

MEANING SHIFT AND THE PURITY OF ‘I’

Edison Barrios

ABSTRACT

In this paper I defend the “Standard View” of the semantics of ‘I’ – according to which ‘I’ is a pure, automatic indexical – from a challenge posed by “deferred reference” cases, in which occurrences of ‘I’ are (allegedly) not speaker-referential, and thus non-automatic. In reply, I offer an alternative account of the cases in question, which I call the “Description Analysis” (DA). According to DA, seemingly deferred-referential occurrences of the 1st person pronoun are *interpreted as* constituents of a definite description, whose operator scopes over an open sentence Rxy – where R is a contextually selected relation ranging over pairs of people and objects. The role of intentions is thus limited to the determination of R , which is posterior to the fixation of the reference of ‘I’. In support of the DA I present evidence that, in the cases in question, the (Determiner) phrase containing ‘I’ behaves in relevant ways like a description. I show that the DA can account for the problematic examples, while preserving the simplicity of the standard semantics of ‘I’. Finally, I examine a rival account of the data, offered by Nunberg (1993), and argue for the superiority of the DA.

1. – The Standard View and a Challenge

The most widespread view of the semantics of the first person singular pronoun contains, among others, these two theses:

- a. *Speaker Referentiality*: every utterance u of ‘I’ refers to the speaker of u .¹
- b. *Purity*: ‘I’ is a pure indexical.

Pure indexicals, such as ‘I’, are “automatic”. That is, their designation “is *fixed* given its meaning and public contextual facts” and “no further intention, [beyond the intention of using it with its ordinary meaning] is relevant to determining the referent” (Perry 2006, 320, italics mine). In this respect, pure indexicals – and ‘I’, in particular – are the opposite of “discretionary” indexicals (such as demonstratives), whose designation depends partly on the speaker’s intentions. In consequence, the second thesis entails that intentions can have no role in fixing the reference of an occurrence of ‘I’.²

¹ For instance according to Kaplan, “In each of its utterances, ‘I’ refers to the person who utters it”. Likewise, for Perry (2006, 318), “a given utterance u of ‘I’ refers to the speaker of u ”.

² “The linguistic rules which govern their use fully determine the reference for each context. No supplementary actions or intentions are needed” (Kaplan, 1989, p. 491). Also, on Kaplan’s view, ‘I’ stands in sharp contrast to “impure” indexicals, such as demonstratives, whose reference cannot be fixed unless some demonstration or intention is taken into account.

Let us call this view of the semantics of 'I' the "Standard View".

Despite its intuitive appeal the Standard View has had its share of doubters. Among the most recent is Mount (2008), who argues that all indexicals, including 'I', are discretionary (i.e. non-automatic). This claim occurs in a more general discussion, aimed at showing that the distinction between "pure" and "impure" indexicals is spurious, since speaker intentions invariably have a crucial role in the determination of reference, indexical or otherwise. Mount's objections to the Standard View are driven by a series of cases in which the designation of (occurrences of) the first person singular pronoun seems to have shifted away from the speaker, thus challenging presumptions of speaker-referentiality and thus of automaticity. This phenomenon – let us collectively refer to its instances as "Agent Shift" cases – is illustrated by the following utterances:³

(1)

a. [Bill Clinton is at a celebrity wax museum, with Hillary. After a quick search, he finally finds the figure that represents him, which stands between Charlie Chaplin's and Groucho Marx's figures. However, Hillary hasn't spotted it yet. To help her find his figure, Bill says to her:]

I am next to Chaplin.

b. [We are playing Monopoly, and I just made a move. I say to my wife, who has just arrived:]

I am on the purple square (adapted from Mount, 2008, p. 200).

c. [A man bets on a horse at a particular race, and the horse happens to finish last in the race. He says to his friend:]

I finished last in the race (Adapted from Smith, 1989, p.183).

d. [There is a summit of world leaders. To ensure the safety of the participants, each politician has been assigned a body double. One morning, Vladimir Putin saunters into the room and, upon seeing his double, he addresses him thus:]

How am I doing today? (Adapted from Akman, 2001, q.b. Perry, 2006).

e. [Borges, the Argentine writer, is asked to comment on his own work, but he refuses. When the journalists inquire about the reason, Borges says:]

I haven't read myself in a long time.

Prima facie, agent shift examples seem to be instances of metonymy, brought about by what has been called "deferred reference".⁴ Deferred reference is a process whereby an expression gets to refer to an object B, where such an expression does not conventionally refer to B but to an object A that is related to, but distinct

³ A sentence like my example (1b) appears in Mount's paper, the rest do not.

⁴ The phenomenon of deferred reference has been widely reported in the literature. An important source in linguistics is Nunberg (1977), and in philosophy discussion goes back at least as far as Quine's (1968) discussion of "deferred ostension".

from B.⁵ Deferred reference has been mostly studied with respect to demonstrative and personal pronouns, where the demonstratum can be different from the referent, e.g.:

(2)

a. [Restaurant patron to valet, holding up a car key:]

This is parked out back. (Nunberg 1995, ex. 1).

b. [Pointing to a portrait of Saul Kripke:]

That is the author of *Naming and Necessity*.

In (1a) the demonstratum is the car key, while the referent is the car. Likewise, in (2b), Kripke himself is being referred to, though the portrait is the object of the demonstration. Other kinds of expressions can also receive deferred readings, such as personal pronouns, proper names⁶ and definite descriptions.⁷ According to Mount (2008) the first person singular pronoun is also susceptible to this general phenomenon, and this would show that ‘I’ is not invariably speaker-referential and thus that the Standard View is false.

⁵ This definition is put forward with an eye to the discussion of cases such as (1), which would appear to be examples of “referential metonymy” (see Stallard, 2003; Warren, 2006; Bezuidenhout, 2008).

However, deferral phenomena have not always been discussed in connection with metonymy. In philosophy, the phrases “deferred reference”, “deferred ostension”, and similar ones have frequently been deployed in discussions in which figurative or nonstandard uses are not the main concern. These include debates about the role of deferral phenomena in indexical reference and about ostension as a method of definition and reference-fixation, as well as arguments about the nature of reference and their attending philosophical implications, such as the inscrutability of reference and ontological relativism (Quine, 1968, 1990).

In linguistics Nunberg (1979, 1993) took deferral to be an essential component of all indexical/demonstrative reference. A consequence of this is that the import of deferral is not exhausted by its role in figuration. For instance, the interpretation of indexicals, according to Nunberg, involves two stages. The first is the recovery of a semantically indicated element in the context, called the “index” (e.g. the speaker, in ‘we’); the second is the identification of the actual referent or interpretation (e.g. a group appropriately related to the speaker, in the case of ‘we’), which depending on the case may or may not be identical to the contextual element. So, we would have a whole kind of cases in which the argument of the deferral function is not itself a referent, either standard or nonstandard.

The characterization in the text could be modified to accommodate these features, but it, as it stands, has the virtue of being both straightforward and sufficient for the goals of this paper.

⁶[Looking at the Chancellor’s personal art collection. No physical demonstration is involved:]

He must be rich!

[At a wax museum, a visitor asks the clerk at the information booth about the location of the president’s wax statue, and the clerk responds:]

Obama is on the 2nd floor, near the elevator.

In many cases, though, deferral for proper names is blocked. See Nunberg (1993) and Powell (1998).

⁷ a. [Server to co-worker in a diner]

The ham sandwich is at Table 7 (Nunberg 1995, ex. 19).

b. [Physician to head nurse, at a hospital]

The ruptured spleen in 250 needs a nurse.

Nunberg (1995) proposes an alternative analysis, in which the element that shifts meaning is the noun phrase (‘ham sandwich’, ‘ruptured spleen’), rather than the whole determiner phrase.

To answer this challenge, the defender of the Standard View must provide a plausible alternative analysis of agent-shift cases that secures reference to the speaker while at the same time accounting for the apparent reference shift. This is the task I take on in this paper. To this end, I provide an account of agent-shift cases, which I call the “Description Analysis” (DA). On the DA, seemingly deferred-referential occurrences of the 1st person pronoun are indeed speaker-referential, but they are interpreted as if they were constituents of an implicit definite description, whose *denotatum* is an object other than the speaker. The net result is that while the (determiner) phrase – i.e. a DP – in which ‘I’ occurs does not designate the speaker,⁸ pronoun ‘I’ keeps its ordinary reference.⁹

In the next section I will introduce a bit of handy terminology. Subsequently, I will motivate the DA (3.1) and provide arguments for it (3.2). The case for the DA mainly consists in showing that agent shift sentences exhibit a description-like behavior – such as scope taking and attributive readings – which follows straightforwardly from the DA, but cannot be accommodated under a deferred-reference view of agent shift. In (3.3) will offer some clarificatory remarks on the nature of the proposal. Finally, in section 4 I will discuss Nunberg’s “Predicate Transfer” approach, and will argue that the DA ought to be preferred.

2. – P(hrasal) Shift vs. R(eferential) Shift

In this preliminary section I will introduce a couple of distinctions that are necessary to appreciate the differences between rival accounts of agent shift cases, in terms of the possible sources of meaning shift.

Limiting ourselves to (complete) declarative sentences, we can understand ‘the ATM swallowed my card’, ‘the White House is worried’ or ‘the ham sandwich wants his check’ as involving a main predicate or relation,

⁸ I’m using the verb ‘designate’ as a general term for semantic relations through which an expression picks an individual, thus it includes both reference and denotation (in the Russellian sense) and is neutral between them.

⁹ There are other apparent counterexamples to the Standard View. Among these are examples based on answering machine messages and post-it notes that contain utterances (including both spoken and written ones) where ‘I’ does not refer to the utterer or writer. These cases arise when there is a discrepancy between the context in which the message was coded and that in which is decoded, or because the transmitter or deliverer of the message is not the same as its originator. In general, these cases can be accommodated within the Standard View, and there are several proposals about how to do it, which may be mutually equivalent. One way is by distinguishing between speaker/utterer and agent/author, and tying the reference of ‘I’ to the second. Another strategy is we can call “context shift” the context relevant for the interpretation is different from the context U of utterance, that is, from the context of material production of the utterance. (For a variety of solutions along these lines see Bianchi, 2001; Corazza, Fish and Gorvett (2001); Romdenh-Romluc, 2006; Dodd and Sweeney, 2009). In any case the essence of the character of ‘I’ is preserved, since the changes are peripheral, or in any case pre-semantic (see Predelli, 2005).

Other interesting *prima facie* counterexamples are (among others) so-called “descriptive uses” (such as the example of the condemned prisoner: ‘I am traditionally allowed to order whatever I like for my last meal’. See Nunberg, 1993) and impersonal uses of the first person singular pronoun, observed, for instance, in languages such as English and German (Zobel, 2010). These cases merit attention, but I won’t discuss them here, as this paper focuses on “deferred reference” cases.

which can be expressed as an open sentence, such as ‘*x* swallowed *y*’s card’, ‘*x* is worried’, ‘*x* wants *y*’, as well as a set of arguments (*The ATM, The White House, The ham sandwich, x*’s check’), including subjects and objects. Meaning shift can occur in the main predicate or in some of the arguments, or both.¹⁰ If the shift takes place in the predicate, we have a case of Predicate Shift. Suppose, for instance, that I utter the sentence, ‘John is at Harvard now’, while my interlocutor and I both mutually know that at the moment John is no more than 20 ft away from us (and that we are in California). In this case it is plausible to say that ‘*x* is at Harvard’, but not ‘John’, underwent a shift in its meaning, where the relation ‘*x* is at *y*’ is interpreted as something like ‘*x* works at *y*’ or ‘*x* is associated with *y*’.

In contrast, the following sentences seem to contain cases of argument shift, where the italicized term is the most obvious shifter:

(3)

- a. *Washington* fears that the enrichment activities, some of them initially conducted in secret, could service a clandestine bomb program (*Time*, 8/21/2010).
- b. *This* is parked out back [holding up a car key].

The most widely held assumption about Agent Shift cases – which is needed for them to constitute counterexamples to the Standard View – is that they are examples of argument shift. I will go along with this assumption. However, Nunberg (1993, 1995) does offer an alternative analysis in which the meaning transfer observed in agent shift cases is due to predicate shift, not argument shift. I will discuss this view in section 4.

Within the category of argument shift, there are further distinctions to make. A particularly important one has to do with the syntactic *locus* of the shift within the phrase that realizes an argument. Some expressions, which we could call “subphrasal terms” (or simply “terms”), are “atomic” from a syntactic point of view. Let us assume that referential expressions such as names and pronouns are in this category. Thus, a shift due to a change in a term’s reference is an instance of ‘R(eferential)-shift’. Deferred reference is the clearest example of R-shift: for instance, in the example of the car keys (number) the reference of the demonstrative term ‘this’ has shifted away from the demonstratum (the keys) towards a suitable related car.

Phrases, of course, are composed out of terms or other phrases. P(hrasal)-shift occurs when the transference originates at the phrasal level, and cannot be traced back to any of the individual terms composing the phrase. Thus, argument shift is not exhausted by R-shift, since phrases constituted by expressions that are not type-referential, such as quantifier phrases, can also shift meaning. For instance in a sentence such as ‘The ham

¹⁰ In some cases, such as deferred equatives (‘I am the pad thai’), the shift occurs in the copula, according to Ward’s (2004) analysis.

sandwich wants his check', the locus of shift is the whole descriptive phrase, since none of its constituents seems to have undergone a shift in meaning.¹¹

The distinction between P-shift and R-shift serves to differentiate my position from those of other authors who agree that agent shift cases are characterized by Argument Shift, but who classify them as instances of “deferred reference”. In particular, the deferred reference challenge proposed by Mount and others assumes that the locus of shift is the pronoun ‘I’, and thus that the cases at issue are examples of R-shift. In this paper I reject this assumption, and argue instead that the shift is localized at a purely phrasal level (P-shift). More specifically, I will argue that the shift originates, not in ‘I’, but in [DP I], a phrase that contains the occurrence of the pronoun. The justification of this claim depends on the Description Analysis of agent shift cases, which will be introduced in the next section.

Before I finish this section, I would like to make a terminological parenthesis. As mentioned earlier, the examples in (1) are *prima facie* cases of metonymy, so for the rest of the discussion it will be convenient to borrow some nomenclature from the literature on this phenomenon. In particular, I will adopt part of Fauconnier’s (1985) terminology (though not his analyses). So, in the case of a particular phrase P used metonymically: the *trigger* is the conventional interpretation of P, i.e. some object, property, relation or event A; the *target* is the intended interpretation of P, i.e. some object, property, relation or event B distinct from, but related to A. In (1), the trigger is invariably the agent of the context, i.e. the speaker of the utterance, whereas the target and trigger-target relation vary from case to case.

2.0 THE DESCRIPTION ANALYSIS

2.1 Agent-Shift Cases and Possessives

There are striking similarities between the interpretation of agent-shift utterances and the interpretation of utterances containing first-person possessives in argument positions. For every example in (1), a speaker can express the same idea by substituting the first person pronoun (the trigger) with a DP of the form ‘My x’, where x stands for an NP containing a relevant predicate satisfied by the target.¹² E.g.:

¹¹ In this example the shifting phrase contains a noun-noun compound, but this is not the only configuration in which P-shift can occur, as we’ll see later.

According to certain analyses (Nunberg, 1993, 1995; Sag, 1981) the shifter in this case is the NP ‘ham sandwich’, rather than the whole DP. In this case, we have an instance of “common noun shift”, which can be regarded as a shift in a predicate that sits within the argument.

¹² The sentences could be paraphrased in different, equally acceptable ways, but at this stage of the discussion this should not be a worry.

(1a) I am next to Chaplin.

(1a') My statue is next to Chaplin.

(1b) I am on the purple square.

(1b') My piece is on the purple square.

(1c) I finished last in the race.

(1c') My horse finished last in the race.

(1d) How am I doing today?

(1d') How is my double doing today?

(1e) I haven't read myself in a long time.

(1e') I haven't read my own work in a long time.

This systematic correspondence suggests that both classes of phenomena are candidates for receiving uniform treatment. A phrase such as 'My x' can be interpreted as picking something (the "possessee") which stands in a relevant relation R to the possessor – the speaker, in this case.

In these phrases the possessee is marked as *definite*.¹³ This can be clearly seen in this non-agent shift example (taken from Barker, 1995, p. 78 ff):

(4)

a. I saw a child.

b. I saw the child.

c. I saw John's child.

In (a), the complement is an indefinite phrase, and there is no presumption that the child in question would be identifiable or contextually unique. However, in (b), the definite version, the child *must* be contextually unique or identifiable. Sentence (c) patterns with (b) rather than (a) in this respect, so possessives seem to be closely associated with definiteness (compare, for instance 'my child' with 'a child of mine').

In fact, possessives have traditionally been analyzed as definite descriptions,¹⁴ so that a sentence such as (5') is interpreted in the manner of (6), where R is a relation holding between the possessor and the possessee:

¹³ That is, as involving uniqueness relevant to the context (or familiarity/salience/givenness/accessibility, etc.). Feel free to substitute with your favorite view about the nature of definiteness.

(5) [To a parking lot attendant]:

(5') My car is parked out back.

(6) [The x : x is a car and x stands in R towards me] x is parked out back.

This makes it explicit that there is a car that uniquely stands in a relation R towards the speaker.

Let us reconsider, then, the sentences in (1). Take (1a'), repeated here for convenience:

(1a') My statue is next to Chaplin.

According to the definite description view of possessives (1a') is equivalent to:

(1a'') The statue that stands in R to me is next to Chaplin.

Now, since (1a) – in the intended interpretation – and (1a') are interchangeable, and since (1a') has the same truth conditions as (1a''), then an utterance of (1a) – again, with the intended reading – will be interpreted as an equivalent of (1a''). Suppose that the value of R in this case is something like ‘ x represents y ’ or ‘ y is represented by x ’, then the interpretation of (1a) would be expressed by (7) or the more English-like (8):

(7) [The x such that x is a statue and I am represented by x] x is next to Chaplin.

(8) The statue that represents me is next to Chaplin.

Similarly, as pointed out earlier, (1b) and (1b') are interchangeable in the context at issue:

(1b) I am on the purple square.

(1b') My piece is on the purple square.

And since (1b') can be interpreted as a description, as in (1b''), the same goes for (1b):

¹⁴ See, for instance, Russell's (1905, p. 484) analysis of the denotation of “my only son”, as well as Donnellan's (1966) famous discussion of cases like “Smith's murderer” and “her husband”.

(1b") The piece that stands in relation R to me is on the purple square.

Now, suppose that in this context R is something like 'x is used by y' or 'y is playing with x'; then the interpretation of (1b) might be something like:

(9) The piece I am using is on the purple square.

The same kind of analysis can be extended to the other sentences in (1). So, the following hypothesis seems attractive: in (at least) some cases of apparent deferred-referential use of first-person pronouns these pronouns can instead be thought of as *speaker-referential indexicals that are embedded in definite descriptions*.¹⁵ An attractive feature of this analysis is that it is eminently conservative, theoretically speaking: 'I' is *interpreted* as a part of an implicit larger phrase, and what 'I' provides to the composition of the meaning of the sentence is, as the Standard View would have it, the sentence's speaker or agent.

A qualification: one aspect that may make the sentences in (1) different from their (1') and (1'') renditions is that the latter two are considerably more precise than the former. The reason is that both the possessive and definite-description paraphrases of (1) explicitly specify a sortal predicate under which the targets fall (my PIECE, the PIECE I'm using), whereas the original first-person pronoun sentences leave that aspect open. In most of the cases we have considered this is, for practical purposes, not an issue, since it is quite clear (for instance) that in (1a) the speaker is talking about statues, and that (1b) is about monopoly pieces. However, for examples such as (1e) – i.e. 'I haven't read myself in a long time' – the audience must pick up more interpretative slack than in other cases: is he talking about his recent writings? Short stories? Poems? Perhaps even the speaker's intuition wouldn't settle the issue.

I hypothesized that in (1) the DP's containing 'I' designated an individual, in a way that is analogous to a *referentially used* definite description, expressing a singular proposition. As long as the intended individual (the target) is introduced into discourse, it does not matter under which mode of presentation it is introduced (for semantic purposes, that is; the mode of presentation certainly has significant pragmatic and cognitive repercussions). Thus, it is not mandatory for the interpretative process that the audience explicitly represent, or be able to articulate an identifying sortal, but it is crucial that the interpreter be able to recognize the

¹⁵ A similar idea can be found (albeit in inchoate form) in Smith (1989). He dubs the relation between the speaker and the actual referent (for instance, the person who bets on a horse and the horse he bets on) as one of "adoption" of the first by the second. Thus, he says "[...] something x is "adopted" by the tokener [i.e. speaker] [...] if and only if the tokener can truly say of x "x is mine" and can refer to x in sentences beginning with "my x"". (p. 186). However, he works under the assumption that 'I' has actually shifted reference in these cases.

relation that will allow him to identify the target (for instance, the extension of the domain of relevant objects may be obvious to the participants in the conversation, even if they are not able to explicitly characterize the domain or the relation). So, bypassing sortals, I propose to characterize targets merely as *relata* of a contextually identified relation. More technically, the restriction in the quantifier phrase would only contain a dyadic open sentence, instead of the conjunction of such a sentence and a monadic one. So, in place of (1b"), we could get the simpler (1b''') as a representation of (1b):

(1b") [The x: x is a piece and x stands in R to me] x is on the purple square.

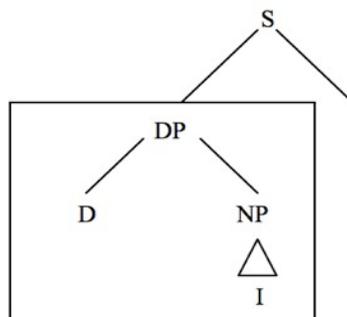
(1b''') [The x: x stands in R to me] x is on the purple square.

Here the audience would only have to identify a contextual value for R, such as ‘...plays with ...’ or ‘... uses ... in the game’, instead of having to identify both an appropriate predicate *and* a relation.

Thus, my proposal, which I call the Description Analysis (DA) is the following: seemingly deferred-referential occurrences of the 1st person pronoun are *interpreted as* constituents of a definite description, whose operator scopes over an open sentence Rxy – where R is a contextually selected relation ranging over pairs of people and objects.

Just for clarity’s sake, this is (broadly speaking) what I take to be the position of ‘I’ in the syntax of an agent shift sentence, according to DA: the whole DP is the subject of the sentence, and ‘I’ is just one of the constituents of the subject DP.¹⁶

Fig. 1



¹⁶ I am also assuming what has been called the “DP hypothesis” (Abney, 1987), according to which determiners (such as ‘every’, ‘the’, etc.) head their own phrasal projections. In consequence, NP’s are complements of – possibly null – determiner heads and so they occur “inside” DP’s, so to speak. I don’t think that this choice affects any of the results in this paper.

2.2 Arguments for the Description View

There is further evidence in favor of the DA. The following arguments depend on this assumption: that if speaker-shift sentences are indeed interpreted as definite descriptions, then they should *behave like* descriptions. Notice, too that the claim is that they are similar in their interpretive possibilities, not that they have a similar underlying syntax.

In the following section I aim to show that sentences like those in (1) can receive both attributive and referential readings (in the sense of Donnellan, 1966), and are also susceptible to scope effects, just as we would expect if they involved implicit descriptions.

(i) *Attributive/Referential*

In the context assigned to (1a), Clinton was talking about a particular statue in his perceptual environment, with which he was causally related through visual contact, and we could think of the subject DP, i.e. [DP I], as being used referentially, that is, as contributing an individual – a lump of wax fashioned Clintonwise – to the proposition being expressed or conveyed.¹⁷

Now, consider this alternative context: Clinton has been told that the National Wax Museum of Baratavia has a gallery in which there is a wax effigy of each and every one of the presidents of the USA, and furthermore, that they are exhibited in one long row, where positions are determined by order of presidential succession. Since he is not an amnesiac, he is aware that he has been a US president, and that the men that preceded and succeeded him in office were George Bush Sr. and George W. Bush, respectively.

After a moment of reflection, he utters:

(10) I am between Bush Sr. and Bush Jr.

In this context, we could interpret (10) as expressing a *general* proposition, such as this: “there is an x such that x is the unique statue of Bill Clinton (in the Wax Museum of Baratavia), and any such x stands between the statues that represent George Bush Sr. and George Bush Jr., respectively”. This corresponds to an attributive reading of [DP ‘I’] in (10).

Continuing with examples involving politicians, we can also get the same attributive/referential duality in (11):

(11) How am I doing today? [=1d]

¹⁷ Or least its contribution is dependent on this object.

A scenario that easily comes to mind is one in which Putin is shaking his double's hand while uttering (11). However, we can also elicit an attributive interpretation. Imagine a context in which Putin has been told that, for his protection, there is going to be a Putin look-alike (not yet chosen) at a given summit he is planning to attend. In this case, Putin is innocent of all contact with or knowledge about the identity of his double. Now suppose that he includes (11) in a note addressed to whomever turns out to have the honor of performing the job. For instance, he might start the letter tongue-in-cheek: "How am I doing today?" and then continue in a more serious tone. In this second context, we can interpret [DP I] attributively.

We could also get an attributive reading of Nunberg's (1993) famous example, 'I'm parked out back', in the following context. In the morning, I leave my car at the repair shop. Later, in the evening, while I'm dining out (I took a cab to "Old Boss McGinty's") I get a call from the shop, telling me that my car won't be ready until the following week. However, they assure me that they will immediately send over a replacement car (about which I get a vague description and a plate number) so that I am able to drive home. They also promise that the car will be left in the parking lot of Old Boss McGinty's. When I leave the restaurant, I notice that the promised car is not in any of the parking spots near the entrance, but knowing that the repair shop is a very trustworthy establishment I infer that the car is parked in a place that is farther away from where I'm standing. I promised my dining companion a ride home, and I reassure her by saying:

(12) [To a parking lot attendant]:

I'm parked out back.

Again, this context seems to license an attributive reading.

Eliciting attributive interpretations for speakers-shift sentences requires a somewhat more elaborate set up than eliciting referential ones. This may be due to the higher salience of referential readings in the cases at issue, because speaker-shift utterances are more commonly used in situations where the target is perceptually accessible or otherwise identifiable in context. Nonetheless, attributive readings *are* available.

(ii) Scope Effects

If speaker-shift utterances are interpreted as definite descriptions, and thus as containing an operator, we expect them to show scope effects when interacting with scope-taking operators, such as modals. Take (13):

(13) I am necessarily going to be next to Chaplin.

This could be interpreted in at least two ways, made salient by contexts 1 and 2, respectively.¹⁸

Context 1:

Earlier on Bill Clinton had been at the wax sculptor's workshop, where he saw a lump of wax being modeled into a figure bearing his likeness. He noticed that the statue was very bulky. Now Clinton is at the museum, surveying the space available for the exhibition. The gallery is almost full with the figures of other celebrities, and, given the size of his statue, he concludes that the only fitting space is next to Chaplin's figure. Worried, Clinton utters (13) to Hillary.

In this case, Clinton would be saying *of* the statue in question that it will necessarily stand next to Chaplin's. That is, he would be expressing a singular proposition, about a particular object. This is a *de re* reading. This reading is available if the interpretation of the utterance is taken to involve a definite description operator taking scope over 'necessarily'. Nevertheless, this is so far consistent with an account in which 'I' – the pronoun – actually *refers* to the statue in (13). But the next context and its intended interpretation will make that position implausible.

Context 2:

Clinton is in his office. In a newspaper he reads of a future exhibition of wax figures of celebrities at Museum X (the exhibit won't be open for some time, though, because the figures have not been made yet). He also reads that the figures will be placed according to the alphabetic order of the celebrities' names, and that all and only those whose names start with the same letter will be in the same row. Later on he comes across a list of the celebrities to be portrayed, and realizes that there are only two people with names starting with 'C': Bill Clinton and Charlie Chaplin. After a moment of reasoning, he utters: 'I am necessarily going to be next to Chaplin' [=13].

In this context, the most appropriate interpretation is one in which the modal operator takes precedence with respect to the descriptive one, and we should take the speaker to be ascribing necessity to the general proposition that the statue representing him (relative to the relevant domain restriction) is going to stand next to the one representing Chaplin. Thus, we could interpret him as saying that, *necessarily, there is (or will be) an x such that x is a statue representing Clinton at place y, and x (will) be standing next to [Chaplin's statue at y]*.

Moreover, the same sentence could be true or false in the same context, and this can be attributed to the availability of different scope relations. To see this let us go back to Context 1. If Clinton's size estimations

¹⁸ I will ignore 'going to', a temporal operator here, because what matters in the example is the relative scopes of the modal and descriptive operators.

are accurate, then a reading of (13) in which the modal takes narrow scope would intuitively be true, as was indicated above. Now let us also assume – as part of Context 1 – that sheer size constitutes the only reason for Clinton’s statue being next to Chaplin’s in this particular exhibit, and also suppose (plausibly enough) that there is no size requirement for qualifying as a Clinton statue. In this case, an interpretation that assigns wide scope to the modal would be false, because of the possibility of relevant situations in which the figures in question are not adjacent.¹⁹

The number of available interpretations for the sentences at issue increases with the number of operators involved, and this provides an additional reason for thinking that the availability of different readings is due to permutations among operators. Suppose, for instance, that Clinton periodically receives news (sometimes true, sometimes false) of exhibitions of his effigy in galleries all over the world. When he reads this kind of news the thought of being next to Chaplin invariably springs to his mind. He describes the situation thus:

(14) I always believe I’m next to Chaplin [adding something like: “but sometimes I’m disappointed”].

These interpretations of (14) are admissible:²⁰

(14a) At every occasion t , Clinton believes at t that the statue x that represents Clinton at t is next to Chaplin (at t). (*De dicto, inside the scope of the temporal operator*)

(14b) At every occasion t the statue that represents Clinton at t is believed at t by Clinton to stand next to Chaplin (at t). (*De re, within scope of the temporal operator*)

(14c) The statue x that represents Clinton at t is such that at every occasion t Clinton believes of x at t that x is next to Chaplin (at t). (*De re, outside the scope of the temporal operator*)

¹⁹ Something similar could be said with respect to Context 2. If the information in the newspaper is accurate, then the interpretation that assigns wide scope to the modal seems true. But suppose that the article is slightly inaccurate, in that the statue corresponding to Clinton (let us call it “Bob”), is already finished and currently in storage at the museum’s basement. Now suppose further – as seems reasonable – that being Clinton-shaped or Clinton-representing is a contingent property of Bob. In this case, had things gone differently, Bob could have been fashioned into an effigy of Nixon, and, likewise, the actual statue of the latter have been molded into a Clinton figure. If this is so, then an interpretation of the utterance of (13) that assigns narrow scope to the modal would turn out false, while the wide scope reading would still be true (matters here could be complicated by delicate issues touching on the individuation and modal properties of statues, but I won’t address them here).

²⁰ For the examples that follow let us suppose that the domain of occasions is suitably restricted, for example, to *occasions in which Clinton reads that a wax figure representing him is being exhibited [at place x]*. Let us also assume that only one Clinton effigy is exhibited in the world at a particular occasion; in this way we can get around the need for a variable for places. Also, to get the desired readings, the temporal operator must invariably take wide scope relative to the propositional attitude verb.

We can obtain many other examples of this kind.²¹ This suggests that in the cases at hand the phrase containing the first person pronoun behaves like a definite description.

Both the reference/attributive alternation and the scope effects are readily explainable on the assumption that in speaker-shift utterances occurrences of [DP I] are interpreted as implicit definite descriptions. However, the phenomena cannot be accommodated within a view according to which agent shift sentences are cases of R-shift (i.e. deferred reference).

2.3 *The nature of the proposal*

According to the DA, a phrase of the type [DP I] occurring in a speaker-shift utterance can be interpreted as an implicit definite description. Thus, occurrences of the first person pronoun in so-called deferred-reference uses can be seen as parts of definite descriptions.

In the cases at issue, then, [DP I] receives an interpretation equivalent to (15):

(15) [ιx: Rxs]

The iota symbol is the definite description operator, followed by a bound occurrence of the variable x. The value of R must be determined contextually. The second occurrence of x is a placeholder for the first argument of R, whereas s is a special variable whose range is restricted in a very specific way. This variable, which constitutes the second argument of R, can only take as value the speaker of the utterance in question. The identity of the speaker, together with the physical context, conversational interests, salience, etc. helps the audience identify the value of R. Thus, R is a contextually salient and relevant relation between the speaker and a certain object or set of objects. Once the speaker is identified, the phrase that is formed by applying the definite operator to the open sentence Rxs is a phrase (term) denoting the unique entity that satisfies Rxs (if there is one).²²

What 'I' contributes to the composition of meaning is the agent, as propounded by the Standard View. This is so because, if 'I' is *interpreted* as embedded in a denoting phrase, then the semantic value of 'I' becomes itself the argument of the function that would be expressed by that phrase. It is the value of this function, and not the value of 'I', that provides an argument to the main predicate of the sentence. How can we account for the *perceived* shift in reference? Agent Shift cases are brought about by a shift at the phrase level (P-Shift), not at

²¹ Consider an utterance of 'I believe someone painted a moustache on me', said by Clinton, talking about his statue.

²² Again, you may modify the formulation to suit the account of definiteness that you think is best.

the level of the pronoun's reference. Although 'I' and [DP I] are syntactically and semantically distinct they are phonetically indistinguishable, so it's easy to misattribute to the first the properties of the second.

In the terminology introduced earlier, the value of *s* is the TRIGGER of the metonymic process, whereas the value of *R* is the relevant RELATION between the trigger and the target, and the unique satisfier of *R* with respect to the speaker in question will be the TARGET.

I would like to make clear that I'm not advancing a syntactic claim. I do not mean that the implicit description is actually articulated or represented at some linguistic level of representation, such as LF. Thus, I am *not* asserting that in speaker-shift utterances the DPs including 'I' contain silent though syntactically articulated operators and variables. As far as my commitments are concerned, the phenomena in question may be the outcome of a pragmatically controlled process of interpretation. *A fortiori*, I do not interpret this phenomenon as a case of (syntactic) ellipsis, but rather, of what could be very broadly described as "utterance ellipsis" (as the expression is used by Neale, 2002, commenting on proposals by Quine and Sellars). In this sense, speaker-shift utterances involve an elliptical *use* of phrases containing the first person pronoun. The speaker leaves out certain elements from the sentence, confident that the addressee will manage to fill in the gaps in the abbreviated version.

At this point somebody may wonder whether these data cannot simply be analyzed as implicatures, or more broadly, whether they can be explained by appeal to familiar Gricean considerations.

I doubt that a Gricean story alone would be satisfactory. Even if the agent shift interpretations are implicatures (and I am not sure that they are) this hardly eliminates the need for further analysis.²³ This is because, although the Gricean story might identify the general mechanisms through which interpretations are arrived at – such as the maxims and the Cooperative Principle – it doesn't address the interesting issue of

²³ They are certainly not conventional implicatures, since they appear to be cancelable. However, although cancelability is arguably an indication that pragmatic inferencing is involved, the phenomenon in question does not have to be an implicature (and even in conjunction with the other customary tests the verdict is not clear). As a matter of fact, in the cases in question *both* the agent ("literal, non-implicated") and the shifted (non-literal, implicated") interpretations are cancelable (for instance, these two kinds of continuations are admissible in wax-statue cases, provided the appropriate contexts: "... but I'm not talking about the statue, I'm talking about myself", "Oh, I'm not talking about myself, I'm talking about my statue!") This is a case of "symmetrical cancelability", as we may call it. In contrast, in a typical implicature cancelability is asymmetrical: only the implicated, "non-literal" interpretation is cancelable (e.g. in Grice's, 1989, famous "There is a garage nearby" example you can cancel the implication that the garage is open, but you cannot consistently backtrack on the existence or location of the garage).

Agent-shift cases may also exhibit variability in terms of the processing behind them: in some cases, a process of inferencing leading to the intended meaning will have to be performed online by the hearer, while in others there may be established "processing shortcuts", where the agent-shift reading will be so standardized that arriving at the intended meaning will be (or will be experienced as) automatic, as is the case with short-circuited implicatures and frozen metaphors, for instance.

which kinds of interpretations are made available by the utterances in question. Also, because a Gricean analysis by itself (in the cases in question, at least) doesn't determine any internal structure, it can predict neither scope interactions nor the availability of attributive uses in agent-shift cases. The DA analysis, in contrast, throws some light on these phenomena.

2.4 *Non-referential Impurity*

A couple of comments are needed before we end this section. I must note that there are, within the space of possibilities, other "impurity" accounts of the phenomena that differ from DA and seem compatible with the evidence. These positions, like the R-Shift one, must place the shift at the subphrasal level – more specifically, in the interpretation of the term 'I' – and must also attribute the shift to the semantic properties of the first person singular pronoun. However, to be taken into consideration, the views must also allow argument phrases to exhibit general interpretations (unlike R-Shift).

At least two possible positions seem to have the requisite characteristics. One is a "hybrid" or "disjunctive" view of the semantics of 'I', according to which it is intrinsically neither referential nor non-referential, but can behave in either way, depending on the circumstances (though it could be said that the referential mode is the default). According to this view, the character or meaning specification of 'I' is such that in non-shift cases (or in typical circumstances) it behaves in accordance to the Standard Theory, whereas in agent-shift cases it behaves semantically like a description.

The second view consists in an entirely non-referential semantics for 'I', according to which it is always descriptive in content. For instance, it would invariably involve a description of the sort 'the unique x that stands in relation R to the speaker', where R would have identity as its default value, but it could also take others (such as 'statue of', 'deputy of') depending on the circumstances.

Both of these views can handle the evidence better than the R-Shift approach. However, they have little to recommend over the combination of the Standard View and the DA that is being considered here.

The hybrid view would be extremely *ad hoc*, and it would postulate a complication of the semantics of 'I', with respect to the Standard View, without any compensating increase in descriptive or explanatory power.

The second view is intriguing, and it has the virtue of providing a unified account of the semantics of 'I'. However, to be considered, it would have to draw support from a wider range of phenomena. Moreover, it remains to be seen whether a suitable version of this view can handle the standard objections against descriptive accounts of the contents of indexicals (such as those in Kaplan's (1989, p. 498 ff., for example).²⁴

²⁴ But see Schlenker (2003) on this issue.

3.0 A Rival Account

3.1 Predicate Shift

Nunberg (1993, 1995) has proposed an influential analysis of agent shift data. This analysis is also compatible with the Standard View, but is quite different from the DA account. In particular, Nunberg’s view denies that agent shift cases involve any argument shift. He uses (3) - repeated here for convenience – as his main illustration:

(16) I’m parked out back [=12].

There are two ways of resolving the apparent categorical clash, or “sortal crossing” between subject and predicate, that occurs when both are given their standard interpretations. The first one is the familiar “deferred reference” reading, in which ‘I’ shifts reference from the speaker to her car, an entity that can appropriately be said to be parked. Here the predicate ‘park out back’ preserves its customary meaning. The other way keeps the agent-reference of ‘I’ unmodified and instead locates the shift in the predicate ‘x is *parked out back*’. Nunberg (*op. cit*) recommends the second approach, which is based on a process called “predicate transfer” Nunberg (1993, 1995). According to Nunberg, predicate transfer is “an operation that takes names of properties into new names that denote properties to which they functionally correspond” (1995, 109).²⁵ In his interpretation of (3), the predicate ‘x is *parked out back*₁’, which takes vehicles as external arguments (subjects) becomes the predicate ‘x is *parked out back*₂’ which instead takes people,²⁶ through a coercion mechanism guided by the salient relation between cars and their owners.

Thus, his proposal is that in (3) “the predicate *parked out back* contributes a property of persons, the property they possess in virtue of the locations of their cars” (1995, 111).

According to Nunberg, predicate transfer is only possible when the property contributed by the new predicate is *noteworthy*. A property is noteworthy if it offers a useful way of classifying its bearer relative to conversational interests (Nunberg, 1995, 114).

²⁵ This also resembles notions such as those introduced by Sag (1981) – “sense transfer” – and by Stallard (1993) – ‘predicative metonymy’. Nunberg’s notion of predicate transfer (as well as Sag’s proposal) includes the transfer in expressions that occupy a predicate position in sentence (predicate transfer proper) or those that attach to common nouns in other positions (for instance, they could be NP’s that complement a DP, for instance, a quantificational one). This last phenomenon, which Nunberg calls “common noun transfer” forms the base of his account of ‘ham sandwich’ cases.

²⁶ The subindex notation is mine, not Nunberg’s.

In the case of (3), the property of being parked is noteworthy because it allows us to classify the members of one set of things (car owners), in virtue of their relation to the properties of the members of a related group of things (their cars).

In the terminology introduced in section 2, Nunberg analyzes the examples in (1) as cases of predicate shift, not argument shift, thus providing another route for 'I' to retain its speaker-referentiality.

3.2 Nunberg's Tests

As noted above, deferred reference (R-Shift) is a type of Argument Shift. Moreover, Nunberg's arguments do not discriminate between the two kinds of Argument Shift we distinguished earlier – R-shift and P-shift – and so it is harmless to treat Argument Shift analyses in general as the target of Nunberg's objections, as I will do in the rest of this essay. This is because, if Nunberg shows that the sentences in (1) are cases of Predicate Transfer then he will have thereby shown that they are not cases of Argument shift. Moreover, to make terminology more uniform, I will use the term 'shift' where he uses 'transfer'. Nothing substantive hinges on this choice of labels.

Nunberg compares (16) with (17) which he takes to be a legitimate example of deferred reference.²⁷

(17) This is parked out back [=2a].

Let us give the name "Standard Argument" to the object that in (16) or (17) would *ordinarily* be designated by the corresponding *subject* phrase²⁸ (i.e. the speaker in (16), and the demonstratum in (17)). Also, let us give the name "Situation-Appropriate Argument" to the object that, in the *context*, is the best candidate to satisfy the *predicate* in its standard or conventional interpretation.

In "normal" instances of true subject-predicate sentences, as well as in cases of predicate shift, the standard argument is identical to the situation-appropriate one, but in cases where there is argument shift they can diverge. Let us look at (16) and (17) again:

(16) I am parked out back

(17) *This* is parked out back

²⁷ With respect to this sentence, Nunberg (1995, 111) says that it is "a case of deferred ostension or deferred indexical reference, a process that allows a demonstrative or indexical to refer to an object that corresponds in a certain way to the contextual element picked out by a demonstration or by the semantic character of the expression."

²⁸ This formulation doesn't quite apply to (1e), because the relevant pronoun there occurs in object position, but that difference is irrelevant for the purposes of presenting Nunberg's views.

Given the presence of the predicate *park*, the situation-appropriate argument must in both cases be a vehicle. The examples differ, however, with respect to their standard arguments: in (17) it is the object demonstrated (a key, in this case), whereas in (16) it is the speaker, therefore a person.

Nunberg's case for the predicate shift analysis is based on tests for pinpointing the location of meaning shift in a particular sentence. The guiding assumption is that the designation of a subject phrase will place restrictions on subsequent reference and predication. The presence of these restrictions is evinced by a series of related phenomena, such as the range of predicates that can be felicitously coordinated with the initial predicate, or appropriately inserted after the subject (as part of an appositive clause).

The tests work in the following way: in cases of argument shift, only predicates applicable to the situation-appropriate argument can be coordinated with the original predicate, since it is the argument that undergoes coercion. For instance, (16) only admits predicates applying to the car (situation-appropriate), but not to the key (standard), e.g.:

(18) This is parked out back and *may not start*. (Nunberg, 1995, Ex. 4)

(19)??This *fits only the left front door* and is parked out back. (Nunberg, 1995, Ex. 5. Judgment in the original).

An analogous reasoning applies in the case of apposition:

(20) This, *which is a blue Chevrolet*, is parked out back (Nunberg, 1993, Ex. 89)

(21) *This, *which I had just duplicated at the locksmith*, is parked out back (Nunberg, 1993, Ex. 90. Judgment in the original).

In predicate shift cases the reverse obtains: only predicates applicable to the standard argument can be felicitously coordinated with the initial predicate, because the predicate, and not the subject, has undergone shift. Thus According to Nunberg, in (16) only predicates applying to the person (standard), but not the car (situation-appropriate), can be successfully coordinated or placed in apposition:

(22) I am parked out back and have been *waiting for 15 minutes* (Nunberg, 1995, Ex. 8)

(23) *I am parked out back and *may not start*. (Nunberg, 1995, Ex. 9. Judgment in the original)

(24) *I, *which is (am) a blue Chevrolet*, am parked out back (Nunberg, 1993, Ex. 87. Judgment in the original).²⁹

(25) I, *who am an excellent tipper*, am parked out back (Nunberg, 1993, Ex. 88).

According to these tests, then, (16) would exemplify predicate shift, whereas (17), would be an instance of deferred reference, and thus (in our terminology) of argument shift.

There are reasons, however, for doubting that these tests establish Predicate Shift. These will be reviewed in the next section. (Nunberg also bases his cases on evidence from agreement and anaphora; I won't discuss anaphora here but agreement will be reviewed in later sections).

3.3 Problems with Predicate Shift

Predicate Conjunction and Apposition

There are agent shift cases in which use of Nunberg's tests appears to lead to a diagnosis of *argument* shift, rather than predicate shift. Consider the following tripartite examples, where the (a) parts are agent shift sentences, the (b) parts are continuations that (if acceptable) favor an argument shift analysis of (a), and the (c) parts favor a predicate shift one (again, if acceptable):

(26)

[Context: Al, a very famous singer, is in a gallery, scrutinizing his new portrait.³⁰ The gallery is also exhibiting paintings by famous artists. Although Al finds his depiction unflattering, and though photographers are continuously pestering him, he is glad that his portrait was given a golden frame and was placed next to a painting by Kandinsky. Al calls his girlfriend on the phone and tells her the following:]

(26a) I'm framed in gold.

(26b) I am framed in gold and [*hung*] *next to a Kandinsky*.

²⁹ But notice that when you eliminate the relative pronoun and use a definite article, the utterance's acceptability seems to increase:

(24d) I, the blue Chevrolet, am parked out back.

(Thanks to Joshua Spencer for pointing this out to me).

We get a slightly degraded, but still acceptable version if we instead substitute an indefinite article:

(24i) I, a blue Chevrolet, am parked out back.

The same result can be obtained with (25).

³⁰ Here we are talking about a portrait *of him*, not a portrait (possibly of someone else) *made by him* or *owned by him*, etc.

(26c) ?I am framed in gold and *annoyed with the paparazzi*.

(27)

[Context: Suppose we walk out of a restaurant and find that a snowstorm broke out while we were having dinner (this example was adapted from Mount, 2008):]

(28a) I'm coated with ice.

(28b) I, *the blue Chevrolet*, am coated with ice.

(28c) ?I, *the excellent tipper*, am coated with ice.

It appears that, in all these examples, the (b) sentences are more felicitous than the (c) ones. This seems to indicate that (a) examples are instances of argument shift, rather than predicate shift.

VP Anaphora

We can also find evidence against predicate shift in cases of VP anaphora. If in (12) the predicate *parked out back* had really undergone transfer, to signify (say) 'being the owner of a car that is parked out back' (that is, a function defined over persons), then the following cases of ellipsis and gapping would *not* be possible, since Beth's car can't itself be the owner of a car (*mutatis mutandis* for the rest).³¹

(29a) I am parked out back and *so* is Beth's car

(29b) I am parked next to the oak tree, and Beth's car, next to the elm.

(29c) I am parked out back, and Beth's car is, too.

Emphatic Reflexives

³¹ The reason for that is the following: suppose that the first conjunct clause undergoes predicate transfer, and thus that the standard predicate of vehicles 'x is *parked out back*' is reinterpreted as 'x is *parked out back*₂', a predicate which applies to people (not to cars) in virtue of having a car that is parked out back. However, assuming that meaning is constant under ellipsis (as is standard practice), 'x is *parked out back*₂' must reappear in the second clause. But that is undesirable, since 'x is *parked out back*₂' cannot be satisfied by inanimate objects (such as Beth's car), and, as a consequence, the semantic mismatch remains unresolved, and predication is thwarted. This is an example of a general phenomenon: Predicate Shift can yield a coherent analysis of conjunctive sentences with predicate anaphora (such as those in (29)) only if the subjects of both conjunct clauses belong, in their standard interpretation, to the *same* category, relevant to the predicate at issue.

In contrast, categorically heterogeneous subjects accompanied by VP anaphora are not problematic for the DA, since the *loci* of shift are argument phrases, not predicates. Thus, the desired readings can be obtained simply by shifting one constituent, namely, the first subject phrase (so as to denote a car, a proper argument for the standard 'x is *parked out back*').

In situations when both the standard and metonymic readings of a phrase are available, we can block metonymic interpretations through an emphatic, pleonastic use of reflexive pronouns. This provides us with yet another argument against predicate shift. But first, to focus on the phenomenon I want of discuss, we must distinguish (non-exhaustively) between these two uses of emphatic reflexives:

(i) A non-contrastive use (or at least not necessarily contrastive), which can be paraphrased with a sentence containing ‘also’ ‘too’, or ‘for my part’, as in:

(30a) I am myself a wine lover.

(30b) I myself have a master’s degree in Film Studies.

(30c) I’m not a fan of Cheney myself, but ...

(ii) A contrastive or exclusive use, which could be paraphrased by sentences containing ‘in person’ (as in “not by proxy/deputy”). In these uses, occurrences of ‘-self’ receive stress, and are thereby focused:

(31a) Eric Clapton *himself* gave this guitar to me.

(31b) The President *himself* will be here tomorrow.

(31c) Why can’t I listen to the artist *himself* when I play his radio?? For example: I play Mitch Hedberg radio, but I get tracks from a lot of artists, none of them is Mitch (Last. fm Forum, 5.16.11).

(31d) The Pope *himself* will deliver the homily.

(31e) I *myself* will challenge the winner.

In the examples in (31) – which do not involve metonymy – felicitous use of the form Px + ‘-self’ seems to presuppose a set of relevant alternatives, containing satisfiers of P that are different from x. So, what ‘-self’ seems to be inducing here is the singling out of an object as a satisfier of P, out a set of other possible P-satisfiers. Presumably, the value of the expression modified by ‘myself’ (plus the context) can give us important clues for determining the members of the set of alternatives (respectively, Clapton’s assistants, the President’s representatives, musicians who play Hedberg-like music, less distinguished ecclesiastics, etc.). It seems that this operation takes place over a domain of objects or groups of objects.

Now let us consider *prima facie* cases of metonymy:

(32a) Clinton *himself* is next to Chaplin.

(32b) Beth *herself* is on the purple square.

(32c) Our Dear Leader *himself* is on the second floor.

(32d) Nixon *himself* dropped the bombs on Hanoi.

These examples do not admit of shifted interpretations. The same principle seems to be at work here: a set of alternatives is evoked, which in this case includes relevant things that could be on the second floor. This set (in the pertinent contexts) will include metonymical referents. So, in (32a) – uttered at Mme. Tussaud’s – the set would include Obama’s statue, for instance, and in (32c) it would contain Air Force pilots.³²

Now, for the cases we are interested about:

(33a) ?I *myself* am next to Chaplin. [Said by Obama at Mme. Tussaud’s, to indicate the location of his statue]

(33b) ?I *myself* will do that [Said by a department chair, at a college meeting, to communicate that the members of his department (*not* him) are going to undertake a certain task].

It looks like the same process is again responsible for forcing a standard reading of the subject: it singles out the *speaker* from among a set of possible alternatives, which includes possible targets of metonymy. Again, this makes sense only if we take the other members of the contrast class to be possible satisfiers of P. This appears to entail that there is a possible case where *Obama’s statue* can be the value of x in Px, and thus that argument shift is possible.

Now, I am not sure how the predicate transfer view would accommodate this phenomenon. Since it is (semantically) impossible for ‘P’ or [DP I] to shift its designation, a predicate transfer account would have to say that the contrast introduced by ‘myself’ does not take place among possible individuals, but (perhaps) among possible predicates.

This seems inadequate for the following reasons:

First: in (34) ‘myself’ is an adjunct to the subject noun phrase,³³ and agrees with it; it doesn’t belong syntactically with the predicate. Thus, it is reasonable to think that it will affect the interpretation of the subject, rather than the predicate. Second: This approach doesn’t apply to any of the other cases of (contrastive) emphatic ‘-self’. It can’t account for non-metonymical cases, and the nonindexical metonymical cases seem to work in the same way as the non-metonymical ones. So, if we desire a uniform treatment, we should reject the predicate shift view.

³² Notice that the point is not that metonymic/nonstandard readings can never be made salient by means of an emphatic reflexive, which is certainly false (witness, e.g. ‘the suits themselves showed up at the meeting’). Rather, the claim is that when *both* the standard and nonstandard interpretations of a DP are in the alternative set, the modification of that DP by an emphatic reflexive decisively favors the standard reading.

³³ Instances of emphatic, contrastive ‘-self’ can also occur in “adverbial” position: ‘I will catch the thief myself’, but the principle at work responsible for the interpretation seems to be the same.

Interlude: Agreement

Nunberg recruits (semantic) constraints on number and gender agreement as evidence for his view. Let us start with *number*:

(34) We are parked out back (Nunberg, 1993, p. 40).

An utterance of (34) is acceptable in a situation in which the speaker has only one car and someone is riding with her, but it can't be used if the speaker is a lone driver who has more than one car parked out back. But, if (34) involved a deferred-referential interpretation of 'we', we would expect it to be admissible in the second case, since the subject would be plural.

As for gender agreement, consider also a language that marks for gender, such as Italian. Here, the equivalent of 'I'm parked in the back' requires agreement with the ("real") gender of the speaker, not with the (grammatical) gender of 'car'. Thus, consider (36), uttered by an Italian man. Note that the noun for 'car' (*macchina*) is feminine.

(35) Io sono parcheggiato (**parcheggiata*, fem, sg.) dietro (Nunberg, 1995, p. 111. Judgment in the original).

Thus, agreement phenomena do seem to behave as Nunberg predicts. Nevertheless, and for reasons to be explained in the following paragraphs, I don't think they make a strong case for the predicate shift view.

In the first place, *number* data do not present a significant obstacle to the DA, since the view can straightforwardly account for Nunberg's purported counterexamples. To see how, notice that the DA can easily be applied to the second person singular pronoun, via a parallel line of reasoning:

(36a) *I* am parked out back → *My* [car] is parked out back → [The x: Rxs] x is parked out back.

(36b) *You* are parked out back → *Your* [car] is parked out back → [The x: Rxa] x is parked out back.

Here 'a' is a variable ranging over individual hearers, analogous to 's' in first person cases. The analysis can also be trivially extended to 3rd person and plural pronouns. What is important to see is that contents of the description within the brackets are systematically dependent on the grammatical features of the corresponding pronoun. In particular, those features determine the nature of the first *relatum*. In this way, the first person and singular feature values of 'I' determine the possible values of 's', and the second person, singular features of 'you' determine the possible values of 'a' (also, the PERSON feature entails ANIMATE, and since both of

these pronouns also have the ANIMATE feature, so does the corresponding variable). The general principle is as in (37), where F is a set of feature values, and the subscript _F, attached to a variable v, means that the domain of v is restricted to those objects satisfying F:

(37) pronoun_F [*predicate*] → [The x: R(x, v_F)] [*predicate*].³⁴

The NUMBER value of ‘we’ set to plural and PERSON is first person. Thus, its DA version must contain a variable v over plural entities, out of which we must select a group that is related in an appropriate way to the *speaker*³⁵ (most likely through a relation such as membership). Finally, the value of v must also be animate, i.e. must be a group composed of animates. Thus, a sentence such as ‘we are parked out back’ [=35] is paraphrased as ‘our [car] is parked out back’, and this can in its turn be taken as equivalent to (39), where u is a variable ranging over relevant speaker-including animate groups:

(38) [The x: Rxu] x is parked out back.

The members of a group comprised of two cars, unlike the group composed of the speaker and a companion, do not satisfy the animacy requirement, and so are not in the domain of u. Thus, in such a case, we have no value for u, and thus ‘we’ cannot be assigned an interpretation. Hence the infelicity of the utterance in the “one person, several cars” interpretation.

Things are muddier when it comes to the gender agreement data, and it is not entirely clear what they indicate. It may be that in sentences like (36) the adjective must have an *overt* agreement controller, such as a noun in object position (or the main verb’s morphology, etc.).³⁶ This would hold even if the overt controller is ultimately *interpreted* as part of a broader implicit phrase.

Thus, in agent shift case such as (36) the pronoun ‘*it*’ is the only overt term in controller position, and so it is *its* reference – i.e. the speaker, as the DA holds – that automatically determines the gender of the adjective. After this grammatical requirement has been met, further interpretive processes could take over, in particular those leading to the interpretation of the pronoun as if it were embedded in a description.

³⁴ Since R is functional, to every value of s, a, or u there corresponds at most one element in the value domain. This element could be an individual, or it could be a collective entity, as would be the case when a general points to the location on his army on a map and says: ‘I’m here’.

³⁵ The deictic component of the pronoun’s semantics, in Nunberg’s (1993) analysis.

³⁶ In this case the agreement is “semantically”, rather than “syntactically” driven, since the first person pronoun lacks a person feature. See Corbett (1991, ch. 2) for a discussion of this distinction.

3.4 *Summing up*

The balance of evidence seems in favor the DA as a general account of agent shift cases. Nevertheless, we must be open to the possibility that, despite initial appearances, agent shift is not a unified phenomenon, just as metonymy certainly is not. Consequently, we cannot rule out the possibility that no single story will be able to cover all the cases I have included under the heading of “agent shift”. These speculations, though interesting and worth looking into, cannot be adequately discussed in this essay, and so I will not attempt pursue them here. In any case, it is safe to conclude that DA accounts for at least some of the cases of agent shift.

4. - Conclusion

In this paper I argued for two conclusions. The first of them, which I regard as the most important, is that agent shift cases do not involve deferred reference and thus do not constitute a threat to the Standard View. This conclusion, which seems to me quite solid, is independent of the choice between the DA and Nunberg’s Predicate Transfer View, since both accounts agree on this score.

The second (somewhat more controversial) conclusion is that the DA gives the best account of agent shift cases.

According to the DA, occurrences of ‘I’ in agent shift cases are tacitly understood as providing one of the arguments for a contextually identifiable relation R, which is itself in the scope of a definite description operator. Thus, what ‘I’ contributes to the composition of propositions expressed by the utterances in which it occurs is the utterer or agent.

For Mount, agent-shift examples show that discretionary intentions are involved in the determination of the *reference* of ‘I’. Mount is right to point out that intentions play an indispensable role in the interpretation of Agent Shift utterances, and, specifically, in the determination of the designation of the shifted phrases. Nevertheless, acknowledging such a role is compatible with denying that discretionary intentions are involved in fixing the reference of ‘I’. In particular, within the DA analysis, speaker’s intentions are relevant because they contribute to the determination of the relation R between the *speaker* – the explicit argument in the relation, consisting in the referent of the occurring first person pronoun – and the target. This determination requires the previous identification of the reference of the pronoun, and so it is a precondition for, not a result of, the operation of intentions. If the DA is true, then, nothing in agent-shift utterances indicates that

discretionary intentions have a role to play in the interpretation of 'I'. Therefore, they give us no grounds for thinking that 'I' is impure.

Acknowledgement

An earlier incarnation of this paper was presented at the conference "Meaning, Context and Implicit Content", at Cerisy-La-Salle, France, 2011. I wish to thank the audience at the talk for their valuable feedback.

I am also grateful to David Braun, Christine Gunlogson, Dave Kush, Ishani Maitra, Tom McKay, Geoff Nunberg, Jay Rourke and Zoltán Gendler Szabó, as well as an anonymous referee for this journal, for their insightful comments, encouragement and advice.

Special thanks are due to Shannon Barrios for her patience in proofreading the manuscript.

REFERENCES

- Abney, S. P. (1987). *The English Noun Phrase in its Sentential Aspect*. Ph. D. thesis, MIT, Cambridge MA.
- Barker, C. (1995) *Possessive Descriptions*. Stanford, CA: Center for the Study of Language and Information.
- Bezuidenhout, A. (2008) Metaphorical Singular Reference. *The Baltic International Yearbook of Cognition, Logic and Communication*. 3, 1-22.
- Bianchi, C. (2001) Context of Utterance and Intended Context, *Lecture Notes in Computer Science*. 2116:73-86.
- Corazza, E., Fish, W., and Gorvet, W. (2001) Who is I? *Philosophical Studies*. 107 (1):1--21.
- Dodd, D. and Sweeney, P. (2009) Indexicals and utterance production. *Philosophical Studies*. 150, 3: 331-348.
- Donnellan, K. (1966) Reference and Definite Descriptions. *Philosophical Review*. 75, 3: 281-304.
- Fauconnier, G. (1994) *Mental Spaces*. Cambridge, MA: The MIT Press.
- Grice, P. (1989) *Studies In The Way of Words*. Cambridge, MA: Harvard University Press.
- Kaplan, D. (1989) Demonstratives. In Almog, J. Perry, J. and Wettstein, H. (Eds.) 1989 *Themes from Kaplan*. Oxford: Oxford University Press.
- Mount, A. (2008) The Impurity of Pure Indexicals. *Philosophical Studies* 138 (2).
- Nunberg, G. (1977) *The Pragmatics of Reference*. City University of New York PhD Dissertation.
- Nunberg, G. (1993) Indexicality and Deixis. *Linguistics and Philosophy* 16 (1993): 1-43
- Nunberg, G. (1995) The Pragmatics of Deferred Interpretation. In Ward G. and L. Horn (Eds.) *Blackwell Encyclopedia of Pragmatics*. Oxford: Blackwell.

- Perry, J. (1997) Indexicals and Demonstratives. In Hale, B. and Wright, C. (Eds.) *A Companion to the Philosophy of Language*, pp. 582-612. Oxford: Blackwell.
- Powell, G. (1998) The deferred interpretation of indexicals and proper names. *University College of London Working Papers in Linguistics* 10.
- Predelli, S. (1998) I am not here now. *Analysis*, 58, 107–115.
- Predelli, S. (2005) *Contexts: Meaning, Truth, and the Use of Language*. Oxford University Press
- Quine, W. V. (1968) Ontological relativity. *Journal of Philosophy* 65, 185-212.
- Quine, W.V. (1990) *Pursuit of Truth*. Harvard University Press.
- Recanati, F. (2004) *Literal Meaning*. Cambridge: Cambridge University Press.
- Romdenh-Romluc, K. (2006) 'I'. *Philosophical Studies*, 128: 257-283.
- Russell, B. (1905) On Denoting. *Mind*, 14, 56: 479-93.
- Sag, I. (1981) Formal semantics and extralinguistic context. In Cole, P. (Ed.) *Radical pragmatics*. New York: Academic Press.
- Schlenker, P. (2003) A plea for monsters. *Linguistics and Philosophy* 26. 29–120.
- Sidelle, A. (1991) The answering machine paradox. *Canadian Journal of Philosophy*, 21, 525–539
- Smith, Q. (1989) The Multiple Uses of Indexicals. *Synthese*, 78: 167-191.
- Stallard, D. (1993) Two kinds of metonymy. *Proceedings of the 31st annual meeting of the Association for Computational Linguistics*, 87-94. Columbus, OH: Association for Computational Linguistics.
- Ward, G. (2004) Equatives and Deferred Reference. *Language*, 80, 2: 262-89
- Warren, B. (2006) *Referential Metonymy*. Stockholm: Almqvist & Wiksell.
- Zobel, S. (2010) Non-standard Uses of German 1st Person Singular Pronouns. *Lecture Notes in Computer Science*. 6284: 292-311