7

Anaphora: The Structure of Token Repeatables

Untwisting all the chains that tie
The hidden soul of harmony.
MILTON, L'Allegro

As fast as thou shalt wane, so fast thou grow'st
In one of thine from that which thou departest,
And that fresh blood which youngly thou bestow'st
Thou mayst call thine when thou from youth convertest
. . . . .
Let those whom nature hath not made for store,
Harsh, featureless and rude, barrenly perish
. . . . .
She carved thee for her seal, and meant thereby
Thou shouldst print more, nor let that copy die.
SHAKESPEARE, Sonnet 11

I. FREGE'S *GRUNDLAGEN* ACCOUNT OF PICKING OUT OBJECTS

1. Introduction

The first step in understanding why and in what sense claims represent or are about objects is to see what sort of conceptual content can be associated with the use of singular terms—the expressions that purport to refer to or represent objects—and correlatively with the predicates that denote the properties of objects and the relations they stand in. Only what has propositional (assertible or believable) content can serve as premise and conclusion—can both be given as a reason and have reasons given for it—and hence play a directly inferential role of the primary sort. But all sorts of conceptual content are essentially inferentially articulated. So the conceptual contents of singular terms and predicates must be understood in terms of their *in*directly inferential role—the contribution their occurrence makes to the inferential potential of sentences containing them. As Frege puts it: "We ought always to keep before our eyes a complete proposition. Only in a proposition have the words really a meaning . . . It is enough if the proposition taken as a whole has a sense; it is this that confers [*erhalten*] on its parts also their content."[1]

One of the principal technical conceptual debts the inferentialist semantic

tradition owes to Frege is the idea of using *substitution* to understand how the directly inferential articulation of what is expressed by sentences induces an indirectly inferential articulation of what is expressed by their semantically significant parts. The previous chapter investigated the substitutional species of inferential relations and showed how to understand singular terms and predicates in terms of the roles that expressions of those categories play in substitution inferences. It showed further the sort of independently characterizable expressive impoverishment to which a discursive structure is doomed unless its subsentential substitution-inferential structure takes the specific form of singular terms and predicates. Subsentential structure may be eschewed entirely, though the cost is substantial. For one then forgoes the expressive empowerment provided by the combinatory generation of novel interpretable sentences from familiar sentence-parts, which looms so large in our actual discursive practice. If semantically significant, essentially subsentential structure is discerned substitutionally, however, it can take a form other than that of singular terms and predicates only by relinquishing the full semantic explicitating expressive resources otherwise provided by sentential logical locutions, paradigmatically the conditional. This, it was claimed, is why there are singular terms (and so predicates, since the two categories come as a package). This same argument provides the ultimate reason why sententially atomic propositionally contentful claims are, or purport to be, about objects, and to represent those objects as having properties and standing in relations. The connection between singular-term usage and purported representation of objects can be filled in a bit by looking at how Frege, in the *Grundlagen,* understands the representation of objects in purely substitutional terms.[2]

2. *Objects Are Given to Us by the Use of Singular Terms*

One of Frege's concerns in the *Grundlagen* is to explain "how numbers are given to us."[3] In order to do that, he must consider the wider question of how particular objects are "given to us" cognitively. On the face of it, explaining what it is for our thought and talk to pick out or be directed at objects seems particularly difficult for the case of numbers, since, as he puts it, "we cannot have any ideas or intuitions of them."[4] Translated from the neo-Kantian idiom he is employing here, this means that the aboutness of numerical thought can be understood neither as derived from the supposedly more primitive aboutness of subjective mental pictures nor as a feature of the way in which thought about numbers is causally influenced by the numbers it is about. In fact in this context the abstractness of number is a philosophical boon because it requires Frege to address in its most general terms the question of what it is to pick out objects with our concepts—undistracted by such ultimately misleading features of some prominent special cases as the presence of mental pictures of or causal commerce with what is

thought and talked about. Instead, the role played by causal contact in the ability to pick out perceivable objects in thought and talk must be understood in terms of some more general conception of object-directedness.

Frege calls the grammatical category of expressions used to talk and think about objects "proper names": "The name of a thing is a proper name [*Eigenname*]."[5] This usage elides the distinction—of the first importance for Frege's project in the *Grundlagen*—between lexically simple singular terms, such as 'Frege', and definite descriptions formed from predicates and sortals, such as 'the author of the *Grundlagen*'. Frege's discussion focuses on the latter for two reasons: numerical expressions are formed in this way, and the definite article makes *explicit* the singular referential purport that is implicit in the use of other singular terms.

> We speak of "the number 1," where the definite article serves to class it as an object.[6]
>
> The definite article purports to refer to a definite object.[7]

The general question Frege is addressing is how expressions must be used for them to *succeed* as singular terms by referring to, picking out, or giving us a cognitive grip on definite objects—as "ways in which objects are given to us," ways of "arriving at determinate" objects, or "symbols signifying objects."[8] The issue of what it is to use an expression as a name of an object is ultimately a normative one; it is to be responded to by specifying proprieties of practice. Since the use of the definite article makes singular referential purport explicit, those proprieties can be brought out into the open by asking (in deontic scorekeeping terms) what sort of *commitment* is expressed by the use of the definite article, and what is required for *entitlement* to that commitment.

Frege insists that the issue of entitlement to singular referential purport is an important one. The use of the definite article stands in need of *justification* [*Rechtfertigung*].[9] The definite article is used in forming definite descriptions from predicates—what he calls "the definition of an object in terms of a concept under which it falls."[10] Frege is explicit about what is required for a justification of such a use of a definite article:

> If, however, we wished to use [a] concept for defining an object falling under it, it would, of course, be necessary first to show two distinct things:
>
> 1. that some object falls under this concept;
> 2. that only one object falls under it.[11]

These are the paired conditions, of existence and uniqueness, on which Russell later erected his theory of descriptions.

It may seem, that however it is with definite descriptions, explaining the object-directedness of thought (the way it puts us in touch with particular

objects our judgments are about) need involve attending only to the first of these. Showing what it means for atomic judgments to be analyzable in terms of the application of predicates would seem to suffice, for objects should then emerge as what the predicates are judged to apply *to* or be true *of*. In fact, however, understanding this 'to' and this 'of' requires mastery of the sort of practical issues of identity and individuation that are appealed to in the second condition. For in the absence of such considerations, one grasps only the use or application of whole sentences (what it is to take them to be true)—not yet what it is to apply them *to* something or take them to be true *of* something. That a judgment is directed toward an *object* is intelligible only in the context of practices of identifying objects as the same again, and individuating them as distinct.

3. *Judgments Expressing Our Recognition of an Object as the Same Again Are Substitution Licenses*

This is to say that the use of expressions as singular terms essentially involves, not only norms that could be made explicit as criteria of *application*, but also norms that could be made explicit as criteria of *identity*. Frege formulates this categorial point as the demand that "if we are to use a symbol *a* to signify [*bezeichnen*] an object, we must have a criterion for deciding in all cases whether *b* is the same as *a*, even if it is not always in our power to apply this criterion."[12] For what an expression makes cognitively available for us to "have a definite character" as an object our judgments are about, it is necessary that "it can be recognized again beyond doubt as the same, and can be distinguished from every other."[13] As indicated by the qualification "even if it is not always in our power to apply this criterion" in the previous passage, in spite of the epistemic flavor of "recognition" and "beyond doubt," the requirement is not that we in fact be able to apply the implicit criterion of identity or be infallible in our recognitions. It is just that a notion of *correctness* of identifications and discriminations must have been settled somehow. The normative status must have been instituted, even though any particular attitudes, attributions, and assessments may get it wrong.

The demand for an implicit criterion of identity associated with the use of a singular term is presupposed by the uniqueness condition on the application of definite descriptions, but it is not restricted to those singular terms in which the singular referential purport is marked overtly by the use of a definite article. In its absence, no sense could be made of the notion that terms (including those that are not definite descriptions) implicitly involve a specifically *singular* referential purport. An implicit criterion of identity provides the "authority to pick out [particulars] as self-subsistent objects that can be recognized as the same again [*selbständige, wiedererkennbare Gegenstände zu unterscheiden*]."[14] What does it mean for such authority or enti-

tlement (which could be made explicit in the form of a criterion of identity) to be in place? The key fact is that "objects too can change their properties without that preventing us from recognizing them as the same [*sie als dieselben anzuerkennen*]."[15] Recognizing an object as the same again is making a certain kind of judgment, what Frege calls a "recognition judgment." Thus "For every object there is one type of proposition which must have a sense, namely the recognition-statement."[16]

Indeed, the use Frege makes of the concept of a recognition judgment shows that he is committed to a much stronger claim. Not only is fixing the sense of recognition judgments *necessary* for entitlement to use an expression as singular term, it is *sufficient.* And once an expression has qualified as entitled to its singular referential purport, it is a way in which a determinate object can be picked out or given to us.[17]

> How, then, are numbers to be given to us, if we cannot have any ideas or intuitions of them? Since it is only in the context of a proposition that words have any meaning, our problem becomes this: To define the sense of a proposition in which a number word appears. That, obviously, leaves us still a very wide choice. But we have already settled that number words are to be understood as standing for self-subsistent objects. And that is enough to give us a class of propositions which must have a sense, namely those which express our recognition of a number as the same again . . .
>
> In doing this, we shall be giving a general criterion for the identity of numbers. When we have thus acquired a means of arriving at a determinate number and of recognizing it again as the same, we can assign it a number word as its proper name.[18]

That an expression is used as a singular term, and so has singular referential purport—that it is a way in which determinate objects can be made available to judgment ("arrived at," "given to us")—is a significance that performances can be accorded in the context of practices of keeping deontic score on special sorts of commitment and entitlement. It emerges from the passages quoted above that the central technical concept Frege employs to explain the commitments and entitlements that define singular term usage is that of *fixing the sense of a recognition claim.* The rest of this section is devoted to exploring how Frege uses this concept to elaborate his understanding of what it is to talk or think about particular objects.

Securing singular reference is for Frege "a matter of fixing the content of a recognition-judgment [*Wiedererkennungsurtheils*]."[19] Recognition judgments have the form of *identity* claims. Identity claims express recognition of an object as "the same again" when given or referred to in two different ways. To establish reference to a particular object by a given expression, one must settle what would make true or false various identities in which that expression occurs (even if one is not in a position to tell of each such identity

whether it is in fact true or false). Since the singular referential purport of terms amounts to claiming that recognition judgments involving those terms have a definite sense,[20] taking it that the significance of asserting an identity involving a term has been settled is treating the term as referring to or picking out an object. That is why "to use the symbol '=' is likewise to designate [something] an object."[21]

Frege's problem was set by the fact that the absence of causal contact with and mental images of numbers made the possibility of picking them out as objects of thought and knowledge seem particularly mysterious. Reconceiving the problem of securing singular reference in terms of recognition judgments yields the result that "to obtain the concept of Number, one must fix the sense of a numerical identity."[22] The general account of what it is to talk and think about particular objects accordingly shows how our cognitive and conceptual grasp on numbers can be made intelligible in terms of our capacity to take or treat sentences involving numerical terms as expressing identity claims.

> Our aim is to construct the content of a judgment [*den Inhalt eines Urtheils zu bilden*] which can be taken as [*auffassen lässt*] an identity such that each side of it is a number.[23]
>
> In the same way with the definitions of fractions, complex numbers and the rest, everything will in the end come down to the search for a judgment-content [*beurtheilbaren Inhalt*] which can be transformed [*verwandelt*] into an identity whose sides precisely are the new numbers. In other words, what we must do is fix the sense of a recognition-judgment for the case of these numbers.[24]

So Frege's explanatory strategy begins with the idea that particular objects are to be distinguished as what can be recognized as the same again—in the sense that the norms governing the use of terms referring to them would be made explicit by associating with them not only criteria of application but also criteria of identity. The recognition judgments that express the applicability of such norms are thus to be construed as identity claims. To carry this strategy through to completion, Frege must address two further issues, one quite general and the other specific to the case of numbers (as abstract objects). The general question is what it is to "fix the sense of an identity": How must a sentence be used, what sort of significance must it be accorded, in order to confer the content of an identity claim? The question specific to numbers is then what is required to confer such content on claims involving numerical expressions.

Frege's answer to the first question is straightforward, and just as it should be from the point of view of the discussion of using expressions as singular terms in Chapter 6. Identity claims make explicit *substitution* licenses. "In universal substitutability [*allgemeinen Ersetzbarkeit*] all the laws of identity

are contained."[25] Since identity claims are the form of recognition judgments, recognizing an object as the same again is itself to be understood in terms of substitutional commitments. "When are we entitled to regard a content as that of a recognition-judgment? For this a certain condition has to be satisfied, namely that it must be possible in every judgment to substitute without loss of truth the right-hand side of our putative identity for its left-hand side."[26] The consequences of application distinctive of identity claims consist in the undertaking of substitution-inferential commitments. What is made assertionally explicit as a claim of the form $a = b$ is commitment to a pattern of inferences requiring doxastic (assertional) commitment to the claim expressed by *Pa* whenever one undertakes doxastic commitment to the claim expressed by *Pb*, and vice versa. Frege understands particular objects as what we get cognitive and conceptual access to by using expressions as singular terms, and he offers a substitutional construal of what it is to use expressions as singular terms. The proprieties governing the circumstances and consequences of their application are those codified explicitly in identity claims, which have the significance of symmetrical substitution licenses.

4. The Maximal Substitutional Requirement on Using an Expression as a Singular Term

There is more to introducing a new term by "fixing the sense of an identity" involving it, however, than just understanding what one is committing oneself to in asserting such an identity. Ordinary cases of term introduction are special in a way that tends to obscure what more (beyond a general understanding of identity) is required to fix the sense of identities in which the new expression occurs. The sort of example that best highlights what else Frege takes to be needed is that of introducing not only new terms but new objects. The lesson appears most clearly from consideration of the role played by *abstraction* in his account of how we can become entitled to use numerals as names of definite objects.

The key point is that to be entitled to introduce a new term as the name of an object, one must settle when it would be *correct* to recognize the object picked out as the same again; in this way one distinguishes it from all other objects.[27] Frege officially insists that to do this one must see to it that the truth or falsity of *all* identities involving it has been settled. Doing so is settling when it would be correct to recognize the object picked out as the same again, and thus distinguishing it from all other objects. When what settles the truth-values of these identity claims involving a term is made explicit, it takes the form of a criterion of identity.[28]

In run-of-the-mill cases of term introduction, this requirement is quite easy to satisfy. For in the central cases a new singular term is being introduced to refer to an object that can already be referred to by using other terms

already available in the language. When a proper name is introduced for a person, place, or perceivable thing, there are typically already-individuating sortals in place appropriate to it, and it can be picked out by definite descriptions using those sortals, combined with specifications distinguishing it from others of its kind (for instance, spatiotemporal ones): 'the person who just came out of the front door of Jay's Bookstall', 'the Northwest corner of the intersection of Forbes and Meyran avenues', 'the black telephone in that corner', and so on. Assuming that the use of these antecedently available terms is already in order (as far as Frege's official requirement is concerned), all that is required to introduce a new term a is commitment to a reference-fixing identity. For under these circumstances, if the term '(the x)[Dx]' is already in use in the language, then by hypothesis the truth-values of all identities of the form (the x)[Dx] = t (where t is another term already in use in the language) have been settled. The introducing stipulation that a = (the x)[Dx] then automatically settles the truth-values of all the identities involving a and antecedent vocabulary: $a = t$ just in case t = (the x)[Dx], and not otherwise. In these cases, then, committing oneself to the truth of a single-identity claim linking the novel term to a familiar one serves to fix the sense of *all* the identities involving the novel term, for it settles all their truth-values.

Clearly this technique is not available for introducing new terms for *new* objects—ones that cannot already be referred to in the antecedent vocabulary.[29] The problem of introducing numerical expressions referring to numbers, Frege says, is the problem of fixing the sense of numerical identities. He does this by the method of *abstraction:* a particular way of explaining the use of novel terms (referring to novel objects) by means of the use of familiar terms (referring to familiar objects). The idea is this: Where $\ulcorner a \urcorner$ and $\ulcorner b \urcorner$ are terms whose use is already established, new terms of the form $\ulcorner fa \urcorner$ and $\ulcorner fb \urcorner$ can be introduced wherever there is an equivalence relation R available defined on the old vocabulary. For one can then define the sense of identities involving f terms by stipulating that

$$fa = fb \quad \text{iff } Rab.$$

In this way, if $\ulcorner a \urcorner$ and $\ulcorner b \urcorner$ are terms designating lines, one can introduce new terms of the form $\ulcorner$direction of $a \urcorner$ and $\ulcorner$direction of $b \urcorner$ (and hence the new sortal or object-kind *directions*) by appeal to the equivalence relation . . . *is parallel to*___ defined on lines:

$$\text{the direction of } a = \text{the direction of } b \quad \text{iff } a \text{ is parallel to } b.[30]$$

In just the same way, if $\ulcorner a \urcorner$ and $\ulcorner b \urcorner$ are terms designating collections of already-available objects, one can introduce new terms of the form $\ulcorner$number of $a \urcorner$ and $\ulcorner$number of $b \urcorner$ (and hence the new sortal or object-kind *numbers*)

by appeal to the equivalence relation . . . *can be put in one-to-one correspondence with*___ defined on collections of objects:

> the number of a = the number of b
> iff a can be put in one-to-one correspondence with b.[31]

The claim that the relevant equivalence relation obtains between the familiar objects accordingly serves as the content of a judgment that can be taken as or transformed into an identity relating numerical (or direction) expressions, as Frege requires in the two passages quoted above.[32] That the judgment *Rab* can be reconstrued as an assertion of identity involving terms referring to novel abstract objects—rather than just as asserting a relation between familiar concrete (relative to this construction) ones—depends just on R being an equivalence relation; to be entitled to the reconstrual of such claims as putting us in cognitive and conceptual touch with abstract objects is just to be entitled to characterize R as reflexive, symmetric, and transitive. For since "in universal substitutability all the laws of identity are contained," it follows that

> in order to justify our proposed definition of the direction of a line, we should have to show that it is possible, if line a is parallel to line b, to substitute
>
> 'the direction of b'
>
> everywhere for
>
> 'the direction of a'.
>
> This task is made simpler by the fact that we are being taken to know of nothing that can be asserted about the direction of a line except the one thing, that it coincides with the direction of some other line. We should thus have to show only that substitution was possible in an identity of this type, or in judgment-contents containing such identities as constituent elements. The meaning of any other type of assertion about directions would have first of all to be defined, and in defining it we can make it a rule always to see that it must remain possible to substitute for the direction of any line the direction of any line parallel to it.[33]

Showing that the relation R on which the abstraction is based is an equivalence relation entitles one to regard *Rab* as an identity relating the new expressions *fa* and *fb* (circumstances of application). Regarding it that way is undertaking a substitutional commitment to the propriety of the inference from $P(fa)$ to $P(fb)$, and vice versa, for any sentential context in which one discerns a primary occurrence of the new terms (consequences of application). The doctrine of abstraction Frege puts forward here is the claim that the significance of attributing this constellation of entitlement and commit-

ment is as taking the subject of those deontic statuses to be in a position to make judgments (to think and talk) about a new range of abstract objects—which may be thought of as equivalence classes of the old ones. This is how objects, paradigmatically mathematical ones, which we do not have causal commerce with (and need not be able to form mental images of) can be "given to us."[34]

As indicated above, Frege's official view is that to introduce a new term one must settle the truth-values of *all* identities relating it to other terms. This requirement leads to disastrous results in the later *Grundgesetze,* and Frege never does find an acceptable way to satisfy it for the introduction of terms referring to abstract objects.[35] In any case, the sort of abstractive definition just considered "fixes the sense of numerical identities" only in the sense of settling the truth-values of identities, both sides of which are numerical expressions—in the general case identities of the form $fa = fb$, but not of the form $fa = c$, where $\ulcorner c \urcorner$ is a bit of antecedent vocabulary, a term referring to an object that is concrete relative to the abstractive method of term-and-object introduction. The significance of the failure of abstractive definitions to meet the very strong condition Frege puts on term introduction—what one must do or show in order to be entitled to use an expression as a singular term—depends not only on whether it is possible to satisfy that condition in some other way but also on the reasons there are for endorsing that condition.

Frege's basic insight is that the essential singular referential purport involved in singular-term usage consists in the role such terms play in identity claims. Since he further analyzes what is expressed by identity claims in terms of the significance of such claims as intersubstitution licenses, this amounts to taking singular referential purport to consist in a structure of symmetric substitutional commitments. It is in terms of the undertaking and attributing of such substitutional commitments that the scorekeeping significance of using a singular term to express a claim is to be understood. Frege takes it that the strong condition he imposes on successful term introduction is a consequence of this substitutional analysis of what it is to use an expression as a singular term. For he takes it that unless the truth-values of *all* identities involving the candidate term have been settled, it has not been settled what one would be committing oneself to by employing it to make claims (for the identities merely make substitutional commitments explicit, that is, assertible). Abstractive definitions settle whether in using one of the new terms to make a claim of the form $P(fa)$ one is thereby committing oneself also to $P(fb)$, but they do not settle for arbitrary c whether one is committing oneself to $P(c)$.

Appealing to symmetric substitutional commitments (a species of inferential commitment) to explain what it is to use an expression as a singular term—the fundamental Fregean insight that is developed in detail in the previous chapter—does not necessitate the maximalist reading of what is

required for successful introduction of terms that Frege thinks follows from it. Frege thinks that there is something wrong with using an expression where it has been settled (whether or not anyone in particular is in a position to tell—a matter of status rather than of attitudes) that in endorsing a sentence in which it appears, one is thereby committed to the claims expressed by *some* substitutional variants of that sentence (and precluded from entitlement to commitments to the claims expressed by various other substitutional variants of that sentence), if there are some *other* substitutional variants on which one is *not* thereby counted as taking up a stance.[36] But what is wrong with its being settled that when I claim that the largest number that is not the sum of the squares of distinct primes is odd, I am thereby in some sense committing myself (whether I know it or not) to the claim that 17,163 is odd, am making a claim incompatible with the claim that 17,163 is even, and am not taking a stand on the question of whether England or the direction of the earth's axis is odd? Why would not such a situation count as one in which it had been settled exactly what I am and am not committing myself to (and similarly for entitlements), and so one in which a perfectly *definite* sense is associated with the numerical expressions involved, even though that sense is not *complete* in the way that Frege wants to require?

For many purposes it may be appropriate to insist on Frege's strong condition that the truth-values of all identities be settled; these may even include the purposes that motivate the development of the formal language Frege uses in the *Grundgesetze*. The issue is not even one of whether, relative to these purposes, a language in which this condition is imposed is *better* than one in which it is not. The question of interest at this point is rather whether there is some way of relaxing Frege's condition while maintaining the features of singular term use that make it appropriate to think of them as purportedly (and in favored cases successfully) picking out particular objects. Furthermore, it would be of interest to know just how weak the condition on the symmetric substitutional commitments associated with an expression could be made without endangering its singular referential purport. What is the *minimal* substitutional requirement (or necessary condition) on using an expression to pick out an object in thought, corresponding to the *maximal* substitutional requirement (or sufficient condition) that Frege imposes?

5. *The Minimal Substitutional Requirement on Using an Expression as a Singular Term*

A good place to begin in addressing this question is to notice that even according to Frege, to fix the sense of a novel term (for instance a numerical expression) it is *not* in fact sufficient merely to settle the truth-values of all the identities it can occur in. For it would not suffice for term

introduction to settle the truth-value of all the nontrivial identities—all those that relate the term to some *other* term—as *false*. No criterion of identity is implicitly associated with the expression ⌜*a*⌝ by stipulating that *a* = *a*, but that if ⌜*b*⌝ is any expression distinct from ⌜*a*⌝, then *a* = *b* is false. Settling the truth-values of all the identities involving the new expression in this way does not even implicitly involve associating with it an object that can be recognized as the same again. "All identities would then amount simply to this, that whatever is given to us in the same way is to be reckoned as the same. This, however, is a principle so obvious and sterile [*unfruchtbar*] as not to be worth stating. We could not, in fact, draw from it any conclusion which was not the same as one of our premises. Why is it after all that we are able to make use of identities with such significant results in such diverse fields? Surely it is rather because we are able to recognize something as the same again even though it is given in a different way."[37]

Objects are essentially things that can be recognized as the same again, even though given in different ways. That is why they are things for which the issue of identity arises—why using the identity sign with an expression is treating it as referring to an object.[38] To be an object is to be something that can be referred to in different ways; to associate an object with an expression as its referent requires settling what would count as *another* way of picking out that same object. Frege's maximalist claim is that introducing a term as picking out a definite object requires settling *every* other way of picking out that same object. The corresponding minimalist claim is that it requires settling at least *some* other way of picking out that same object.

The thought can be put more clearly by shifting from material mode to formal mode: from talk of objects to talk of the substitutional significance of singular terms by means of which talk of objects is officially to be understood. The basic idea is that unless the occurrence of a candidate term in the expression of a claim has *some* substitution-inferential significance (unless it commits one to *some* further claim that is expressed by a sentence resulting from the first by substitution of another term for the candidate), then the candidate is not functioning as a singular term at all. Its occurrence is not semantically significant in the way terms are; it is substitutionally idle, thus inferentially idle, and therefore semantically idle to discern its occurrence at all.[39] The minimal condition on using an expression as a singular term that emerges from understanding the characteristic substitutional role terms play is just that it must have been settled that the occurrence of the putative term have *some* (symmetric) substitution-inferential significance. As elaborated in the previous chapter, for the occurrence of an expression to have a significance of the kind characteristic of singular terms, its use must be governed by some simple material substitution-inferential commitments (SMSICs)—commitments of the sort that can be expressed explicitly as nontrivial identity claims or recognition statements. Where Frege demands a *complete* set of substitutional commitments associated with each term, the

minimal demand compatible with a substitutional understanding of singular terms (motivated by the observation that Frege would not permit all the nontrivial identities to be settled as false) is that a *nonempty* set of substitutional commitments be associated with each term; at least one nontrivial identity must be settled as true.[40]

Talk of objects as what can in principle be recognized as the same again—what can be given to us or referred to in different ways—reflects the structure of substitutional significance that the occurrence of bits of subsentential vocabulary must have for them properly to be understood as having the indirectly inferential content characteristic of the use of singular terms. The singular referential purport of such vocabulary consists in the fact that the deontic scorekeeping significance of its use is to be determined by symmetric substitutional commitments that link it to other vocabulary. These are the commitments that are made explicit by the nontrivial identity claims that Frege calls "recognition judgments"—which he takes to express the recognition of an object as the same again, though given in two different ways. This much of Frege's thought in the *Grundlagen* can be taken over without a consequent commitment to the requirement that the truth-values of *all* nontrivial identities must be settled in order for a singular term to have been properly introduced.

Even the minimal claim that settling the truth of some nontrivial identities involving a candidate singular term is a necessary condition for using it as a name of an object, however, has consequences that can seem mysterious unless the substitutional gloss on that claim is kept firmly in mind. For it follows that the idea of an object that can be picked out or referred to only in one way is not an idea of an *object* at all. (Recall the discussion above in 6.4.) A language cannot refer to an object in one way unless it can refer to it in two different ways. This constraint will seem paradoxical if referring to an object by using a singular term is thoughtlessly assimilated to such activities as using a car to reach the airport or using an arrow to shoot a deer: even if only one car or one arrow is available and impossible to reuse, what one is doing can still genuinely be driving to the airport or shooting the deer. Why should referring be different, something that cannot be done one way unless it can be done two ways? Understanding an expression's purporting to refer to an object in terms of its use being governed by proprieties articulating its significance according to substitution-inferential commitments dispels the puzzlement that can otherwise attend this phenomenon. An object that can be referred to in only one way is the sound of one hand clapping.

So for an expression to be used as a singular term, there must be *some* substantive substitutional commitment undertaken by the one who uses it. It is not necessary that either the one who undertakes that commitment or the one who attributes it—by attributing a doxastic commitment that would be avowed by the assertion of a sentence containing the singular term—be able to specify just what the content of that commitment is. But it is only

where the interpreter takes it that there *is* some such substitutional commitment included in the significance of the underlying doxastic commitment that the one who undertakes that commitment is interpreted as using a singular term to make a claim about an object. This is just the conclusion that was drawn in Chapter 6: the category of singular terms should be understood as comprising expressions whose proper use is governed by simple material substitution-inferential commitments (SMSICs) linking them to other such expressions. Taking an expression to be a singular term—taking it to purport to pick out a particular object—just is taking its use to be governed by some such SMSICs. When such a simple material substitution-inferential commitment linking two expressions is made propositionally explicit (as an assertible), it takes the form of a nontrivial identity claim. That is why to introduce an expression as a singular term, one must somehow settle the truth-value of at least one such identity (of what can be so expressed in an idiom with suitable explicitating resources—that is, logical vocabulary). Purported reference to objects must be understood in terms of substitutional commitments linking diverse expressions.

6. *Substitutional Triangulation*

This *substitutional holism*—according to which mastery of the use of one expression as a singular term involves mastery of the use of many—is the reflection at the subsentential level of the *inferential holism* according to which mastery of the use of one expression as a sentence (even one that can be used to make noninferential reports) involves mastery of the use of many.[41] Carving up sentences according to their substitutional relations to one another is just a method for extending the notion of content-conferring, inferentially articulated deontic significance to the subsentential level—to expressions that cannot themselves play the directly inferential roles of premises or conclusion of inferences. The conceptual content expressed by a sentence depends on its place in a network of inferences relating it to other sentences; the conceptual content expressed by a singular term depends on its place in a network of substitutions relating it to other terms. The substitutional roles that determine the pragmatic significance of the occurrence of singular terms are a kind of indirectly inferential role because substitutional commitments are a kind of inferential commitment.

Another topic this minimal substitutional requirement for using an expression as a singular term illuminates concerns picking out objects by conceptual *triangulation*. Triangulation strategies arise from consideration of a fundamental problem concerning the discrimination of a particular *stimulus* to which some sort of response is reliably keyed. In his discussion, Davidson introduces the familiar point this way: "Why say the stimulus is the ringing of the bell? Why not the motion of the air close to the ears of the dog—or even the stimulation of its nerve endings? Certainly if the air were

made to vibrate in just the way the bell makes it vibrate it would make no difference to the behavior of the dog. And if the right nerve endings were activated in the right way, there still would be no difference."[42] Typically, there is a whole causal chain of covarying events culminating in a response. In the standard case, the occurrence of one is accompanied by the occurrence of all the rest. Under these circumstances, the response being keyed to one of the event kinds is its being keyed to all the rest. How is one element of the chain to be singled out as the stimulus? What is the nature and source of the privilege that distinguishes one element from another?

One strategy for assigning such privilege, and therefore picking out as the stimulus one element from the whole chain of covarying event types that culminates in a response of the specified type, is (as Davidson goes on to suggest) to look to *proximity* to the eventual response. The justification for seizing on causal proximity of stimulating event to the response as what matters is maximizing the relative *reliability* of the connection between the occurrence of events of the distinguished stimulus type and the occurrence of events of the distinguished response type. The proximal element of the chain is the one that most reliably brings about the response. For prior occurrences in the chain elicit the response only in the cases where they succeed in bringing about an event of the proximal type, while events of that type can elicit the response regardless of whether they have themselves been brought about in the standard way. The trouble is that such a proximal theory of stimuli will always yield the result that the stimuli being responded to are at the sensory surfaces or within the nervous system of the responding organism.

In the context of the project of using reliable differential responsive dispositions as a model to understand which objects basic empirical concepts are being applied to, the adoption of such a policy for the discrimination of stimuli is disastrous. For what is classified by the protoconcepts that repeatable responses are going proxy for is not bells and tables and rabbits but only states of the responding organism. Nothing that looks like one of our ordinary empirical concepts, applying to ordinary observable objects, is within reach of such an approach. A distal strategy is required in order to get the protoconcepts represented by reliably differentially elicited noninferential response types to count as classifying and so applying to ordinary observable objects and properties. Understanding them this way involves respecting the language-learning situation in which these reliable differential responsive dispositions are established.

The most popular approach to identifying distal stimuli as what is classified by the exercise of reliable differential responsive dispositions is to appeal to *triangulation*. This is a strategy for picking out or privileging one bit of the causal chain of covarying event types that reliably culminates in a response of a distinguished type, by looking at the *intersection* of *two* such chains. The insight it develops is that the best way to pick a single *point* (the

stimulus) out of a *line* (the causal chain of covarying event-types that reliably elicit a response of the relevant type) is to *intersect* it with another line—another causal chain corresponding to another reliable differential responsive disposition.

One writer who employs such a triangulation strategy to address the problem of picking out distal stimuli as what a response is *about* is Dretske.[43] In order to pick out the distal stimulus he looks to the upstream intersection of two distinct "flows of information" (or causal chains of reliably covarying event-types) that reliably culminate in responses of the same type. A simple example of the sort of system he has in mind would be a thermostat that keeps the temperature of a room within a certain range by turning a furnace on and off. If the thermostat has only one way of measuring temperature—for instance by the bending of a bimetallic strip until it touches either the left electrical contact (too cold) or the right one (too warm)—there is no way, Dretske acknowledges, to say that what the system is responding to is the temperature of the room, rather than the temperature of the bimetallic strip or the curvature of the bimetallic strip or the closing of the circuit between the bimetallic strip and one or the other of the contacts. Notice that a pragmatist appeal to practical *consequences* of the response in question is of no help here; turning the furnace on affects not only the temperature of the room but also that of the bimetallic strip, its curvature, and so its relation to the electrical contacts.

The idea is that one can be entitled to such a description if the thermostat is slightly more complicated and has another causal route to the same response (turning the furnace on or off). If the thermostat has a second sensor—for instance a column of mercury supporting a float with an electrical contact that completes one circuit to turn the furnace on whenever the float is below one point (too cold) and turns it off whenever the float is above another point (too warm)—then the system has two ways of responding to the change in temperature in the room. Although for this second route by itself (just as for the first by itself) there is no feature of the system that entitles one to say it is responding to changes in the temperature of the room rather than to the temperature of the mercury or the length of the mercury column or the closing of the switches, when the two routes are considered together, they intersect in just two places—upstream at the change of temperature in a room (which is included in the "flow" or causal chain corresponding to each route) and downstream in the response of turning the furnace on or off.[44] Dretske shows how the general strategy of looking to the *intersection* of two reliable differential responsive dispositions might be funded from the resources of the responding system itself.

One might worry that Dretske has not in fact succeeded in responding to the general worry about how to justify describing the system as responding to a distal stimulus rather than a proximal one. For there is an objection available to his strategy that seems to reinstate the original worry. Why, it

might be asked, ought we not to conclude that even in the two subsystem case, what is responded to is a proximal stimulus, but a *disjunctive* one? The system turns on the furnace just in case *either* the temperature of the bimetallic strip is low enough *or* that of the mercury column is low enough, or alternatively, in case the curvature of the bimetallic strip is far enough to the left *or* the mercury column is short enough. (Again, pragmatic appeal to the practical consequences of entering this state will not solve the problem.)

This worry is connected to the complaint voiced already in Chapter 2, to the effect that mere differential responsiveness is not sufficient for identifying the responses in question as applications of *concepts*. The rationalist supplementation suggested there—that what is distinctive of the conceptual is the *inferential* role played by the responses that stimuli differentially elicit—is also what is required to exploit the triangulation strategy in connection with genuine concepts in a way that responds to the worry about disjunctive proximal stimuli.

Consider a man who reliably responds (as one wants to say) to the visible presence of rabbits by saying "Gavagai." Suppose further that he is reliably differentially responding not just to rabbits, but to the presence of the distinctive (according to him) rabbit flies that are for him decisive evidence of the presence of rabbits, or that the visual cue he is using, as determined by a physiologist of perception, is a glimpse of the fluff around the tail of the rabbit. What is it about the situation in virtue of which he can be said nonetheless to be reporting not the presence of the rabbit flies or of the fluffy tail but the presence of a rabbit? The inferentialist response is that the difference is not to be found in the reliable differential responsive dispositions, not in the causal chain of covarying events that reliably culminates in the response 'gavagai', to which not only the rabbit but the flies or the fluffy tail belong. It lies rather in the inferential role of the response 'gavagai'. For instance, does the commitment undertaken by that response include a commitment to the claim that what is reported can fly? Or is the claim expressed by 'gavagai' incompatible with the further characterization of the item reported as flying? If it is incompatible, then it is not the flies that are being reported. What determines which element of the causal chain of covarying events that reliably elicit the report is being reported is the *inferential* role of the report, what it *entails*, what is *evidence* for it, what it is *incompatible* with.

Assuming that the observable predicate corresponding to 'flying' has already picked out the things that fly, noticing that the report 'gavagai' could mean rabbit flies in case its applicability entails the applicability of 'flying' and could not mean rabbit flies in case its applicability is incompatible with the applicability of 'flying' is just what is wanted to pick out the distal stimulus the concept expressed by 'gavagai' is being applied to or is classifying. But the appeal to inference and incompatibility may seem just to put off the issue. How does 'flying' get to apply properly to flying things, and not

to whatever cues we in fact use in discriminating flying things—in short to one element of the causal chain of covarying event types that reliably culminate in its application? The answer must be that what the appeal to inferential role does is establish a sort of *triangulation*, or intersection of flows of information or reliable differential responsive dispositions. If 'gavagai' is used so as to entail 'flying', then whatever is properly responded to by the former expression must be properly responded to by the latter, so what is classified as gavagai must also be classified as flying, so 'gavagai' must apply to rabbit flies, and not to the rabbits that are their invariable (we are supposing) concomitants. In short, the appeal to inferential role, in addition to reliable differential responsive dispositions, involves triangulation of the sort that Dretske invokes, where two (or more) different reliable responsive dispositions of the system are invoked, so that their intersection can pick out a unique element of the causal chain of covarying events as the stimulus being classified by a response of a certain type. Because 'flying' will *not* be taken to apply to lots of things that merely hop, we can be sure that it does not mean *flying or hopping*, and so that 'gavagai' does not mean something disjunctive like *rabbit or rabbit-fly*.[45]

In sum, to make the triangulation approach to distinguishing distal stimuli work, one needs to look further 'downstream' from the response, as well as 'upstream'—just as orthodox functionalism would lead one to expect. What picks one kind of thing out as what is being *reported*, from among all those that are being differentially *responded* to, is a matter of the inferential commitments that response is involved in. These inferential consequences of going into a state make it clear that what is being classified is something outside the system. They are what determine that a physicist is reporting the presence of a mu-meson in a bubble chamber, and not simply a large hook-shaped pattern. For the consequences of classifying something as a microscopic mu-meson are quite different from those of classifying something as a macroscopic hook-shaped trace. It is the lack of such consequences that makes Dretske's dual thermometer liable to a disjunctive proximal interpretation. The conclusion is that *causal triangulation* by intersecting causal chains associated with reliable differential responsive dispositions must be supplemented by *inferential triangulation* associated with different concepts.

The minimal condition on singular reference that has been extracted from Frege in this section amounts to the demand that objects be picked out by *substitutional triangulation*. Taking it that an expression is being used to pick out an object is taking it that that *same* object could be picked out in some *other* way—that some commitment-preserving substitutions involving that expression are in order. Substitutional commitments are compound inferential commitments, corresponding to *patterns* of simple inferential commitments. Substitutional articulation is a kind of inferential articulation, and substitutional triangulation is a kind of inferential triangulation. The notion of substitutional commitments is what is needed to explain what

it is to take two distinct claims (whether responsively elicited or not) to be applications of concepts to the *same* object. The significance of causal triangulation is to be understood in terms of the supporting role it can play in this sort of substitution-inferential triangulation. It cannot by itself provide an analysis of picking out objects. And as Frege's discussion of picking out abstract objects shows clearly, however important a role it plays in the way perceivable objects are given to us, causal triangulation is not even a necessary component of the substitution-inferential triangulation that is what our cognitive grip on objects in general consists in.

7. *Conclusion and Prospectus*

This is by no means to say that the discussion of substitutional triangulation here and in the previous chapter suffices to understand what our talking and thinking about objects consists in. That discussion addresses primarily the issue of what it is for it to be *objects* (and their properties and relations) that our talk (and so our thought) purports to be about. To understand fully what it is for our thought and talk to purport to be *about* them requires an account of the crucial *social* dimension of the substitutional triangulation that structures the contents expressed by the use of singular terms (and predicates) and of the inferential triangulation that structures the contents expressed by the use of sentences. The way in which the social structure of the broadly inferential articulation of discursive practice bears on the nature of the conceptual contents that practice confers on the intentional states it institutes (and on the performances that express them) is already implicit in the discussion of discursive practice in terms of deontic scorekeeping, in Chapter 3. It is the task of Chapter 8 to make it explicit, and thereby to show how the representational dimension of conceptual content arises out of, and essentially depends on, differences in social perspective among the various discursive practitioners.

A further shortcoming in the account of picking out objects in terms of substitutional triangulation as adumbrated so far is that it is primarily addressed to the phenomenon of *purported* singular reference. Although general reasons have been offered motivating a direction of explanation that begins with the notion of representational purport, it remains to say something also about the *success* of such purport. To this end, the next section discusses what we are doing when we take it that a singular term succeeds in referring, in that the object the term purports to refer to actually *exists*. An account is offered of *existential* commitment as a kind of *substitutional* commitment. This story in turn permits an analysis of the commitments characteristic of the use of expressions as definite descriptions, and so shows how to extend the deontic scorekeeping model from languages with predicates and lexically simple singular terms to ones that contain definite descriptions as well.

The rest of the chapter then addresses the structures that make substitu-

tional triangulation (and so recognizing an object as the same again) possible when *unrepeatable* expression tokenings are involved. This means above all the *deictic* or *demonstrative* use of terms that is so important for understanding the role causal triangulation can play in substitutional and inferential triangulation, and so ultimately for understanding what is distinctive about *empirical* knowledge. Moving to the level of unrepeatable tokenings requires discerning a finer structure of *token recurrence* below that of *substitution*, just as the finer substitutional structure had to be discerned below that of *inference*. The key concept in this account is that of *anaphora*. Explaining the anaphoric linkage of tokenings in terms of the inheritance of the determination of substitutional commitments provides an official account in deontic scorekeeping terms of the phenomenon by means of which the traditional semantic vocabulary, 'true' and 'refers', was explained in Chapter 5, redeeming the promissory note issued there. So by the end of this chapter the full three-leveled structure of fundamental concepts in terms of which conceptual content is to be understood in the semantic portion of the present account will have been made available: *inference*, *substitution*, and *anaphora*. At that point the semantic raw materials will be on hand to be combined with the underlying pragmatics to yield in Chapter 8 an account of representation by conceptual contents, on the semantic side, and objectivity of conceptual norms, on the pragmatic side.

II. DEFINITE DESCRIPTIONS AND EXISTENTIAL COMMITMENTS

1. *Forming Singular Terms from Predicates*

To talk about the singular referential purport of singular-term usage is to talk about what kind of substitutional commitments one must attribute (and, as will emerge in the next chapter, acknowledge) in order for what one is doing—the practical deontic scorekeeping attitude one is adopting—to count as taking someone to be using an expression as a singular term.[46] To be doing that, one must treat the use of the expression as governed by proprieties determined by symmetric simple material substitutional commitments—commitments that in languages with sufficient logical expressive power are made explicit in the form of the nontrivial identity claims that Frege calls "recognition judgments." The substitutional commitments involving a singular term that a scorekeeper attributes and undertakes determine the pragmatic significance, for that scorekeeper, of each use of that term. That the significance for deontic scorekeeping of its occurrences is to be determined in this way is what treating it as a singular term (as purporting to pick out an object) consists in.

The referential purport that is in this way acknowledged or attributed concerns the committive antecedents and consequences of application of singular terms as such. Chapter 6, some of the points of which were reca-

pitulated in Frege's terminology in the previous section, showed how those committive circumstances and consequences of application can be understood substitutionally. But referential *purport* is one thing, referential *success* is another. What is the difference between taking it that an expression has been introduced as purporting to refer to a definite object and taking it that it in fact picks out or gives us a cognitive or semantic grasp on such an object? This is a question about a certain kind of *entitlement* to the substitutional commitments in which singular referential purport consists, and so in a broader sense about the appropriate circumstances of application of singular terms as such.

The deontic attitudes that constitute taking the singular referential purport characteristic of singular terms to be successful emerge most clearly from consideration of what is involved in taking someone to be entitled to use a definite description formed from a predicate. This is what Frege calls "the definition of an object in terms of a concept under which it falls."[47] As an example of a definite description that is defective—whose referential purport is not successful because it involves substitutional commitments the user cannot in the relevant sense be entitled to—he considers the expression 'the largest proper fraction'. The predicate that description is formed from is one that can be used to express commitments with appropriate entitlements.

> The expression "the largest proper fraction" has no content, since the definite article purports to refer to a definite object [*der bestimmte Artikel den Anspruch erhebt, auf einen bestimmten Gegenstand hinzuweisen*]. On the other hand, the concept "fraction smaller than 1 and such that no fraction smaller than 1 exceeds it in magnitude" is quite unexceptionable: in order, indeed, to prove that there exists no such fraction, we must make use of just this concept, despite its containing a contradiction. If, however, we wished to use this concept for defining an object falling under it, it would, of course, be necessary first to show two distinct things:
>
> 1. that some object falls under this concept;
> 2. that only one object falls under it.
>
> Now since the first of these propositions, not to mention the second, is false, it follows that the expression "the largest proper fraction" is senseless.[48]

Suppose that a predicate $P\alpha$ has been introduced and is in use. (This is a supposition that has been given definite content in substitutional terms by the discussion of Chapter 6.) The problem Frege is addressing is to make explicit what else is required for it to be proper to take someone to be entitled to use a definite description formed from it—a singular term of the form 'the P', or as it may be expressed more generally, $!x(Px)$. Treating $!x(Px)$ as a singular term is taking it that its use is governed by symmetric simple

material substitution-inferential commitments, that is, that there is some true nontrivial identity involving it. Thus, there must be some true recognition statement of the form $!x(Px) = a$.

As already indicated, the two conditions Frege imposes are that there be something that is *P*, and that it be unique:

1. *Pa* and
2. For any *y*, if *Py*, then $y = a$.

The second condition—uniqueness—can be straightforwardly parsed in the substitutional idiom already available. For it just amounts to saying that for any terms *t*, *t'*, if *Pt* and *Pt'*, then $t = t'$; the substitutional commitments associated with claims of these forms have already been explained. The first condition requires more discussion, however. In the *Grundlagen* Frege frames his dispute with the formalists in terms of the necessity of proving the existence of an object falling under the concept, or as he puts it, "*producing* something that falls under it."[49] What must one do to satisfy this requirement?

2. *Substitutional Commitments Expressed by Quantifiers*

It too can be understood in substitutional terms. Existential commitments are a kind of substitutional commitment, related to, but not identical with, the substitutional commitments involved in the use of quantifiers. As the discussion of the formation of complex predicates in the previous chapter indicates, universal and particular quantifiers are logical locutions that have the expressive function of making propositionally explicit conjunctive and disjunctive substitutional commitments. Attributing commitment to a claim of the form $(x)Px$ is attributing commitment to *all* claims of the form *Pa*. Such a substitutional rendering of quantification has been criticized as inadequate in cases where, for cardinality reasons, there are not enough singular terms to pick out all the objects one is quantifying over. It is very important that we be able to make claims about all real numbers—for instance that every one can be represented by converging sequences of rational numbers—even though we are in principle limited to the use of at most a countable number of singular terms referring to them.[50] In fact, however, that the stock of available substituends is in this way *limited* threatens a substitutional construal of quantifiers only if that stock is in addition conceived of as being *fixed*.

It is of the essence of singular-term usage that new terms can always be introduced—both new terms for familiar objects and terms that introduce unfamiliar objects, paradigmatically by description. We cannot indeed extend our language so as to have separate terms for *all* real numbers at once, but *each* real number can be picked out. For there is no real number that we cannot specify—for instance by a definite description in terms of converging

sequences of rational numbers. The substitution instances *Pa* one becomes jointly and severally committed to by committing oneself to a claim of the form $(x)Px$ include, not only those formed from terms *a* that are currently in the language, but all those that *could* be introduced (not necessarily simultaneously). The substitutional construal of quantificational commitments requires that the expressive powers of a set of discursive practices be conceived in the wider sense that takes account of the possibility of introducing novel expressions, rather than in the narrower sense that restricts attention to locutions already actually in use. This latter view amounts to freezing an idiom: taking a snapshot of it and evaluating its expressive capacities in abstraction from the process by which it develops. It is encouraged by thinking of languages as formal objects (perhaps set-theoretic structures) that have fixed vocabularies. If languages are instead conceived as living practices, then the ways in which new vocabulary is introduced take their place as fundamental aspects of those practices—as central as the ways in which new claims are made. Frege is the father of the formal approach to languages, but his project in the *Grundlagen* leads him to be vitally concerned with the process of introducing novel expressions functioning as singular terms, not only for unfamiliar objects of familiar kinds (by description), but even for unfamiliar objects of unfamiliar kinds (by abstraction).

Similarly, the use of a particular quantifier in connection with a complex predicate makes explicit a disjunctive substitutional commitment to the effect that for some term *a*, *Pa*. To be entitled to such a claim one may, but need not, be able to produce the relevant substitution instance. The vindicating substituend *a* need not even already be in the language; one is committed only to the possibility of introducing such a term. The point of the existence requirement Frege imposes on entitlement to introduce definite descriptions is that a certain kind of bare stipulation is not in general enough to entitle one to such term introduction. One is not permitted without further ado to introduce the expression $!x(Px)$ and then, relying on the fact that $P(!x(Px))$ (whenever use of the definite description is appropriate), to use that description as the substituend that vindicates the claim made by use of a particular quantifier. The large question of interest in this section is precisely what that existential condition on the employment of definite descriptions comes to.

Though it is common to do so, it is not necessary, however, to extend the existential condition Frege imposes on the use of definite descriptions to whatever counts as a vindicating substituend for a particular quantification. It is for this reason that the general formal notion of particular (that is, disjunctive) quantification should be distinguished from the substantive notion of specifically existential quantification. Free logics distinguish particular quantificational commitments from existential commitments so as to allow an idiom in which 'Pegasus is a winged horse' can count as true, even though Pegasus does not physically exist, and so in which 'Pegasus' can serve

as a substituend that vindicates a particular quantificational commitment to there being some winged horses. Frege, of course, does not want to talk this way, taking it that because 'Pegasus' has sense but no referent, 'Pegasus is a winged horse' cannot be true. For many purposes (certainly for Frege's), this policy is no doubt the best. Nonetheless there is nothing incoherent about scorekeeping practices that permit particular quantificational commitments to be vindicated by term substituends with respect to which the scorekeeper does not undertake existential commitments, and considering such ontologically relaxed idioms highlights certain important features of genuinely existential commitments.

The substitutional significance of particular quantification is entirely determined by features of discursive scorekeeping practices that have already been discussed if it is stipulated that a particular quantificational commitment with respect to the predicate $P\alpha$ is to be equivalent to the commitment expressed by $\sim(x)\sim Px$. The negation of a claim p was defined in Chapter 2 as its minimum incompatible: the inferentially strongest claim that is commitment-entailed by every claim incompatible with p. Thus the claim that for some x, Px is incompatible with any claim that for every term a entails some claim Qa that is incompatible with Pa. So the particular quantificational claim that for some x, Px is both commitment- and entitlement-entailed by any claim of the form Pa (but not necessarily just by these).

3. *Sortally Restricted Substitutional Commitments*

In his technical systems (both that of the early *Begriffsschrift* and that of the late *Grundgesetze*) Frege offers a substitutional account of the formation of complex predicates and of the formation of sentences from them by the application of quantifiers. In each case the scope of the quantifiers is unrestricted: *every* well-formed singular term can serve as the substituend that vindicates a particular quantification and can serve as a potential counterinstance to a universal quantification. One consequence of running these systems with their quantifiers wide open is that in order to give quantificational claims the force Frege wants them to have—above all for the assertion of claims formed by particular quantification to involve the undertaking of specifically *existential* commitment—Frege must ensure that it can be proven that every well-formed singular term has a referent. As Russell notoriously showed, another consequence—in the context of the expressive power provided by unrestricted formation of complex predicates or sentential functions by substitution—is that the resulting systems are inconsistent. That unpalatable result has prompted the investigation of how the various theoretical commitments that conspire to produce it might be relaxed so as to avoid it. One popular candidate is Frege's insistence on unrestricted quantification; the strategy of placing restrictions on the substitution instances relevant to the semantic evaluation of claims formed by the application of

particular and universal quantifiers is what lies behind the theory of types Russell develops in *Principia Mathematica* to evade Frege's difficulty.

Restricted quantification, however, is not of merely technical interest. Indeed, Frege's unrestricted version appears as an artificial extrapolation once it is realized that in natural languages ordinary quantificational tropes are sortally restricted.[51] The central uses of quantifiers are to make claims such as:

> Every integer is the sum of nineteen or fewer fourth powers.[52]
> Some nineteenth-century German philosophers did not care about ontological issues.
> All bank employees must wear neckties.
> A deer made those tracks less than an hour ago.
> Each of us has intentional states.

The central quantificational construction is *every K* or *some K*, where *K* is a sortal expression such as 'dog' or 'book'. 'Everyone' and 'someone' have the sense of 'every person' and 'some person', and even the apparently wide open 'everything' usually carries some restriction, either explicitly, as in

> Everything the author says about propositional content is confused,

or implicitly, as in

> Everything is a disappointment in the end.

The sortal restriction puts conditions on allowable substituends, so that even though 'the author of *The Stones of Venice*' is a perfectly good singular term, substitution instances formed from it are not relevant to the semantic evaluation of "Every integer is the sum of nineteen or fewer fourth powers."

As Frege indicates in the *Grundlagen*, sortals are like predicates, except that they have not only criteria and consequences of *application* but (like singular terms) also criteria (and so consequences) of *identity*. For many purposes, '. . . is a dog' functions predicatively, just as '. . . is large' does. But if *a* is a dog and *b* is a dog, it makes sense to ask whether *a* is the *same* dog as *b*. Sortals have associated with them practices of identifying and individuating the things they apply to, as nonsortal predicates do not. So in order to introduce the sortal 'number', Frege insists on "a general criterion for the identity of numbers [*Kennzeichen für die Gleichheit von Zahlen*]."[53] When made explicit in the form of a claim, such a criterion has the form:

> If *x* is a *K* and *y* is a *K* and *Rxy*, then *x* is the same *K* as *y*.

Introducing a sortal, like introducing a predicate, requires fixing the sense of claims formed by substitution into sentence frames of the form "α is a *K*"

("α is P"), but it requires in addition establishing a criterion of identity. Satisfaction of this additional constraint ensures that K's can be *counted*.

In fact, establishing a criterion of identity (and so a sortal) is not only sufficient for countability, it is necessary as well. Unsortalized 'things' or 'objects' cannot be counted. There is no answer to the question how many things there are in this room; there is one number of books, another of molecules, another of atoms, another of subatomic particles. As Frege says: "If I place a pile of playing cards in [someone's] hands with the words: Find the Number of these, this does not tell him whether I wish to know the number of cards, or of complete packs of cards, or even say of honour cards at skat. To have given him the pile in his hands is not yet to have given him completely the object he is to investigate; I must add some further word—cards, or packs, or honours."[54] Counting is intelligible only with respect to a sortal concept.

'Thing' and 'object' are pseudosortals. They can occupy the syntactic positions occupied by sortals, but they do not individuate as sortals must. They are mere placeholders for sortals, used when for some reason—often the clumsy disjunctiveness of the sortal that would be required—one does not want to specify the relevant sortal explicitly. When we say something like "Put everything that is on top of the desk into the drawers," we usually mean all the middle-sized bits of dry goods: books, papers, pens, paper clips, and so on. We do not mean 'things' such as designs in the dust, cool spots, drops of water, and so on. One of the central uses of 'one' in English is as an *anaphoric prosortal*—an anaphoric dependent standing in for a sortal that is its antecedent—as in "John quoted an English philosopher, and I quoted a German one," or "Eric wants an ice cream cone, and Russell wants one too." In these examples 'one' is used in the 'lazy' anaphoric way, where it is replaceable by its antecedent sortal. Like the pronoun 'it', however, it is promiscuously available to stand in for a wide variety of antecedents. 'Thing' and 'object' are what one gets if one misunderstands this grammar and instead construes 'one' as expressing a genuine sortal.

Frege in fact makes exactly this objection to the attempt to press the term 'unit' (or 'one') into generalized duty in place of substantive sortals in theorizing about counting. His own view is that the invocation of substantive sortal concepts cannot in this way be avoided; he takes it rather that "a concept [is] the unit relative to the Number which belongs to it."[55] Not all concepts will do; only those expressed by sortals (rather than predicates without individuating criteria of identity): "The concept 'syllables in the word "three"' picks out the word as a whole, and as indivisible in the sense that no part of it falls any longer under the same concept. Not all concepts possess this quality. We can, for example, divide up something falling under the concept 'red' into parts in a variety of ways, without the parts thereby ceasing to fall under the same concept 'red'. To a concept of this kind no finite number will belong. The proposition asserting that units are isolated

and indivisible can, accordingly, be formulated as follows: Only a concept which isolates what falls under it in a definite manner, and which does not permit any arbitrary division of it into parts, can be a unit relative to a finite Number."[56] This insight ought to have led Frege to see quantifiers as coming with sortal restrictions on the admissible term substituends. For quantifiers quantify; they specify, at least in general terms, *how many*, and how many there are depends (as Frege's remarks about playing cards indicate) on *what* one is counting—on the sortal used to identify and individuate them. As Frege saw clearly, the use of quantifiers depends on the use of the singular expressions that provide their substitution instances. It is best therefore to think not only of quantifiers but of singular terms as properly introduced only in connection with some at least implicit sortal.

Definite descriptions should be explicitly sortally restricted: 'the *man* in the brown suit', 'the *book* that Carlyle had to rewrite because of Mill's maid', and so on. Individual proper names and demonstratives and other indexical expressions cannot properly be understood except in terms of their associated sortals. Thus one cannot simply point in the direction of a statue of a man on a horse and christen ^{s}it^{s} 'Lumpl'. It matters whether one is naming the statue or the lump of clay it is made of. If the former, reshaping it into a statue of a mother with a child destroys Lumpl; if the latter, not—for the transformed figure is the same lump, but a different statue.[57] If I hold up my copy of Kant's first Critique and ask "Has Eric read this?" my remark is susceptible to two different sorts of readings, depending on whether the demonstrative is associated with the sortal that individuates books according to the content of the text or rather (as might arise if I have just discovered a large jelly stain defacing the page that sets out the Table of Judgments) the sortal that individuates them according to particular physical copies. 'This' or 'that' used by itself should on these grounds always be understood as elliptical for 'this *K*' or 'that *K*'. Again, it is important that 'I' implicitly invokes the sortal that individuates persons—it is a *personal* pronoun. For I, who am flying to London, am the same person who last month flew to Philadelphia, while I am *not* the same passenger who did so.[58] The discussion in this vicinity about "relative identity" is prompted by this sort of observation. But it often takes the form of a mysterious thesis about *things*, rather than a clear one about the conditions that ought to be met to count as having introduced (or understood) a singular term (even a tokening of a demonstrative) as having a definite reference.[59] Such a confusion is the result of thinking of sortally unrestricted quantification and singular-term usage as conceptually fundamental, and seeing sortal restrictions as optional additions—rather than seeing the restricted case as fundamental, and unrestricted quantification as a dangerous and often unwarranted extrapolation based on a misunderstanding of the way pseudo- and prosortals such as 'thing', 'object', 'one', and 'item' function. Frege's requirement that to introduce a referring expression one must fix the sense of identities involving it—settle how it is

to be distinguished from other objects—turns out to require that it be associated, implicitly or explicitly, with an individuating sortal concept.

4. *Existential Commitments*

Existential commitment is a species of substitutional commitment. It can be thought of as a particular quantificational commitment that involves a special sort of restriction on the vindicating substituends that determine the content of that commitment. Generically, the restriction on substituends characteristic of existential commitments is akin to the sortal restriction involved in quantificational commitments—and more generally in the use of any expressions conveying singular referential purport. The structure of the restriction is different, however. What is distinctive of specifically existential commitments is the special role played by a distinguished class of admissible substituends, here called *canonical designators.* The difference between the substitutional function of canonical designators and that of sortally qualified substituends in ordinary quantification is what stands behind the principle that existence is not a property—that existent objects are not a *kind* of object.

The best way to appreciate the role an expression must play to be functioning as a canonical designator with respect to a species of existential commitment is by examples. Three different sorts of existential commitment are considered here, corresponding to *numerical* existence, *physical* existence, and *fictional* existence: the sense in which there is a number such that every number greater than it is the sum of distinct primes of the form $4n + 1$, the sense in which there is a pen on my desk, and the sense in which there is someone who keeps house for Sherlock Holmes, respectively. By looking at these different kinds of existential commitment, it is possible to see what they all have in common, in virtue of which they deserve to be understood as species of a genus—so that 'exists' can be understood as univocal, in spite of the important differences between commitment to the existence of particular numbers, of physical objects, and of fictional characters. The claim is that what these different sorts of existential commitments share—what makes them all varieties of *existential* commitments—is the way in which their pragmatic significance is determined by a set of expressions playing the role of canonical designators. The manifest differences between them are consequences of the very different sorts of expressions that serve as canonical designators in each case.

It is clear how the sense of the expression "the smallest natural number such that every larger one is the sum of distinct primes of the form $4n + 1$" is to be determined. The predicate this definite description is formed from has clear circumstances and consequences of application, and its inclusion of the specification 'the smallest' ensures that if that predicate applies to any natural number, it is to a unique one. (In this way it can be contrasted with

the predicates 'largest possible fraction' or 'most rapidly converging sequence', which Frege considers.) What does it mean to say in addition that in fact there is a number to which the predicate applies—that the smallest natural number such that every larger one is the sum of distinct primes of the form $4n + 1$ *exists*? What is it for the definite description not only to *purport* to refer to or pick out a particular number but actually to *succeed* in doing so?

In this case its success consists in the truth of this identity:

> The smallest natural number such that every larger one is the sum of distinct primes of the form $4n + 1 = 121$.

Entitlement to the existential commitment implicit in the use of the definite description, and so entitlement to use that description, can be secured by entitlement to *any* identity of this form. Of course the existential claim is not equivalent to the substitutional commitment that is made explicit by this particular identity—or indeed, any other of this form. The existential commitment is rather equivalent to the disjunctive claim that some identity of this form is true; the significance of the existential commitment is determined by its being incompatible with any claim that is incompatible with all claims of the form of this identity. The significance of such existential commitments is accordingly to be understood, and their propriety assessed, in terms of the class of vindicating substituends supplied by identities of the form of the one above.

What is the relevant form? Not just any identity will do. For instance

the smallest natural number such that every larger one is the sum of distinct primes of the form $4n + 1$	=	40 less that the smallest natural number such that every larger one is the sum of distinct primes of the form $6n - 1$

is an identity (substitution license) that does not, like the one above, settle it that the smallest natural number such that every larger one is the sum of distinct primes of the form $4n + 1$ *exists*. It just links that question to the question of whether the smallest natural number such that every larger one is the sum of distinct primes of the form $6n - 1$ exists. If the latter number exists, then so does the former. The claim is, however, that in identifying the first number with 121 (or the second with 161), one is doing more than merely settling this sort of conditional existential question. One is in that case settling the categorical existential question of whether the existential commitments implicit in the use of these definite descriptions are in order, whether those descriptions are successful singular referring expressions, whether the numbers they purport to specify exist.

To say this is to say that the issue of the success of their singular referential purport does not arise for expressions such as '121' and '161' in the same way that it does for expressions such as 'the smallest natural number such

that every larger one is the sum of distinct primes of the form $4n + 1$'. It is to take a frankly inegalitarian approach to referential purport and its success.[60] Numerals are semantically privileged ways of picking out numbers. By contrast to definite descriptions of numbers, the well-formedness of numerals suffices for their referential success, guaranteeing that they pick out a corresponding object. Furthermore, distinct numerals are guaranteed to correspond to distinct objects. According this privileged status to a class of singular terms is treating them as *canonical designators* of a kind of object.

In the paradigmatic case of natural numbers, numerals such as '121' can serve as canonical designators because they are systematic abbreviations of successor numerals: elements of the sequence

$$0, 0', 0'', 0''', 0'''' \ldots .$$

Peano's axiomatization ensures that numerals of this form cannot fail to refer to numbers, and further that their lexical distinctness (marked by the number of successor marks they bear) suffices for the distinctness of the numbers they refer to. Claiming that some noncanonical numerical expression succeeds in referring to a number is just claiming that it is intersubstitutable (saving discursive commitments) with some element in the canonical sequence of successor numerals. Existential commitment with respect to this sort of object, natural numbers, is a disjunctive substitutional commitment linking the candidate numerical expression to some canonical substituend. Saying which number a numerical expression refers to is producing the canonical designator that is intersubstitutable with it. (In this sense one has not yet said how many seconds there are in a century when one asserts the identity

$$\text{the number of seconds in a century} = 100 \cdot 365\tfrac{1}{4} \cdot 24 \cdot 60 \cdot 60.$$

One has only given a recipe that would make it possible, with some work, to say which number the definite description picks out—a recipe that guarantees at least that it does pick out some definite number.) Once the use of some expressions as canonical designators has been established, Frege's requirement that entitlement to use an expression as a singular term depends on its having been settled which object it refers to in a sense that includes distinguishing that object from others—the requirement that becomes explicit in the introduction conditions he imposes on definite descriptions—is satisfied for numerical expressions by settling it that there is some canonical designator linked to the expression in question by a true recognition statement: a nontrivial identity claim making explicit a simple material substitution-inferential commitment. Ensuring that novel singular terms are suitably substitutionally linked to canonical designators establishes both the existence and the uniqueness of the objects they pick out, and so secures the

success of the singular referential purport that distinguishes them as singular terms.

So the notion of canonical designators makes it possible to think of existential commitment as a kind of substitutional commitment. It can be thought of as a particular quantificational commitment in which the vindicating commitments that determine its content are restricted to canonical designators. It is important to notice, however, how differently such a restriction functions from the sortal restrictions associated with quantificational commitments generally. The sortal restriction is a restriction to kinds of *objects*, while the restriction to canonical designators is a restriction to kinds of *expression*. In substitutional terms, this means that if a sortal qualifies one of an equivalence class of intersubstitutable (that is, coreferential) singular terms, it qualifies all of them. If Kant is a person and Kant is the author of "Was ist Aufklärung?" then the author of "Was ist Aufklärung?" is a person. So the sortal restriction does not discriminate among different ways of referring to the same thing. Not so for the restriction to canonical designators. '9' is a canonical designator of a natural number,[61] and 9 is the number of solar planets, but 'the number of solar planets' is not a canonical designator of a natural number. It is of course a designator of a natural number: 'natural number' specifies a sortal, picks out a kind of thing. But 'canonical designator' picks out a kind of expression, not a kind of thing. So the structure of the restriction on admissible substituends involved in existential commitment is quite different from that involved in ordinary sortally restricted quantification. Existence is not a predicate or property, and existing things are not a *kind* of thing.

5. *The Role of Canonical Designators*

For a class of singular terms (for instance successor numerals) to have the status of canonical designators with respect to a kind of objects (for instance natural numbers) is a matter of the significance their use has according to the relevant discursive scorekeeping practices. The institution of that significance presupposes, rather than establishes, entitlement to use those expressions as singular terms, however. As Frege would be the first to insist, one cannot simply stipulate that the use of successor numerals as singular terms is in order. Like all expressions, their use must be governed by some nontrivial identities for it to count as the use of expressions as singular terms at all. The point of Frege's disagreement with the formalists is that merely laying down the Peano axioms is not enough to satisfy this requirement. Identities of the form

$0''' =$ the successor of $0''$,
$0''' =$ the successor of the successor of $0'$,

and so on will not do because such identities are in the relevant sense trivial;

they do not link two different ways of picking out an object because the expressions flanking the identity sign are merely notational variants of each other. Thus they cannot serve as genuine recognition statements; they do not express substantive material substitution-inferential commitments. Frege's response to this situation in the *Grundlagen* is to link the use of the successor numerals to that of other expressions already in use, by means of the method of abstraction. In this way the use of successor numerals is determined by its relation to the process of counting the previously discriminated objects that fall under some sortal. His definition then permits versions of the Peano postulates to be *proven* to hold for expressions introduced in this way. Thus entitlement to the use of successor numerals as singular terms is secured, and they become available for duty as canonical designators. One must be entitled to use expressions as designators first, and only then can they serve as *canonical* designators, which can be appealed to in explaining what existential commitment consists in.

As Frege and others have shown, once one understands existence claims regarding expressions that purport to pick out natural numbers, one can systematically extend that understanding to existence claims regarding rational, real, and complex numbers, and so on. The canonical designators that give sense to the notion of existence for rational numbers can be pairs of successor numerals (corresponding to ratios of natural numbers), for real numbers they can be converging sequences of canonical specifications of rational numbers, for complex numbers pairs of such canonical specifications of reals, and so on. (Notice that in order to be entitled to use a sequence of canonical rational number designators as a canonical designator of a real, one is obliged to *prove* the convergence of the sequence of numbers those designators pick out.) The idea behind the use of successor numerals as canonical designators in explicating existential commitments as a species of substitutional commitment is that to say that some numerical expression succeeds in referring—to say that a number corresponding to it *exists*—is to say that it has some address in the structured space mapped out by the successor numerals. This idea can be carried over, with some differences, to the case of *physical* existence.

To say that some physical object expression succeeds in referring, that the object it designates exists, is to say that it exists *somewhere* in space and time, that it occupies some spatiotemporal region. This is to say that it has some address in the structured space of spatiotemporal coordinates centered on the speaker.[62] The speaker who takes it that Pegasus does *not* (and never did) exist, while P. T. Barnum's elephant Jumbo *does* (or did) is claiming that a continuous spatiotemporal trajectory cannot be traced out connecting the region of space-time occupied by the speaker[63] to one occupied by Pegasus, while such a trajectory can be traced out connecting the speaker's region with that occupied by Jumbo. It is not that Pegasus must be conceived as not taking up any space or surviving for any time; it follows immediately from

his being a horse that he does both. It is that the region he occupies is inaccessible from here and now—"You can't get there from here." He does not exist in *our* space and time, the one that defines physical existence for us. The analog in the case of physical existence of the structured address space defined by the successor numerals is the structured address space defined by egocentric spatiotemporal coordinate descriptions. Thus the term 'Pegasus' is not properly intersubstitutable with any expression of the form 'the (winged) horse located at ⟨x, y, z, t⟩ from here, while the term 'Jumbo' is intersubstitutable with an expression of the form 'the elephant located at ⟨x, y, z, t⟩ from here'.[64] Thus, like numerical existential commitments, physical existential commitments can be understood as substitutional commitments involving a class of canonical designators (again a kind of *expression*, not a kind of *thing*).

Of course there are also disanalogies between the way the spatiotemporal designators that are canonical for physical existence work and the way the successor numerals that are canonical for numerical existence work. Here-now centered coordinate specifications of accessible spatiotemporal regions are, like successor numerals, guaranteed to succeed in their referential purport. But the canonical designators of physical *objects*, as opposed to the regions they occupy, must include sortal information as well: the statue and the lump of clay *may* occupy just the same spatiotemporal region over the whole course of their existence. The sortals relevant in this case are those where identity (or difference) of the spatiotemporal regions occupied guarantees identity (or difference) of the objects within the sortal (as opposed to across sortals, as in the lump/statue case). Thus if horse_1 occupies region r and horse_2 occupies region r, then $\text{horse}_1 = \text{horse}_2$ (and if not, not).[65] The individuation of horses is parasitic on spatiotemporal individuation, in that if one has used the horse-specific criteria of application of 'horse' to stick labels only on horses, one then uses spatiotemporal coincidence or divergence to decide how many horses have been labeled and how many labels each horse has.

As in the case of numerical existence, these existential commitments can be understood as substitutional commitments involving physically canonical designators only where one is entitled to apply those canonical designators—which in this case are formed from sortals plus specifications of accessible space-time regions. In the case of successor numerals, this could be done wholesale—'producing' objects for them to refer to by abstraction, so that the Peano postulates could be proven. As in the case of using converging sequences of canonically designated rational numbers as canonical designators of real numbers, however, not only must one pick out a privileged general form as canonical, one must also settle which of the designators of that form are suitable for endorsement as canonical. Just as one must prove the convergence of each sequence of rational numbers that is put forward as a substituend that could vindicate a commitment regarding the existence of a

real number, so one must show in the physical case for each sortal-plus-region pair that the sortal properly applies to the region—that the region specified is occupied by an elephant. The variety of spatiotemporally individuating sortals means that there is nothing useful and general to say about how one becomes entitled to claims applying a sortal to a region. The appropriate circumstances of application for applying the sortal-derived predicate '. . . is an elephant', or '. . . is occupied by an elephant' to a particular spacetime region are quite different from those of '. . . is (or is occupied by) an electromagnetic force field'. But these details concern the use of these particular sortals and predicates, not the notion of existence in general. The surplus significance of a commitment to physical existence lies in the *accessibility* to the one undertaking the commitment (via a continuous trajectory from here-now) of a spatiotemporal region to which the sortal (or its derived predicate) is properly applicable. For that reason appealing to the notion of a predicate or sortal being applicable to a region does not make this way of thinking about physical existence circular.

As a final example, *fictional* existence, existence in or according to a story, can be understood as having the same shape as that common to physical existence and the various sorts of numerical existence. To say that in or according to the Sherlock Holmes stories Holmes's housekeeper exists (or that the expression 'Holmes's housekeeper' succeeds in referring to an individual) is to say that that expression is intersubstitutable with some singular term that actually appears in the story (in this case a tokening of 'Mrs. Hudson'). The singular terms that appear in the text that defines the fictional context can be considered as the canonical designators. Thus the claim that according to those stories Holmes's archenemy exists but his fairy godmother does not involves undertaking a substitutional commitment regarding a canonical designator in the first instance, and a commitment incompatible with it in the second. Even if the phrase 'Holmes's archenemy' does not ever appear in the text, the fact that 'Professor Moriarty' does occur there and that it can be deduced from what is said about him that this term is intersubstitutable with 'Holmes's archenemy' ensures that 'Professor Moriarty' is a canonical designator that can vindicate the substitutional existential commitment. That no such term plays this role with respect to 'Holmes's fairy godmother' is the thrust of the denial of even fictional existence to that creature.

For some purposes it is useful to consider as canonical designators not only expressions that actually occur in the stories but also those, like 'Sherlock and Mycroft Holmes's maternal grandmother', whose applicability is entailed by what is said there, though they are never actually used. These boundaries are hazy, for it is not clear what auxiliary hypotheses one is entitled to appeal to in extracting the consequences of what we are told in the story. In most settings regularities of nature, even if not explicitly mentioned, seem safely carried over, but beyond that the matter seems one for decision rather than

discovery. To say this is just to say that the notion of fictional existence is itself hazy outside those objects actually mentioned in the text in question. Another regard in which fictional existence is ontologically indeterminate as far as singular reference goes is that a difference in canonical designators need *not* here entail a difference of objects. The story may simply be silent (even by implication) on the subject of whether the person who had last hired the hansom cab Holmes is riding in at a certain point was Mrs. Hudson (or Sherlock and Mycroft Holmes's maternal grandmother) or not. One can constrain such individuative issues by importing the physical spatiotemporal individuating apparatus into the fictional context; Pegasus was spatiotemporally accessible to Perseus, according to the story. Like the invocation of regularities of nature, these constraints only go so far, and certainly in typical cases far underdetermine the identity and individuation of the fictional objects referred to by canonical designators.

The point of this discussion does not reside in the particular choices that have been offered here as candidates for sets of canonical designators corresponding to different sorts of existence. It is that existential commitments can be understood as a special kind of substitutional commitment (akin to but distinct from sortally restricted particular quantificational commitments) by using the notion of a privileged set of substituends playing the special substitutional role of canonical designators. What one is doing in claiming that the largest number that is not the sum of the squares of distinct primes exists (its canonical designation is '17,163') is different in specific ways from what one is doing in claiming that the tallest edifice in Washington, D.C., exists; numerical existence is different from physical existence. But these different sorts of existence, or even senses of exist, have a structure in common that qualifies them both as notions of existence. It is that common structure that the notion of existential commitments as substitutional commitments restricted to canonical designators seeks to capture.

If something like this account of existential commitments is right, then kinds of existence are to be individuated by the associated sets of canonical designators. Picking different sets of canonical designators gives different senses of existence. So one who treats specifications of real numbers by *pairs* of sequences of specifications of rationals such that sufficiently late elements of one sequence are arbitrarily close to those of the other would, strictly speaking, mean something different by saying that a certain real number *exists*. Of course, as long as each canonical designator in the one set is intersubstitutable (coreferential) with a designator in the other, this difference would not make a difference. It would be significant only when, perhaps against all expectation, apparently equivalent specifications turn out to diverge in hitherto unconsidered cases.[66]

A set of canonical designators has a *structure*—paradigmatically that of the successor numerals or coordinatized specifications of spatiotemporal regions—which systematically provides *addresses* for all the existing objects

of the class in question. As was indicated above, this need not be a totality that is present in the language at any given moment; the structure may provide only recipes for producing suitable canonical designators, or even just criteria for recognizing them, as is the case for some definitions of real numbers in terms of converging sequences of rational numbers. It remains to say what it is to treat such a set of designators as privileged in the particular way necessary for them to be functioning as canonical designators defining a sort of existential commitment. This is a matter of the consequences of undertaking and attributing existential commitments. One of the consequences, of course, is that being entitled to an existential commitment regarding a definite description is a necessary condition of being entitled to use that description at all. Only scorekeepers who attribute an entitled commitment to the existence of some x such that Fx to an interlocutor take the use of a definite description of the form $!x(Fx)$ by that interlocutor to be appropriate.

More important, anyone who does *not* undertake a commitment to the physical existence of the object referred to by a term t cannot endorse any claims of the form Pt, where P is a physical predicate. Under these circumstances one can be entitled only to endorsements of the fictional, 'according to the story' truth of a claim, which differ from endorsements that are not in the fictional mode—for instance in the unavailability as premises for practical reasoning of the claims that appear inside the scope of the explicitating 'according to the story' operator. Similarly, one does not endorse numerical claims whose expression essentially involves terms with respect to which one is not prepared to undertake numerical existential commitments, and one does not endorse fictional claims whose expression essentially involves terms with respect to which one is not prepared to undertake fictional existential claims. It is this intimate connection between existential and doxastic commitments that led Frege to forbid nonreferring terms in his ideal languages.

Picking out the set of descriptions of accessible spatiotemporal regions as playing the role of canonical designators with respect to claims of physical existence includes offering a gloss on what someone is saying who denies that Pegasus (or any winged horse) existed, or ever flew over Greece. Negative existential judgments—claims to the effect that something or other does not exist—have been the source of considerable philosophical confusion over the years. (A paradigmatic example is the doctrine that nonexistent entities must at least have some sort of 'subsistence' for us to be able to refer to them in denying that they exist.) But understanding existential claims as expressing substitutional commitments with respect to a class of canonical designators yields a straightforward reading of the sense of such negative existentials. Denying numerical existence to the largest integer, or physical existence to Bellerophon's flying horse, is committing oneself to something incompatible with *all* of the identities, one side of which is a canonical designator of the

relevant sort and the other side of which is the expression 'the largest integer' or 'Bellerophon's flying horse'—or expressions anaphorically dependent on them.

On this relaxed account, there is no reason to boggle at claims that numbers or other abstract objects exist. One must insist only that a determinate sense have been given to such claims, by specifying the relevant class of canonical designators. Once it has been settled that a class of expressions functions as singular terms, if some of the sentences in which they have primary occurrences are true, then it is a criterion of adequacy on the specification of a class of terms as canonical designators relative to such claims that there be canonical substituends for them. Where those conditions are satisfied, corresponding existential claims have been given definite sense and are themselves true. (Of course, how *interesting* they are is another matter.)[67]

III. SUBSTITUTION, TOKEN RECURRENCE, AND ANAPHORA

1. *Inference, Substitution, and Anaphora*

Discursive practice has at its center the game of giving and asking for reasons; social deontic scorekeeping practices confer specifically *propositional* contents because they are *inferentially* articulated. The previous chapter examined the *substitutional* fine structure discernible within that inferential articulation—the substructure in virtue of which subsentential expressions can play an indirectly inferential role by making a systematic contribution to the propriety of inferences in which the sentences they occur in serve as premises or conclusions. That claims are articulated according to substitution inferences in turn presupposes a further level of structure. For substitution is not definable for individual unrepeatable expression tokenings. It requires some notion of token *repeatability*.

To begin with, in order to be available as reasons, sentences used to express claims must be at least in principle repeatable—both within and across interlocutors. Furthermore, for an expression to be used as a singular term (or predicate), it must be possible for it to occur in different sentences, combined with different predicates (or terms).[68] The definition of substitution inferences requires that occurrences of the same expression, whether term or predicate, be identifiable in both premise and conclusion. The first section of this chapter showed how Frege's triangulation principle (according to which purporting to pick out a definite object depends on its having been settled what would count as recognizing that object as the same again) can be understood in terms of the substitutional construal of what it is to use an expression as a singular term. That discussion simply presupposed the availability of repeatable terms.

When analytic focus is sharpened from repeatables such as 'Benjamin Franklin' and 'the inventor of bifocals' to particular tokenings of such types

as 'this man' and 'it', however, more must be said about how such unrepeatables are grouped together into term repeatables. The structure that emerges as crucial to generalizing substitutional considerations so as to encompass expressions of this sort is just that appealed to in Chapter 5 in explaining the use of the traditional semantic vocabulary 'true' and 'refers'—namely *anaphora.* The rest of this chapter investigates this phenomenon, adding a final level of semantic analysis to the two considered already. The result is a tripartite theoretical semantic structure whose key concepts are *inference, substitution,* and *anaphora.*

2. *Substitution and Repeatability*

What sort of relations do the repeatable terms, predicates, and sentences that have been under discussion thus far stand in to the *un*repeatable tokenings that instantiate them in the actual performances that are accorded pragmatic significance by discursive scorekeepers? What is it for two tokenings to be occurrences of the *same* term or sentence? Up to this point the examples considered have done nothing to discourage the supposition that those repeatables are just lexical *types:* equivalence classes of lexically cotypical tokenings. But cotypicality is neither necessary nor sufficient for tokenings to be occurrences of the same term in the semantically relevant sense of 'same term'. It is not sufficient because cotypicality cannot guarantee the correctness of substitution inferences of the form:

> This organism is a mammal,
> therefore this organism is a vertebrate.

For the two tokenings of the type ⟨this organism⟩ may involve different demonstrations, and so be governed by different material substitution-inferential commitments. In that case they are not guaranteed to be coreferential (= intersubstitutable), as different (primary) occurrences of a single term are.

Cotypicality is not necessary for tokenings to be occurrences of the same term in the semantically relevant sense of 'same term' because the relation between the predicates *does* guarantee the correctness of substitution inferences of the form:

> This organism is a mammal,
> therefore it is a vertebrate.

For the tokening of the pronoun is guaranteed to be governed by the very same material substitution-inferential commitments as the tokening that is its anaphoric antecedent. In the sense required by the discussion of the previous chapter, the latter of these is a proper substitution inference, and the former is not. The latter involves what are in the semantically relevant sense two occurrences of the same term, while the former does not.

These examples emphasize that the idea of substitution inferences (and

hence the assimilation of sentences as substitutional variants of one another) presupposes that repeatable expressions can *recur*—occurring now in one context, now in another. This makes no sense applied directly to unrepeatable tokenings. One might, of course, literally cut a token out of one written context and paste it into another, but this is just the sort of case that points to the need to focus on token*ings* or particular unrepeatable *uses* of tokens, rather than on the tokens themselves. For the same token may be used to perform various different speech acts at different times. One sign saying "Dig here" may be moved from place to place on the campus during a treasure hunt; the different tokenings of 'here' have different referents (for instance, they are intersubstitutable with different definite descriptions of locations) at different times, even though just one token is involved. Token recurrence may be determined by cotypicality, but that is not the only structure of token recurrence. Indeed, as will emerge, if what it is for a term to have a cotypicality recurrence structure is for *all* tokenings of a given type to be tokenings of that same term, and hence be guaranteed to be coreferential, then *no* sort of expression has such a recurrence structure—not even proper names or definite descriptions.

Practical acknowledgment or attribution of expression recurrence (treating some tokenings as tokenings of the same term or sentence) is an attitude that is implicit in the adoption of substitutional, and hence of inferential commitments. A particularly vivid reminder of the implicit presupposition of the reidentifiability of *terms* that stands behind the substitutional reidentifiability (and hence identifiability) of *objects* is provided by consideration of the way identity claims work in making substitutional commitments explicit. Identity locutions make explicit the claim that two terms pick out the same object. Their defining use is accordingly in explicit substitution inferences of the form:

Φa,
$a = b$
therefore Φb.

The correctness of an inference of this form depends on the tokening of the type ⟨a⟩ occurring in the first premise and the tokening of that type occurring in the second premise being tokenings of the same term, and similarly for the tokening of type ⟨b⟩ occurring in the second premise and in the conclusion. Clearly identity locutions cannot be used to make *this* sort of implicit presupposition explicit in turn, on pain of an infinite regress. For the use of identity locutions to license substitutions presupposes the possibility of tokening recurrence between elements of identity claims and elements of the sentences substitutionally governed by them. Of course, like the substitutional commitments made explicit by identity locutions, commitments to treating some tokenings as recurrences of others can be made explicit if suitable further vocabulary is introduced. If the two premises and

the conclusion of the inference above are denominated (i), (ii), and (iii), respectively, then the metalinguistic apparatus for talking about tokenings that was introduced in Section IV of Chapter 5 can be used to do just that.

The first tokening of type ⟨a⟩ can then be picked out as $/a/_{i}$, the second as $/a/_{ii}$, and the claim that they are tokenings of the same term (not just of the same type) expressed by saying "$/a/_{ii}$ is a recurrence of $/a/_{i}$," abbreviated as Recur($/a/_{i}$, $/a/_{ii}$). The reason the previous substitution inference is a good one is then that Recur(/This organism/$_{premise}$, /it/$_{conclusion}$), and the reason the prior inference need not be a good one is that it can happen that ~Recur(/This organism/$_{premise}$, /This organism/$_{conclusion}$). So it is possible to express the sort of inference that is licensed by explicit identity claims by saying that the inference from

Φa, and
$a = b$, to
Φb

is a good one provided that Recur($/a/_{i}$, $/a/_{ii}$) and Recur($/b/_{ii}$, $/b/_{iii}$). Of course, such a further explicitation is of use only insofar as one can implicitly recognize the different occurrences of 'a', '$/a/_{i}$', and so on as recurrences of the same expressions. It is not in principle possible to use explicit stipulations to eliminate the need for reliance on implicit capacities to recognize recurrences.[69] For this reason, an implicit token recurrence relation is appealed to in what follows (much as an implicit substitutional variation relation was appealed to in the previous chapter).

3. *Token Recurrence*

There are two varieties of substitutional equivalence. These are *intra*term and *inter*term, or de jure and de facto equivalences of tokenings. The former are (taken to be) binding on all interlocutors; the latter vary from doxastic repertoire to doxastic repertoire, according to the particular substitutional commitments undertaken by or attributed to an individual. Each attributor takes recurrence to bind all, in keeping track of significances of identificatory commitments and invocations of them by term use. But the identificatory or substitutional commitments themselves vary from individual to individual.

Substitutional structure requires both sorts. They cannot be defined separately, apart from their role in such a structure; one cannot have the one sort of equivalence without the other. What the intraterm equivalences are *for* is to be the vehicles of interterm substitutional commitments. These latter in turn presuppose them, in that they could not otherwise have content. The model of invariance of something under substitution involves changing something and preserving something else. The changes that do preserve the appropriate something define interterm equivalences.

The notion of recurrence or repetition without change that is presupposed by (is part of, as complementary to) the notion of change invoked in that definition is intraterm or de jure intersubstitutable equivalence. The fact that tokenings can have deictic and anaphoric significances means that the notion of recurrence presupposed by substitutional (and so inferential) relations cannot be reduced to that of lexical cotypicality. This raises three fundamental questions. First, once this issue has been distinguished from that of their being of the same lexical type, what does it mean to say that two tokenings are tokenings of the *same* term or sentence? Second, granting that expressions may exhibit a recurrence structure that can be represented by an equivalence class of cotypical tokenings, what *other* sorts of recurrence structure are there? Third, what difference does it make to the expressive power and function of a repeatable term (or expression of another grammatical category) *which* sort of recurrence structure it exhibits? That is, what sort of expressive impoverishment does an idiom suffer from if it lacks one or another of the different sorts of recurrence structure?

The terms 'term' and 'sentence' are usually thought of as picking out items that are lexically individuated—that is, by character rather than content, in Kaplan's typology. The *semantically* relevant recurrence classes of tokenings correspond not to this sort of repeatable but to what particular tokenings *express*. At the propositional level, it is possible to talk about different sentence tokenings that (in different contexts) express the same *claim*. It is in this sense, of shared content rather than character, that tokenings of "You are tired," uttered by me, and "I am tired," uttered by you, can under suitable circumstances be said to express the same claim. In the next chapter, where same-claiming is investigated in more detail in connection with the use of the 'that' clauses that specify the propositional contents of explicit ascriptions of propositional attitude, this semantic equivalence relation over sentence tokenings is grounded in practical proprieties of *deference* (see 8.4.3). This is a central structure of giving and asking for reasons as a *social* enterprise—of assertion as making premises available for *others* to use in inferences.

The semantically relevant recurrence classes of term tokenings considered here similarly correspond not to the characters of terms as lexical items but to the contents their tokenings express. Awkwardly, there is in ordinary philosophical parlance no generally agreed-upon expression that stands to 'term' as 'claim' or 'propositional content' stands to 'sentence'. 'Individual concept' and 'singular sense' are sometimes used, but each carries theoretical baggage it is best not to import here. The more neutral notion of the *conceptual contents* expressed by term use will be used here to indicate what is wanted. The use of such substantives is misleading in the present context in any case, for it suggests that what links the tokenings that make up the recurrence class of a term is their *shared* relation to some one content. This presupposes that recurrence classes are equivalence classes, and that is im-

portantly not always the case. It is best instead to look at how recurrence classes work and then to adjust talk of contents to fit their expressive function.

For an expression to be used as a singular term is for the inferential potential of sentences containing it to be determined in part by a set of material substitution-inferential commitments linking that term to others. This principle obviously applies to *repeatable* expressions, the tokenings of which can recur from sentence to sentence. Sharpening the focus to the level of unrepeatable tokenings requires that the idea of one term appearing in different sentences—presupposed by its being linkable to others by substitution-inferential commitments—be translated in terms of tokenings that count as recurrences of other tokenings. An unrepeatable tokening can occur in only one sentence tokening. So what it is for an unrepeatable *tokening* to have the significance of an occurrence of a singular term must be cashed out in terms of the inferential potential of the sentence tokening in which it occurs, and of other sentence tokenings in which tokenings occur that qualify as recurrences of the original tokening. For two tokenings $/t/_i$, $/t'/_j$ to be linked by a substitution-inferential commitment means that if $/\Phi t/_k$ is any sentence tokening containing a singular term tokening $/t/_k$ that is a recurrence of $/t/_i$, what it expresses has an inferential consequence that can be expressed by a sentence tokening $/\Phi t'/_l$, where $/t'/_l$ is a recurrence of $/t'/_j$. Thus substitution-inferential commitments should be thought of as linking tokening recurrence structures. For a tokening to be used or treated as an occurrence of a singular term is accordingly for its significance to be determined by SMSICs relating the recurrence structure to which it belongs to the recurrence structures to which other tokenings belong.

These recurrence structures *may* be equivalence classes of term tokenings all of which share a single lexical type. In that case a substitution-inferential commitment linking an equivalence class of tokenings of the type ⟨Benjamin Franklin⟩ and an equivalence class of tokenings of the type ⟨the inventor of bifocals⟩ licenses inferences from what is expressed by tokenings of the form $/\Phi$(Benjamin Franklin)$/_i$ to what is expressed by tokenings of the form $/\Phi$(the inventor of bifocals)$/_j$, and things look just the way one would have expected from the discussion in the previous chapter. But recurrence structures not only need not be restricted to lexically cotypical tokenings; they need not be equivalence classes at all. Given the role that recurrence structures of tokenings play in defining substitution-inferential relations, recurrence must be a reflexive relation; each tokening must (trivially) count as a recurrence of itself. After all, the significance of each tokening is guaranteed in advance to be governed by the same set of substitutional commitments that it is itself governed by. More substantively, it must be a transitive relation: if $/t/_k$ is a recurrence of $/t/_j$ and $/t/_j$ is a recurrence of $/t/_i$, then $/t/_k$ is a recurrence of $/t/_i$. For from the point of view of substitution, to say that $/t/_k$ is a recurrence of $/t/_j$ is to say that $/t/_k$ inherits the SMSICs determining its significance from

$/t/_j$. If $/t/_j$ in turn inherits the SMSICs determining its significance from $/t/_i$, then $/t/_k$ inherits the SMSICs determining its significance ultimately from $/t/_i$.

But it does not follow from the role of recurrence relations in defining substitution-inferential commitments that recurrence must be a *symmetric* relation. $/t/_j$ may inherit the SMSICs determining its significance from $/t/_i$ without its being the case therefore that $/t/_i$ inherits the SMSICs determining its significance from $/t/_j$. If the inheritance runs in either direction, both tokenings have their significance determined in the end by the *same* SMSICs. Nonetheless, as will become clear, the direction of inheritance can make a difference in counterfactual situations: if the SMSICs governing $/t/_i$ *were* different, so *would* those governing $/t/_j$ be, but not in the same sense vice versa.[70] Where the recurrence structure is not symmetric, it has the form of a *chain* or *tree*, rather than of an equivalence class.

4. Anaphoric Recurrence Structures

This asymmetric structure of recurrence (inheritance by one token of the substitution-inferential potential of another) is *anaphora*. For one tokening to be anaphorically dependent on another is for it to inherit from that antecedent the substitution-inferential commitments that determine the significance of its occurrence. Different interlocutors may disagree about what those commitments are, but to take it that there is an anaphoric connection is to take it that the use of the anaphoric dependent is correctly evaluated according to whatever substitutional commitments govern the use of its antecedent. One cannot settle the substitutional commitments determining the significance of the occurrence of a singular-term tokening without computing the recurrence class to which it belongs.[71] If one takes a certain tokening to be an anaphoric dependent, attributing a determinate substitutional significance to it accordingly requires identifying some other tokening as its antecedent.

Thus

> (p) Carlyle wrote his brilliant satire of Hegel, *Sartor Resartus*, in part to show that he was an important thinker

has two readings, depending on whether $/he/_p$ is taken to be anaphorically dependent on the tokening $/Carlyle/_p$ or on the tokening $/Hegel/_p$. Those who understand the claim as involving the latter anaphoric commitment thereby take it to entail and be entailed by its substitutional variant

> (p′) Carlyle wrote his brilliant satire of Hegel, *Sartor Resartus*, in part to show that Hegel was an important thinker,

for both of these tokenings of ⟨Hegel⟩ are, as $/he/_p$ is on this reading, recurrences of $/Hegel/_p$. It does not follow from this account, however, that dox-

astic commitment is always preserved by "replacing pronouns by their antecedents."

Whether it is or not depends precisely on how that phrase is understood. If "replacing a pronoun by its antecedent" means generating a substitutional variant in which another tokening of the same *type* as the antecedent is put in the place originally occupied by the pronoun, the principle is false. It is true only if it means generating a substitutional variant in which a tokening that is a *recurrence* of the antecedent is put in the place originally occupied by the pronoun. Thus

> (q) An influential British author wrote his brilliant satire of Hegel, *Sartor Resartus,* in part to show that he was an important thinker

is not, even on the first reading of the anaphoric commitment it involves, equivalent to

> (q′) An influential British author wrote his brilliant satire of Hegel, *Sartor Resartus,* in part to show that an influential British author was an important thinker.

For the latter, but not the former is entailed by

> (q″) An influential British author wrote his brilliant satire of Hegel, *Sartor Resartus,* in part to show that his friend John Stuart Mill was an important thinker.

In light of the discussion (in Chapter 5, Section IV) of definitization transformations in connection with Chastain's similar examples, it is clear that what is wanted is rather

> (q‴) An influential British author wrote his brilliant satire of Hegel, *Sartor Resartus,* in part to show that *the* influential British author [or *that* author] was an important thinker.

Though grammatically a definite description, ⟨the influential British author⟩ is not semantically a definite description, failing as it does of unique designation; and /the influential British author/$_{q''}$ should be understood as anaphorically dependent on, and hence as a recurrence of /an influential British author/$_{q'''}$ (itself a tokening of a type other tokenings of which function quantificationally, rather than as singular terms). Similar remarks apply to the variant in which a tokening of ⟨that author⟩ is used as the dependent (again a tokening of a type other tokenings of which function differently, in this case deictically rather than anaphorically).

According to this account, understanding one tokening as anaphorically dependent on another is attributing (or in one's own case, acknowledging) a certain kind of commitment. In the language employed here, it is a commitment to the dependent tokening being a *recurrence* of the antecedent tokening. Recurrence commitments of this sort can be understood in terms of the

inheritance of substitutional commitments, which has already been discussed.[72] It should be emphasized that this is an account of what anaphoric relations consist in—or better, given the methodological phenomenalism about normative statuses that governs the theoretical idiom employed here, about the practical attitudes that constitute *taking* or *treating* two expressions as anaphorically linked. It does not pretend to address the questions about anaphora that linguists and cognitive psychologists have been most concerned with—namely, when it is *correct* to adopt this attitude and treat one expression rather than another as the anaphoric antecedent of another, what *lexical* or *syntactic cues* there are for adopting this attitude, or how audiences *in fact* go about deciding which of various possible readings to adopt. The question of interest here is what it is to do the trick—what counts as doing that trick—rather than when it is called for or how it can be brought off.[73]

The basic claim is that tokenings can have their recurrence classes determined in different ways. The recipe for calculating the recurrence class may be one that looks first to the lexical *type* of the tokening, including as basic recurrences other possible tokenings of the same type. Anaphora, by contrast, is a way of computing the recurrence class of a tokening by reference to another *tokening*. It is from that tokening that the anaphoric dependent inherits or borrows its recurrence class (which in turn determines the substitutional commitments that determine its significance). The first recurrence mechanism is *symmetric*. It results in a core class of cotypical tokenings, each of which is related to the others in the same way: by means of their type similarity. The second is *asymmetric*. It depends on the distinct roles played by the anaphoric antecedent and the anaphoric dependent tokenings. The distinction crucial to understanding the way in which the notion of *substitutional* structure makes it possible to extend the notion of *inferential* significance from the sentential to the subsentential level turned out to be the distinction between the *symmetric* substitutional commitments that govern the use of singular terms and the *asymmetric* substitutional commitments that govern the use of predicates. Just so the distinction crucial to understanding the way in which the notion of *token-recurrence* structure makes it possible to extend the notion of *substitutional* commitment from the level of repeatable types to that of unrepeatable tokenings turns out to be the distinction between the *symmetric* recurrence structure governing the use of expressions such as (some occurrences of) genuine definite descriptions and the *asymmetric* recurrence structure governing the use of anaphoric dependents.

5. *The Significance of Asymmetric Recurrence*

The asymmetry of recurrence reflected in the fact that interchanging the expressions playing the roles of anaphoric antecedent and dependent in general preserves neither the identity of claims nor doxastic commitment

pays substantial expressive dividends. Discursive practices that lack this structure are expressively impoverished in a variety of ways. The two most important concern the *empirical* dimension of discourse and its social or *communicational* dimension. Under the first heading, it is because they encompass noninferential reporting practices (with their attendant reliability structure of authority, as introduced above in 4.2–3) that languages can express empirical claims. Indexical or token-reflexive constructions, particularly deictic or demonstrative ones, play an essential role in such reports. But no language can have such constructions unless it also has asymmetric recurrence structures: *deixis presupposes anaphora.*

Under the second heading, the capacity to pick up another interlocutor's reference by using a pronoun is one of the central mechanisms by which communication is secured across the interpersonal gap created by differences in doxastic commitments (which induce differences in the inferential circumstances and consequences of application expressions are taken to have by different interlocutors). The significance of anaphora in this context is that it permits each interlocutor to produce utterances employing tokenings that have been stipulated to be recurrences of arbitrary tokenings by others. At both the level of sentences and the level of terms, such recurrence provides the basic points of contact between different repertoires of commitments (including inferential ones). The rest of this chapter is concerned with exploring these two sorts of expressive capacity that anaphora bestows.

The definition of anaphoric dependence allows many tokenings to be (treated as) anaphorically dependent upon the same antecedent. As the mention of the significance of anaphoric links among tokenings for securing communication among interlocutors suggests, the antecedent tokening may be uttered by someone other than the one who produces the tokenings that are anaphorically dependent upon it.[74] It is also possible for that antecedent itself to be anaphorically dependent on some prior antecedent. Since recurrence and inheritance of substitutional commitments is transitive, so is anaphoric dependence. It is in this way that anaphoric chains or trees are formed. They can be anchored or initiated by tokenings that are not themselves anaphorically dependent on other tokenings. These are *anaphoric initiators.*

This role can be played by expressions of two different kinds. Tokenings that acquire their substitutional significance symmetrically (recurring by type), such as some uses of definite descriptions and proper names, can serve as initiators; they are potential anaphoric antecedents that need not themselves in turn be anaphoric dependents. Also, indexical and demonstrative tokenings—which acquire their recurrent tokenings asymmetrically in the form of anaphoric dependents, without themselves being recurrences of any prior tokenings—can serve as initiators; indeed, if their occurrence is not to be cognitively and semantically idle, they *must* so serve. Being an anaphoric antecedent or dependent is a role that individual tokenings can play. Though

their lexical and syntactic types can suit them for those roles, tokenings are not *compelled* to adopt one or the other role by those types. It is crucial to the communicative function of anaphora that any term tokening (or sentence tokening—indeed any tokening that is not syncategorematic) whatsoever can function as an anaphoric antecedent.

Less obviously, with the exception of indefinite descriptions, there seem to be no expression types that preclude their tokenings from functioning as anaphoric dependents. That is, there are no other types *all* of whose tokenings function as anaphoric initiators—expressions *none* of whose tokenings can have the significance of anaphoric dependents. Definite descriptions, for instance, can be used anaphorically—typically as the result of definitization transformations of indefinite descriptions. Demonstrative constructions can also be used to form anaphoric dependents, as in:

> Kant admired Rousseau, but that writer admired only himself

and

> Fichte fought for political tolerance in Germany, and partly because of his efforts this precious end was eventually achieved.

Anaphorically dependent uses of proper names are discussed below in 8.5.6. Since indefinite descriptions always also have quantificational and predicative uses (as in "Carlyle was an influential British author"), there are *no* expression types *all* the tokenings of which are recurrences of one another—or indeed (for that reason) even all coreferential. There can, of course, be sets of cotypical tokenings all of which are recurrences of each other (and hence coreferential). But they all belong to recurrence structures that can also contain anaphoric recurrences of other lexical types, and they all exclude some (possible) tokenings of that type.

IV. DEIXIS AND ANAPHORA

1. Demonstratives

Substitutional commitments relate token repeatables. So unrepeatable tokenings must be sorted accordingly as some count as recurrences of others in order for any of them to have the sort of indirectly inferential significance in virtue of which their production can contribute to making a move in the language game. That sorting may be based only on the lexical or syntactic type instantiated by the tokenings, or it may be based on further features of the individual tokening. Within tokening-based approaches to explaining the significance in discursive practice of expressions and complex components of expressions, the two central phenomena appear to be anaphora and deixis.

In recent years demonstrative and indexical constructions have received a

lot of attention. Perry and Lewis and others have shown that such expressions are not eliminable in favor of expressions that recur by type, such as standard uses of definite descriptions (those that do not themselves contain indexical elements).[75] Evans has provided an elegant account of how the capacity to have demonstrative thoughts and the practical capacity to locate one's actions and perceptions simultaneously in egocentric and public space and time are interdependent.[76] The result of these discussions (and of the many others addressing the same topic) is a robust sense that deictic tokenings provide expressive resources that are essential to our conception of ourselves as empirically situated knowers and agents.

Discussions of the use of token-reflexive or indexical expressions usually do not include treatments of anaphoric dependency. Indeed, pronouns do not fit well with the paradigms that lead to this terminology for tokenings whose use depends on features other than their types. The idea behind talk of *token reflexiveness* is that what tokenings of types such as 'I' refer to depends on who utters the tokening, for types such as 'now' and 'here' it depends on when and where the tokening was produced, and so on. What a tokening of 'it' refers to indeed depends on features of the tokening itself (just which features depending on the type or character of the pronoun). But the relevant feature of the tokening is just its anaphoric antecedent, and this is not a feature that can be specified independently and in advance of semantic and pragmatic interpretation, the way speaker, time, and place can. (Indeed, whether it is to be counted a semantic or pragmatic feature on traditional ways of dividing up these issues is itself a nice question.) Similarly, thinking of pronouns as *indexical* demands that one be able to specify the index to which their semantic evaluation is relative, and once again the anaphoric relation of a tokening to an antecedent tokening is not happily assimilated to general indices such as speaker, time, and place, which are specifiable in nonsemantic terms. Anaphoric chains running through bits of discourse are not naturalistic features of them like which organism produces the tokening, or when or where it is produced. They are *normative* features attributed to the discourse by deontic scorekeepers, matters of conditional commitment or commitment inheritance—of the obligation that the significance assigned to, or score kept on, one part of the discourse answer in systematic ways to the significance assigned to, or score kept on, another.

It is interesting to consider the use of demonstratives in connection with this contrast. A common strategy is to assimilate demonstratives to the indexical paradigm. The semantic interpretant correctly assigned to a demonstrative tokening depends not only on the lexical type it instantiates but also on a further index, which is taken to function in a way analogous to the way the indices of speaker, time, and place function in the semantic interpretation of tokenings of 'I', 'now', and 'here'. That additional index in the case of demonstratives is the object indicated or demonstrated by the one using the

demonstrative. What sort of an index is this? Is what a speaker is indicating or demonstrating a naturalistic fact, specifiable in nonnormative, nonsemantic terms, as speaker, time, and place are? Or is it a normative, scorekeeping matter of how various commitments should be understood as related to one another, as anaphoric dependence is? Or is the dependence of the significance of demonstratives on what is demonstrated of yet a different sort, not to be assimilated to either sort of relation?

What is it to indicate or demonstrate an object? The idea that animates discussions of demonstratives in terms of the indexical paradigm is that the core phenomenon is *pointing:* the demonstrated object is the one pointed at. This idea encourages a picture of indication as a physical matter of picking out an object by extending the line formed by the knuckles of the index finger out until it intersects something opaque. Of course things are not so simple. Wittgenstein reminds us that even such a practice of pointing requires a great deal of social stage-setting—the untrained may be unable to transfer their attention beyond the tip of the pointing finger, or may perversely trace the line of indication in the wrong direction, from finger tip to base, and so take it that something behind the one pointing has been singled out. Again, he reminds us of the emptiness of 'bare' demonstration. The use of 'this' or 'that' must at least implicitly be connected with some sortal, for the same physical gesture can have the significance of pointing to a book or to its cover, its title, its color, its shape, and so on. Although everyone would concede that for reasons such as these the actual practice of pointing is a complex affair, the picture of the virtual line extending from the finger nonetheless exercises considerable force. It seems to promise a way of picking out objects by causal, rather than conceptual triangulation (though as the point about sortals indicates, one of the intersecting beams is quite hazy in the picture).

But this is a false promise. The one using a tokening of 'that pig' demonstratively in many cases need not *do* anything in order to have indicated or demonstrated one particular pig in the barnyard, provided that the unique salience of that pig has somehow already been established—whether through the efforts of the one employing the demonstrative or not. The requisite sort of salience is a motley; it can consist in distinguishing features of the perceptible environment, properties that are highlighted by their relation to psychological factors, conceptual stereotypes, background beliefs, the previous course of the conversation, and perhaps much else. As with anaphora, it is helpful to put to one side the difficult psychological question of how scorekeepers in fact determine salience and so often correctly discern what object is being indicated or demonstrated, and also the difficult conceptual question of what makes it *correct* to take or treat one object rather than another as the one being demonstrated. For the fundamental question concerns the scorekeeping attitude itself: What is it to *take* or *treat* some object *as* what is being indicated or demonstrated in connection with a demonstrative use

of some expression? What does taking or treating a singular-term tokening as having a demonstrative significance that ties it to a particular object consist in?

The function of demonstration is to pick out an object and attach a tokening to it—to settle that the tokening refers to that object (notice that the second 'that' is functioning anaphorically, not demonstratively in this sentence). In light of the general discussion of picking out or using terms to refer to objects, this means that the effect achieved by successful demonstration must be understood in terms of the *substitutional* role of the demonstrative tokening.[77] But a demonstrative tokening as such is unrepeatable; substitutional commitments govern repeatable expressions—those that can occur in more than one sentential context and that can be replaced in each context by others. Only accidentally and in very special circumstances would a speaker be in a position to repeat a demonstration of an object. Even where repeated demonstration is possible, it results simply in the production of another unrepeatable tokening, one that can be seen not to have the significance of a recurrence of the original by the fact that there is always the possibility that, unbeknownst to its author, it in fact picks out a different object. If demonstrative tokenings could not recur, then they could play no substitutional role, hence no inferential role, and so would be semantically and cognitively idle. They would in that case not be ways of picking out or talking about objects at all, but mere noises.

Of course the recurrence structure in virtue of which demonstrative tokenings can play a conceptual role is not far to seek. Demonstrative tokenings can be picked up anaphorically. Because they can have anaphoric dependents, demonstratives can figure in substitution inferences:

> *That pig* is grunting, so *it* must be happy. I'm glad, because *it* is our champion boar, Wilbur.

Anaphoric chains that the demonstrative initiates are available to figure in substitutional commitments and the inferences they govern just as repeatable term types such as ⟨Wilbur⟩ are. Because they are, uttering the demonstrative can be understood as contributing to making a move in the language game, in particular as indicating an object. It follows that the capacity of pronouns to pick up a reference from an anaphoric antecedent is an essential condition of the capacity of other tokens (which can serve as such antecedents) to have references determined deictically. Deixis presupposes anaphora. No tokens can have the significance of demonstratives unless others have the significance of anaphoric dependents; to use an expression as a demonstrative is to use it as a special kind of anaphoric initiator.

One might choose to assimilate the use of indexicals generally to deictic uses in this regard. I can make a claim by uttering a token of the type ⟨John should be leaving the house *now*⟩, only because I and another can later utter tokens of types such as ⟨If he had left *then*, he would have been at the

meeting on time⟩, in which the tokening of ⟨then⟩ should be understood as anaphorically dependent on the earlier tokening of ⟨now⟩.[78] The claim would be that in such a context 'then' functions much as it does in contexts such as

> Clothes were made differently in *the seventeenth century* than they are these days; *then* they were made by hand, while now they are made by machines.

The tokening /then/ here clearly is used anaphorically, with /the seventeenth century/ as its antecedent. It could be claimed that 'then' is *always* used in this way (though only at the price of allowing virtual or merely possible antecedent tokenings). In the same way, tokenings of 'you' could be understood as anaphorically dependent on (perhaps merely possible) tokenings of 'I' by others.

In these cases, however, by contrast to that of demonstratives, one can understand how recurrences of indexical tokenings are possible without invoking specifically anaphoric connections. The character ⟨I⟩, for instance, is systematically linked to ⟨you⟩ (and ⟨he⟩ or ⟨she⟩), in such a way that tokenings of the one can count as recurrences of tokenings of the other (so that the same content is expressed by their use) without that linkage having to be understood anaphorically. The way in which uses of ⟨here⟩ and ⟨now⟩ can be picked up by uses of ⟨there⟩ and ⟨then⟩ respectively seems to work in a way that is intermediate between the systematic interchange of speaker and audience that governs ⟨I⟩ and ⟨you⟩, on the one hand, and the demonstrative element that is the essence of ⟨this⟩ and ⟨that⟩, on the other. But in these cases, too, index-matching rules connecting the contexts of utterance can be formulated that specify which tokens of the one type count as recurrences of which tokens of the other. Nothing like this, though, is possible in the case of purely deictically significant tokenings. There recurrence can be understood *only* on the anaphoric model.

It has been argued that deixis presupposes anaphora. But if anaphora could be given a deictic analysis, then it would not follow that anaphora is the more fundamental phenomenon. Here one would look for a converse of the familiar cases in which apparently deictic or demonstrative expressions are actually playing the role of anaphoric dependents, as in "Kant was a Prussian Pietist, but *that philosopher* did not always think like one." It is hard to assess the promise of this idea in the absence of a detailed working-out of it. One idea would be to assimilate the relation between an anaphoric dependent and its antecedent to the relation between a demonstrative and what is demonstrated. Anaphoric dependents would be understood as indexical tokenings that referred to their antecedents. In a tokening /Hegel understood Kant's argument, but he did not refute it/$_i$, the token /it/$_i$ would be understood as meaning what a token /that/ would mean, if it could be arranged that what was demonstrated was the antecedent tokening /Kant's argument/$_i$.

But as it stands this cannot be how such a story would go. For in that case the tokening /it/$_i$ would be conceived of as intersubstitutable, not with other tokenings co-identified with /Kant's argument/$_i$, but with tokenings that (presystematically) would be said to refer to the *tokening* /Kant's argument/$_i$—such as /the very tokening of type ⟨Kant's argument⟩ that was just uttered (or tagged with the index *i*)/. That is not what Hegel understood but failed to refute, for he never heard of that tokening. Some mechanism would be required to get us from the demonstrated anaphoric antecedent as referent of the supposed demonstrative-anaphoric dependent to the referent of that tokening. It should be obvious from the discussion of Chapter 5 what sort of mechanism that is. It is precisely this job that defines some locution as meaning what 'refers' means. An operator that forms indirect definite descriptions of objects from direct definite descriptions (or other specifications) of term tokenings is just what 'refers' is. If the natural anaphoric analysis of such locutions is to be avoided (as pursuing this strategy would require), then an alternate demonstrative or more broadly deictic reading must be offered of 'refers' and locutions performing a cognate function, in order to offer a deictic analysis of other anaphoric locutions. If this is not forthcoming, then anaphora must indeed be seen as presupposed by but not presupposing deixis.

To summarize, then. The recent tradition has focused on the cognitive centrality and irreducible importance of deixis. Concern with "direct reference" has been developed in part by exploiting intuitions about the basic nature of the word-world link established by the demonstrative use of expressions—what we used to think of as the use of demonstrative expressions, before it became apparent that anything can be used demonstratively (just as anything can be used anaphorically). From this point of view it has seemed natural to make the point that linguistic significance is always the significance of a (possible) event or uttering, that is, of a tokening—which may then be connected to semantically relevant types according to various models, by insisting on the priority and irreducibility of unrepeatable demonstrative and indexical uses over descriptive repeatable ones. Anaphora, as another tokening-based phenomenon, has not seemed of the essence, for it deals only with intralinguistic continuations or preservations of something one must in principle be able to understand already without it. Thus if one cared about tokenings as well as types, it seemed natural to care about deixis first, and about anaphora, if at all, only later. It now appears, however, that such an attitude, natural as it seems, is not only strategically wrong-headed but actually incoherent.

2. *Deictic Mechanisms Presuppose Anaphoric Mechanisms*

Deictic uses presuppose anaphoric ones. One cannot coherently describe a language in which expressions have demonstrative uses but no pronominal uses (although the converse is entirely possible). For indexical

uses generally, like deictic ones, are essentially unrepeatable according to types. Different tokenings of 'this' or 'here' or 'now' are not in general recurrences of each other, or even co-identifiable. Yet it is only as repeatable (that is, as elements of recurrence classes) that they can be substituted for. Recurrence is presupposed by the possibility of substitution, and the possibility of substitution is presupposed by picking out occurrences as semantically significant (that is, as indirectly inferentially and so assertionally significant). Since deictic uses as such are not type-recurrent, that recurrence must be understood as token recurrence—in particular, as anaphoric.

In short, unless one could pick deictic uses up anaphorically to generate recurrence classes, one would not be able to involve such deictic tokenings in (undertaken or attributed) identificatory substitutional commitments, and so could not treat them as involving occurrences of singular terms. Without the possibility of anaphoric extension and connection through recurrence to other tokenings, deictic tokenings can play no significant semantic role, not even a deictic one. Deixis presupposes anaphora. Anaphora is the fundamental phenomenon by means of which a connection is forged between unrepeatable events and repeatable contents. No semantically significant occurrence of a subsentential expression can be discerned unless it is governed by substitution inferences, which requires token recurrence: no (semantically significant) *occurrence* without (the possibility of) *recurrence.*

One consequence of this claim is an argument that is similar in many ways to one of the most important arguments of "Empiricism and the Philosophy of Mind." There Sellars argues that it is not possible to conceive of a language consisting only of noninferential reports. Noninferential responsive reporting does not form an autonomous language fragment, does not constitute a set of practices one could engage in without also engaging in specifically inferential practices. Not every claiming or sentence tokening can be a noninferential reporting, the exercise of a reliable disposition to respond differentially to features of the nonlinguistic environment. For what distinguishes reports from any other response produced according to a reliable responsive disposition is precisely its possession of propositional content, that it means something that can be understood by the responder and by others. Possession of such content is in turn a matter of the inferential significance of the reporting response. Understanding reports requires being able to distinguish what further claims follow from them and what claims would provide evidence for them. This inferential articulation, which is the possession of conceptual content by the responses produced noninferentially, presupposes the possibility of making the inferences involved, drawing the conclusions or offering the justifications indicated. These will not be reporting uses of sentences, since they are arrived at inferentially. So not all claimings can be reportings, if even the noninferential language entries are to have the significance of reportings.

The conceptual dependence of deictic mechanisms on anaphoric ones

suggests a way of extending this argument from the level of sentences to that of singular terms. Just as reportings as events are not autonomously significant (they depend for their semantic significance on the possibility of connection with other claimings by inference), so demonstrations are not autonomously significant (they depend for their semantic significance on the possibility of connection with other term tokenings by anaphora). The same sort of mistake that Sellars diagnosed in those who were tempted to see claimings acquiring their content solely through the circumstances in which they are appropriately produced as noninferential reports or *Konstatierungen*—a content that is then only reflected in (consumed or presupposed by) inference—can then be discerned in those who would assign a similar priority (presupposing the possibility of autonomous significance) to deictic mechanisms of "direct reference" in securing references that are then merely reflected in or preserved by (and so are consumed or presupposed by) anaphora. The phenomenon here is exhibited for terms, but the significant extension of Sellars's point is better thought of as one that moves from lexical-syntactic types to tokenings:

inference : anaphora
::
semantic significance of types : semantic significance of tokenings
::
reporting uses of sentences : deictic uses of terms.

Just as it is their potential for inferential involvements that makes sentence repeatables bearers of contents, so it is the potential for anaphoric involvements that makes unrepeatable tokenings bearers of contents.

Conceptual articulation, it was claimed in Chapter 2, is in the first instance, *inferential* articulation. *Sub*sentential expressions, it was claimed in Chapter 6, can be understood as *indirectly* inferentially articulated, and hence as conceptually articulated—in spite of the fact that they cannot play the role of premise or conclusion in inference—in virtue of the significance of their occurrence for *substitution* inferences. Unrepeatable tokenings, paradigmatically demonstratives, can now be seen to be *conceptually* articulated, for they can stand in *anaphoric* relations to other tokenings, and the chains thus formed can be involved in substitutional, and hence inferential, commitments. The use of a demonstrative may be elicited noninferentially as a response to an environing stimulus. What makes it a term referring to an object—rather than a mere conditioned response like "Ouch"—is its role as an anaphoric initiator of chains that can be the subjects of substitutional commitments. It is in virtue of those anaphoric connections that a demonstrative tokening can play a conceptual role.

Equipped with this thought, it is possible to address the question of what it is to take it that some particular object has been demonstrated (whether by an actual gesture or by implication). For now this attitude of taking it that

some particular object has been demonstrated appears just as a special case of taking the use of any expression at all to have picked out a particular object. It is to be understood in terms of the sort of cognitive, conceptual, or semantic triangulation discussed in Section I—recognizing an object as the same again by undertaking a substitutional commitment. Such commitments are made explicit in the form of identities employing expressions that are anaphoric dependents of the demonstrative tokening. The anaphoric relations may be implicit, as in:

> *It* is Wilbur (or the big one),

or they may involve explicit use of what might be called (after the terminology of Chapter 5) anaphorically indirect demonstrative descriptions, as in:

> *The pig he pointed to (indicated, demonstrated) saying 'that pig'*, is Wilbur (or the big one).

Both sorts of tokenings, /It/ and /The pig he pointed to (indicated, demonstrated) saying 'that pig'/, should be understood as pronouns, anaphorically dependent on a prior utterance /that pig/.[79] Thus the concept expressed explicitly by locutions such as 'what one is *pointing* to', like what is expressed by locutions such as 'what one is *referring* to', must be understood anaphorically. Thinking of referring in the way that has been developed here (by appeal to the concepts of inference, substitution, and anaphora), one could say that referring cannot be explained in terms of pointing, because pointing must be understood in terms of referring.

Anaphora has been presented here as a kind of token recurrence—a relationship among tokenings that is presupposed by, and hence not analyzable in terms of, substitutional commitments. Taking one individual's tokening to be anaphorically dependent on another is not attributing a substitutional commitment; it is attributing a more primitive sort of commitment, one that determines which substitutional commitments regarding other tokenings are relevant in assessing the substitutional significance of the one treated as anaphorically dependent. It may seem that this is an unnecessary shuffle, that a further level of analysis need not be broached. For anaphorically related tokenings are coreferential, and treating two expressions as coreferential has been explained in terms of the practical deontic scorekeeping attitude of attributing substitutional commitments. The reason such an account will not do has already been indicated: substitutional commitments govern the use of *repeatable* expressions. Anaphora is required to generate repeatables from unrepeatable tokenings, paradigmatically deictic ones, where cotypicality does not carry even a defeasible presumption of coreference, hence not of (co-)recurrence.

Identity claims make substitutional commitments explicit as the contents of possible assertions and judgments. It was pointed out above that the use of these explicitating locutions as inference licenses evidently depends on a

prior notion of term recurrence, since terms occurring on the two sides of the identity must be reidentifiable in the other premise and in the conclusion (respectively) of such inferences. It is only by means of anaphoric recurrence that deictic tokenings become accessible to identity claims at all:

> That cat is watching television.
> It is (=) the most spoiled cat in the room.
> Therefore the most spoiled cat in the room is watching television.

The tokening /it/ here is not replaceable with another tokening of the type ⟨that cat⟩, saving the goodness of the inference. The closest thing along these lines that will do is to replace /it/ with a tokening of the type ⟨that *same* cat⟩ or ⟨that *very* cat⟩. These look like deictic uses but are in fact anaphoric dependents, as tokenings of ⟨the same cat⟩ or ⟨the cat just mentioned (or pointed at)⟩ would be in this context. The function of such locutions is precisely to make available tokenings that are recurrences of the one they (anaphorically) refer to.

3. *Rigidity is an Anaphoric Phenomenon*

Recognizing that locutions such as 'that very *K*' play the role of operators that form anaphoric dependents sheds light on another aspect of demonstratives: their modal *rigidity*. Kripke introduced the term "rigid designator" to distinguish expressions (such as proper names) that pick out the same individual in all possible worlds, from those (such as definite descriptions) that do not. The pretheoretical phenomenon this theory-laden description addresses is the observation that even if Archie is the most spoiled cat in the room, Archie *might not* have been the most spoiled cat in the room (one even more indulged may just have stepped out for a snack); but it is not possible that Archie not have been Archie. Evaluated with respect to the actual situation, the claim that Archie is the most spoiled cat in the room is true. The coreference (that is, intersubstitutability) of these terms is a commitment actually undertaken by the evaluating scorekeeper. Evaluated with respect to other possible situations, the claim that Archie is the most spoiled cat in the room would not be taken to be true, since that description would then pick out another cat, say Ana; a different set of commitments undertaken on the part of the scorekeeper would entail that Ana, not Archie, is the most spoiled cat in the room. Against either background set of commitments, however, Archie would still be Archie. For the coreference of different tokenings of the proper-name type ⟨Archie⟩ is guaranteed by their being recurrences of one another. Definite descriptions are not rigid designators in that when modal contexts (which make explicit the inferential potentials of expressions when evaluated with respect to diverse sets of background commitments) are taken into account, different tokenings of the same definite description type

are not guaranteed to be coreferential. It follows that at least in these contexts, cotypicality of definite descriptions is not sufficient for tokenings to belong to the same token recurrence structure. Of course, cotypical tokenings of grammatically definite descriptions are not guaranteed to be coreferential in any case, because such descriptions have uses in which they play the role of anaphoric dependents; two tokenings of the type ⟨the author⟩ would not be expected to be coreferential if they were drawn from different anaphoric chains appearing in reviews of different books.

Rigidity is an anaphoric phenomenon. Instead of repeating the proper name, the modal claim above can be expressed by saying:

> Archie is the most spoiled cat in the room, but *he* (or *that cat*) might not have been the most spoiled cat in the room.

The pronoun has as its antecedent, and so is a recurrence of, the tokening /⟨Archie⟩/ that appears in the claim about how things in fact are. The anaphoric chain to which it belongs is then available to specify that same cat in other possible situations. The anaphoric chain, in other words, denotes rigidly.[80]

Kaplan introduces a rigidifying operator 'dthat' with the stipulation that while *!xDx* is not in general a rigid designator (since its denotation varies from world to world), dthat(!*xDx*) is to be rigid, picking out in each world whatever !*xDx* picks out in the actual world.[81] His operator in effect forms a type all of whose tokenings are stipulated to be anaphoric dependents of a tokening in the actual world. Thus 'dthat' does systematically what expressions like 'he' and 'that very cat' do informally in specifying other possible situations involving the same objects that are picked out in the actual situation by contingently associated expressions. The idea that rigidity has something special to do with the use of demonstratives—which lies behind the choice of 'dthat' (a homonym of the demonstrative 'that') to express the rigidifying operator—arises precisely because demonstrative tokenings can recur *only* anaphorically, and hence rigidly.[82] Thus someone who says

> *That bar* is the standard meter stick. If *it* had been heated, *it* would have been more than one meter long.

is using the anaphoric dependent to make the demonstrative recur, thereby displaying the same rigidity that would be made explicit in saying

> *That bar* is the standard meter stick. If *it* had been heated, *Dthat(that bar)* (or *Dthat(the bar the speaker indicated)* would have been more than one meter long.

Kripke's original discussion pointed to the distinction between the modal rigidity of proper names and the nonrigid behavior of the definite descriptions that often serve to fix the reference of those names.[83] In Kaplan's terms, proper names can be understood as rigidified, 'dthat'ed descriptions or dem-

onstratives. Understanding rigidity in anaphoric terms accordingly suggests that tokenings of proper names be themselves understood as anaphoric dependents—elements in an anaphoric chain that is anchored in some name-introducing tokening.[84] Causal-historical theories of proper names then appear as dark ways of talking about the sorts of anaphoric chains that link tokenings of proper names into recurrence structures. Concern about the nature of baptism or name introduction is concern about how this special sort of anaphoric chain or tree can properly be initiated, so that all the anaphorically linked tokenings it contains pick out one object. Concern about transmission of names is concern about what is involved in earlier tokenings being picked up as anaphoric antecedents by later ones. It is clear in these terms what is happening when different tokenings of the same lexical type are used as names of different people—how there can be more than one person called 'George' or 'Aristotle'. In such cases there are just multiple anaphoric chains; the multiplicity of people who can be referred to as 'George' is a phenomenon to be understood by analogy to the way in which many people can be referred to as 'she'. Investigations of the roles played by social-linguistic context and practices, or conventions, on the one hand, and individual name-user's intentions, on the other, in determining what previous uses a particular tokening ought to be considered beholden to should be understood as investigations of which anaphoric chain a particular tokening ought to be considered to be part of.

The claim is not that assimilating these questions about the use of proper names to scorekeeping questions about recurrence commitments (commitments regarding the inheritance of substitutional commitments) solves at a stroke all the questions that have vexed causal-historical theories. On the contrary, those questions, when transposed into the anaphoric framework, typically address issues about what determines when it is correct to treat one tokening as anaphorically dependent on another. The present discussion pretends to address only the issue of what it is for a scorekeeper to take one tokening to be anaphorically dependent on another, not the specific practices by which such scorekeeping attitudes are assessed as correct or incorrect. Nonetheless, seeing these issues concerning proper names as special cases of general issues concerning anaphoric links pays certain explanatory dividends. (Some of these are exploited below in 8.5.5–6, in the discussion of Kripke's puzzle regarding the behavior of proper names in ascriptions of propositional attitude.) For instance, it becomes clear that the issue of whether a name user remembers where the name was picked up from or is disposed to defer to some other individual more knowledgeable in its use is relevant to assessing uses of the name only insofar as they bear on the *commitments* that name user has undertaken; one may be obliged (according to a scorekeeper) by the circumstances under which one acquired a name even if one is ignorant or mistaken about them, and may be committed to defer to various authorities without being disposed to do so.

Fans of causal-historical theories of reference do not typically restrict these accounts to proper names. Natural-kind terms and many predicates are thought to function according to this model as well. Indeed, it is important to remember that the opposition between causal and descriptivist theories of reference that Kripke appealed to in his influential arguments against the latter arises specifically in addressing the reference of proper names. No one was ever a descriptivist about the reference of the predicates used to form descriptions—on pain of an obvious regress. Some other mechanism of reference was always envisaged for those predicates.

Thus it is natural to understand expressions such as ". . . is red" and ". . . has a mass of twelve grams" as having the denotations they do in virtue of their links to authoritative episodes of calling things red and measuring their mass in grams. According to the present suggestion, then, these expressions ought also to be understood as functioning anaphorically. Insofar as this is the right way to look at things, then, there is in fact only one primitive recurrence structure, namely the anaphoric one. Apparently type-recurrent expressions such as proper names and basic predicates should in fact be understood as having their tokenings linked by relations of anaphoric dependence. The relations between tokenings in these structures can be symmetric, by contrast to the asymmetry of paradigmatic pronoun-antecedent links, because and insofar as those tokenings owe allegiance to a common antecedent.

Such an approach, according to which *everything* works anaphorically, may seem to explain too much. If all these sorts of expressions have anaphoric recurrence structures, and such structures act rigidly in modal context, what room has been left for expressions that are modally flaccid, varying from context to context in the substitutional commitments determining their significance? Although basic predicates should be understood to recur anaphorically, compounds of them—in particular definite descriptions formed from them—need not. Although the recurrence structures determining the inheritance of substitutional commitments of their basic parts link them to antecedents in the actual world—that is why the answer to Lincoln's question "If we agree to call the tail a 'leg', how many legs would horses have?" is still "Four"—the existential and uniqueness commitments involved in the use of definite descriptions are in modal contexts evaluated with respect to alternative situations.

Archie, the cat in front of the television, does exist; but he, that very cat, might not have—even though another cat much like him might have been named 'Archie' in circumstances much like those in which Archie was baptized, and might be sitting in front of the television. Evaluated with respect to a different set of commitments undertaken by the one keeping score (that is, with a different set of claims being taken-true, and so functioning as the facts), a description might pick out a different object than it does according to the actual commitments of the scorekeeper. For what it is

intersubstitutable with depends on what other substitutional commitments are in play. This is compatible with there being an underlying anaphoric recurrence structure governing the use of the components of the description, which determines which of the substitutional commitments in play are in fact relevant to determining the proper use of the description in question.

The foregoing remarks are not intended to present a theory of the use of demonstratives, proper names, and definite descriptions. They are meant rather to suggest reasons for recasting the standard problems regarding such expressions into the terms of the common framework provided by the notion of anaphoric structures of tokening recurrence. According to the picture being presented, taking someone to have used an anaphorically dependent tokening in making a claim is attributing an *anaphoric commitment.* An anaphoric commitment is a commitment to treating the dependent tokening as a recurrence of the tokening that is taken (by the one attributing the commitment) to be its antecedent. Recurrence commitments, the genus of which *anaphoric* commitments are a species, are commitments regarding the *inheritance* of *substitutional* commitments. In the same way, *inferential* commitments, the genus of which *substitutional* commitments are a species, are commitments regarding the *inheritance* of *doxastic* or assertional commitments.

The story accordingly has three layers. At the top, sentences can be understood as propositionally contentful in virtue of their use in expressing *claims*—that is, assertional commitments. The key concept at this level is *inference,* for what makes the contents expressed *propositional* is the role of sentences in giving and asking for reasons. Inferential connections among claims are understood in turn pragmatically, in terms of *consequential* relations among the attitudes by means of which score is kept on commitments and entitlements to commitments—how attributing one commitment entails attributing others, precludes entitlement to others, and so on. At the next level, subsentential expressions can be understood as indirectly inferentially contentful, in virtue of the significance their occurrence has for the inferential involvements of the sentences in which they occur. The key concept at this level is *substitution,* for taking subsentential expressions to be contentful consists in distinguishing some inferences as substitution inferences, some inferential commitments as substitutional commitments. The substitutional structure of the inferences sentences are involved in is what the contentfulness of their subsentential components consists in. At the lowest level, unrepeatable tokenings (paradigmatically deictic uses of singular terms) can be understood as involved in substitution inferences, and so as indirectly inferentially contentful, in virtue of their links to other tokenings in a recurrence structure. The key concept at this level is *anaphora.* For taking an unrepeatable tokening to be contentful requires associating it with a repeatable structure of the sort that can be the subject of substitutional commitments. Anaphoric inheritance by one tokening of the substi-

tution-inferential potential of another does just that. The articulation characteristic of specifically *discursive* commitments is to be understood most broadly in terms of *inference,* the details of which require attention to *substitution,* the details of which in turn require attention to *anaphora.*

V. INTERPERSONAL ANAPHORA AND COMMUNICATION

1. Communication

Linguistic studies of anaphora typically distinguish between intrasentential anaphora and discourse anaphora.[85] The study of discourse anaphora is addressed not only to inter*sentential* anaphora, where an anaphorically dependent tokening occurs in a sentence different from that in which its antecedent occurs, but also to inter*personal* anaphora, where those tokenings are uttered by different interlocutors. The account of anaphora in terms of token recurrence—that is, as consisting in inheritance by one tokening from another of the structure that determines which substitutional commitments are relevant to its semantic assessment—applies to interpersonal (and hence intersentential) anaphora as well as to the fundamental kind of intrasentential anaphora. Indeed, certain features of that account stand out more sharply in the interpersonal case. Thus looking at examples of this kind a bit more closely provides an opportunity to clarify the sort of scorekeeping involved in attributing recurrence commitments, by specifying further the sort of inheritance they involve. More important, a new dimension of the expressive role of anaphoric connections among tokenings comes into play in the interpersonal case.

Anaphora as here conceived contributes two crucial sorts of expressive power to the idioms in which it is operative. First, as discussed in the previous section, anaphora makes possible the construction of repeatable expressions from unrepeatable tokenings. It is the mechanism by means of which unrepeatable tokenings are picked up and made conceptually—that is, ultimately, inferentially—significant. It is only because deictic and other indexical tokenings can recur anaphorically that their occurrence contributes to the inferential role played by sentences containing them, and hence that their occurrence can be counted as semantically significant at all; no occurrence without recurrence. Even where no overtly indexical expressions occur, this expressive capacity is crucial to the functioning of empirical languages.

Empirical languages are those that include noninferential reporting practices, and the authority of such reports is essentially tied to the particular unrepeatable tokening that is elicited as a response by the exercise of a reliable differential disposition. For such tokenings to have cognitive authority (for them to be available as expressing commitments to which interlocutors can be entitled and to which they can appeal in entitling themselves to further conclusions), their significance must be governed by substitutional

commitments. Such commitments relate recurrence repeatables, which in this case must be constructed anaphorically. Thus the report

The traffic light has just turned red

can serve as a premise from which to draw conclusions according to inferential patterns that can be made explicit in the form of such conditional principles as:

If it was red then, it will be green soon.

This expressive capacity is important even if it is the same interlocutor who makes the noninferential report and draws conclusions ("It will be green soon") from it; after the first glimpse the reporter may no longer be in a position to report the color noninferentially. But there is a second sort of expressive power anaphora contributes to discursive practice that arises *only* in the social context of inter*personal* communication of information.

For information (whether true or false) to be communicated is for the claims undertaken by one interlocutor to become available to others (who attribute them) as premises for inferences. Communication is the social production and consumption of reasons. So communication (giving and asking for reasons) involves the interaction of the *inferential* articulation of *contents* that is at the center of the *semantics* presented here and the *social* articulation of discursive *commitments* that is at the center of the *pragmatics* presented here. The nature and significance of this interaction of the inferential and the social dimensions of discursive practice is a large and important topic. It is the subject of the next chapter, which argues that the representational dimension of propositional contents is a reflection of the essential role played in their specifically inferential articulation by differences of social perspective—that is, differences between the point of view of the one who undertakes a commitment and the points of view of those who attribute it. As a result, the contents of the claims that are deployed monologically in intrapersonal reasoning in soliloquy must be understood as having been conferred by public practices of deploying claims dialogically in interpersonal reasoning in conversation. Meditation is made possible by disputation.

In advance of that fuller discussion of communication, the treatment here of the specifically social dimension of the expressive function of anaphora can only be preliminary. Nonetheless, some of the cardinal points are sufficiently detachable to be available already at this point. Interpersonal anaphora plays an important role in securing the possibility of communication across the doxastic gap created by the differing commitments of speaker and audience. The capacity of those in the audience to pick up a speaker's tokening anaphorically, and so connect it to their own substitution-inferential commitments, is part of what makes it possible for them to understand the speaker's utterance by extracting information from it. Anaphoric connec-

tions among tokenings that are utterances by different interlocutors provide a way of mapping their different repertoires of substitutional commitments onto one another—a structure scorekeepers can use to keep track of how each set of concomitant commitments relates to the others.

Such correlation of the substitution-inferential commitments (and hence doxastic commitments) undertaken by a scorekeeper with those attributed to others is a necessary part of the *interpretation* that is the uptake by a scorekeeping audience of some speaker's claim. It is an essential part of being able to use others' judgments as reasons, as premises in the scorekeeper's own inferences (even just hypothetically) to assess their significance in the context of those collateral commitments. Interpretation in this sense is necessary even in the case where all parties share a language.[86] The reason communication requires interpretation of this sort is twofold. First, speaker and audience typically have *different* sets of collateral commitments—if they did not, communication would be superfluous. Second, the inferential significance of a claim (what its consequences are and what would count as evidence for it) depends on what *auxiliary hypotheses* are available to serve as collateral premises. So differences in background beliefs mean that a remark may have one inferential significance for the speaker and another for each member of the speaker's audience.

2. *Frege and Kant on Fruitfulness*

This point reaches deep into inferentialist approaches to semantic content. It will be recalled that in the definition of conceptual content that opens the *Begriffsschrift,* Frege acknowledges the role of collateral commitments serving as auxiliary hypotheses: two judgments have the same content if and only if "all inferences that can be drawn from the first judgment when combined with certain other ones can always also be drawn from the second when combined with the same other judgments." The 'always' here signifies universal quantification over auxiliary hypotheses. It is not enough if there is *some* set of further judgments that yields the same set of consequences when combined with each of the candidates whose contents are being assessed. Such a requirement would obliterate distinctions of content, since for *any* two claims such a set of auxiliary premises can be found—in Frege's systems, for instance, any two claims have the same consequences when conjoined with a logical contradiction.[87] This quantification over possible sets of background beliefs accordingly is an acknowledgment that what follows from a claim depends on which further claims one is allowed to assume in extracting those consequences. (The dual point also holds, of course; for what constitutes evidence for a claim—its inferential circumstances, rather than consequences, of application—equally depends on the available auxiliary hypotheses.)

Frege makes more of this relativity of inferential significance to available

auxiliary hypotheses in the famous opening paragraph of "On Sense and Reference": "$a = a$ and $a = b$ are obviously statements of differing cognitive value [*Erkenntniswerte*]; $a = a$ holds *a priori* and, according to Kant, is to be labelled analytic, while statements of the form $a = b$ often contain valuable extensions of our knowledge and cannot always be established *a priori*. The discovery that the sun is not new every morning, but always the same, was one of the most fertile [*folgenreichsten*] astronomical discoveries." *Folgenreichsten* here is literally *richest in inferential consequences* (in what follows from it). The cognitive value of a statement is to be assessed by its inferential significance, by the difference that adding it to one's repertoire of endorsed judgments makes to what else one is committed or entitled to.[88]

This idiom, and in particular the understanding of analytic identities it expresses, is borrowed directly from Kant. In the section of his *Logik* entitled "Logical Perfection of Cognition as to Quantity,"[89] he says: "The magnitude of cognition may be understood in a twofold way, either as extensive or as intensive magnitude. The former refers to the extension of cognition and therefore consists in its volume and manifoldness; the latter refers to its content [*Gehalt*], which concerns the manifold validity [*Vielgültigkeit*] or logical importance and fruitfulness [*Fruchtbarkeit*] of a cognition, as far as it is considered as a ground for many and great consequences [*großen Folgen*] (non multa sed multum)." Here content is understood in terms of fruitfulness in the sense of leading inferentially to many consequences. Kant's definition of analyticity to which Frege is appealing is similarly couched in these terms:

> The identity of concepts in analytic judgments can be either explicit [*ausdrückliche*] (explicita) or non-explicit [*nicht-ausdrückliche*] (implicita). In the former case analytic propositions are tautological.
>
> > Note: Tautological propositions are virtualiter empty or void of consequences [*folgeleer*], for they are of no avail or use. Such is, for example, the tautological proposition Man is man. For if I know nothing else of man than that he is man, I know nothing else[90] of him at all.
>
> Implicitly (implicite) identical propositions, on the contrary, are not void of consequences or fruitless [*folge- oder fruchtleer*], for they clarify the predicate which lay undeveloped (implicite) in the concept of the subject through development (explicatio).
>
> > Note: Propositions void of consequences must be distinguished from propositions void of sense.[91]

Implicitly identical propositions have an expressive role—namely developing the content of a term by making explicit some of its inferential consequences, as in analytic claims such as "The oldest living mammal is a vertebrate."

The circumstances under which claims of the form $a = b$ "contain valuable extensions of our knowledge" are those in which, first, they are not

analytic (and hence a priori) in the explicative sense and, second, the knowledge they are extending includes further claims expressed using the terms *a* or *b*. For it is only in the presence of such auxiliary hypotheses that the identity licenses nontrivial substitution inferences. What it is for a claim to have a nontrivial cognitive value or content is accordingly defined by a *particular* quantification over possible sets of collateral commitments; there must be *some* context in which adding the claim has nontrivial inferential consequences. So the notion of content is being defined in terms of a more basic notion of the inferential significance of adding a claim to a set of antecedently endorsed claims. For a claim to have a nontrivial content at all is for the inferential significance of its endorsement to include nontrivial inferential consequences in *some* doxastic context, and for two claims to have the *same* cognitive or conceptual content is for their significances to comprise the same inferential consequences in *all* doxastic contexts. The primitive notion of inferential significance (of what follows from a claim and what is evidence for it) is explicitly relativized to a set of background claims—namely the set of those that are available as auxiliary hypotheses or collateral premises in extracting inferential consequences.

So even though Kant and Frege do not talk about the social dimension of inferential articulation, their elaborations of inferential conceptions of conceptual or cognitive content implicitly acknowledge that the inferential *content* of a claim manifests itself in different inferential *significances*—different claims counting as its consequences and potential evidence—from the perspectives provided by various sets of concomitant commitments. When Frege speaks of "extensions of our knowledge," he is comparing the perspective available before a claim is added to the repertoire of commitments we undertake (and take ourselves to be entitled to) with the perspective available afterward. What is a fruitful (inferentially significant) addition from one point of view may not be so from another. Though Frege does not discuss the consequences this observation has for understanding synchronic communication connecting different doxastic perspectives, looking at that case is helpful in understanding the diachronic cases he does appeal to.

3. *Quine, Communication, and Reference*

The underlying point is that what a given endorsement of claim commits one to, is entitled by, and is incompatible with depends on what else one is committed to, on what collateral information is available as auxiliary hypotheses for the inferences in question. Quine appeals to this Duhemian relativity of evidential significance to total evidential context in the closing sections of "Two Dogmas" to enforce a constraint on theoretical concepts of meaning (that is, claim content). Transposed into the idiom in use here, his holist argument for relativizing the meaning of a claim to the "total theory" of which it is a part is that:

—the meaning of a claim is what must be grasped to understand it, and

—what is understood must at least determine the inferential significance of endorsing what is understood, but

—what follows from a claim depends on what other claims are available as auxiliary hypotheses, so

—any difference in collateral commitments means a difference in inferential significance, hence meaning.

The fact that the inferential significance of endorsements is always and in principle relative to collateral commitments available as auxiliary hypotheses shows just what Quine wants it to, and thereby gives a definite sense to the claim that "the unit of meaning" is the whole theory or set of concomitant beliefs.

What effect does this relativity of inferential significance have for understanding *communication*? Quine does not explicitly raise this issue because he systematically waffles on the question of whether his "webs of belief" or "total theories" are individual or communal, whether we each have a different "total theory" or all share one. The account being unfolded here of the social-practical structure of inferential articulation—and hence of propositional contents—is one way of trying to take account of the motivations that push him now to talk one way, now another. This is an important issue in the context of an argument for not distinguishing changes of meaning from changes of belief; what is one to make of the consequence Harman extracts from this theory, that when I notice a cloud pass in front of the sun, the meaning of all my words changes? It must be granted that the noninferential addition of this new commitment alters (at least slightly) the inferential significance of all the claims that I do endorse, and all those I might. At the very least, conditionals of the form "If there is a cloud in front of the sun, then *p*" clearly would come in this way to have a different potential for transforming my commitments, and this would in turn alter the inferential significance of any claim that could appear as the consequent of such a conditional (or the conclusion of the inference it makes explicit)—and that is any claim whatsoever.

But must this alteration of the inferential *significance* different claims have for me be understood as involving an alteration in the inferential *content* they express? The view developed in the next chapter is one according to which the *inferential* holism that requires the pragmatic significance of doxastically endorsing a propositional content to be relativized to a repertoire of concomitant commitments must be understood in the context of a *social* holism. That social holism requires the grasp of the semantic content whose endorsement has such a significance to depend on scorekeepers' abilities to exploit relations among the different perspectives constituted by the different commitments undertaken by and attributed to those whose deontic scores

they keep track of. The significance for the understanding of communication of a holism that relativizes to a repertoire of background commitments either the inferential significance or the conceptual content of claims depends, like the significance of all commitments, on the auxiliary hypotheses that are available to serve as collateral premises in drawing inferential consequences from it. In this case, the background commitments it is important to be aware of take the form of a model of communication. Holism about inferential significances has different theoretical consequences depending on whether one thinks of communication in terms of *sharing* a relation to one and the same *thing* (grasping a common meaning) or in terms of *cooperating* in a joint *activity* (coordinating social perspectives by keeping deontic score according to common practices).

Communicating is naturally conceived of as *conveying* something. According to such a conception, before an episode of communication takes place only the communicating agent possesses what is to be conveyed; after successful communication the recipient possesses it as well. Overt performances serve as the *vehicles* by which what is communicated is transported from speaker to audience. In the Lockean version of this transportation model of communication, what is transferred is ideas—which are related to words by conventions, which are in turn reflected in the associations of the various interlocutors. Upon having an idea, the speaker associates a spoken or written word with it, and upon hearing or seeing the word, the audience associates the corresponding idea. Communication is a way for speaker and audience to achieve a shared idea.

The framework conception of communication as conveyance of something can be filled in by various particular notions of what is conveyed. Rather than ideas, it might be propositions, meanings, or information that speaker and audience are understood as sharing. Of course not every sort of performance that brings about a similarity between its maker and its taker provides a candidate for a conception of communication. The concept of communication involves that of *understanding*. What is to be communicated by an utterance is what its audience is to understand by it. What the producer of a meaningful performance has initially and what in the case of successful communication its consumers eventually acquire is something—a content or meaning determining the significance of the remark—that is understood by both parties. How the details of the conveyance model are filled in depends on a further conception of what it is to grasp or understand what is conveyed.

A problem arises if this commonsensical model of communication is combined with the inferentialist account of discursive practice in terms of deontic scorekeeping presented in these pages. According to that account, the fundamental communicative performance (making a claim) is acknowledging or undertaking a doxastic commitment. The sort of understanding or uptake of such a performance required for successful communication is for

the audience to figure that performance correctly in its score: to attribute the right commitment to the one making the claim. What makes the commitment a discursive commitment is its inferential articulation. In particular, the propositional contentfulness of a doxastic commitment consists in the material norms governing its role as premise and conclusion in inferences.[92] In the paradigmatic case of communicating by claiming, the audience's understanding of a claim must determine the inferential significance that adopting or believing that claim would have—that is, what one would be committing oneself to by endorsing it, what other commitments might entitle one to that endorsement, what other commitments are incompatible with it (and so preclude being entitled to such an endorsement), and so on. It is only insofar as the audience assigns some such significance to an utterance that a claim content is communicated or conveyed thereby.

Given the relativity of the inferential significance of a claim to the context of concomitant commitments available to serve as auxiliary premises, it follows that inferential significance is not preserved in communication—is not conveyed or transported from producer to consumer of communicational performances. For any difference in collateral commitment may involve a difference of inferential significance, understanding, and appropriated meaning. If I believe that Zoroaster is the sun and that its shining is his beatitude, then an utterance of "The sun is shining" means something different in my mouth than it does in your ears. If it is nonetheless possible for us to agree or disagree about that claim, that cannot be because it has the same significance for us. Inferential significance can be determined only relative to a total belief-set, so if what audiences understand must determine such significances, it cannot be independent of the context of collateral commitments. Since, as pointed out above, communication is superfluous in the case in which all commitments are shared (which alone would guarantee sameness of inferential significance), if inferential significances were what needed to be conveyed for communication to take place, communication would be impossible in all cases in which it was not otiose.

This is the line of thought that led theorists (such as Feyerabend[93]) who took Quine's inferential holism seriously to worry about the incommensurability of different theories or sets of commitments. Corresponding to the transportation model of communication is an accumulation model of progress—acquiring epistemic rights to more and more true claims fabricated out of a common stock of meanings or candidate belief contents. If these must at least include inferential significances, then since those significances can depend upon any collateral commitments, meanings are not shareable across theories, and so not establishable cumulatively as theory develops and changes. Given the very different background beliefs quantum theory has given us to govern our inferences involving the word 'electron', how can we so much as understand Rutherford's turn-of-the-century claim that electrons are particles with definite boundaries, orbiting atomic nuclei with definite

boundaries? Given that 'electron' meant something so different for him than it does for us, how is it possible for us to deny the very claim that he was making—as opposed to denying some *other* claim whose content appeals to *our* concept of electrons? 'Incommensurability' is the name given to this threat to our understanding of what communication is and how it is possible. It is a threat that arises for inferentialist approaches to meaning and understanding once the sensitivity of inferential significances to background beliefs is appreciated. Although this challenge is more pointed in the case of diachronic conceptual change in the history of scientific theories, the corresponding difficulty evidently confronts inferential role theories of the sort of content that is grasped and conveyed in synchronic, face-to-face, intralinguistic communication among interlocutors with different repertoires of doxastic commitments.

Quine himself quickly drew the conclusion that what matters semantically is not *meaning* but *reference*—what we are representing or talking about rather than just what we are saying about it. Although he does not put the point in the context of communication, this move reflects the realization that even if (in virtue of my Zoroastrian beliefs) the observation that the sun is shining means something different in my mouth than it does in your ears, you can still learn something from me that you can use in your own inferences—if and insofar as you understand me to be talking *about* the sun, and saying *of* it that it belongs in the class of shining things, that I am representing that thing as being in that class. Again, though Rutherford's many false background beliefs make his claim that electrons orbit around an atomic nucleus mean something to him that is unintelligible in the context of post-quantum-theory background beliefs, we can still understand him to have been talking *about* electrons (the same things we refer to) and to have been representing them as having certain properties and standing in certain relations. The information communicated consists in the purely extensional content of the claim made. A difference in inferential significance and commitment is compatible with identity of referential commitment and achievement. Where common reference of terms and extensions of predicates can be secured in spite of inferentially different employments, progress is comprehensible as talking about more and more objects, invoking more and more predicate-extensions, and coming to say more and more true things about those objects, for instance by classifying them under the predicate extensions. This is the conclusion that comes out of the debate between Feyerabend and Sheffler.[94] It is clearly the lesson that the early, realistic Putnam drew and conveyed to students and admirers such as Field, Boyd, and Devitt. It even sets the terms for Putnam's recantation of realism in "Realism and Reason." Thus an important motivation for the emphasis on semantic extensions—the referential dimension of discourse—can be found in the concern with making intelligible the possibility of communication.

This strategy gives up on contents as inferential roles, in favor of a differ-

ent sort of primitive. Inference can then be reinstated at two levels. First, some inferential proprieties can be read off of inclusion relations among the extensions of expressions. More important, recognition of the relativity of extensions to various elements of context yields the notion of intensions, as functions from indices to extensions. Such intensions are a more robust sort of content, which can be seen to be shared by speaker and audience in favored cases. Also, more finely grained inferential proprieties can be read off of inclusion relations among the sets that serve as domains and ranges of the intension functions. This is a two-leveled scheme, starting with extensions and ascending to intensions as functions defined on them. Inferential significances (the inferential potentials of particular claims in particular doxastic contexts) play no systematic role. Instead of inferential significances varying from speaker to speaker, there are extensions varying from possible world (together perhaps with other indices) to possible world.

4. Intensions

Analogy with this appeal to functions suggests that in the inferential case one might treat the inferential content expressed by a sentence tokening as a function, assigning to each repertoire of concomitant commitments an inferential significance. Such significances could be (crudely) thought of as ordered pairs of circumstances and consequences of application. The first element then might consist of sets of inferentially sufficient antecedent claims (those from which the claim in question can be inferred) and the second of a set of inferentially necessary consequent claims (those that can be inferred from the claim in question).[95] Since what is evidence for or commits one to a claim, and what it is evidence for or commits one to, depends on what background commitments are available as auxiliary hypotheses, inferential *contents* could then be thought of as *functions*. The content of each claim would be represented by a function that takes sets of concomitant background commitments as arguments and yields inferential significances as values.

The theoretical advantages of such a picture would accrue from taking inferential contents so construed as what is shared and communicated within a discursive community. Differences between the inferential significance that a claim has in the mouth of a speaker and the ears of an audience would then be compatible with a common understanding of what is being claimed. It would not then be necessary to concede the counterintuitive claim that the meanings of all one's words change (at least slightly) whenever one acquires a new belief. This explanatory advantage would be bought at a significant price, however. Unless the theorist is content with stipulative semantics—associating intension functions with expressions by fiat—an account must be offered of what it is about the way expressions are *used* that *confers* such contents on utterances and the states and attitudes

they express. What is required in this case is an answer to the question, What is it for an expression to be so used as to have associated with it one rather than another intension determining a function from doxastic context to inferential significance?

The generic difficulty with answers to this question stems from the very features that make an intensional response attractive in the first place. For functions of the sort in question are individuated so finely that it is hard to see how the use of an expression could determine that one rather than a slightly different one should be associated with it. In different forms this is the worry underlying Quine's rejection of intensions, Lewis's discussion of the relation of linguistic behavior to formal semantics for artificial languages in "Languages and Language,"[96] and Kripke's "finiteness" version of Wittgenstein's skeptical arguments concerning the underdetermination of use by meaning if meanings are conceived in standard ways. In each case the difficulty arises because one can in general construct a function that differs from a given one only for arguments that are in one way or another beyond the reach of behavioral dispositions. Where this is so it becomes difficult to see what is being envisaged (never mind how one could know that it is true) when it is said by the theorist that one rather than another of these behaviorally indistinguishable functions is nonetheless to be associated with a particular claiming.

This difficulty is particularly pressing in the case of communication across generations. Thinking of the communication of content in terms of shared intensions, functions from context of collateral commitments to significance, is most plausible as a response to worries about incommensurability for a synchronic linguistic community. Sharing intensions is speaking the same language in a strong sense. It is not clear how plausible such an account is in the diachronic case, where what is at issue is the possibility of incommensurability produced by conceptual change within a scientific tradition. Surely Rutherford or even Bohr did not and could not have shared the intensions contemporary physicists associate with such expressions as 'electron', 'mass', 'particle', and so on. It is not just that our views have changed substantially during the twentieth century, but that they have changed in ways unforeseeable by our conceptual ancestors of a few generations ago. It would require considerable argument to show that they had nonetheless used their expressions according to intensions that left room for all of our radical rethinkings, which could accordingly be represented just by differences in the context of assertional commitments with which each claim is conjoined. It is not possible to rule out such an approach a priori, but it is not surprising that it is hard to find a champion for an intensional transportation model of diachronic communication.

These cases provide one of the strongest motivations for adopting a different strategy: one that breaks with the conclusion Quine arrives at in "Two Dogmas" by distinguishing between a kernel of inferential (and perhaps also

doxastic) commitments that must be shared by those who count as grasping the content, concept, or meaning in question, and a shell of peripheral beliefs, which could differ without alterations of content. One privileges *some* of the inferences a concept is involved in as constitutive of it, treating the rest as warranted by collateral information. Grasping the concept then involves mastering only these essential inferences, and these are what interlocutors must share on pain of misunderstanding one another. Acquiring a new peripheral belief—for instance that a cloud now obscures the sun—would not then count as altering the concepts expressed by such words as 'sun'.

The difficulty faced by this approach is just the one Quine emphasized: saying what it is about the practices of using expressions that deserves to be characterized as treating some claims and inferences involving a concept as essential to it, and others as providing merely ancillary information about what it applies to. In constructing artificial languages, one might simply stipulate that some commitments are to be in the first class, while others are in the second. Even then one would be obliged to say how the proprieties of using expressions then differ depending on how particular commitments are classified. But insofar as this apparatus aspires to contribute to the analysis of natural languages or languages in use, those features of discursive practice that confer such a distinction of status between conceptual and merely empirical commitments must be specified. In the present context, one would need to explain in scorekeeping terms the different roles played by the practical attitudes of taking or treating commitments as conceptual and empirical. Of course, the fact that Quine can find no trace in our discursive practice of an analytic/synthetic distinction by looking at such candidate attitudes as treating as unrevisable or as a priori is hardly decisive. Other possibilities are not far to seek. (One that has not gotten the attention it deserves is Sellars's suggestion that the practical status that privileges concept-constitutive inferences is their counterfactual robustness.)[97]

5. *A Three-Leveled Approach*

Nothing rules out such a strategy, but it is not the one pursued here. The present account substitutes a three-leveled approach for the standard two-leveled one.[98] Instead of beginning with extensions and defining intensions as functions from indices (including possible worlds, which provide a background of endorsed claims serving as the facts) to extensions, the story begins with the inferential significances of claims. The theory then moves down, defining the extensional dimension of discourse in terms of substitution-inferential commitments. Those commitments in turn determine equivalence classes of expressions corresponding to what is represented—what is talked and thought *about*. Various features of the interpretive scorekeeping practices appealed to in this move *down* from perspectival inferential significances to *extensions* then make it possible to

move *up* from those significances to propositional or conceptual *contents* (corresponding in some ways to *intensions*), which systematically relate the distinct perspectives responsible for the different significances claims have to different interlocutors. What has been left out of the traditional formalism of extension and intension, from this point of view, is precisely the *interpersonal communicational* dimension. Yet it is this that gives the discerning of extensions and intensions its connection to discursive practice, and hence the explanatory role in virtue of which alone it can be appropriate to call what is discerned *semantic* correlates.

The way in which concern with what is talked about arises in the process of mapping the repertoire of commitments of an interpreted interlocutor onto the repertoire of commitments of an interpreting interlocutor is discussed in the next chapter. That chapter also seeks to explain the sort of perspectival propositional contents that coordinated scorekeeping practices confer. The paradigm of communication as joint possession of some common thing is relinquished in favor of—or modified in the direction of—a paradigm of communication as a kind of cooperation in practice. What is shared by speaker and audience is not a *content*-as-function but a scorekeeping *practice*. Contents as functions from repertoires to inferential significances can be seen as implicit in such practices, but the practice can retain its identity even though the functions implicit in it are different (at different times, and from different doxastic points of view).

For what is implicit can be made explicit in various, not always compatible ways. From each doxastic point of view on a speech act there can be a content common to the one undertaking a commitment and the scorekeepers attributing it, but what is taken to be shared may be different from the points of view of different scorekeepers. Thus inferential *contents* are essentially perspectival—they can in principle be specified only from a point of view. What is shared is a capacity to navigate and traverse differences in points of view, to specify contents from different points of view.[99] Explaining this capacity is explaining what it is to take or treat (understand or interpret) someone's remark as representing or being about one thing rather than another. So what appear theoretically as distinct moves *down* from inferential significances to *extensions* by assimilating expressions as intersubstitutable (= coreferential), on the one hand, and *up* from those significances to *intensions* by relativizing them to repertoires of background commitments, on the other, correspond to aspects of a single interpretive activity of understanding, grasping a meaning—the cognitive uptake of communication that is deontic scorekeeping.

The perspectival nature of propositional contents and the way in which their essential representational dimension emerges from communicative scorekeeping practice is approached in the next chapter by considering what is made explicit in *de dicto* and *de re* ascriptions of propositional attitude. The role of anaphora in securing coreference across differences in perspective

can be considered here as an introduction, however. Anaphora serves to link the equivalence classes of expressions that are intersubstitutable according to one interlocutor to the classes generated by the substitutional commitments of others. The need for such a mechanism arises in the interpersonal context because the speaker may have different substitutional commitments from the audience. If the speaker believes that the first postmaster general of the United States is the inventor of bifocals, and the audience does not, the inferential significance of the claim "The inventor of bifocals spoke French well" is different from their various perspectives. The question then arises how those in the audience can manage to have an attitude toward the same claim the speaker is making, can agree or disagree with *it,* rather than some variant of it that they associate with the same noises or inscriptions. Given that speaker and audience disagree about whether the claim is about the first postmaster general of the United States, how can they nonetheless secure a common topic of conversation in order to argue about whether or not *he* spoke French well?

This way of putting the question contains the answer. (Compare the way the fact that one can assert the modal nonrigidity of the description 'the first postmaster general' by saying "Benjamin Franklin was the first postmaster general, but *he* [the man just referred to] might not have been" points to the central expressive function played by anaphoric relations in understanding that phenomenon.) Use of an anaphoric proform implicitly *stipulates* coreference with the anaphoric antecedent upon which it is semantically dependent. Thus differences in the substitutional commitments that determine the propriety of inferences involving 'the inventor of bifocals' according to speaker and audience can be bracketed and a common topic of conversation secured by using a tokening that is anaphorically dependent on the speaker's tokening. To respond to the speaker by saying "*He* did not speak French well" is to disagree with the claim made, *whoever* the inventor of bifocals might turn out to be. Indeed, if more than the object referred to is in question, the claim can be affirmed or denied by using an anaphoric dependent on the whole sentence, rather than just picking up one of its singular terms: the audience can say "That is true" or "What you claim is false." Interpersonal anaphora achieves just the effect that matters for securing communication in the face of differences in collateral commitments.

The capacity to use a pronoun that anaphorically picks up another's tokening is also a cardinal component of another important ability, one whose cognitive significance is often underrated. For pronouns enable us to talk without knowing what we are talking about. Thus a speaker can come late into a conversation in which someone is already being referred to as 'he' and can jump in—continuing that conversation with a remark such as "If he did that, he deserves whatever he gets." The speaker may under such circumstances have no idea at all of who it is that is being talking about. The form in which the later claims are expressed nonetheless commits the speaker

anaphorically to their being about whoever it is the others were already talking about. That is, a scorekeeper will assess the doxastic commitment the latecomer has undertaken according to whatever substitutional commitments that scorekeeper takes to govern the antecedent of the anaphoric tokening 'he'. Anaphora is a mechanism that permits undertaking and attributing commitments concerning objects that one need not be able to specify (nonanaphorically) if challenged. Thus one is not obliged to know or accept the descriptions by means of which the utterer of the anaphoric initiator might pick out the subject with respect to which both are undertaking and attributing commitments.

6. *Speaker's Reference*

These cases of interpersonal anaphora show that one must be careful in thinking of anaphora as inheritance of substitutional commitments by one tokening from another. The anaphoric antecedent is what determines the substitutional commitments relevant to the assessment of the significance of its dependents. But in using a pronoun that is anaphorically dependent on a tokening uttered by another, one is not thereby bound by whatever substitutional commitments the other happens to acknowledge as governing that tokening. An interlocutor who disagrees with the speaker's assertion "The inventor of bifocals spoke French well" by saying "*He* did not speak French well" is not making an incompatible assertion by adding "And he was not the first postmaster general, either," even though the utterer of the antecedent of those pronouns *is* committed to the intersubstitutability (that is, coreference) of 'the inventor of bifocals' and 'the first postmaster general.' Although the divergence of perspective that makes the point evident did not arise in the case of intrapersonal anaphora,[100] the substitutional commitments to be inherited anaphorically by one token from another are assessed by the scorekeeper who attributes the anaphoric commitment, that is, who takes or treats the one tokening *as* anaphorically dependent on another. To take one tokening to be anaphorically dependent on another is to take it that it should be understood as governed by whatever substitutional commitments govern its antecedent.

Different scorekeepers may disagree about what these are, and they may disagree even with the ones producing the performances whose significance they are assessing. They may nonetheless all agree in attributing an anaphoric commitment, that is, in interpreting one tokening as being anaphorically dependent on (hence a recurrence of) the same antecedent tokening. A scorekeeper who takes it that the inventor of bifocals is the inventor of the lightning rod will take it that the first speaker claimed *of* the inventor of the lightning rod that he spoke French well, and that the second speaker claimed *of* that same individual that he did not speak French well. That is, a scorekeeper who undertakes such a substitutional commitment and attributes

that anaphoric commitment is obliged to take it that what the first speaker said is *true* just in case the inventor of the lightning rod spoke French well, and that what the second speaker said is true just in case he did not. For what a scorekeeper takes to be true is just what that scorekeeper endorses. The scorekeeping significance of attributing an anaphoric commitment is accordingly just that the significance of the dependent tokening is to be assessed according to the *same* substitutional commitments by which its antecedent tokening is assessed—whatever those are. Where the scorekeeper is concerned with when a given claim is true, it is the substitutional commitments that scorekeeper *undertakes* that matter, rather than those attributed to the utterer of the antecedent.

This is to say that according to a scorekeeper who undertakes a commitment to the intersubstitutability of 'the first postmaster general' and 'the inventor of the lightning rod,' one who asserts "The first postmaster general spoke French well" has thereby in a certain sense undertaken a commitment to the claim that the inventor of the lightning rod spoke French well. And this is true even in the case where the one making the original assertion would deny that the first postmaster general is the inventor of the lightning rod. The speaker is, according to such a scorekeeper, committed to that further claim just in the sense that what he has said is true if and only if the inventor of the lightning rod spoke French well. In this sense, what someone is committed to may (according to a scorekeeper) not only outrun, but even conflict with, what that interlocutor is prepared to acknowledge. The scorekeeper must keep two sets of books.

The necessity for this dual score follows from the fact that there are in principle two places a scorekeeper can draw auxiliary hypotheses from in extracting the inferential consequences of (and so the commitments consequentially undertaken by) a set of commitments some individual is taken to acknowledge. Those auxiliary hypotheses may be other commitments the individual acknowledges, or they may be commitments the scorekeeper *undertakes* (acknowledges), rather than attributing as acknowledged. Since these latter represent the *facts* (facts being just true claims), according to the scorekeeper, these latter consequences are those that, according to the scorekeeper, *actually* follow from the claims made (given how things really are), regardless of whether the one making the claims realizes that they follow or not. The relations between these two sets of books, and the way their interaction constitutes the representational dimension of propositional content, is the topic of the next chapter. As an introduction to the perspectival character of claim contents that is investigated there, it is helpful to consider the phenomenon of *speaker's reference* in terms of interpersonal anaphora.

What Kripke called "speaker's reference," by contrast to "semantic reference"—a distinction closely allied to Donellan's distinction between "referential" and "attributive" uses of definite descriptions[101]—is a phenomenon that depends on the possibility of taking up identificatory or substitutional attitudes toward a tokening that is not treated as functioning in a type-recur-

rent way. It is a matter of the significance (substitution-inferential potential) that an audience attributes or ought to attribute to a particular tokening, by contrast to the significance that would otherwise be associated with it on the basis of its type. What is fundamental is the way an audience interprets or keeps score on the tokening. Once what it is to *take* someone to be speaker referring is understood, it will be possible to understand what it could be for some situation conventionally to call for or warrant the use of this sort of interpretation. So the account is in terms of audience uptake, not what the speaker does or intends. Like the cases just considered, the situation in which an audience counts as treating someone as having "speaker-referred to someone other than the one semantically referred to" by a remark is always one in which the identificatory commitments in the vicinity of the recurrence class of the uttering to be interpreted that the audience attributes to the speaker are different from those that the audience undertakes itself.

Adopting the subsentential forms a bit so as better to accord with tradition, a case might go like this. The speaker, Fred, says, "The man in the corner with champagne in his glass is very angry." According to Wilma, in the audience for this remark, Fred claims that Barney is the man in the corner with champagne in his glass. So according to Wilma (that is, the commitments she *attributes*), Fred might just as well have expressed his claim by saying, "Barney is angry." But according to Wilma (that is, the commitments she *acknowledges*), Barney is the man in the corner with ginger ale in his glass, and the man in the corner with champagne in his glass is Nelson. (She takes it that Fred does not see Nelson and does not believe that he is in the corner at all.) Then we can say that, according to Wilma, Fred has speaker-referred to Barney and attributed anger to him but has semantically referred to Nelson and attributed anger to him.

These two different ways of interpreting the claim that Fred has expressed by his utterance correspond to assessing his assertional commitments with respect to the identificatory commitments that Wilma attributes to him, and to assessing those commitments with respect to the identificatory commitments that Wilma herself undertakes. It is essential that, according to Wilma, there is some expression that Fred could (compatibly with the commitment Wilma attributes to him) have used to semantically pick out his referent, in order that he be able to speaker-refer to it by another expression. For what Wilma is doing when she assesses his remark as true by taking him to have speaker-referred to Barney is treating his tokening /the man in the corner with champagne in his glass/ as an anaphoric dependent whose antecedent is another tokening that Fred *could* have used (and would have used had he realized that there was a dispute about the matter): perhaps a tokening of the type ⟨Barney⟩, or ⟨the man in the corner with bubbly liquid in his glass at whom I am looking⟩, or just ⟨that man⟩. Attributing speaker-reference rather than semantic reference is assessing the substitutional commitments a tokening owes its allegiance to anaphorically rather than by type.

Having this interpretive or scorekeeping strategy available is useful for

reasons of charity. Wilma can make more of what Fred says come out *true* (according to her) by taking some of his remarks this way. Charity of this sort is necessary only where there is a relevant difference in perspective between audience and speaker—that is, where it makes a difference whether the commitments taken to be available as auxiliary hypotheses in drawing inferential consequences from a claim are those *undertaken* by the scorekeeper or those *attributed* to the one whose performances are being assessed. When such a difference in social perspective becomes explicit in ascriptions of doxastic commitment (the fundamental propositional attitude), it appears as the difference between ascriptions *de re* and ascriptions *de dicto*.

Appendix: Other Kinds of Anaphora—Paychecks, Donkeys, and Quantificational Antecedents

Section III offers an account of the practical attitude a discursive scorekeeper must adopt in order to count as treating one tokening as anaphorically dependent on another in the most basic sense—that is, taking the substitutional commitments that determine the significance of the dependent to be inherited from those that determine the significance of its antecedent. This suffices to show how anaphoric relations can be introduced into or diagnosed in the simplified discursive practices described here. In actual natural languages, anaphora is an immensely complex phenomenon; many more sophisticated tropes have been built up around the asymmetric token-recurrence structures identified here as the fundamental anaphoric phenomenon. Discussion of these goes beyond the scope of the present project, but perhaps a few signposts are in order.

One important issue that is put to one side here concerns the thorny problem of paycheck cases. Understanding a sentence like

> (a) The man who gives his paycheck to his wife is wiser than the man who gives it to his mistress

requires treating the anaphor /it/$_a$ as replaceable by another tokening of the same type as its antecedent, /his paycheck/$_a$, even though these tokenings will be governed by different SMSICs (have different referents). /His paycheck/$_a$ will not have the same referent as /it/$_a$, any more than the two tokenings of ⟨the man⟩ do. So this sort of anaphora cannot be dealt with in terms of the establishment of token-recurrence structures. Nonetheless, it is clear enough how to understand this sort of lexically 'lazy' anaphora: /it/$_a$ is replaceable by another token of type ⟨his paycheck⟩, and the antecedents of the two tokens of type ⟨his⟩ are the different, noncoreferential tokens of type ⟨the man⟩. The hard question (which is important to linguists and for some projects in artificial intelligence) is not understanding the correct reading but

telling when that sort of reading is called for. It is how to tell when one ought to understand anaphoric dependence in terms of token-recurrence, and when it should be understood rather in terms of the sort of type-recurrence that paycheck cases demand. The explanatory task undertaken here is finished, however, when the differences between the two sorts of readings have been made clear.

Another large issue passed by here concerns anaphoric dependents whose antecedents are *quantificational* expressions. In the simplest cases, the interpretation of claims formed in this way follows from the general account of quantifiers; the anaphoric chains determine what count as substitution instances of particular and universal quantifiers, the significance of those substitution instances is determined by the token-recurrence model, and the significance of the quantificational claim is determined disjunctively or conjunctively by those instances. This is how anaphoric dependents on quantificational initiators should be understood when those dependents behave like the bound variables of the predicate calculus. As Evans has pointed out, however, not all anaphoric dependents on quantificational antecedents are happily assimilated to this model.[102] Thus on the most natural reading,

> John bought *some donkeys*, and Harry vaccinated *them*

entails that Harry vaccinated *all* the donkeys John bought, whereas the bound reading

> [some *x*: donkeys *x*] (John bought *x* & Harry vaccinated *x*)

requires only that there be some donkeys that John both bought and Harry vaccinated. In his excellent discussion, Neale points out further that

> John bought *exactly two donkeys*, and Harry vaccinated *them*.
> *Few politicians* came to the party, but *they* had a good time.
> *Just one man* drank rum, and *he* was ill.

entail that John bought exactly two donkeys, few politicians came to the party, and just one man drank rum, respectively—consequences that are lost on the bound reading.[103] As he concludes: "The upshot . . . is that among pronouns anaphoric on quantifiers we need to distinguish between those that function as bound variables and those that do not."[104]

The interpretation wanted has already been alluded to, in the discussion above of 'definitization transforms' in connection with Chastain's treatment of anaphoric chains (see 5.4.2).[105] The quantificationally unbound anaphoric dependents of quantificational expressions in examples such as those above go proxy for definite descriptions formed from their antecedents. So the sentences above make the same claims as:

> John bought *some donkeys*, and Harry vaccinated *those donkeys*.

John bought *exactly two donkeys*, and Harry vaccinated *those donkeys*.
Few politicians came to the party, but *those politicians* had a good time.
Just one man drank rum, and *the man who drank rum* was ill.

Here the quantificational antecedent determines the class of relevant substitution instances, and the significance of the clauses in which the anaphoric dependent appears is determined by that class. The bound cases differ just in that the clauses in which the anaphoric dependent appears also function to constrain the class of substitution instances with respect to which *both* clauses are evaluated. The difference between the bound and unbound cases accordingly corresponds to a difference in the order of application of the two operations of determining a class of substitution instances and making anaphoric connections. In the bound case, the anaphoric connections govern the inheritance of substitution-inferential significance by one clause from another within each quantificational substitution instance; in the unbound case, they govern rather the inheritance of a class of quantificational substitution instances by one clause from another.

Thus what might be called 'definitizing' anaphora—which governs the inferential significance of quantificationally unbound anaphoric dependents having quantificational antecedents—is another sort of sophisticated anaphora. It is distinct both from lazy, type-recurrent ('paycheck') anaphora and from the basic case of token-recurrent anaphora discussed in the body of the text. It is clear, at least in broad outlines, how such anaphora should be understood in the discursive scorekeeping idiom developed here. As before, the difficult task is formulating rules codifying when it is appropriate to adopt one sort of reading rather than another. As before, no stand is taken here on this difficult problem—and so none on Evans's suggestion that a pronoun anaphorically dependent on a quantificational expression behaves like a variable bound by it just in case the pronoun is c-commanded by the quantifier.[106] In the idiom suggested in the text, these issues are all taken to concern when it is appropriate to do the trick (construe the substitution-inferential significance of anaphoric dependence according to one model rather than another) rather than what it is to do the trick (keep score according to one reading rather than another), which is all that is of concern here.

Geach's original donkey sentence was

Any man who owns a donkey beats it.[107]

Here one does not want the definitized reading, for those who own two donkeys are being accused of beating both of them. The trouble is that apparently then 'a donkey' must be understood as expressing a particular quantifier relativized to the universal quantifier expressed by 'any man'. But

neither of the two ways of putting this relation in standard first-order quantificational language seems right. Unlike the original,

[every x: man x] [some y: donkey y] (Owns (x, y) → Beats (x, y))

is compatible with

Some man who owns a donkey does not beat it,

while

[every x: man x] ([some y: donkey y] (Owns (x, y)) → Beats (x, y))

is syntactically incoherent, having the second quantifier, which occurs in the antecedent of a conditional, binding variables that occur in the consequent of that conditional.

This causes a problem, however, only for those concerned to provide a uniform way of mapping quantificational expression-types in natural languages onto operators in the first-order predicate calculus. Those not concerned with rules determining when it is appropriate to interpret tokens of the type ⟨a K⟩ or ⟨some K⟩ one way rather than another can rest content with understanding Geach's donkey sentence as having the inferential role of

[every x: man x] [every y: donkey y] (Owns (x, y) → Beats (x, y)).

Further anaphoric dependents on these quantificational expressions may then act either as quantificationally bound anaphors or as definitized ones.

A final sort of example that deserves mention is Bach-Peters sentences, such as

A boy who was fooling her kissed a girl who loved him.
The pilot who shot at it hit the MiG that chased him.

The difficulty here is that the anaphoric chains cross; each dependent inherits its substitution-inferential role from an antecedent that inherits its role in turn from the original dependent. So these anaphoric circles do not settle what is to be counted as the anaphoric initiator. As has long been recognized, however, these surface forms are ambiguous; they have two nonequivalent readings, depending on which expression is treated as an initiator (which includes a dependent), and which as a dependent.[108] With definitization, quantificational cases such as the first example reduce to those involving definite descriptions, like the second, and the two readings of those are not far to seek. As Neale puts it:

> If 'the pilot who shot at it' is given wider scope, 'him' is bound and 'it' is D-type:

[the *x*: pilot *x* & [the *y*: MiG *y* & *y* chased *x*] (*x* shot at *y*)] ([the *y*: MiG *y* & *y* chased *x*] (*x* hit *y*)).

If 'the MiG that chased him' is given wider scope, 'it' is bound and 'him' is D-type:

[the *y*: MiG *y* & [the *x*: pilot *x* & *x* shot at *y*] (*y* chased *x*)] ([the *x*: pilot *x* & *x* shot at *y*] (*x* hit *y*)).[109]

What Neale calls 'D-type' anaphoric dependents are those to be interpreted by definitization transforms of their antecedents. Once again, there is no special problem with interpreting each of these readings in discursive scorekeeping terms, so long as care is taken to distinguish anaphora determining the inheritance of substitution-inferential significance within quantificational substitution instances from anaphora determining the inheritance of classes of quantificational substitution instances (in the case of definite descriptions, singletons). The defining symmetry of the Bach-Peters sentences ensures that in this case there is no residual problem of determining when one reading rather than the other is appropriate; the *only* task is making sense of the two readings.

In conclusion, although there are other sorts of anaphora in play in natural languages besides the one taken as fundamental in the discursive scorekeeping semantics, there are strategies available for making sense of them within the model as developed in the text.