

The Predicate View of Names

Aidan Gray

March 28, 2008

There is, and has for some time been, a consensus among philosophers interested in language that proper names in natural language are properly thought of on the model of individual constants in a formal language. This is to say, at least, that their semantic function is to pick out individuals. This as opposed to, e.g., intransitive verbs, whose semantic function is to pick out properties of individuals (let's say) or quantifiers phrases whose function is to pick out properties of properties (let's say). There has been disagreement as to the details of this approach — for instance, in grasping the meaning of the name must I simply know which object it refers to or must I also know the manner in which the name *presents* the object? — but the basic premise has rarely been questioned.

What reasons have philosophers had to assimilate names to individual constants? The question immediately strikes one as odd. The oddness is, in part, an expression of fact the picture is so deeply held as to constitute part of our framework for asking questions about language. More interestingly, perhaps, the question strikes one as odd because it seems like, as philosophers, we understand individual constants on the model of names *just as much* as we understand names on the model of individual constants. Individual constants were introduced into formal languages to do the job that names do in natural languages.

Though this may well be true, it has little bearing on the question at hand. The fact that individual constants were 'invented' with the semantic function of names in mind — or the fact that when teaching a student a formal language, one is prone to explain the role of individual constants in terms of the antecedently understood role of names — has very little bearing on the question of the proper understanding of the semantics of proper names. Consider an analogy: it seems equally likely that the symbol '∀' was introduced in quantificational logic on the model of expressions like "every" in natural languages. This has not stopped a broad consensus from emerging to the effect that the semantic function of "every" should not be assimilated to that of '∀'.

It is worth asking, though, what sort of concerns led philosophers to introduce individual constants as a way of making explicit the semantic function of proper names; so we should stop speaking of 'philosophers' in the abstract and start naming names (so to speak). We might as

well focus on Frege and Russell. The first thing to note is that it is widely agreed among historians that neither Frege nor Russell was primarily concerned with natural language¹; they took themselves to be investigating a constellation of logical, metaphysical, and epistemological problems. In this investigation, natural language was often thought to be an *impediment*. A well-constructed logical notation was not designed to properly capture the semantic potential of expressions in natural language it was designed to *improve* on natural language. It was thought that a formal language more systematic than natural language would be a better medium in which to investigate, say, the structure of thought or that of the world.²

The 'translation' of proper names into a formal language was therefore constrained in various ways. In the first instance, theorists only sought to capture the semantic features of proper names that are revealed in a narrow class of linguistic constructions. If the basic question is, say, how to represent deductive inference, one might be content to focus on the occurrence of names in sentences like

- (1) Socrates is mortal

If one has metaphysical puzzles in mind, one might be inclined to focus on sentences like

- (2) Pegasus does not exist

If one had epistemology in mind, one might focus on sentences like

- (3) Roger believes that Cicero denounced Cato

or even

- (4) Cicero is Tully

With sentences like this in mind, it seems perfectly natural to think of the semantic function of names as simply to pick out an object for the rest of the sentence to say something about.

Not only did Frege and Russell have a restricted range of linguistic constructions in view, they had strict constraints on the shape of the formal system into which those constructions were translated. Both Frege and Russell thought that natural languages were ill-suited to philosophical investigation; natural languages are rife with ambiguity, vagueness and context-sensitivity. Formal languages ought to be free of these defects. This meant that even if names in natural languages were context-sensitive devices, an appropriate formal system would not capture this context-sensitivity but rather eliminate it.

¹See ... For a dissenting view see...

²this is the part where i take it back

If, unlike Frege and Russell, *our* task is systematic theorization about natural language we will have both a broader range of linguistic constructions to account for and more latitude with respect to the shape our theory will take. If we use formal machinery, it will not be as a more perfect language to translate our confused natural language into, it will be as a tool to represent what words in our natural languages *actually* mean. Our target will be those properties our words have that allow us to say what we actually say in our language.

If we look beyond the kind of constructions in (1)-(4), what other uses do we find for names? Are those uses consistent with the picture according to which names are semantically analogous to individual constants? Beginning with Tyler Burge's "Reference and Proper Names" in 1973, a view of proper names has emerged which focuses on constructions in which names seem to have a semantic function similar to that of common nouns. Names can combine with definite and indefinite articles, as in

(5) The Alfred I insulted last week is the Mayor

and

(6) In every country club there is always an Alfred

Names can occur with quantificational determiners, as in

(7) Most Alfreds are shy

(8) Every Alfred I've met has been English

and

(9) Two Alfreds entered the ring but only one shall leave

Names have bare plural uses, as in

(10) Throughout the ages, Alfreds have been proud of their name

(5)-(10) are all perfectly acceptable English sentences. It is clear, though, that whatever 'Alfred' is contributing to the meanings of (5)-(10) it is not some individual. The truth of (5)-(10) does not depend upon how any particular individual is. The names in these sentences occupy a place usually occupied by a common noun; (5)-(10) appear to be structurally equivalent to

(11) The man I insulted last week is the Mayor

(12) In every country club there is always a doctor

- (13) Most librarians are shy
- (14) Every cricket-player I've met has been English
- (15) Two fighters entered the ring but only one shall leave
- (16) Throughout the ages, kings have been proud of their name

in which the name 'Alfred' is replaced by a common noun. Common nouns are typically thought of as predicates - that is to say that their semantic function is to introduce a condition that individuals can satisfy or not. Burge argued that we ought to take appearances at face value and say that in (5)-(10) 'Alfred' is a predicate - that it contributes a set of individuals (or, perhaps, a function from individuals to truth-values) to the meaning of a sentence containing it. It is clear from the meaning of (5)-(10) that the condition on individuals which 'Alfred' introduces into the meaning of a sentence is that of being called 'Alfred'. It is clear, I think, that (7), for instance, is equivalent to

- (17) Most things called 'Alfred' are shy

So people named 'Alfred' literally are Alfreds - that is one of the things, along with say cricket-players and mayors, that they are.

If we say that 'Alfred' is a predicate in (5)-(10) we will need to tell a different story about bare occurrences of names (i.e. singular occurrences with no overt determiner). We cannot, it seems, treat 'Alfred' as a predicate in a sentence like

- (18) Alfred is the mayor of Newark

If we did, we would have a type mismatch when we tried to combine the predicate 'Alfred' with the predicate 'is the mayor of Newark'. More than that, the truth of an utterance of (18) does not depend on how the set of Alfreds is but on how some particular, contextually relevant, Alfred is. Burge proposes that bare singular occurrences of proper names abbreviate the semantic role of predicate and demonstrative.³ So in (18) 'Alfred' has a semantics equivalent to something like 'That Alfred'. 'Alfred', in (18), does contribute an individual to the proposition expressed by the sentence. It does not do so because it operates as an individual constant but because it has hidden semantic complexity. Reference to an individual is not direct but is the result of the semantic interaction of demonstrative and predicate. In (9), the semantic contribution of 'Alfred' is the contextually relevant Alfred just as in

- (19) That man is the mayor of Newark

the semantic contribution of 'That man' is the contextually salient man.

³Burge 1973. 428

(**** say something about 'Trieste is no Vienna'*****)

One thing to note here is that bare occurrences of proper names are treated as context-sensitive referring devices. We can see, then, how far we have come from the kind of semantic investigation that preoccupied Frege and Russell. We have been focusing on a class of constructions that simply were of no interest to them, and doing so has led us to posit a meaning for (bare occurrences of) proper names that would have been unsuitable for the kind of formal analysis of language that they were after. Burge's view treats an utterance of (18) used to say something about, say, Alfred Hitchcock and one used to say something about, say, Alfred Nobel, as two utterances of the *same sentence*. Such an analysis would be anathema to Frege and Russell because, for example, in such a system sentences containing proper names would not wear their inferential properties on their sleeve. From a true utterance of (18) and a true utterance of

(20) Alfred is the mayor of Saskatoon

one could not infer

(21) Someone is the mayor of both Newark and Saskatoon

It is, of course, true that from a true utterance of (18) and a true utterance of (20) one *cannot* infer (21) but Frege and Russell would have seen this as a *defect* of natural language. Given their philosophical purposes, they were interested in developing a formal analysis of language in which sentences *do* wear their inferential properties on their sleeve.

A question immediately arises: given that we cannot give a completely unitary account of predicative and bare occurrences of proper names, why should we complicate the account of bare occurrences beyond the individual-constant picture? To put this question another way, why should we think of predicative uses of words like 'Alfred' as uses of *names* at all? After all, there is deep sense in which Burge's account agrees with the individual constant picture when it comes to bare occurrences of proper names. According to both accounts, bare occurrences of proper names have individuals as their semantic contributions.⁴ Given this fact, why not think

⁴The story is not quite as simple as this. Strictly speaking, according to the individual constant view, the meaning of a name is given by a clause of the form

'Alfred' refers to Alfred

while on Burge's account the meaning of a bare occurrence of a name is the complex singular term

$$\iota y(Alfred(y) \wedge y = x)$$

where x is a free variable assigned an individual in context. There is a sense, then, in which on Burge's view the semantic contribution of bare occurrences of 'Alfred' is *both* the descriptive condition of being an Alfred and the individual assigned to the free variable. The descriptive material is truth-conditionally inert in that it only constrains the valuation of the variable. ELBOURNE

of 'Alfred' simply as ambiguous between a referential and a predicative use? Why think that the two kinds of uses should have anything semantically in common at all?

This is a fundamentally important question, and no definitive response is available to us at this juncture. We can, though, offer a few remarks. Given that we are engaged in an empirical investigation of the human language faculty, ambiguity ought to be avoided where possible. The guiding goal of our inquiry is to give as economical as possible a representation of the rules and structures that generate the meanings of sentences in natural language. As in all empirical investigation, simplicity is a virtue. If two competing accounts generate the same set of predications and one does so with fewer explanatory mechanisms, this is *prima facie* evidence for the more parsimonious account. But, of course, it is not anything like *conclusive* evidence. There is no *a priori* reason to suppose that the human language faculty could not be as the individual constant view claims it is - there is no reason to suppose, for instance, that names could not be ambiguous between referential and predicative meanings - and so the mere fact of simplicity is not anything close to being decisive. A more fundamental worry is that evaluations of simplicity are notoriously agent-relative and thoroughly theory-laden. The mere fact that on the individual constant picture, names would be ambiguous between referential and predicative uses is not enough to establish that Burge's account is simpler. After all, in order to account for bare occurrences of names, Burge needs to posit hidden semantic complexity.

So appeals to simplicity will not be much help here, but there can be other reasons to prefer one of two extensionally equivalent theories over another. Two extensionally equivalent theories can differ with respect to their explanatory potential relative to some important empirical generalizations. To take an example from semantic theory: two different semantic theories might generate the same truth-conditions for a sentence like

(22) John buttered the toast

but one might do so in a way that makes intelligible the fact that from (22) one can infer

(23) John buttered the toast somewhere

This might be taken by some to be evidence in favour of the theory that made the inference intelligible.⁵ In discussing a set of examples similar to (5)-(10), Burge writes,

⁵One could, of course, describe this example in a different way. If one thought that a certain class of entailments was part of the explanandum of a semantic theory, then one would not think of the two theories in question as extensionally equivalent. This would, of course, amount to a different criterion of material adequacy on semantic theorizing. It is my view that it is an essential feature of scientific investigation that a phenomenon that was initially taken as an external empirical generalization can become part of the conditions of material adequacy for an explanation of the phenomenon. This is one of the senses, I think, in which the 'subject matter' of any particular empirical investigation is not fixed antecedently to that investigation.

The modified proper names in the examples just given have the same conditions for literal application to an object that singular, unmodified proper names have. This point is confirmed by such sentences as

(24) Jones is a Jones

which is an obvious truth under normal conditions of use.⁶

Burge can be seen here, I think, as making the kind of argument just alluded to. The motivation to give a unified account of bare and predicative occurrences of 'Alfred' is not simply that they each involve the same string of letters (or phonemes) and so doing otherwise would involve positing an ambiguity. There is an important semantical connection between the bare and predicative occurrences - they share "the same conditions for literal application to an object".

Burge is making an important point here but it is not immediately clear how it tells against the individual constant view. As we have seen, Burge explains the connection between the application conditions of the bare and predicative uses by positing hidden semantic complexity in the bare occurrences. At first blush, it seems that the individual constant view could offer the same type of explanation by positing semantic complexity in the *predicative* occurrences — perhaps the individual constant theorist could posit a function that took an individual as argument and returned the set of similarly called individuals as value. Doing so would explain the semantic connection between the bare and predicative occurrences in the same way that Burge does. This strategy, however, will not work. Let's try to define the function in question

$F(x) =_{def}$ the set of things called by x's name

But, of course, things can have more than one name so 'x's name' in the definition does not pick out a unique name (relative to an object) and so 'the set of things called by x's name' does not pick out a set. This is just a reflection of the familiar point that there is no 'road back' from an object to its name. Even if we restricted ourselves to an idealized community in which everything was only called by one name, this strategy would still be unacceptable. For any utterance involving a predicative occurrence of 'Alfred' we would have to pick some particular Alfred as argument to the function **F**. That is to say, if I uttered

(25) All Alfreds are uptight

Its meaning would be given by a clause of the form

(26) [All x : $F(\text{Alfred}_3) x$] x is uptight

⁶number change. 426-427

where the subscripted '3' after 'Alfred' is a disambiguating device in the meta-language used to make clear which Alfred I'm talking about. So in uttering (28) I would be making a claim about some *particular* Alfred. Or to put the point another way, the meaning of (28) will stand in some relationship to, say, Alfred Hitchcock (if he is denoted Alfred₃) that it does not stand in to say, Alfred Nobel. To assert something about all Alfreds, I must first pick some particular Alfred to work with. But, of course, which Alfred I pick will be truth-conditionally irrelevant. This arbitrariness, I submit, is deeply troubling.⁷

It is clear, then, that the individual constant view will not be able to capture the sameness in conditions of application in the bare and predicative occurrences. This should not come as a surprise given that, according to the individual constant view, bare occurrences of names *do not have* conditions for application to their objects. They might be able to capture the truth of (24) by employing some mechanism for quotation; but this would introduce mild non-compositionality into the meaning of the predicative occurrence. The meaning of the predicative occurrence would not be a function of the meaning of the bare occurrence (compare this to Burge's account where the meaning of the bare occurrence *is* a function of the meaning of predicative occurrence). In each account, there is an asymmetry between the bare and predicative occurrences. In Burge's account the bare occurrences are understood in terms of the predicative ones, in the individual constant view the predicative occurrences are understood in terms of the bare ones. The order of explanation associated with Burge's account allows us to see the connection between the two kinds of use as *semantic* while the opposite order can only see the connection as involving something like semantic ascent. But this is not, after all, obviously wrong. According to the individual constant view, the truth of (24) is of the same sort as the truth of, say,

(27) Dogs are called 'Dogs'

⁸ And even on Burge's view the predicative occurrence of names is essentially metalinguistic; there remain two differences though. First, for Burge, predicative occurrences are metalinguistic in virtue of their content not in virtue of their form; and second, in Burge's view, since the meaning of the bare occurrences is a function of the meaning of the predicative occurrences, the bare occurrences themselves are essentially metalinguistic.

It seems that no decisive objection to the individual constant view is available here. Each view generates the same predications; each *explains* the data differently but we are, so far, unable to see why we should prefer one explanation over the other (in our darker moments, per-

⁷Thanks to Rachel Goodman for helpful discussion here

⁸Burge would think of the truth of (24) as being of the same sort as

(28) That dog is a dog

haps, we worry that the explanation are, at some deep level, the same and that the two views are simply notational variants of each other). Ideally, we will be able to elaborate each theory to a level at which they *do* generate different predictions from each other. This would give us reason to prefer one over the other.

As I see it, the essential features of Burge's account are as follows:

1. The meaning of predicative uses of proper names is metalinguistic
2. The meaning of bare occurrences of proper names is a function of the meaning of predicative occurrences
3. Bare occurrences are context-sensitive in some important sense

Let's call any view with these features a Predicate View of proper names (henceforth PV). We can think of Burge's own proposal as amounting to particular determinations of features (2) and (3). Burge thinks that the meaning of a bare occurrence of 'Alfred' is a function of the meaning of a predicative occurrence of 'Alfred' in the same way that the meaning of 'that dog' is a function of the meaning of 'dog'; let's say that, on Burge's version of PV, the meaning of a bare occurrence of a name is a *demonstrative* function of the meaning of a predicative occurrence of that name. Given a certain view of demonstrative constructions, this gives us a determination of the third feature of PV - namely, that bare occurrences of proper names are basically context-sensitive *referring* devices and that the contribution that a bare occurrence makes to the meaning of an utterance is some contextually salient individual who bears the name in question.⁹

If we see (1)-(3) as representing the essential features of Burge's proposal, then we can see the details of Burge's proposal as being, to some extent, arbitrary. If we think of the meaning of 'Alfred' on the model of a common noun, and want to say that bare occurrences of proper names have the kind of meaning a noun-phrase has, there are a variety of potential functions available to us. This is, of course, simply because noun-phrases can have different sorts of meanings. After all, 'That Alfred', 'Every Alfred', 'An Alfred', etc, are all perfectly legitimate noun phrases. Burge's own proposal (that bare occurrences are demonstrative functions of predicative occurrences) is, of course, not completely arbitrary. Bare occurrences of proper names are *definite* noun phrases so it would clearly be inappropriate to analyze them on the model of say, 'An Alfred' (say more about definiteness). More intuitively, it seems clear that the point of using a bare proper name is (almost always) to say something about some contextually salient individual; so analyzing bare proper names on the model of a noun phrase the point of which was not to do *that* would have seriously counter-intuitive consequences. Having said all of that, we are not forced to think of bare proper names as demonstratives. For all that we have said, proper names might be more like definite descriptions; definite descriptions are, not surprisingly, semantically definite and it is clear that, at least as far as untutored linguistic intuitions

⁹with the same qualifications as above

go, definite descriptions are often used to say something about some contextually salient individual.¹⁰

In fact, the few philosophers and linguists who have adopted PV since Burge's initial publication have tended to move away from the details of his view. For example, Larson and Segal and Elbourne have both argued that it would be better to understand bare occurrences of names on the model of definite descriptions. In support of this amendment, Larson and Segal and Elbourne draw on data from languages other than English. Elbourne notes that in some dialects of German there are no bare occurrences of names - "we always have *der* Hans 'the Hans', where in English we would just expect Hans."¹¹ Similar patterns exist in other languages and to the extent that one believes that at some important level all natural human languages share a structure, this presents significant support for the idea that apparent bare occurrences of names involve covert definite determiners. Let's say that on this such a view, the meaning of a bare occurrence of a name is *definitive* function of the meaning of its predicative occurrence.

We should note here that on the definitive proposal, bare proper names will be semantically analogous to *incomplete* definite descriptions. It is a bit tricky to give a precise definition of the distinction between complete and incomplete definite descriptions, but the following will serve as a rough and ready characterization.

A description 'The **F**' is complete relative to a conversational situation iff it is common knowledge among the conversational participants that one and only thing is **F**. It is incomplete iff it is not complete.¹²

To say that bare occurrences of names are incomplete definite descriptions is to say, then, that in most conversational contexts it is not common knowledge that only one individual bears the name in question.

So the definitive determination of (2) entails a corresponding determination of (3). If bare occurrences are definitive functions of predicative occurrences then bare occurrences are context-sensitive in whatever way that incomplete definite descriptions are. It needs to be said at this point that claiming that bare occurrences of names are definitive functions of predicative occurrences still leaves (2) underspecified in important ways. This is simply because there is substantial disagreement over how to understand the semantics of definite descriptions.¹³ There are at

¹⁰We might say that the fact that Burge adopts the demonstrative proposal without considering the definitive proposal reflects the hold that the individual constant view still had over his thinking. This seems right, but the situation is somewhat more complicated than that. For one thing, Burge thought that many *definite descriptions* were better thought of on the model of demonstratives. Source

¹¹Elbourne. 173.

¹²It is important that this proposal does not say 'one and only salient thing is **F**'. Presumably, in the case of a felicitously used incomplete definite description it will be common knowledge among conversational participants that one and only one salient thing is **F** and *not* common knowledge that one and only thing in the world is **F**.

¹³Note that, as a matter of fact, the same thing is true of the demonstrative proposal. Cf King

least four proposals alive in the literature today for how to treat definite descriptions. There is the traditional Russellian analysis according to which definite descriptions are basically quantifier phrases which assert the existence and uniqueness of an object satisfying the restrictive material. There are approaches which treat definite descriptions simply as existential quantifiers, but also attach some pragmatic condition of *identifiability* or *givenness* to the satisfier of the restrictive material. There are the approaches that we can broadly identify with Frege and Strawson; according to this approach, definite descriptions are thought of as complex singular terms which presuppose that a unique object satisfies their descriptive material. Lastly, there is the approach associated with the tradition of dynamic semantics. According to this tradition, definite descriptions are thought of as free variables restricted by various conditions; uttering a sentence containing a definite description effects certain changes on the conversational participants representation of salient individuals and their properties.¹⁴ For the purposes of this project, I will attempt to remain neutral with respect to the question of which account is the right account of definite descriptions. As a matter of fact, the dynamic account falls so far outside the framework of the other three accounts, that I will not be able to consider its relation to this work. This attempted neutrality should not be taken as a belief that what the correct view of definite descriptions is *does not matter* with respect to the Predicate View of names. To a large extent, the kinds of considerations that will decide amongst the various accounts of definite descriptions are orthogonal to what is at issue here. But this is not universally the case, and we will at times, have to stop and consider the matter from the point of view of different theories.

In moving to the determinative version of PV we have opened ourselves up to a range of important criticism. Philosophers have a range of arguments that they take show that names should not be assimilated to definite descriptions. It is close to philosophical orthodoxy that names and definite descriptions do not exhibit the same range of semantic potential. A famous example of such a claim is Kripke's modal argument in *Naming and Necessity*. According to Kripke, names do not exhibit scope interactions with modal operators.¹⁵ An utterance of a sentence like

(29) George Bush might have been a Democrat

can only have truth conditions according to which the same person's party affiliation is relevant to evaluation of the contained claim at each possible world. This is in contrast to a sentence like

(30) The president might have been a Democrat

¹⁴For the Russellian approach see Neale; for the existential see Ludlow, Szabo, Hawthorne; for the Frege/Strawson approach see Heim, Elbourne; for the dynamic approach see Heim.

¹⁵Of course, Kripke's modal arguments were meant to show more than just that names do not exhibit scope interactions with modal operators. The thesis that names are rigid designators is not reducible to that claim. See Kripke contra Dummett. Recanati. This will become important later on. Let's leave it aside for now.

which *does* have a reading according to which different people's party affiliation is relevant with respect to evaluation of the contained claim at each possible world.

So it looks as though the interpretation of definite descriptions, and not that of names, is relativized in some important way to possible worlds. The same thing can be said of relativization to *time*. The sentence

(31) Next year the president will be a Democrat

has reading according to which it is asserting that the current president will change parties and a reading according to which it is making a prediction about the outcome of the next election. Sentences involving names do not seem to have the same potential for ambiguity.

The interpretation of a definite description can also be relativized to objects introduced by quantifiers. For example, in

(32) Every man remembers the first woman he kissed

the interpretation of 'the first woman he kissed' can be relativized to the objects introduced by the quantificational expression 'Every man'.¹⁶

Lastly, there are sentences in which the interpretation of a definite description is relativized to objects (or situations) introduced by a quantifier in the sentence despite not containing a pronoun bound by that quantifier. Consider the sentence

(33) Every farmer who owns a donkey and a mule beats the donkey

Here, the interpretation of 'the donkey' is relativized to the the farmers introduced by 'every farmer who owns a donkey and a mule'.¹⁷

All of these examples have something in common. In every case the interpretation of the definite description is *subordinated* to the interpretation of some other element of the sentence. It is somewhat difficult to make this sense of *subordination* more precise without going in for one theory of definite descriptions or another. At a general level, we can say that in the subordinated readings, the descriptive material introduced by the description is interpreted relative to some other semantic machinery - be it a possible world, time, situation, or variable assignment. In non-subordinated cases, the descriptive material is simply interpreted relative to the circumstance of evaluation. That proper names seem not to exhibit the possibility of subordinated readings is *not* strong evidence that they are not associated with descriptive material¹⁸; it

¹⁶I say 'can be relativized' because, of course, the sentence has a reading in which 'he' is treated referentially

¹⁷One might think that at the level of logical form, something like quantifying in is going on here too. This would amount to positing an unpronounced anaphoric pronoun (Cf King). I will not take a stand on this issue here.

¹⁸After all, philosophers who have thought that complex demonstratives are rigid designators have presumably not thought that they were devoid of descriptive material

is strong evidence, though, that *if* names are associated with descriptive material, that material does not interact with semantic machinery in the way that the descriptive material associated with definite descriptions does. And this is good evidence that names and definite descriptions are not semantically analogous.

We should pause, though, to wonder just how robust the above generalizations about definite descriptions are. For instance, does just any sentence involving a definite description and a temporal operator have a semantically distinct subordinated reading? We should especially focus on incomplete definite descriptions because on the determinative version of PV, *that* is what bare proper names are. It turns out that if we shift our attention to incomplete definite descriptions, the apparently robust pattern we canvassed in the last paragraph begins to break down. Consider the sentence

(34) In ten minutes, the table will be covered in food

It is not immediately clear that this sentence has a reading in which the definite description ('the table') has a semantically distinct temporally subordinate reading. Part of the problem here is that the description is incomplete and so it's hard to see what uttering (34) would amount to absent a description of the context of utterance. Imagine, then, that you and I are hosting a dinner party; the guests arrived an hour ago, and the food is still not prepared. Standing in front of the empty table, trying to assure you that dinner is imminent, I utter (34). It is clear that in this case, the definite description does not receive a temporally subordinate reading. Or perhaps it is better to say that it is not at all clear that there is a semantic difference between a temporally subordinate and a non-subordinate reading. If I am not making a predication about the very table in front of us, what else could I be doing, given what I said. The expression I chose, 'the table', does not contain enough descriptive content, relative to our conversational purposes, for me to expect you to be able to identify any *other* table as the relevant one ten minutes from now.

It is tempting to explain what is going on here by relying on some account of how incomplete definite descriptions are 'completed' in context. On one prominent story, associated with Stephen Neale, descriptions can be completed with singular material.¹⁹ I will not go into the details of how this is supposed to work here; the basic idea is that, in the context just imagined, (34) is somehow semantically equivalent to

(35) In ten minutes, the table over there will be covered in food

where 'over there' is an expression that picks out a certain region of space/time. If this were true, then the whole description would be rigid (or at least inflexible), and so a semantically distinct temporally subordinate reading would be ruled out.

¹⁹source

This may very well be the correct account of what is going on in the case we imagined. The point at hand, though, does not depend on any particular analysis of the phenomenon. The thing to note is that the temporally subordinate reading is not always available with incomplete definite descriptions and so we need to investigate a little further before we decide that names are not incomplete definite descriptions.

In the context we imagined, it was clear that (34) did not have a distinct reading in which the incomplete definite description had a temporally subordinate interpretation; but this does not show that there is no context in which an utterance of (34) could have that reading. In explaining the absence of a distinct subordinate reading, I claimed that the description did not contain enough descriptive content, relative to our conversational purposes, for me to expect you to be able to identify any *other* table as the relevant one ten minutes from now. If this is the right explanation of what was going on, then *if* we find a context in which the description does contain enough descriptive material relative to our conversational background, *then* we should see a subordinate reading. Imagine, then, that you and I are inspectors at a furniture factory. We are at the end of an assembly line, and we inspect tables and desks as they pass on the conveyor belt. Given the way the factory operates, every five minutes a desk and a table pass by us on the conveyor belt. On this particular day, some of the workers further up the assembly line are not taking their responsibilities particularly seriously; in fact it appears as though they are having a party. First a desk comes down the conveyor belt with a couple of empty bottles on it; then a table comes by with a half-eaten pizza on it. Things are obviously getting worse (the last desk had a keg on it) and so I utter (34). Even though there might be a table right in front of us on the conveyor belt, it seems clear that the natural interpretation of what I said is one on which the definite description has a temporally subordinate reading. It seems clear that I am making a predication about the table that will come down the conveyor belt in ten minutes. Given that the conversational background includes such information as *every ten minutes a table comes down the conveyor belt* and *the workers upstream seem to be putting more and more inappropriate stuff on the table* the descriptive material associated with the description I uttered is sufficient to make it clear which table in the future I mean to be saying something about.

A note on the example: it was not merely poetical flourish that made me include the desk alongside the table on the conveyor belt. If we had imagined the example with only tables coming down the conveyor belt, I think that it would have been noticeably harder to hear the relevant reading of (34) (I'll leave it as an exercise for the reader to run through the example again without the desks). This is not an accidental feature of the case but, I think, is a symptom of a fairly general feature of sentences involving definite descriptions. Sentences involving definite descriptions are naturally read with contrastive stress on the restrictive material in the description. We'll call such a case a 'contrastive use' of a definite description. A contrastive use of a description 'The **F**' presupposes that there is a contextually relevant object that is not **F**. In *Complex Demonstratives*, in the course of distinguishing demonstratives from definite descrip-

tions, Jeff King imagines the following example.²⁰ A group of senators hears that an unnamed senator will be testifying to the senate ethics committee “regarding campaign finance irregularities”. One of them utters

(36) I’m going to kill that senator

This is a perfectly felicitous utterance and it is clear that the speaker means to be saying something about whichever senator is the one who will be testifying. Note that uttering

(37) I’m going to kill the senator

would have been infelicitous. King claims that this reveals that “the semantics of ‘that’ phrases allows for supplementation by properties of speakers’ intentions whereas the semantics of ‘the’ does not.”²¹ The idea is supposed to be that ‘that senator’ in (36) can be interpreted relative to the speakers clear intention to be speaking of whichever senator will be testifying and that ‘the senator’ in (37) cannot be so interpreted. The problem is, though, that the infelicity of (37) has nothing to do with the relation between definite descriptions and speakers’ intentions. The problem with the example is that (37) is naturally read with contrastive stress on the restrictive material in the definite description (i.e. ‘senator’). Given this, there should be two ways to fix the infelicity. The first way, would be to alter the example to satisfy the contrastive presupposition. Consider a slightly modified version of the example. A group of senators hear that an unnamed senator *and* an unnamed congressman will be testifying to the senate ethics committee. Feeling betrayed by one of his own, a senator utters (37). Now (37) is clearly felicitous and not because it stands in some different relationship to the intentions of the speaker (the speaker still means to be saying something about whichever senator will be testifying). The difference is that, in this case, the contrastive presupposition of the definite description is satisfied. Among the contextually relevant individuals (people testifying before the house senate committee) there is something that is not a senator (namely, the congressman). The other way to make (37) felicitous would be to read it with contrastive stress on some other element of the sentence (and ensure that the contrastive presupposition triggered by that reading is satisfied). Imagine if, before the utterance of (37), someone had uttered.

(38) I’m going to hurt that senator

Even in the initial scenario (with only a senator and no congressman) if (38) was followed by an utterance of (37) with contrastive stress on ‘kill’, (37) would be perfectly felicitous. This is simply because putting the stress on ‘kill’ invites the audience to consider a range of possible alternative actions (for instance, mere harm).

²⁰70-71

²¹71

An in-depth analysis of this phenomenon would take us well beyond the scope of this project. I bring it up merely as a caution; it will be important as we test for different readings of definite descriptions (and names) that we control for this effect. I think, though, there is a further reason why this phenomenon has a particularly strong effect on sentences like (37). Notice that in the first imagined senatorial scenario, but not in the second, it would be felicitous to say

(39) I'm going to kill him

The use of the pronoun is felicitous because there is only one contextually relevant individual (the person testifying before the committee). One needs to move to the definite description (or the demonstrative) in the second case because one needs to distinguish among *different* contextually relevant individuals. In this contrast we have the resource for a pragmatic explanation of the strong infelicity of (37) in the first imagined scenario. If one attempts to use a definite description in a context where there is only contextually relevant individual, that fact that one *could* have used a pronoun (which, anyway, would have been quicker) but *chose* not to, suggests to the audience that you take there to be some need to distinguish among different individuals. But *ex hypothesi* there is no such need, and the audience knows it, so the utterance triggers a search for an implicature (you have flouted the maxim of manner). If no reasonable implicature can be found, then the utterance is infelicitous.²² The point of this sketchy analysis is just to draw the connection between when definite descriptions *are not* felicitous and when pronouns *are*. This will give us another resource to make sure that extraneous factors are not interfering when we are searching for subordinate readings of definite descriptions and names. To return to tables and desks: the example works better with both tables and desks because the presence of desks satisfies the contrastive presupposition of 'the table'. In the case with only tables, the subordinate reading is slightly harder to hear because I could just have easily said

(40) In ten minutes, it will be covered in food

So we have some reason to suppose that our generalizations about the possibility of subordinate readings of definite descriptions need to be significantly more nuanced than we had originally supposed. In particular we need to pay close attention to the relationship between the descriptive material associated with a definite description and the common ground in a conversational situation. Daniel Rothschild, in "Presupposition and Scope", undertakes exactly such a nuanced discussion of when definite descriptions can have modally subordinate readings.²³

²²The final explanation of the phenomenon will have to be significantly more complicated than this, and perhaps categorically different. In the first place, nothing I've said here predicts the difference in behaviour between definite descriptions and complex demonstratives that King noted. Most importantly talk of contextually relevant individuals is going to have to be cashed out in some detailed way.

²³Rothschild actually talks about 'narrow scope' reading; I prefer 'subordinate' because I think that 'narrow scope'

²⁴ He distinguishes between *role-type* and *particularized* definite descriptions. A role-type definite description is such that it is “part of the common ground that there is exactly one person (or one salient person) satisfying the descriptive content across a range of relevant metaphysically possible situations and that the satisfier sometimes varies from situation to situation.”²⁵ A particularized description is simply one that is not role-type. He goes on to argue that only role-type descriptions can have modally subordinate readings. Let’s look at an example.

Suppose that throughout an entire dinner party Siegfried does not eat anything, and is unique in this regard. Suppose that I have another friend, say Siegmund, who also would not have eaten anything if he had been at the dinner. Now, suppose I say something like this:

(41) I might have enjoyed talking to the person fasting through the dinner

It does not seem like I could mean anything but that I might have enjoyed talking to Siegfried by an utterance of (41). This is true even if it is possible that Siegmund could have come and Siegfried not come. In this possible situation, of course, Siegmund would have been the only person fasting. Nonetheless, it does not seem like (41) can easily express the proposition that there is a possible situation in which I would have enjoyed talking to *whoever* was unique in fasting at the dinner.²⁶

Rothschild goes on to rightly point out that if we change the situation slightly by adding that the organizers of the dinner party, because of some idiosyncrasy, are always sure to invite one and only one person who they know will fast throughout dinner, then the subordinated (in his words ‘narrow-scope’) reading becomes available. This is, of course, because altering the context in that way has changed ‘the person fasting through dinner’ from a particularized to a role-type description.

So we can see that just as it was in the temporal case, the conditions under which definite descriptions can have modally subordinate readings are somewhat more delicate than we might have initially thought. Rothschild’s notion of a role-type description is a tool for making precise the relationship a description must have to the conversational common ground to

implies a privileged sort of explanation for the readings (one associated with a quantificational account of definite descriptions). I think that ‘subordinate’ does a better job of giving a theory-neutral characterization of the phenomena.

²⁴It is important to mark here the extent of my debt to Rothschild. I began the present project, and developed the main strand of argument, before having read Rothschild’s work. However, Rothschild has done such an admirable job of showing the contours of the phenomenon in question that I’m sure that my thinking has been greatly shaped by his work. In any case, it is important to note that he has done the pioneering work in this region

²⁵6

²⁶10. numbers

admit of modally subordinate readings.²⁷ We can imagine generalizing his idea of role-type descriptions to encompass subordination of other kinds, but we should wait until we examine quantificational subordination to see if any new complexities arise there.

What are the conditions under which descriptions can have semantically distinct quantificationally subordinate readings? The first kind of case is one in which the description contains a pronoun which is interpreted as anaphoric on a quantificational antecedent. Recall that in

(32) Every man remembers the first woman he kissed

if 'he' is interpreted as anaphoric on 'Every man' the description has a subordinated reading. Contrast (32) with

(42) Every man remembers the woman he kissed

Absent a context of utterance, the subordinated reading of the description seems equally available here. But imagine the following scenario: you and I are at a country fair. In front of us is a row of kissing booths, each staffed by a different man. We watch the proceedings for an hour, seeing multiple women partake of the services of each man. Commenting on the situation, I utter (42). Certainly there is no way to hear the pronoun as anaphoric on the quantifier. So unless there is some contextually salient male I might be taken to be referring to with 'he', the utterance is simply infelicitous. If I had uttered (32) instead, everything would have gone smoothly. The explanation is easy to find. It is part of the conversational common ground that for every relevant man, there is a unique first woman he kissed; it *is not* part of the common ground that for every relevant man, there is a unique woman he kissed.

What about cases in which there is no (overt) anaphoric pronoun in the definite description? Consider a sentence like

(43) Every man remembers the woman

Absent the right kind of context, it's difficult (more difficult than with (42), I think) to see what a subordinated reading of the description would be. The sentence is most naturally read as simply making a quantificational claim about a single, contextually salient woman. Given the right context, though, the subordinated reading is easily available. By this point, the procedure is getting to be fairly rote. We just need to find a context in which for each relevant man, it is part of the common ground that there is an identifiable (potentially different) woman. Imagine the following scenario: you and I are psychologists studying memories from early childhood. Due to our Freudian leanings, we suspect that men will remember their female grade-school

²⁷Rothschild goes on to explain the pattern in term of the presuppositions associated with definite descriptions. His analysis is plausible, but I want to leave it aside because it relies on a particular account of definite descriptions

teachers more than their male grade-school teachers. As an experiment, we bring in a group of adult men, and one grade-school teacher of theirs of each sex. Reporting on the results of the experiment, I utter (43). ‘The woman’ clearly has a subordinated reading here. Notice that part of the strategy to bring out the subordinated reading in this case involved satisfying the contrastive presupposition of ‘the woman’; I did this in the example by relating each man to both a male and a female teacher. If the example had only included female teachers, the reading would have been somewhat harder to get.²⁸

One way to ensure a subordinated reading quantificational cases, is to build the relevant material into the sentence itself. Recall

(34) Every farmer who owns a donkey and a mule beats the donkey

Here ‘the donkey’ gets a subordinated reading because the quantifier it is subordinated to explicitly introduces the identifying conditions necessary to get a subordinated reading of ‘the donkey’. It is clear that with respect to every farmer who owns a donkey, there will be a contextually identifiable donkey - namely the donkey that that farmer owns.²⁹

What can we say, in general, about when a definite description can admit of a subordinated interpretation? First of all, let’s call an entity introduced by either a modal operator, temporal operator or quantifier phrase a *subordinating condition*. Our generalization will take the form of a condition that the descriptive material associated with a definite description must stand in to each subordinating condition introduced by some other element. At a completely general level, the condition would like this:

(C) A description ‘The F’ can have a reading on which it is subordinated to some operator, *O*, iff there exists an identifiable *F* relative to every subordinating condition introduced by *O*^{30 31}

²⁸Again, I’ll leave it to the reader to do the exercise. I don’t think that it’s impossible to get the reading. It helps, for instance, to put contrastive stress on ‘remembers’. This introduces the contrast class of other relevant relations each man might stand in to the woman (e.g. forgetting, resenting etc) rather than the non-existent (or at least very difficult to find) contrast class of other relevant people every man might remember.

²⁹Donkey anaphora

³⁰It is important to note here that what subordinating conditions are in fact introduced by *O* will be context-sensitive in various ways. This is simply the result of, say, quantifier domain restriction etc.

³¹What ‘relative to’ in ‘there exists an identifiable *F* relative to every subordinating condition’ comes to will be slightly different with respect to different kinds of subordinating condition. For instance, for there to be an identifiable *F* relative to a possible world *w* is just for there to be an identifiable *F* in *w*. Depending on how one thinks about temporal operators, something similar can be said in that case. When the subordinating condition is quantificational, the situation is a little more delicate. What it is for there to be an identifiable *F* relative to some object, *o* (introduced by a quantifier phrase)? In the case of a description with an overt anaphoric pronoun, the situation will be relatively straightforward. In such a case, *F* instead of being a condition on individuals will be a function from individuals to conditions on individuals (type <e,et>, as they say). For there to be an identifiable *F* relative to *o*

So, after a moderately thorough investigation of the conditions under which definite descriptions can have subordinated readings, we are ready to return to names. Recall our problem: we have adopted a version of PV according to which bare occurrences of proper names are semantically analogous to incomplete definite descriptions. Philosophers have taken the lack of semantically distinct subordinated readings of names as strong evidence that names and definite descriptions do not have the same range of semantic potential. We have come to see that a description ‘The **F**’ can have a subordinated reading only if the property of being an **F** is salient enough in the conversational context to allow the audience to identify a unique **F** with respect to every subordinating condition. From the perspective of this generalization, it is not at all surprising that it is difficult to find subordinated reading of names. According to PV, the descriptive material associated with a name simply is the property of being called by that name. In most situations, knowing someone is called by a name is not knowing very much about them. More precisely, their name does not figure in explanations, predictions or generalizations of any useful sort. That being the case, one cannot assume that the property of bearing a name can bear the weight of picking out a relevant individual relative to a subordinating condition.

If we want to test whether or not bare proper names really are semantically analogous to definite descriptions, we need to contrive conversational situations in which the property of bearing the name in question has the kind of salience necessary to underwrite subordinated readings. Such contexts will necessarily be somewhat odd; but, again, this is simply an expression of the typically minimal role that name-bearing properties play in our informational transactions with one another.³² It turns out, that if we find such situations we also find bare proper names behaving exactly like definite descriptions.

Let’s start with temporal subordination. Imagine the following scenario: we have a monthly dinner engagement with an eccentric friend of ours who strives always to have one girlfriend named ‘Helen’ and one named ‘Jane’. During dinner, he expresses to us his dissatisfaction with his current relationships. When you and I are alone afterward, we come to suspect that what our eccentric friend, who we know to be rather shallow, really wants is to have one rich girlfriend

is just for there to be an identifiable object that satisfies the condition that is the the value of $F(o)$. For example : for their to be an identifiable *woman he kissed*, relative to some *man, M*, is for there to be an identifiable *women M kissed*. The situation is more complicated in the case of descriptions without overt anaphors. Here, at least on the surface, **F** is simply a condition on objects. We must introduce a contextually salient relation, *R*. There will be an identifiable **F** relative to some object *o* iff there is an identifiable **F** that is *R*-related to *o*

³²This is too crude. There is a sense in which name-bearing properties play a *fundamental* role in our informational transactions with one another: they ground our most common mode of reference to individuals (this is their role in non-subordinated bare occurrences). The sense of ‘playing a role in our informational transactions’ at issue here, though, is somewhat different. It is just that name bearing properties do not play a constitutive role in the assertions, explanations, and generalizations that form the basis of our informational commerce with each other.
SAY SOMETHING ABOUT STALNAKER

and one interesting girlfriend. I utter

(44) Next month, Helen will be an investment banker and Jane will be an artist

The salient reading of (44) is one in which ‘Helen’ and ‘Jane’ receive a temporally subordinate interpretation. The most natural reading is not one on according to which I am predicting a change in career for our friend’s current girlfriends, but one on which I am making a prediction about our friend’s next girlfriends, whoever they may be.³³ I can do this by using ‘Helen’ (and ‘Jane’) because it is part of the common ground of the conversation that in a month’s time there will be an identifiable person named Helen (and one named ‘Jane’) who is not necessarily the woman our friend is currently dating. So given the hypothesis that bare proper names are semantically analogous to definite descriptions, and the account of the conditions under which descriptions can have subordinated readings, we would predict that ‘Helen’ and ‘Jane’ in (44) should admit of temporally subordinate readings. And that is exactly what we found.

Moving to the modal case, we find exactly the same pattern. Imagine that because of some bizarre bureaucratic condition, the White House must hire an intern named Dave after each election. It is common knowledge that the winning party always hires a Dave from their own party and that, because of other eccentric rules, no two people with the same first name are allowed to work in the White House at the same time (perhaps because the President is easily confused). Discussing how tight the last election was, I utter

(45) If things had been different, Dave might have been a Democrat

In this sentence, ‘Dave’ clearly has a modally subordinate reading. Again, this is exactly what the determinative version of PV predicts.³⁴

³³With the qualification that they be named ‘Helen’ and ‘Jane’. See section**** for an amplification of this point.

³⁴The modal case has received the most attention; here are some more examples from the literature. Elbourne, reporting an example from Heim, writes

John has four sons: Primus, Secundus, Tertius, and Quartus. Detecting a pattern here, and knowing that the four sons are very overbearing, we say

(46) If John had had five sons, Quintus would have been bullied by the others

This sentence is perfectly grammatical under the interpretation whereby it claims that John’s fifth son would have been bullied by the others.(181-182.number)

In “Presupposition and Scope”, Rothschild writes of the sentence

(47) Aristotle could have been a sailor

“If there is a Greek law stipulating that there is always one and only Aristotle at any given time, then (47) seems like it has a reading on which the name gets narrow scope.”(41)

The case of quantificational subordination requires a few comments. It is important to note that we cannot look for cases where bare occurrences of names contain pronouns anaphoric on quantificational antecedents. This is simply because names cannot contain pronouns. Or more precisely, bare names cannot be modified by restrictive relative clauses. That is, I can say ‘The N who I met last week . . .’ and have ‘who I met last week’ function to help identify the N in question. If I say ‘N, who I met last week, . . .’ ‘who I met last week’ can only function as a kind of comment on the rest of sentence

But if we look to cases where we would expect definite descriptions which do not contain anaphoric pronouns to have quantificationally subordinate readings, we find that names have the same possibilities. You and I are discussing names that can be given to either males or females. Having noticed an interesting pattern I utter,

(48)

Every time that I’ve met a Leslie who was married to a Hilary, Hilary has been the husband

Here, ‘Hilary’ clearly has a quantificationally subordinate reading. I am not asserting something about some particular contextually relevant Hilary. I am making a general claim about Hilary and Leslie pairs; namely that in each such pair, Hilary is the man.³⁵

We can also find cases in which the identifying condition is not explicitly contained in the sentence. You and I have made a bet about the relative appeal of the names ‘Hugo’ and ‘Bernard’. We decide to settle the bet by conducting an experiment. We send different pairs of Hugos and Bernards to approach different women and ask them out (we have, of course, controlled for the relative attractiveness of each pair). The women overwhelmingly chose Bernard; Consoling you, in a slightly mocking way, I say

(50) Don’t feel so bad. One woman chose Hugo

Here ‘Hugo’ receives a subordinated interpretation. I am not asserting that one woman chose some contextually salient Hugo, but that one woman chose whichever Hugo was sent to her.

One can imagine someone objecting at this point that the cases we have canvassed here are so contrived so as to not reveal anything interesting about the semantics of proper names. Why should we think that the kind of contribution a name makes to the meaning of a sentence in these bizarre circumstances should shed any light on the kind of contribution a name makes in

³⁵A similar example from Elbourne (181):

(49)

Every woman who has a husband called John and a lover called Gerontius takes only Gerontius to the Rare Names convention

more mundane circumstances? This amounts to claiming that bare occurrences of names are semantically ambiguous between referential uses and role-type uses. In the context of the disagreement between the individual constant view and the predicate view, this must be seen as a somewhat desperate strategy. The individual constant theorist was already forced to concede that names were ambiguous between predicative and bare uses; she is now forced to concede that bare occurrences themselves are ambiguous. Moreover, claiming that a role-type use of a bare name is semantically anomalous is simply and clearly begging the question against the predicate view. On the assumption that names are semantically analogous to different definite descriptions one predicts that, under certain conditions, they should have a use that is comparable to role-type uses of definite descriptions. Unless one is willing to deny that there is a unified semantical treatment of definite descriptions, it is simply begging the question against PV to claim that existence of role-type uses for names suggests that no unified treatment of bare occurrences of names is desirable.

We ended our discussion of Burge's version of PV by noting that his theory seemed to be extensionally equivalent to the individual constant view; and though there was some reason to suppose that his account had more explanatory scope, there certainly was no decisive reason. The determinative version of PV is in a very different boat. In investigating the determinative version of PV, we found that it made substantially different predictions from the individual constant view. In particular, the determinative version of PV predicted that bare occurrences of names should exhibit subordinated readings. Philosophical orthodoxy holds that names exhibit no such readings. In following where the determinative version of PV lead us we found that names *do* exhibit subordinated readings, but only in slightly contrived contexts. But this is exactly what we would predict if bare occurrences were associated with the kind of descriptive material that PV claims them to be. This, I submit, is strong evidence in favour of the determinative version of PV. It did what a good scientific theory does. It made predictions that were clearly falsifiable. In testing those prediction, we discovered a class of hitherto unnoticed constructions. This is what progress looks like.

1 Objections

1.1 The Pope and the Condemned Man

In the body of the paper I imagined a worry of the following form: "aren't the cases in which you claimed to have found subordinated uses of names so contrived so as to not reveal anything interesting about the semantics of proper names?" At the time I claimed that such a response simply begs the question against the determinative version of the predicate view. I think, though, that there is a more sophisticated version of this worry. If I am claiming that names are semantically analogous to definite descriptions, and I argue for that by showing that, in certain

contexts, names can have subordinated interpretations — it would help to see the force of the argument if it was clear what sort of expressions *did not* exhibit subordinated interpretations given the right sort of context. It would be clearer what the content of my claim about names was if it was clearer what they were being *opposed to*.³⁶ If we look at other kinds of referring expressions, though, it is not clear that *any* of them do not exhibit the kinds of readings I found with names. Consider

(51) [Said while pointing at the Pope] He is usually Italian

(52)

[Said by a condemned prisoner] I am traditionally allowed to order whatever I like for my last meal

(53) Every father dreads that moment when his oldest child leaves home

In (51)³⁷, (52) and (53)³⁸ respectively, a pronoun, indexical and demonstrative have subordinated readings. Insofar as names, definite descriptions, pronouns and demonstratives exhaust the class of referring expressions, it appears that I am committed to the claiming that *all* referring expressions are semantically analogous to definite descriptions.

The first thing to note is that it is not completely implausible to suggest that pronouns and demonstratives *are* semantically analogous to definite descriptions. Paul Elbourne has argued that pronouns have the same syntax and semantics as descriptions and Jeff King has argued that complex demonstratives are quantificational expressions.³⁹ I think that both of these fews are plausible, but am certainly in no position to offer a serious defense of them. The point is that it is not as implausible as one might think to suggest that all referring expressions were semantically analogous to definite descriptions (at least in the sense that they allow subordinated readings). As a matter of fact, though, I do want to deny that indexicals are semantically analogous to definite descriptions; so in this section, I will set myself the modest task to showing how the arguments I made in the body of the paper will not work for indexicals.

Do indexicals exhibit the range of readings that names do? (52) is supposed to be a case in which 'I' has a quantificationally subordinate reading. Are there cases in which indexicals have modally subordinate readings. A putative example from the literature⁴⁰ is

(54) [Explaining why I waited to answer the door] You might have been an armed robber

³⁶After all, a big part of the intuitive appeal of the individual constant view of names lies in name's supposed *contrast* with descriptions

³⁷Nunberg. SOURCE. Thanks to Zoltan Szabo for pointing out the relevance of these sentences

³⁸King. 10

³⁹See...and...

⁴⁰King 3

This utterance is clearly not asserting of the addressee that there is a possible world in which he is an armed robber. The thought is that the pronoun is somehow interpreted as within the scope of the modal operator and so the sentence ends up asserting that there is a possible world in which some other addressee is an armed robber (although, as we will see, this can't be quite right either). To me, the most obvious reading of (53) is one on which the kind of modality involved is epistemic and *not* metaphysical. Couldn't I have said much the same thing by saying

(55) For all I knew, you were an armed robber

and here there clearly is no need to give 'you' a quantificationally subordinate reading (though there may well be some delicate *de re/de dicto* issues).

Something similar can be said for (52). It is not entirely clear to me what the force of 'traditionally' is in this sentence - whether we should really think of it as an adverb of quantification. It seems natural to say that (52) is equivalent to something like

(56) According to tradition, I am allowed to order whatever I like for my last meal

and here 'I' does not receive a subordinated reading⁴¹; this sentence merely makes a claim about the relationship between some particular individual and the norms set down by tradition. If we alter the sentence to read

(57) I am usually/always allowed to order whatever I like for my last meal

my sense is that it is significantly harder to hear the intended reading.

But I do not want to spend too much time squabbling about the readings of (52) and (55). Because even if I am wrong about their natural readings, I do not think that they are an example of the sort of phenomenon we examined in the case of names. Let's recall our analysis of subordinated readings; the possibility of a subordinated reading of a definite description (or name) consisted in a certain relationship between two things: a set of subordinating conditions introduced by a quantifier phrase or temporal/modal operator and a descriptive condition associated with a description. If there is an identifiable satisfier of the descriptive material relative to every member of the set of subordinating conditions then a subordinated reading is possible. The utility of this analysis depends on the two elements (the set of conditions and the descriptive material) being, to some extent, independently identifiable. Recall the context which generated a subordinated reading of

(50) One woman chose Hugo

⁴¹Thanks to Irene Heim for pointing this out to me

it was clear that, given contextual restriction of the quantificational domain, the set of subordinating conditions contained every woman in the experiment (and no one else). The name admitted of a subordinated reading because relative to each woman in the experiment, there was an identifiable Hugo (namely, the one who approached her).

Something very different is going on in the Nunberg cases. What these cases seem to show is that if the role that an individual plays (e.g. being the Pope) is especially salient in a context, then referring to that individual can be a way of conveying information about the role. So the explanatory focus is in a very different place. Nunberg readings are available if the right relationship holds between an individual referred to in a context and the conversational purposes — the relationship being that the individual bears a role that is contextually salient in the right way.⁴² In the subordinated readings there is no individual referred to with the name. The important relationship is that between the descriptive condition (context-independently) associated with the name and the conversational purposes.

Notice that in the Nunberg cases, it simply does not matter *which* expression is used to pick out the relevant individual. Imagine the prison guard saying to the condemned man

(58) You are traditionally allowed to order whatever you like for your last meal

or the warden saying to the guard

(59) He is traditionally allowed to order whatever he likes for his last meal

These, I take it, would all be different ways of saying the *same thing*. There is no equivalent possibility with subordinated readings of names. No individual is referred with ‘Hugo’ so there are no *other*, equally felicitous, way of referring. (50) could not express what it does without containing ‘Hugo’ because the meaning of ‘Hugo’ is (part of) what is doing the work. We can start to see how we might insist that subordinated readings of names are revelatory of the meaning names in a way that Nunberg readings are not revelatory of the meaning of indexicals.

Notice that in the case of subordinated readings of names, not only is the meaning of the name essential to explaining the availability of the subordinated interpretation, but is part of the content of the subordinated interpretation. Remember our eccentric friend who strives always to have a girlfriend named Helen and a girlfriend named Jane. When we utter

(44) Next month, Helen will be an investment banker and Jane will be an artist

it is not simply that the property of being a Helen and the property of being a Jane are salient enough in the context to allow a subordinated reading. The reading itself contains those properties. We can see this in the difference between three possible continuations of the story. Imagine

⁴²This is all, of course, hopelessly vague. The point is just to gesture at the structure here.

that in three months from now 1) Our friend has a girlfriend named Helen, a girlfriend named Jane but both are philosophy professors 2) Our friend has a girlfriend named Helen, a girlfriend named Jane but Helen is an artist and Jane is an investment banker 3) Our friend has a girlfriend named Julia and a girlfriend named Mary (i.e. he has abandoned his strange predilection) and Julia is an investment banker and Mary is an artist.

It is clear that these three cases are different ways that what I predicted could fail to be true
⁴³ This shows that what is expressed by (44) is very different from what would be expressed by, say,

(60)

Next month, one girlfriend will be an investment banker, and the other one will be an artist

This is strong evidence, I think, that the subordinated readings of names I have are partly constituted by the *meanings* of the names in them. Nothing equivalent can be said about the Nunberg cases; nothing turns on whether I say (56), (58) or (59). It is not as if the privilege enjoyed by the condemned men which is asserted in (56) has anything whatsoever to do with the fact that he is the speaker of the utterance. The Nunberg readings simply do not result from any substantive interaction between the *meaning* of the referring expressions in them and the conversational common ground. This is strong evidence that they do not have the same import for our semantical theorizing as do subordinated readings of names.⁴⁴

⁴³My own intuition is that in cases (1) and (2) what I predicted was false but that in case (3) something like presupposition failure is going on. Nothing turns on that here though

⁴⁴Interestingly, both King (source) and Nunberg (source) assert that Nunberg-style readings are not available with names. If this were true, I'm not really sure what the consequences would be for my argument. On one hand, this might be evidence that subordinated readings of names and Nunberg-style readings of pronouns are completely different phenomenon — which is half of what I was trying to argue for. On the other hand, this might suggest that the *meaning* (read: character) of pronouns is, in some sense, important to generating the Nunberg-style readings — which is what I was trying to argue against. Luckily, I do not think that King and Nunberg are right to suggest that names can never have such readings. The following example is unique in this chapter for being a real example from the corpus (thanks to Chris Kennedy for recording it). Two Boston Red Sox fans are discussing an enthusiastic jig that star closer Justin Papelbon performed after the Sox won the ALCS. Remarking on how the curse that, for years, had plagued the Sox seems to have lifted, one utters

(61) Twenty years ago, Papelbon would have pulled his hamstring

This, I would suggest, is a Nunberg style reading of 'Papelbon'. The role that Papelbon is playing (i.e. being an important player on the Red Sox) is salient enough that one can convey something about the role by asserting something about Papelbon. (61) is roughly equivalent to

(62) Twenty years ago, the Red Sox' star player would have pulled his hamstring

Notice that the availability of the reading does not depend on the conversational participants assuming that there

1.2 Alternation

Another worry one might have is that the subordinated readings of names that we have canvassed are somehow awkward or forced. This sort of claim might mean one of two things. Firstly, it might simply be an expression of the fact that the objector does not have share the intuitions that I have reported about the relevant sentences in the paper. I will confess to not having much sympathy for this sort of response. The overwhelming majority of informants I have asked have shared the intuitions I report above. I am inclined to think that an inability to hear the subordinated readings of names reveals more about a firmly held philosophical theory than it does about the nature of semantic competence. But this is just *ad hominem* speculation.

The second version of the worry is more nuanced. It would admit that the sentences have the readings that I claim them to have; it would suggest, though, that those readings represent some sort of seamless reconstruction strategy on the part of the audience. We are, after all, clever creatures and we can quickly and seamlessly grasp what a speaker meant to express with all kinds of ill-formed constructions (*an example here would be good*). A way of making this plain, would be note that in nearly all of the examples of subordinated readings of names that we have canvassed, substituting ‘The N’ for a name ‘N’ makes the subordinated reading noticeably easier to hear. Recall

(44) Next month, Helen will be an investment banker and Jane will be an artist

some informants, though they accept and understood (44), also report that a sentence like

(63) Next month, the Helen will be an investment banker and the Jane will be an artist

seems more ‘normal’, or even more ‘correct’.⁴⁵ The suggestion would be, then, that the subordinated readings do not reveal anything about the semantics of names but merely about the audience’s ability to seamlessly reconstruct a (correct) sentence like (60) when they hear an (incorrect) sentence like (44).

I take myself to have shown that the phenomenon of subordinated uses of names is systematic and integrated with other semantic machinery in a way that renders implausible the idea that it is a mere reconstructive strategy. However, unless I have an explanation for the perceived difference between (44) and (60), substantial doubts may remain. There are two broad strategies for accounting for the difference: semantic and pragmatic. The semantic strategy

was a unique salient Papelbon twenty years ago; nor is being a Papelbon, any part of the meaning of what is conveyed by (61). Even if there were a Papelbon on the Red Sox twenty years ago but he was an insignificant player (say, a bullpen catcher), the truth of (61) does not depend on whether *he* would have sprained his ankle twenty years ago but on whether, say, Wade Boggs would have.

⁴⁵I, myself, do not share this intuition. But I have spent so much time with these sentences that I’m not sure my intuitions can always be relied upon.

would involve claiming that a bare occurrences of name ‘N’ does not make exactly the semantic contribution to a sentence containing it as does the definite description ‘The N’.⁴⁶ Absent any strong independent reason to adopt this approach, I think it would be a mistake for a proponent of PV. In the first place, it dramatically reduces the simplicity and economy of the account. The simplicity of the account lies in its ability to explain a bare occurrence of name in terms of two antecedently understood expressions — predicative occurrences of that name and the determiner ‘the’. If we retreat from this claim, bare occurrences of names start to look mysterious; for example, are to take it that they involve a determiner that has no overt expression in English or any other natural language?

Fortunately, I think the pragmatic strategy is promising. Notice that on the present account ‘N’ and ‘The N’ already differ in the following respect: ‘The N’ can be followed by a restrictive relative clause. So I can say ‘The N who I met last week ...’ and have ‘who I met last week’ function to help identify the N in question. If I say ‘N, who I met last week, ...’ ‘who I met last week’ can only function as a kind of comment on the rest of sentence. How could this difference explain the perceived difference between (44) and (60)? Note that given that the bare use cannot be followed by a restrictive relative clause, ‘N’ can only be used in a context in which the property of being an N is salient enough to bear the weight of identifying a relevant individual. Note that, as we saw, *bearing the weight of identifying a relevant individual* can take two very different forms. It can (in non-subordinated cases) mean *bearing the weight of identifying a relevant individual relative to the context of utterance* or (in subordinated cases) it can mean *bearing the weight of identifying a relevant individual relative to each subordinating condition*. The distribution of name-bearing properties and their informational value is, as a matter of fact, such that conversational circumstances in which the property of being an N are capable of bearing identificatory weight relative to the context of utterance are overwhelmingly common while conversational circumstances in which the property of being an N is capable of bearing identificatory weight relative to a set of subordinating conditions are exceedingly rare. This is not a fact about the meaning of words but merely about the epistemological relevance of certain properties. This means that bare occurrences of names are overwhelming associated with non-subordinated uses and that if the property of bearing a name is involved in a subordinated use it is usually because it is explicitly restricted with other material — for example, in a sentence like

(64) Every boy resents the Dave who kicked him

(64) has a subordinated reading without need to appeal to some unusual conversational situation.

⁴⁶One such account, of course, is simply the individual constant approach. On this approach ‘N’ and ‘The N’ make categorically different semantic contributions.

So as a matter of distribution, bare occurrences of names are associated with non-subordinated readings and subordinated readings involving name bearing properties are associated with overt definite articles. This, in itself, I think is enough to explain the sense in which (60) sounds more ‘natural’ than (44). Although ‘N’ and ‘The N’ mean the same thing, the presence of the definite article in ‘The N’ primes the audience for a subordinated interpretation. What is satisfying about this interpretation is that it explains the perceived difference between (60) and (44) without claiming that (44) is *defective* in any way; the explanation relies merely on the statistical abnormality of (44) — a property that we already knew (44) had anyway. This kind of analysis is preferred not simply because it accords with PV; more importantly, I think it captures the intuitions about the sentences better. It is simply not true that sentences involving subordinated reading of bare proper names are heard as defective; an explanation that validates that intuitive judgment is, to that extent, preferable.

1.3 Romanovs and Kennedys

In “The Predicate View of Proper Names” Reinaldo Elugardo considers an objection that is originally due to Steven Boer.⁴⁷ He writes:

Assume, for the sake of discussion, that ‘Romanov’ refers to the famous Russian family dynasty. Then, on that interpretation, sentence (65) can come out true:

(65) Wally Cox (my gardener) is a Romanov⁴⁸

where (65) is to be understood as reporting ‘an exciting fact revealed by recent historical investigations’ (Boer, 1975. 390). By hypothesis, Cox never got, and never will get, the name ‘Romanov’. So (65) comes out false on Burge’s analysis, which is the wrong result. This a [sic] case, then, where a name is true of an individual who is not given the name in any appropriate way.

The idea here is that there clearly is a predicative use of proper names that is not metalinguistic in the way that Burge identified. When one says, for example, that Maria Shriver is a Kennedy, one is certainly not claiming that she is, or has been, called ‘Kennedy’ by anyone (though, of course, she may have been). Does this throw Burge’s analysis into question? Is this evidence that bare occurrences of proper names are not essentially meta-linguistic? One might think that the proper response from PV would be concessive: admit that meaning of proper names should be broadened to include people who are descended from people who bear the name in question.⁴⁹ But this would be exactly the wrong thing to do. It is perfectly plain that one cannot

⁴⁷source

⁴⁹Elugardo seems to adopt this strategy. SOURCE

use a sentence like

(66) Romanov is a gardener

to say something true of Wally Cox after learning that he descended from the Romanovs. Whatever we end up saying about predicative occurrences, it is clear that bare occurrences *are* essentially metalinguistic. But this means that if we complicate the account of predicative occurrences beyond the simple metalinguistic story, we will lose one of the central features of the predicate account. Recall that it was central to the explanatory power of the account that the meaning of bare occurrences is treated as a function of the meaning of predicative occurrences. Broadening the meaning of predicative occurrences would be abandoning that central explanatory feature.

A response left to PV is to claim that predicative occurrences of, say, 'Romanov' are *ambiguous* between meta-linguistic and genealogical uses. This would be to reject the demand for a unified semantic treatment of 'Romanov' as it appears in (65) and (66). We can retain the tight connection between the meaning of a metalinguistic predicative use of Romanov and a bare use of Romanov, but allow that there is another predicate one can express with the 'Romanov'.

It would be fair to ask, at this point, whether this analysis is simply *ad hoc*. Positing ambiguities is often a way of trying to save a failing theory, and the need to posit an ambiguity was one of the faults that we originally found with the individual constant view. This is a serious worry, but I think that there is *independent* evidence that predicative occurrences of proper names are ambiguous between a metalinguistic and a genealogical uses. One can truly assert

(67) You cannot become a Kennedy simply by changing your name

but this suggests that the meaning of a predicative occurrence of Kennedy is not some disjunction of meta-linguistic and genealogical properties. Unfortunately, the intuitive truth of (67) is not, by itself, particularly strong evidence for the idea that 'Kennedy' is ambiguous between a meta-linguistic and genealogical meaning. After all, one can say things like

(68) You cannot become an American simply by changing your citizenship

and this is not taken to obviously show that 'American' is ambiguous.⁵⁰ It seems as though the truth of (67) might be taken to reflect a metonymic or metaphorical use.⁵¹ My impression is that (67) is not most naturally read as anything other than literally true but this is difficult to argue for. I am not, for instance, hopeful that one could alter the sentence to clearly force a literal meaning (for instance by adding 'literally'). A more promising, though somewhat vague,

⁵⁰Although, it may be taken to show this. Say something about Chris's stuff about gradable adjectives.

⁵¹Burge himself notes the need to distinguish literal and metaphorical predicative uses of names. SOURCE

strategy would be to identify some feature that metaphorical uses of language have and see if on a natural reading (67) has that feature. For example, it seems that on a pretty standard story about how what metaphors mean, a metaphor invites the speaker to draw an open-ended number of comparisons with the literal reference of the metaphorically used expression. So, for instance, hearing (68) as metaphorical would mean considering some open-ended list of 'typically' American qualities. Importantly, then, what one would understand (68) as expressing would depend heavily on what qualities one thinks of as typically American. My sense is that there is a reading on which (67) is true and does not have this feature. My sense is that there is a reading of (67) on which simply asserts that one cannot become part of a certain genealogical line simply by changing one's name. This is not a completely satisfactory response, but is the best I can do for now.

1.4 Russell's Principle

My impression is that some philosophers might take the following sort of claim as an argument in favour of the individual constant view.⁵²

The meaning of a name is different from that of, e.g., a definite description in that one can understand the meaning of a sentence containing a definite description without knowing which object satisfies that description whereas one cannot understand the meaning of a sentence containing a name without knowing who the name is a name of.

It strikes me that this is just a *restatement* of the individual constant view rather than an argument in its favour. If the meaning of a name, O, just is the object X, not knowing that O refers to X amounts to not knowing the meaning of O. Combined with the unobjectionable premise that one knows the meaning of a sentence only if one knows the meaning of all of the words that comprise it, **Russell's Principle** simply follows. It's worth noting that Russell took the above condition *not* as a claim about (what we call) proper names in natural languages but as a *definition* of 'logically proper names'.

Even if the above principle is simply a reformulation of the individual constant view, we might think that reformulating the view in this way points us toward another line of argument for it. If we, as language-users, have strong intuitions about the conditions under which someone could be said to understand a sentence that line up with **Russell's Principle**, this might be taken as evidence in favour of the individual constant account. Russell himself thought that our intuitions about proper names do not line up with this condition and this is, of course, why he thought that ordinary proper names were *not* logically proper names. Still, given a limited

⁵²Recanati? Neale?

range of examples, the principle might seem to cut in the way that the individual constant theorist claims it does. Compare

(69) Alfred is a Democrat

with

(70) The current President of the United States is a Democrat

There is clearly a sense in which one can understand (70) without knowing who the current president is and one cannot understand (69) without knowing who Alfred is. It is not clear, though, that his sense of ‘understand’ is of particular theoretical interest. The obvious thing to say here is that the source of this intuition is the fact that ‘Alfred’ is context sensitive in a way that ‘the current president of the United States’ is not. If we contrasted (69) with

(71) The man is a democrat

I doubt that we could use our pre-theoretical intuitions about **Russell’s Principle** to provide evidence for the individual constant view. In fact, I don’t think that our pre-theoretical intuitions about when someone does or does not understand a sentence are going to be particularly helpful here at all. It’s just that our intuitions about understanding track importantly different features of language. For example, it seems natural to say that the intuition we have about the difference between (69) and (70) is a result of the fact that it is pretty clear, absent a context of utterance, what the *point* of uttering (70) might be, while it is not at all clear, absent a context of utterance, what the point of uttering (69) would be. (THIS SHOULD BE CLEARER). It is our job as semantic theorists to disentangle the different sources of these intuitions but this task just *is* the task of, among other things, figuring out what the semantic contribution of a name is. Different accounts of the meaning of names will entail different accounts of the systematicity of our semantic intuitions, but there is no reason to suppose that the individual constant view has some privileged diagnosis of the source our intuitions about understanding.

There is, though, another worry behind the scenes here. The main part of the argumentative strategy of this chapter has been to argue that bare names exhibit all of the ambiguities associated with definite descriptions. I have focused, though, only on those ambiguities that can be captured with the subordinated/non-subordinated distinction.⁵³ I have thus ignored one of the most celebrated sorts of ambiguity thought to be associated with definite descriptions — namely ambiguity with negation. Russell in ‘On Denoting’ famously argued that

(72) The present King of France is not bald

⁵³And not even with all of those. I have not mentioned readings on which names are subordinated to quantificational adverbs (e.g. ‘usually’). These behave the same way as other subordinating devices. I’ll leave a demonstration of this as an exercise for the reader

is ambiguous depending on the relative scope of the negation and the definite description. On one reading (where the negation gets wide scope) the sentence can be true despite France not having a reigning monarch. On the other reading (where the description gets wide scope) the absence of a reigning monarch renders the sentence false.

If bare occurrences of names do not exhibit scope interaction with negation, this might be taken as evidence that they were not properly assimilated to definite descriptions. The problem with this argument is that, *pace* Russell, there is no real evidence that definite descriptions display scope interactions with negation. I do not have space to argue this point here, suffice it to note that it is not generally accepted that (72) is ambiguous in the way that Russell imagined.⁵⁴ The fundamental worry in the background here, though, is about the relationship between non-denoting descriptions and vacuous names. There is a basic thought that sentences containing non-denoting expressions are not thereby rendered meaningless and can even be true; in contrast, sentences involving vacuous names are thought to be meaningless. Part of this thought comes from the misunderstanding of the status of **Russell's Principle** that I noted in the previous section. The other source of this thought is the idea that definite descriptions are quantificational. If they are quantificational, then there should be sentences in which they get narrow scope relative to negation; and in such sentences, the fact that the description is non-denoting will render the sentence true (provided there are no interfering factors, like belief reports, etc).⁵⁵ If there are no sentences in which names get a similar treatment, this might be taken as evidence that names are semantically analogous to definite descriptions. This objection raises a number of complicated issues that we are not yet in a position to tackle. We will return to this issue in a later chapter, when we have the resources to fully explore it. Suffice it to say for now: try to imagine a context in which

(74) The man didn't come to the party

might be truly uttered despite there being no existing salient man. If you can do that, you should have no problem imagining a true utterance of

(75) Dave didn't come to the party

despite there being no existent salient Dave.

⁵⁴Cite pro...con...

⁵⁵It should be noted that this line of thought is already misguided. There are quantifiers which clearly do not exhibit this kind of ambiguity. The sentence

(73) Some girl didn't come to the party

clearly does not have a wide-scope reading for the negation. The sentence cannot mean that *no girl went to the party*. Thanks to Chris