

This is an EARLY ACCESS version of

Kirk-Giannini, Cameron Domenico. 2024. Covert mixed quotation. *Semantics and Pragmatics* 17(5). <https://doi.org/10.3765/sp.17.5>.

This version will be replaced with the final typeset version in due course.
Note that page numbers will change, so cite with caution.

Covert mixed quotation*

Cameron Domenico Kirk-Giannini
Rutgers University–Newark

Abstract The term *covert mixed quotation* describes cases in which linguistic material is interpreted in the manner of mixed quotation — that is, used in addition to being mentioned — despite the superficial absence of any commonly recognized conventional devices indicating quotation. After developing a novel theory of mixed quotation, I show that positing covert mixed quotation allows us to give simple and unified treatments of a number of puzzling semantic phenomena, including the projective behavior of conventional implicature items embedded in indirect speech reports and propositional attitude ascriptions, so-called ‘c-monsters,’ metalinguistic negation, metalinguistic negotiation, and ‘in a sense’ constructions.

Keywords: mixed quotation, slurs, c-monsters, metalinguistic negation, metalinguistic negotiation, counterfactual conditionals, conventional implicature

Ordinary examples of mixed quotation — the variety of quotation wherein quoted material is both used and mentioned — are marked explicitly using quotative punctuation, as in:

- (1) Quine said that quotation ‘...has a certain anomalous feature’ (Davidson 1979: 28)
- (2) Bush is proud of his ‘eckullectic’ reading list. (Shan 2010: 417)

(1) and (2) differ from examples of pure quotation — the more familiar construction in which a quotation functions syntactically as a determiner phrase

*I am grateful to Chris Barker, Bob Beddor, Kyle Blumberg, Sam Carter, Simon Goldstein, Ben Holguín, Arc Kocurek, Harvey Lederman, Matt Mandelkern, Kate Hazel Stanton, and two anonymous referees for *Semantics and Pragmatics* for their feedback on earlier versions of this material. I also benefitted from discussions with the Virtual Language Work-in-Progress Group (VirLaWP) and the Logic, Language, and Cognition Seminar at the University of Turin. This paper grew out of a version of what is now Section 3, which I wrote under the supervision of Simon Charlow at Rutgers–New Brunswick.

and semantically as a device for referring to the expression quoted — not only in that the quotations they contain do not have the syntactic type of determiner phrases, but also in that the expressions quoted appear to contribute to composition something like the semantic values they had in the mouths of those who originally produced them.

Ordinary examples of pure and mixed quotation are by now reasonably well described.¹ Yet it would be narrowminded in developing an account of quotation — and especially of mixed quotation — to limit our theoretical attention to cases with explicit quotative punctuation. Natural language abounds with examples in which, despite the superficial absence of quotation marks, expressions appear to be neither simply used nor simply mentioned. Indeed, this feature unifies a number of puzzling semantic phenomena which have so far been discussed in isolation from one another. Consider:

- (3) I didn't manage to trap two MONGEES — I managed to trap two MONGOOSES. (Horn 1989: 371)
- (4) Pluto could have easily been a planet. But that one stubborn scientist voted for the current definition, so it is not. (Kocurek et al. 2020: 6)
- (5) My father screamed that he would never allow me to marry that Kraut Webster, who always stank of sauerkraut and Kölsch, or any other damned Kraut. (Kirk-Giannini 2019: 18)
- (6) A: Secretariat is an athlete.
B: No, Secretariat is not an athlete. (Ludlow 2008: 118)
- (7) You are right to say that viruses are alive in a sense, but not usually classified in those 5 kingdoms.²

(3-7) resist treatment as cases involving nothing more than the ordinary use of expressions. If the word 'mongeese' in (3) were used simply to express the property *mongoose*, the resulting conjunction would be a contradiction and therefore unsuitable for assertion. Given that, since the International Astronomical Union's 2006 decision, the word 'planet' has had in its extension only celestial bodies which clear their orbital neighborhoods, interpreting the token of it that occurs in (4) as a normal use would render a perfectly

¹For a survey of work on pure quotation, see Maier (2014b). For recent accounts of mixed quotation, see e.g. Potts (2007), Shan (2010), and Maier (2014a).

²Shannon DeVaney, MadSci Network: General Biology forum post, March 3, 2005. URL: www.madsci.org/posts/archives/2005-03/1109887785.Gb.r.html.

Covert mixed quotation

reasonable discourse incoherent: no stubborn scientist could have changed whether Pluto does or does not clear its orbital neighborhood. Analyzing (5) as involving nothing more than indirect quotation — and thus no genuine mention of the expressions embedded under ‘screamed’ — would give us no explanation of why the pejorative content of the slur and the propositional content of the non-restrictive relative clause do not project out of this embedding environment, so that the speaker is understood to endorse her father’s unsavory perspective on Webster. A normal-use account provides no explanation of the possibility that the interlocutors in (6) might be disagreeing about whether Secretariat *should be* classified as an athlete rather than about whether Secretariat *is* classified as an athlete. And interpreting (7) as a case of ordinary use leaves us nothing in its logical form for the existential quantifier ‘in a sense’ to bind, and thus no hope for a satisfactory analysis.

At the same time, each of (3-7) has something of the flavor of quotation. (3) and (4) contain expressions or clauses which seem, in some sense, to characterize how language is or might be used. (5) seems to involve the attribution of a certain way of speaking to someone other than the speaker. The interlocutors in (6) seem in the first instance to be disagreeing not about what is or ought to be the case regarding Secretariat, but rather about what is or ought to be the case regarding the world ‘athlete.’ And (7) seems to involve quantification over ways in which the expression ‘alive’ could be interpreted. Yet if the expressions in these examples were mentioned in the manner of pure quotation, the sentences containing them would be defective both syntactically and semantically.

To be clear, I do not wish to suggest that (3-7) are *obviously* quotative. One could treat them theoretically in a piecemeal fashion, positing quotation in only some, or perhaps none — and indeed this is the approach which has so far been adopted in the literature. But a unified treatment is preferable to a piecemeal one, and the suggestion I develop in what follows is that, superficial appearances notwithstanding, the quotative flavor of (3-7) is best explained by positing quotation: all five phenomena involve mixed-quoted material. Since the sentences in question contain no overt elements indicating that quotation is in play, the mixed quotation I posit is covert. I understand covert mixed quotation as a semantic phenomenon involving a syntactically present but phonologically unrealized quotation operator. In this respect the view I develop in what follows departs from those of most others who have taken covert quotation seriously, since they have generally understood it to be a pragmatic phenomenon rather than a semantic one.

It will emerge that there must be more to (3-7) than covert mixed quotation. I suggest that, in addition to covert mixed quotation, (3) and (6)

involve other covert material contributing the content that it is *appropriate* to describe mongooses using the expression ‘mongeese’ or Secretariat using the expression ‘athlete,’ (4) involves covert material which performs a function structurally similar to Stalnakerian diagonalization, and (7) involves the interaction between a covert quotation operator and the overt quantificational phrase ‘in a sense’. Nevertheless, in each case, my analysis appeals essentially to the presence of covert mixed quotation: the extra material I posit could not, by itself, account for the phenomena in question. Where this extra material does appear, moreover, it is always independently motivated either by existing work on mixed quotation (as in the diagonalizer required to analyze (4)) or by the nature of the phenomenon in question (as in the appeal to appropriateness to explain (3) and (6) and the quantificational structure required to explain (7)). For this reason, I regard all five examples as manifestations of the same underlying semantic phenomenon: the presence of mixed quotation in natural language in the absence of quotative punctuation.

Two notes on the relationship between my proposal and existing work. First, I am not the first to appreciate the potential explanatory power of covert mixed quotation. Potts (2007), Maier (2015), and Stokke (2021) appeal to covert mixed quotation to offer accounts of metalinguistic negation, free indirect discourse, and protagonist projection and character focus, respectively. My goal in what follows is rather to build on this earlier work to motivate and significantly extend the empirical coverage of the hypothesis of covert mixed quotation.

Second, I am not the first to offer accounts of (3-6). (3) is an example of metalinguistic negation, a phenomenon which has generated an extensive literature and received a variety of theoretical treatments, including, as just mentioned, one (Potts 2007) which appeals to mixed quotation. Kocurek et al. (2020) account for (4) by positing the existence of ‘c-monsters’ — operators which shift the linguistic conventions relative to which embedded material is evaluated. Harris & Potts (2009) suggest a non-quotative explanation for cases like (5). And Plunkett & Sundell (2013) have taken (6) and related cases as evidence that disagreement does not require the expression of incompatible semantic contents.

To my knowledge, however, the account developed below is the first to treat (4-6) as involving mixed quotation, the first account of any kind of sentences like (7), the first empirically adequate account of (3) involving mixed quotation, and the first account to subsume (3-7) into a unified

semantic category.³ Because it engages with a range of topics which have not traditionally been regarded as quotative, I hope the account will be of interest to an audience in linguistics and the philosophy of language beyond those whose research already concerns the semantics of quotation.

Section 1 below outlines the evidence for the existence of covert mixed quotation and sketches some considerations in favor of treating it as a semantic phenomenon. Section 2 introduces a theory of covert mixed quotation and shows how it accounts for some of the main examples of covert quotation discussed in Section 1. The theory is novel in that it explains the semantic properties of mixed quotation in terms of a compositional interaction between ordinary pure quotation and adjacent covert material. Sections 3, 4, 5, 6, and 7 extend the basic theory to account for the projective behavior of conventional implicature items, ‘c-monsters,’ metalinguistic negation and related phenomena, metalinguistic negotiation, and ‘in a sense’ constructions, respectively. Section 8 responds to an objection to the thesis of covert mixed quotation and compares my treatments of the phenomena in Sections 3 and 5 with some alternatives in the literature. Section 9 concludes. The formal theory developed in the paper is presented in detail in the appendix.

1 Evidence for (a semantic treatment of) covert mixed quotation

The existence of covert mixed quotation is accepted by most linguists and philosophers, and my primary focus in what follows is to show that covert mixed quotation can form the basis of an analysis of (3-7) rather than to respond to any remaining skeptics.⁴ Nevertheless, it is worth beginning by briefly summarizing the evidence for covert mixed quotation and the motivation for a semantic treatment thereof.

In spoken English, the presence of quotation marks need not be marked prosodically. The word “Boston” in “Boston’ has six letters’ may be pronounced in exactly the same way as the word ‘Boston’ in ‘Boston is a city’ without loss of intelligibility. Our interpretive faculties deal handily with articulations which might or might not be quotative.

Saka (1998: 118) argues that this lesson applies equally to written language:

³For discussion of the empirical issues facing Potts’s account of metalinguistic negation, see Section 8.

⁴Shan (2010: 440) even goes so far as to suggest that “mixed quotation subsumes all of language except coinage.” This is not, however, a widely endorsed position.

“...it is downright *normal*, outside of scholarly writing, to exclude quote marks, especially in constructions like ‘The word cats is a noun’...” (emphasis in original)

These seem to be examples of covert pure quotation. But there are also apparent examples of covert mixed quotation.

[Recanati](#), for example, cites the following sentence from Dickens’s *Martin Chuzzlewit*:

- (8) To which Mr Bailey modestly replied that he hoped he knowed wot o’clock it wos in gineral. (2001: 675).

Here the material “knowed wot o’clock it wos in gineral” appears to characterize the words Mr. Bailey used in the course of his reply as well as the content of that reply. In other words, (8) appears to be semantically equivalent to:

- (8’) To which Mr Bailey modestly replied that he hoped he “knowed wot o’clock it wos in gineral.”

Further evidence for covert mixed quotation can be found in Relevance Theorists’ discussions of the phenomenon they call “echoic use.” For example, [Carston \(2002\)](#) draws attention to examples like:

- (9) Americans eat tom[eIrouz] and Brits eat tom[a:touz].
(Carston 2002: 299)

(9) contains no overt quotation marks, but succeeds in conveying something metalinguistic: that the American and British pronunciations are as the speaker indicates.

There are, broadly, three sorts of responses available to the observation that covert mixed quotation appears in spoken and written language. The first is to attempt a unified treatment of overt and covert mixed quotation by taking covert mixed quotation to indicate that mixed quotation as a phenomenon does not depend on the presence of quotation marks at any level of linguistic description. According to this strategy, metalinguistic reference is achieved pragmatically or speech-act theoretically, and overt quotation marks serve only as convenient but dispensable guides indicating that a speaker intends to use some material metalinguistically.

The second strategy, and the one I explore here, is to attempt a unified treatment of overt and covert mixed quotation by positing phonologically

Covert mixed quotation

unrealized quotation operators in all cases of covert quotation. On this strategy, both overt mixed quotation and covert mixed quotation are genuine semantic phenomena.

The third strategy is to resist a unified treatment: overt quotation marks play a semantic role, but covert cases are to be explained pragmatically. Cappelen & Lepore (2007), for example, adopt this strategy when it comes to covert pure quotation. Following García-Carpintero (1994), they explain some cases of covert pure quotation in terms of syntactically elided quotation marks and others in terms of pragmatic reinterpretation.

I think there are clear methodological reasons to prefer a unified treatment of overt and covert quotation to a patchwork one. The question, then, is whether the unified pragmatic strategy is viable. It is beyond the scope of this article to refute the unified pragmatic strategy, but I will briefly discuss a consideration which I think counts in favor of a semantic approach.

Mixed-quoted material seems able to interact with other material in a sentence to generate bound-type readings. Potts (2007) gives us the following two sentences involving overt mixed quotation, arguing that they suggest a semantic account of mixed quotation:

- (10) When in Santa Cruz, Peter orders ‘[eI]pricots’ at the local market.
- (11) When in Amherst, Peter orders ‘[æ]pricots’ at the local market.

Here it seems that the implication that Peter uses the pronunciation indicated by the quoted material is bound by the event quantifiers ‘when in Santa Cruz’ and ‘when in Amherst’.

Similarly, we have:

- (12) If Peter ordered ‘[eI]pricots’ at the local market, he was in Santa Cruz.

(12) has a reading on which its antecedent requires that Peter ordered apricots by uttering ‘[eI]pricots’. Thus there are situations in which (12) might be true even though it is false that *if Peter ordered apricots at the local market, he was in Santa Cruz*. These are situations in which the closest apricot-ordering worlds are not Santa Cruz worlds, but the closest apricot-ordering and ‘[eI]pricots’-uttering worlds are Santa Cruz worlds.

We can easily generate similar cases involving covert quotation. Thus

- (13) When in Santa Cruz, Peter orders eggplants at the local market; when in London, however, he orders aubergines at the local market.
- (14) If Peter ordered aubergines at the local market, he was in London.
- (15) Oh, now I remember that the first floor is the ground floor in London. (Noh 2000: 141)

In (13), the contrastive ‘however’ suggests a reading on which ‘eggplants’ and ‘aubergines’ characterize not only what produce Peter orders but also the lexical items with which he orders it, and we have the same binding relationship between this characterization and the event quantifiers. Similarly, (14) has a reading on which its antecedent takes us not to the closest world in which Peter ordered eggplants at the local market, but rather to the closest world at which he ordered eggplants by calling them aubergines. In (15), the modifier ‘in London’ serves to specify where the first floor is characterized using the expression ‘the ground floor.’

Naturally, I do not intend these remarks to be a decisive refutation of a unified pragmatic account of mixed quotation. There are various moves which might be made in defense of the pragmatic strategy, including perhaps appealing to optional pragmatic strengthening or “narrowing” of the meanings of lexical items (see, for example, Geurts (2010)). Without a clearer explanation of how this kind of pragmatic process might work in the case of mixed quotation or what might drive its application, however, it is difficult to assess the merits of such a proposal.

In any case, the existence of sentences like (13)–(15) motivates developing the semantic approach in order to better appreciate its explanatory potential. Moreover, if (as I argue below) a semantic account of covert mixed quotation can be extended to provide a unified treatment of a range of seemingly unrelated semantic phenomena, and if pragmatic accounts cannot so easily be extended to cover the same range of cases, this is additional evidence for a semantic account.

2 A theory of mixed quotation

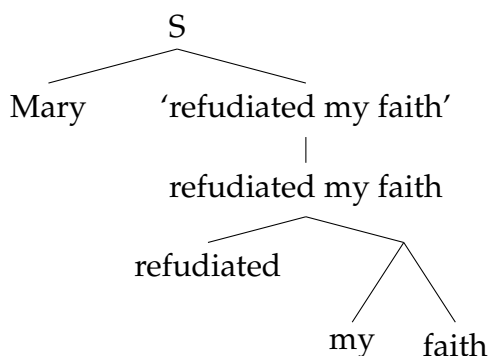
In this section, I outline a theory of overt and covert mixed quotation. The theory I present draws on the work of Chris Potts (2005, 2007), Chung-chieh Shan (2010), and Emar Maier (2014a), though it ultimately departs from each of them. From Potts, I take the idea that the mixed quotation operator takes some text and contributes to a peripheral dimension of composition

Covert mixed quotation

that some actual or hypothetical person uttered it verbatim. From Maier, I take the idea that the content a mixed-quoted expression contributes to at-issue composition is whatever it contributed on this actual or hypothetical occasion of utterance. Drawing on Shan, I implement this proposal about the at-issue contribution of mixed-quoted items using a purpose-built quotative interpretation function (*). The theory I sketch here differs from other proposals, however, both in its details and, more generally, in understanding mixed quotation as arising compositionally from the interaction between pure quotation and nearby covert material. My focus in this section is to present the theory in as informal a way as possible without glossing over essential details. A complete formal exposition can be found in the appendix.

Existing characterizations of mixed quotation have generally taken it to be syntactically no more complex than pure quotation. Thus Maier (2014a), to take one recent example, holds that there is a mixed quotation operator which can be inserted into the syntactic tree representing a sentence above any node, creating a new node of the same syntactic category. On this sort of proposal, sentence (16) below is assigned the following structure:

(16) Mary ‘refudiated my faith.’ (Maier 2014a)



Accounts of this sort have a ready explanation for what Cappelen & Lepore (2007) evocatively call the ‘syntactic chameleonism’ of mixed quotation: the fact that mixed quotations inherit their syntactic properties from the material quoted. Despite their simplicity, however, such accounts have the disadvantage that they posit two unrelated quotative mechanisms in natural language — pure quotation and mixed quotation — without providing any underlying story about how they are related. In contrast, the account I favor modifies the syntax slightly in a way which allows us to understand the device of mixed quotation as arising compositionally from the interaction between pure quotation and surrounding covert material. As we will see, this

more developed syntactic proposal is also useful in giving an explanation of sentences like (4) and (5) in terms of covert mixed quotation.

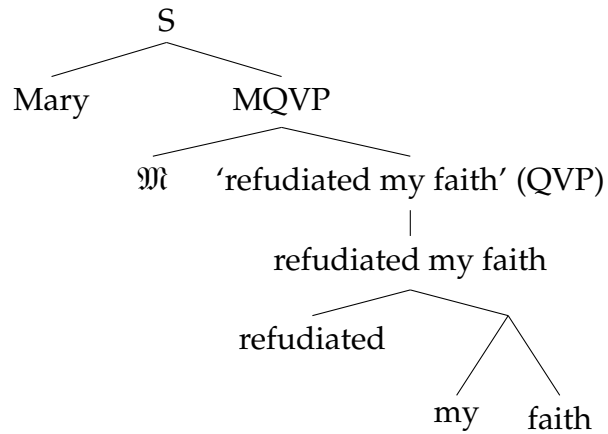
Since my proposal treats mixed quotations as arising from an interaction between pure quotations and adjacent covert material, it is no longer possible to have mixed-quoted expressions directly inherit their types from the nodes they dominate, as on Maier's proposal. Instead, I analyze each mixed quotation as having two daughter nodes, one a pure quotation and the other the (phonologically unrealized) mixed quotation operator \mathfrak{M} .

The connection between the syntactic type of a mixed quotation and the syntactic type of the material quoted, which comes for free on Maier's proposal, is enforced on my account by introducing a set of pure-quotational syntactic types corresponding to the more familiar non-quoted types. Thus, in addition to determiner phrases (DP), verb phrases (VP), and so forth, I admit pure-quoted determiner phrases (QDP), pure-quoted verb phrases (QVP), and so forth. The result of applying the pure quotation operation to a DP is a QDP, and this holds generally for the set of non-quoted syntactic types. Pure-quoted expressions, regardless of their syntactic type, function syntactically like DPs.⁵ I then introduce a family of syntactic types for mixed-quoted expressions in addition to the types for pure-quoted expressions. Since on my view a mixed quotation is formed from the interaction between a pure quotation and adjoining covert material, it is only Q-type syntactic expressions which can be mixed-quoted. We thus have types MQDP, MQVP, and so forth. An expression of type MQDP behaves syntactically like a DP, an expression of type MQVP behaves syntactically like a VP, and so forth. Thus, though all pure-quoted expressions regardless of type behave syntactically much like DPs, the syntax does not "forget" the original types of pure-quoted expressions, and these are then inherited indirectly by mixed quotations.⁶

⁵This proposal is inspired by Cappelen & Lepore (2007), who also introduce a set of quotational syntactic types corresponding to the non-quoted types. I depart from Cappelen & Lepore, however, in distinguishing between pure-quotational types and mixed-quotational types.

⁶Admittedly, the syntax I propose here is complicated, especially as compared to Maier's alternative proposal. I offer three considerations in its defense, in increasing order of importance. First, the syntactic complexity is driven by the fact that the theory treats mixed quotation as arising from a compositional interaction between pure quotation and nearby covert material. Second, the proposal allows us to explain the phenomenon of covert mixed quotation in terms of the phenomenon of covert pure quotation. Third, we will see below that the fact that mixed-quoted expressions are first pure-quoted, as well as the more articulated syntax, play important explanatory roles in some of the applications of the theory. For example, it is the idea that mixed-quoted expressions are first pure-quoted that explains why the peripheral content of covertly mixed-quoted items does not project in Section 3. Similarly, the fact that on my proposal the mixed quotation operator \mathfrak{M} is an independent

An example will clarify the proposal. According to the view I propose, (16) is analyzed syntactically as:⁷



Here the mixed-quoted verb phrase ‘refudiated my faith’ arises from the interaction between the pure-quoted verb phrase ‘refudiated my faith’ and the mixed quotation operator \mathfrak{M} . \mathfrak{M} heads the MQVP in which it is contained. Note again that \mathfrak{M} has no overt realization on this proposal. Overt mixed quotations are overt because the pure quotations they contain are overt; covert mixed quotations are covert because the pure quotations they contain are covert. Thus the existence of covert mixed quotation is predicted to follow directly from the existence of covert pure quotation.

Turning next to semantics, there is widespread agreement among those working on mixed quotation that a mixed-quoted expression simultaneously contributes some at-issue content to semantic composition — this is usually held to be something like whatever the quoted material meant as used in the instance being quoted — and the not-at-issue proposition that (roughly) the quoted material was uttered verbatim by some salient individual. Different accounts of mixed quotation differ with respect to the precise nature of the contents they hold to be contributed, as well as with respect to how the not-at-issueness of the second proposition is to be understood. While many understand the not-at-issueness in question in terms of presupposition,

element in the syntactic tree allows it to be modified by further operators, which is a crucial part of the explanation of alleged c-monsters in Section 4.

⁷I suppress the fact that the lexical item ‘refudiated’ is arguably not in the syntax of the language of the speaker of (16). To handle cases like this, it is necessary to allow the pure quotation operation to operate on subtrees of syntactic languages other than that of the speaker. See [Shan \(2010\)](#) and [Maier \(2014a\)](#) for discussion. My account also sets aside the problem of ‘non-constituent quotation’; for discussion, see [Maier \(2014a\)](#).

I follow Potts (2007) in handling the not-at-issue contribution of mixed quotation within a multidimensional semantic framework based on the one introduced by Potts (2005) to handle conventional implicatures.

Examples like (10), (11), and (12) make it clear that certain operators can take as argument not only the normal at-issue content of a mixed-quoted expression but also its peripheral content. For example, recall that the intuitive truth conditions for (12) are such that the antecedent of the conditional takes us not to the nearest world where Peter orders apricots at the local market, but rather to the nearest world where Peter orders apricots at the local market *by calling them* ‘[e]l[pricots’]. It is straightforward to account for this sort of behavior in a multidimensional framework in terms of “shunting” between the at-issue and peripheral semantic dimensions (see, for example, McCready 2010), and for this reason I have developed my account within that framework.

Note, however, that I do not claim that the not-at-issue content of a mixed quotation is conventionally implicated. Conventional implicatures are taken to have a number of features, such as anti-backgrounding, which do not fit well with the data regarding mixed quotation.⁸ A complete semantic account of natural language will likely need to posit multiple peripheral dimensions of meaning and explain how, if ever, they interact. For simplicity, I make do in what follows with a single peripheral dimension.

Note also that my choice to work in a multidimensional framework in the tradition of Potts may ultimately be no more than a matter of theoretical taste. It is possible that the analyses I give below could be reproduced in a presuppositional framework like DRT. Maier (2014a), to take one example, develops a presuppositional account of mixed quotation within a DRT framework. My goal is to provide *a* formal treatment of the target phenomena in terms of covert mixed quotation, not the only such treatment.⁹

It is commonly observed that the implication that mixed-quoted material has been uttered verbatim by some salient individual should be understood in terms of discourse anaphora rather than an existential quantifier. For this reason, I take the peripheral content a mixed-quoted expression contributes to introduce two free variables, u_x and s_x , taking utterances and individuals as their values, respectively. These are understood to function as discourse anaphors: mixed-quoting ‘refudiated my faith’, as in (16), is felicitous only if there is an identifiable utterance u_x of the quoted material by speaker s_x which is anaphorically retrievable. On the present account,

⁸See Potts (2005) for discussion.

⁹Thanks to an anonymous referee for pressing me to clarify this point.

Covert mixed quotation

the full peripheral content introduced by ‘refudiated my faith’ is then that s_x produced utterance u_x of ‘refudiated my faith’, or, if we let R be a three-place predicate expressing this relation between speaker, utterance, and expression, $R(s_x, u_x, \text{‘refudiated my faith’})$. If we let the variable q range over expressions, we can derive this peripheral content compositionally by associating the mixed quotation operator \mathfrak{M} with the function $\lambda q.R(s_x, u_x, q)$.¹⁰ This function is saturated by the pure quotation which is its syntactic sister, yielding the proposition that s_x produced utterance u_x of q .

It is worth noting that this account of the peripheral content of mixed quotations glosses over two issues (neither of which will be central in what follows). First, there need not be a unique anaphorically retrievable utterance or speaker.¹¹ A mixed quotation can pick up on generic uses of language in a community, as in:

- (17) Young writers often make ‘noticable’ typographical mistakes.

Second, the utterance picked out by a mixed quotation may be merely hypothetical, as in:

- (18) The next person who claims to have a ‘pacific’ point rather than a ‘specific’ point will be ridiculed.¹²

Generic and hypothetical uses of mixed quotation can be combined. For example, imagine that a mean-spirited British linguistic prescriptivist to utters the following sentence to an American child who has just offered her a tomato:

- (19) That isn’t a ‘tomaydo,’ or indeed a ‘tomootoo’ or a ‘tomaitai.’ That is a ‘*tomahto*’ — end of story.

Here, while the source for the first quoted item is a generic American speaker and the source for the last is a generic British speaker, the sources for the two others are hypothetical generic speakers of imaginary comical English dialects.

¹⁰Strictly speaking, lambda terms like this one are expressions in the formal system’s intermediate logical language, not functions. When interpreted by the interpretation function $\llbracket * \rrbracket^{c,w,g}$, however, they denote functions, and for this reason I will sometimes refer to them as functions in what follows.

¹¹See Geurts & Maier (2003: 120-1) for discussion.

¹²Thanks to Sam Carter for suggesting this example.

A complete theory of mixed quotation would need to account for these complexities. Though I will write informally of generic and hypothetical utterances in what follows, however, in the formalism I will suppress these issues, treating all utterances as if they are non-generic and non-hypothetical.

Drawing on Shan (2010), I present the at-issue contribution of a mixed-quoted item in terms of a (curried) interpretation function $\llbracket * \rrbracket$, which takes an expression, a world of utterance w_1 , and a speaker or linguistic community, and, relative to a context c and a point of evaluation w_2, g , returns the at-issue meaning (extension) at w_2 of that expression as uttered by that speaker at w_1 . Thus $\llbracket \llbracket \text{‘refudiated my faith’} \rrbracket (w_1)(\text{Sarah Palin}) \rrbracket^{c, w_2, g}$ is the extension at w_2 of the intension contributed by an utterance of ‘refudiated my faith’ by Sarah Palin in w_1 — which, if w_1 is the actual world, is perhaps something like the extension at w_2 of standard English ‘repudiated my faith.’

I take the at-issue contribution of a mixed-quoted item q (relative to context c and point of evaluation w, g) to be $\llbracket \llbracket q \rrbracket (w_c)(s_x) \rrbracket^{c, w, g}$, with the second argument of $\llbracket * \rrbracket$ occupied by the world of the context w_c and its third argument occupied by the same variable s_x introduced in \mathfrak{M} ’s peripheral content. We can derive this contribution compositionally by associating \mathfrak{M} with the function $\lambda q. \llbracket q \rrbracket (w_c)(s_x)$. This function is saturated by the pure quotation which is its syntactic sister, yielding the meaning of the pure-quoted expression as uttered at the world of the context by the speaker who produced it during the anaphorically retrieved utterance.

We now have two lexical entries associated with the mixed quotation operator \mathfrak{M} , and something must be said about how to combine them within a multidimensional framework. Following Potts (2005), I employ a two-stage semantics. First, the English expressions at the terminal nodes of the syntactic tree representing the logical form of an English sentence are translated into an intermediate logical language \mathcal{L}_{MQ} . A single English lexical item might be translated into more than one expression of \mathcal{L}_{MQ} . Expressions of \mathcal{L}_{MQ} are then directly interpreted by the function $\llbracket * \rrbracket^{c, w, g}$. The semantic well-formedness of the sentence as a whole is governed by a set of *tree admissibility conditions*, which play the role of semantic composition rules.¹³

Following Potts (2005), I use the bullet (\bullet) to indicate that a single node in a syntactic tree is associated with multiple formulas in \mathcal{L}_{MQ} . The fact that \mathfrak{M} is semantically associated with two different functions is thus represented

¹³Unlike Potts, who interprets semantic parsetrees as a whole using a rule which is sensitive not only to the semantic content of their root nodes but also to peripheral contents which may occur lower in the tree without being ‘passed up’ to the root node, I adopt the more traditional approach of identifying the interpretation of a parsetree with the set of semantic values of its root node.

Covert mixed quotation

as: $\mathfrak{M} \rightsquigarrow \lambda q. (\lambda q. (w_c)(s_x)) \bullet \lambda q. R(s_x, u_x, q)$. Both the lambda expression to the left of the \bullet and the lambda expression to its right are genuine meanings of \mathfrak{M} , and both are (independently) interpreted by the model theory. Note that the bullet indicates only that more than one \mathcal{L}_{MQ} expression is associated with a node; it does not signify anything about whether the expressions in question belong to the at-issue dimension of meaning or the peripheral dimension of meaning. The symbol \rightsquigarrow indicates translation from the object language (English) to \mathcal{L}_{MQ} .

Both meanings of \mathfrak{M} take the pure-quoted expression ‘refudiated my faith’ as argument; the first, however, outputs a meaning in the at-issue dimension, while the second outputs a meaning in the peripheral dimension. To explain how this might be so, we must complicate the usual system of semantic types to distinguish between at-issue types (superscripted a) and peripheral types (superscripted p).¹⁴ Then we can say that one meaning of \mathfrak{M} is of type $\langle q^a, \chi^a \rangle$ — that is, that it takes an at-issue expression-denoting expression and returns an at-issue expression of the type appropriate to the quoted material — and the other is of type $\langle q^a, t^p \rangle$ — that is, that it takes an at-issue expression-denoting expression and returns a peripheral proposition.

The introduction of semantic multidimensionality necessitates a corresponding elaboration of the rules of semantic composition. Both dimensions of the meaning of \mathfrak{M} , for example, compose with the meaning of its type q^a sister. How does this work? And what happens to the peripheral content of a node when only its at-issue content has the right type to compose with its sister? My answer to the first of these questions is that the meaning of a node can compose with more than one dimension of the meaning of its sister. This means that the meaning of the type q^a expression which goes sister to \mathfrak{M} is not “used up” when it composes with the at-issue content of \mathfrak{M} . My answer to the second question is that a peripheral content is simply passed along unmodified if it is not of the right type to compose with any meaning of its sister node.¹⁵

We are now in a position formally to describe the examples of covert quotation introduced in Section 1. For clarity, I will represent covert quotation

¹⁴Officially, our basic type system will be defined as follows: e^a , q^a , u^a , s^a , and t^a are the basic at-issue types; t^p is the basic peripheral type. (Type q^a expressions refer to quoted items; type u^a items refer to utterances.) The system also incorporates an at-issue type χ^a , distinct from the basic types, which is the type of the output of the $(*)$ function. For further details, see the appendix.

¹⁵Formally, these two aspects of the semantic system are captured by introducing two tree admissibility conditions: At-Issue Function Application and a -to- p Shunting. See the appendix for details.

using the following convention: for a covertly mixed-quoted expression ϕ (henceforth, an *m-quoted* expression), I will write ${}^m\phi^m$; for a covertly pure-quoted expression ψ (henceforth, a *p-quoted* expression), I will write ${}^p\psi^p$. Starting with a simple example, we have:

- (9) Americans eat tom[eIrouz] and Brits eat tom[a:touz].

I analyze this as:

- (9') Americans eat ${}^m\text{tom}[\text{eIrouz}]^m$ and Brits eat ${}^m\text{tom}[\text{a:touz}]^m$.

On the natural assumption that the discourse anaphors introduced by the *m-quoted* expression ‘tom[eIrouz]’ are understood in context to pick out a generic American English utterance with that pronunciation, the first conjunct of (9) expresses the at-issue content that Americans eat tomatoes and the peripheral content that Americans (generically) utter ‘tom[eIrouz]’. In slightly more detail: both dimensions of the meaning of \mathfrak{M} take as argument the covert pure-quotation ${}^p\text{tom}[\text{eIrouz}]^p$. One dimension (the one that accounts for the peripheral content of mixed quotations) then introduces the not-at-issue content that individual or community s_x (here resolved in discourse to the set of Americans) has produced a (generic) meaningful utterance u_x of ‘tom[eIrouz]’: $R(s_x, u_x, \text{‘tom}[\text{eIrouz}]’)$. Relative to context c and point of evaluation w, g , the other dimension returns the extension at w of the intension contributed by an American speaker uttering ‘tom[eIrouz]’ in w_c : $(\text{‘tom}[\text{eIrouz}]’)(w_c)(s_x)$. Thus $\llbracket \lambda q. (q)(w_c)(s_x) \rrbracket^{c, w, g} (\llbracket \text{‘tom}[\text{eIrouz}]’ \rrbracket^{c, w, g}) = \llbracket \text{tomatoes} \rrbracket^{c, w, g}$. The same story applies, mutatis mutandis, for the second conjunct of (9'). This captures the intuitive appropriateness conditions for the sentence.

Of course, what has been said so far does not yet explain examples like (10) and (12), which I took above to be evidence for a semantic treatment of mixed quotation. The problem is that the proposition that the relevant speaker produced an utterance of the quoted material, which I have located in a peripheral dimension of meaning, appears to interact compositionally with higher at-issue material in these sentences. To explain how such examples work, we need to introduce into our semantics an operator ‘ \downarrow ’ which optionally occurs in the syntax and shunts propositional content from the peripheral dimension into the at-issue dimension, conjoining it with at-issue content.¹⁶ We can then explain (10) by saying that the \downarrow operator occurs at

¹⁶More carefully, we need both an operator \downarrow and a new tree admissibility condition. The translation of \downarrow into \mathcal{L}_{MQ} is simple enough:

Covert mixed quotation

the top of the clause which is modified by the adverb of quantification, so that ‘When in Santa Cruz’ goes sister to the conjunction of the proposition that Peter orders apricots and the proposition that Peter utters ‘[eI]pricots’ in so doing. Schematically:

- (10′) When in Santa Cruz, \downarrow (Peter orders ‘[eI]pricots’ at the local market).

Similarly, in (12), we can hold that \downarrow takes scope over the conditional’s antecedent to explain why it intuitively takes us not just to the closest worlds where Peter orders apricots but to the closest worlds where he does so by saying ‘[eI]pricots’.¹⁷

The following five sections apply the account of mixed quotation just described to develop novel analyses of five puzzling semantic phenomena.

3 First application: Conventional implicature items

Conventional implicature items like expressives and non-restrictive relative clauses (NRRCs) project out of a wide range of embedding environments, including negation, polar questions, and many indirect speech reports and propositional attitude ascriptions:

- (20) I didn’t see that bastard Jones on my way to work this morning.
 \hookrightarrow *the speaker has a negative attitude toward Jones*

$$\downarrow \rightsquigarrow \lambda p. \overleftarrow{p} : \langle t^p, t^a \rangle$$

Here, \overleftarrow{p} is constrained to have the same truth-value as p , but is of type t^a rather than type t^p .

But we now require a new tree admissibility condition which explains what to do when a node is associated with an expression in \mathcal{L}_{MQ} which has only peripheral content. The condition I propose, which I call *p-to-a Shunting*, stipulates that when \downarrow goes sister to a node hosting expressions of both type t^a and t^p , it shunts the peripheral proposition into the at-issue domain and conjoins it with the at-issue proposition. Details can be found in the appendix.

¹⁷The \downarrow operator also provides an explanation for readings of quantified sentences like ‘Whenever a student from Santa Cruz visits, she orders ‘[eI]pricots’ at the market’ on which they require that each student order apricots by uttering ‘[eI]pricots’. For if the peripheral content of ‘she orders ‘[eI]pricots’ at the market’ is shunted into the at-issue dimension, the speaker and utterance variables it contains can then be bound by higher quantifiers. This is what enables the account to avoid the kinds of binding problems which Maier (2021) uses to argue against two-dimensional treatments of mixed quotation. Thanks to Matt Mandelkern for suggesting this example.

- (21) Did you see Smith, who has eleven toes, on your way to work this morning?
 ↪ *the speaker believes that Smith has eleven toes*
- (22) Smith said that he saw that bastard Jones on his way to work this morning.
 ↪ *the speaker has a negative attitude toward Jones*
- (23) Jones thinks that I saw Smith, who has eleven toes, on my way to work this morning.
 ↪ *the speaker believes that Smith has eleven toes*

In each of (20–23), the most natural interpretation is one on which the speaker has the attitude signaled by the expressive or believes the propositional content of the NRRC. Potts (2005) calls these *speaker-oriented* interpretations of the relevant material. The data just canvassed provides motivation for holding, along with Potts, that the implicated content of CI items projects out of *all* embedding environments (except quotation). This has been thought to distinguish them semantically from presuppositions, which do not generally project out of indirect speech reports or nonfactive propositional attitude ascriptions.¹⁸

Slurs appear to pattern with conventional implicature items when it comes to projection:

- (24) Did you see two Krauts on your way to work this morning?
 ↪ *the speaker has a negative attitude toward German people*
- (25) Jones thinks that I saw two Krauts on my way to work this morning.
 ↪ *the speaker has a negative attitude toward German people*

This sort of observation has led many theorists, including Potts (2005), Williamson (2009), McCready (2010), and Kirk-Giannini (2019) to analyze slurs as conventional implicature items. On the basis of these similarities, in

¹⁸Perhaps drawing this distinction is not so simple. For example, *de re* readings of definite descriptions appear to involve presuppositions projecting out of indirect speech reports and propositional attitude ascriptions. My remarks in this section are directed primarily to those who believe there is an important theoretical distinction between presuppositions and conventional implicatures and hold that conventional implicatures generally project out of all embedding environments. Thanks to an anonymous referee for pressing me to clarify this point.

Covert mixed quotation

what follows I will use the term ‘conventional implicature items’ (CI items) to refer to slurs, NRRCs, and expressives.

Interestingly, there are contexts in which projection of the implicated content of CI items seems not to occur.

- (26) Jones complained all day that he’d had to spend an hour that morning looking for his goddamned keys.
 ↗ *the speaker has a negative attitude toward Jones’s keys*
 ↖ *Jones has a negative attitude toward Jones’s keys*
- (27) My father screamed that he would never allow me to marry that Kraut Webster, who always stank of sauerkraut and Kölsch, or any other damned Kraut. (Kirk-Giannini 2019)
 ↗ *the speaker believes that Webster always stinks of Sauerkraut and Kölsch*
 ↖ *the speaker’s father believes that Webster always stinks of Sauerkraut and Kölsch*
- (28) John thinks that the Krauts will have taken over the whole neighborhood in another couple years. But of course, I think it’s great that we’re developing such a vibrant German community. (After Camp 2018)
 ↗ *the speaker has a negative attitude toward German people*
 ↖ *John has a negative attitude toward German people*

These are cases in which the attitude signaled by the expressive, propositional content of the NRRC, or intended pejorative effect of the slur intuitively attach not to the speaker but to the subject of the sentence (that is, these are cases in which the natural interpretation of the material is *non-speaker-oriented*).

There is thus a tension between the treatment of CI items suggested by the initial projection data and the fact that there are isolated cases in which projection fails and the resulting readings are non-speaker-oriented. Potts (2005), for example, achieves the projective profile for CI items suggested by (20–23) by handling conventionally implicated content in a peripheral dimension of semantic composition with an enforced separation from the at-issue dimension, so that it is in principle not possible for embedding expressions semantically to prevent its projection.

To explain the failure of conventional implicature items to project out of certain constructions, I motivate the idea that, while mixed-quoted material contributes its at-issue semantic content to composition, it does not contribute any other dimension of its content. The fact that sentences involving

overt mixed quotation are especially liable to generate non-speaker-oriented readings lends support to this idea:

- (26') Jones complained all day that he'd had to spend an hour that morning looking for his 'goddamned keys.'
- (27') My father screamed that he would never allow me to marry 'that Kraut Webster, who always stank of Sauerkraut and Kölsch, or any other damned Kraut.'
- (28') John thinks that 'the Krauts' will have taken over the whole neighborhood in another couple years. But of course, I think it's great that we're developing such a vibrant German community.

On the assumption that expressives and NRRCs contribute exclusively to peripheral content (defended by Potts (2005)) and that slurs are "mixed-content" expressions in the sense that, apart from their peripheral pejorative content, they contribute the same property as their neutral counterparts (endorsed by e.g. Williamson (2009), McCready (2010), Camp (2013, 2018), and Kirk-Giannini (2019)), sentences like (26') to (28') constitute empirical evidence for the thesis that the quotational interpretation function ($\llbracket \cdot \rrbracket$) returns only the at-issue content of quoted expressions. If we adopt this view of overt mixed quotation, we can hold that the at-issue content of (27'), for example, is that *the speaker's father said (in the manner of screaming) that he would never allow the speaker to marry that German Webster or any other German*, while the sentence peripherally entails that some salient individual (here the speaker's father) literally uttered 'that Kraut Webster, who always stinks of Sauerkraut and Kölsch, or any other damned Kraut.' This explains why the content of the NRRC (inter alia) is not entailed by the sentence, as it is in (21) and (23).

In slightly more detail, the idea is that the material 'that Kraut Webster, who always stinks of Sauerkraut and Kölsch, or any other damned Kraut' is first pure-quoted, resulting in an expression of type q^a . This process eliminates all semantically encoded peripheral content associated with the slurs, the expressive, and the NRRC. The pure-quoted expression then serves as argument for \mathfrak{M} , which, *ex hypothesi*, contributes to further at-issue composition only its at-issue content. \mathfrak{M} also contributes the peripheral proposition that the speaker's father uttered 'that Kraut Webster, who always stinks of Sauerkraut and Kölsch, or any other damned Kraut'. Given interlocutors' background knowledge of English, it will be clear to them that anyone who sincerely utters such a thing (i) harbors negative attitudes towards German people and (ii) believes that Webster always stinks of

Covert mixed quotation

Sauerkraut and Kölsch. This is why (27') conveys that the speaker's father (rather than the speaker) has the relevant attitudes and beliefs.

If this theory of the interaction between overt mixed quotation and CI items is correct, we can explain non-speaker-oriented interpretations of slurs and other CI items by appealing to covert mixed quotation in the sentences which embed them, as in:

(26'') Jones complained all day that he'd had to spend an hour that morning looking for his ^mgoddamned keys.^m

The prediction would then be that only the at-issue content of the quoted material would be speaker-oriented, while the not-at-issue content would be associated with whatever individual is presupposed to have literally uttered the quoted material. This does indeed seem to be what is going on in non-speaker-oriented interpretations of CI items.

Thus a theory of covert mixed quotation which incorporates the independently motivated idea that not-at-issue content does not project out of mixed quotation offers an elegant explanation of non-speaker-oriented interpretations of CI items.

4 Second application: C-monsters

If the world's English speakers collectively agreed to use the word 'water' to pick out gold, water would (still) be wet. This is because a conditional's consequent is interpreted using the actual conventions of English, even if its antecedent introduces a situation in which different conventions prevail. Or, at least, so it has generally been assumed. Kocurek et al. (2020; hereinafter *KJR*) label this thesis about the way the actual conventions feature in the interpretation of sentences *Conventional Wisdom*:

(Conventional Wisdom): Truth at a scenario (counterfactual or otherwise) is evaluated relative to our (or the speaker's) actual linguistic conventions, even if those conventions diverge from the ones adopted in that scenario. (2020: 2)

KJR argue that (Conventional Wisdom) is false: certain embedded clauses can, when uttered in the right contexts, characterize conditions on the linguistic conventions in force rather than on the way the non-linguistic world is. When this happens, they call the embedding expressions *c-monsters* — expressions which shift the conventions relevant for interpreting embedded material.¹⁹

¹⁹This terminology was first introduced by Einheuser (2006).

As an example of the failure of (Conventional Wisdom), KJR remind us of the controversy that followed Pluto’s 2006 recategorization as a dwarf planet by the International Astronomical Union. In the context of that controversy, the following sentences all seem acceptable:

- (4) Pluto could have easily been a planet. But that one stubborn scientist voted for the current definition, so it is not. (2020: 6)
- (29) If Pluto were a planet, there would be dozens of planets in the solar system. (2020: 5)
- (30) Alpha thinks that Pluto is a planet. (2020: 6)
- (31) Pluto used to be a planet, but it isn’t any more. (2020: 6)
- (32) Beta wishes Pluto were a planet (but they couldn’t care less about whether it clears its orbital neighborhood). (2020: 7)

The key motivation behind Pluto’s recategorization was that it does not ‘clear its orbital neighborhood.’ We can imagine that (4) and (29–32) are acceptable even though there is no uncertainty about whether Pluto clears its orbital neighborhood. In such cases, (4) seems to characterize the modal fragility of the linguistic conventions. Similarly, (29) seems to convey that, if the conventions governing the use of ‘planet’ were such as to include Pluto in the extension of that term, they would also be such as to include many other objects in the solar system, (30) seems to attribute to Alpha belief that the conventions governing the use of ‘planet’ *are* such as to include Pluto, (31) seems to characterize the temporal contingency of the conventions, and (32) seems to characterize a wish about the conventions.²⁰

KJR suggest responding to evidence like (4) and (29–32) by substantially overhauling both the semantics of natural language and the theory of communication: worlds as points of evaluation and objects of the propositional attitudes used to define the common ground are replaced with world/convention pairs, and normal subject-predicate sentence semantically express something very similar to what in the Stalnakerian tradition is known as *diagonal content* — the proposition semantically associated with ‘armadillos are edentate’ is not that armadillos are edentate but rather that whatever

²⁰At least, this is one available reading of (30). I find it more natural interpret (30) as indicating that Alpha believes that the meaning of ‘planet’ *should* be such as to include Pluto in its extension. The account of metalinguistic negotiation I present in Section 6 predicts the availability of this reading.

Covert mixed quotation

individuals are picked out by ‘armadillos’ according to the conventions have whatever property is picked out by ‘edentate’ according to the conventions. KJR then provide a semantics for the conditional which predicts its c-monstrous behavior in sentences like (29).

The overall package of views KJR propose is fairly radical: not only do we give up on (Conventional Wisdom), we must also revise our intuitive beliefs about the truth conditions of all natural language sentences and the objects of the propositional attitudes. So there is reason to explore ways to account for (4) and (29–32) without giving up quite so much.

By treating (4) and (29–32) as involving covert mixed quotation, it is possible to explain their intuitive truth conditions without giving up (Conventional Wisdom) or radically altering our semantics or theory of communication. In particular, I suggest that c-monstrous behavior is to be explained by positing covert mixed quotation together with an independently motivated second covert lexical item † which performs an operation somewhat like the diagonalization operation familiar from two-dimensional theories of assertoric content.²¹

Recall that, according to the theory of mixed quotation we have adopted, a mixed quotation contributes to at-issue composition whatever semantic value the quoted material had when it was uttered by the anaphorically retrieved individual or linguistic community. This is not quite what we want for an analysis of alleged c-monsters, since (for example) the counterfactual morphology in (4) and (29) indicates that Pluto is not in fact a planet in whatever sense of ‘planet’ is operative, and a straightforward mixed-quotative analysis on which ‘planet’ was used with its pre-2006 meaning would make it true in the actual world that Pluto is a ‘planet’. What is wanted is instead something closer to *thing which counts as a planet given the operative conventions*. Assigning this meaning to ‘planet’ explains why (i) Pluto is not actually in its extension, but (ii) there are nearby possible worlds where Pluto is in its extension, as suggested by (4).

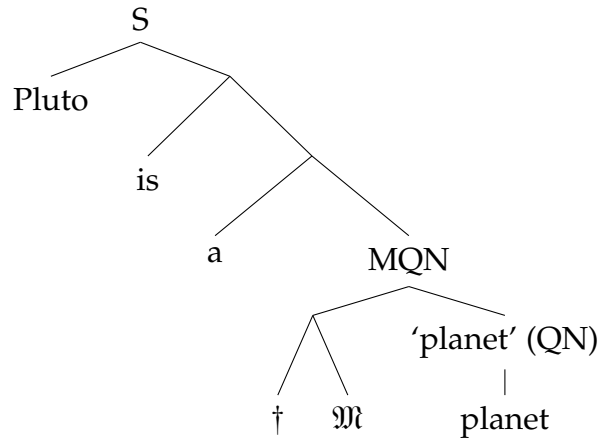
There is evidence independent of the phenomenon of alleged c-monsters that overt mixed quotations can sometimes have the desired sort of meaning. Shan (2010: 428), for example, points out that:

²¹The operation performed by † differs from diagonalization in the traditional sense (as defined in Stalnaker 1978) in that, whereas diagonalization in the traditional sense targets an assertoric utterance and considers what the horizontal proposition of that utterance would have been if it had been produced at other possible worlds, † makes use of Shan’s quotative interpretation function (∗) to consider what the semantic content of the quoted material would have been if it had been produced by the anaphorically retrieved speaker at other possible worlds.

(33) To be ‘eckullectic’ is to have never been seen by Bush.

has a reading on which it is true at a world of evaluation w just in case *at* w (rather than at the world of utterance) Bush uses ‘eckullectic’ to express the property of having never been seen by him. This is exactly the sort of meaning which is needed to explain alleged c-monsters: the idea is that ‘eckullectic’ contributes at a world of evaluation w the property a thing has just in case it falls in the extension of ‘eckullectic’ as Bush would use it at w . The question is simply how to account for readings of this kind within a broader theory of mixed quotation.

Shan’s observation motivates positing a covert lexical item which ‘diagonalizes’ the content of a mixed quotation. This is the function of the \dagger operator, which, I suggest, can optionally occur in the syntax as a sister to \mathfrak{M} :



Informally, the role of \dagger is to change the argument structure of the $(*)$ function so that, relative to a world of evaluation w , it takes an expression and an individual or linguistic community and returns the extension at w of that expression as uttered by that individual or linguistic community at w .²²

²²Formally, we can first define $(*)^*$ as the function such that, for any expression q and any speaker or community s , $\llbracket (q)^*(s) \rrbracket^{c,w,g} = \llbracket (q)(w)(s) \rrbracket^{c,w,g}$. Then, for any function f such that there is a speaker or community s such that $\llbracket f \rrbracket^{c,w,g} = \llbracket \lambda q.(q)(w_c)(s) \rrbracket^{c,w,g}$, let $\llbracket f^* \rrbracket^{c,w,g} = \llbracket \lambda q.(q)^*(s) \rrbracket^{c,w,g}$. Now let $\dagger \rightsquigarrow \lambda f_{\langle q^a, \lambda^a \rangle} | \exists s : f = \lambda q.(q)(w_c)(s).f^*$.

In order to secure the well-definedness of f^* , I assume here that there are no two speakers or linguistic communities which assign the same extensions to all expressions in w_c but assign different extensions to some expressions in other worlds. This assumption is will likely be plausible to those who think that there are pervasive individual differences in the extensions of everyday predicates like ‘chair’ and ‘cup.’ One could in principle give up this assumption, but doing so would require complicating the semantics.

Covert mixed quotation

Applied to, for example, (29), this proposal predicts that the antecedent of the conditional takes us to the nearest world w at which Pluto has whatever property is picked out by the word ‘planet’ as uttered by the anaphorically retrieved speaker or community (here assumed to be a generic speaker of English) at w . In the context of discussion of the 2006 recategorization, the world in question will be one at which the meaning of ‘planet’ never changed. This vindicates KJR’s intuition that (29) in some sense shifts the conventions relative to which the word ‘planet’ is interpreted, but does so without requiring substantive changes to the theories of meaning or communication or to the lexical entries for the various embedding expressions KJR identify as c-monsters.²³

Thus a theory of covert mixed quotation which incorporates the independently motivated operator \dagger offers an elegant explanation of alleged c-monsters.

5 Third application: Metalinguistic negation and related phenomena

In this section, I argue that the theory of covert mixed quotation I have developed can be extended to account for cases of metalinguistic negation, as in (3) and (34), as well as related phenomena, as in (35) and (36):

- (3) I didn’t manage to trap two MONGEES — I managed to trap two MONGOOSES. (Horn 1989: 371)
- (34) I don’t eat tom[eirouz]; I eat tom[a:touz]. (Carston 2002: 298)

²³KJR (2020, 18) consider and reject diagonalization strategies for accounting for alleged c-monsters on the basis that they cannot account for the truth of sentences like ‘If Pluto were a planet, then there would be dozens of planets in the solar system even if everyone were dead.’ The problem here is that the antecedent takes us to a world where there are no people, so *a fortiori* there is nothing they mean by the word ‘planet’ for the diagonalization operator to pick up on. But I believe this conclusion is too hasty when it comes to theories, like mine, which combine diagonalization with mixed quotation. We saw in Section 2 that the source of a mixed quotation can be merely hypothetical. In cases where the source is hypothetical, the at-issue content of a mixed quotation cannot be $\llbracket (q)(w_c)(s_x) \rrbracket^{c,w,g}$, since the source s_x does not exist at w_c . Instead, we do something like consider the closest world to w_c at which s_x exists and utters q ; what is contributed to at-issue composition is what s_x means by uttering q in that world. Taking the source in KJR’s problem case to be a hypothetical English speaker allows us to apply the same strategy, evaluating the mixed quotation in the sentence by moving to the closest world to the world of evaluation where this hypothetical English speaker exists and utters ‘planet’. Thanks to an anonymous referee for pressing me to clarify this point.


- (35) If you use a sidewalk, you must be American. (Carston 2002: 299)
- (36) It's highly unlikely (I very much doubt that) Johnny's eaten SOME of the cakes; he will have eaten ALL of them. (Carston 2002: 300)

(3) and (34) have the characteristic property of metalinguistic negation: the negated sentence is entailed by the affirmed sentence. As a result, the negation seems to convey not that the negated sentence is false, but rather that it is in some sense inappropriate. Theorists who have drawn a connection between metalinguistic negation and inappropriateness or objectionability include Horn (1985, 1989), Seuren (1990), Carston (1996, 2002), and Noh (2000). (35) and (36) similarly convey something about appropriateness, but without involving negation. In (35), the consequent seems to be evaluated at worlds where it is appropriate in the dialect of the generic individual introduced in the antecedent to use the word 'sidewalk', while (36) conveys that the speaker thinks it unlikely that Johnny has eaten some but not all of the cakes, and thus unlikely that it is appropriate to describe the situation using 'some,' as this would generate the scalar implicature *not all*.

The fundamental challenge for theories of metalinguistic negation is to explain how examples like (3) and (34–36) come to be associated with the metalinguistic contents they intuitively convey. (3), for example, conveys that it is or would be inappropriate to refer to mongooses by uttering 'mongeese'. Simply introducing some covert operator contributing the meaning *it is appropriate that* does not explain how the sentence comes to express a metalinguistic content about utterances of 'mongeese'. But it is possible to combine the intuitive idea that metalinguistic negation involves attributions of appropriateness with a theory of covert mixed quotation in order to explain sentences like (3) and (34)–(36). On the resulting account, it is covert mixed quotation which generates the metalinguistic contents that then interact with an appropriateness operator to produce the desired readings.

To account for metalinguistic negation, then, let us introduce a covert lexical item which composes with the peripheral meaning of covert mixed-quoted material so that what is contributed is not the proposition that the quoted material has been uttered verbatim, but rather the proposition that its verbatim use is or would be appropriate. This proposition concerning appropriateness can then be shunted into the scope of the negation operator

by \downarrow . If we symbolize this covert material using \mathfrak{A} , the idea is that the logical form of the first part of (3) is:²⁴



On this analysis, (3) contributes truth-conditionally the proposition that (i) it is not the case that (I managed to trap two mongooses and it is appropriate to express this fact using the word ‘mongeese’), and (ii) I managed to trap two mongooses. In this way, the apparently metalinguistic character of metalinguistic negation is explained by the presence of covertly quoted material in the scope of the negation rather than by positing anything unusual about the negation operator itself, as suggested by Horn (1989) and Potts (2007).

What, formally, is the meaning of \mathfrak{A} ? The idea is that \mathfrak{A} operates on peripheral propositional material. In particular, it takes the peripheral proposition contributed by the mixed-quotation operator \mathfrak{M} — the proposition that the anaphorically retrieved speaker uttered the quoted material in the course of the anaphorically retrieved utterance — and modalizes it to yield the proposition that it is or would be appropriate for the relevant speaker to do so. If we introduce the existential modal \blacklozenge into \mathcal{L}_{MQ} in order to express

²⁴To avoid overgeneration issues, we can stipulate that \mathfrak{A} can occur only as the specifier of a mixed-quotational phrase.

this modality, we can assign \mathfrak{A} the \mathcal{L}_{MQ} translation $\mathfrak{A} \rightsquigarrow \blacklozenge$ (type $\langle\langle s^a, t^p \rangle, t^p \rangle$). Since \blacklozenge is the first \mathcal{L}_{MQ} expression we have considered which operates on peripheral content without shunting it into the at-issue dimension, we will also need to add in a corresponding composition rule.²⁵

On the present proposal, then, the semantic story about the first conjunct of (3) is as follows. The covertly pure-quoted expression ‘mongeese’ serves as argument to the two meanings of \mathfrak{M} , yielding the at-issue property *mongoose* and the peripheral proposition that the relevant speaker uttered that expression verbatim. As we move up the tree, \mathfrak{A} composes with this peripheral meaning to produce the proposition that it is or would be appropriate for the relevant speaker to utter ‘mongeese’; at this point, this proposition remains in the peripheral dimension of meaning. Composition in the at-issue dimension then occurs normally until the t^a -type meaning corresponding to the proposition that I managed to trap two mongooses is produced. At this point, \downarrow shunts the peripheral proposition about appropriateness into the at-issue dimension and conjoins it with the t^a -type meaning which arose from normal composition. It is this conjunctive proposition which is then negated, yielding the final interpretation that it is not the case that: I managed to trap two mongooses and it is or would be appropriate to describe this state of affairs using ‘mongeese’. In the presence of the second conjunct of (3), which entails that I managed to trap two mongooses, the first conjunct is interpreted as entailing that it is not appropriate to describe mongooses using ‘mongeese.’ This result corresponds to the desired truth-conditions for the sentence as a whole.

(34) receives a similar treatment to (3). It is worth noting that in (34) it is somewhat unclear whether what is being covertly quoted is just ‘tom[eirouz]’ or the entire verb phrase ‘eat tom[eirouz]’. On my proposal, this is a distinction without a difference: if we analyze the first clause as *I don’t eat $\mathfrak{A}(\text{tom[eirouz]}^m)$* , we generate a reading on which it is not the case that (the speaker eats tomatoes and it is or would be appropriate for her to refer to them as ‘tom[eirouz]’), but she does eat tomatoes. If we analyze the first clause as *I don’t $\mathfrak{A}(\text{eat tom[eirouz]}^m)$* , we instead get a reading on which it is not the case that (the speaker eats tomatoes and it is or would be appropriate

²⁵The rule in question is Peripheral Intensional Function Application — see the appendix for details. Note that the behavior of \blacklozenge might be thought to be unusual in so far as not-at-issue content is usually held to project out of operators like modals. In the context of the present system, however, this generalization is better understood as holding that not-at-issue content projects out of at-issue operators. \blacklozenge does not threaten the generalization when it is construed in this way. Thanks to an anonymous referee for pressing me to clarify this issue.

Covert mixed quotation

for her to refer to this act by uttering ‘eat tom[eirouz]’), but she does eat tomatoes. Both readings capture the intuitive truth conditions of (34).

This leaves us with (35) and (36), the examples not containing negation. It is relevant in discussing these examples that, as in the case of overt mixed quotation, the peripheral entailment that someone uttered the quoted material verbatim can be satisfied by introducing a hypothetical or generic individual. This is especially clear if, for example, we introduce overt mixed quotation into (35):

(35′) If you use a ‘sidewalk,’ you must be American.

On its generic reading, (35′) is not directed to a particular addressee. What the sentence seems to be conveying is that it is true of a generic individual that, if it is or would be appropriate for them to utter ‘sidewalk’ and thereby express the property *sidewalk*, they must be American. In my view, this is also what is happening in (35):

(35″) If ↓ (you use a \mathfrak{A}^m sidewalk m), you must be American.

Examples like (35) highlight the semantic flexibility of the appropriateness modal \blacklozenge . In some cases, like (3) and (34), it purports to pick out some objective standard of appropriateness on the basis of which the speaker objects to some feature of an earlier utterance; in (35), on the other hand, it picks out a notion of appropriateness relativized to the speech community of a particular individual. In this respect, \blacklozenge patterns with other evaluative modals, and indeed other modals generally.

Similarly, in (36), I suggest that the material in the scope of the probabilistic operator ‘it is highly unlikely that’ is not just that Johnny has eaten (at least) some of the cakes, but rather that he has eaten (at least) some of the cakes *and* it is or would be appropriate to describe this situation by uttering ‘some of the cakes’. Given that the second clause in (36) entails that Johnny has probably eaten some of the cakes, the earlier conjunction is understood to be unlikely because, if Johnny has eaten all of the cakes, it is not appropriate to describe this situation using ‘some of the cakes’ because this would generate a misleading scalar implicature.

One attractive feature of a quotational analysis of metalinguistic negation is that it provides a straightforward explanation of certain syntactic properties observed by Horn (1985, 1989) and Burton-Roberts (1989). First, Horn notes that in cases of metalinguistic negation, it is not possible for the negation to be morphologically incorporated into a predicate:

- (37) She's not happy, she's ecstatic.
 (38) #She's unhappy, she's ecstatic.

The present proposal explains this nicely, since, while 'unhappy' and 'not happy' are arguably semantically equivalent, the material negated in (37) is not simply 'happy' but a complex of overt and covert material.

Second, Horn suggests that cases of metalinguistic negation do not license negative polarity items:

- (39) John didn't manage to solve some of the problems — they were quite easy for him to do.
 (40) #John didn't manage to solve any of the problems — they were quite easy for him to do.

On the assumption that 'manage to solve some of the problems' is *m*-quoted in (39) (and anaphorically retrieves part of a prior utterance of 'John managed to solve some of the problems'), we have an explanation for why its 'some' cannot be replaced with 'any': this would violate the requirement on mixed quotations that the quoted material have been uttered verbatim by the anaphorically retrieved speaker.

Third, [Burton-Roberts \(1989\)](#) points out that cases of metalinguistic negation do not obey the rule of double-negation elimination:

- (41) She's not not happy, she's inconsolable.
 (42) #She's happy, she's inconsolable.

On the assumption that 'not happy' is *m*-quoted in (41), we again have an explanation for why it is not equivalent to (42).

Thus a theory of covert mixed quotation can be combined in a straightforward way with the independently motivated idea that metalinguistic negation and related phenomena involve attributions of appropriateness to predict the intuitive truth conditions of sentences like (3) and (34–36) and explain the syntactic peculiarities of metalinguistic negation.

6 Fourth application: Metalinguistic negotiation

[Plunkett & Sundell \(2013: 3\)](#) introduce the notion of a *metalinguistic negotiation*: "a dispute... that employs competing metalinguistic usages of an expression, and that reflects a disagreement about the proper deployment of linguistic representations." For example, a dispute about whether the

Covert mixed quotation

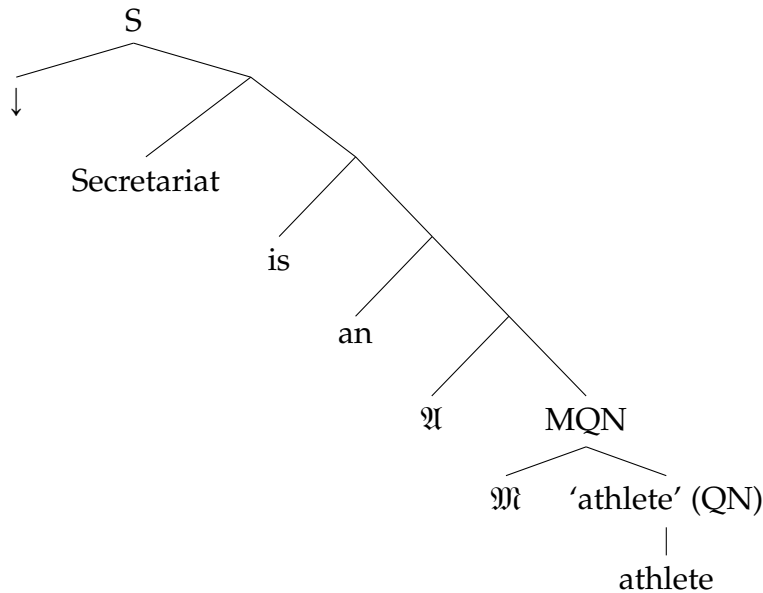
racehorse Secretariat is an athlete might constitute a metalinguistic negotiation if the disputants do not take themselves to have a factual disagreement about whether the conventions of English settle that Secretariat is in the extension of ‘athlete’. If we imagine that (6) is such a case, A seems in some sense to be making a normative rather than a descriptive claim about language — Secretariat *should* count as an athlete, whether or not this is already so. And B seems to be disputing this claim.

- (6) A: Secretariat is an athlete.
B: No, Secretariat is not an athlete.

Plunkett & Sundell claim, regarding (6), that A and B ‘literally express mutually consistent contents... *because* they do not mean (in the relevant sense) the same things by their words’ (2013: 18). The word ‘athlete’ has a different extension at the actual world for A than it does for B, so that Secretariat is in the extension of ‘athlete’ as A uses the word but not as B uses it. They conclude that “the connection between genuine disagreement and sameness of meaning is broken” in the sense that two speakers can disagree even while the contents they express are compatible (2013: 18).

Without presuming to pronounce on the virtues of the broader theoretical package Plunkett & Sundell offer for thinking about normative disagreement, I would like to offer an alternative diagnosis of cases like (6) which preserves the intuition that the speakers are disagreeing in virtue of expressing incompatible contents. In our discussion of metalinguistic negation, the propositions about appropriateness generated by \mathfrak{A} took scope under negation. But there is nothing which requires that this always be the case. Indeed, in positing \mathfrak{A} , we predict the existence of sentences where it occurs outside the scope of negation. And this is exactly what appears to be happening in cases of metalinguistic negotiation.

Let us follow Plunkett & Sundell in holding that A and B express different properties when they utter ‘athlete’ — say that A expresses the property of *broad athleticism* and B expresses the property of *narrow athleticism*. Now suppose that A’s utterance in (6) contains covert mixed quotation around ‘athlete’, with the anaphorically retrieved utterance and speaker understood in discourse to be a hypothetical use of ‘athlete’ (by an individual who agrees with A concerning the meaning of the word) to describe a horse or other nonhuman animal. The system developed so far then predicts that the following underlying structure is available for A’s utterance:



On this analysis, A's utterance is true just in case (i) Secretariat instantiates the property of broad athleticism (which, as Plunkett & Sundell point out, is uncontroversial) *and* (ii) it is or would be appropriate to characterize a horse or nonhuman animal using 'athlete.' Intuitively, it is this second conjunct with which B takes issue.

We can then take B's utterance to be a straightforward case of metalinguistic negation: it contains covert mixed quotation around 'athlete', the anaphorically retrieved utterance and speaker are understood to be the same as those picked out by A's utterance, and the appropriateness content is first shunted into the at-issue dimension and conjoined with the proposition that Secretariat instantiates broad athleticism by ↓ and then negated as part of the resulting conjunction. Since it is uncontroversial that Secretariat instantiates broad athleticism, the negated conjunction will be understood as a denial that it is or would be appropriate to characterize a horse or nonhuman animal using 'athlete'. Thus, while metalinguistic negotiations may be metalinguistic in the sense that they involve quotation, on the present analysis they do not support Plunkett & Sundell's claim that there may be genuine disagreement without the expression of incompatible contents. To the extent that the idea that disagreement requires the expression of incompatible contents is intuitively plausible, this fact constitutes a reason to prefer a quotational account of (6) and related cases over their account.

The fact that the mechanisms posited in Section 5 to account for metalinguistic negation offer us such a straightforward explanation of metalinguistic

Covert mixed quotation

negotiation constitutes an important kind of indirect evidence for the correctness of both analyses. Moreover, the fact that what is at issue in metalinguistic negotiations is so clearly the appropriateness of different ways of using language should assuage any concerns that the introduction of \mathfrak{A} into our total theory is undermotivated.

7 Fifth application: ‘In a sense’

Recall sentence (7) above:

- (7) You are right to say that viruses are alive in a sense, but not usually classified in those 5 kingdoms.

What does it mean to say that viruses are alive *in a sense*? The claim that viruses are alive in a sense should be distinguished from the claim that viruses are alive. When one claims that viruses are alive, one uses the sentence ‘viruses are alive’ with its normal meaning in one’s language. When one claims that viruses are alive in a sense, on the other hand, one conveys only that the sentence ‘viruses are alive’ expresses a truth when used with some (not completely outlandish) meaning or other. That is, in claiming that viruses are alive in a sense, one conveys that there is a (not completely outlandish) interpretation of ‘viruses are alive’ such that the sentence is true when it has that interpretation.

The intuitive truth conditions for ‘in a sense’ constructions are thus metalinguistic at least in the sense that they are most easily stated by mentioning expressions. My suggestion is that they are also metalinguistic in the stronger sense of involving existential quantification into covert mixed quotation. In particular, though in ordinary cases of mixed quotation, \mathfrak{M} ’s at-issue contribution $\lambda q.(q)(w_c)(s_x)$ contains a free variable s_x which discourse-anaphorically picks out the speaker of the salient utterance, nothing prevents us from binding this variable to produce a predicate of speakers (that is, of individuals or linguistic communities). Since different individuals and linguistic communities have different idiolects, quantifying over them is a way of quantifying over interpretations of the mixed-quoted material.

For simplicity, let us substitute for (7) the simpler:

- (7’) There is a sense which viruses are alive in.

We will build up to a full analysis of (7’) by first considering a non-quantificational analogue:

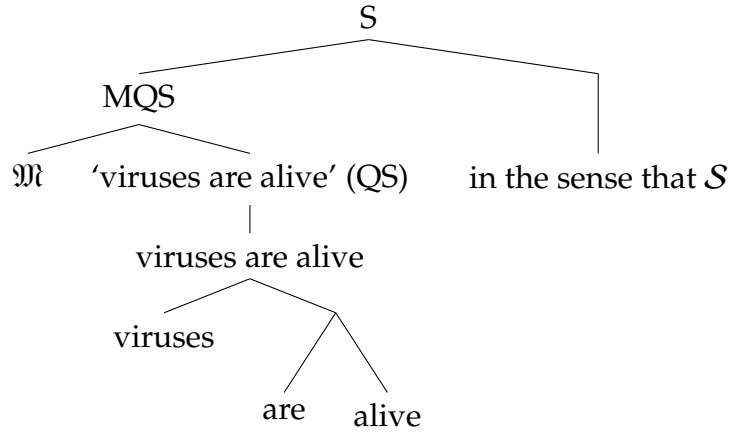
- (43) Viruses are alive in the sense that they have genetic material.

Intuitively, ‘that they have genetic material’ in (43) semantically contributes the sense (intension) of a sentence, and (43) as a whole is true just in case the sentence ‘Viruses are alive’ is true when it is understood as contributing this intension. Since taking the covert quotation in (43) to be pure quotation would result in syntactic and semantic type mismatches, it must be mixed quotation; in this case, the anaphorically retrieved speaker appears to be a generic speaker of English. Drawing on the idea of binding the speaker variable s_x , I analyze (43) as:

$$\exists s_x[(\ulcorner \text{‘viruses are alive’} \urcorner)(w_c)(s_x) \wedge ([\hat{\ulcorner \text{‘viruses are alive’} \urcorner}](w_c)(s_x)) = [\hat{\ulcorner \text{viruses have genetic material} \urcorner}]]^{26}$$

Here, the \mathcal{L}_{MQ} function $[\hat{\quad}]$ intensionalizes the material which occurs inside it. Thus $[[\hat{\phi}]]^{c,w,g}$ is the function which maps a world w' to the extension of ϕ at point of evaluation w',g .

More generally, if \mathcal{S} is an English sentence with intension μ , my suggestion is that $\ulcorner \text{Viruses are alive in the sense that } \mathcal{S} \urcorner$ has the underlying syntactic structure:



And that its root node is translated into \mathcal{L}_{MQ} as:

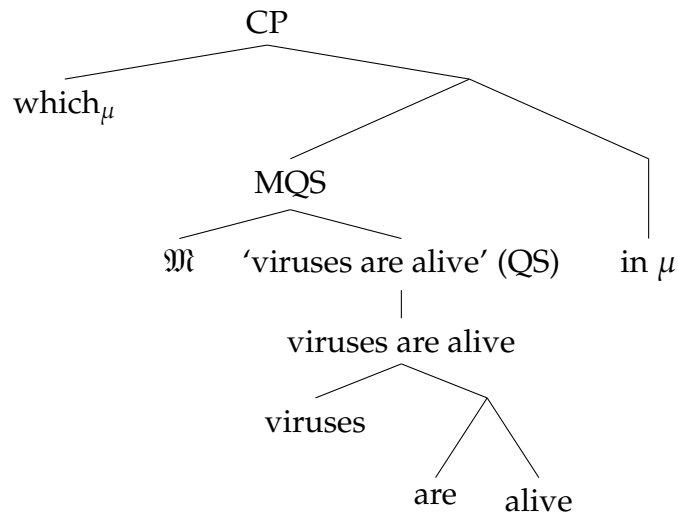
$$\exists s_x[(\ulcorner \text{‘viruses are alive’} \urcorner)(w_c)(s_x) \wedge ([\hat{\ulcorner \text{‘viruses are alive’} \urcorner}](w_c)(s_x)) = \mu]$$

²⁶This is a slight abuse of notation: strictly speaking, the material which falls within the scope of $[\hat{\quad}]$ must be the \mathcal{L}_{MQ} translation of ‘viruses have genetic material,’ not the sentence itself.

Covert mixed quotation

Here I assume that the \mathcal{L}_{MQ} translation of ‘in the sense that \mathcal{S} ’ is a variable μ of type $\langle s^a, t^a \rangle$ which picks out the same function as $[\hat{\cdot}(\mathcal{S})]$. We can integrate this analysis into the formal system developed in the preceding sections by introducing a syncategorematic meaning rule for sentences of the form $\ulcorner^m \mathcal{S}^m$ in the sense that \mathcal{S}' .

We are now in a position to analyze the quantified example (7'). Central to my explanation is the idea that the relative ‘which’ in (7') functions, like other relatives, to introduce lambda abstraction — in this case, abstraction on a covert sense variable μ . Adopting the standard assumptions that English relative pronouns come along with unpronounced indices and that English relative clauses contain unpronounced syntactic traces, I take the structure of the relative in (7') to be as follows:



I take the \mathcal{L}_{MQ} translation of the English relative ‘which _{μ} ’ to be a corresponding \mathcal{L}_{MQ} relative ‘which _{μ} ’ and the \mathcal{L}_{MQ} translation of ‘in μ ’ to be a variable μ of type $\langle s^a, t^a \rangle$. Following a version of the predicate abstraction rule familiar from Heim & Kratzer (1998) appropriate for our system, we can then interpret this structure by having ‘which _{μ} ’ introduce a lambda expression binding the occurrence of μ lower in the tree:

$$\lambda\mu.\exists s_x[(\ulcorner^m \text{viruses are alive}'\urcorner)(w_c)(s_x) \wedge ([\hat{\cdot}(\ulcorner^m \text{viruses are alive}'\urcorner)(w_c)(s_x)]) = \mu)]$$

Thus in the at-issue dimension the complete relative clause expresses a property of intensions: the property had by an intension μ at a world of evaluation w just in case there is some individual or community s_x such that

the content s_x would express at the world of the context by uttering ‘viruses are alive’ is μ , and μ is true at w .

We can then treat ‘there is a sense’ as quantifying over intensions: $\exists\mu$. This gets us the following complete \mathcal{L}_{MQ} translation for (7’):

$$\exists\mu\exists s_x[(\ulcorner\text{‘viruses are alive’}\urcorner)(w_c)(s_x) \wedge (\lceil\text{‘viruses are alive’}\rceil)(w_c)(s_x)] = \mu)$$

We have vindicated the idea that (7’) contributes truth-conditionally that there is a meaning of ‘viruses are alive’ on which that sentence expresses a truth. But we have not yet said anything about how to ensure that the meaning in question is not too outlandish. Indeed, if we freely allow ourselves to quantify over hypothetical intensions and hypothetical speakers, it seems all but guaranteed that we will find one which secures the truth of the existential. But it does not seem that, for any sentence ϕ , ‘There is a sense in which ϕ ’ is true. To avoid this unpalatable consequence, I appeal to the idea that in all normal contexts of utterance, the existential quantifiers in sentences like (7’) are restricted so that they quantify over not-too-outlandish speakers and intensions.²⁷

In sum, it is apparent from the overt structure of ‘in a sense’ constructions that they are quantificational in some way, but without positing any covert structure it is mysterious how to derive their intuitive truth conditions. Combining an independently motivated theory of mixed quotation with the idea that ‘in a sense’ constructions involve covert mixed quotation provides us with the resources to extend standard accounts of quantification to capture their truth conditions.

8 Objections and comparisons

This section defends the theoretical package consisting of my account of covert mixed quotation and its five applications from two different sorts of worries: first, it responds to a methodological concern targeting the idea of covert mixed quotation; second, it compares my treatments of two of the semantic phenomena discussed above to some salient alternatives and argues that my treatments should be preferred.

²⁷On quantifier domain restriction, see e.g. Stanley & Szabó (2000).

8.1 An objection to the thesis of covert mixed quotation

Any semantic treatment of covert quotation faces the worry that freely allowing covert quotation operators into the syntax will result in an implausible multiplicity of possible interpretations even for simple sentences. There is an obvious overgeneration worry here. If nearly any part of any sentence could in principle contain a covert mixed quotation operator, how are we to explain the fact that, in normal cases, speakers are taken not to be quoting?

I offer three observations in an effort to make this overgeneration worry less acute.

First, it is not clear that pragmatic accounts of covert quotation do better than semantic accounts in this area: they simply face a pragmatic overgeneration problem rather than a semantic one. As soon as it is admitted that covert quotation is possible and occurs in a wide variety of environments, there is a worry that *whatever* mechanism is invoked to account for it will overgenerate quotative readings. Without a detailed comparison of theories, there is no reason to think that a pragmatic account will do better with respect to overgeneration than a semantic one.

Second, we can say that there is a defeasible preference for attributing simpler logical forms to speakers, so that interpreters will in general posit covert quotation only if there are special reasons for them to do so. For example, one would be attributing a category mistake to a speaker if one interpreted 'Boston has six letters' as predicating something of the city, and, except in special contexts, one would be attributing a strange thought indeed to a speaker if one interpreted 'Boston is lovely' as concerning the word rather than the city.

Finally, I would like to suggest that interpretations involving covert mixed quotation are dispreferred under indirect speech verbs. Because indirect speech verbs signal a paraphrase rather than verbatim reporting of an utterance, interpreting the complement of an indirect speech report as containing a covert mixed quotation operator is only expected when failing to do so would result in an intuitively less plausible interpretation. For example, in (8), the covert mixed quotation operator is recoverable even in the scope of the indirect speech operator 'replied that' because the alternative to positing covert mixed quotation would be to believe that the narrator's idiolect has shifted dramatically mid-sentence.

8.2 Conventional implicature items

An alternative way of accommodating non-speaker oriented readings, explored by Harris & Potts (2009), is to posit a covert free ‘orientation’ variable associated with each occurrence of a CI item and say that its value is flexibly resolved in discourse. Harris & Potts suggest that non-speaker-oriented interpretations should be explained in this way, rather than by modifying the semantics of various embedding expressions, because non-speaker-oriented readings can occasionally occur even with unembedded CI items. This proposal is, however, subject to at least two difficulties not faced by an analysis in terms of covert mixed quotation.

First, positing a free variable which is contextually resolved in discourse offers no explanation of the fact that speaker-oriented interpretations are almost always preferred over non-speaker-oriented ones. On this subject, Harris & Potts write only that “appositives have many of the morphosyntactic and intonational properties of regular asserted declaratives, which are also overwhelmingly speaker-oriented, so perhaps it is unsurprising that appositives are generally speaker-oriented as well” (547). This explanation does not extend to the observation that expressives and slurs are also usually interpreted in a speaker-oriented manner.

Second, if the semantics of CI items is imagined to work in the same way when the orientation variable takes the speaker as its value as when it takes some other individual, Harris & Potts’s proposal makes it difficult to accommodate certain theories of speaker-oriented uses of CI items. Kirk-Giannini (2019), for example, develops a view according to which the not-at-issue contribution of slurs is directive, where the semantic function of directives is understood, following Portner (2004, 2007), in terms of updates to interlocutors’ To-Do Lists. Accounting for non-speaker-oriented readings of slurs by incorporating orientation variables is in tension with this sort of view insofar as it is not straightforward in the To-Do List framework to model genuinely directive meanings that are not speaker-oriented. Instead, we get a collapse of the view that slurs are directives into a version of a propositional conventional implicature view: speaker-oriented uses of slurs involve the speaker *reporting* that he or she directs his or her audience to adopt the associated perspective rather than actually *directing* them. While this sort of worry does not arise when it comes to NRRCs, which are the focus of Harris & Potts’s discussion, one might reasonably balk at ruling out non-propositional treatments of slurs and expressives in the same way. The problem is perhaps most acute in the case of expressives, where it seems

Covert mixed quotation

clear that speakers are in some sense directly communicating their attitudes rather than describing them.

Accounting for non-speaker-oriented occurrences of CI items in terms of covert mixed quotation thus yields a simpler, more independently motivated, and less restrictive theory.

8.3 Metalinguistic negation

It is important to note that the idea that metalinguistic negation has something to do with quotation does not originate with my proposal. Early discussions of presupposition involve the suggestion that ‘presupposition-cancelling negation’ is quotative in some way. Fillmore (1969: 122), for example, proposes that such cases involve ‘semi-quotations’ — as an example, he offers the sentence ‘I didn’t ‘chase’ the thief; as it happened he couldn’t get his car started.’ Fillmore’s quotative proposal is, however, left undeveloped. More recently, Relevance Theorists like Carston (2002: 301) have suggested that the proposition pragmatically communicated by (34) is: \neg (I eat what is properly called ‘tom[eirouz]’); I eat what is properly called ‘tom[a:touz].’ Unfortunately, Relevance-Theoretic accounts are not specific concerning how the relevant quotative readings are generated. What sets the account I propose apart from these alternatives is that it shows how the desired readings can be derived semantically while departing minimally from an independently motivated theory of mixed quotation.

Two other quotation-based accounts of metalinguistic negation bear mentioning here. First, Seuren (1990) suggests that metalinguistic negation involves (i) the introduction of an underlying cleft structure by contrastive prosody, and (ii) some underlying constituent predicating *being properly said* of some quoted material. For example, on Seuren’s account, the logical form of (3) is (3’):

- (3’) \neg [the x [such that ‘I managed to trap two x ’ is properly said] is ‘mongeese’]; the x [such that ‘I managed to trap two x ’ is properly said] is ‘mongooses’]

It is doubtful, however, that contrastive prosody alone could explain how an underlying structure like (3’), which includes a variable ranging over expressions and quantification into pure quotation, could correspond to (3). Indeed, as Noh (2000) points out, contrastive prosody does not always indicate metalinguistic negation, and not all cases of metalinguistic negation involve contrastive prosody.

The most recent attempt to explain metalinguistic negation in terms of covert quotation is due to Potts (2007). Potts's account is similar to the one I offer below in that it builds on a theory of mixed quotation. It differs from mine, however, in holding that natural language negation is semantically ambiguous between at-issue-content denying and peripheral-content denying meanings. The account I propose has the attractive feature of allowing a univocal account of natural language negation. In treating his peripheral negation as a simple truth-function, Potts also derives implausible truth conditions for simple cases of metalinguistic negation. For example, his account predicts that the metalinguistic negation reading of (3) entails at-issue that the speaker trapped two mongooses while peripherally denying that the speaker has uttered the word 'mongeese'. But of course the speaker *does* utter that word in producing a token of (3). Another issue is that taking metalinguistic negation not to involve any claims about appropriateness makes it impossible to account for its intuitively most important function: that of correcting a prior utterance.

There is thus currently no quotational analysis of metalinguistic negation which is both fully semantically explicit and empirically plausible. It is this gap which my proposal fills.

9 Conclusion

We began with the observation that certain natural-language constructions have the flavor of mixed quotation — that is, of simultaneous use and mention — though they lack any overt indication that they are quotative. My aim has been to vindicate this intuition by constructing a theory of mixed quotation and showing how it can be used to analyze five apparently quite different phenomena in terms of the presence of a covert mixed quotation operator, sometimes in conjunction with other independently motivated covert material. The theory of mixed quotation I have presented above is novel in that it shows how the semantic properties of mixed quotations can be compositionally derived from the semantic properties of pure quotations. It is worth noting that the interest of this analysis, and of the theoretical program of accounting for various natural language constructions by positing covert mixed quotation, is independent of the plausibility of any of the particular applications of the theory suggested above. Whether or not those applications are ultimately vindicated, it is likely that a variety of puzzling semantic phenomena will be amenable to treatment in terms of covert mixed quotation.

Covert mixed quotation

Appendix: The language \mathcal{L}_{MQ}

1 Syntax

1.1 Types

- i. e^a, q^a, u^a, s^a , and t^a are the basic at-issue types.
- ii. t^p is the basic peripheral type.
- iii. If σ and τ are basic or derived types, then $\langle \sigma, \tau \rangle$ is a derived type.
- iv. χ^a is a (non-basic, non-derived) at-issue type.
- v. If σ and τ are types, then $\langle \sigma, \tau \rangle$ is a type.
- vi. The set of types is the smallest set containing the types described in (1)–(5).

Remark 1 *The basic and derived types form the subset of all types which do not contain χ^a , the type of the output of the function $(*)$. χ^a must be distinguished from the other types because the output of $(*)$ can in principle belong to the domain of any basic or derived type. (See Section 2.1 below.)*

1.2 Meaningful and vacuous expressions

Let ME_τ be the set of all meaningful expressions of type τ .

- i. If c is a constant of type τ , then $c \in ME_\tau$.
- ii. If x is a variable of type τ , then $x \in ME_\tau$.
- iii. If f is a variable of type $\langle q^a, \chi^a \rangle$, then $f^* \in ME_{\langle q^a, \chi^a \rangle}$.
- iv. If t is a variable of type t^a , then $\overleftarrow{t} \in ME_{t^a}$, and if t is a variable of type t^p , then $\overleftarrow{t} \in ME_{t^p}$.
- v. If $\alpha \in ME_{\langle \sigma, \tau \rangle}$ and $\beta \in ME_\sigma$, then $\alpha(\beta) \in ME_\tau$.
- vi. If $\alpha \in ME_\tau$, $\beta \in ME_{t^a}$, and x is a variable of type σ , then $\lambda x.\alpha \in ME_{\langle \sigma, \tau \rangle}$ and $\lambda x|\beta.\alpha \in ME_{\langle \sigma, \tau \rangle}$.
- vii. If $\alpha \in ME_\tau$, then $[\hat{\alpha}] \in ME_{\langle s^a, \tau \rangle}$.
- viii. If $\alpha, \beta \in ME_{t^a}$, then $\neg\alpha, \alpha \wedge \beta \in ME_{t^a}$.
- ix. If $\alpha, \beta \in ME_{t^p}$, then $\neg\alpha, \alpha \wedge \beta \in ME_{t^p}$.

- x. If $\alpha, \beta \in ME_\tau$, then $\alpha = \beta \in ME_{t^a}$.
- xi. If $\alpha \in ME_{t^a}$, and x is a variable, then $\exists x[\alpha] \in ME_{t^a}$.
- xii. If $\alpha \in ME_{t^p}$, and x is a variable, then $\exists x[\alpha] \in ME_{t^p}$.
- xiii. $w_c \in ME_{s^a}$.

Remark 2 The vertical stroke in $\lambda x|\beta.\alpha$ indicates that β is to be understood as a condition which restricts the domain of the function to those elements x which satisfy β .

The set ME of all meaningful expressions is the union of ME_τ for all τ .

Let Vac_τ be the set of all vacuous expressions of type τ .

- xiv. If x is a variable of type τ , then $\text{which}_x \in Vac_\tau$.

The set Vac of all semantically vacuous expressions is the union of Vac_τ for all τ .

Remark 3 Meaningful expression of \mathcal{L}_{MQ} are those which are interpreted by the interpretation function $\llbracket * \rrbracket^{c,w,s}$. Vacuous expressions are not interpreted, serving only to make it possible to incorporate predicate abstraction into the system.

1.3 Partial lexicon

Noteworthy \mathcal{L}_{MQ} constants:

- (\ast) (type $\langle q^a, \langle s^a, \langle e^a, \chi^a \rangle \rangle \rangle$).
- $(\ast)^*$ (type $\langle q^a, \langle e^a, \chi^a \rangle \rangle$).
- R (type $\langle e^a, \langle u^a, \langle q^a, t^p \rangle \rangle \rangle$).
- \blacklozenge (type $\langle \langle s^a, t^p \rangle, t^p \rangle$).

English translations (Overt items):

- Peter $\rightsquigarrow p$ (type e^a).
- tomato $\rightsquigarrow \lambda x.\text{tomato}(x)$ (type $\langle e^a, t^a \rangle$).
- $\text{which}_x \rightsquigarrow \text{which}_x$ (type of x).

Covert mixed quotation

English translations (Covert items):

- $\mathfrak{M} \rightsquigarrow \lambda q. \langle q \rangle (w_c)(s_x)$ (type $\langle q^a, \chi^a \rangle$) $\bullet \lambda q. R(s_x, u_x, q)$ (type $\langle q^a, t^p \rangle$).
- $\downarrow \rightsquigarrow \lambda p. \vec{p}$ (type $\langle t^p, t^a \rangle$).
- $\dagger \rightsquigarrow \lambda f | \exists s : f = \lambda q. \langle q \rangle (w_c)(s). f^*$ (type $\langle \langle q^a, \chi^a \rangle, \langle q^a, \chi^a \rangle \rangle$).
- $\mathfrak{A} \rightsquigarrow \blacklozenge$ (type $\langle \langle s^a, t^p \rangle, t^p \rangle$).

1.4 Parsetrees

1.4.1 Definition

A *parsetree* is a connected, rooted, acyclic directed graph containing no nodes with more than two daughters, to each node of which is assigned one or more meaningful or vacuous expressions of \mathcal{L}_{MQ} subject to the following admissibility conditions.

1.4.2 Admissibility conditions

Remark 4 *Five notes to improve legibility:*

- Expressions in this subsection are indicated along with their types. Expressions are separated from their types with a colon. Thus ‘ $\beta : \sigma^a$ ’ indicates a \mathcal{L}_{MQ} expression of an at-issue type.
- Material in parenthesis is optional; it may or may not be present.
- With the exception of ‘*which _{μ}* ’ in the Predicate Abstraction admissibility condition below, all metavariables should be understood to quantify over meaningful expressions only.
- Since the only type of peripheral content in the system is propositional, I have simplified the tree admissibility conditions accordingly.
- I use the bullet (\bullet) to indicate that a single node in a syntactic tree is associated with multiple formulas in \mathcal{L}_{MQ} . The bullet indicates only that more than one \mathcal{L}_{MQ} expression is associated with a node; it does not signify anything about whether the expressions in question belong to the at-issue dimension of meaning or the peripheral dimension of meaning.

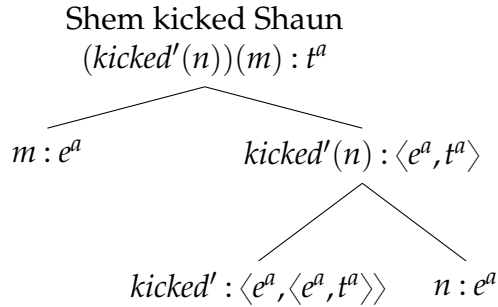
At-Issue Function Application

$$\alpha(\beta) : \tau^a \bullet (\gamma : t^p) \bullet (\delta : v^a) \bullet (\epsilon : t^p)$$

$$\alpha : \langle \sigma^a, \tau^a \rangle \bullet (\gamma : t^p) \quad \beta : \sigma^a \bullet (\delta : v^a) \bullet (\epsilon : t^p)$$

Remark 5 This rule corresponds to normal (at-issue) function application; it says that if a node has two daughters hosting expressions with at-issue types such that one (α) could take the other (β) as argument, that node hosts the complex expression $\alpha(\beta)$. In this case, any other content which might be hosted by the daughter nodes is simply passed up the tree to the mother node without modification.

Example 1 A simple example making use of At-Issue Function Application.



Shem $\rightsquigarrow m$ (type e^a).

Shaun $\rightsquigarrow n$ (type e^a).

kicked $\rightsquigarrow kicked'$ (type $\langle e^a, \langle e^a, t^a \rangle \rangle$)

Pure Quotation

$$\begin{array}{c}
 \gamma : q^a \\
 | \\
 \alpha : \sigma^a \bullet (\beta : t^p)
 \end{array}$$

(For γ a constant of type q^a picking out the English string syntactically associated with the daughter node.)

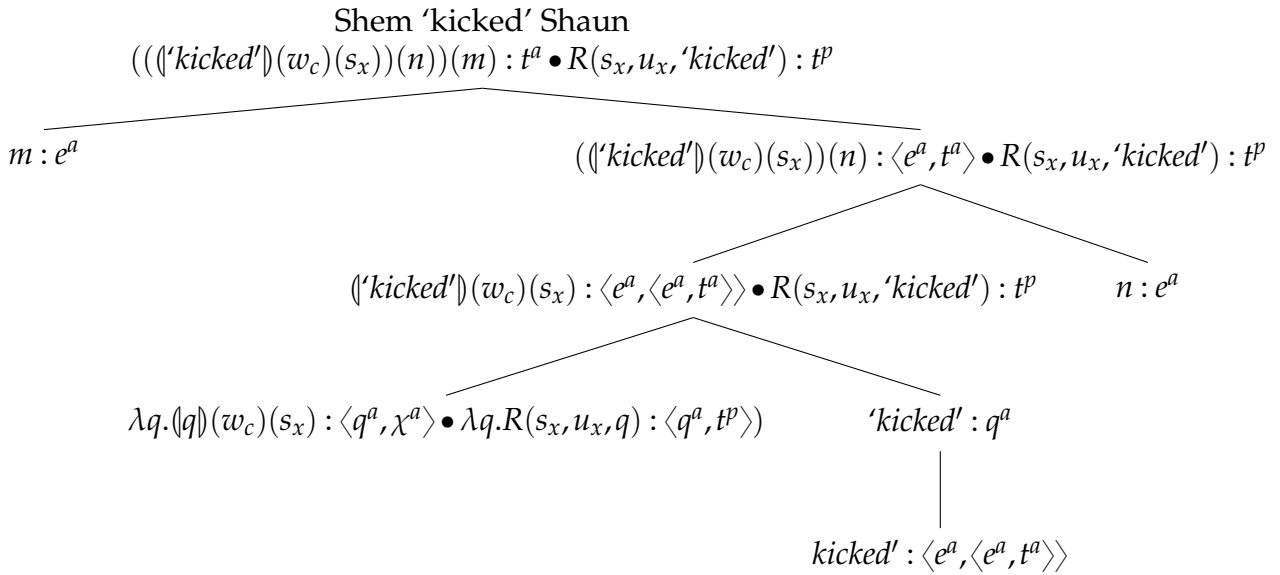
Remark 6 Pure quotation blocks the passage of peripheral content up the tree.

$$\begin{array}{c}
 \text{a-to-p Shunting} \\
 \alpha(\gamma) : \tau^a \bullet \beta(\gamma) : t^p \bullet (\delta : t^p) \bullet (\epsilon : t^p) \\
 \swarrow \quad \searrow \\
 \alpha : \langle \sigma^a, \tau^a \rangle \bullet \beta : \langle \sigma^a, t^p \rangle \bullet (\delta : t^p) \quad \gamma : \sigma^a \bullet (\epsilon : t^p)
 \end{array}$$

Covert mixed quotation

Remark 7 This rule characterizes the behavior of \mathfrak{M} : both of its at-issue meanings (α and β) operate on the at-issue meaning of its sister node (γ), while any other incidental peripheral content is passed up the tree without modification.


Example 2 An example involving mixed quotation. The speaker is attributing the choice of the word ‘kicked’ to some other source. β -reduction has been applied at some points for readability.



Shem $\rightsquigarrow m$ (type e^a).

Shaun $\rightsquigarrow n$ (type e^a).

kicked \rightsquigarrow kicked' (type $\langle e^a, \langle e^a, t^a \rangle \rangle$)



Remark 8 This rule characterizes the behavior of \downarrow , which is the only expression whose \mathcal{L}_{MQ} translation is of type $\langle t^p, t^a \rangle$. Note that the node which goes sister to \downarrow is required to have an at-issue propositional meaning as well as a peripheral propositional meaning — the two are then conjoined. This ensures that the presence of \downarrow will induce composition failure unless it occurs at the boundary of a clause. \downarrow can be iterated if it goes sister to a node with multiple peripheral contents of the

right type. In such cases, iterating \downarrow will incrementally conjoin these peripheral propositions with the at-issue content of its sister. In situations where the rule can be applied to more than one peripheral content, any may be shunted.

Peripheral Intensional Function Application

$$\begin{array}{c} \alpha([\hat{\beta}]) : t^p \bullet (\gamma : \sigma) \\ \swarrow \quad \searrow \\ \alpha : \langle \langle s^a, t^p \rangle, t^p \rangle \quad \beta : t^p \bullet (\gamma : \sigma) \end{array}$$

Remark 9 This rule, which is a version of the more familiar Intensional Function Application restricted to apply only in the peripheral dimension, characterizes the behavior of \mathfrak{A} .

'In The Sense That'

$$\begin{array}{c} \exists s_x[\lambda q. (q)(w_c)(s_x)(\beta) \wedge ([\hat{\lambda} q. (q)(w_c)(s)(\beta)] = \mu)] \\ \swarrow \quad \searrow \\ \begin{array}{l} \lambda q. (q)(w_c)(s_x)(\beta) : t^a \bullet \\ \lambda q. R(s_x, u_x, q)(\beta) : t^p \end{array} \quad \mu : \langle s^a, t^a \rangle \\ \swarrow \quad \searrow \\ \begin{array}{l} \lambda q. (q)(w_c)(s_x) : \langle q^a, \chi^a \rangle \bullet \\ \lambda q. R(s_x, u_x, q) : \langle q^a, t^p \rangle \end{array} \quad \beta : q^a \\ | \\ \alpha : t^a \end{array}$$

Remark 10 Here β should be understood as being derived from α using the Pure Quotation rule. The material which goes sister to β is simply the \mathcal{L}_{MQ} translation of \mathfrak{M} . On the right side of the tree, μ should be understood as the \mathcal{L}_{MQ} translation of an English prepositional phrase of the form 'in the sense that \mathcal{S}^\top ', which is assumed to have type $\langle s^a, t^a \rangle$ in accordance with the discussion in Section 7. Note that this rule eliminates the peripheral content introduced by \mathfrak{M} .

Predicate Abstraction

$$\begin{array}{c} \lambda \mu. \alpha : \langle \sigma, \tau^a \rangle \bullet (\beta : t^p) \\ \swarrow \quad \searrow \\ \text{which}_\mu : \sigma \quad \alpha : \tau^a \bullet (\beta : t^p) \end{array}$$

Covert mixed quotation

Remark 11 *The expression ‘which _{μ} ’ is a vacuous expression of \mathcal{L}_{MQ} : it is not interpreted by the function $\llbracket * \rrbracket$, serving instead merely to mark where Predicate Abstraction will apply.*

2 Semantics

2.1 Intended model

Let D be the intended set of domains, defined as follows:

- i. The domain of e^a is D_e , a set of entities.
- ii. The domain of u^a is D_u , a set of utterances, disjoint from D_e .
- iii. The domain of q^a is D_q , a set of expressions, disjoint from D_e .
- iv. The domain of s^a is D_s , a set of worlds, disjoint from D_e .
- v. The domain of t^a , D_{t^a} , and the domain of t^p , D_{t^p} , are identical to each other and to the set $\{0,1\}$.
- vi. If σ and τ are types, the domain $D_{\langle\sigma,\tau\rangle}$ of functional type $\langle\sigma,\tau\rangle$ is $\{f \mid f : D_\sigma \rightarrow D_\tau\}$.
- vii. The domain of χ^a , D_χ , is the union, for all basic and derived types τ , of D_τ .

Let C be the intended set of contexts, subject to the constraint that if $c \in C$, then there is a unique $w \in D_s$ such that c belongs to w .

Let \mathcal{V} be a function which maps each ordered pair of a \mathcal{L}_{MQ} constant and a world to the intended element of the corresponding domain at that world.

2.2 Interpretation of \mathcal{L}_{MQ}

Define the interpretation function $\llbracket \cdot \rrbracket^{c,w,g}$ over the set of meaningful expressions such that:

- If α is a constant of \mathcal{L}_{MQ} , then $\llbracket \alpha \rrbracket^{c,w,g} = \mathcal{V}(\langle\alpha,w\rangle)$.
- $\mathcal{V}(\langle\llbracket * \rrbracket, w\rangle)$: the function which maps an expression, a speaker, and a world w' to the extension χ at w of that expression as uttered by that speaker at w' .

- $\mathcal{V}(\langle \langle (*), w \rangle \rangle)$: the function which maps an expression and speaker to the extension χ at w of that expression as uttered by that speaker at w .
- $\mathcal{V}(\langle \langle R, w \rangle \rangle)$: the function which maps an individual, an utterance, and an expression to 1 at w iff that individual produces an utterance of that expression at w .
- $\mathcal{V}(\langle \langle \diamond, w \rangle \rangle)$: the function which maps an intension p of type $\langle s^a, t^p \rangle$ to 1 at w iff for every world w' in the set which characterizes what is or would be appropriate at w , $p(w') = 1$.
- If α is a variable of \mathcal{L}_{MQ} , then $\llbracket \alpha \rrbracket^{c,w,g} = g(\alpha)$.
- For α of type $\langle \sigma, \tau \rangle$ and β of type σ , $\llbracket \alpha(\beta) \rrbracket^{c,w,g} = \llbracket \alpha \rrbracket^{c,w,g}(\llbracket \beta \rrbracket^{c,w,g})$.
- For variable t of type t^a or t^p , $\llbracket \overline{t} \rrbracket^{c,w,g} = \llbracket t \rrbracket^{c,w,g}$.
- For x of type τ , $\llbracket \lambda x. \alpha \rrbracket^{c,w,g}$ = the function f such that, for every element d of D_τ , $f(d) = \llbracket \alpha \rrbracket^{c,w,g[x:=d]}$. ($g[x:=d]$ is the assignment function which maps x to d and every other variable y to $g(y)$.)
- For x of type τ , $\llbracket \lambda x | \beta. \alpha \rrbracket^{c,w,g}$ = the function f such that (i) f 's domain is the set of elements d of D_τ such that $\llbracket \beta \rrbracket^{c,w,g[x:=d]} = 1$, and (ii) for every element d of f 's domain, $f(d) = \llbracket \alpha \rrbracket^{c,w,g[x:=d]}$.
- For α of type τ , $\llbracket [\hat{\alpha}] \rrbracket^{c,w,g}$ = the function from D_s to D_τ which maps each world w' to $\llbracket \alpha \rrbracket^{c,w',g}$.
- For α, β of type τ , $\llbracket \alpha = \beta \rrbracket^{c,w,g} = 1$ iff $\llbracket \alpha \rrbracket^{c,w,g}$ is the same element of D_τ as $\llbracket \beta \rrbracket^{c,w,g}$.
- $\llbracket \neg \alpha \rrbracket^{c,w,g} = 1$ iff $\llbracket \alpha \rrbracket^{c,w,g} = 0$.
- $\llbracket \alpha \wedge \beta \rrbracket^{c,w,g} = 1$ iff $\llbracket \alpha \rrbracket^{c,w,g} = 1$ and $\llbracket \beta \rrbracket^{c,w,g} = 1$.
- If x is a variable of type σ , $\llbracket \exists x(\alpha) \rrbracket^{c,w,g} = 1$ iff there is some element d of D_σ such that $\llbracket \alpha \rrbracket^{c,w,g[x:=d]} = 1$.
- If $f^* \in ME_{\langle q^a, \chi^a \rangle}$, and if $\llbracket \exists s(f = \lambda q. \langle q \rangle (w_c)(s)) \rrbracket^{c,w,g} = 1$, then $\llbracket f^* \rrbracket^{c,w,g} = \llbracket \lambda q. \langle q \rangle^*(s') \rrbracket^{c,w,g[s':=d]}$, where d is any element of D_e such that $\llbracket f = \lambda q. \langle q \rangle (w_c)(s) \rrbracket^{c,w,g[s:=d]} = 1$. Otherwise $\llbracket f^* \rrbracket^{c,w,g}$ is undefined.
- $\llbracket w_c \rrbracket^{c,w,g} =$ the $w \in D_s$ such that c belongs to w .

References

- Burton-Roberts, Noel. 1989. On Horn's dilemma: Presupposition and negation. *Journal of Linguistics* 25(1). 95–125. <https://doi.org/10.1017/S0022226700012111>.
- Camp, Elisabeth. 2013. Slurring perspectives. *Analytic Philosophy* 54(3). 330–349. <https://doi.org/10.1111/phib.12022>.
- Camp, Elisabeth. 2018. A dual act analysis of slurs. In David Sosa (ed.), *Bad words: Philosophical perspectives on slurs*, 29–59. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780198758655.003.0003>.
- Cappelen, Herman & Ernest Lepore. 2007. *Language turned on itself: The semantics and pragmatics of metalinguistic discourse*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199231195.001.0001>.
- Carston, Robyn. 1996. Metalinguistic negation and echoic use. *Journal of Pragmatics* 25(3). 309–330. [https://doi.org/10.1016/0378-2166\(94\)00109-X](https://doi.org/10.1016/0378-2166(94)00109-X).
- Carston, Robyn. 2002. *Thoughts and utterances: The pragmatics of explicit communication*. Oxford: Blackwell Publishing. <https://doi.org/10.1002/9780470754603>.
- Davidson, Donald. 1979. Quotation. *Theory and decision* 11(1). 27–40. <https://doi.org/10.1007/BF00126690>.
- Einheuser, Iris. 2006. Counterconventional conditionals. *Philosophical Studies* 127. 459–482. <https://doi.org/10.1007/s11098-004-7790-5>.
- Fillmore, Charles J. 1969. Types of lexical information. In F. Kiefer (ed.), *Studies in syntax and semantics*, 109–137. Dordrecht: D. Reidel Publishing Company. https://doi.org/10.1007/978-94-010-1707-7_6.
- García-Carpintero, Manuel. 1994. Ostensive signs: Against the identity theory of quotation. *The Journal of Philosophy* 91(5). 253–264. <https://doi.org/10.2307/2940753>.
- Geurts, Bart. 2010. *Quantity implicatures*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511975158>.
- Geurts, Bart & Emar Maier. 2003. Quotation in context. *Belgian Journal of Linguistics* 17(1). 109–128. <https://doi.org/10.1075/bjl.17.07geu>.
- Harris, Jesse A & Christopher Potts. 2009. Perspective-shifting with appositives and expressives. *Linguistics and Philosophy* 32. 523–552. <https://doi.org/10.1007/s10988-010-9070-5>.
- Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Malden, MA: Blackwell Publishing.
- Horn, Laurence R. 1985. Metalinguistic negation and pragmatic ambiguity. *Language* 61(1). 121–174. <https://doi.org/10.2307/413423>.

- Horn, Laurence R. 1989. *A natural history of negation*. Chicago, IL: University of Chicago Press.
- Kirk-Giannini, Cameron Domenico. 2019. Slurs are directives. *Philosophers' Imprint* 19(48). 1–28.
- Kocurek, Alexander W, Ethan Jerzak & Rachel Etta Rudolph. 2020. Against conventional wisdom. *Philosophers' Imprint* 20(22). 1–27.
- Ludlow, Peter. 2008. Cheap contextualism. *Philosophical Issues* 18. 104–129. <https://doi.org/10.1111/j.1533-6077.2008.00140.x>.
- Maier, Emar. 2014a. Mixed quotation: The grammar of apparently transparent opacity. *Semantics and Pragmatics* 7(7). 1–67. <https://doi.org/10.3765/sp.7.7>.
- Maier, Emar. 2014b. Pure quotation. *Philosophy compass* 9(9). 615–630. <https://doi.org/10.1111/phc3.12149>.
- Maier, Emar. 2015. Quotation and unquotation in free indirect discourse. *Mind & Language* 30(3). 345–373. <https://doi.org/10.1111/mila.12083>.
- Maier, Emar. 2021. Mixed quotation. In Daniel Gutzmann, Lisa Matthewson, Cecile Meier, Hotze Rullmann & Thomas E. Zimmerman (eds.), *The Wiley Blackwell companion to semantics*, Malden, MA: Wiley-Blackwell. <https://doi.org/10.1002/9781118788516.sem080>.
- McCready, Elin. 2010. Varieties of conventional implicature. *Semantics and Pragmatics* 3(8). 1–57. <https://doi.org/10.3765/sp.3.8>.
- Noh, Eun-Ju. 2000. *Metarepresentation: A relevance-theory approach*. Amsterdam: John Benjamins Publishing Company. <https://doi.org/10.1075/pbns.69>.
- Plunkett, David & Tim Sundell. 2013. Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint* 13(23). 1–37.
- Portner, Paul. 2004. The semantics of imperatives within a theory of clause types. In *Proceedings of semantics and linguistic theory*, vol. 14, 235–252. Ithaca, NY: Cornell University Press. <https://doi.org/10.3765/salt.v14i0.2907>.
- Portner, Paul. 2007. Imperatives and modals. *Natural Language Semantics* 15. 351–383. <https://doi.org/10.1007/s11050-007-9022-y>.
- Potts, Christopher. 2005. *The logic of conventional implicatures*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199273829.001.0001>.
- Potts, Christopher. 2007. The dimensions of quotation. In Chris Barker & Pauline Jacobson (eds.), *Direct compositionality*, 405–431. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780199204373.003.0012>.
- Recanati, François. 2001. Open quotation. *Mind* 110(439). 637–687. <https://doi.org/10.1093/mind/110.439.637>.
- Saka, Paul. 1998. Quotation and the use–mention distinction. *Mind* 107(425). 113–135. <https://doi.org/10.1093/mind/107.425.113>.

Covert mixed quotation

- Seuren, Pieter AM. 1990. Burton-roberts on presupposition and negation. *Journal of Linguistics* 26(2). 425–453. <https://doi.org/10.1017/S0022226700014730>.
- Shan, Chung-chieh. 2010. The character of quotation. *Linguistics and Philosophy* 33. 417–443. <https://doi.org/10.1007/s10988-011-9085-6>.
- Stalnaker, Robert C. 1978. Assertion. In Peter Cole (ed.), *Pragmatics*, vol. 9 Syntax and Semantics, 315–332. New York, NY: Academic Press. https://doi.org/10.1163/9789004368873_013.
- Stanley, Jason & Zoltan Gendler Szabó. 2000. On quantifier domain restriction. *Mind & Language* 15(2-3). 219–261. <https://doi.org/10.1111/1468-0017.00130>.
- Stokke, Andreas. 2021. Protagonist projection, character focus, and mixed quotation. In Emar Maier & Andreas Stokke (eds.), *The language of fiction*, 379–404. Oxford University Press. <https://doi.org/10.1093/oso/9780198846376.003.0015>.
- Williamson, Timothy. 2009. Reference, inference, and the semantics of pejoratives. In Joseph Almog & Paolo Leonardi (eds.), *The philosophy of David Kaplan*, 137–158. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195367881.003.0009>.

Cameron Domenico Kirk-Giannini
 Department of Philosophy
 Rutgers University–Newark
 175 University Ave
 Newark, NJ 07102
camerondomenico.kirkgiannini@gmail.com