists x$(John believes that x is happy)' will be true only if there is some x such that John believes a singular content about x, to the effect that x is happy. (However, as we will see, some will want to deny the further step that being belief-related to a singular content is sufficient for having a genuinely singular *thought*.)

This route to the notion of singular contents has certain advantages over our earlier characterization of them as contents of the sort that are expressed by sentences containing referential terms. First, the latter idea leans very heavily on the locution 'of the sort'. After all, there may be plenty of objects in the world for which we do not have referential terms, but we may not want to disallow singular contents about those things. This problem is less serious for the variable-theoretic formulation, since for any object in the world there will be an assignment that assigns the variable in 'x is happy' to that object.

For some readers there may also be a second advantage. We have in mind those who are unsure about whether any of the putatively paradigmatic referential devices—like names—actually function as such. Such readers may feel more secure about the ternary relation holding between an object, a variable, and an assignment than about the alleged binary relation holding between an object and a referring expression. For them, variables may be preferable as a point of departure for theorizing about singular content.[39]

(iv) *The metaphysics of propositions*. Discussions of reference and singular thought are typically undertaken within a framework of realism about propositions, according to which there are abstract objects that people believe and that utterances express. This seems to open up another way to characterize singular contents. Perhaps they are propositions with a special kind of metaphysical structure; indeed, following Russell's lead, perhaps they literally contain objects as constituents.[40] Given that the semantic life of a referential term is especially bound up with a particular object, it is tempting within this framework to say that all and only expressions that refer to Socrates will contribute Socrates himself to a proposition, and the resulting proposition will contain Socrates as a constituent.

One who tries to provide a direct characterization of singular contents in metaphysical terms has a number of complications to confront. First, constituency cannot be thought of merely as parthood, because parthood is transitive. If Socrates is part of a proposition, then so his left kidney. But proponents of the relevant vision do not think

---

[39] However, it is not inevitable that a language use variables, nor even that a lack of variables would make its expressive power impoverished. (See Quine, 'Variables Explained Away', pp. 343–7.) Suppose it were to turn out that natural languages do not in any sense contain variables, and that variables occur only in artificial languages. It would then be *prima facie* embarrassing to rely on variables to anchor our theorizing about singular content.

[40] See for example: Russell, 'Knowledge by Acquaintance and Knowledge by Description', p. 116; *My Philosophical Development*, p. 152.

of the proposition that Socrates is tall as, in the relevant sense, containing Socrates' left kidney as a constituent. This can *perhaps* be handled by thinking of propositions set-theoretically and of constituency as set membership, though some will wonder if it can ultimately be metaphysically satisfying to think of propositions as sets.[41] Second, a thing can appear as a constituent in the content of a sentence even if it is not *referred* to by any element in that sentence. For example, it is natural to think of tallness as a constituent of the content of 'Someone is tall' as well as of the content of 'Tallness is a property'. But only the second sentence is usually treated as singular in content, because only the second contains an element that refers to tallness. Thus, the simple constituency test is too blunt an instrument to characterize singular content, especially if we hoped to characterize referring terms in turn as those that semantically contribute objects to singular contents.[42]

Perhaps these challenges can be met. We have no principled aversion to realism about propositions or to the possibility that they may be structured in interesting ways.[43] But we doubt that metaphysical speculations about the structure of propositions are a promising place to start when thinking about the referential structure of language and singular thought. Perhaps, with the best final theory in hand, a directly metaphysical characterization of singular contents will suffice. But we do not want to begin with any such expectation. Still less do we want to rely on any particular metaphysical picture of singular content.

1.4 Singular thought after Russell

We have encountered several interconnected themes about referential expressions. In what follows we will begin by adopting—along with philosophical orthodoxy—the working hypothesis that there is a single semantic natural kind on which most or all of these themes converge. (By the end of the book we will be in a position to reconsider this hypothesis.) Accordingly, 'referential expression' will serve as a placeholder term for expressions of the envisaged kind—and 'reference' for the relation that holds between such terms and the objects they are about.

---

[41] We say 'perhaps' because the reduction of constituency to facts about membership is not straightforward. Suppose one construed 'Socrates teaches Plato' as expressing a set that is given by some set-theoretic reduction of the ordered triple <Socrates, teaching, Plato>. Standard set-theoretic reductions of ordered triples will not make either philosopher a member of the resulting set. So if constituency is membership we get the wrong results. Suppose instead we make constituency the ancestral of membership. Then we risk making the number five a constituent of the proposition that Socrates adores Josephine, where 'Josephine' is a name for the set of natural numbers. (Thanks to Tim Williamson here.)

[42] One fix would be to claim that a proposition p is a singular proposition about o iff (i) o is a constituent of p and (ii) p is the content of some open sentence relative to an assignment of o to some first-order variable in that open sentence. But then we would be no longer characterizing singular propositions simply in terms of their distinctive metaphysical structure: in fact, condition (i) becomes redundant given usual assumptions.

[43] For an engagement with issues of this sort, see King, *The Nature and Structure of Content*.

sample content of The Reference Book

- **download online Plutarch: How to Study Poetry (De audiendis poetis) (Cambridge Greek and Latin Classics)**
- read online A Privilege to Die: Inside Hezbollah's Legions and Their Endless War Against Israel pdf, azw (kindle), epub
- download online Incendies
- *download online The Uplift War (The Uplift Saga, Book 3) online*

- http://test.markblaustein.com/library/Plutarch--How-to-Study-Poetry--De-audiendis-poetis---Cambridge-Greek-and-Latin-Classics-.pdf
- http://junkrobots.com/ebooks/A-Privilege-to-Die--Inside-Hezbollah-s-Legions-and-Their-Endless-War-Against-Israel.pdf
- http://junkrobots.com/ebooks/The-Melancholy-Science--An-Introduction-to-the-Thought-of-Theodor-W--Adorno--Radical-Thinkers-.pdf
- http://xn--d1aboelcb1f.xn--p1ai/lib/The-Humanure-Handbook--A-Guide-to-Composting-Human-Manure--3rd-Edition-.pdf