Pronominal Anaphora, Coreference, and Closed Quotation Marks

Luca Gasparri, IJN Paris | April 26, 2018

(penultimate draft, final version published in *Mind & Language*)

Abstract

Consider the following sentence: "Mary meditated on the sentence 'Bill is a good friend' and concluded that he was a good friend". It is standardly assumed that in sentences of this sort, containing so-called "closed" quotations, the expressions occurring between quotation marks are mentioned and do not take their ordinary referents. The quoted NP "Bill" refers, if anything, to the name 'Bill', not to the individual Bill. At the same time, the pronoun "he", apparently anaphoric on quoted "Bill", refers to the individual Bill. The case seems thus to invalidate the intuitive principle that pronouns anaphoric on referential expressions inherit their reference from their antecedents. The paper formulates the argument, argues that sentences exhibiting the described pattern do not constitute evidence against the intuitive principle, and proposes an alternative account of the anaphoric relation involved.

1. SETTING THE STAGE

Consider (1) and (2).

1) Jane was at the party but Peter didn't see her.
2) Adam thinks that Murakami is a stuntman, but in fact he is a famous writer.

(1) and (2) share a common pattern. First, in both sentences there is a free pronoun ("her", "he") standing in a relation of anaphoric dependency with a referential antecedent ("Jane", "Murakami"). Second, in both sentences the free pronoun appears to inherit its reference from its antecedent. Witness (3), where the substitution of the antecedent originally featuring in (2) yields a change in value of the pronoun itself.

2) Adam thinks that Murakami is a stuntman, but in fact he [= Haruki Murakami] is a famous writer.

3) Adam thinks that DeLillo is a stuntman, but in fact he [= Don DeLillo] is a famous writer.

Examples like (1), (2) and (3) seem thus to encourage a very intuitive way of thinking about the interplay between anaphora and coreference in cases of pronominal anaphora on referential antecedents. The thought is summarized by the following principle.

(R-Inheritance) Pronouns anaphoric on referential expressions inherit their reference from their antecedents.

R-Inheritance leads to good predictions if tested against (1)-(3). However, it is natural to wonder whether the principle generalizes beyond the simple or "unproblematic" (King & Lewis, 2016) instances of anaphora exemplified by (1)-(3). The task of this paper is to evaluate an argument against R-Inheritance motivated by the behavior of a peculiar environment, which appears to generate a strong case for the rejection of R-Inheritance: *closed quotation*. The paper will review the signature properties of closed quotations, introduce the target case of sentences featuring apparent instances of pronominal anaphora on referential antecedents reaching across closed quotation marks, illustrate why their behavior threatens R-Inheritance, and determine whether sentences exhibiting the described pattern actually speak against R-Inheritance. The take-home message will be that even if sentences apparently allowing pronominal anaphora to reach across closed quotation marks raise significant concerns about the tenability of R-Inheritance, a more careful analysis of the anaphoric relation involved reveals that such sentences can be accounted for in a way that is consistent with the intuitive principle.

The discussion will proceed as follows. Section 2 will make some preliminary considerations about the interplay between coreference and pronominal anaphora on referential antecedents. Section 3 will formulate the quotation-based argument against R-Inheritance. Section 4 will describe two possible ways of rejecting the argument, based on two alternative analyses of the problematic pattern: one appealing to deixis, one appealing to linguistic salience. Section 5 will spell out and reject the analysis appealing to deixis. Section 6 will spell out and defend the analysis appealing to linguistic salience. Section 7 will conclude.

2. QUESTIONING R-INHERITANCE

As we have seen, R-Inheritance is consistent with the "unproblematic" instances of pronominal anaphora found in (1)-(3). Yet, as soon as one steps beyond the comfort zone of simple examples, the idea that pronominal anaphora on referential antecedents systematically entails coreference starts to look suspicious. There is consensus on this point. Neale (2005) has suggested that talk of coreference is ultimately misplaced in the attempt to produce a general theory of (pronominal) anaphora, which should rest on a weak notion of referential dependency rather than on the strong notion of sameness of reference. Safir (2005) has argued that the notion of coreference should play no role in theorizing about the foundations of anaphoric dependency, since only requirements of non-covaluation, rather than coreference, are primitive and represented in linguistic form. Similarly, Reuland (2011) has claimed that canonical Binding Theory should be replaced by a derivational approach focusing on the interaction between chain formation and the lexical properties of anaphors, and has defended a view where even requirements of non-covaluation are a syntactically derivative phenomenon.

To make a concrete example, assume the file terminology of Heim (1998), and suppose that the interpretation of a free pronoun proceeds by redeploying the same file initially recruited to process its antecedent. Then recall (1).

1) Jane was at the party but Peter didn't see her.

On the file-based explanation, the JANE file associated to the (disambiguated) name 'Jane' is first deployed to process the NP "Jane", and then redeployed to interpret the pronoun. The process can, and does in (1), end up attributing the same reference to the expressions involved. But there are cases in which this mechanism of reference inheritance does not obtain, or at least does not obtain so unambiguously. Co-predication sentences such as (4) and (5) (inspired by Chomsky, 2000 and Asher, 2011) illustrate the point.

4) The book George took from the library is on Pam's desk. It was originally written in Chinese.
5) I felt disappointed by Harry's report and Sue found her cat sitting on it.

Take (4). The entity predicated of the property of "having been originally written in Chinese" cannot be the same entity that "George took from the library" and "is on Pam's desk". The first clause of (4) predicates a property of a physical object, whereas the second is clearly describing an abstract type. Since these are different entities, the natural conclusion would seem that we should drop R-Inheritance.

However, the advocate of R-Inheritance might not be impressed. Cases like (4) and (5) place an additional explanatory burden on the systematic intertwinement between pronominal anaphora and coreference envisioned by R-Inheritance. Yet, it is not obvious that the bill cannot be paid. In particular, two immediate lines of reply come to mind. One goes through an extension of the descriptive scope of the notion of coreference. Call it, for brevity, the Descriptive Option. The other appeals to pragmatics. Call it, for brevity, the Pragmatic Option.

The Descriptive Option would be to claim that the speaker-accessible difference in reference between the pronoun and the antecedent of (4) does not suffice to invalidate the claim that the two terms are coreferential. Think of the way the physical object vs. abstract type alternation surfacing from the example is captured within frameworks such as Pustejovsky (1995). The basic idea is that the perceived systematicity of the relation between "book" qua physical object and "book" qua abstract type is best illuminated by seeing the two meanings as specifications of a unique, complex semantic entity of type INFORMATIONAL OBJECT • PHYSICAL OBJECT whose facets can be selectively picked out by the predicates "being on Pam's desk" and "having been originally written in Chinese", and thereby make co-predication possible. Once the conditions of application of the notion of coreference are relaxed by stipulation so that expressions picking out specifications or facets of a unique dotted type qualify as coreferential, the pronoun-antecedent pair of (4) corefers and R-Inheritance is safe (for more on this line of argument, see Recanati, 2018).

The Pragmatic Option would be to construe R-Inheritance as a principle narrowly about semantic structure and argue that the reading explicitly entertained by subjects while interpreting the antecedent-pronoun pair of (4) does not reflect the underlying semantics of the sentence. In other words, "the book" and "it" may be fully coreferential for the semantic machinery, and the speaker-accessible difference in reference between the two terms might arise only when the output of grammatical calculus is interfaced with probabilistic reasoning and general intelligence, which modulate the values initially assigned to the two expressions by the semantics and yield the referential asymmetry consciously entertained by pragmatically competent interpreters of the sentence. Supposing that R-Inheritance is threatened only if the referential discrepancy surfacing from (4) is a faithful image of the semantics of the sentence, and that the perceived difference in reference between "the book" and "it" could instead be generated pragmatically, the argument against R-Inheritance via (4) is not guaranteed to succeed.

The provisional moral is simple: no matter how coarse-grained, crude or simplistic it might sound, R-Inheritance is not going to be taken down easily. However, it seems that a stronger case against R-Inheritance is available. The next section presents the case and illustrates the reasons why it undermines the principle.

3. THE ARGUMENT FROM CLOSED QUOTATION

It is generally accepted that there are two basic varieties of quotation: *closed* and *open* quotation (e.g., Recanati, 2001; Potts, 2007; Cappelen & Lepore, 2012; Maier, 2014). Consider (6), an example of closed (pure) quotation, and (7), an example of open (mixed) quotation (adapted from Reinhart, 1983).

6) The phrase "exploits the secretary who works for him" expresses a predicate.

7) Martin said that each manager "exploits the secretary who works for him".

The key difference between the two examples is that while the quotation of (6) presents a string of terms while preventing their linguistic features from interacting with grammatical structure at the matrix level, in (7) the quoted predicate contributes to grammatical structure just as it would do in a non-quotational context, and responds to the principles of general compositionality. Approximating a bit, while the closed quotation of a VP prevents the quoted VP from being interpreted as a VP by grammatical structure at the matrix level, the open quotation of a VP allows the quoted VP to be interpreted as a VP by grammatical structure at the matrix level (the same line of reasoning extends, with the due adjustments, to other grammatical categories).

Let us leave open quotation aside and focus on the contrast between sentences containing closed quotations, such as (6), and sentences containing indirect reports, such as (2). Here is a representative sample of their differences.[1]

The first difference is that while there can be grammatical dependencies between constituents in an indirect report and its embedding clause, closed quotations generally inhibit grammatical dependencies between quoted and used linguistic material. While in indirect reports the semantico-syntactic attributes of the reported material contribute to the grammatical structure of the rest of the sentence, closed quotations operate as indecomposable units whose internal semantico-syntactic features are segregated from the rest of the sentence. Witness the difference in well-formedness between (8) and (9), where the quotation marks affect movement, the difference in meaning between (10) and (11), where the quotation marks prevent the NPI from being licensed by "nobody", and the difference in felicity between (12) and (13), where quotation marks absorb the failure in double access reading of the predicate under the verb of saying.

8) [ok] Who did John think Mary loves?

9) * Who did John think "Mary loves"?

10) [ok] Nobody said Franz has ever been to Kathmandu.

11) [ok] Nobody said "Franz has ever been to Kathmandu".

12) [#] Ten years ago, Lea said Lucy is pregnant.

13) [ok] Ten years ago, Lea said "Lucy is pregnant".

1 In what follows, I'll be relying on written sentences, where the boundary between quoted and used linguistic material is noticed straight off due to the presence of quotation marks. However, the analysis extends to spoken language, where the same boundary is marked, albeit less determinately, through prosody and voice quality: closed quotations are usually preceded and followed by a pause, and are often pronounced with the intonation pattern that the quoted sentence would have in isolation. See, e.g., Günthner (1999).

A second difference is that while indirect reports must be grammatical, the material contained in a closed quotation does not have to make up a well-formed sentence: it can be an incomplete sentence fragment, a pseudoword, an utterance in a different language, or a string of non-linguistic sounds. See (14)-(19).

14) [ok] William said “à l’oeuvre, on connaît l’artisan”.

15) * William said à l’oeuvre, on connaît l’artisan.

16) [ok] William said “ixtrit”.

17) * William said ixtrit.

18) [ok] William thought she was smart, that is, intelligent.

19) [?] William thought “she is smart”, that is, “intelligent”.

Third, while it is possible to quantify into an indirect report, it is not possible to quantify into a closed quotation. The pronoun remains trapped inside quotation marks and yields a formula which mentions the pronoun rather than expressing its value. For example, (21) follows from (20), but (23) does not follow from (22). The same goes for syntactic binding, as shown by the lack of entailment between (24) and (25). See also the lack of entailment between (26) and (27). Co-indexing “her” and “every girl Carl met” seems initially possible. However, allowing for covariation in substitution instances among the quantified expression and the in-quote pronoun would invalidate the initial description of the words uttered by Carl.

20) Julie believes that Scotland is beautiful.

21) ∃x (Julie believes that x is beautiful)

22) “Scotland” has eight letters.

23) ∃x (“x” has eight letters)

24) [ok] To every girl he met, Carl said he loved her.

25) [ok] To every girl he met, Carl said "I love her".

26) * The sentence "I love her$_1$" was uttered by Carl about [every girl Carl met]$_1$.

27) [ok] The sentence "I love Melissa" was uttered by Carl about Melissa.[2]

All this is taken to indicate that sentences containing closed quotations describe a relation between a subject and a black box of linguistically inert or "merely mentioned" expressions, rather than a relation between a subject and a string of linguistically active or "used" expressions. Accordingly, that closed quotations should be understood as devices whose primary function is to picture or demonstrate linguistic items (Clark & Gerrig, 1990 is perhaps the *locus classicus*). The idea has received formal implementation in recent years. Potts (2007) proposes to characterize (closed) quoting as a function that derives a singular-term type whose denotation is a linguistic entity from any well-formed expression in a language. If expressions are modeled as triples of form $\langle \Pi, \Sigma, \alpha: \sigma \rangle$, where Π is a phonetic representation, Σ is a syntactic representation, and α is a

2 The asterisk before (26) is intended to signal that "her" and "every girl Carl met" cannot bind, not that the sentence is ungrammatical per se. To understand why "her" and "every girl Carl met" cannot bind, consider a world W in which the set of girls met by Carl includes Melissa. Suppose also that "The sentence 'I love her' was uttered by Carl about every girl Carl met" is true in W. Now assign the quantified expression "every girl Carl met" the value "Melissa". We obtain "The sentence 'I love her' was uttered by Carl about Melissa". If "her" and "every girl Carl met" were eligible to a bound reading, the in-quote pronoun should be allowed to co-vary with the value assigned to the quantified expression. Hence, if "her" and "every girl Carl met" were eligible to a bound reading, then, since "The sentence 'I love her' was uttered by Carl about every girl Carl met" (i.e., (26)) is true in W, "The sentence 'I love Melissa' was uttered by Carl about Melissa" (i.e., (27)) should also be true in W. However, "The sentence 'I love her' was uttered by Carl about every girl Carl met" and "The sentence 'I love Melissa' was uttered by Carl about Melissa" give two competing descriptions of the exact words pronounced by Carl about Melissa ("I love her" in the former case, "I love Melissa" in the latter case). Hence, "her" and "every girl Carl met" cannot bind: in a world W where Carl meets Melissa, the truth of (26) does not entail the truth of (27). Thanks to Mikhail Kissine for pressing me to make this point explicitly.

semantic representation of type σ, quoting can be defined as $f(\langle \Pi, \Sigma, \alpha: \sigma \rangle) = \langle \Pi, \Sigma, \ulcorner \langle \Pi, \Sigma, \alpha: \sigma \rangle \urcorner: u \rangle$ (where u is the singular-term type outputted by the function), with $[\![\ulcorner \Pi, \Sigma, \alpha: \sigma \urcorner]\!] = \langle \Pi, \Sigma, \alpha: \sigma \rangle$. Similarly, de Vries (2008) maintains that (closed) quoting can be modeled as a function f(S) that converts a string of linguistic items S into a syntactically nominal category [$_N$ "S"] which operates as an atomic element and whose semantic content is not analyzed as part of the sentence embedding "S". Pagin & Westerståhl (2010) also present extensive evidence that closed quotations do not behave like ordinary constituents of quoting sentences.

This body of evidence appears to license a broad and uncontroversial principle about closed quotation. Call it, for brevity, CQ-Shift (where 'CQ' abbreviates "closed quotation" and the term "shift" signals that referential expressions, once positioned within closed quotation marks, undergo a shift that makes them drop their ordinary referential properties).

(CQ-Shift) Referential expressions in closed quotations do not contribute their ordinary referents.

For instance, tokens or occurrences of "exploits" normally denote the action or the event of exploiting. However, in the quotation environment of (6), "exploits" contributes the word 'exploits', not the action or the event of exploiting.[3]

6) The phrase "exploits the secretary who works for him" expresses a predicate.

3 For example, this is the main observation fueling Kaplan's (1989) claim that (closed) quotations in English are not monsters. In a sentence like "In some contexts 'I am tired' is a truism", the quoted indexical "I" is not used and does not refer to any speaker in particular; it is merely mentioned or "referred to". Thus, closed quotations are not a counterexample to the thesis that indexicals like 'I', 'you', 'here', and 'today', when used, always take "primary scope" and refer to the actual speaker, addressee, location, and time of the utterance. For more on quotation and monsters see, e.g., Maier (2016).

With CQ-Shift in mind, consider (28) (inspired by examples (11a-c) in Partee, 1973).

28) Mary meditated on the sentence "Bill is a good friend" and concluded that he was a good friend.[4]

In (28), the free pronoun "he" seems to be anaphoric on the in-quote NP "Bill". The case should therefore be captured by R-Inheritance. Now, in order for R-Inheritance to be consistent with the example, the pronoun "he" should be coreferential to "Bill". Yet, it seems uncontroversial that while the pronoun "he" refers to the individual Bill, the in-quote NP "Bill" responds to CQ-Shift and evaluates to something other than the individual Bill, presumably the proper name 'Bill' itself (nothing crucial hinges on this). Clearly enough, the phenomenon raises important concerns about the tenability of R-Inheritance. In (28), the free pronoun "he" appears to be anaphoric on quoted "Bill". Based on R-Inheritance, the two expressions should be coreferential. But the two expressions cannot be coreferential, since "Bill" is subject to CQ-Shift and does not contribute the individual Bill, whereas the pronoun is readily interpreted as contributing Bill himself. Once again, the natural conclusion seems that we should abandon R-Inheritance.

Let us call this line of attack to R-Inheritance the Argument from Closed Quotation. Here is a more orderly reconstruction of how the argument runs ('P' abbreviates "Premise", 'S' abbreviates "Step").

4 In the event that the example sounds odd to your ears and, e.g., you wonder how meditating on a single short sentence from an unknown source may lead anyone to conclude that someone is a good friend, think of (28) in the following context. Mary has been having aching doubts about the loyalty of her friend Bill and decides to reach out for help from her aunt Anna, whose advice and opinion Mary values immensely. Aunt Anna is certain of Bill's good intent and tells Mary that she shouldn't worry because "Bill is a good friend". On her way home, Mary ponders aunt Anna's sentence and convinces herself that Bill is in fact a good friend.

Argument from Closed Quotation

(P1) In (28), "he" is anaphoric on the quoted NP "Bill".

(P2) The quoted NP "Bill" of (28) does not refer to the individual ordinarily designated by the proper name 'Bill'. [= CQ-Shift]

(S3) The pronoun "he" of (28) does not refer to the individual ordinarily designated by the proper name 'Bill'. [P1-2, R-Inheritance]

(S4) The clause "he was a good friend" featuring in (28) predicates the property of being a good friend of something other than the individual ordinarily designated by the proper name 'Bill'. [S3]

(P5) S4 is undesirable.

(S6) S3 should be rejected. [P5, S4]

(P7) Rejecting S3 requires rejecting either P1, or P2, or R-Inheritance.

(P8) P1 and P2 are uncontroversial.

(C) R-Inheritance should be rejected. [S6, P7-8]

The argument, a version of which is endorsed by Smit & Steglich-Petersen (2010), generates a strong case against R-Inheritance. Notice that the issue appears bound to surface no matter what specific approach to the nature of unbound anaphora one chooses to adopt. E.g., whether one chooses to model unbound pronouns in line with a non-descriptive approach such as classical Discourse Representation Theory (Kamp, 1981; Heim, 1982) or to view them as proxies for definite descriptions à la Elbourne (2005). Furthermore, a version of the issue extends to cases of non-pronominal anaphora that uncontroversially do not entail coreference, such as modal anaphora (Stone & Hardt, 1999). Consider (29).

29) Rebecca meditated on the sentence “A puma might come in” and concluded that she would be terrified.

Interestingly, (29) is felicitous even if, on account of the grammatical inertness enforced by the closed quotation, the quoted linguistic material should not be able to introduce an antecedent for the conditional mood.

Notwithstanding all this, I shall argue that the argument should be resisted. Investigating whether the Argument from Closed Quotation really invalidates R-Inheritance is important for two reasons. First, I believe that accepting the argument leads to a sub-optimal explanation of the way pronominal anaphora operates in sentences exhibiting the pattern of (28). In other words, even if R-Inheritance were ultimately untenable, believing it is because of sentences like (28) would be a mistake. Second, to the best of my knowledge, on the market there is no satisfactory account of how anaphora can reach across quotation marks in cases like (28), where it is plausible to assume that the quoted material is opaque to the compositional structure of the rest of the sentence, and yet manages to provide an antecedent for a pronoun outside the quotation. Questioning the Argument from Closed Quotation will bring us closer to such an account.

4. NEUTRALIZING THE ARGUMENT

There are three possible ways of neutralizing the Argument from Closed Quotation. Call them, for brevity, Option 1, Option 2, and Option 3.

Option 1: Block the argument at P5.

Option 2 : Go through P7-8 and block the argument at P2.

Option 3 : Go through P7-8 and block the argument at P1.

Option 1 (i.e., welcoming as acceptable the idea that the clause "he was a good friend" featuring in (28) predicates the property of being a good friend of something other than the individual ordinarily designated by the proper name 'Bill') is a reinstatement of the Pragmatic Option described in Section 2. The strategy would be to claim that for the semantics the clause "He was a good friend" does in fact predicate the property of being a good friend of something other than the individual normally contributed by 'Bill', and that the speaker-accessible reading of the pronoun obtains thanks to further pragmatic inferencing. In a nutshell, the pronoun "he" in (28) is first interpreted by the semantics in the metalinguistic sense (i.e., as referring to the proper name 'Bill'), this interpretation is then fed into the proposition recovered by the grammar, tested for contextual likelihood, and finally replaced via post-semantic inferential work with the speaker-accessible value of the pronoun.

There is nothing wrong, in principle, with appealing to this line of thinking. However, the hypothesis faces two immediate problems. First, the reply would be satisfactory on condition that an account of the computation involved is given rather than merely postulated, which seems far from an unproblematic task. Second, there is an issue of agreement. Consider (28), (30), and (31).

28) Mary meditated on the sentence "Bill is a good friend" and concluded that he was a good friend.

30) John pictured the word 'eureka' and screamed it.

31) Dave mumbled "croissant..." and decided to have one.

In cases where it is uncontroversial, as in (30), or at least possible, as in (31), that the metalinguistic reading of the quoted NP plays the role of propositional constituent paired to the pronoun by the semantics, the pronoun has the appropriate morphology (e.g., the pronoun of (30) can refer only to entities with no natural gender). However, the pronoun of (28) is masculine. Thus, it is more likely

that in (28) the interpretive process is geared from the start to the recovery of a proposition where the pronoun is assigned its speaker-accessible value (i.e., the individual Bill), rather than arriving at the desired outcome through some two-step reanalysis of an intermediate propositional template pairing the pronoun with a linguistic object (i.e., the name 'Bill').

Option 2 (i.e., reject the premise that the quoted NP "Bill" of (28) does not contribute the individual ordinarily designated by the name 'Bill') is a de facto suspension of CQ-Shift. Interestingly, approaches to closed quotation such as the Identity Theory originated from Washington (1992) do pursue the hypothesis that the proper name quoted by (28) is actually used rather than mentioned, contra CQ-Shift. A somewhat less revisionary strategy would be to appeal to a variant of Davidson's (1984) "double reference" account, and hypothesize that the expressions between the quotation marks of (28) are somehow both mentioned and used. On this analysis, the Argument from Closed Quotation fails because it mistakenly assumes that the NP "Bill" featuring in the quotation of (28) evaluates only to the name 'Bill', whereas in reality it evaluates both to the individual Bill *and* to the proper name 'Bill', and the pronoun corefers to "Bill" with respect to the former of its two contributed referents.

Once again, there is nothing wrong, in principle, with pushing this line of reply. CQ-Shift rests on the widely well-regarded assumption that open and closed quotation must be given different theoretical treatments, and this assumption can be denied. One could even conjecture that precisely because there are independent motivations in favor of R-Inheritance and anti-CQ-Shift views allow for a painless refutation of the Argument from Closed Quotation, the present discussion provides reasons to think that such views are on the right track. However, there are two drawbacks. First, there is no agreement on whether anti-CQ-Shift theories provide a viable explanation of linguistic data. For example, they fare pretty poorly when it comes to illuminating in any sensible way the phenomena affecting binding highlighted by (20)-(27), and seem to be committed to the wrong prediction that quoted indexicals receive their normal evaluation (recall fn. 3). Second,

considerations of reflective equilibrium suggest that the ideal reply to the Argument from Closed Quotation should not appeal to unorthodox approaches to the nature of quotation. More precisely, since the foreground of our analysis centers on the interplay between pronominal anaphora and coreference rather than on the nature of quotation, our task is best addressed by trying to make sense of the problem raised by (28) in line with a vanilla view of quotation, rather than by appealing to a theory that goes against mainstream agreement on the properties of quotation environments.

We are then left with Option 3: abandoning the initial premise that in (28) "he" is anaphoric on the quoted NP "Bill". There are two basic ways in which the rejection of P1 can be enforced, depending on the scope one chooses to assign to the negation of the premise.

Wide: In (28), "he" is *not* [anaphoric on in-quote "Bill"]

Narrow: In (28), "he" is anaphoric on *not* [in-quote "Bill"]

The difference between the two proposals is clear enough. On the former diagnosis, the unbound pronoun of (28) is not an anaphor. Hence, it should be construed as a deictic particle. For brevity, call this the Deixis Hypothesis. On the latter diagnosis, the pronoun of (28) is in fact an anaphor. However, the expression with which the pronoun stands in the anaphoric relation is not the quoted NP "Bill", but a contextually supplied antecedent. For brevity, call this the Contextual Antecedent Hypothesis. Both the Deixis Hypothesis and the Contextual Antecedent Hypothesis neutralize the Argument from Closed Quotation. Yet, they give two sharply different accounts of the way value assignment to the unbound pronoun of (28) operates. Our task is to determine which of these two accounts shows more promise.

5. COREFERENCE BY DEIXIS

Let us start by examining the Deixis Hypothesis. On this view, the pronoun "he" of (28), rather than being anaphoric on "Bill", operates as a deictic particle whose value, the individual Bill, is determined relative to a referent raised to salience by the in-quote NP "Bill". Smit & Steglich-Petersen (2010) object that embracing this line of thinking would force one to uphold that even the pronoun occurring in (32) is deictic, which, they argue, is problematic.

32) Mary said that Bill was not a good friend, but in fact he was a good friend.

But the parallel pressed by the objection is debatable, for two reasons.

The first reason is that it is not clear why the choice of characterizing the pronoun of (28) as a deictic particle interpreted in virtue of a raising-to-salience mechanism should care to avoid the consequence that the pronoun of (32) is not anaphoric. After all, as we saw in Section 3, closed quotation marks make a difference in terms of grammatical integration, as they curtail the set of semantico-syntactic relations available between material inside and outside the quotation. Thus, it appears perfectly sensible to argue that the absence of closed quotation marks allows the pronoun-antecedent pair of (32) to result in a full-fledged anaphoric relation, whereas the putative pronoun-antecedent pair of (28) is better described as linked by a deictic dependency which happens to have a surface realization that resembles a genuine case of pronominal anaphora.[5] Overall, coupling the premise that the pronoun of (28) is deictic with the premise that the distribution of linguistic items of (32) patterns along the surface structure of (28) appears insufficient to press the conclusion that the pronoun of (32) should also be classified as deictic.

5 For example, some have maintained that referential expressions have to be "syntactically explicit" in order to function as antecedents and take part in relations of bona fide anaphoric dependency (e.g., Heim, 1990). From this standpoint, the reply would be that while in (32) "Bill" is syntactically realized as an NP and can be recruited to play the antecedent-role, in (28) the closed quotation prevents "Bill" from being syntactically realized as an NP, and anaphora fails.

The second reason is the following. On the one hand, the objection accepts that closed quotations act as indecomposable units whose internal semantico-syntactic properties are segregated from matrix grammatical structure. On the other, the objection seems to suggest that, despite the semantico-syntactic segregation imposed by the closed quotation, (28) still allows for the structural dependency that should obtain between "Bill" and "he" in order for the pronoun to be classified as a canonical anaphor. The quotation marks, hence, should both segregate the quoted sentence when it comes to allowing the penetration of syntactic indices for binding (recall (24)-(27)) and leave it open when it comes to allowing the penetration of syntactic indices for unbound anaphora. But once we agree that in cases like (28) the quotation enforces the structural segregation exemplified by (8)-(13) and (20)-(27), and that co-indexing for bound anaphora cannot travel across closed quotation marks, why exactly should one allow free anaphora to penetrate the closed quotation?

Suppose, furthermore, we do grant that the Deixis Hypothesis inevitably leads to the conclusion that the seeming pronoun-antecedent pair of (32) does not result in anaphora. There is room to argue that the upshot would not be as bad as the objection wants us to believe. For example, Schlenker (2005) would be comfortable to say that since in (32) there is no "formal connection" (i.e., binding) between the pronoun "he" and its alleged antecedent "Bill", the pronoun should be treated as a deictic particle from the point of view of syntactic typology. In general, any view of pronominal anaphora broadly in agreement with the orthodoxy of Reinhart (1983) (roughly: any view of pronominal anaphora maintaining that talk of anaphora should be restricted to variable binding) would be happy to uphold that the pair of dependent terms appearing in (32) should not be viewed as engaging into the same variety of uncontroversially anaphoric dependency involved, e.g., in cases of binding to quantified antecedents.

There is, however, one more serious problem for the hypothesis that the pronoun of (28) should be treated as a particle demonstrating a salient referent. The issue is that the hypothesis fails to generalize to neighboring cases. To understand the point, consider (33) and (34).

33) If Kate encountered in the newspaper an announcement saying "A great Indian guru is coming to town", she would want to go meet him.

34) The sign at the zoo read "The Bengal tiger is in danger". However, it forgot to mention that humans are its most significant predator.

Suppose, for the sake of argument, that the pronoun "he" of (28) operates in line with the Deixis Hypothesis: the quoted NP "Bill" introduces (or raises to salience) a referent which is then demonstrated by the pronoun. The explanation fits (28), but does not extend to (33) and (34). The indefinite and the generic opening the quotations of (33) and (34) do not raise to salience any individual that may be picked out deictically by "him" and "its". Yet, e.g., (33) gives us a felicitous instance of conditional donkey anaphora. Since it stands to reason to assume that our account of the mechanisms licensing the use of pronoun in (33) and (34) should run on par with our account of the mechanisms licensing the use of pronoun in (28), (33) and (34) raise an issue for the Deixis Hypothesis. Our current hypothesis would predict that sentences exhibiting the pattern of (28) are acceptable because the putative anaphor "he" manages to take a salient referent by deixis, but (33) and (34) show that pronominalization is licensed even if the expressions relative to which the pronoun is evaluated do not contribute anything that could possibly be demonstrated by an unbound pronoun.

6. CONTEXTUALLY SUPPLIED ANTECEDENTS

Let us thus examine the Contextual Antecedent Hypothesis. This analysis of (28) accepts that the pronoun "he" appearing in the sentence is in fact an anaphor, but argues that its antecedent, rather than "Bill" qua constituent occurring in the quotation environment, is a contextually supplied expression. More specifically, the proposed diagnosis of (28) is the following. The quotation of (28)

raises to salience *the quoted sentence itself*, and "he" is anaphoric on "Bill" qua constituent of the sentence so given contextual prominence by the closed quotation rather than on "Bill" qua constituent of (28). (35) provides an intuitive representation of the point.

35) Mary meditated on the sentence **"Bill is a good friend"** and concluded that <u>he</u> was a good friend.

Makes salient in context: <u>Bill</u> is a good friend

The key difference between the analysis of (28) produced by the Deixis Hypothesis and the analysis of (28) produced by the Contextual Antecedent Hypothesis is that while on the Deixis Hypothesis value assignment to the pronoun proceeded by inspecting the worldly referents raised to salience by the referential expressions in the quotation, on the Contextual Antecedent Hypothesis the quotation is understood as a device for raising to salience *the quoted linguistic material itself*, and the non-quoted pronoun is anaphoric on the appropriate constituent of the sentence raised to salience in the context of interpretation, rather than on the NP "Bill" qua constituent of (28). In short, the closed quotation pushes the sentence "Bill is a good friend" to gain prominence on the contextual blackboard, and the linguistic material so given contextual prominence becomes available to undergo semantico-syntactic inspection in cases where, as in (28), the operation is necessary for the interpretation of the constructions flanking the quotation. So there is a pronoun-antecedent coreference relation, just like R-Inheritance wanted. But the coreference relation, rather than reaching across quotation marks and obtaining between "he" and "Bill" qua constituents of (28), obtains between "he" qua constituent of (28) and "Bill" qua constituent of the sentence given prominence on the contextual blackboard by the closed quotation. The appropriate distribution of syntactic indices for (28) is therefore the one given in (37), rather than the one in (36). At best, the indexing pattern of (36) can be used to signal that the in-quote term putatively operating as the

antecedent of the unbound pronoun is phonologically equivalent to the contextually supplied expression that actually enters in the anaphoric relation with the pronoun. Correlatively, (38) gives the appropriate distribution of the donkey indices for (33).

36) * Mary meditated on the sentence "Bill$_1$ is a good friend" and concluded that he$_1$ was a good friend.

37) Mary meditated on the sentence **"Bill is a good friend"** and concluded that he$_1$ was a good friend.

Makes salient in context: Bill$_1$ is a good friend

38) If Kate encountered in the newspaper an announcement saying **"A great Indian guru is coming to town"**, she would want to go meet him$_i$.

Makes salient in context: A great Indian guru$_i$ is coming to town

Besides neutralizing the Argument from Closed Quotation, the account of (28) produced by the Contextual Antecedent Hypothesis seems to exhibit a few immediate virtues. First, it aligns nicely with the mainstream agreement that quotations should be understood as devices for the demonstration of linguistic material (recall Section 3). Second, it generalizes smoothly to the donkey anaphora of (33) and other instances of non-pronominal anaphora: e.g., (29) is felicitous because the quoted sentence introduces an antecedent for the conditional mood by being salient on the contextual blackboard. Third, the proposed treatment of (28) licenses pronominal anaphora while not forcing indices to travel across quotation marks: throughout the interpretive procedure, the closed quotation remains isolated from the compositional structure of the rest of the sentence.

Maier (2014) gestures at a similar pragmatic mechanism of raising to salience of linguistic material, which, building on Grice, he refers to as "bridging". However, Maier adds that any account arguing along (in my terminology) the Contextual Antecedent Hypothesis is bound to face

an important complication. Consider (28), (39) (adapted from Recanati, 2010), and VPC (from Hankamer & Sag, 1976).

28) Mary meditated on the sentence "Bill is a good friend" and concluded that he was a good friend.

39) Emma said to herself "I am calling the dean right now". And she did [*call the dean*].

(VPC) VP-ellipsis cannot get an antecedent derived pragmatically from context.

Faced with (39) and VPC, it is natural to wonder whether evoking the pragmatic mechanism envisioned by the Contextual Antecedent Hypothesis to account for value assignment to the pronoun of (28) is still a viable move. The reasoning is the following. According to the proposed story, the pronoun of (28) is assigned an antecedent thanks to the contextual prominence of the quoted sentence "Bill is a good friend". Now, if the Contextual Antecedent Hypothesis provides the correct analysis of (28), it stands to reason to hypothesize that a similar pragmatic mechanism should be responsible for the interpretation of the ellipsis of (39). Yet, VPC states that the ellipsis of (39) cannot be resolved by appeal to contextual inferencing. Absent the possibility of accounting for the resolution of the VP-ellipsis featuring in (39) in pragmatic terms, (39) would then seem to substantiate the idea that the quoted material of (39) is not linguistically inert. More precisely, if the copy of the VP "call the dean" into the ellipsis site of (39) cannot be explained by appeal to contextual inferencing, it seems to follow that it should rely on the intra-sentential inspection of the grammatical structure of the quotation, a possibility that would obtain only if the grammatical structure of the quotation were in fact readable by the constituents outside the quotation (contra the "black box" orthodoxy). Were that the case, the same could hold for (28). So perhaps, after all, we should entertain the possibility that the semantico-syntactic structure of the closed quotation of (28) is not segregated from the rest of the sentence, and that "Bill" and "he" stand in a relation of

canonical anaphoric dependency giving rise to the same dynamics of reference inheritance found in unproblematic examples like (1)-(3). Cases such as (40), where the nominal attached to the non-predicative "most probable" seems to be extracted from the quotation context, and (41), where the complement reading of "the others" is inferred relative to "a few students", raise a similar worry.

40) Megan mumbled "Mark might have left for a million reasons", but did not think that the most probable was Claudia's annoying attitude.

41) "Only a few students left in the afternoon", said David, implying that the others had all left in the morning.

There is, however, one simple reply to offer on behalf of the Contextual Antecedent Hypothesis. Namely, (39) does raise an issue for the Contextual Antecedent Hypothesis under the assumption of VPC. But VPC is dispensable. As it turns out, in fact, the ellipsis of (39) can be modified into an example that backfires at the VPC-based objection. Compare (39) and (42).

38) Emma said to herself "I am calling the dean right now". And she did [*call the dean*]

42) [*Context: Emma and Alex are reading a hand-written post-it which contains the sentence "Paul will see the dean today". Emma turns to Alex and says*:]
He did [*see the dean*], eventually.

On the Contextual Antecedent Hypothesis, the dynamics underlying the interpretation of VP-ellipsis in (39) and (42) are fundamentally the same. In both cases, VP-ellipsis is resolved by inspecting the structure of a sentence raised to salience in the context (by the quotation in (39), by the extra-linguistic situation in (42)) and extracting from the contextually prominent sentence the chunk of phrase structure that needs to be copied into the ellipsis site. Accordingly, the lesson we should

draw from (28)'s comparison to (39) and (42) is not that the quoted "Bill" of (28) is grammatically integrated to the rest of the sentence, but simply that, pace VPC, context *can* in fact feed the resolution of VP-ellipsis or provide a referential antecedent to an anaphoric pronoun, provided that some specific conditions are met (see, e.g., Merchant, 2010). In cases like (42), the conversational participants need to access some salient linguistic material whose inspection suffices to ensure a coherent reconstruction of the elided VP. In cases like (28), the conversational participants need to access some salient linguistic material which contains at least one plausible candidate for the role of antecedent of the pronoun. In other cases, the non-linguistic information provided by context might be robust enough to suspend both requirements, as shown by (43) (adapted from Miller & Pullum, 2013).

43) [*Context: It's dark outside. Fred and Sue are about to leave the house for a bike ride. Sue hands Fred a cycling vest with reflective stripes, clearly implicating that Fred should wear it. Fred stares at Sue with an upset look and says*:]
Do I really have to [*wear the cycling vest*]?

Even (44) (inspired by Gaskin & Hill, 2013) is acceptable if the interpreter is fluent in Pig Latin and can easily reverse the procedure responsible for the derivation of the words in the quotation from standard English.[6]

44) Dustin said "Iyay amyay oinggay otay alktay otay ayay inguistlay", and he did [t*alk to a linguist*].

6 That is, if the English word starts with a consonant, take the first consonant or consonant cluster, move it to the end, and add "ay"; if the English word begins with a vowel sound or a silent letter, just add "yay" to the end. The source sentence of the quotation of (44) is "I am going to talk to a linguist".

The viability of the Contextual Antecedent Hypothesis is thereby not affected by the VPC-based objection. Let me now provide two final clarifications. The first clarification has to do with how the Contextual Antecedent Hypothesis is supposed to take care of examples like (45) and (46).

45) Mary heard Sue whisper "Laura got married yesterday", and wondered whether he was a doctor.

46) "The renovation of Building 38 will be completed by the end of October", said the sign. However, the windows were still missing and the walls not yet painted.

According to the Contextual Antecedent Hypothesis, the quotations of (45) and (46) give contextual prominence to the respective quoted sentences. However, the salience of the quoted sentences per se is insufficient to generate the results of having the pronoun of (45) pick out the individual married by Laura and of having the bridging anaphors of (46) be semantically related via the appropriate parthood relation to "Building 38". How come that the examples are acceptable, then? Saka (1998) helps us here. The examples are felicitous because as soon as the relevant sentence has gained contextual prominence, other things become salient in the context, so to speak, for free, with no need of additional raising-to-salience operations: the graphic features of the sentence, its phonological form, its division in lexical units, its surface syntax, and, crucially, its semantic properties. (45) is thus explained in line with the proposed account by positing that as soon as the sentence "Laura got married yesterday" is raised to salience by the closed quotation, the entailments of "Laura got married yesterday" gain contextual accessibility as well. Since the hearer can assume that getting married means taking part to a ceremony involving two partners, the proposition that some individual married Laura becomes contextually accessible, and the pronoun "he" is interpreted as conveying information about that individual. In Saka's terminology, the closed quotation, by "directly" giving contextual prominence to the sentence "Laura got married

yesterday", "deferringly" gives contextual prominence to the proposition that some individual married Laura, and the contextual accessibility of this proposition licenses the use of "he" after the quotation.

The second clarification concerns the contrast between (47) and (48).

47) [?] We wondered whether Andrew was hallucinating when he said "A woman has walked into the bar". She was ordering a drink.

48) [ok] We wondered whether Andrew was hallucinating when he said "A woman has walked into the bar". But he wasn't. She was ordering a drink.

It could be objected that the Contextual Antecedent Hypothesis does not have the resources required to predict the difference in degree of acceptability between (47) and (48) and, therefore, that it overgenerates. More precisely, if the correct explanation for the viability of the anaphoric pattern exhibited by (48) is that the quotation gives contextual prominence to the quoted sentence and the pronoun is anaphoric on the indefinite qua constituent of the sentence so raised to salience, why is it that (48) is fine and (47) sounds odd?

Here I would like to argue that the contrast between (47) and (48) is best construed, rather than as symptomatic of a deficiency in the proposed hypothesis, as indicating that the recruitment of the linguistic material made contextually accessible as per the Contextual Antecedent Hypothesis responds to additional, general-purpose focus constraints that filter how the prominent sentence can be accessed for anaphora-making (cf. De Brabanter, 2010). The reason why (47) and (48) differ in acceptability has nothing to do with the Contextual Antecedent Hypothesis per se (in other words, nothing to do with the idea that the quotations of (47) and (48) give prominence on the contextual blackboard to the sentence "A woman has walked into the bar"). Rather, the point is that only (48), thanks to the mediation of "But he wasn't [*hallucinating*]", allows for a smooth transition in focus

from the exact wording of Andrew's utterance to the issue of its factual adequacy and, in turn, to a stage of processing where the content of the quotation, already contextually prominent, is evaluated "linguistically enough" for the indefinite to supply an antecedent for "she". To further illustrate the point, witness the difference between (50) and (52), which are felicitous, and (49) and (51), which sound degraded. For example, (49) is odd because its first clause leads the interpreter to expect a follow-up on the question of the referentiality of an NP, a question that only (50) is able to move to the background in a way consistent with the preservation of discourse coherence.[7]

49) [?] Maggie wonders whether the word 'Neptune' refers. It's a huge, blue planet.

50) [ok] Maggie wonders whether the word 'Neptune' refers. It does. It's a huge, blue planet.

51) [??] Kim wonders whether "Tom is great at chess" is a fair thing to say. He sure is.

52) [ok] Kim wonders whether "Tom is great at chess" is a fair thing to say. It is. He sure is.

Similarly, the quality of sentences containing variants of the pattern found in (28) appears to be sensitive to specific properties of the clause used to introduce the quotation. In particular, while the clauses used to introduce the quotation in (28) and (55) place focus on the whole quoted sentence, (53), (54), and (56) appear less felicitous because the clauses introducing the quotation place focus only on some specific properties of the quoted sentence (phonetic and phonological properties in (53) and (54), syntactic properties in (56)).[8]

28) [ok] Mary meditated on the sentence "Bill is a good friend" and concluded that he was a good friend.

7 Examples (49)-(52) are indebted to Smit & Steglich-Petersen (2010) and to conversation with Jack Spencer.

8 Examples (53) and (54) are indebted to conversation with Salvador Mascarenhas.

53) ? Mary thought about the sequence of sounds “Bill is a good friend” and concluded that he was a good friend.

54) ?? Mary counted the number of vowels in the sentence “Bill is a good friend” and concluded that he was a good friend.

55) ok “A woman has walked into the bar”, whispered John. As it turns out, she was ordering a drink.

56) ?? Consider the use of indefinites in the sentence “A woman has walked into the bar”. As it turns out, she was ordering a drink.

Once again, the correct explanation for this pattern seems to be that it does not suffice for an expression to be contextually accessible in order to be recruited to play the antecedent-role: in order to be recruited to play the antecedent-role, the expression must also be in focus. Just as in (42) the use of VP-ellipsis appears felicitous only if the sentence is uttered in a context where Emma and Alex are actually reading the post-it (rather than in a context where the post-it is somewhere in the vicinity of the two interlocutors, or in a context where Emma and Alex are counting how many token-letters it contains while paying no attention to its meaning), so in (28) and (53)-(56) the contextual accessibility of the quoted linguistic material needs to be supplemented by the appropriate kind of focus to issue in the provision of an antecedent for the pronoun. However, it seems fair to say that the existence of this constellation of additional pragmatic effects, whatever their ultimate nature and workings, complements rather than affecting the overall viability of the proposed account.

7. CONCLUSION

The discussion has proceeded as follows. Section 1 introduced the interplay between pronominal anaphora on referential antecedents and coreference, and formulated R-Inheritance. Section 2

offered some preliminary considerations about the viability of R-Inheritance, highlighted some of its immediate controversial features, and called for a stronger case against the principle. Section 3 presented (28), formulated the Argument from Closed Quotation, and illustrated the reasons why it threatened R-Inheritance. Section 4 described two competing ways of neutralizing the Argument from Closed Quotation. One, the Deixis Hypothesis, based on the notion that the pronoun of (28) is a demonstrative particle rather than an anaphor. One, the Contextual Antecedent Hypothesis, based on the notion that the pronoun of (28) is anaphoric on the constituent of a sentence given contextual prominence by the closed quotation. Section 5 discussed and argued against the Deixis Hypothesis. Section 6 discussed and defended the Contextual Antecedent Hypothesis.

Along the way, I have focused on two main tasks. First, I have attempted to show that sentences exhibiting the anaphoric pattern of (28) cannot be used to adjudicate the viability of R-Inheritance, contra the Argument from Closed Quotation. Second, I have proposed that the anaphoric pattern of (28) can be accounted for under the hypothesis that closed quotations raise to salience the linguistic material their contain, and that the linguistic material so raised to salience provides a contextually supplied antecedent for the pronoun. I hope that the analysis presented in this paper will help shed new light and raise new questions on the complex interplay between quotation, anaphora, coreference, and neighboring linguistic phenomena.

ACKNOWLEDGMENTS

I am indebted to Claire Beyssade, Philippe De Brabanter, Paul Egré, Irene Heim, Mikhail Kissine, Salvador Mascarenhas, Michael Murez, François Recanati, Jack Spencer, and Josef Stern for discussion and/or feedback on earlier versions of this paper, and to three anonymous reviewers for detailed comments on the submitted manuscript. All errors are mine. The research that led to this paper was supported by grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL*.

References

Asher, N. (2011). *Lexical meaning in context: A web of words*. Cambridge: Cambridge University Press.

Cappelen, H., & Lepore, E. (2012). Quotation. In *The Stanford Encyclopedia of Philosophy (Spring 2012 Edn.)*, ed. E. N. Zalta. http://plato.stanford.edu/archives/spr2012/entries/quotation/

Chomsky, N. (2000). *New horizons in the study of language and mind*. Cambridge: Cambridge University Press.

Clark, H. H., & Gerrig, R. J. (1990). Quotations as demonstrations. *Language*, 66, 764–805.

Davidson, D. (1984). Quotation. In his *Inquiries into truth and interpretation*, 79–92. Oxford: Oxford University Press.

De Brabanter, P. (2010). Constraints on metalinguistic anaphora. In *Constraints in discourse 2*, ed. P. Kühnlein, A. Benz, & C. L. Sidner, 141–162. Amsterdam: Benjamins.

De Vries, M. (2008). The representation of language within language: A syntactico-pragmatic typology of direct speech. *Studia Linguistica*, 62, 39–77.

Elbourne, P. (2005). *Situations and individuals*. Cambridge, MA: MIT Press.

Gaskin, R., & Hill, D. J. (2013). Reach's puzzle and mention. *dialectica*, 67, 201–222.

Günthner, S. (1999). Polyphony and the 'layering of voices' in reported dialogues: An analysis of the use of prosodic devices in everyday reported speech. *Journal of Pragmatics*, 31, 685–708.

Hankamer, J., & Sag, I. (1976). Deep and surface anaphora. *Linguistic Inquiry*, 7, 391–428.

Heim, I. (1982). *The semantics of definite and indefinite noun phrases*. Dissertation, University of Massachusetts, Amherst.

Heim, I. (1990). E-type pronouns and donkey anaphora. *Linguistics and Philosophy*, 13, 137–178.

Heim, I. (1998). Anaphora and semantic interpretation: a reinterpretation of Reinhart's approach. *MIT Working Papers in Linguistics*, 25, 205–246.

Kamp, H. (1981). A theory of truth and semantic representation. In *Formal methods in the study of language*, ed. A. G. Groenendijk, T. M. V. Janssen, & M. B. J. Stokhof, 277–322. Amsterdam: Mathematical Center Tract 135.

Kaplan, D. (1989). Demonstratives. In *Themes from Kaplan*, ed. J. Almog, J. Perry, & H. Wettstein, 481–563. New York, NY: Oxford University Press.

King, J. C., & Lewis, K. S. (2016). Anaphora. In *The Stanford Encyclopedia of Philosophy (Summer 2016 Edn.)*, ed. E. N. Zalta. https://plato.stanford.edu/archives/sum2016/entries/anaphora/

Maier, E. (2014). Mixed quotation: The grammar of apparently transparent opacity. *Semantics & Pragmatics* 7, Article 7, 1–67.

Maier, E. (2016). A plea against monster. *Grazer Philosophische Studien*, 93, 363–395.

Merchant, J. (2010). Three kinds of ellipsis: Syntactic, semantic, pragmatic? In *Context-dependence, perspective, and relativity*, ed. F. Recanati, I. Stojanovic, & N. Villanueva, 141–192. Berlin: de Gruyter.

Miller, P., & Pullum, G. K. (2013). Exophoric VP ellipsis. In *The core and the periphery: Data-driven perspectives on syntax inspired by Ivan A. Sag*, ed. P. Hofmeister & E. Norcliffe, 5–32. Stanford, CA: CSLI Publications.

Neale, S. (2005). Pragmatism and binding. In *Semantics versus pragmatics*, ed. Z. G. Szabó, 165–285. Oxford: Oxford University Press.

Pagin, P., & Westerståhl, D. (2010). Pure quotation and general compositionality. *Linguistics and Philosophy*, 33, 381–415.

Partee, B. (1973). The syntax and semantics of quotation. In *A Festschrift for Morris Halle*, ed. S. R. Anderson & P. Kiparsky, 410–418. New York, NY: Holt, Reinhart and Winston.

Potts, C. (2007). The dimensions of quotation. In *Direct Compositionality*, ed. C. Barker & P. I. Jacobson, 405–431. Oxford: Oxford University Press.

Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge, MA: MIT Press.

Recanati, F. (2001). Open quotation. *Mind*, 110, 638–687.

Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.

Recanati, F. (2018). Fictional, metafictional, parafictional. *Proceedings of the Aristotelian Society*, 118, 25–54.

Reinhart, T. (1983). *Anaphora and semantic interpretation*. Chicago, IL: University of Chicago Press.

Reuland, E. (2011). *Anaphora and language design*. Cambridge, MA: MIT Press.

Saka, P. (1998). Quotation and the use-mention distinction. *Mind*, 107, 113–135.

Safir, K. (2005). Abandoning coreference. In *Thought, reference, and experience: Themes from the philosophy of Gareth Evans*, ed. J. L. Bermudez, 124–163. Oxford: Oxford University Press.

Schlenker, P. (2005). Non-redundancy: Towards a semantic reinterpretation of binding theory. *Natural Language Semantics*, 13, 1–92.

Smit, J. P., & Steglich-Petersen, A. (2010). Anaphora and semantic innocence. *Journal of Semantics*, 27, 119–124.

Stone, M., & Hardt, D. (1999). Dynamic discourse referents for tense and modals. In *Computing Meaning*, ed. H. Bunt & R. Muskens, 302–321. Dordrecht: Kluwer.

Washington, C. (1992). The identity theory of quotation. *Journal of Philosophy*, 89, 582–605.