

# 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 acclimated 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.<sup>1</sup> And aboutness seems to connect straightforwardly with truth, at least for simple subject-predicate sentences.<sup>2</sup> 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.<sup>3</sup>

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.<sup>4</sup> For the moment, let us call

<sup>1</sup> 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).

<sup>2</sup> See Evans, *The Varieties of Reference*, p. 49.

<sup>3</sup> At least reading the ‘iff’ as a material biconditional.

<sup>4</sup> 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.<sup>5</sup> 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.<sup>6</sup> 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.<sup>7</sup> 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.<sup>8</sup> In short, as Gareth Evans put it, Russell ‘challenged the unity of the intuitive category of referring expressions’.<sup>9</sup>

(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.<sup>10</sup> 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.<sup>11</sup> 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

<sup>5</sup> 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.

<sup>6</sup> See ‘The Philosophy of Logical Atomism’, p. 201; *My Philosophical Development*, p. 168.

<sup>7</sup> *Ibid.*, p. 46.

<sup>8</sup> 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.

<sup>9</sup> Evans, *The Varieties of Reference*, p. 42.

<sup>10</sup> 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.

<sup>11</sup> ‘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.<sup>12</sup> 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*.<sup>13</sup> 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.<sup>14</sup> 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.<sup>15</sup> 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.<sup>16</sup> As a result, if two genuine names 'a' and 'b' refer to the same particular, 'a is b' has no epistemic significance.<sup>17</sup> (And as an added benefit, Russell's requirement on understanding rules out the epistemic possibility that someone understands a logically proper name and yet it fails to refer.)<sup>18</sup>

<sup>12</sup> 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.*'

<sup>13</sup> '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).

<sup>14</sup> See for example *ibid.*, pp. 22–7.

<sup>15</sup> 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).

<sup>16</sup> 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.

<sup>17</sup> Russell, 'The Philosophy of Logical Atomism', pp. 244–6.

<sup>18</sup> 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.<sup>19</sup>

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.<sup>20</sup> (We will return to such arguments in Chapter 3.)<sup>21</sup>

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.<sup>22</sup> They posit a kind of meaning ('primary meaning') such that it is always available to a *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.<sup>23</sup>

<sup>19</sup> Thanks to David Kaplan for this way of expressing it.

<sup>20</sup> 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.)

<sup>21</sup> 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*.

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

<sup>23</sup> 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.<sup>24</sup> (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.<sup>25</sup>

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.’<sup>26</sup> 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’.

<sup>24</sup> For exploration of the relevant issues see Williamson, *Identity and Discrimination*.

<sup>25</sup> 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’.

<sup>26</sup> *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*?<sup>27</sup> 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.<sup>28</sup> 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,<sup>29</sup> 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.

<sup>27</sup> 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*.

<sup>28</sup> 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.

<sup>29</sup> 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.<sup>30</sup> And what this shows is that

<sup>30</sup> 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.<sup>31</sup> 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.<sup>32</sup> 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.<sup>33</sup> 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 ‘ $\exists x$  (*x* is not the president of the U.S. but could have been)’ is true but ‘ $\exists x$  (*x* is not Nixon but could have been)’ is false—though they will articulate that distinction in rather different ways.

<sup>31</sup> See Kripke, *Naming and Necessity*, p. 48ff, Kaplan, ‘Demonstratives’, pp. 492–3, and fn. 16; Kaplan, ‘Afterthoughts’, p. 569ff.

<sup>32</sup> 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.

<sup>33</sup> 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.<sup>34</sup> 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.<sup>35</sup>

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 'possibly', which behaves like 'possibly' except that in the scope of 'possibly', any name refers to Plato. The truth-conditions of 'Possibly Socrates is identical to Plato' could be generated in a way that is perfectly compositional, and it would turn out true even though the

<sup>34</sup> 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.

<sup>35</sup> 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.<sup>36</sup> But on standard assumptions about the semantics of natural language, ‘possibly’ 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 ‘possibly’: there can be no special principles governing the reference of referring expressions within their scope.<sup>37</sup>)

(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 ‘ $\exists x(x \text{ is happy})$ ’ 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.<sup>38</sup>

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

<sup>36</sup> Note that this would not make ‘possibly’ 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).

<sup>37</sup> 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.

<sup>38</sup> 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.<sup>39</sup>

(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.<sup>40</sup> 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

<sup>39</sup> 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.

<sup>40</sup> 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.<sup>41</sup> 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.<sup>42</sup>

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.<sup>43</sup> 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.

<sup>41</sup> 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.)

<sup>42</sup> 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.

<sup>43</sup> For an engagement with issues of this sort, see King, *The Nature and Structure of Content*.

We intend to approach the special category of *singular thought* in a similar fashion, so let us begin by sketching some of the usual ideas with which it is associated.

(i) *Content*. One common desideratum for the relevant psychological kind is parasitic on the notion of a singular *content*—a singular thought involves (at least) bearing a cognitive attitude to a singular content. A singular content, in turn, can be understood as a content of the sort expressed by sentences containing referential terms, or by open sentences relative to assignments. This idea is connected with some standard ways of *detecting* the presence of a cognitive attitude towards a singular content. First, if someone understands and sincerely utters a sentence containing a referential term, that is *prima facie* reason to think she believes the singular content that she expresses. Second, suppose we assume that an ordinary attitude report expresses (at least) a cognitive relation between the subject and the proposition expressed by the complement clause. Then a cognitive attitude towards a singular content would be signaled by a true attitude report with a referential expression in its complement clause, and also by a true report that ‘quantifies in’ (where a variable within the report is bound by a quantifier outside it). That is, we regiment ‘Someone is believed by John to be happy’ as  $\exists x(\text{John believes that } x \text{ is happy})$ , which in turn is true iff the open sentence is true relative to some assignment.

(ii) *Tags or files*. Another way to characterize singular thought involves appealing to the cognitive mechanisms whereby we entertain contents.<sup>44</sup> If one adopts a model on which our token thoughts have linguistic structure, it is natural to hold that singular thoughts are those involving mentalese analogues of referential expressions—call them ‘tags’. On such a picture there are simple unstructured object-representations in our language of thought that refer to particular objects. Another common picture ascribes far more internal structure to the object representations: an object is represented by a ‘file’ that contains various items of information associated with that object.<sup>45</sup> On this view, a file differs from a mere conjunctive description in that a file can represent an individual even if much of the associated information is incorrect. As John Perry puts it: the file gets to be about the source of the relevant information even if it does not satisfy

<sup>44</sup> Appeals to mental files or ‘dossiers’ go back at least to Grice, ‘Vacuous Names’, see pp. 141–4. See also Bach, *Thought and Reference*, pp. 34–44; Perry, ‘A Problem About Continued Belief’, pp. 70–4; and *Reference and Reflexivity*, especially pp. 54–7, Forbes, ‘Cognitive Architecture and the Semantics of Belief’; Recanati, *Direct Reference: From Language to Thought*, chs. 5–6; and Lawlor, *New Thoughts About Old Things*, ch. 4. The metaphor of files has also been put to a distinct but related use in theories of dynamic semantics as a way of modeling how an object may be represented over an entire discourse. (See for example, Heim, ‘File Change Semantics and the Familiarity Theory of Definites’.)

<sup>45</sup> Not only does the ‘file’ view differ from the ‘tag’ view in characterizing object representations as complex; it might also be less sanguine about the linguistic structure that the tag view seems to impose on thought tokens. However, these theories are typically entertained and discussed in very schematic form. For example, it remains very unclear whether, for example, there is anything at stake in the decision whether to regard a complex file itself as the object representation, or instead some simple tag within the file. And one might also posit unstructured object representations (with the tag view as we have described it) and yet (against that view) deny that thought tokens have anything like syntactic structure.

that information.<sup>46</sup> If there turns out to be an important natural kind of representation here, then arguably that which best deserves to be called a ‘singular thought about  $x$ ’ will be a cognitive file dedicated to representing  $x$ .

There are various motivations for characterizing singular thoughts by means of tags or files. For some, the appeal to cognitive mechanisms simply takes the place of an appeal to singular contents. One might, for example, have a non-hyperintensional picture of contents that does not distinguish the content of a belief that 1 is odd from the content of a belief that the smallest prime is odd. On such a view it is natural to hold that what makes a thought singular is a matter of how the content is represented, as opposed to the nature of the content itself. Others might accept a singular/non-singular distinction at the level of content but hold that the most interesting distinctions concerning object-directed thought are marked by the presence or absence of the key cognitive mechanism and not by the bearing of propositions attitudes to singular contents.<sup>47</sup> Here the motivation might be the idea that something like tags or files will have to be appealed to for the purposes of perspicuous psychological explanation: inspired by Frege puzzles, one might hold that the best such explanations will often employ a very fine-grained way of individuating object-directed thoughts. Thus it is not sufficient for two distinct thought-tokens to count as ‘the same singular thought’ that they have the same semantic content. They must also involve the same mental tag or file.<sup>48</sup>

(iii) *Satisfactional vs. relational*. Still another idea associated with singular thoughts—one that can be combined with one or both of the foregoing ideas—involves the contrast between *satisfactional* and *relational* object-representations.<sup>49</sup> This contrast has been characterized by reference to the behavior of indexicals, the use of which can exploit a relation between an indexical token and the referent without the speaker having to represent that the relevant relation holds. Kent Bach puts it this way: ‘*De re* modes of presentation . . . determine the contextual relation that something must bear to a thought to be the object of that very thought’.<sup>50</sup> They determine ‘the object one is thinking of’ not by some imposed descriptive condition, but by means of a relation to that very token representation, whether immediate or remote. Of course, we may be able to provide a descriptive *characterization* of the token-reflexive constraints at work: for example, we might say that the internal analogue of ‘here’ is constrained to pick out the place of the thought token. But, crucially, in exploiting this relation to pick out its

<sup>46</sup> See ‘A Problem about Continued Belief’, p. 73.

<sup>47</sup> One might deny that all singular thoughts have singular content if one holds that a thought can be made singular by containing the mental analog of an empty referential term. (See Jeshion, ‘Acquaintanceless De Re Belief’.)

<sup>48</sup> For more on the importance of such semantic interconnections among thoughts, see Taylor, ‘The Psychology of Direct Reference’ and Fine, *Semantic Relationism*, chs. 3–4.

<sup>49</sup> Bach is explicit that the presence or absence of relational representation of an object does not line up with the truth or falsity of belief reports where a referential device to that object occurs in that that-clause. (In particular, he rejects the principle of HARMONY, discussed in our Chapter 2; see ‘Getting a Thing into a Thought’, p. 53.)

<sup>50</sup> Bach, *Thought and Reference*, p. 13.

object, the thought token ‘does not have to represent its [own] being in that relation to the object but merely has to be in that relation’.<sup>51</sup> By contrast, a satisfactoral object representation involves thinking of something ‘under a description,’ the satisfaction of which determines the object of the representation.<sup>52</sup>

Bach goes on to insist that the relation exploited by a relational representation must be ‘causal-historical’ and involve ‘a chain of representations originating with a perception of the object’.<sup>53</sup> One way of understanding this causal condition is that it is a mere stipulation about Bach’s use of ‘relational’ or ‘de re’. In that case, the only question is whether this way of dividing things up is explanatorily useful, or explicates a distinction that was pre-theoretically important and intuitive. As will become clear in what follows, we are doubtful on both points. Alternatively, Bach may be making a substantive claim to the effect that all object-representations that exploit relations (in the way sketched in the previous paragraph) must as a matter of fact exploit *causal-historical* relations. In that case, we see no reason to accept the claim. For example, one might well imagine that our primary internal representations of *here* and *now* are anchored by spatio-temporal rather than causal-historical relations. Or we might imagine that an internal representation can be *of* a particular place by disposing its bearer to move to that place given various conative attitudes. And so on. (We discuss causal constraints on singular thought further in the next two chapters.)<sup>54</sup>

Here is another concern. It is far from clear whether all object-representations fit smoothly into the relational/satisfactoral taxonomy as described. First, consider a mental name introduced by way of a reference-fixing description. This name will pick out the object it does thanks to the description’s being satisfied, and it does not exploit a relation, so it will not count as ‘relational’. On the other hand, arguably a given thought token that *deploys* that name need not itself represent the crucial descriptive condition—so in that sense it is not clearly ‘satisfactoral’ either. And if we go on to suppose that the relevant description actually involves a relation to tokens of the name in question—but one that is not causal-historical—then it seems to have even more claim to being relational, but would not quite fit the bill if the ‘causal-historical’ nature of the relation is *critical* to being ‘relational’.

One could say instead that relational object-representations are those that pick out objects without the need for *any* descriptive conditions represented by the thinker of the thought. But it is far from clear that any of our object-directed representations pass this test. Consider our object-directed representations of Socrates. It may be that causal-historical relations play a crucial role in explaining how such representations pick out Socrates and also that thinkers need not represent such causal-historical relations as obtaining in order for them to do their work. Still, it is plausible that

<sup>51</sup> ‘Getting a Thing into a Thought’, p. 55.

<sup>52</sup> *Thought and Reference*, p. 12.

<sup>53</sup> ‘Getting a Thing into a Thought’, p. 55.

<sup>54</sup> Thanks to Rachel Goodman for helpful conversations about this contrast.



such tags pick out Socrates rather than Socrates' flesh owing to the categorial constraints governing them that are internally represented: we intend to be thinking about the man and not the lump of flesh. More generally still, many metasemantic theorists hold that referential facts are fixed by a delicate balancing act of descriptive associations, causal connections, the relative naturalness/non-gerrymanderedness of candidates, and so on. Such factors together will determine, for example, when and whether any reference shift occurs in a given chain of object-representations. If this is right, failure of any given descriptive association may be quite compatible with referential success. But at the same time, referential success is hardly ever purely relational.

(iv) *'Believing of'*, etc. A final approach characterizing singular thought is one with which we have very little sympathy: namely, that of invoking locutions like 'believing of  $x$  that it is  $F$ ' or 'having  $x$  in mind'.<sup>55</sup> Insofar as these express ordinary notions, their acceptability does not track the presence or absence of the kinds of thoughts that theorists typically take to be singular. For example: using the ordinary sense of 'having in mind' it made perfect sense for Thurgood Marshall to say: 'When the Founding Fathers used this phrase ["the People"] in 1787, they did not have in mind the majority of America's citizens.'<sup>56</sup> But he was hardly pointing out the fact that the Founding Fathers did not have singular thoughts (plural or otherwise) about everyone in the country. Likewise, if he had said 'The Founding Fathers did not believe of all Americans that they deserved equal rights', the truth of his claim would have nothing to do with the whether the founders had singular thoughts about all Americans.

Unfortunately, as we will see, philosophers often surreptitiously employ such locutions as terms of art tied to the notion of singular thought. The result is that the line between the ordinary and theoretically-laden uses of these locutions has become hopelessly blurred.<sup>57</sup> (For more on this, see §§2.5 and 3.1.)

## 1.5 Acquaintance after Russell

The history of the notion of acquaintance has, since Russell, been marked by steady liberalization. The original picture was that genuine reference requires a kind of revelation or unmediated presentation. This extreme requirement is usually weakened in one way or another by contemporary theorists, though most still posit a special relation of epistemic or causal *rapport* that is necessary for reference and/or singular

<sup>55</sup> For particularly heavy use of the notion of 'having an object in mind' see Devitt, *Designation*. For an account that leans on 'believing of', see Soames, *Philosophical Analysis in the Twentieth Century, Vol 2: The Age of Meaning*, pp. 414–6.

<sup>56</sup> From Marshall's speech, 'Reflections on the Bicentennial of the United States Constitution'.

<sup>57</sup> When these terms are not explicitly introduced as terms of art, the suggestion is that they are simply used as expressions of natural language, and that readers should exercise their pretheoretic understanding of those expressions. But in fact it often seems that terms of art are in play whose artificiality is masked by the existence of similar expressions in ordinary speech. We recommend assessing the relevant philosophical texts while replacing these expressions with complete neologisms.

thought. (Dissenters include Ernest Sosa, Robin Jeshion, and David Kaplan in some of his later work.<sup>58</sup>) We have been told, variously, that the right relation to some object *o* requires that one ‘knows who/what *o* is’;<sup>59</sup> that one can recognize or individuate *o*,<sup>60</sup> that *o* has left a ‘vivid mark’ on one;<sup>61</sup> or that one can employ a singular term whose reference to *o* is ‘determined solely . . . by (appropriately shaped) causal chains’ grounded in *o*.<sup>62</sup> Unfortunately, characterizations of these relations seldom go beyond the off-hand and picturesque.<sup>63</sup> One gets the sense that theorists are not quite happy to leave such a significant notion unexplored, but at the same time they have not found anything very precise to say about it.

Still, we can divide these requirements into two broad categories: those that are epistemic, and those that are causal.

(i) *Epistemic acquaintance*. One leading proponent of an epistemic constraint has been Gareth Evans, who argued that singular thought requires *discriminating knowledge*: ‘the subject must have a capacity to distinguish the object of his judgment from all other things’.<sup>64</sup> Such a capacity might take the form of being able to identify an object in one’s perceptual field, being able to recognize the object if presented with it, or else (perhaps) being able to distinguish the object descriptively from everything else. An inability to discriminate the relevant object from every other would rule out having a singular thought about that object (even if, for example, the subject had a memory deriving from the relevant object).<sup>65</sup> And while one can master the use of a proper name merely by picking up a social name-using practice, one does not thereby ‘come to a proper understanding of what is said’ in one’s own remarks using the name. Acquiring the object’s name and thereby acquiring a distal causal link to the object is not sufficient for entertaining singular propositions about it.<sup>66</sup> (We focus on the epistemic approach to acquaintance in Chapter 3.)

<sup>58</sup> Jeshion, ‘Acquaintanceless De Re Belief’; ‘Donnellan on Neptune’; Kaplan, ‘Dthat’; Sosa, ‘Propositional Attitudes De Dicto and De Re’. Though see also Kaplan, ‘Afterthoughts’, pp. 604–7, where he hedges his view. Mark Crimmins also rejects acquaintance at least in certain cases: e.g. Crimmins, *Talk About Beliefs*, §3.3.

<sup>59</sup> The ‘knowing who’/‘knowing what’ test is common (though as we shall see in Chapter 3, the relevant knowing-wh expressions are sometimes being avowedly deployed with a quasi-technical meaning). For examples of this test see Russell, ‘Knowledge by Acquaintance and Knowledge by Description’, p. 113; Strawson, *Subject and Predicate in Logic and Grammar*, p. 47; Evans, *The Varieties of Reference*, ch. 4; Soames, ‘Beyond Singular Propositions?’, p. 518, fn. 5; Recanati, ‘Singular Thought: In Defense of Acquaintance’, p. 151. Others, such as Kaplan in ‘Quantifying In’, speak of ‘knowing *x*’ rather than ‘knowing who *x* is’.

<sup>60</sup> Evans, *op. cit.*, ch. 4.

<sup>61</sup> Kaplan, ‘Quantifying In’, p. 384.

<sup>62</sup> Boër and Lycan, *Knowing Who*, p. 128. This view is quite common: see fn. 67 below for some examples.

<sup>63</sup> One exception is the work of Igal Kvat: see Kvat, ‘A Theory of Thinker Reference’; ‘Mediated Reference and Proper Names’.

<sup>64</sup> Evans, *Varieties of Reference*, 89.

<sup>65</sup> *Ibid.*, Chapters 4–6 and 9, and especially pp. 89–92, 115–19, 132–5. Evans argues that the subject cannot *exploit* the uniqueness condition afforded by his single memory unless he explicitly thinks of the object as ‘the one from which his current memory derives’ (90). We explore this and related issues in Chapter 3.

<sup>66</sup> *Ibid.*, pp. 398–404.

(ii) *Causal acquaintance*. By far the most common way of characterizing acquaintance is that it involves some appropriate causal connection with the object of thought.<sup>67</sup> Obviously one could be more or less strict: perhaps only perception and memory suffice. But having allowed singular thought to extend beyond the Cartesian theatre, few philosophers want to deny that we can all have singular thoughts about (say) Kurt Gödel. As a result, many allow that the right sort of connection can be supplied by competence with a name that is linked causally, by way of a chain of name-users, to an initial baptism in the presence of the object.<sup>68</sup> Here Kaplan famously introduced the analogy of photographs: for a photo to count as *of* a particular person, ‘the person must serve significantly in the causal chain leading to the picture’s production’.<sup>69</sup> The idea is that names and even mental episodes are like photographs in this way: to count as representations *of* an object, there must be some causal connection to that object, however mediated.

On the other hand, it seems too strong to insist that one’s use of a name be linked by earlier uses of that name to an initial baptism in the presence of the named individual. Take ‘Homer’: for all we know, this name was introduced as a name for the author of the two epic poems, at a time when they were his only remaining traces. But it is counterintuitive to think that whether or not we are actually referring with ‘Homer’—or having singular thoughts about the poet of the *Iliad*—depends on whether someone actually named him ‘Homer’ in his presence before he died.<sup>70</sup>

As a result, many theorists have proposed even more liberal variants of the causal-acquaintance idea. For example, Nathan Salmon holds that, coming upon the scene of a murder, a homicide detective can form singular thoughts about the murderer. This is because the detective forms her thoughts using genuine causal traces of the entity in question, so that the ‘entity enter[s] properly into the “genetic” account of how the speaker came to learn the term he/she uses to refer to it.’ On the other hand, if Ralph

<sup>67</sup> Among many others, this view has been expressed by Donnellan, ‘The Contingent a Priori and Rigid Designators’; Davies, *Meaning, Quantification, Necessity*, p. 97; Salmon, *Frege’s Puzzle*, pp. 179–80, fn. 19; ‘How to Measure the Standard Metre’, pp. 199–214; Bach, *Thought and Reference*, part I; Neale, *Descriptions*, p. 18; Fitch, ‘Thinking of Something’; Recanati, *Direct Reference: From Language to Thought*; Soames, ‘Beyond Singular Propositions?’, p. 518, fn. 5; *Beyond Rigidity*, pp. 92, 329, 339.

<sup>68</sup> While many of those who play up a causal chain picture of acquaintance are inspired by Kripke’s view of the way in which names refer, there is no evidence of a *general* acquaintance constraint on reference in *Naming and Necessity* (Oxford: Blackwell, 1980).

<sup>69</sup> Kaplan, ‘Quantifying In’, pp. 198–201. Kaplan later changed his view: see Kaplan, ‘Demonstratives’, pp. 536, 560 n.76; ‘Afterthoughts’, pp. 603–7; and the hedged position in ‘Afterthoughts’ pp. 603–7.

<sup>70</sup> Kaplan makes a similar point about ‘Aristotle’ in ‘Afterthoughts’, p. 605. Of course, there are more complicated cases. Suppose, for instance, that the poet left his work unsigned and an otherwise unknown prankster inserted his own name at the head of each work. This is similar to the Ibn Khan case in Evans, ‘The Causal Theory of Names’, p. 19. Our intuitive verdict for these cases arguably contrasts with our verdict for Kripke’s Gödel/Schmidt case (Kripke, *Naming and Necessity*, pp. 83–4). Perhaps this is because, regardless of who proved the incompleteness of arithmetic, all the ordinary conventions that tie a name to an individual are still in place between the name ‘Kurt Gödel’ and the individual who, in Kripke’s example, is the imposter. (For example, one could look up official birth records for the individual with that name.) In the envisaged ‘Homer’ scenario, there is no intact and independent chain of name-usage connecting us to the imposter.

believes ‘solely on the basis of reflection on the concepts, that whoever is shortest among spies is a spy’, he is blocked from having singular thoughts involving the shortest spy, even if he introduces a name with a reference-fixing description.<sup>71</sup> In a similar vein, François Recanati allows that we can have singular thoughts about items known only by description, as long as the descriptions are *informative* as opposed to *reference-fixing* descriptions. For a description to be informative requires that it has a ‘target’ object that is ‘determined independently of its satisfying the description’. The information (or misinformation) encoded in an informative description is ‘anchored’ in the target object by way of a causal relation, however indirect. It is for this reason that an informative description may *incorrectly* describe its purported target.<sup>72</sup>

This kind of view only posits a weak causal constraint; and the weaker the constraint, the fewer the potential counter-examples from which it will suffer. But two points are immediately worth noting.

First, to motivate a very weak constraint, one cannot appeal to the original motivations that were offered for acquaintance constraints. To begin with, one must obviously give up the Russellian motivations: certainty of successful reference and discrimination of co-reference. But other more popular motivations must also be abandoned. For instance, many appeal to a special kind of belief-reporting data to motivate acquaintance constraints. In the case just sketched involving Ralph, even if Ralph introduces the name ‘Vladimir’ for whoever happens to be the shortest spy, it sounds wrong to say, ‘There is someone that Ralph believes is the shortest spy’. This fact could be explained by combining the idea that Ralph lacks acquaintance with the thesis that this sort of attitude report ascribes a singular thought. But weakening acquaintance will lead to many cases where the same problem arises. For example, suppose that in the detective case we allow the detective to have singular thoughts about the murderer through the latter’s traces. We will then need some other explanation of the fact that it still sounds wrong to say, ‘There is someone that the detective believes is the murderer’. We will examine these data carefully in the next chapter—but it is worth keeping in mind that as one liberalizes acquaintance in order to handle intuitions about some cases, one will thereby give up the ability to appeal to acquaintance when it comes to explaining our intuitions in other cases.

Another problem for weak acquaintance constraints is this. It is far from clear that thoughts connected to ‘traces’ or ‘informational anchors’ form a theoretically significant category of cognitive states or modes of epistemic access. Indeed, in many cases it is difficult to see how such a taxonomy would even apply. Consider ‘Julius’, whose

<sup>71</sup> Salmon, ‘The Good, the Bad, and the Ugly’, pp. 198–9. The original Smith’s murder example is from Donnellan, ‘Reference and Definite Descriptions’, and the ‘shortest spy’ case is from Quine, ‘Quantifiers and Propositional Attitudes’ and Kaplan, ‘Quantifying In’.

<sup>72</sup> Recanati, *Direct Reference: From Language to Thought*, pp. 111–2. This relatively liberal restriction also appears to be at work in the notion of a ‘prompt’ discussed by R. M. Sainsbury in ‘Reference and Anaphora’, §3. Sainsbury’s plausible examples of anaphoric reference appear to generate pressure to reject any strict acquaintance condition on reference.

reference we stipulate to be fixed by ‘the individual who invented the zip’.<sup>73</sup> Can we have singular thoughts about the inventor of the zip (assuming there is only one) because we have encountered a trace in the form of his creation? If not, would it be sufficient if we had learned that there was a unique inventor of the zip from someone acquainted with the inventor who nevertheless did not mention any name? Or consider Ralph and the shortest spy. Suppose that Ralph had learned that there is a unique shortest spy from the Spy Census, which would not have indicated this were it not for Orcutt (there being a tie for second-shortest)? In this case, the source of Ralph’s information is counterfactually dependent on Orcutt and his height; but we doubt that proponents of even the weakest acquaintance relation would allow Ralph to have singular thoughts about him.

In short, the notion of informational anchoring is shaky at best. At worst it is an *ad hoc* gesture, invoked whenever a theorist wishes to deny the presence of singular thought, towards whatever tenuous causal relation would explain its absence. But why think there is any such relation in the first place?<sup>74</sup> We will focus on that question in the next chapter.

(iii) *Split acquaintance*. Another way of liberalizing acquaintance constraints is to apply them only to thought and not to language.<sup>75</sup>

On one way of fleshing out such a view, there are cases in which one can refer to an individual even if one is not in a position to grasp singular contents about that individual. For example, perhaps one can use ‘Julius’ to refer to the inventor of the zip without being able to token singular thoughts about Julius, because one does not bear the right causal relationship to him or her. But this has some initially unappealing features. To begin with, it appears to require that we give up the intimate tie between sincere assent and belief—a connection adumbrated by Kripke as follows:

If a normal English speaker S, on reflection, sincerely assents to ‘p’, then S believes that p.<sup>76</sup>

In defense of split acquaintance, perhaps this could be replaced by a modified principle, in which the antecedent involves not only assenting to ‘p’ but *fully understanding* it. The

<sup>73</sup> Evans, *Varieties of Reference*, p. 31.

<sup>74</sup> After all, there may be no feature of object-representations at all (let alone a causal feature) that explains the range of intuitions related to singular thought. In particular, our intuitions about the acceptability of exported attitude ascriptions often seem more attuned to features of the ascriber’s situation than features of the subject’s situation. If they are context-dependent in this way, there may be no particular feature of the subject’s thought that will suffice to explain the range of intuitions at play. (More on this in Chapter 2).

<sup>75</sup> See e.g. McGinn, ‘The Mechanism of Reference’, p. 167; Evans, *Varieties of Reference*, §11.3 and §11.5, Salmon, ‘The Good, the Bad, and the Ugly’, p. 247; Soames, ‘Beyond Singular Propositions?’, p. 518, fn.5. Kaplan also expresses some sympathy towards a view like this in Kaplan, ‘Afterthoughts’, pp. 606–7.

<sup>76</sup> That is, as Kripke puts it, where ‘p’ is to be replaced, inside and outside quotation marks, by any appropriate standard English sentence’, and where ‘the sentence replacing ‘p’ is to lack indexical or pronominal devices or ambiguities that would ruin the intuitive sense of the principle’ and where a normal English speaker ‘uses the sentence to mean what a normal speaker should mean by it.’ Kripke, ‘A Puzzle About Belief’, §2.

idea would be that one can utter ‘Julius was inventive’ without understanding it in an appropriately robust way, and nevertheless succeed in expressing a singular proposition. (We will have more to say in Chapter 3 about the notion of robustly understanding a name.)

A related problem is that this view complicates a natural picture of communication on which one expects one’s interlocutors to understand the propositions one semantically expresses. Presumably, if the speaker does not fully understand ‘Julius’, then neither will interlocutors to whom he or she introduces the name. In fact it seems that a whole community of users might successfully employ a name without the requisite level of acquaintance. But the claim that people commonly express propositions that no one can grasp generates tension with the natural idea that meanings of utterances are—in some fairly direct way—parasitic on the contents they are used to communicate.<sup>77</sup>

Here is another variant on the split acquaintance view. Perhaps believing singular *contents* about Julius is easy but not sufficient for having *bona fide* singular *thoughts*. The latter require something additional, such as grasping a singular content *by way of* the right kind of cognitive vehicle—a file or tag, perhaps. The question then becomes: why would this latter step require causal *acquaintance* with Julius? Surely one can form the relevant cognitive vehicle, individuated internally, without the right causal connection to the purported referent. (This presumably happens in cases of illusion.) But it seems at least *prima facie* odd to deny that someone who both believes the relevant singular content about Julius and associates it with the right kind of internal vehicle to refer to Julius, nevertheless is in no position to have a genuine singular representation about Julius. We will return to this point in §2.4. For the most part, however, we will set aside concerns that are local to the split acquaintance view.

(iv) *Liberalism*. The thesis that there is no general acquaintance restriction on reference or singular thought we will call *liberalism*. Three points of clarification are in order.

First, liberalism is compatible with a variety of views about what sorts of cognition best deserve to be called ‘singular thought’. Some liberals will hold that a singular thought is simply the bearing of a propositional attitude to a singular content. Others will hold that the most natural kind in the neighborhood is a type of thought involving a special class of internal vehicles like ‘tags’ or ‘files’. Still others will hold that these distinctions coincide, or that they correspond to different joints of equal theoretical interest. What liberals have in common is that they think none of the important joints in the vicinity are governed by general acquaintance constraints. (That said, liberalism is

<sup>77</sup> Evans takes note of this issue in Evans, *The Varieties of Reference*, p. 400. For a useful discussion of the connection between the meanings of words and the meanings of speakers, see Grice, *Studies in the Way of Words*, ch. 6.

perfectly consistent with—and is indeed likely to accompany—the thesis that reference and singular thought *often* involve exploiting causal relations.)

Second, liberalism is compatible with various other constraints on singular thought in cases where one is *not* exploiting causal connections. For example, some may hold that one must in such cases exploit a reference-fixing description; others may allow the use of referential vehicles that do not exploit causal relationships or function like descriptive names; and still others may deny that singular thoughts must be realized by states involving any mental analog of a referential expression.<sup>78</sup>

Third, liberalism is compatible with various views on the truth of exported belief reports, such as ‘The shortest spy is such that Ralph believes that he is a spy’. Thus, a liberal may insist that this report can be true only if Ralph has a singular thought about the spy. (Indeed, the liberal may think that the best candidate for ‘singular thought’ is a phenomenon *defined by* the truth of such ascriptions.) Or a liberal could deny any interesting connection between singular thought and such belief reports. We will look more closely at exported ascriptions in the next chapter.

## 1.6 Should auld acquaintance be forgot?

Our first goal in this book is to argue that there is no good reason to impose acquaintance constraints of any sort on reference or singular thought. (This will pave the way for our second goal, which is to promote a certain approach to the semantics of descriptions, demonstratives, and names.) Acquaintance, in our view, is a dispensable relic of a bygone era in the philosophy of language and mind.

In arguing for liberalism we will adopt the working hypothesis that the terms ‘reference’ and ‘singular thought’ do pick out a semantic and a psychological natural kind, respectively—and that each exemplifies to some extent the confluence of themes discussed in §1.3 and §1.4 above. (We will revisit this working hypothesis in the book’s concluding remarks.) Of course, some theorists will emphasize one or another of these themes to a greater degree.<sup>79</sup> But our aim in the next few chapters is not to take a stand on which of these themes deserves emphasis. Our aim is to challenge the idea that acquaintance constraints apply to any kind of object–representation that is of sufficient theoretical interest to deserve the title ‘singular thought’.

(i) *Caveats.* Here are some related claims that we are *not* making. (a) We are not denying that, for example, linguistic reference often exploits a causal or epistemic connection.

<sup>78</sup> A liberal could, for example, take her cue from Salmon (who is not a liberal) in holding that if one has a general belief of the form, ‘The murderer is insane’, one will automatically have a second and distinct belief in the singular proposition involving the murderer himself. (For him, this is modulo the requisite acquaintance relation.) See ‘The Good, the Bad, and the Ugly’.

<sup>79</sup> Thus, in postulating constraints on singular thought, some are primarily making claims about the conditions for standing in various attitudes to singular contents, while others are primarily concerned with the conditions for tokening inner representational vehicles of a special type, such as ‘tags’ or ‘files’.

(b) We are not denying that there are any acquaintance constraints when it comes to various subclasses of object-representation, delineated by a taxonomy of mechanisms or modules. For example, perhaps object-representations that occur in a particular vision module in humans always represent an object due, in part, to causal relations between it and the object represented. Or perhaps, more generally, causal interaction is a necessary condition on successful *perception*, and therefore on referential intentions mediated by perceptions.<sup>80</sup> We have no desire to argue with claims made at this level of specificity. (c) We are not rejecting anyone's right to announce that they will reserve the term 'singular thought' for cases where an appropriate causal or epistemic connection is at work.<sup>81</sup> We would simply point out that such a stipulation would trivialize the claim that singular thought requires acquaintance, and we see no reason to think that the stipulated term would delineate an interesting kind from the point of view of a general theory of meaning and representation.

Our target, then, is acquaintance proposed as a substantive and general constraint on meaning, thought, or understanding. This is the idea often expressed with (albeit vague) slogans like 'Reference and singular thought require acquaintance'.

(ii) *Burden of argument.* We are liberals primarily because we see no reason to accept an acquaintance constraint, and we have never encountered a specific proposal that seemed at all plausible.<sup>82</sup> In general, we are completely puzzled by the kind of policing or gate-keeping that goes on when philosophers makes speeches of the form, 'In order to have a genuine singular thought/have a genuine *de re* thought one must satisfy constraints X, Y, and Z.' Such claims are invariably either true and uninteresting—because some special use of 'singular thought' has been stipulated—or else false on any good precisification.

Those who postulate substantive acquaintance constraints are taking on an argumentative burden, and we will be examining their positive arguments in Chapters 2 and 3. However, we suspect that some of those who find acquaintance constraints plausible are not leaning very heavily on those positive arguments. They may simply find it natural to generalize from certain paradigm cases of reference in which causal or epistemic relations do play a significant role in securing reference: and certainly many ordinary uses of names and demonstratives are like this. But one should not fixate on a restricted range of cases. When one thinks through a sufficient variety of examples, a

<sup>80</sup> There may be exceptions. Perhaps one perceives a hole when only its lining is causally active; or perhaps one perceives a statue when only the lump is causally active; or perhaps one can perceive an absence of bananas without the absence exercising a causal influence.

<sup>81</sup> A natural way of reading Bach, for example, is that a causal acquaintance constraint is simply built into his conception of singular thought: see §1.4.ii above. Consequently, to avoid a terminological dispute, our disagreement with Bach may be thought of as concerning whether his taxonomy of thoughts is more theoretically useful and better cuts psychological reality at the joints than one on which the key notion is not burdened with a causal condition.

<sup>82</sup> We do offer some positive considerations in favor of liberalism in Chapter 2, at least for those who wish to retain a close connection between exported belief reports and singular thought.



general acquaintance constraint is not *prima facie* compelling after all. To illustrate this point, we offer in what follows a sample of cases where standard heuristics suggest that reference and singular thought are present, but where this or that candidate acquaintance relation is absent. For now we will be focusing on causal acquaintance constraints: we will set aside epistemic acquaintance until Chapter 3.

To detect the presence or absence of singular thought in the following cases, some readers will proceed via judgments about the truth of exported belief reports, or reports with referential terms in their complement clauses.<sup>83</sup> (We shall, for now, assume with semantic orthodoxy that names, demonstratives, and simple indexicals are referential terms.)<sup>84</sup> Others will proceed via judgments about whether subjects can use a referential expression and grasp the propositions they express with it. And still others will ask themselves whether tags or files are present in the subject's psychology, and whether these succeed in referring. Readers are free to consider their own preferred heuristics, as long as those heuristics do not import an assumption of causal acquaintance: we continue to assume that those who postulate an acquaintance constraint on singular thought intend to put forward a substantive claim. For this reason we recommend against diagnosing the presence or absence of singular thought by asking oneself whether the speaker 'knows who *x* is' or 'believes of *x* that she is *F*' or 'has *x* in mind'. As we mentioned above and will be arguing further in later chapters, these expressions have been hijacked for a theoretical use in the philosophy of language—and terms of art cannot be mined for diagnostic, pretheoretical intuitions. (Meanwhile, the original ordinary expressions homonymous with these terms of art are of no special use when it comes to identifying singular thought.)<sup>85</sup>

We have grouped our discussion of such cases, rather arbitrarily, into eight broad ways of achieving reference: by reverse causal chain, by decomposition, by character, by convention, by conspicuity, by proxy, by conjecture, and by depiction.<sup>86</sup> (As will be clear, many of our examples are variations on cases from the literature.) While we do not pretend that every case causes trouble for every causal acquaintance constraint that has been proposed, we can think of none so lenient it can properly handle them all.

(iii) *Reverse causal chains*. Those who take there to be a causal acquaintance constraint on singular thought typically suppose that the object of thought must in some way causally

<sup>83</sup> A special word of warning concerning belief reports in these cases. As we argue in Chapter 2, the acceptability of exported belief descriptions is heavily context dependent. As a result, in many of the cases that follow, it will be possible to construct exported belief ascriptions that are not acceptable as well as those that are. It serves our purposes well enough to notice that it is easy to flesh out the context in ways that generate acceptable reports. As we shall see, the negative data are of no special importance since they can just as easily be generated in cases where candidate acquaintance relations are present.

<sup>84</sup> The second half of this book will complicate this assumption significantly, but not in a way that threatens the thesis operative in the following examples; *viz.* that singular thoughts are typically expressed when sentences contain proper names and demonstratives, and that singular thoughts are typically ascribed when proper names and demonstratives appear in the complement clauses of attitude ascriptions.

<sup>85</sup> See ch. 2, §2.5, ch. 3, §3.1, ch. 4, §4.1.

<sup>86</sup> We by no means suppose that this taxonomy is exclusive or exhaustive.

impact the subject. But as Ernest Sosa and Robin Jeshion have pointed out, sometimes the causal chain runs in the opposite direction.<sup>87</sup> For example, in 1512, Henry VIII ordered the construction of a great warship to be called *Henry Grace à Dieu*. Suppose he knows that his order will be carried out on time, but is given no reports of the progress of the ship on the principle that no news is considered to be good news. After the expected date of completion, he makes plans to visit the harbor. Surely he succeeds in referring to the ship if he says ‘I plan to see *Henry Grace à Dieu*.’ Moreover, we would not withhold an exported ascription: there is a warship such that King Henry believes he will see it.<sup>88</sup> But his causal link to the ship leads in the wrong direction for him to satisfy a standard causal acquaintance constraint. Neither is any of his information about the ship ‘derived’ from it.

(iv) *Reference by decomposition.* In the ordinary sense of ‘perceives’, one can surely perceive an object even if parts of it are hidden. One can perceive a dog without perceiving its organs or hidden surfaces, and one can see a city without seeing many of the buildings of which it is composed. Moreover, when one is perceptually confronted with an object, it often seems that one can name—and have singular thoughts about—both the parts that are perceived and the parts that are not. For example, suppose a mechanic points at a car and says, ‘Let me see that engine’. We would not want to deny that there is a particular engine such that the mechanic wants to see it. And yet there may be no causal relation here to speak of.

Might one suggest that it is sufficient for acquaintance with *o* to be causally connected to something that is mereologically related to *o*? This threatens to make the condition quite vacuous: if we are all both causally and mereologically related to the biosphere, it will follow that we are all acquainted with each other.

(v) *Reference by character.* Suppose we adopt a roughly Kaplanian account of indexicals. Then there is a linguistic rule for every indexical—its ‘character’—that determines its reference with respect to a context. Must this kind of rule be restricted to specifying some *causal* relation between the referent and the speaker or the token of utterance?<sup>89</sup> If not, the anti-liberal who proposes a causal constraint will have to maintain that we are somehow barred from properly using these terms as vehicles for expressing singular thoughts in cases where we are not causally connected to the object determined by the character.

<sup>87</sup> They both describe examples where one names something before bringing it into existence. (Sosa, ‘Propositional Attitudes De Dicto and De Re’, p. 889; Jeshion, ‘Descriptive Descriptive Names’, p. 609; ‘Ways of Taking a Meter’, p. 300). Kaplan even mentions such examples in ‘Quantifying In’, 214. However, we do not wish to take a stand on whether it is possible to have singular thoughts about as yet non-existent objects: see the following footnote.

<sup>88</sup> Intuitions may waver a little as to whether Henry can refer to the ship *before* it is built. This is because many of us have presentist intuitions according to which, in the broadest natural sense of ‘exists’, only the present exists. (See see ch. 2, §2.5.) But presentism is no threat to liberalism. Meinongians aside, we can all agree that only existing things can serve as relata to the reference relation.

<sup>89</sup> For an affirmative answer, see Devitt, *Designation*.

The problem, as Colin McGinn pointed out decades ago, is that some indexicals refer to objects that have no causal effect on us.<sup>90</sup> For example, whatever linguistic rule determines the referent of a given utterance of ‘tomorrow’ cannot involve the speaker or the token being causally affected by the interval of time referred to. (In conversation, a proponent of acquaintance suggested—half in jest—that we can refer to tomorrow because it ‘touches’ today.)<sup>91</sup> One would need substantial positive motivation to accept ‘today’ but not ‘tomorrow’ as a pure indexical that can express singular contents.<sup>92</sup> Moreover, we have no hesitation using ‘tomorrow’ in the complement clauses of attitude ascriptions, or endorsing the relevant exported ascriptions. Related points can be raised for terms referring to other spatial or temporal regions.<sup>93, 94</sup>

At any rate, the proponent of a causal acquaintance constraint faces a difficult choice: she must either (a) cast about for an additional non-causal relation that suffices for acquaintance; (b) let semantic theory dictate on issues in speculative metaphysics, such as whether we are causally related to the current moment; or else (c) suspend judgment on whether we have thoughts about the present time or place, pending a consensus on the metaphysical issue. None of these options seems particularly appealing.

Finally, let us extend our inquiry to merely possible indexicals. Must the linguistic rules for indexicals *in principle* operate via those relations that are allegedly necessary for acquaintance? Imagine a language community containing many tribes, each of which has a fairly well delineated geographical region to call its own. They have a term, ‘Chief’, which they use in a paradigmatically referential way. Its rules determine that it refers to whatever individual is chief of the location in which the utterance occurs. Just as ‘today’ and ‘tomorrow’ can suddenly shift their referents at midnight without one’s knowing it, ‘Chief’ shifts its referent as soon as one crosses the boundary from one tribe to another. (We can suppose that various tests on the linguistic intuitions of the community, having to do with embedding in modal and attitude contexts, bear out our classification of the term.) Is this *in principle* a defective term, so that such a community could not succeed in their referential intentions?

<sup>90</sup> McGinn, ‘The Mechanism of Reference’, p. 160. However, McGinn himself is a proponent of the split acquaintance view: see p. 167.

<sup>91</sup> Obviously, acquaintance with *o* had better not be conferred by the fact that *o* touches something with which one is acquainted, lest we thereby count as acquainted with everything touching the Earth.

<sup>92</sup> One might be attracted to treating ‘yesterday’, ‘today’, and ‘tomorrow’ as partly descriptive but involving reference to the present moment; e.g. they are elliptical for ‘the day before the day it now is’, ‘the day it now is’, and ‘the day after the day it now is’, respectively. However, we think it would be misguided to motivate this by the idea that ‘yesterday’ decomposes into ‘yester-’ and ‘-day’; compare the archaic use of ‘yesteryear’.

<sup>93</sup> Do the referents of ‘here’ and ‘now’—assuming these function as referring devices—have causal effects on us? Even the case of ‘I’ is not straightforward: what causal relationship does one bear to oneself? Both McGinn (op. cit.) and Wettstein have emphasized the oddness of the view that the linguistic rules that govern ‘I’ and ‘now’ would determine referents by way of causal relations (Wettstein, *Has Semantics Rested on a Mistake?*, p. 73).

<sup>94</sup> Consider also the second-person pronoun. Suppose that Smith knows on general grounds that he has a local Member of Parliament and writes her letters. Among other things he writes, ‘You are not doing enough to prevent the takeover of small shops by large chains.’

(vi) *Reference by convention*. A similar point can be made about proper names. Consider a community that assigns every newborn a number, by way of a secret algorithm that takes time and place of birth as inputs.<sup>95</sup> Thereafter one is referred to by the corresponding numeral: say ‘9200386’.<sup>96</sup> Moreover, if the number is miscalculated it is assumed that one’s *true name* would be the correct output of the algorithm. In short, the intention of the community is to exploit the algorithm to generate names, in *place* of baptism. There is no such thing as baptism, by their lights: there is a fact of the matter about one’s true name, which is settled by the community’s strong convention plus facts about the time and place. Surely a community with such linguistic conventions is *possible*. Now suppose a member of such a community is discussing a distant relative with whom she has had no direct causal contact. Instead of exploiting a causal chain of name-uses—as we would—this individual attempts to exploit a convention involving an algorithm. Does she succeed? On the one hand, it is implausible to insist that she cannot have a singular thought about that individual. On the other hand, it is implausible to suppose that she succeeds *despite* her linguistic intentions and those of her whole community, by unwittingly tapping into some kind of causal chain.

(vii) *Reference by conspicuity*. Paradigmatically, a speaker accompanies a demonstrative with an act of *pointing* at an object or otherwise rendering it salient by shaking it, gazing at it, or aiming a flashlight beam at it. But in a wide range of contexts no complement gesture is required. As Evans puts it, ‘a speaker can exploit some extreme or heightened salience which an object has anyway (without his bringing it about)’.<sup>97</sup> An object can be made conspicuous in any number of ways.<sup>98</sup>

We need take no position on the question whether gestures and other salience cues themselves ‘determine reference in the sense of making it the case that a certain [object] is being referred to’ or whether they merely enable the audience to identify the referent that has been determined by the intentions of the speaker.<sup>99</sup> What concerns us is whether the mechanisms by which the reference of a demonstrative expression is fixed must involve a special *causal* relation to the object. We see no *prima facie* reason to think that they must.<sup>100</sup>

<sup>95</sup> A similar example—used to make a different point—can be found in Wettstein, *op. cit.*, p. 5.

<sup>96</sup> A given numeral, of course, could serve both as the name of a number and the name of an individual, just as many individuals are named ‘John’.

<sup>97</sup> This quote and flashlight example are from Evans, *The Varieties of Reference*, p. 312.

<sup>98</sup> Consider these pointing-free demonstratives:

- i. This is beautiful (said with closed eyes during a musical performance).
- ii. It is more impressive than I imagined! (as the speaker comes within view of Mount Everest).
- iii. He really beat me (Jimmy Carter to Rosalyn, waking up in the night after his debate with Reagan).

The last example is from Wettstein, who notes that relevant cues ‘are not limited to those provided by the context of utterance. . . . [and] can be provided by what the addressee knows about the speaker’s interests, desires, history, and so on’. (Wettstein, *Has Semantics Rested on a Mistake?*, pp. 78–9.)

<sup>99</sup> Bach, ‘Intentions and Demonstrations’, p. 144. Bach defends the latter view, as does Kaplan in ‘Afterthoughts’. See *inter alia* Reimer, ‘Do Demonstrations Have Semantic Significance?’ for the former view.

<sup>100</sup> *Contra e.g.* Devitt, *Designation*, pp. 42–3.

Suppose Jones learns in some indirect fashion that his old friend Smith is now engaged to a woman (perhaps via learning the general fact that everyone in a group to which Smith belongs is engaged) and that there is a virus afflicting women in Smith's town. Jones calls Smith and (before Smith speaks) says: 'Is she all right?' The simple fact that Jones knows that Smith's fiancée will be foremost on his mind at this time appears *prima facie* sufficient for Jones to be able to refer to her.

Here is a second case. John crosses over a summit to the other side—a location David has never seen. David calls out: 'Can you see the city from there?' John replies that he can indeed see the city 'from here', and they go on talking about that location. Must David be within sight of the location to have referred to it, or to grasp what John is saying? (What if he is blind: what would count as sufficient aural or tactile contact with a location?) A causal acquaintance constraint yields counterintuitive results.<sup>101</sup>

(viii) *Reference by proxy*. Deferred demonstratives—whereby one refers to an object absent from the perceptual scene by exploiting some salient prop—appear every bit as referential as ordinary demonstratives. On the face of it, a speaker expresses the same proposition when she says, 'That is my friend' whether she is (a) pointing at Bob; (b) pointing at a door through which Bob has just left; or (c) pointing at a photograph of Bob. Even if one somehow convinces oneself that the relevant linguistic device refers to the prop in (b) and (c), it remains natural—by the lights of standard heuristics—to think that there are associated singular thoughts directed at Bob that are being communicated.<sup>102</sup>

Does this phenomenon ever involve flouting acquaintance constraints? As we have seen, only the weakest sort of acquaintance restriction allows that the homicide detective is acquainted with Smith's murderer, or that we can come to have singular thoughts about someone just by seeing his footprints in the snow. But pointing to the grisly scene and the footprints, the detective can say 'He's insane,' referring to the murderer. (And as we will see in the next chapter, there will be plenty of contexts where subsequent exported belief ascriptions will be acceptable.) Similar examples abound. One points to a beautiful office and says, 'That person is lucky to have such a nice office'. Or one points to a heavily vegetated island and says, 'Those people never go hungry'. Again, *prima facie*, there seems no reason to deny the semantic well-functioning of referential terms and intentions in these cases. But, as we will see, the project of weakening acquaintance in such a way that it both accommodates these cases

<sup>101</sup> These cases bear some similarity to those in McGinn, 'The Mechanism of Reference', pp. 60–4.

<sup>102</sup> We are not ourselves very partial to the view that linguistic reference is to the prop. Here is Emma Borg: '[T]he sort of sharp distinction between deferred and nondeferred case . . . never really existed. Once we recognize that I can refer to you by pointing at your arm, or at just that part of you that is visible through a door that is slightly ajar, or at the tail of your coat as you leave the room, or at your image in the mirror, or at your photograph, or at your shadow, etc., the idea of drawing a semantic distinction at any point on this scale comes to seem quite hopeless' (Borg, 'Deferred Demonstratives', p. 226). Along these lines, one might say that reference by proxy is just a special case of reference by conspicuity: the prop helps make the relevant object conspicuous.

and still does the theoretical work that motivates it (see Chapter 2) appears quite hopeless.

(ix) *Reference by postulation.* We can begin with an example famously discussed by Kripke. In 1846, Urbain Leverrier reasoned that certain perturbations in the orbit of Uranus could be elegantly explained by postulating the presence of a new planet in a specific celestial location. He sent his results to the Berlin Observatory, which discovered the planet with a telescope in less than half an hour: Leverrier had predicted its location to within 1 degree.<sup>103</sup> Now let us add the fictitious supposition that before the planet had been seen, Leverrier stipulated that the name ‘Neptune’ would refer to whatever planet was the perturber. We might also imagine him pointing at the location where he had determined the planet should be, though unable to see it with his weak telescope, and saying, ‘That is the cause of the perturbations!’ There seems no *prima facie* reason to deny that Leverrier succeeds in referring to the planet with ‘Neptune’ or ‘that’, despite the lack of perceptual contact. As John Herschel put it in 1846, speaking before Neptune had been found with a telescope:

We see it as Columbus saw America from the shores of Spain. Its movements have been felt, trembling along the far-reaching line of our analysis, with a certainty hardly inferior to that of ocular demonstration.<sup>104</sup>

The acquaintance lover might object that in this case the reference has been fixed in a way connected to causal traces, and insist that the perception of perturbations in the orbit of Uranus put Leverrier and Herschel in sufficient causal contact with Neptune. But there are other cases of reference by postulation where this point cannot be made.

Consider, for example, a variant on an example due to Robin Jeshion.<sup>105</sup> In 1870, Mendeleev stated that it was possible to predict the properties of certain elements even though no samples of those elements had yet been discovered. He made predictions for, among others, an element he called ‘eka-aluminum’. Five years later, samples of eka-aluminum were discovered by Lecoq de Boisbaudran and named ‘gallium’. Here again, it seems that Mendeleev succeeded in referring to the element in speech and in thought, and we are happy to say (for example): ‘In 1870 there was a substance such that Mendeleev thought samples of it would soon be discovered’. Apparently, reference to a substance-kind requires no prior acquaintance with its terrestrial instances.<sup>106</sup>

<sup>103</sup> Similar calculations were made independently by John Couch Adams.

<sup>104</sup> From a speech given on the 10 September 1846, to a meeting of the British Association in Southampton. For more on Herschel, see Standage, *The Neptune File*.

<sup>105</sup> See ‘Donnellan on Neptune’, p. 113; ‘Ways of Taking a Meter’, p. 300.

<sup>106</sup> This example is more problematic for the anti-liberal than the fact that we can refer to abstract objects like numbers even though they have no causal connection to us; see e.g. Putnam, ‘What Is Mathematical Truth’, p. 59.) The anti-liberal’s typical response to that point is to posit a kind of mental connection with abstract entities in order to promote an appearance of unity in the theory. But usually reference to natural substance-kinds like water and gold is thought to operate very much like reference to ordinary objects—paradigmatically, someone perceives an instance of the kind and introduces a referential device by ostension, and then others who pick up that term by way of a chain of uses can do so as well.

Finally, consider ‘God’, which is arguably a proper name. (If it were short for some definite description like ‘the all-powerful, all-knowing . . .’ it would be hard to explain the controversy among theists about which among these attributes God enjoys.<sup>107</sup>) Atheists hold that ‘God’ fails to refer, while theists typically hold that humans have causally interacted with its bearer. But suppose it turns out that there is exactly one all-powerful, all-knowing, and completely benevolent spirit whom we will meet in the afterlife; and yet this spirit has not causally interacted with the universe. (Perhaps the universe has always existed, or was created by much lesser beings who have themselves have always existed.) It seems that on this scenario, ‘God’ refers to the great spirit, and theists have correctly held that God exists, but wrongly thought that God has interacted with the universe. Once again we have referential success unmediated by causation.

(x) *Reference by depiction.* Imagine Leverrier with a map of the fixed stars like that on the cover of this book, working out the path of the postulated planet and adding a mark with the intention of representing it in its current location. Suppose Leverrier points to the mark and says ‘That is the perturber.’ The mark then is a perfectly adequate prop for demonstrative reference. If there is no reason to deny that Leverrier’s mark represents the planet, why deny that his demonstrative does so?

Here again the acquaintance lover may hold that this case involves causal traces that are sufficient for acquaintance. But consider another example. The seventeenth-century cartographer John Speed mapped the coastline of the British Isles. Suppose he is told by sailors about the coastline of a large, unexplored island. He draws an accurate map of the coastline of this island, then points at an area on his map representing an area within the island that is well beyond sight of the coastline and says, based on purely general considerations: ‘It is windy here’ or ‘It is windy at that location’. Even if no one has had any causal contact with the place he’s pointing to, it would seem odd to deny that part of Speed’s map represents that very place, that he has a belief about it, or that he has just referred to it.<sup>108</sup> And yet this part of the island has not had even an indirect causal effect on him.

Of course, maps need not be set down on paper or pencil. People frequently make rough egocentric mental maps, using a few proximate landmarks to set the axes. We cannot think of any reason to deny that such representations could serve as devices allowing a sort of deferred demonstration.

<sup>107</sup> See, for example, Whitehead’s *Process and Reality*, Hartshorne’s *Omnipotence and Other Theological Mistakes*, and Kushner’s *When Bad Things Happen to Good People*. One might propose that the definite description involved itself makes reference to descriptions or attributes, as in ‘the best candidate for satisfying the description, “the all knowing, etc.”’, or ‘the being that has the most of the following attributes: being all-knowing, etc’. Such a proposal appears implausible enough for us to ignore it here.

<sup>108</sup> On using maps as props for deferred demonstration where acquaintance is not at issue, see e.g. Evans, *Varieties of Reference*, pp. 161, 162; Kaplan, ‘Demonstratives’, p. 491; Recanati, ‘Are Here and Now Indexicals?’, §3.2.

An acquaintance constraint on reference can seem natural if one assumes that *any* representation of an object must be causally connected with that object somehow. And this in turn can seem tempting if one takes the case of photographs as paradigmatic. A photograph only counts as *of* so-and-so if that individual plays an important causal role in its production.<sup>109</sup> (Thus Kaplan called his view the ‘photograph model’ of naming.<sup>110</sup>) But why model mental or semantic representation after photography? Even among less technological types of graphic representation, there doesn’t seem to be any corresponding constraint.

Consider what it takes for a given painting to count as a depiction *of* someone. An artist can, apparently, choose to depict anyone that she can uniquely identify. There need not be any particular causal link or relation of similarity between the portrait and its subject. Take, for example, Salvador Dalí’s portrait of Juan de Pareja, assistant to Velázquez. What makes it the case that his painting is a depiction of Juan de Pareja rather than, say, Charles V? It is certainly not a matter of *resemblance*. Neither do we worry about the robustness of Dalí’s informational link to the subject. The painting counts as a portrait of Juan de Pareja—simply because that is what Dalí intended it to be.

Just as a cartographer can represent regions with which he had no causal interaction, it seems that a painter can depict things based only on general information or even speculation. A forensic artist might piece together evidence to create a portrait of the culprit. Seeing a mountain, an artist can imagine and depict—however inaccurately—the other side. Another might portray an event from prehistory such as the discovery of fire,<sup>111</sup> or the emergence of a galaxy whose existence is inferred from a general theory of the formation of the universe. With the right intentions, these portrayals will all count as depictions *of* their intended subjects. And we see no reason to deny that the representational contents of these images—assuming they have contents—will involve the objects that are depicted, or that they can be used as props for deferred demonstratives.<sup>112</sup> If an artist can represent a subject on a cloth canvas, it would seem, she can also represent that subject on the canvas of imagination.

Why hold that our powers of representation are limited by the causal connections that photography exploits? The relatively artificial fashion in which objects are depicted by

<sup>109</sup> It would be a mistake to think that light bouncing off the relevant object is literally required: think of a silhouette in which the represented object serves only to block light from a distant source, but reflects no light of its own in the direction of the camera. The precise causal role that an object must play in producing photographs of it is unclear.

<sup>110</sup> Kaplan, ‘Quantifying In’, §9.

<sup>111</sup> While Magritte did not seriously intend to depict the historical event with his *The Discovery of Fire*, Piero di Cosimo’s *The Forest Fire* has been argued to have a more serious intent in that respect. (Thanks to Tim Williamson here).

<sup>112</sup> Of course, since photographic representation involves a causal link, even a liberal should acknowledge that to demonstrate a person in absentia by way of a photograph is to exploit a causal relationship between the referent and the photograph; and reference would fail if there were no such causal link.



film is unlikely to be the best model for the way in which objects are represented by mental states—and the expressions we use to give voice to them. As Dalí expressed it: ‘What is a television apparatus to man, who has only to shut his eyes to see the most inaccessible regions of the seen and the never seen, who has only to imagine in order to pierce through walls?’<sup>113</sup>

## 1.7 Gameplan

At the beginning of this chapter we presented a number of features associated with referential terms. Those familiar with recent traditions in analytic philosophy will recognize that these features are standardly associated with a cluster of putatively special expressions in natural language—among them names, demonstratives, and pure indexicals. Insofar as one takes the various features we have discussed to converge on these terms, one can very easily feel that a semantic natural kind is coming clearly into view. And many see an even richer and more dramatic confluence of themes: a natural psychological kind corresponding to the semantic kind, and an acquaintance constraint that governs both the understanding of referential expressions as well as the grasping of singular thoughts. To what extent is such optimism warranted? Our aim in this book is to undertake some critical self-examination.

(i) Procedural points. Let us be clear about certain ways that our inquiry is limited or conservative.

First, with the semantic mainstream, we assume that intuitions about the truth and falsity of sentences at contexts are a reasonable guide to their semantic content. We thus will be setting aside views that utterly divorce the truth conditions of sentences at contexts from ordinary judgments of truth and falsity: for example, views on which ‘There is nothing in the fridge’ is always false when there are atoms in the fridge.<sup>114</sup>

Second, we assume for the most part that declarative sentences, as uttered at contexts, are typically assessable for truth and falsity. We thus do not in any systematic way explore views that claim that the semantic contents of such utterances are fragmentary or skeletal, requiring extra pragmatic supplementation in order for questions of truth-value to arise.<sup>115</sup> However, much of what we say could be adapted to such frameworks. (In particular, our search in Part II for the semantically referential

<sup>113</sup> Janet Wasko, *A Companion to Television*, 9.

<sup>114</sup> See e.g. Cappelen and Lepore, *Insensitive Semantics*. For critical discussion see Cappelen and Hawthorne, *Relativism and Monadic Truth*, ch. 2.

<sup>115</sup> As Scott Soames puts it, the idea is that the semantic content of an utterance is frequently ‘only a skeleton, or partial specification’ of a proposition; but ‘whether it is a complete proposition or not’ it ‘interacts with an expanded conception of pragmatics to generate a *pragmatically enriched proposition* that it is the speaker’s primary intention to assert’ Soames, ‘The Gap between Meaning and Assertion’, p. 5 of the ms. See also Bach, ‘Conversational Implicature’ and the literature arising in connection with Sperber and Wilson, *Relevance*. There are many related views in the neighborhood, depending on how exactly one delineates the line between semantics and pragmatics.

expressions of natural language could arguably be adapted to a search for those expressions that we typically use to communicate singular contents: see §4.10.i.)

Third, we will be setting aside certain questions in foundational semantics that are pertinent to a fully satisfying story about reference. These include: What is the difference between the way a referential term is related to the object to which it refers, and the way a predicate is related to a property it expresses? And how must the metaphysical ground floor of reality be configured for a binary reference relation to arise between an object and an occurrence of a term? These are good questions, but they lie outside our domain of inquiry.

Instead, we aim to accomplish three tasks: (a) to exhibit the inadequacy of modern successors to Russell's ideas about acquaintance; (b) to challenge the standard taxonomy of referential terms in natural language by exploring a more unified semantic approach to several of the expressions we use to talk about particular individuals; and (c) to reconsider the categories of reference and singular thought in light of the foregoing.

(ii) *Synopsis.* Here is our plan for the rest of the book. In the next chapter we defend liberalism in the face of two general types of argument. The first of these claims that an acquaintance constraint is required to explain certain kinds of *belief-reporting data*, and the second uses acquaintance to appease our intuitions about certain cases of apparently unwarranted *epistemic advance*. In both cases we argue that upon closer examination the relevant data are actually better explained by liberalism. In Chapter 3 we turn our attention to various forms of epistemic acquaintance, arguing that none of them provides a plausible condition on reference or singular thought.

The liberalism defended in Part I does not provide much positive insight into which natural semantic relation best deserves to be called 'reference'. But it does fit nicely with a certain semantic picture of the various expressions that we use to talk about particular individuals—a picture we present in Part II. In Chapters 4–6 we examine a range of noun phrases in natural language, inquiring whether there is an interesting referential/non-referential distinction among them. In particular, we look at specific indefinite descriptions, definite descriptions, demonstratives, and names. We suggest that liberalism allows for a unified picture of all these expressions on which they single out objects by way of restricted quantification that is sometimes covert. If we are right, none of these noun phrases paradigmatically fits the profile of a referential term, though they all admit of *uses* that exhibit many of the traits associated with reference. On our preferred framework, this phenomenon is due to the presence of what we call a *singular restriction* on the existentially quantified domain.

In the Afterword, we re-evaluate our central questions, examining the possibility that no expressions of natural language are paradigmatically referential—the only dedicated vehicles of reference are cognitive object representations of a special kind. If this view is correct, it lends additional theoretical significance to the category of singular thought.

# 2

## A defense of liberalism

Liberalism about singular thought rejects the following principle:

CONSTRAINT: To have a singular thought about an object, one must be acquainted with it.

This chapter addresses two important arguments for CONSTRAINT. The first appeals to the intuitive truth conditions of certain belief reports whose complement clauses contain either singular terms or variables bound from outside the clause. The second claims that denying CONSTRAINT leads to an absurd conclusion: namely, that introducing a name by way of a definite description can yield illicit epistemic advantages such as new contingent *a priori* knowledge. While these arguments are closely related, they are best addressed individually.

### 2.1 The spy argument

The first argument can be illustrated with an example of Kaplan's from 'Quantifying In'.<sup>1</sup> The ubiquitous Ralph—as he appears in this example—simply believes the general truth that there are spies, without suspecting anyone in particular. Moreover, he believes that there are finitely many spies and that no two spies are exactly the same size, since that would be very unlikely. He then deduces that there is a unique shortest spy who is a spy. Now compare:

1. Ralph believes that at least one person is a spy.
2. There is at least one person that Ralph believes to be a spy.

<sup>1</sup> See pp. 192–3. In that article, Kaplan does not himself endorse an argument exactly like the one that follows. For one thing, the expression 'singular thought' did not yet have a life of its own. The issue was simply to account for belief ascriptions that involve quantification-in. In explaining them, Kaplan appeals to 'genetic' connections between names and their bearers, a constraint that he later modified. (See fn. 69 of our ch. 1.) Note that Kaplan does not here treat 'There is someone Jones believes to be a spy' as having the logical form ' $\exists x$  (Ralph believes  $x$  is a spy)'. Rather his treatment involves quantifying over both an *object* and a *name* that (in a technical sense) represents that object to Jones, where this requires bearing the right sort of causal connection to the object. (We will set aside the Quinean approach challenged by 'Quantifying In', which relies on a purported ambiguity between notional and relational senses of 'believes', and which has—justly, in our view—fallen into disfavor.)

As Kaplan put it, an exported report like (2) seems to express a fact ‘which would interest the F.B.I.’ but ‘we would not expect the interest of that organization to be piqued by Ralph’s conviction that no two spies share a size’.<sup>2</sup> Likewise, we would not say to the shortest spy

3. Ralph believes that you are a spy.

Nor would we say, if the spy’s real name were ‘Ortcutt’, ‘Ralph believes that Orcutt is a spy’. And things do not seem much different if we consider an example where Ralph goes on to introduce ‘Sam’ as a name for the shortest spy, and then exclaims, ‘Obviously, Sam is a spy’. Even in that case, (2) and (3) seem unacceptable.

Another example illustrates the same phenomenon. Donnellan imagines that a detective at the scene of a murder introduces ‘Vladimir’ as a name for the murderer.<sup>3</sup> Presumably, even if the detective says to herself, ‘Vladimir is the murderer,’ she cannot, upon tracking down the murderer, say

4. I knew you were the murderer.

Our target is an explanation for this phenomenon that appeals to CONSTRAINT along with two other principles:

**HARMONY** Any belief report whose complement clause contains either a singular term or a variable bound from outside by an existential quantifier requires for its truth that the subject believe a singular proposition.

**SUFFICIENCY** Believing a singular proposition about an object is sufficient for having a singular thought about it.

Of course, there are anti-liberals who reject one or both of these principles; but such theorists are not in a position to make use of the argument we will be considering. For our part, we are happy to grant both principles for the sake of argument: indeed, they have a good deal of initial plausibility. Concerning HARMONY: it is natural to hold that the truth of (3) requires its complement to express a singular proposition (since it contains a singular term) and the subject to stand in the belief relation to the proposition so expressed. Likewise for (2): relative to an assignment, the open sentence ‘Ralph believes  $x$  is a spy’ relates Ralph by the belief relation to a singular content involving the assigned object. If we can regiment (2) as ‘ $\exists x$  (Ralph believes  $x$  is a spy)’, then by the standard treatment of the existential quantifier, the result is true iff the open sentence is true relative to some assignment. Meanwhile, SUFFICIENCY is simply a consequence of a common characterization of singular thought, on which it is just a matter of bearing a cognitive attitude towards a singular proposition; see §1.4.

<sup>2</sup> Op. cit., p. 192.

<sup>3</sup> Donnellan, ‘The Contingent a Priori and Rigid Designators’, p. 20.

The anti-liberal can now combine these principles with CONSTRAINT to explain why (2) and (3) are unacceptable in the cases described above. By HARMONY and SUFFICIENCY, the truth of (2) and (3) will require that Ralph believe a singular proposition, which in turn will entail that he have a singular thought. Thus (2) and (3) are false if Ralph simply has a general thought about spies. And given CONSTRAINT, Ralph will be unable to form singular thoughts about the shortest spy (even assuming there is one) in the absence of acquaintance. So, for example, neither the belief that there is a unique shortest spy nor the act of introducing a name for the shortest spy will put him in a position to have a singular thought. Thus, the unacceptability of (2) and (3) in cases where Ralph is musing about the shortest spy is happily explained; and likewise for the detective and (4).

In short, the idea is that the *semantics* of exported belief reports (and those containing singular terms in their complement clauses) suggest that they are true of a thinker only if the thinker has a certain kind of singular thought. Meanwhile, the intuitive *acceptability conditions* of these reports require the thinker to be acquainted with the object of thought. But this suggests a nice explanation for the unacceptability of the relevant reports. Thus, CONSTRAINT is vindicated by a sort of inference to the best explanation; or so the reasoning goes. We shall call this ‘the spy argument’.<sup>4</sup>

Aside from arguments of this kind for CONSTRAINT,<sup>5</sup> it is very common to encounter *appeals* to CONSTRAINT in explaining the unacceptability of the relevant attitude reports in such cases. The implication of such appeals is that CONSTRAINT provides a good explanation for this phenomenon. In what follows, we will argue that—even assuming HARMONY and SUFFICIENCY—CONSTRAINT has no explanatory value whatsoever when it comes to the acceptability of the relevant attitude reports.

<sup>4</sup> Note that an analogous argument can be run in the plural case. Suppose Ralph believes there are spies and attempts to introduce a plural demonstrative ‘them’ to plurally refer to all of the spies. Nevertheless, it would be unacceptable to use the exported belief report ‘There are some people that Ralph believes to be spies’. This would be explained, given HARMONY and SUFFICIENCY, if attempting to introduce the demonstrative term does not put Ralph in a position to have a ‘plural singular thought’ about all the spies.

<sup>5</sup> For variants of this argument, see for example Donnellan, ‘The Contingent a Priori and Rigid Designators’, pp. 23–5; Devitt, *Designation*, pp. 222–9; Soames, *Philosophical Analysis in the Twentieth Century, Vol 2: The Age of Meaning*, pp. 414–6. In his famous paper, Donnellan argues that Kaplan is right that using a name to assert or form attitudes towards singular propositions about an object requires the object to be ‘(in the right way) a part of the history of one’s use of the name’, and that this would ‘account for why the sort of stipulations we have been discussing do not put us in a position to assert and, thus, to know anything about [the object]’ (25). We take it there are two strains of argument in this paper. One emphasizes the alleged ability to acquire (contingent *a priori*) *knowledge* by way of such stipulations; we consider this argument in §2.6ff. The other simply emphasizes the unacceptability of various attitude reports: ‘Essentially the same considerations that were adduced for denying that there was knowledge of an entity just in virtue of the sort of stipulation that introduces a rigid designator by means of a description can be applied to the other propositional attitudes’ (23). He goes on to say that in the case of Newman1, it seems to him just as incorrect to ascribe to the subject *beliefs* about the individual who satisfies the description as it does to ascribe *knowledge* about that individual. That is the sort of consideration we are addressing here.

## 2.2 Acquaintance and attitude reports

Our response to the spy argument is not particularly original. As we will see, a number of thinkers have offered examples of a sort that, taken together, undermine the explanatory power of CONSTRAINT. And some have also made the point that the standard ways of handling Frege puzzles can be used to explain much of the relevant data. But given the tenacity of appeals to CONSTRAINT, it seemed worth belaboring these points a bit. In this section we argue that these considerations are adequate to respond to the spy argument. In the next section we will argue that, in fact, those who wish to retain HARMONY and SUFFICIENCY have good reason to *reject* CONSTRAINT.

Let us begin by considering a case of Stephen Schiffer's. Someone snatches Thelma's purse, but Thelma does not catch a very good look at his face.<sup>6</sup> As she is chasing him, it is acceptable to point at him and say either of the following:

5. Thelma thinks that he is a thief.
6. That man is such that Thelma believes he is a thief.

So far, so good: after all, most would allow that Thelma has singular thoughts about the man. However, suppose that later Thelma is unable to pick him out of a line-up at the police station. In this context, even if we know that the third man in the line-up is the man she was chasing, we could not point at him and comfortably utter either (5) or (6). So what governs the varying acceptability of these sentences? Clearly, Thelma has singular thoughts about the third man in the line-up, so the problem is not that she does not satisfy CONSTRAINT. (If one doubts that simply glancing at the third man in the line-up is sufficient, we may let her poke and prod the suspects as much as is necessary.) It might be argued that by this time she has lost her singular thought about the man *that he is a thief*; but this does not square with the fact that in other contexts, even at this later time, (5) and (6) are felicitous. (For example, we would happily utter them while looking at a film of the incident.<sup>7</sup>)

Of course, many theorists will approach this kind of case with whatever tool-kit they prefer to employ to handle apparent failure of inter-substitutivity of names in belief contexts. Here are two standard options:

(a) *The naïve approach.* Take the familiar 'naïve' view according to which the sentence 'Bill thinks that Jill is ill' *simply* ascribes to Bill the belief in a singular structured proposition built out of Jill and the property of being ill. On this view, the contribution of a simple proper name, even in a propositional attitude context, is merely to supply

<sup>6</sup> Schiffer, 'Naming and Knowing', pp. 65–6.

<sup>7</sup> Note that even if the line-up situation is very salient, it is fine to say 'There is someone such that Thelma is wondering whether he is in the line-up', and once one has said that it is not misleading to respond to 'And what does Thelma know about him?' by 'Well, she remembers that he stole her handbag'.

an object.<sup>8</sup> This in turn entails that co-referential names are interchangeable in propositional attitude contexts.<sup>9</sup> Meanwhile, apparent substitutivity failures are handled pragmatically by pointing out that there are various ways of apprehending the same structured proposition, and a speaker may *call attention* to one of these modes of presentation by using one name rather than another in her belief report.<sup>10</sup> While this will make no difference to the truth-value of that report, it can make a great deal of difference to what is implied or communicated. Consider Lois Lane, of whom we are happy to say ‘Lois believes that Superman can fly’ but not ‘Lois believes that Clark Kent can fly’. While these reports express the same proposition, the second can misleadingly suggest that Lois would accept the sentence ‘Clark Kent can fly’, or that she connects the ability to fly with the property of being a journalist.<sup>11</sup>

Proponents of this approach will also want to allow ways for a belief report to be misleading quite apart from misleading deployments of a proper name. Suppose someone points at Clark Kent at a time when he is working as a journalist in the office and says ‘Lois believes he can fly’. On most ways of fleshing out the scenario, the report will be heard as infelicitous. The current approach will deem the report true but misleading, though its misleading character will not be rooted in the use of a name. Rather, the setting of the report gives the hearer expectations about a way that Lois believes the proposition. Suppose for instance that for the conversational participants it is common ground that Superman can fly. We expect speakers to be attempting to communicate useful information to us by their assertions. So on the assumption that the speaker is not violating conversational norms, the hearer can assume that the speaker is trying to communicate information along the following lines: If you were to point to this person in this kind of context and ask Lois ‘Does he fly?’, Lois would say ‘Yes’.<sup>12</sup>

<sup>8</sup> Moreover, according to the naïve approach, there are no unarticulated constituents at play or context dependence in the verb ‘thinks’ that would allow neo-Fregean elements to be snuck in. According to the naïve approach, then, ‘Bill thinks Jill is ill’ will betray no context dependence that would be of interest to theorists of belief ascription. (Of course, it may inherit whatever context dependence is associated with unembedded uses of ‘is ill’.)

<sup>9</sup> But the interchangeability-in-a-context thesis does not entail the naïve view. One might hold that, at a particular context of utterance, there is a certain mode of presentation requirement on Lois imposed by ‘Lois believes Clark Kent wears glasses’, but that, holding the context fixed, the same requirement would be imposed by ‘Lois believes Superman wears glasses’. In short, one might hold that this requirement is not semantically generated by the name but by some other aspect of the context of utterance.

<sup>10</sup> The approaches described here have natural extensions to the predicative case: are ‘ketchup’ and ‘catsup’ interchangeable in belief contexts? See e.g. Salmon, ‘How to Become a Millian Heir’.

<sup>11</sup> See Soames, ‘Lost Innocence’; ‘Direct Reference, Propositional Attitudes, and Semantic Content’, §5; *Beyond Rigidity*, chs. 6–7; Salmon, *Frege’s Puzzle*. See also Barwise and Perry, *Situations and Attitudes*, pp. 253–65, though the semantic framework of their ‘innocent’ theory is different in ways that are unimportant here.

<sup>12</sup> There are, however, ways of fleshing out this scenario in which the report is *not* misleading. Suppose there is a debate rife among the population as to whether ‘Superman flies’ is true. Some people say, ‘Superman flies’. Some people say, ‘Superman is held up by strings’. Suppose further that it is common knowledge among speaker and hearer that the journalist is the superhero, and also that Lois does not know that the journalist is the superhero. In this context, ‘Lois believes he flies’, even used at a time where a bespectacled journalist is demonstrated, may naturally be heard as a report on Lois’s views on the debate as to

(b) *The notional approach.* Another standard approach, and one we prefer to naïve accounts, is exemplified by what we will call ‘notional’ theories. According to notionalists, a belief report often does more, semantically, than relate a subject to whatever complex of objects and properties is expressed by its complement clause. For instance, the ascription ‘Lois believes that Clark Kent can fly’ tells us *how* Mr. Kent is represented by her in the relevant belief. Thus in some settings we would speak falsely were we to say, ‘Lois believes that Clark Kent can fly’ or ‘Lois believes that he can fly’ even though she believes the proposition expressed by ‘Clark Kent can fly’ in relation to certain guises. On one standard version of such a view, different mental representations of Superman are actually ‘unarticulated constituents’ of what is semantically expressed by the belief reports—and so belief becomes a relation with an argument place for the subject, another for the proposition, and a third that provides information about the guise of representing.<sup>13</sup> On a second version of this view, the verb ‘believes’ (along with other attitude verbs) is itself context-dependent, expressing different binary relations in different contexts, each encoding different guise-theoretic information. On yet another version of this view, one that we shall revisit later, a that-clause expresses a semantic tree with both semantic objects and ‘labels’ at nodes, where each label plays a guise-indicating role (that is, it encodes information about how the object at that node is represented).

Now what bears emphasis is that both the naïve account and the notionalist account extend altogether naturally to the case of Thelma. She does have a singular thought about the man she chased to the effect that he is a thief; but she also has a singular thought about him, to the effect that he is the third man in the line-up; and she has not connected these two thoughts together. She has two modes of presentation of him, and the context of the police line-up makes one of these modes salient. Whether Thelma thinks of him as a thief *under that mode* appears to be what is governing the acceptability of the ascription in that context. Naïve theorists will treat this as a pragmatic phenomenon, and notional theorists will treat it as a semantic one. Either style of explanation will also have the resources to explain why many other singular-thought-requiring ascriptions to Thelma at later times seem perfectly felicitous. By contrast, an account of the infelicity data about Thelma in terms of CONSTRAINT fails to handle the felicity data.

whether Superman is held up by strings. Various people have noted this sort of case; e.g. Joseph Moore, ‘Propositions without Identity’.

<sup>13</sup> Crimmins and Perry, ‘The Prince and the Phone Booth: Reporting Puzzling Beliefs’; Crimmins, *Talk About Beliefs*. Thus, the role of ‘Superman’ in ‘Lois believes Superman can fly’ is twofold: semantically, it supplies an object, and pragmatically, it indicates the unarticulated constituent that is semantically in play. (Note, then, that this approach can agree with the previous one regarding the semantic contribution of a name in an attitude context.)



Crucially, this pattern generalizes to cases that motivate the spy argument. This point has recently been emphasized in particular by Robin Jeshion and Jonathan Sutton—concerning the naïve and notional accounts, respectively.<sup>14</sup>

(i) *The murderer*. As we have seen, even if he introduces the name ‘Vladimir’ for the murderer, the detective cannot later say:

7. From the first, I knew you were the murderer.

But this can be explained with standard resources that have nothing to do with acquaintance. To say ‘I knew you were the murderer’ would be to suggest that the detective was able to identify him in a canonical way *other* than as being the murderer—by sight, perhaps, or by name. Moreover, and crucially, proponents of CONSTRAINT can hardly complain about this explanation, because a slight shift in the example will force most of them to say the very same thing about (7). Suppose the detective had caught the murderer in the act, touselled with him, and even caught a good look at his face. And suppose that somehow the detective failed to recognize the murderer, who was in fact the Vice President. When, much later, the detective tracks the Vice President to his office and suddenly realizes that the murderer is the Vice President, it would be highly misleading for the detective to utter (7). (This case is in all relevant respects like the Thelma case.) Since most acquaintance theorists will have to admit that the detective is acquainted with the villain, they will have to explain the infelicity of (7) by appeal to the very kinds of resources that the liberal will use in the initial ‘Vladimir’ example.

(ii) *The shortest spy*. Again, there are variants of the case on which acquaintance constraints are clearly satisfied but where the felt infelicity of the singular-thought-ascribing report is equally strong. Suppose, for example, that Ralph was kidnapped, taken to an undisclosed location, and introduced to a masked man—‘Ralph, meet the shortest spy’. All but the most stringent acquaintance constraints will allow that Ralph can now have singular thoughts about that individual. However, there will clearly be plenty of contexts in which it is unacceptable, and at best highly misleading, to say, ‘There is someone Ralph believes to be the shortest spy’. After all, we may suppose that Ralph still does not have information about the shortest spy that would be of interest to the F.B.I. (He cannot even pick out the man in a line-up.) Here again the presence or absence of acquaintance is far too blunt an instrument to account for the range of intuitions about the acceptability of reports about Ralph. But both the liberal and the anti-liberal can employ the kinds of tools described earlier.

<sup>14</sup> See Jeshion, ‘Donnellan on Neptune’; and Sutton, ‘The Contingent a Priori and Implicit Knowledge’, p. 276. They raise these points largely in connection with the Leverrier–Neptune case and the contingent *a priori* (see §2.6 below); but given our taxonomy of motivations for acquaintance, these points apply with special force to the ‘spy argument’. (As we stress below, they do not by themselves rebut the idea that Leverrier should not be in a position to learn anything new from his stipulation, regardless of the truth or acceptability of particular belief reports about him.) See also Ernest Sosa, ‘Propositional Attitudes De Dicto and De Re’.

The relevant belief ascriptions about Ralph are unacceptable because they suggest that Ralph has a more useful mode of presentation of the shortest spy than he does—just as, in the context of the police line-up, it would be unacceptable to say ‘There is someone in the room that Thelma believes to be a thief’.

Admittedly, the original case involving Ralph is more dramatic because the sentence ‘There is someone Ralph thinks to be a spy’ is *almost inevitably* misleading in a case where the basis for this report is nothing more than a recognition of the fact that Ralph thinks that the shortest spy is a spy (and perhaps has named the spy ‘Sam’). Why is this? A natural explanation suggests itself. When a belief report postures as providing useful information about someone’s state of mind, we do not naturally hear it as reporting a trivial belief. So, for example, it would in nearly all contexts be highly misleading to say ‘There is someone that Jones thinks is George Harrison’ if the only basis for saying that is that Jones gives lip service to the trivial identity ‘George Harrison is George Harrison’. Relatedly, it is rare to find a context where ‘George Harrison’ counts as a serious answer to the question ‘Who does he think George Harrison is?’ (The same points apply if we replace ‘George Harrison’ with ‘the youngest of the Beatles’ throughout these examples.) The presence or absence of acquaintance is simply not relevant to this phenomenon.

Now, the claim ‘The shortest spy is a spy’ is not quite trivial. After all, on standard assumptions it makes an existential commitment that may not be satisfied. But it is quite natural to think of this commitment as having a presuppositional part—that there is such a thing as the shortest spy, as well as an assertoric part—that the relevant thing is a spy. Let us say that a claim is *incrementally* trivial when the assertoric part is satisfied whenever the presuppositional part is satisfied.<sup>15</sup> While the claim ‘The shortest spy is a spy’ is not logically trivial, it is incrementally trivial. This gives it a genuine feel of triviality and, as such, it feels like an odd and pointless thing to assert or report as a belief. Correlatively, insofar as we try to give information about someone’s state of mind it is very rare indeed to report beliefs of incrementally trivial claims. In a normal context, a hearer would hardly envisage that ‘There is someone that Ralph believes to be a spy’ is being used to communicate the presence of an incrementally trivial thought.<sup>16</sup>

In short, whether or not one accepts CONSTRAINT, one will have to appeal to these sorts of considerations for many cases of exported belief. So the appeal to CONSTRAINT for a subset of these cases is not only unmotivated, but amounts to a far less unified account of the data.

<sup>15</sup> We picked up the notion of incremental content from Stephen Yablo (in conversation).

<sup>16</sup> This is not intended as an analysis of ‘trivial belief’. Suppose someone, having decided to think about the shortest spy, thinks to herself that the person she is thinking about is a spy. Her epistemic achievement is too trivial to form a natural basis for the report that there is someone that the person thinks to be a spy. But the content of the thought is not incrementally trivial in this sense. Thanks to Elia Zardini here.

## 2.3 Turning the tables

We have seen that *CONSTRAINT* is irrelevant to the unacceptability of the belief reports that motivate the spy argument. Thus, the spy argument is a failure. In this section we will argue that, for those who accept *HARMONY* and *SUFFICIENCY*, the best explanation for many attitude reports actually requires us to reject *CONSTRAINT*. In the last section we focused on cases where certain attitude reports are unacceptable. But we will also need an explanation for the fact that, in some contexts, exported attitude reports are *acceptable* even though acquaintance is lacking.

Let us begin by illustrating this phenomenon with some examples.

(i) *Earlier examples.* In §1.6 we offered a number of examples of apparent singular thought without acquaintance. It is clear that—at least in plenty of contexts—it is acceptable to report the relevant beliefs using exported ascriptions and ascriptions containing singular terms in their content clauses. Here is a sample of such ascriptions:

8. King Henry expects to visit *that* ship (pointing).
9. There is an engine that the mechanic wants to see.
10. Yesterday, Fred thought the meeting would be held today.
11. David knows that I can see the city from here.
12. One substance is such that Mendeleev believed in 1870 that samples of it would soon be discovered.
13. There is a location on the island that John Speed believes will be windy.

The combination of *HARMONY* and *SUFFICIENCY* will require that in each case the subject has a singular thought.

(ii) *The murderer and the shortest spy.* As we saw, when the detective finally tracks down the murderer, she cannot say

7. From the first, I knew you were the murderer.

But a number of thinkers have pointed out that the acceptability of exportation in such cases is highly context-dependent.<sup>17</sup> Even in the same context, it is far more acceptable for her to say

14. From the first, I knew you were insane.

But however one explains the infelicity of (7), one is also faced with the challenge of explaining the relative felicity of (14). And this will be especially difficult for a proponent of all three principles, which together appear to entail that (14) is *false*.

Or consider Kaplan's example of Ralph and the shortest spy. Suppose Ralph pays a hit man to find and kill the shortest spy before the weekend. In that case, it is easy to imagine conversational contexts where it is acceptable to say

<sup>17</sup> Two early examples are Schiffer, 'Naming and Knowing', and Searle, 'Referential and Attributive'.

15. There is someone that Ralph wants dead.

Alternatively, suppose we are debating about whether the shortest politician, whoever that is, is likely to be a spy. Ralph thinks that the shortest politician is not only a spy but also the shortest spy. Here, ‘Someone is believed by Ralph to be the shortest spy’ is acceptable.

Such examples are easy to contrive if we note the connection between the exported description and various ‘wh’-questions. Assume it is known that there is a shortest politician. If it is acceptable in a context to answer the question ‘Who does Ralph believe to be the shortest spy?’ by saying, ‘The shortest politician’, then it will also be acceptable in that context to claim ‘There is someone Ralph believes to be the shortest spy’. (Likewise, if it is acceptable in a context to answer the question ‘Who does Ralph believe to be a spy?’ with a list that includes ‘The shortest politician’, it will be acceptable in context to report that there is someone that Ralph believes to be a spy.)<sup>18</sup>

(iii) *The lion*. Consider Quine’s original example of the apparent ambiguity in

16. Ernest wants to find a lion.

Is there a particular lion he is after, or does he want to find just any lion? It is often thought that the two standard interpretations of this sentence involve permutations of quantifier scope, though we will question this assumption in Chapter 4. Our point here is that the analogs of HARMONY and SUFFICIENCY applied to desire will require that if ‘a lion’ takes wide scope in the relevant way, Ernest must have a singular desire about a particular lion. And there are cases in which a ‘particular reading’ appears *required*, even though Ernest does not have any acquaintance with the relevant lion. Suppose Ernest has just learned that exactly one lion escaped from the zoo in Chicago, and he forms a desire to catch that very lion. A speaker aware of this fact may use (15) in an attempt to get across that there is a particular lion Ernest is after—in which case the ‘any old lion’ reading is simply not the intended reading. Or to remove even the most tenuous informational link, suppose that rather than getting his information from the newspaper, Ernest knew yesterday that some villains were going to release exactly one lion from the Chicago zoo at midnight. And he now wants to catch whatever lion they released. Here again only the ‘particular’ reading of (15) will get the speaker’s message across.<sup>19</sup>

<sup>18</sup> Parallel considerations would tell against a version of the spy argument that focuses on plural reference. It will inevitably be infelicitous to report ‘There are people that Ralph believes to be spies’ if one’s sole basis is Ralph’s assent to ‘There are spies’, but it is easy to contrive contexts where it is felicitous to report ‘There are people that Ralph believes to be evil’ when one’s only basis is Ralph’s assent to ‘There are spies and they are all evil.’

<sup>19</sup> Or suppose Ernest is a time-traveler who knows that the world will end in the distant future and wants to meet the very last lion. Here it is again easy to generate contexts where the ‘particular’ reading is natural, even if Ernest fails in his ambition and so never achieves contact with the very last lion.

(iv) 'Newman1'. We can now turn to a more difficult case. Kaplan imagines the introduction of 'Newman1' by way of the reference-fixing description 'the first child born in the twenty-second century'.<sup>20</sup> Suppose Ralph attempts this dubbing and is confident that there will be a unique first child born in the twenty-second century. He then says to himself, 'Newman1 is the first child born in the twenty-second century.' Still, it seems entirely unacceptable to say, 'There is someone Ralph believes is the first child born in the twenty-second century' or to say to that child when he or she arrives, 'Ralph believed you would be the first child born in the twenty-second century.'

This example introduces two additional complications. First, many philosophers have intuitions according to which reality consists in the present only ('presentism') or in the present and past only ('growing block theory'). On these pictures, Ralph cannot even truly think the existential proposition that there is someone who will be born in the twenty-second century. But then the data could not, even *prima facie*, motivate a constraint stronger than that which the liberal accepts. (Even the liberal can agree that there is no *x* such that *x* does not exist and we think about *x*.)<sup>21</sup>

The second complication is that there may be extra noise created by the choice of tense for the copula. Even someone who believes in the reality of the future may have difficulty with an utterance of 'There is someone Ralph believes to be ...' especially with focus placed on 'is'. (This effect is intensified if some but not all of the candidate babies have already been born.) The examples are ameliorated if we use: 'There will be someone that Ralph now believes will be ...'<sup>22</sup>

If we control for these factors we can generate contexts where the relevant exported reports are acceptable. Suppose an institute wants to incentivize higher birth-rates and plans to give away a large sum of money as a prize for a child born in the future. However, they cannot decide on the criterion to employ in deciding upon a winner. Some argue that the billionth child in the nation should win the prize. Some argue that the largest baby in the next ten years should win a prize. Now, suppose that Ralph's solution is that the first child to be born in the twenty-second century should win the prize. And suppose we want to convey, in a general fashion, that Ralph has an opinion on the matter of the prize-criterion. Assuming metaphysical eternalism, we could easily utter

17. There is someone that Ralph believes should win the prize.

or answer 'Yes' to the question, 'Is there someone that Ralph believes should win the prize?'<sup>23</sup> All that we have done in specifying this context is removed the impediments discussed in §2.2 that usually make such a claim bizarre or misleading.

<sup>20</sup> Kaplan, 'Quantifying In', pp. 201–3.

<sup>21</sup> See McGinn, 'The Mechanism of Reference', pp. 184–5, fn. 31.

<sup>22</sup> As Sosa puts it, 'It can certainly be false that there is (now) a house that I hope will please us, even when there is to be such a house' Sosa, 'Propositional Attitudes De Dicto and De Re', p. 889.

<sup>23</sup> If that judgment needs reinforcing, imagine that (17) is uttered in a context where we have a list of candidate prize-winners: 'The largest baby in the next ten years,' 'The first child in the twenty-second century', and so on.

The examples just discussed all illustrate that exported reports can be acceptable even in cases where *CONSTRAINT* is not satisfied. Can a proponent of *CONSTRAINT* appeal to the naïve or notionalist tool-kits to explain our intuitions in these cases? Let us examine how things look from the perspective of each view of attitude reports.

(a) *The notional approach.* One important benefit of notionalism over the naïve theory is that on the latter view, people often sincerely utter false belief reports in order to communicate truths. A discomfort with this phenomenon has been one of the primary motivating factors for notionalism over and against the naïve theory.

Consider some examples. The naïve theorist must admit that ‘Lois does not believe that Clark Kent is Superman’ is false; that ‘Thelma does not think *he* is the thief’ (pointing at the third man in the line-up) is false, and so on. These utterances strike us as true. As we have seen, there is a pragmatic story that the naïve theorist will tell about why these are acceptable: for example, the speakers are conveying information to the effect that Lois does not think any one person is both a journalist and a superhero.

Granted, we sometimes utter falsehoods in order to communicate truths: for example, when we exaggerate, speak sarcastically, and so on. But in such cases we typically do not believe the falsehoods that we are uttering. At times we are not being sincere, at times we are not being careful, and at times we employ devices for rhetorical effect. But many of the utterances that the naïve theorist must evaluate as false are quite different. We utter them in complete sincerity, without any of the phenomenology or intent that accompanies exaggeration, sarcasm, and so forth. It is therefore not surprising that notionalists are driven to complicate the semantics of belief attribution in part to avoid just such a result.

However, this motivation is lost if we combine notionalism with *HARMONY*, *SUFFICIENCY*, and *CONSTRAINT*. For all of (8)–(17) strike us as literally true in their respective contexts above. Moreover, they are uttered with all sincerity, and their contents are apparently believed by those who utter them. But on this combination of views, they cannot be true. According to *HARMONY* and *SUFFICIENCY* they require certain singular thoughts, and according to *CONSTRAINT*, those singular thoughts are unavailable. Notionalism does nothing to block the conjunction of the three principles from yielding falsehoods in these cases. Of course, the notionalist can in these cases appeal to the naïve theorist’s resources, but in doing so she significantly complicates her view and abandons a primary incentive of notionalism.<sup>24</sup>

<sup>24</sup> Should notionalists say that claims about whether Thelma has a singular thought about the thief are context-dependent because ‘There is someone that Thelma believes to be a thief’ is context-dependent? There are various options here (see §2.4 below). It is certainly possible for notionalists to hold that ‘x has a singular thought’ is not context-dependent. For example, if the belief relation is a three-place relation between a subject, proposition, and a guise, one might say that Thelma has a singular thought about x iff there is some singular proposition about x and some guise such that the three place relation *believing* holds between her, that proposition, and that guise. (Insofar as we accept this type of notionalism, we arrive at our notion of singular thought by existentially generalizing on the third argument place).

(b) *The naïve approach.* How do things look from the perspective of the naïve theorist? As we have seen, the naïve theorist must learn to live with the fact that we assert and even believe many false attitude reports, and with the need to explain away our propensity for false belief. (Perhaps it is just that the folk have the wrong theory of their own language.) Whatever explanation is forthcoming will have to include an account of why the various false utterances are communicatively useful despite being false. Such theorists will also want to explain why we are disinclined towards various candidate utterances that are true ('Lois believes Clark Kent can fly'), and why engaging in such utterances would be communicatively disruptive.

This is not the setting to evaluate the naïve theorist's explanatory tool-kit. We have urged that, assuming the naïve view, the data that motivated the spy argument are most elegantly explained without invoking CONSTRAINT. But the situation may be worse than that: sentences (8)–(17) sound felicitous, but also violate the conjunction of the three principles. Thus, combining those principles with the naïve theory generates an additional stock of sentences that are deemed false but felicitous. Unless the resources originally invoked to explain the false yet felicitous use of the old problem cases—such as 'Lois does not believe Clark can fly'—can be extended to handle these extra cases, there will be need for a new exercise in pragmatic apologetics, specially designed for the latter.

So, are the old tricks adequate to the new task? The standard stock of false but felicitous sentences are all *negative* attitude reports. But the new stock would be replete with false yet felicitous *positive* attitude reports, such as (8)–(17). It is not immediately clear how the naïve theorist's typical form of explanation would go in such cases. To grant that the folk believe what they are saying when they utter 'Lois does not believe Clark can fly', the naïve theorist admits in effect that the folk are tacit Fregeans.<sup>25</sup> But to grant that the folk believe what they are saying with this new stock of belief reports, the naïve theorist will have to admit that the folk are also tacit liberals about acquaintance. This additional convolution further compromises our ability to draw on ordinary linguistic intuitions about sentences when doing systematic semantics. In contrast, naïve theorists who reject CONSTRAINT can simply avoid the need to explain away this new stock of positive belief reports. Admittedly, the benefits of liberalism arising from these reporting data are not as decisive for the naïve theorist as they are for the notionalist. But it is clear that, as far as these cases go, the naïve theorist who accepts HARMONY and SUFFICIENCY has nothing to gain from CONSTRAINT, and may have plenty to lose.

<sup>25</sup> For discussion, see Braun, 'Russellianism and Explanation'.

## 2.4 HARMONY, SUFFICIENCY, and impoverished cases

Our limited offensive in the previous section is, of course, based on the assumption that our opponent holds both HARMONY and SUFFICIENCY. But acquaintance theorists who do not rely on the spy argument may of course reject one or the other of those principles. In fact, it could be argued that regardless of one's position on CONSTRAINT, some of the examples just adduced provide reason to reject at least one of them.

Here is one way to advance such an argument. As we have seen, there are contexts where it seems we can make true claims of the form 'There is someone that S believes to be F' simply because S accepts a claim of the form 'There is a unique G that is F'. For example, in the hit-man case, the escaped-lion case, and the baby-prize case, it does not make any difference to the acceptability of exported belief reports whether the subject has introduced a special name, or uses any other term typically considered to be referential. (Call versions of these examples in which it is stipulated that no such referential vehicles are in play the 'impoverished' cases.) But suppose we have a picture of singular thought on which having a singular thought about *x* requires employing something like a referential term for *x*, whether in language or in some system of internal representations. Then surely it will be tempting to deny either that these exported reports ascribe cognitive attitudes to singular propositions (i.e. HARMONY) or that such attitudes are sufficient for singular thought (i.e. SUFFICIENCY).

We think these considerations do generate pressure for the proponent of CONSTRAINT to abandon one of the other two principles. But not so for the liberal—as we will see, she has no reason to reject the natural semantics of attitude reports enshrined in HARMONY or the natural connection between singular contents and singular thought enshrined in SUFFICIENCY.

(i) *Options for the liberal.* What then should the liberal who wishes to maintain these two principles say about the argument just sketched? There are a few options:

(a) First, she can argue that impoverished cases are far more rare than one might think: when people think there is a unique F, they routinely form a singular tag in thought that refers to the satisfier of the description. Such a mental tag could be held to function as the vehicle of a singular thought even if the thinker does not give voice to it with a name or singular term in the sphere of public language. Supposing this to be true, impoverished cases are not worth worrying about.

(b) Less contentiously, the liberal can stress the *easy availability* of referential vehicles in the impoverished cases. The idea is that if Ernest has a definite description that designates the lion in question, he is only a stipulation away from having a referential vehicle. (He need not actually name the lion—he may simply raise it to salience in his imagination and then employ the mental equivalent of a demonstrative: 'I want to catch *that* lion'.)<sup>26</sup> It could then be argued that this kind of state is good enough for the

<sup>26</sup> Similarly, in the plural case one can easily introduce a plural demonstrative on the heels of a plural definite description.



acceptability of an exported belief ascription. By analogy, we frequently ascribe to others beliefs in obvious consequences of their ‘occurrent’ beliefs. Whether one treats these ascriptions as actually true or just near enough to the truth will rest on one’s general approach to the existence of ‘tacit’ beliefs. (Note that even if the liberal treats exported ascriptions in our impoverished cases as technically false, she nevertheless has a more straightforward account of their acceptability than does the acquaintance-loving naïve theorist.)

(c) A third approach denies any generalization to the effect that bearing an attitude towards a singular proposition requires publicly or privately accepting a sentence-like vehicle that makes singular reference to that object. Perhaps one can count as belief-related to a singular proposition even if one is not in a position to introduce such a vehicle. On the assumption that there can be thoughts not realized by any linguistic vehicle in the mind, there is no obvious incoherence in such a suggestion.

A variant of this view—one with which we are fairly sympathetic—allows for context-dependence in what sort of internal representation it takes to satisfy reports that ascribe cognitive attitudes towards singular propositions.<sup>27</sup> For example, the requirement that one accept sentence-like vehicles in order to count as grasping a singular proposition might hold in certain kinds of contexts, but not in others. It is also open to the proponent of such a view to allow that some contexts are *extremely* relaxed. For example, perhaps there are contexts where ‘Ralph believes that they are F’ and ‘N is one of them’ entails ‘Ralph believes N is F’. Certainly there are contexts where such transitions seem felicitous. For example, there are contexts where one of us can say ‘Phil believes I am engaged in a worthless activity’ on the basis of Phil saying ‘Philosophers are engaged in a worthless activity’.<sup>28</sup> It is worth emphasizing that this view is compatible with maintaining that both HARMONY and SUFFICIENCY are true at every context. Thus, for example, in the very relaxed context just described, one should also be willing to say ‘Singular thought is very easy to come by’. (Note that if there is contextual variation of this sort, there will be no straightforward answer to questions like ‘What does it take in general for a singular belief ascription to be true?’)

We do not wish to adjudicate between these options here. If one has a range of intuitions about the acceptability of exported attitude reports (and those involving singular terms in their complement clauses), one might want to provide different explanations for some cases than for others. For example, one might choose to appeal to option (a) or (b) for the lion case—for which our intuitions of acceptability are quite strong—and appeal to (c) in the ‘worthless activity’ case—for which our intuitions are somewhat weaker. It may be that *typically* context requires that the subject has formed some kind of referential tag in thought (or could easily do so), but in special loose contexts

<sup>27</sup> For a detailed exploration of context dependence in singular thought ascriptions, see Dorr, ‘De Re Apriori Knowledge’.

<sup>28</sup> Similar examples can be found in Sosa, ‘Propositional Attitudes De Dicto and De Re’.

that is not required. We will touch on these issues again at a later point.<sup>29</sup> Suffice it to say that liberals are well-equipped to account for the exportation data while embracing both HARMONY and SUFFICIENCY.<sup>30</sup>

(ii) *Revisiting the principles.* The situation is different for anti-liberals. Because of the lack of acquaintance in the cases at issue in §2.3, the three explanatory options open to the liberal are not available to those who would add CONSTRAINT to HARMONY and SUFFICIENCY. After all, the first two options involve granting that singular thought is typically present or easily available in such cases. And even an appeal to context-dependence for exported attitude reports will not help unless one gives up CONSTRAINT for some contexts. Of course, in response to impoverished cases, the anti-liberal could give up HARMONY or SUFFICIENCY, and with them the spy argument. However, both of these principles have their own theoretical appeal.

To reject HARMONY involves departing from the natural semantics described earlier: a belief report ascribes (at least) a belief in the proposition expressed by its complement clause, and when the complement clause is an open sentence with a free variable in it, that open sentence expresses a singular proposition under any assignment of values to the variable. Those who reject this picture need an alternative account of how the truth-conditions for exported belief reports are compositionally generated.<sup>31</sup> Are we to deny that ‘believes’ expresses the relation *believing*, or are we to deny that ‘Bill is happy’ expresses a singular content? Pending some worked-out semantics of belief ascriptions or some novel criterion for when sentences express singular contents, HARMONY is an attractive principle.

What about SUFFICIENCY? One might hold that singular thought requires a special kind of inner object representation—but believing a singular proposition does not.<sup>32</sup> One could then retain HARMONY while denying that the truth of exported attitude reports is a good guide to the presence of singular thought. Such a view, of course, faces an initial argumentative burden: there is no *prima facie* reason for thinking these special representational vehicles would be governed by acquaintance. (Again, we are assuming that the notion of singular thought at issue is not stipulated to involve acquaintance—if that were the case, CONSTRAINT would turn out to be trivial.) To begin with, it seems implausible to hold that a cognitive token of the relevant internal type could only be formed in the presence of acquaintance.

<sup>29</sup> See the Afterword, and also §4.12.

<sup>30</sup> One fallback for a proponent of CONSTRAINT is to argue that there are some special contexts where the favored version of acquaintance governs *all* exported belief reports at that context. We have encountered no reason to think that there are such contexts.

<sup>31</sup> One might try to revive the currently unfashionable view that belief ascriptions with a free variable in the that-clause express not a two-place relation to a proposition, but rather a three-place relation between a subject, an object, and a concept. This view depends on an unappetizing ambiguity thesis for ‘belief’ that is not likely to fare well by standard ambiguity tests.

<sup>32</sup> This approach might seem especially natural for those who think it is context-dependent whether one counts as believing a singular proposition—it would seem preferable for singular thought not to be shifty in this way if it is to bear significant theoretical weight.

After all—setting aside Russellian acquaintance—a thinker may be in a situation internally just like one in which she is acquainted with *x*, but due to some unlucky external circumstances the acquaintance relation does not actually hold. Perhaps the idea would be that one may form the relevant internal token and also grasp the relevant singular proposition, but acquaintance is required to grasp the proposition *with* the internal token. But why hold that, in the Vladimir case, the detective forms the right kind of internal token for the singular thought, and believes the relevant singular proposition, but still does not have a singular thought? Acquaintance theorists are left with the question of how to detect the successful use of such vehicles, in order to test the hypothesis that their use is governed by *CONSTRAINT*. But as illustrated at the end of Chapter 1, all the typical divining rods for the presence of singular thought suggest that its presence does not require acquaintance. One begins to suspect that either (i) the relevant intuitions have been schooled by theory to be sensitive to acquaintance, or (ii) the very notion of singular thought is being stipulated to involve acquaintance, in which case *CONSTRAINT* turns out to be trivial.

## 2.5 ‘Believing of’

Here is a fall-back argument involving attitude ascriptions that could be endorsed by anti-liberals who reject either *HARMONY* or *SUFFICIENCY*. They could claim that while the truth of exported attitude reports is not governed by the presence of singular thought, there is another kind of attitude report whose truth is so governed: namely, what Kenneth Taylor has called ‘fulsomenly *de re*’ locutions like ‘David believes *of John* that he is *F*’.<sup>33</sup> After all, such locutions are commonly used by philosophers and semanticists when they wish to emphasize that a singular thought is at issue. The anti-liberal might claim that reports of *this* kind always ascribe singular thoughts, and yet are intuitively acceptable only when the subject is acquainted with the alleged object of thought. She might then claim that *CONSTRAINT* best explains these facts.<sup>34</sup>

In response to this line of reasoning, we would begin by asking whether the relevant ascriptions are being used as ordinary language locutions, or as expressions of art. It is not always easy to distinguish these possibilities, especially when an expression homonymous with (if not identical to) an ordinary language expression has been put to extensive theoretical use. The problem is particularly acute in this case because ascriptions like this are rare in ordinary language. In fact, we suspect that philosophers of language will encounter them far more often in philosophical settings than in ordinary settings.

<sup>33</sup> Taylor, ‘Singular Beliefs and Their Ascriptions’

<sup>34</sup> Thanks to Nate Charlow for discussion here.

To illustrate this point, consider the fact that in the entire *Corpus of Contemporary American English* there are only twelve related uses of ‘believe(s) of’ and four related uses of ‘believe(s) that of’. In the entire *British National Corpus*, the numbers were eight and four, respectively.<sup>35</sup> Moreover, taking both sources together, in only four cases is ‘believe(s) of’ followed by a sentential clause, as in the paradigmatic ‘fulsomely *de re*’ locution—and in three of these the setting involves a statement of religious belief. In contrast, ‘believe(s)’ has more than 180,000 occurrences in the two sources combined. Simple Google searches also illustrate this point: for example, the first fifty relevant, novel search results for the exact phrase ‘believes of ★ that she’ contained 49 results from work in contemporary analytic philosophy or semantics—the vast majority of those directly addressing the issues at hand.<sup>36</sup> In contrast, there was only one result from analytic philosophy or semantics in the first 50 results for the search ‘believes that she’.<sup>37</sup>

We will not undertake a semantics of ‘fulsomely *de re*’ ascriptions here. But insofar as they genuinely belong to ordinary language, we see no reason to think their acceptability is governed by any acquaintance constraint. In fact, though some have claimed it is the purpose of these expressions to semantically ascribe singular thoughts, we very much doubt this is the case in ordinary language. In the corpus results, ‘believe(s) of NP’ is usually used without specifying the relevant belief, as in ‘be careful what you believe of NP’ and ‘what can we believe of you?’ Such expressions can be useful if one wants to make clear that certain beliefs involve a particular person, without stating what the beliefs are. Alternatively, the relevant belief has recently been specified in the discourse, and one picks it out anaphorically with a bare ‘that’, but wishes to emphasize that it involves NP, as in ‘I can’t believe that of you’. In both cases the relevant belief must be in some loose sense *about* the object(s) specified, but it is not obvious that they must be *singular* beliefs. For example, one of the few COCA results mentioned above involves a discussion of a survey finding that 1 in 5 librarians claim to have had sex between the stacks. This is followed by the comment:

18. I cannot believe that of librarians.

Other uses show up in Google searches and seem perfectly natural: ‘Actually I believe that of every Republican alive’, and so on. One might argue on theoretical grounds that these are ‘loose’ uses and that such expressions always semantically ascribe singular thoughts—that would be the analog of accepting HARMONY and SUFFICIENCY for these locutions. Still, we see absolutely no reason to think their acceptability is governed by any kind of acquaintance constraint.

<sup>35</sup> We have omitted one result of the BNC search, which was in a scholarly book by an analytic philosopher.

<sup>36</sup> We set aside irrelevant cases, like typos and ‘believe, of course, that’ and we stopped looking at 50. (Admittedly, academic websites are given high status in Google’s search algorithm, so this result is not a reason to think that nearly all English uses of “believes of ★ that she” appear in works of philosophy or semantics.)

<sup>37</sup> Again, we stopped looking at 50.

There is, we suspect, one feature of ordinary uses of these locutions that can be misleading. In the majority of cases we found, the speaker is drawing attention to an assessment of the *character* of the object of thought. Here is a typical example from our searches. Having been accused of infidelity, the speaker says

19. How could you believe that of me, Ruth?

He could easily have left off 'of me', but it serves to emphasize that it is because of her view of his character that Ruth should not have the relevant belief, rather than (for example) because he has an alibi. Or again, consider

20. I could not believe it of Jesus, that he would say to little children, 'Depart, ye accursed, into everlasting fire in hell!'

Parsimony would favor 'I could not believe that Jesus would say . . .', but the more elaborate locution emphasizes the role of the speaker's view of the character of Jesus. Of course, in many of the cases discussed in §2.3 the speaker does not have any relevant views about the character of the object of thought. And in such cases this kind of use would not be appropriate. But this is only accidentally related to the lack of acquaintance in those cases.<sup>38</sup> For example, the detective in the 'Vladimir' case may well have formed some opinion about the murderer's character. And if someone points to evidence that Smith's murderer committed some other atrocious crime with all the hallmarks of insanity, the detective might say

21. I can believe that of Smith's murderer.

Likewise, one might use 'believe of' to emphasize the relevance of one's beliefs about the character of a class of people. In fact, example (18) gets across precisely that the speaker's reluctance to believe the survey stems from her views about the character of librarians.

Another option for the anti-liberal is to acknowledge that 'believes of', as employed in this new version of the spy argument, is an expression of art. (And indeed, the way in which these locutions are introduced without argument as devices for explicitly ascribing singular thought often has the feel of stipulation. This coupled with the relative paucity of ordinary uses suggests that these expressions have a specialized theoretical use in semantics.) In that case, however, there is no bite to this version of the spy argument because there is no relevant pre-theoretical intuition for the theory of singular thought to explain. To claim that (21) is only true if the

<sup>38</sup> Another reason for this kind of locution appears to be to emphasize reflexivity. Thus, 'Everyone can truthfully believe of himself that he is a contemptible wretch' emphasizes the reflexive belief more than does 'Everyone can truthfully believe that he is a contemptible wretch'. The former also avoids any possible misinterpretation if there is another possible antecedent for 'he'. Obviously, it will turn out that in these cases the subjects are all acquainted with the objects of thought as well.

detective is acquainted with Smith's murderer is simply to register a theoretical commitment to CONSTRAINT.

In short, we find this version of the spy argument to be just as hopeless as the original.

## 2.6 The Neptune argument

A second major motivation for CONSTRAINT begins with an example from Kripke. Kripke's reflections on contingent apriority arguably raise a special puzzle. Assuming he could fix the referent in this way, it seems that Leverrier could know that (N) is a true sentence:

(N) If a unique planet is the perturber, it is Neptune

If Leverrier also grasped the proposition expressed by (N), he was not only in a position to know that it expressed a truth, but also to know that truth.<sup>39</sup> Given the semantic framework in which the puzzle is raised, the relevant truth is a singular and contingent proposition involving Neptune—knowledge that on some accounts would even count as a *priori*. But how could Leverrier gain a piece of knowledge like this merely by performing an act of linguistic stipulation? A very natural conclusion to draw is that he must *not* have been able to grasp the proposition that (N) expresses: he just was not acquainted with Neptune in the right way to have a singular thought about it. This type of argument has been offered by a number of thinkers including Donnellan, Blackburn, and Soames.<sup>40</sup>

It is important to be clear about the assumptions needed to generate this argument.<sup>41</sup> First, we need to assume that if Leverrier can introduce the name 'Neptune' by way of a definite description, he is in a position to know that (N) expresses a truth. Second, we need a premise to the effect that in the relevant cases, if 'S' means that *p*, then Leverrier knows that 'S' means that *p*. Third, we need to assume that in the relevant cases Leverrier knows that if 'S' means that *p* and 'S' is true then *p*. And fourth, we need a suitable epistemic closure principle for knowledge: in the relevant cases, propositions deduced from known propositions are also known. (We shall take this kind of closure principle for granted in what follows.) Finally, if the *reductio* is supposed to involve apriority, Leverrier's initial knowledge must be a *priori*, and that status must be retained through all of these transitions. Taken together,

<sup>39</sup> As Kripke puts it, for Leverrier 'such statements as "if such and such perturbations are caused by a planet, they are caused by Neptune" had the status of *a priori* truths'. Kripke, *Naming and Necessity*, p. 79, fn. 33.

<sup>40</sup> Donnellan, 'The Contingent *A Priori* and Rigid Designators'; Blackburn, *Spreading the Word*, pp. 333–6; Soames, *Philosophical Analysis in the Twentieth Century, Vol 2: The Age of Meaning*, pp. 403–16. In Blackburn's version, the subject knows *a posteriori* that the reference-fixing description denotes. This way of setting up the puzzle does not involve *a priori* knowledge but avoids the red herring of existential commitment, as discussed below.

<sup>41</sup> The auxiliary assumptions are not always brought out in expositions of this argument. For an exception, see Ray, 'Kripke and the Existential Complaint', especially pp. 128–30. See also the related principles in Soames, *The Age of Meaning*, pp. 403–16.

these premises will give us the counter-intuitive result, which is taken as a *reductio* either of the claim that Leverrier can successfully introduce a name with the relevant description, or of the idea that introducing that name puts him in a position to have singular thoughts about Neptune.

We want to begin by examining four replies to this argument that we think are not fully persuasive.

(i) *Naïve and notional accounts*. It may be tempting to assimilate the Neptune example to those we discussed in the previous section—so that the Neptune argument is simply an instance of the spy argument. Thus we could admit that Leverrier comes to know the proposition expressed by (N), but point out that it would be false or misleading to report this knowledge by saying

22. Leverrier knows that if a unique planet is the perturber, it is Neptune.

For, one might argue, this sentence suggests (or on the notional theory, entails) that Leverrier believes the relevant proposition under some guise associated with ‘Neptune’ that is in fact unavailable to him. Meanwhile, it seems far less infelicitous in certain contexts to say of Leverrier that he believed that Neptune had such-and-such a diameter, and this will be explained by claiming that in such contexts the guise constraint is less demanding.<sup>42</sup>

This reply is certainly helpful to a point—and it has been persuasively made by Jeshion and Sutton.<sup>43</sup> But those who focus on the Leverrier case often emphasize that Leverrier should not have ability to acquire *a new piece of knowledge* by stipulation. Nothing like that was at stake in the original spy argument. In other words, even if we can explain why (22) is invariably false or misleading, there remains the problem that a stipulation apparently placed Leverrier in a position to know a contingent planetary fact that ought to be a matter of substantial empirical inquiry. And this result, specified with whatever circumlocution, appears *prima facie* unacceptable.

(ii) *Assumption of free logic*. Here is a second unsuccessful reply to the Neptune argument. Gareth Evans and others have pointed out that in Kripke’s example, Leverrier did not know whether there was a unique perturber: so he could not know whether any referent had been fixed for ‘Neptune’. On the assumption that sentences with empty names are not truth evaluable, it follows that he was not in a position to know that (N) was true. (For all he knew, (N) lacked truth-value.) But then, Evans concludes, the puzzle does not get off the ground:

<sup>42</sup> To say, ‘Leverrier believes of some planet that it is the perturber,’ in a context where we are considering this case, will almost always be at least misleading. This is likely due to our awareness of the *a priori* connection (for Leverrier) between ‘Neptune’ and ‘the perturber’ coupled with our reluctance to ascribe thoughts that are (or are derived from) incrementally trivial beliefs. (See §2.2.ii.)

<sup>43</sup> See fn. 14. Jeshion implements the naïve strategy, and Sutton the notional strategy.

*There simply is no puzzle unless the use of free logics is accepted.* Unless a sentence concerning the name 'Julius' can be formulated which is free of existential commitment, there is not even a candidate for the status of the contingent *a priori*.<sup>44</sup>

Now Evans himself advocates free logic for names that are descriptively introduced, thereby denying a rule of Existential Generalization for proper names. (We shall return to his views on that topic in Chapter 3.) But is he right that the puzzle dissolves on the assumption that free logic is rejected?

There are two reasons that the proponent of classical logic cannot rest content with rejecting the puzzle in this way.

First, there is a closely related puzzle concerning uncontroversially *a posteriori* knowledge. Suppose that Leverrier performed all kinds of reliable calculations on the orbits of other planets that amounted to conclusive evidence for the existence of a unique perturber. Then, given our assumptions, he could know—albeit *a posteriori*—that (N) is a true sentence. Moreover, he could know the proposition that it expresses, and since he would know the antecedent, he could conclude that Neptune is the perturber. Those who feel uncomfortable ascribing Leverrier singular knowledge in the original case should also be ill at ease in ascribing singular *a posteriori* knowledge to Leverrier in this case.<sup>45</sup> The idea that stipulation can be a source of astronomical knowledge is puzzling, whether or not the resulting knowledge counts as *a priori*. (This point has also been compellingly argued by Jeshion and Sutton.<sup>46</sup>) As Kripke put it, it is hard to allow that Leverrier 'thereby *learned* some (contingent) information about the world, some new *fact* that he did not know before'.<sup>47</sup>

Second, there are other examples that overcome the issue of reference failure. Here is one due to David Cowles.<sup>48</sup> Suppose we introduce 'Alpha' by way of the reference-fixing description 'the number of planets'. There is no fear that the term will fail to refer, since if there are no planets, 'Alpha' refers to 0. We can now reconstruct the original puzzle: we appear to be in a position to know that the sentence 'Alpha is the number of planets' is true, and thus, arguably, we can know the truth it expresses. But this does not appear to be a way of knowing the proposition that nine is the number of planets.

(iii) *Empirical linguistic knowledge*. A third response to the Neptune argument involves denying that even the Alpha case is an example of the contingent *a priori*. For it would appear that one can only know that the sentence 'Alpha numbers the planets' is true by

<sup>44</sup> Evans, 'Reference and Contingency', p. 172. Emphasis in the original.

<sup>45</sup> We might reinforce this unease by imagining that his charts were inconclusive on the question whether the perturber was a planet invisible to him (as Neptune was) or a visible planet like Mars or Venus. It sounds terrible to say that Leverrier knew it was *Neptune* causing the perturbations.

<sup>46</sup> See Jeshion, 'Ways of Taking a Meter', p. 302; 'Donnellan on Neptune', p. 113, and Sutton, 'The Contingent a Priori and Implicit Knowledge', p. 257.

<sup>47</sup> Actually, he is discussing the meter-stick case in this quote: see *Naming and Necessity*, p. 63, n.26. More on that case, which involves certain complications, later.

<sup>48</sup> Cowles, 'The Contingent a Priori: An Example Free of Existential Worry', p. 140. See also McGinn, 'Review of *Studies in the Philosophy of Language*'.



remembering an *empirical* fact about the word ‘Alpha’, namely that it was introduced by way of a certain description. Since the memory of this stipulation enters into the justification of the belief that the sentence is true, any knowledge that results must be *a posteriori*.<sup>49</sup>

So goes the objection. But it is worth stressing in this connection the distinction between experiences that play a *causal* role in one’s knowledge of a proposition, and those that play a *justificatory* role. While one’s ability to grasp the proposition that no bachelor is married may derive in part from one’s learning (through experience) how to use ‘bachelor’ and ‘married’,<sup>50</sup> that sort of reliance on experience has not traditionally been taken to impugn the *a priori* status of one’s knowledge that no bachelor is married. Might it not be argued that the experience of fixing the referent of ‘Neptune’ serves in a *causal* but not a *justificatory* role for Leverrier’s knowledge? We may suppose that his reference-fixing act inculcates a disposition to believe that Neptune is the perturber, without the belief being justified in any way by his memory of the stipulation. (The belief may or may not be justified by a premise about his current dispositions.)<sup>51</sup> What this all suggests is that there are perhaps versions of the Neptune example where metalinguist facts about the relevant words or representations do not serve as premises in *a priori* arguments, but instead form the causal basis for an *a priori* known belief.

At any rate, as we have seen, the issue of apriority is not essential. There remains the problem that his act of dubbing seems to have given Leverrier an epistemic advantage. In what follows we will be concerned with this puzzle, which sets aside apriority, and the best-explanation argument it allegedly provides for CONSTRAINT. Call this the ‘epistemic advantage argument’.

(iv) *Linking requirements*. A fourth response is due to Thomas Ryckman, who offers the following variant on the original example. Suppose Leverrier calculates that a planet is causing perturbations in the orbit of Mars—and in fact the planet is Earth, though Leverrier does not realize this. Next, Leverrier introduces the term ‘Spock’ by way of

<sup>49</sup> See McGinn, ‘Review of *Studies in the Philosophy of Language*’ and Ryckman, ‘Contingency, a Priority and Acquaintance’, p. 342. Soames agrees but considers the point about a priority inessential to the motivation these examples provide for acquaintance: Soames, *Philosophical Analysis in the Twentieth Century, Vol 2: The Age of Meaning*, pp. 407–9, 414–8.

<sup>50</sup> We are not suggesting that it would have been impossible for one to grasp that proposition without learning any words; just that as a matter of fact one’s grasp of that proposition is causally dependent on one having learned those words. (And surely, left to our own devices like Mowgli, we would individually never have achieved the conceptual richness afforded by our participation in the vast and ancient practice of speech.)

<sup>51</sup> Thomas Ryckman agrees that when it comes to experiences required to *understand* sentences like ‘all bachelors are unmarried’, we often exempt them from counting against our *a priori* knowledge (Ryckman, ‘Contingency, a Priority and Acquaintance’, p. 342). But he argues that since Leverrier’s experiential knowledge of his stipulation is not knowledge of the *meaning* of the name ‘Neptune’ (on the Millian theory), it *should* count against the apriority of any knowledge for which it is required. However, we think that the distinction between causal and justificatory role best explains which experiences are exempted from undermining apriority.

the reference-fixing description ‘the unique Mars-perturbing planet’. Clearly the very same puzzles arise in this case as in the Neptune case; and yet Leverrier is very well acquainted with Earth. So, Ryckman concludes, appealing to CONSTRAINT does not help.<sup>52</sup> (A similar point can be made about the ‘Alpha’ case: surely the subject was already sufficiently acquainted with the number nine, whatever that requires.)

At least two replies are available to the proponent of CONSTRAINT. First, it could be argued that, while Leverrier has successfully introduced the name ‘Spock’, he is not in a position to *know* that he has done so. Perhaps knowing that ‘Spock’ refers requires him to know that he satisfies CONSTRAINT. Or, more plausibly, perhaps knowing that ‘Spock’ refers requires him to satisfy CONSTRAINT in a fashion that is non-accidental. But the way in which he satisfies CONSTRAINT is highly accidental. Thus, he is not in a position to know that ‘Spock is the unique Mars-perturbing planet’ is a true sentence, and the puzzle does not get off the ground.

Another reply is to modify CONSTRAINT in a way designed to fix the problem. Roughly: to understand a name introduced by a description, the subject must already be acquainted with the object in a manner somehow *linked* to that description. In a paradigm case of perceptual acquaintance, the mode of presentation under which an object is mentally represented simply *arises* from one’s acquaintance with it. Perhaps some such link between acquaintance and guise is required for the successful use of descriptive names. Scott Soames has implemented this idea as follows:

To successfully introduce a name with the reference-fixing description ‘the F’, one must have an *antecedent* singular belief involving the object denoted, to the effect that it is the F (Soames 2003: 414).<sup>53</sup>

Call this the ‘linking requirement’. It blocks Leverrier from introducing ‘Spock’ as a genuine name, because he does not independently believe of the Earth that it causes the perturbations in Mars. Similarly for the ‘Alpha’ case: the subject does not antecedently believe of nine that it numbers the planets, so the introduction of the name is blocked by the linking requirement.

There are plenty of potential complications here. To begin with, there are related cases in which the linking requirement (as stated) does not block the introduction of a name. Suppose that, unbeknownst to us, ‘Twoface’ and ‘Harvey Dent’ corefer. Suppose we assent to ‘Spiderman is the killer of Twoface’. But we still proclaim, ‘It is a mystery who killed Harvey Dent’. The linking requirement does not block us from introducing the

<sup>52</sup> See *ibid.*, 330.

<sup>53</sup> Soames is not here especially concerned with the kind of case Ryckman raises, but finds this principle compelling in its own right: ‘A description cannot semantically fix the referent of a name for a speaker unless the speaker independently believes of the object which is denoted by the description that the description applies to it’ (Soames, *Philosophical Analysis in the Twentieth Century, Vol 2: The Age of Meaning*, pp. 414). And, of course, the requisite ‘believing of’ requires ‘sufficient contact with the object’, whatever that amounts to (416).

term ‘Bill’ by way of the reference-fixing description ‘the killer of Harvey Dent’. (After all, we have known all along of Spiderman that he killed Twoface—that is to say, from the naïve theorist’s point of view, we have known of him that he killed Harvey Dent.) But this seems like the sort of introduction that proponents of the Neptune argument would not wish to allow, because of the potential for deriving new knowledge about Spiderman; to begin with, that he is denoted by ‘the killer of Harvey Dent’.

In response to these concerns, one could jettison the linking requirement and try another way to implement the idea that, in order to introduce a name by way of a description, the subject must already be acquainted with the object in a way that is linked to that description. For instance, one could semantically ascend: to introduce ‘N’ by way of ‘the F’, the subject must have an antecedent singular belief of the would-be referent of ‘N’ that it satisfies ‘the F’. (On the other hand, we suspect the counterexamples will also semantically ascend.)<sup>54</sup> We will suppose for the sake of argument that a successful ‘linking’ requirement can in the end be crafted to handle iterations of the Ryckman case.

## 2.7 The irrelevance of CONSTRAINT

We must, then, directly address the felt oddity that a stipulative act should bring knowledge of new astronomical facts. The structure of our response will be similar to our response to the spy argument: (i) we will illustrate some equally puzzling cases of apparent epistemic advance that arise even for the proponent of CONSTRAINT; (ii) we will explore a range of explanations for the various puzzling cases of epistemic advance; and (iii) we will argue that in fact the best suite of explanations is not even open to proponents of CONSTRAINT.

To begin, let us consider some cases of epistemic advance in the field of astronomy where acquaintance is not at issue—ones that ought to raise a suspicion that acquaintance is not at the heart of what is puzzling about the Neptune case.

(i) *Conjunction*. Consider a community that possesses the connectives ‘or’ and ‘not’ but does not—even in thought—possess the connective ‘and’. However, they go on to introduce ‘and’ via stipulating the standard introduction rules and elimination rules for that connective. With the connective in place, they infer the truth of ‘There are galaxies and there are black holes’ from the truth of ‘there are galaxies’ and ‘there are black holes’. On the plausible assumption that the content of a conjunction is different

<sup>54</sup> Suppose John sees a description inscribed on a page but does not discern its lexical structure. David demonstrates and says, ‘That description is true of Spiderman’. John trusts David and comes to believe of that description that it is true of Spiderman. Suppose in fact the description is ‘the killer of Harvey Dent’. Then John *does* believe, of the description ‘the killer of Harvey Dent’, that it is true of Spiderman. This yields the unwanted result that John can introduce a genuine name for the killer of Harvey Dent after all. The challenge of stating the relevant ‘linking’ requirement in a way that generates the desired result is thus rather more challenging than it may initially appear.

to that of a complex of negation and disjunction, it appears that their new stipulation has put them in a position to know a contingent astronomical fact that was previously beyond their ken.

It might be objected that in this case there is no epistemic advance in the following sense: every intension that they now accept was also an intension that they originally accepted. (By ‘intension’ we mean a function from worlds to truth-values.) And, the objection goes, there is nothing objectionable about stipulative epistemic advances that are intensionally equivalent to something already known. But this fails to get to the heart of the matter. After all, there are plenty of Leverrier-style scenarios where the advance is merely hyperintensional. Suppose Leverrier already had an expression like ‘actual’ in his language—with the semantics usually given to it by philosophers. Then the intension of

23. Neptune is the perturber.

is equivalent to the intension of

24. The actual perturber is the perturber.<sup>55</sup>

Thus, if Leverrier had ‘actual’ as a term in his language, the original Neptune example is one in which Leverrier does not make an intensional advance on what he knew before. So the proponent of the Neptune argument must tell us what makes it so bad to pull this particular astronomical rabbit out of the stipulative hat.

(ii) *Yesterday*. Next, imagine a community living on a very slowly turning planet: indeed, the whole community has evolved in a single day. Now presumably such a community might find enough in the way of causal traces about the previous day to count as acquainted with it. (We have avoided using ‘tomorrow’ because some diehard causal acquaintance lover might deny that anyone can refer to future times.) Next, suppose a community has no device for directly referring to the day before the day of a given circumstance of utterance, though they do have a device—‘Today’—for referring to the day of the circumstance of utterance. Thus, if *y* is the day before *x*, they do not have the resources to affirm singular propositions about *y*. Now suppose that on day *x* they acquire the knowledge that on the day before *x* a star exploded, and they subsequently utter ‘The day before today a star exploded.’ They then introduce a device ‘Yesterday’ for directly referring on any given day to the day before that day: that is, they stipulate the standard Kaplanian character for ‘Yesterday’. This puts them in a position to know that the sentence ‘Yesterday a star exploded’ expresses a truth, and presumably also know the truth it expresses, which is an astronomical fact they were not previously in a position to know.

<sup>55</sup> There is some wiggle room if one thinks that these claims behave differently relative to worlds where Neptune does not exist. Obviously this wiggle room disappears in the Cowles case.

(iii) *Distances*. Let us turn to Kripke's other famous example of the contingent *a priori*. Someone says: let the length of stick S be named 'one meter': can he thereby come to know the length of stick S? Kripke's discussion is short on details, and in particular he does not say whether the individual is in sight of the stick upon dubbing its length 'one meter'. Perhaps the implication is that it does not matter much whether the individual can see the stick—and we agree. But the acquaintance lover should think it makes all the difference. According to Soames, in the version of the puzzle where the subject does not see the stick, the knowledge is blocked because the subject is not acquainted with the length of stick S. And in the version of the puzzle where the subject does see the stick, there is no puzzle because the individual *already knew* the precise length of the stick by being visually acquainted with its length. The relevant piece of knowledge is *a posteriori* and is acquired by perception before the introduction of the term.<sup>56</sup>

We agree with Soames that this is not a case of the contingent *a priori*. But a puzzle remains. To illustrate, suppose Leverrier wants to know the current distance between Mars and Venus. Here are two continuations of the story: (a) He looks briefly at the planets and the distance between them, and he dubs that distance 'one schmile'. As it turns out, the distance he has dubbed a schmile is 102 million miles.<sup>57</sup> (b) Alternatively, suppose Leverrier does not have perceptual contact with the distance. Instead, he asks a friend who in fact knows the relevant distance—not only perceptually but also under the canonical guise '102 million miles'. (She has done all the necessary astronomy to achieve this knowledge.) However, she is not feeling very helpful, so she dubs that length 'one schmile', tells Leverrier about her dubbing procedure, and then adds, 'Of course, the current distance between Mars and Venus is one schmile.'

Now, a causal version of CONSTRAINT does nothing to block Leverrier's ability to know—in either of these cases—the precise distance between the planets. He will not, of course, accept the sentence 'The distance is 102 million miles'; so our theory of attitude reports may tell us that it would be misleading or false to say, 'Leverrier knows that the distance is 102 million miles'. But he does—regardless—know the relevant proposition, whether he perceptually dubbed the length, or his friend dubbed it and passed on its name.

<sup>56</sup> Here is Soames' discussion of the second case: 'In this scenario, we have seen stick *s* at *t*, and formed an idea of its length, prior to introducing the term 'one meter' with the reference-fixing description 'the length of stick *s* at *t*'. As a result, we have a perceptually justified true belief, of a certain length *l*, that it is the length of stick *s* at *t*. Call this proposition about *l* that we believe on the basis of perceptual evidence '*p*'. When we introduce the term 'one meter' with the reference-fixing description, our knowledge of *p* is part of the knowledge required to understand the sentence 'The length of stick *s* at *t* is one meter (if *s* exists at *t*)'. Since we have this knowledge, we understand the sentence. Moreover, our knowledge of the sentence guarantees not only that we know that it expresses a truth, but also that we grasp the proposition it expresses, and know it to be true... Nevertheless, this knowledge is *not a priori* knowledge, since it rests on our knowledge of *p*, which is justified perceptually. Note, it is not just that we happened to acquire our belief in *p* as a result of perception; *p* is a proposition that can only be known on the basis of such evidence' Soames, *Philosophical Analysis in the Twentieth Century, Vol 2: The Age of Meaning*, pp. 416–7.

<sup>57</sup> The example could easily have involved, say, the length of a campus seen from a distance or the area of a city seen from the air.

(iv) *Magic*. Here are two more variations on the original Leverrier case: (a) Leverrier becomes a magician and casts a spell that causes his arm to extend magically into space and his hand to come to rest on the perturber. After having prolonged tactile contact with Neptune, Leverrier says, ‘I hereby dub thee “Neptune”’. (b) Leverrier summons a magical fairy and orders her to fly to the perturber, acquaint herself with it, and call it ‘Neptune’. She returns and confirms that the dubbing has been performed, adding: ‘Neptune is the perturber’. She offers no further information.

Again, in both cases, proponents of a causal acquaintance constraint will want to allow that Leverrier is in a position to acquire the relevant piece of astronomical knowledge: that is, he can know the very proposition we express by ‘Neptune is the perturber’. But while these stories add certain causal channels absent from the original, they do very little to assuage whatever original discomfort one feels about Leverrier’s ability to achieve epistemic advance by linguistic stipulation.

## 2.8 Sources of confusion

What these cases have in common with the original Leverrier cases is that they all involve epistemic advance due to a kind of semantic good fortune that might seem like it should be irrelevant to one’s stock of knowledge. Why do such cases seem puzzling? There are at least four possible sources of confusion at work.

(i) *Knowing which*. A tempting line of thought about the original example is this. If Leverrier did not know *which planet* was the perturber, then he could not know *which planet* is referred to by ‘Neptune’ or even *which proposition* is expressed by ‘Neptune is the perturber’. But if he does not *know which* proposition that is, he can hardly *know* that proposition.

It is worth emphasizing that the same doubts will arise in the magical cases. All we need to do is maintain a cognitive setting where appropriate answers to ‘Which planet is he referring to?’ involves some way of picking out the planet that would be useful for him to write down in his astronomy book.<sup>58</sup> Likewise, in the distance case we will deny that Leverrier knows *how long* a schmile is. Clearly then, as we will stress further in the next chapter, ordinary intuitions about *knowing which* do not in any way line up with any going proposal about acquaintance. In fact, unless we are proponents of a direct revelation constraint, like Russell, we will have to accustom ourselves to the idea that a subject may have singular thoughts about *x*, even though there are contexts where we cannot say that the subject knows which object *x* is. (For example, although Lois Lane has singular thoughts about Clark Kent, there are plenty of contexts where she does not count as knowing which person he is.) We shall in the next chapter look to other and possibly more promising ways of putting *knowing which* to theoretical use in this area.

<sup>58</sup> After all, we can imagine that he does not even know whether the planet is the one usually called ‘Mars’ the one usually called ‘Venus’, and so on—supposing that his arm is insensitive to massive differences in temperature.

(ii) *Close possibilities and phenomenal duplicates.* Here is another temptation that should be resisted. It is tempting to suppose that if *S* knows that *p* then there cannot be close and phenomenally indistinguishable worlds in which *not-p*. For example, in the distance case, Leverrier is supposed to know the proposition concerning 102 million miles, that it is the current distance between Mars and Venus. But there are close possibilities—ones not ruled out by Leverrier's narrow phenomenal state—where 103 million miles is the current distance between Mars and Venus. Thus assuming the epistemic principle connecting knowledge to phenomenal indiscriminability, we are forced to conclude that Leverrier does not, after all, know the relevant singular proposition. (Neither will Leverrier know that his utterance means *p*, since there are close and phenomenally matching worlds where the utterance does not mean *p*.)<sup>59</sup>

Now it is clear that if this principle is accepted, the 'epistemic puzzle' that results has nothing to do with acquaintance and everything to do with semantic externalism. Indeed, given semantic externalism this kind of principle had better be false. Suppose Ned is on the demonstration platform, 'nude, clean shaven, and bathed in light' (to borrow an image from Kaplan).<sup>60</sup> David points at Ned and says 'He is on the platform'. There may be a close, experientially indistinguishable world in which Ned's identical twin Ted is on the platform instead. This means that in some contexts a belief report like 'David thinks that *Ned* is on the platform' will be inappropriate, because it may at least convey that David is aware there are twins and he is trying to distinguish between them. But we should not deny that David has a singular belief—and even singular knowledge—about Ned. David is not in any danger of error, because in the nearby world where Ted is on the platform instead, David has a true belief about *him*.

Admittedly, this diagnosis is not very charitable, since it suggests that the puzzlement with the Leverrier cases arises because of a derelict epistemic principle. Nevertheless, we do suspect that principles of this sort can be seductive if not brought to light, and are often tacitly in play.

(iii) *Linguistic stipulation.* A more important source of confusion is this. It is often tacitly assumed that Leverrier's alleged epistemic advance required the introduction of a proper name.

<sup>59</sup> In this connection it is worth looking at a case of Donnellan's:

I close my eyes and say (pointing), 'I will call the color of that 'Murple'. I do not know what I am pointing to, if anything. Let us suppose I am pointing to something of a definite color. Have we not set up an indefinite number of sentences, for example, 'Murple is the color of my true love's hair,' each of which expresses something true or false? But while my eyes remain shut we do not believe we know what they express (op. cit., p. 57).

It is true that, with his eyes closed, he cannot individuate these propositions in any useful way; i.e. by associating them with a phenomenal feel or with a canonical color name. But there are plenty of contexts in which he would not count as knowing which color 'Murple' refers to, even if his eyes were open! After all, he may not be able to distinguish the color from various other salient colors, the light may be a little off, he may not know its canonical name, etc.

<sup>60</sup> Kaplan, 'Dthat', p. 390.

But why? We have been given no reason to think that in order to grasp the singular proposition in question, Leverrier had to introduce a name. Leverrier could have brought the perturber to cognitive salience and referred to it by a demonstrative instead. Or, if there are such things, he could have used a privately introduced tag in Mentalese. If this is right, none of the linguistic stipulations we have been considering produced a capacity for singular thoughts about Neptune that was previously unavailable. Leverrier needed no extension of vocabulary if he had referring devices as flexible as demonstratives to hand—whether audibly or mentally employed.<sup>61</sup> He had only to think about the perturber and then demonstrate it.

Of course, as we discuss further below, there is an important sense in which knowing the relevant proposition in this way is epistemically uninteresting. But at least we can avoid the strange idea that use of a voiced linguistic item was necessary for Leverrier to have this piece of knowledge, however boring. In fact, the liberal is actually better off than the proponent of CONSTRAINT in this respect. For example, consider the case where Leverrier uses ‘Neptune’ in a way that is parasitic on the fairy’s act of dubbing. Suppose we insist that Leverrier could not grasp the relevant proposition before picking up the name. On this view Leverrier could know (presumably in a boring way) singular propositions about Neptune only *after* hearing the fairy use the name ‘Neptune’. On the liberal view, he could know those propositions in a boring way before *and* after picking up the name. If anything, it is CONSTRAINT that gives the acquisition of the name a counterintuitive epistemic significance.

Relatedly, it is open to the liberal to say that a new astronomical belief is not an interesting epistemic advance if it is *a priori* accessible in an utterly obvious way to anyone who has the original astronomical knowledge. By this criterion neither the ‘and’ introducers nor the ‘Neptune’ introducers make an interesting epistemic advance—that is, given liberalism about singular thought. But given the typical requirements for acquaintance endorsed by proponents of CONSTRAINT, the conversation with the dubbing fairy really did put Leverrier in a position to know a proposition he could not otherwise have known.

Having said all this, two qualifications are in order:

First, there may be contexts where the use of ‘Neptune’ in an attitude ascription communicates or requires that the subject actually thinks about Neptune *under the guise of that proper name*. In that case, of course, the availability of the proper name ‘Neptune’ is important to the truth of ‘Leverrier knew that *Neptune* was the perturber. And so for example, if we were considering an overseas astronomer who had introduced the name ‘Poseidon’ for the perturber, there may be contexts where it is at least misleading to say ‘That astronomer was thinking that *Neptune* was worth investigating’, even if it was acceptable to say that there was a planet that the astronomer thought was worth investigating. (To make such a context vivid,

<sup>61</sup> One could go so far as to argue that his stock of singular terms—private or public—is irrelevant, because he could think the relevant singular thought without the help of any singular vehicle whatsoever.



imagine that he has read that Leverrier is interested in something he has dubbed ‘Neptune’ but has a sufficiently low opinion of Leverrier as to doubt whether anything Leverrier dubs is worthy of investigation.)

Second, we have been proceeding as if thinking singular propositions about Neptune requires at least some referential device in thought. But if ascriptions of singular thought are context-dependent in the fashion discussed at the end of §2.4, then at some contexts far less still might be involved. To take an example, one might hold that there are contexts in which ‘belief’ expresses a relation that can hold between members of the community and various conjunctive propositions even if the community lacks a conjunction operator in their system of representations. In such relaxed settings it will be wrong to say ‘The capacity to mentally token an operator like “and” is required to know that A and B’. There may similarly be relaxed settings where ‘S knows that the actual perturber is F’ entails ‘S knows that Neptune is F’. At such contexts it will be false to say ‘The presence of a referential device for Neptune is crucial to the astronomical knowledge that Neptune is F.’

If this is right, then at contexts where referential devices are irrelevant, we can without hesitation say, ‘There is something that Leverrier knew all along to be the perturber—namely, the planet whose size and mass he had been calculating’. But at more demanding contexts we may say, ‘Leverrier did not come to know that Neptune was the perturber until he exercised his ability to refer to Neptune.’ For in such contexts, Leverrier’s mind must be suitably configured to token referential vehicles before he counts as thinking or knowing the relevant truths. Likewise, *mutatis mutandis*, for the other examples described in §2.7.

(iv) *Epistemic momentousness*. One lesson to draw from the examples of §2.7 is that even supporters of causal acquaintance will have to tell some story about why *some* token pieces of knowledge that *p* may be significant epistemic achievements, while others are not. After all, even a proponent of CONSTRAINT will have a hard time avoiding the conclusion that for various astronomical propositions *p*, Leverrier can come to know *p* in *one way* without making an interesting astronomical discovery, while knowing *p* in certain other ways *would* count as a significant astronomical discovery.<sup>62</sup> Suppose Leverrier can order his fairy to go and dub whatever object he needs dubbed. He then goes on to learn all kinds of singular propositions in this way: for example, he comes to know the proposition we express by ‘Venus is the hottest planet’ by asking the fairy to dub the hottest planet ‘Fornax’. The fairy goes off, acquaints herself with the hottest planet, and performs the dubbing, returning only with the predicable report: ‘Fornax is the hottest planet’. Clearly, this procedure does as little for genuine discovery on Leverrier’s part as the procedure of directly dubbing the planet by using a reference fixing description.

What, then, explains the difference between cases where coming to believe a singular proposition counts as a momentous achievement, and cases where it does

<sup>62</sup> This point is well made in Jeshion, ‘Ways of Taking a Meter’.

not? We will not dwell on detailed explanations here, but the answer in these cases will have nothing to do with causal connections. It is tempting to begin with the fact that Leverrier's knowledge in these cases is largely *useless* due to the observational selection effect inherent in the procedure used to introduce the name. Relatedly, Leverrier's new belief about Venus (for example) lacks interesting *connections* with his other beliefs about Venus, both singular and otherwise.<sup>63</sup> At any rate, surely anyone who claims that 'Now is now' expresses the same proposition as 'Now is Tuesday' must acknowledge at least the need for an explanation of the vast difference in cognitive payoff between the two ways of accessing that proposition. (The same goes for other pairs, like 'Hesperus is Hesperus' and 'Hesperus is Phosphorus'.) Whether or not such an explanation can easily be provided,<sup>64</sup> the need for one is not a special problem for the liberal.

## 2.9 Conditional reference fixers

Let us consider a final salvo from the acquaintance lover. We have so far been considering cases where names are introduced by unconditional reference-fixing descriptions like 'the shortest spy' or 'the perturber'. But one might argue that the case against the liberal is especially forceful in cases where the reference-fixing description itself is conditional.<sup>65</sup>

First, suppose David is very confident that the horse Archipenko won the Epsom Derby, but in fact the winner was Authorized. David says 'Let 'Dobbin' name Archipenko if he won; otherwise let it name my sick dog Rover.' Since David is convinced that Archipenko won, he now asserts

25. Dobbin won the Epsom Derby.

If we allow this kind of stipulative dubbing, 'Dobbin' refers to the sick dog. But it certainly sounds terrible in this case to say:

26. David believes that Rover/his sick dog won the Epsom Derby.<sup>66</sup>

<sup>63</sup> See Fine, *Semantic Relationism*, chs. 3–4.

<sup>64</sup> The matter is far from trivial. One wants an explanation of the proposed epistemic benefit of believing a proposition *p* under a guise associated with a sentence *S* (for example, 'Venus is the hottest planet'). But given our desiderata, the Millian explanation cannot be of the form 'Accepting *p* in this way involves believing *q*', since whatever the *q*, one might already believe *q* in an unhelpful way, even if it is a metalinguistic/metaconceptual proposition! (See fn. 54.) As in the case of Achilles and the Tortoise, adding further beliefs that are subject to the original problem will not help. This is one reason to prefer accounts that appeal to interconnections among beliefs, such as that given in Fine, *op.cit.*

<sup>65</sup> Soames uses such cases to mount a *reductio ad absurdum* in his *Philosophical Analysis in the Twentieth Century*, Vol. 2, pp. 412–3.

<sup>66</sup> Admittedly, as a matter of human psychology, it is very hard indeed to allow certain conditional reference fixers to govern one's use of a name for very long. Even supposing that 'Dobbin'—true to the stipulation—initially names Rover, in many cases semantic drift would quickly occur. (Consider, by analogy, the semantic drift that arguably occurs in baby switching cases where, unbeknownst to the parents, a newly christened baby is switched for a lookalike.) None of this shows, of course, that a genuine name cannot be introduced by a conditional reference fixer.

Does not liberalism of the sort we have been defending lead to the truth of (26)?<sup>67</sup>

We do not think such examples raise important concerns beyond those already considered. Clearly nothing about CONSTRAINT will keep David from introducing ‘Dobbin’ after casting a conditional arm-extending spell that causes his hand to be placed on Archipenko if he won, but otherwise on his sick dog. (Imagine he cannot distinguish their coats by touch.) Nor will it keep David from asking the dubbing fairy to dub Archipenko ‘Dobbin’ if he won, but otherwise dub his sick dog with that name. Neither procedure helps our intuitions about (26).<sup>68</sup> And again our misgivings about those sentences can easily be mollified if we reflect on the general mechanisms of belief reporting. Naïve theorists will say that these sentences are true but wildly misleading in nearly every context; notional theorists can insist that (26) expresses a falsehood in nearly every context. But these strategies involve no appeal to CONSTRAINT.

Alternatively, the anti-liberal might stress the unbridled epistemic advance that is apparently achievable through conditional reference fixing. Suppose David knows that either Archipenko or Authorized won the race, but has no opinion on which. So he stipulates that ‘Donkey’ will refer to Archipenko if Archipenko won, and otherwise it will refer to Authorized. Surely, one might complain, this does not put David in a position to know that Authorized won the race. But again, CONSTRAINT is irrelevant, because we get the same counterintuitive result if the dubbing results from conditional arm-extension spells or conditional instructions to dubbing fairies. (Moreover, there is nothing special about the conditional nature of the introductions, since David could simply have said, ‘Let the winner be named “Donkey”’ or ‘Let whatever I’m touching be named “Donkey”’).<sup>69</sup> So again, in defusing the concern of epistemic advance, both

<sup>67</sup> It would be ad hoc to simply ban conditional reference fixers while keeping everything else in place, allowing ‘Let “Dobbin” name the winner of the race’ as an acceptable means of introducing the name, while disallowing ‘Let “Dobbin” name Archipenko if he won or Authorized if he won’. Modulo the concerns of the previous footnote, the liberal should stand her ground or else engage in a more thoroughgoing retreat.

<sup>68</sup> Again, semantic drift of the sort mentioned in fn. 66 might eventually override the causal-based dubbing—a fact that reinforces the irrelevance of CONSTRAINT.

<sup>69</sup> If one can also introduce *predicates* via conditional reference fixers, knowledge can run more rampant still. Suppose John does not know whether ‘Shaq is tall’ is true. He introduces ‘is schmall’ by the reference fixer ‘Let “is schmall”’ express the property of being tall if Shaq is tall and the property of being not tall if Shaq is not tall. He then thinks to himself, ‘Shaq is schmall’. Assuming a naïve theory extended to predicates (something we do not ourselves recommend any more than we recommend it in the case of names), coupled with liberalism about reference fixers for predicates, we get the result that John now knows that Shaq is tall. It is worth noting that it is not entirely clear how even CONSTRAINT and the linking requirement would apply to a case like this: after all, (i) we still do not know what being acquainted with a property like tallness amounts to; and (ii) it is unclear what sort of belief one would have to have about the property of being tall to satisfy the linking requirement. (Where does the conditional element come in?) At any rate, since there are varying modes of presentation of properties, we suspect that regardless of how one interprets the linking requirement, its proponents will also face a predicative version of this puzzle.

Meanwhile, if the Millian accepts that knowing  $p$  iff  $q$  and knowing  $p$  always puts one in a position to know  $q$ , there is a slingshot argument to make knowledge run rampant (thanks to Tim Williamson here). Suppose we stipulate that ‘Jones’ is to name 1 if  $p$  and 0 if not- $p$ . Then we know that Jones = 0 iff  $p$ . Suppose in fact that  $p$  is true. Then, in fact, Jones is identical to zero. But we already know that zero is zero. So assuming the relevant Millianism, we know that Jones is identical to zero. Further, assuming the relevant closure principle,

proponents and opponents of CONSTRAINT will have to appeal to considerations like those adduced in §2.8.

In conclusion, then, we doubt that there is much prospect of giving new life to the spy argument or the epistemic advantage argument by playing up the example of conditional reference-fixers.

we are in a position to know *p*. However, this closure principle is one that the Millian will not typically accept. One may know that Hesperus is big and Phosphorus is pretty via being told 'Hesperus is big and Phosphorus is pretty', which by Millian lights is tantamount to knowing that Hesperus is big and Hesperus is pretty. And one already knows that if Hesperus is big and Hesperus is pretty, then something is big and pretty. But one is not, by standard Millian lights, in a position to deduce and come to know that something is big and pretty.