Descriptions as Variables

Paolo Santorio
UNIVERSITY OF LEEDS/AUSTRALIAN NATIONAL UNIVERSITY

1 Introduction

In *On Denoting* (1905), Russell proposed a view of definite descriptions that has stayed alive and healthy for over a century. This view is effectively summed up in four words: *the* is a quantifier. In its modern embodiments (based on Barwise and Cooper's (1981) Generalized Quantifier Theory) the view has a syntactic and a semantic core. On the syntax side, it says that a sentence involving a definite description like (1) shares its structure with (say) a universally quantified sentence like (2):

(1) The richest man on Wall Street has no moral scruples.

(2) Every stockbroker has a low moral character.

Roughly, this structure involves *the* and *every* (lumped together under the category of *determiners*) working as two-place functions and taking as arguments two predicates, one of which works as the *restrictor* of the domain of quantification:

> The [richest man on Wall Street] [has no moral scruples]
>
> Every [stockbroker] [has a low moral character]

On the semantics side, it says that the truth-conditions of a statement involving a definite description are existential. For example, for the case of (1):

> There is a unique richest man on Wall Street, and that man has no moral scruples.

This view of definite descriptions naturally dovetails with a quantifier analysis of indefinite ones. On the latter, the indefinite article *a* is a natural language counterpart of the existential quantifier in first-order logic. Hence a sentence like (3) shares its syntactic structure with (2) and has existential truth-conditions:

(3) A stockbroker fled with all my money.

I have an argument against Russell's view. The argument supports a different picture: descriptions, definite and indefinite alike, can behave syntactically and semantically like

variables. This basic idea can be implemented in very different systematic analyses: whichever way one goes, the Russellian view is going to fail in at least one of its two components, if not both.

Over the last decades, versions of the variable view of descriptions have been developed and defended within a semantic tradition that rivals classical truth-conditional frameworks, that of Discourse Representation Theory. So my claim isn't new. What is new is the way of defending it. In the past, the variable view has been sold as part of a package deal encompassing a general switch of semantic frameworks. My argument does not depend on endorsing any such package deal. The argument centers rather on attitude reports: I argue that we should recognize a new reading of descriptions under attitude reports, which I call 'singular opaque'. In turn, the existence of this reading cannot be explained on the traditional Russellian view of descriptions, and demands a switch to the variable view. So I'm going to use arguments from one traditional domain of philosophy of language, attitude reports, to establish a claim about another traditional domain, descriptions. Along the way, some interesting new facts about the former domain will come to light.

Here is the structure of the paper. In section 2, I briefly summarize the state of the debate on *de re* and *de dicto* reports, showing how the quantifier view of descriptions fits in with current accounts. In section 3, I argue for the existence of singular opaque readings. In section 4, I sketch informally two different directions for developing and implementing the variable view. Finally, in section 5, I consider an alternative proposal for handling similar data, recently advanced by Zoltan Gendler-Szabó.

Before proceeding, a terminological note: in line with the linguistics literature, I will use 'definite phrase' and 'indefinite phrase' (or more simply, 'definite' and 'indefinite') to pick out, respectively, definite and indefinite descriptions. I will also use the blanket label 'descriptive phrase' (or 'description') to pick out indiscriminately descriptions of both kinds.

2 Setup: attitude reports and descriptive phrases

2.1 Quantifiers and propositional attitudes

The quantifier analysis of descriptions dovetails with traditional views of attitude reports.[1] Consider:

[1] I assume in the background a view on which *the* and *a* are generalized quantifiers in the style of Barwise & Cooper (1981). Hence *the* and *a*, on a par with *every*, *some*, *most*, *many*, and so on, are two-place functions taking as argument two predicates and returning a proposition. Notice that *the* is a controversial case: several linguists who buy into generalized quantifier theory still want to treat *the* as an operator which, combined with a predicate, refers to an individual, in the style of Frege (1892/1997). (For example,

(4) Stacey intends to marry a stockbroker.

At least since Quine (1956), philosophers and linguists agree that sentences like (4) have two readings. The first—the *de dicto* reading—is true if Stacey intends to marry some stockbroker or other, with no need of her having settled on a particular contender. The second—the *de re* reading—is true if Stacey has a marriage scheme concerning a particular individual who happens to be a stockbroker, whether or not Stacey is aware of their profession.

The *de re*/*de dicto* ambiguity thus characterized is elegantly accounted for by the quantifier view. I'm going to argue that, nevertheless, this account is too simplistic. (4) has readings that cannot be captured by a quantificational semantics. As a result, we must switch to a referential analysis. Before doing so, it will be useful to do some groundclearing to show exactly how I diverge from the orthodoxy.

2.2 Attitude reports: three standard readings

Early analyses of attitude reports resorted to the tools of intensional logics, and especially modal logic. On these analyses, the interpretation of a sentence is relativized to one or more parameters: for example, a possible world. Attitude verbs, on a par with boxes and diamonds of modal logic, have the effect of shifting the world at which a clause is evaluated:[2]

> ⌜*S* believes that ϕ⌝ is true iff for all worlds w' compatible with what *S* believes, ϕ is true at w'

Interestingly, the quantificational view of descriptions, when combined with the intensional setup, predicts exactly the *de re*/*de dicto* distinction as it is described by Quine. Consider again:

(4) Stacey intends to marry a stockbroker.

(4) can be syntactically disambiguated in two ways. The indefinite *a stockbroker* may take narrow scope with respect to the attitude verb:[3]

> S intends [to marry [$\exists_x$ stockbroker(x)]]

see the treatment of definites in Heim & Kratzer (1998).) All my main points apply to a view of this sort; I stick to a uniform quantificational treatment just for expository convenience.

[2]The idea of treating attitude verbs as modal quantifiers is due to Hintikka; see his (1962) and (1969).

[3]In line with Heim & Kratzer (1998), I'm assuming a background syntax on which quantified phrases like *a stockbroker* undergo movement leaving behind a trace (which I represent as 'x'). But I take the liberty of not representing this kind of movement when it's not important for my purposes (as in the schematic representation just below this footnote).

Or it may take wide scope with respect to it:

$[\exists_x$ stockbroker$(x)]$ [S intends [to marry x]]

These two syntactic possibilities give rise, respectively, to the classical *de dicto* and *de re* readings. If the indefinite has narrow scope with respect to the verb, then it is evaluated at the world parameter introduced by the latter. This generates the classical *de dicto* reading. If the indefinite has wide scope, it is evaluated at the world of the context. Then we get the *de re* reading.

It is known that the classical view is problematic. The intensional paradigm doesn't give us enough expressive power.[4] To see this, notice that (4) has a further reading (as was pointed out by Fodor (1970)). Consider the following scenario:

> Stacey, a resident of Wall Street, is fascinated by a group of men who gather regularly in a local bar and that she takes to be lawyers. She forms the plan of marrying one of them, though she doesn't settle which one. Unbeknownst to her, all the men in the group are actually stockbrokers.

In this scenario, (4) has a true reading. But this is not one of the classical *de dicto* or *de re* readings. Stacey doesn't intend to marry some stockbroker or other, since she takes the individuals in the relevant group to be lawyers. She also fails to have marriage plans concerning a particular individual: anyone in the group would do, as far as her desires are concerned. What she wants is to marry an individual within a certain group, though not a particular person; and it just so happens that all people in that group are stockbrokers. The existence of a third reading, together with a number of other puzzles[5], has produced a shift in the way modal talk is modeled in semantics. Rather than relativizing the interpretation of a sentence to a world parameter, we introduce reference to worlds directly in the object language. This generates a shift to an extensional framework for modeling modal talk, on which natural languages like English are assumed to contain covert variables ranging over possible worlds.

I mention the third reading and the switch to extensional frameworks only to set them aside. (For completeness, I show how extensional systems account for the third reading in a footnote.[6]) What matters for me is not discussing the three readings that are already acknowledged, but rather arguing that there is yet another reading.

[4] At least, the intensional paradigm as it stands. Cresswell (1990) points out that, by supplementing a language with enough intensional operators, we obtain a system with expressive power analogous to extensional systems.

[5] For a useful summary of the problems for the intensional theory, see chapter 1 of Keshet (2008). For discussions and criticisms of the extensional framework, see also (among many) Percus (2000) and Romoli & Sudo (2009).

[6] Using variables ranging over worlds allows us to let the indefinite have narrow scope with respect to the attitude verb, and at the same time evaluate the noun *stockbroker* with respect to the actual world. In schematic terms, here is the logical form of (4):

2.3 The claim: a fourth reading

These are the three readings of (4) expected on current accounts:

(a) *Classical de dicto.* Stacey intends to marry some stockbroker or other.

(b) *Classical de re.* There is a particular individual who is a stockbroker, and Stacey intends to marry that individual.

(c) *Narrow-scope de re*: Stacey intends to marry someone within a certain group, though not a particular person. Those people all happen to be stockbrokers.

Notice that these readings can be classified along two dimensions. First, they differ on whether the predicate within the description (for example, *stockbroker* in *a stockbroker*) receives a so-called opaque or transparent interpretation—i.e., whether it is evaluated at the world introduced by the attitude verb or the actual world. Second, they differ in whether they ascribe the subject an attitude about a particular individual or not. Following Fodor (1970), I call readings of the former kind 'specific', and readings of the latter kind 'unspecific'. Here is how the three acknowledged readings fare with regard to these two axes:

	OPAQUE	TRANSPARENT
UNSPECIFIC	classical *de dicto*	narrow scope *de re*
SPECIFIC	?	classical *de re*

As the table shows, there is one possible combination that is not exemplified by current views about attitude reports. There seems to be no readings of descriptions that are opaque—the predicate is evaluated at the world introduced by the attitude verb—and at the same time specific—the attitude ascribed is about a specific individual. But why can't we get that reading?

Why not, indeed? In the next section, I'm going to argue precisely that such a reading exists. To see what this reading amounts to, consider the following scenario:

> Jason sees a man wearing an elegant business suit jump off the edge of a tall building overlooking Wall Street. Jason assumes that the individual is a suicidal stockbroker. But the man is actually a professional stuntman who's training for his next performance.

And now take the report:

(5) Jason believes that a stockbroker jumped off a building.

Stacey intends [to marry w_1[a stockbroker $w_@$]]

I claim that (5) has a true reading with the following truth-conditions:

> Jason has a belief concerning a certain man; the belief is that that man, whom he takes to be a stockbroker, jumped off a building.

The existence of the fourth reading—let me call it 'singular opaque' reading—has important consequences for the semantics of descriptions. The claim is that *a stockbroker* in (5) is somehow semantically linked to actual world individuals. Moreover, this link is fully independent of the descriptive content of the description. This cannot be captured by the traditional quantifier view. On this view, the fact that *stockbroker* receives an opaque interpretation forces the whole description to take narrow scope with respect to the attitude verb; schematically:

> Jason believes [$\exists_x$ stockbroker(x) jumped-off-a-building(x)]

In this situation, the quantifier simply cannot range over actual individuals.

My diagnosis is that the quantifier picture needs amendment. The fourth reading shows that descriptions can have a semantic link to individuals, one that is independent of their descriptive content. This is captured in a simple and straightforward way by the variable view. I will say more about how this view accounts for the fourth reading in section 4, but it's helpful to give a preview. Schematically, (5) is represented as follows:

> Jason believes [x_1(stockbroker) jumped off a building].

Crucially, *a stockbroker* is analyzed as involving a free variable. Given that it's free, this variable can be semantically linked to objects that are outside Jason's attitude worlds, including actual stockbrokers. Simplifying, *a stockbroker* behaves similarly to a deictic occurrence of *he* (i.e., an occurence of *he* accompanied by pointing) in the same position, as in

> Jason_1 believes that he_5 jumped off a building.

The precise nature of the semantic link between the description and the individual Jason saw will vary depending on the specific analysis that one picks. What matters for now is the observation that the variable view has the resources to do very simply something that is impossible on the traditional quantifier view.

But is there no fix for the quantifier view? There might be one. In recent work that starts from the analysis of closely related data, Zoltan Gendler Szabó (2010) has suggested that descriptive phrases like *a stockbroker* can have a 'split' underlying syntax. This means that, at the level of logical form, the indefinite article *a* can separate from the predicate *stockbroker*. The former takes wide scope with respect to the attitude verb, the latter narrow scope:

$\exists_x$ Jason believes [stockbroker(x) jumped-off-a-building(x)]

The split quantifier view is an interesting possibility, but (I claim) suffers from severe problems, and ultimately should be discarded. For the moment, I will set it aside; I will come back to discussing it in section 5.

In passing, let me notice that I'm not the first to argue for the existence of a fourth reading. Fodor (1970) presented arguments for what is essentially the same claim, though couched in different terms. (Though Fodor doesn't seem to worry that the fourth reading requires some overhauling of our standard semantics for descriptions.) Fodor's arguments have turned out to be inconclusive and a consensus has emerged that there is no singular opaque reading. My aim in this paper is to reverse this consensus and draw some consequences.[7]

3 Singular opaque readings

In this section I argue that sentences like (5) can have a singular opaque reading, i.e., they can be read as having the truth-conditions in (6):

(5) Jason believes that a stockbroker jumped off a building.

(6) Jason has a belief concerning a certain man; the belief is that that man, whom he takes to be a stockbroker, jumped off a building.

Proving that the fourth reading actually exists is hard. Notice that the singular opaque reading entails the classical *de dicto* reading. If the singular opaque reading is true, then the subject believes the proposition expressed by the that-clause; but this is enough to make the classical *de dicto* reading true as well. Hence one cannot argue for the fourth reading simply by producing a scenario where that reading is true. Any such scenario won't be able to distinguish between the fourth and the first reading.

Hence an argument for the fourth reading must take a circuitous route. This is exactly the kind of argument I have. I consider the interaction of attitude reports and anaphora. Anaphora is a traditional battlefield for theories of descriptions and reference, but the connection between anaphora and attitude reports has been overlooked so far. I will try to show that we can learn a lot by investigating this area.

[7] Fodor connects the two axes of variations to the two standard tests for opaque contexts, namely failure of substitutivity of coreferential terms and and failure of existential generalization. Here I want to remain neutral on this point.

3.1 The argument

Consider again the stuntman scenario:

> Jason sees a man wearing an elegant business suit jump off the edge of a tall building overlooking Wall Street. Jason assumes that the individual is a suicidal stockbroker. But the man is actually a professional stuntman who's training for his next performance.

The following discourse is true in this scenario:[8]

(7) Jason believes that a stockbroker jumped off a building. But he was actually not a stockbroker, but a stuntman.

Here is my basic argument. The indefinite *a stockbroker* in (7) must be given an opaque reading. That is, the noun *stockbroker* must be interpreted with respect to the world introduced by the attitude verb, rather than the actual world. This should be obvious, since no actual stockbrokers have anything to do with Jason's belief. At the same time, the indefinite refers to a specific individual. If it didn't, then the speaker couldn't use the anaphoric pronoun *he*. So, on the true reading of (7), the indefinite *a stockbroker* has a singular opaque reading.

It's useful to spell out the argument in detail:

(P1) The indefinite *a stockbroker* in (7) has an opaque reading.

(P2) *He* in (7) is an anaphoric pronoun referring to the man Jason saw, and whose antecedent is the indefinite.

(P3) If *a stockbroker* did not refer to the man Jason saw, the anaphoric link between it and *he* would not be available.

(P4) *A stockbroker* in (7) refers to the man Jason saw. (From (P2), (P3))

(C) *A stockbroker* in (7) has a singular opaque occurrence. (From (P1), (P4))

The argument is simple. The bulk of the work consists in defusing alternative strategies of accounting for (7). Two such resistance strategies are suggested by current literature about anaphora. Each of them focuses on denying one between (P2) and (P3), which are the controversial premises. In the remainder of this section, I discuss these strategies.

3.2 The deictic strategy

The first resistance strategy involves denying that *he* in (7) is a genuinely anaphoric pronoun. Pronouns like *he*, when they're not bound by a quantifier, can work in two ways:

[8]For present purposes, I take discourses to be just concatenations of sentences.

they may be devices of anaphora, or they may be used deictically (often accompanied by pointing) to refer to elements of the context. The claim is that, in (7), *he* works in the latter way. The report in (7) only has a (true) classical *de dicto* reading, the argument goes. At the same time, it raises to salience the fact that there is a specific individual that is causally responsible for Jason's belief. *He* refers to the individual made salient in this way.

The deictic strategy fails, for two reasons. The first was originally pointed out by Heim ((1982) and (1990); see also Kadmon (1987)). She observed that pronouns like *he* in (7) require linguistic antecedents of a very specific kind. Even when certain individuals are made very salient in the context, pronouns cannot pick up on them in absence of the appropriate linguistic antecedent. In short, pronouns like *he* in (7) require a *formal link* obtaining between them and their antecedent.[9] To see how this applies to our case, consider the following variants of (7):

> Jason sees three men wearing business suits jump off a building; two are wearing a parachute. Jason assumes that they're all stockbrokers and that the one with no parachute must have died. But they are all stuntmen who manage to land safely.

(8) Jason believes that three men have jumped off a building and only one has died. He has actually survived too.

(9) #Jason believes that three men have jumped off a building and only two have survived. He has actually survived too.

The opening clauses in (8) and (9) are truth-conditionally equivalent. Yet the anaphora in (8) is successful, while the one in (9) isn't. This shows that a pragmatic notion of salience that doesn't track syntactic facts cannot explain the functioning of anaphora. *He* in (8) and (9), as well as in (7), requires a linguistic antecedent.

The second reason to discard the deictic strategy is that there are examples analogous to (7) where the relevant pronoun is bound. Consider the following scenario:

> We have decided to run a semantics experiment on Wall Street. We take a number of subjects and put them through a Jason-style scenario: each of them observes a stuntman dressed up as a stockbroker jumping off a building. (To make the case cleaner, pretend that we use a different stuntman for each subject.) We succeed in fooling all of our subject.

[9] Heim's discussion is focused on the following examples, which she attributes to Barbara Partee:

(a) I dropped ten marbles and found all of them, except for one. It is probably under the sofa.

(b) ??I dropped ten marbles and found only nine of them. It is probably under the sofa.

We can report the results of our experiments as follows:

(10) Our subjects always believed that a stockbroker had just committed suicide in front of them, though he was only a stuntman performing a trick.

The deictic strategy is a non-starter for an analysis of (10). Clearly, *he* in (10) is bound and ranges over the stuntmen that experimental subjects have seen. Hence this occurrence of *he* is not a referential expression at all. But it is essential to the deictic view that the relevant occurrences of *he* be used referentially; thus the deictic strategy has no way to account for cases like (10).

3.3 The D-type strategy

The second resistance strategy involves claiming that *he* in (7) is a so-called D-type pronoun.[10] D-type accounts have emerged in the semantics literature to deal with particularly problematic instances of anaphora. The essential idea behind D-type approaches is that anaphoric pronouns are a kind of covert definite descriptions. For example, the pronoun *it* in (11) is shorthand for (something like) *the donkey the man owns*:

(11) Every man who owns a donkey beats it.

With regard to our case, the strategy is again to claim that (7) has a true classical *de dicto* reading. The anaphora is explained by the fact that *he* is a covert description involving reference to the individual causally responsible for Jason's beliefs. In essence, the idea is that (7) is synonymous with something like:

(12) Jason believes that a stockbroker jumped off a building. But *the person responsible for Jason's belief/the person Jason saw/etc.* was actually not a stockbroker, but a stuntman.

The D-type strategy is currently a very live contender in debates about anaphora. But it lends no help in this case. In fact, just cases like (7) seem to pose a new problem for D-type theories in general. Let me explain.

D-type theories are designed to respond to the problem of the formal link pointed out by Heim. Consider the pair:

(13) Jason's personal banker has a wife. She is a stockbroker.

(14) #Jason's personal banker is married. She is a stockbroker.

An account of anaphora must predict that the occurrence of *she* is felicitous in (13) but not in (14). Now, if we treat *she* as just any covert description, we seem to have the

[10] See, among many, Heim (1990), Neale (1990), and Elbourne (2005).

resources for predicting that both (13) and (14) are good sentences, contrary to fact. For example, we could get that in both cases the covert description is *Jason's wife*:

(15) Jason's personal banker has a wife. Jason's wife is a stockbroker.

(16) Jason's personal banker is married. Jason's wife is a stockbroker.

Hence, if we want to predict the asymmetry between (13) and (14), we must impose some extra constraints preventing (14) to be interpreted as (16). The basic idea shared by all D-type accounts is that the description must be reconstructed from the linguistic material of an antecedent descriptive phrase. For example, *she* in (13) must pick up the noun phrase from the indefinite *a wife*. Conversely, (14) is infelicitous because there is no appropriate descriptive phrase to work as the antecedent for *she*.

I've been speaking informally of pronouns 'picking up' linguistic material from their antecedent. The details of this process are actually quite complex and make an important difference to the final form of D-type theories.[11] But these differences are irrelevant for my purposes. All D-type accounts have the same problem in accounting for (7). The reason is simple: the linguistic material of the antecedent is just the wrong thing to use to build a description that denotes the relevant object. All we could pick up from the antecedent of *he* is the noun phrase *stockbroker*. As a result, by using a D-type account we would make (7) synonymous with:

(17) Jason believes that a stockbroker jumped off a building. But *the stockbroker* was not a stockbroker, but a stuntman.

(17) is nonsense. The problem is that the D-type pronoun picks up linguistic material that is evaluated at the world parameter introduced by the attitude verb and uses it in a position where it must be evaluated at the actual world. This creates trouble: *he* in (7) cannot be construed as *the stockbroker* or a similar description, simply because the relevant person isn't a stockbroker. Hence just the fact that the indefinite phrase has an opaque reading prevents it from providing linguistic material to construct a suitable description.

Perhaps there are ways out for the D-type theorist. I cannot see any promising ones. Let me go through the three best options I can envisage and show why all of them have serious problems.

(1) The first option is claiming that the pronoun can pick up the predicate from the verb, rather than from the indefinite. So *he* in (7) should be construed as *the jumper*, or something along those lines. But this won't work; we could tweak the example in a way that neither the indefinite nor the verb actually apply to the individual, and still get

[11] In particular, D-type approaches diverge on whether the 'picking up' process is semantic (as in Chierchia (1992)) or syntactic (as in Heim (1990) and Elbourne (2005)).

successful anaphora:[12]

(18) Jason believes that a stockbroker jumped off a building. But he was not a stockbroker, but a stuntman; and he didn't even jump, he was hanging via some thin wires from a helicopter.

(2) The second option involves relaxing the constraints attached to the D-type view. One idea would be to preserve the connection with the antecedent description, but allow that some extra material be added. Hence *he* in (7) might be interpreted as, say, *the alleged stockbroker*. This is also problematic. In the first place, it seems just *ad hoc*. Moreover, the resulting theory seems again too unconstrained. Consider yet another variant on the stuntman scenario:

> While walking along Wall Street, Jason has a veridical-seeming hallucination of a stockbroker jumping off the edge of a building. At the same time, and by a striking coincidence, a stuntman dressed up as a stockbroker jumps off a building in a way such that, if Jason had been actually seeing, the scene in front of his eyes would have been exactly the same.

In this modified scenario, (7) is inappropriate. Yet surely we can plug in the place of *he* some description involving the predicate *stockbroker* that would make the sentence appropriate. *The alleged stockbroker who jumped*, or *the man Jason would have taken for a stockbroker if he had seen it*, are two good candidates. Any plausible D-type view which opted for this solution would owe us an explanation of what makes the two cases different.

(3) The third option involves being faithful to the syntactic material of the sentence, but allowing a more complex algorithm to generate the relevant description. The idea is that we can build a description that involves the attitude verb: for example, for the case of (7), *the individual that Jason takes to be a stockbroker*.[13]

The main objection to this view is that it requires the deployment of heavy-handed resources just to handle the specific kind of anaphora exemplified by (7). To see this, notice that we will need a specific grammatical rule to license anaphora in cases when attitude reports are involved. This rule might take the following form, on a first stab:

> When you have a sentence of the form

[12] Also, Heim (1990) notices that mention of the noun phrase earlier in the sentence is not sufficient for successful anaphora; it is necessary that the noun phrase be part of an indefinite or a definite noun phrase. This also concerns noun phrases embedded under attitude verbs; consider:

> ??Jason believes that the husband-wife relation requires trust. As a husband, he always trusts her.

It's perfectly clear what is meant, yet the occurrence of *her* is still infelicitous.

[13] This idea was strongly pushed by Jenn Wang (in written comments) and Mark Richard (in Q&A). Many thanks to both of them.

⌜Subj AttVerb [. . . [*a*/*the*[Predicate]] . . .]⌝, you may:

(a) move the indefinite or definite article in the descriptive phrase out of the scope of the attitude verb;

(b) supplement it with a 'generic' restrictor (like *thing* or *individual*);

(c) use the resulting phrase as antecedent for D-type pronouns to follow.

This rule seems to yield a suitable linguistic antecedent for the D-type pronoun. But notice what we had to do: we had to postulate a new kind of grammatical rule involving a special syntactic transformation (movement plus addition of a restrictor) just to account for anaphora in a very specific class of discourses.[14] This seems obviously *ad hoc*, unless it can be motivated on independent grounds. Why should grammar contain special complex rules for anaphora in a particular category of sentences?

I also have a second objection. There seems to be something conceptually confused with the proposal. Recall: the D-type view assumes that the report in (7) only has a classical *de dicto* reading. Hence the report doesn't say or imply that there actually is an individual that Jason's belief is about. For all it says, Jason might be hallucinating and there might be no object out there. Yet the new stipulated rule allows us to derive a description that presupposes the existence of such an individual. It would quite strange if grammar somehow made available a description denoting an individual, in a context in which nothing previously said entails that that individual exists.

I conclude that the D-type theory—at least, the D-type theory in its present form or extended in straightforward ways—cannot handle the kind of anaphora exemplified by (7). Perhaps more sophisticated descendants of the view will be successful. But, at the very least, these arguments place the burden of proof on the D-type theorist. Until we're shown that there is a viable, non-stipulative D-type account of (7), we should prefer an alternative analysis.

3.4 Recap

I have argued that descriptions in attitude contexts have singular opaque readings. The argument is simple: we can only make sense of the anaphoric link in (7) if we assume that the indefinite occurring in it has this reading. Other routes—based on denying that there is anaphora, or accounting for the anaphoric link in other ways—are unsuccessful.

[14]Incidentally, the rule as I stated is actually still in simplified form. We would also need more explicit provisions about how the binding works, once we move the article. (Essentially, we will need exactly the moves that Gendler Szabó (2010) makes in trying to specify the split quantifier account.) Moreover, some further provisions seem needed for the case of counterfactual attitude verbs like *hope* or *wish*, which aren't accommodated properly by a rule of this kind.

At the end of the last section, I claimed that singular opaque readings have important repercussions for the semantics of descriptions. These readings show that descriptions may be semantically linked to an object, and that this link is independent of their descriptive content. Let me sketch how this idea can be implemented in a semantics.

4 Descriptions as variables

4.1 The variable idea

A variable view of description will naturally take its cue from standard accounts of pronouns like *she* and *he*, which are also treated as variables. More precisely: syntactically, *she* and *he* are taken to be variables equipped with a subscripted index; semantically, they are assigned a referent via an assignment, namely a function which pairs indices with objects. As a sample, here is the denotation of *she*:[15]

$$[\![\text{she}]\!]^g = [\![x_j]\!]^g = g(j)$$

(Read: *she*, namely the variable with subscript *j*, denotes the individual that the assignment pairs with the subscript *j*. The double brackets '$[\![\cdot]\!]$' denote the interpretation function, namely the function that maps object language expressions into their denotations.) The suggestion is that we construe descriptions on an analogous model.

There are many routes to a variable-based account of descriptions, yielding very different syntactic and semantic treatments. But all of them share a claim about syntax: descriptions like *a stockbroker* in (7) involve a variable ranging over individuals. This common core is sufficient to accommodate the puzzle I've focused on in this paper. Consider again (7), repeated below:

(7) Jason believes that a stockbroker jumped off a building. But he was actually not a stockbroker, but a stuntman.

Schematically, (7) gets the following structure:

> Jason believes that [x_1(stockbroker) jumped off a building]. But x_1 was actually not a stockbroker, but a stuntman.

The explanation for the felicity of the anaphora is very simple. The anaphora obtains because *a stockbroker* and *he* are just variables carrying the same subscript. Hence the anaphoric link in (7) is no more problematic than the one in (19) below (where the first *she* is used deictically, and the second anaphorically):

[15] For simplicity, I ignore the part of the meaning of *she* that concerns gender.

(19) She is a very rich woman. In fact, she is the richest stockbroker on Wall Street.

In both cases, what we have is a sequence of variables that share the same subscript. Anaphora happens via this simple mechanism.

For concreteness, let me also show how the variables view is implemented in a specific framework. I choose the framework of File Change Semantics, developed by Heim (1982). (On File Change Semantics, see also Heim (1983). Heim's framework is notably similar to the Discourse Representation Theory framework developed by Kamp (1981); see Geurts & Beaver (2011) for an overview of implementations, applications, and developments of this family of semantic theories.) In Heim's framework, descriptions just are variables. Their semantic value is an individual, namely (as it happens for all variables) the individual assigned to the variable by the assignment;[16] the content provided by the predicate is analyzed as a presupposition on the denotation of the variable.[17] For example, *a stockbroker* denotes its referent, and presupposes that that referent is a stockbroker; in symbols:

$$[\![x_2(\text{stockbroker})]\!]^{w,g} = g(x_2)\colon g(x_2) \text{ is a stockbroker in } w$$

The difference between definites and indefinites, Heim argues, is that the former introduce a referent that is already familiar in the discourse context, while the latter introduce a new, unfamiliar referent. Cashing out the notions of familiarity and novelty in a precise way is nontrivial and requires rethinking our formal models of communication;[18] but this part of Heim's theory is not central to current purposes and can be set aside.

Let me emphasize again that this is just an illustration and different treatments are possible. In the limiting case, the variable view might also be combined with a kind of quantificational account, giving rise to a hybrid between the Russellian and the pronominal view.[19] But, whatever way one goes, either the syntactic or the semantic core of the Russellian view, or both, will have to be amended.

[16]Notice that his won't automatically give rise to singular truth-conditions. For Heim, definites and indefinites are bound a the discourse level by an existential quantifier, hence (at least at one level of analysis) the final truth-conditions associated to a statement involving a description will be still existential.

[17]For current purposes, just think of presuppositions as the information that is 'backgrounded' by an utterance, without being part of what is asserted.

[18]In short, it requires supplementing our model of shared information in communication with a set of *discourse referents*. Following Heim's metaphor, you can think of discourse referents as 'file cards' that speakers use to track which objects are being talked about. An utterance involving a definite works as an instruction to add information to an existing card. An utterance involving an indefinite works as an instruction to open a new file card.

[19]The basic idea in this case would be that definites and indefinites are indeed quantificational, but have a silent restrictor that (a) includes a covert variable ranging over individuals and (b) restricts the domain of quantification to a singleton set. For a well-developed version of this view, see Schwarzschild (2002).

4.2 Extending the variable view

My whole discussion has been centered about descriptions, both definite and indefinite. But the basic puzzle that led me to argue for the fourth reading can be replicated for virtually all quantified phrases in natural language. Consider for example:

(20) Jason believes that **all/few/several/many/**. . . stockbrokers jumped off a building. But they were not actually stockbrokers, but stuntmen.

These data suggest that, if the variable view has to work as a credible solution to the puzzle, it should be generalized to all determiner phrases. The defense and the development of this idea go beyond the purposes of this paper. But here I can flag that this is not a far-fetched possibility. On the contrary, there are independent reasons to think that this is the right direction for the variable view. Recent versions of Discourse Representation Theory (for example, the one developed in Nouwen (2003)) propose exactly this generalization, treating all quantified phrases on the basic model proposed by Heim for descriptions. Eventually, the arguments presented in this paper, appropriately extended to cover the extra data, might be turned into an additional argument for the new breed of Discourse Representation Theory views.

4.3 Two dimensions to attitude reports?

Before moving on to discussing Gendler Szabó's view, let me draw a suggestive connection between the issues discussed in this paper and some general claims about attitude reports. I've said that attitude verbs are normally taken to manipulate a world parameter. This is what generates the distinction between opaque readings of descriptions (obtained when the world of evaluation is the one introduced by the verb) and a transparent one (obtained when the world of evaluation is the actual one). Recently, several philosophers have suggested that attitude verbs manipulate, in addition to the world parameter, an assignment parameter.[20] In other words, in the same way as attitude verbs introduce a new value for the world parameter, they also introduce a new value for the assignment parameter. Schematically:

> ⌜S believes that ϕ⌝ is true iff for all worlds w' and for all assignments g' compatible with what S believes, ϕ is true at w' and g'

For reasons of space, I skirt over the motivations and the exact functioning of this innovation. All I want to do is highlight that, by combining the descriptions-as-variables insight with this enriched view of attitude verbs, we might be able to capture in an elegant way the functioning of the four readings. As I mentioned, whether a description

[20] See Cumming (2008), Ninan (2012), Santorio (2012).

receives an opaque or a transparent reading depends on which world parameter we use to evaluate it. Similarly, it might be that specific and unspecific readings depend on what assignment parameter we use. We might get specific readings by using the assignment of the context; and we might get unspecific readings by using the assignment introduced by the attitude verb. This would yield an exhaustive and elegant taxonomy of the four readings:

	OPAQUE	TRANSPARENT
UNSPECIFIC	shifted assignment; shifted world	shifted assignment; actual world
SPECIFIC	actual assignment; shifted world	actual assignment; actual world

This, of course, is just a suggestive possibility. An argument for this conclusion would take me too far off my current aims and must wait for another occasion.[21]

5 The Gendler Szabó account: split quantifiers

As I anticipated, Gendler Szabó (2010) wants to retain the quantificational picture of determiner phrases (and descriptions in particular), but appeals to a nonstandard syntax. On the resulting view, the 'split quantifier view', all determiners in natural language (*a* and *the*, but also *all*, *some*, *most*, etc.) can take wide scope with respect to attitude verbs, while leaving their overt restrictor behind. Thus the problematic report (5) (repeated here) has the schematic form below:

(5) Jason believes that a stockbroker jumped off a building.

$\exists_x$ Jason believes [stockbroker(x) jumped-off-a-building(x)]

When a determiner takes wide scope in this way, it is supplemented with a nonspecific silent restrictor, which means roughly *thing*. As a result, the truth-conditions assigned to (5) are approximately the following:

> There is something that Jason believes to be a stockbroker and to have jumped off a building.

Thus Gendler Szabó agrees with me that there is a fourth reading of descriptive phrases in attitude reports. But he disagrees that this forces us to revise the semantics of de-

[21] Also, this view is incompatible with the general account in Santorio (2012), according to which attitude verbs obligatorily bind singular terms in their scope. I still find arguments for that account compelling. Hence I'm genuinely uncertain that the suggestion made in this section is on the right track, at least in the present form.

scriptive phrases. Rather, we should amend their syntax. Notice that, once this move is made, all problems generated by instances of anaphora like (7) disappear, since standard versions of the D-type theory are able to handle them.

Gendler Szabó's view is developed as an account of a wide array of data, which includes instances of anaphora like (7), but also cases of different kind. Here I lack the space for an appropriate discussion of all his data, so I just won't engage with them. What I will do, rather, is mention two reasons for preferring the variable approach to the split quantifiers view.

Proportional determiners. The first and main concern is that the split quantifier view undoes one of the central innovations of our modern theory of quantification in natural language, Generalized Quantifier Theory (henceforth, GQT; see Barwise & Cooper (1981)). GQT improved crucially on existing treatments of quantifiers by providing an account of a whole class of recalcitrant expressions, i.e. so-called proportional determiners. These are determiners that compare the cardinality of two sets, like *most*, *less than half*, or *one third*. The split quantifier view undoes this innovation and reintroduces the problem.

To appreciate the problem, consider the following scenario:

> Samantha believes that all stockbrokers are rich. She also believes that stockbrokers are a rather tiny minority of existing things: there are way more things in the world who are not stockbrokers than stockbrokers. Moreover, she also believes that there are way more non-rich things than rich ones.

(21) Samantha believes that less than half of the stockbrokers are rich.

The split quantifier view predicts that (21) has a true reading in the scenario described. How is this possible? The problem is that, by allowing the determiner to take wide scope with respect to the attitude verb while leaving the restrictor behind, the split quantifier view predicts that (21) gets a reading with the schematic form below, which is truth-conditionally equivalent to (22)

Less-than-half$_x$ Samantha believes [stockbroker(x) rich(x)]

(22) Less than half of the things are such that Samantha believes of them that they are stockbrokers and that they are rich.

But this is obviously not something that (21) can mean in natural language.

The problem has deep roots and connects to general issues in the modeling of quantifiers in language. Start from quantifiers in first order logic. On one natural way of

thinking of them, descending from Frege, quantifiers predicate properties of sets. For example, the formula

$$\exists x\, F(x)$$

predicates of the set of *F*s the property of being non-empty. Initial treatments of natural language determiners were built just on this model. By contrast, GQT treats determiners as stating a relationship between two sets. For example, in

A stockbroker fled with my money.

the determiner takes as argument two sets, the set of stockbrokers and the set of things that fled with my money, and says that they have a non-empty intersection. While this difference is not crucial for the treatment of a determiner like *a*, it is central for proportional determiners like *less than half*. The comparative aspect is essential to these determiners: their job is just to compare the cardinality of two sets. If we divorce syntactically the determiner from its restrictor, like the split quantifier view does, we end up getting dramatically different truth-conditions.[22]

Thus, by allowing the determiner to be separated from its restrictor, the split quantifier view reintroduces a major problem that our standard theory of determiners was designed to solve. This seems a very high cost and I would take it to be sufficient to prefer an alternative account. But, to bolster my case, let me give an extra argument.

Indefinites in conditionals. Consider the following scenario:[23]

> Jason has friend, Maurice, who is a professional stuntman. Maurice is a compulsive liar and has managed to convince all his friends that he does stunts just for fun, as a side career, and that he mainly works as a stockbroker on Wall Street. As a result, Jason believes that Maurice is both a stockbroker and a skilled stuntman. Hence the following is true:
>
> (23) Jason believes that, if a stockbroker jumped off the tallest building on Wall Street, he would survive. He is not a stockbroker, despite what Jason believes.

[22] Moreover, there is no easy patch. We know (via a formal result in Barwise & Cooper (1981)) that there is no way to define the meaning of a quantificational phrase like *More than half of the Fs* in terms of *More than half of the things*, plus first-order logic resources. So the comparative understanding of proportional determiners is necessary to capture the meaning of *most*, *more than half*, and the like; moreover, there is no way to recover it starting from the non-comparative one.

[23] Examples of this sort have been discussed in the debate on indefinites and scope islands: see in particular Reinhart (1997), and Schwarz (2011) for an overview of this debate. The basic observations about the problems one get when widescoping existential operators out of antecedents of conditionals goes back to Heim (1982).

The usual anaphora trick ensures that the indefinite in (23) gets a fourth reading. Hence, according to the split quantifier view, the schematic form of the report in (23) is:

(24) $\exists_x$ [Jason believes that, if stockbroker(x) jumped off the tallest building on Wall Street, x would survive.]

But now, the problem is that the truth-conditions of (24) are way too weak. For (24) to be true, all that is required is that there be *some person* (or thing) such that Jason believes that, if that person (or thing) were a stockbroker and jumped off the tallest building on Wall Street, that person (or thing) would survive. And for those truth-conditions to hold, we don't need the kind of scenario that I gave. For example, any stuntman Jason knows might act as a witness for the existential statement. After all, Jason might think that his stuntman friend Fred, whom he believes not to be a stockbroker, is such that, if he were a stockbroker, would still possess his stuntmanlike agility and hence would not die from jumping off the tallest building on Wall Street.

My discussion of the split quantifiers view has been compressed, but I hope that the concerns I have raised are enough to make the variable view preferable. At least I hope that, as I tried to do for D-type theories, I have shifted the burden of proof. We need to be showed that there is a viable version of the split quantifier view before preferring it to the variable view.

6 Conclusion

There is good evidence that descriptions must be analyzed on the variable model. We can see this by recognizing that descriptions have a singular opaque reading under attitude verbs. In turn, we can see this reading by investigating anaphoric links between descriptions occurring in attitude contexts and pronouns occurring outside those contexts. My argument is based on very specific examples, but it is solid and its conclusions are far-reaching. As it often happens in philosophy of language, specific phenomena can give us general lessons.[24]

[24]For lively discussions and feedback thanks to Raul Saucedo, Jenn Wang, and audiences at the 2012 AAP conference in Wollongong and the 2012 BSPC. Thanks also to Paul Elbourne for a brief but useful email correspondence on D-type theories. Finally, special thanks to Zoltan Gendler Szabó for extended email exchanges on descriptions and the fourth reading.

References

Barwise, Jon, and Robin Cooper (1981). "Generalized Quantifiers and Natural Language." *Linguistics and Philosophy*, 4(2): pp. 159–219.

Chierchia, Gennaro (1992). "Anaphora and Dynamic Binding." *Linguistics and Philosophy*, 15(2): pp. 111–183.

Cresswell, Max J. (1990). *Entities and Indices*. Kluwer Academic Publishers.

Cumming, Sam (2008). "Variabilism." *Philosophical Review*, 117(4): pp. 525–554.

Elbourne, Paul (2005). *Situations and Individuals*. MIT Press.

Fodor, Janet D. (1970). *The Linguistic Description of Opaque Contexts*. Ph.D. thesis, MIT.

Frege, Gottlob (1892/1997). "On *Sinn* and *Bedeutung*." In M. Beaney (ed.) *The Frege Reader*, Blackwell Publishers.

Gendler Szabó, Zoltan (2010). "Specific, Yet Opaque." In M. Aloni, H. Bastiaanse, T. de Jager, P. van Ormondt, and K. Schulz (eds.) *Proceedings of the Amsterdam Colloquium 2009*. Berlin Heidelberg: Springer-Verlag.

Geurts, Bart, and David I Beaver (2011). "Discourse Representation Theory." In E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy (Fall 2011 Edition)*, URL = <http://plato.stanford.edu/archives/fall2011/entries/discourse-representation-theory/>.

Heim, Irene (1982). *The semantics of definite and indefinite noun phrases*. Ph.D. thesis, University of Massachusetts, Amherst.

Heim, Irene (1983). "File Change Semantics and the Familiarity Theory of Definiteness." In R. Bäuerle, C. Schwarze, and A. von Stechow (eds.) *Meaning, Use, and the Interpretation of Language*, de Gruyter.

Heim, Irene (1990). "E-Type Pronouns and Donkey Anaphora." *Linguistics and Philosophy*, 13(2): pp. 137–77.

Heim, Irene, and Angelika Kratzer (1998). *Semantics in Generative Grammar*. Oxford, UK and Malden, MA: Blackwell Publishers.

Hintikka, Jaakko (1962). *Knowledge and Belief*. Cornell University Press.

Hintikka, Jaakko (1969). "On the Logic of Perception." In *Models for Modalities*, Dordrecht-Holland: Reidel.

Kadmon, Nirit (1987). *On Unique and Non-unique Reference and Asymmetric Quantification*. Ph.D. thesis, University of Massachusetts, Amherst.

Kamp, Hans (1981). "A theory of truth and semantic representation." In T. J. J.A.G. Groenendijk, and M. Stokhof (eds.) *Formal Methods in the Study of Language*, Amsterdam: Mathematical Centre Tracts 135, pp. 277–322.

Keshet, Ezra (2008). *Good Intensions: Paving Two Roads to a Theory of the* De re /De dicto *Distinction*. Ph.D. thesis, MIT.

Neale, Stephen (1990). *Descriptions*. MIT Press.

Ninan, Dilip (2012). "Counterfactual attitudes and multi-centered worlds." *Semantics and Pragmatics*, 5(5): pp. 1–57.

Nouwen, Rick (2003). *Plural Pronominal Anaphora in Context : Dynamic Aspects of Quantification*. Ph.D. thesis, University of Utrecht.

Percus, Orin (2000). "Constraints on Some Other Variables in Syntax." *Natural Language Semantics*, 8(3): pp. 173–229.

Quine, Willard V. O. (1956). "Quantifiers and Propositional Attitudes." *Journal of Philosophy*, 53: pp. 177–187.

Reinhart, Tanya (1997). "Quantifier Scope: How labor is Divided Between QR and Choice Functions." *Linguistics and Philosophy*, 20(4): pp. 335–397.

Romoli, Jacopo, and Yasutada Sudo (2009). "De Re/De Dicto Ambiguity and Presupposition Projection." In *Proceedings of Sinn und Bedeutung 13*.

Russell, Bertrand (1905). "On Denoting." *Mind*, 14(56): pp. 479–493.

Santorio, Paolo (2012). "Reference and Monstrosity." *Philosophical Review*, 121(3): pp. 359–406.

Schwarz, Bernhard (2011). "Long distance indefinites and choice functions." *Language and Linguistics Compass*, 5(12): pp. 880–897.

Schwarzschild, Roger (2002). "Singleton Indefinites." *Journal of Semantics*, 19(3): pp. 289–314.