Indirect Reports and Pragmatics
Nellie Wieland

Abstract: An indirect report typically takes the form of a speaker using the locution "said that" to report an earlier utterance. In what follows, I introduce the principal philosophical and pragmatic points of interest in the study of indirect reports, including the extent to which context sensitivity affects the content of an indirect report, the constraints on the substitution of co-referential terms in reports, the extent of felicitous paraphrase and translation, the way in which indirect reports are opaque, and the use of indirect reports as pragmatic vehicles for other speech acts such as humor, insult, or irony. Throughout I develop several positions: (i) that a semantic analysis of indirect reports is insufficient, (ii) that the distinction between direct and indirect reports is not clear and that indirect reports are the predominate way of reporting while direct reports may be a para-linguistic variation on them, (iii) that most questions about the semantics and pragmatics of indirect reports will rely on a full understanding of the nature of what is reported and how it gets reported, (iv) that an analysis of reporting requires the pragmatic tools of metarepresentation and a social, inter-personal understanding of relevance and shared knowledge.

1. Introduction: Direct and indirect reports

The philosophical literature is rich with discussion of quotation and indirect reports. It has played a starring role in the philosophy of language at least since Donald Davidson's (1968) "On Saying That." Much of this literature has concluded that the analysis of indirect reports should be largely pragmatic. Unfortunately, there has been a scarcity of accounts of what a pragmatic theory of indirect reports should look like, how, if at all, it should complement a semantic account of indirect reports, and how such a pragmatic account can respond to some of the most interesting philosophical questions about reporting. Whereas quotations are used to directly report the utterances of another, the indirect report has far more latitude. The result is that indirect reports have most of the same features that make direct reports so philosophically interesting and many unique features that make them especially pragmatically interesting.

An indirect report typically takes the form of a speaker using the locution "said that" to report on an earlier utterance (e.g. "Galileo said that the earth moves"). The original utterance could come from the speaker herself or a third party; the contexts of the original and reported utterances could be identical or distinct. Indirect reports differ from direct reports in that they typically do not use quotation (in written form with quotation marks or in spoken form with implied quotation marks) (e.g., "Galileo said, 'the earth moves'"). Or, if they do, they take the form of mixed-quotation where part of the report is direct and part indirect (e.g., "Galileo said that the earth 'moves'").

In one notable philosophical discussion of indirect reports, Cappelen & Lepore (1997b) argue that semanticists will have little to say about the most linguistically interesting features of indirect reports. The reason for this is that semanticists are interested in the most general and systematic features of linguistic expressions, whereas speakers, when they report, are interested in reporting something narrow and particular about a context and an utterance in that context. It is no surprise to Cappelen & Lepore that the aims of a semanticist in analyzing 'to say' and a reporter in using 'said that' differ. It's worth considering their claim at length:

> [Semanticists] tend to agree that semantics is a discipline that aims to characterize systematically certain features of linguistic expressions and to do so in a way that captures general truths about our linguistic practice, not just truths about particular speakers in particular contexts. Indirect speech, on the other hand, is a device reporters use for characterizing acts (utterances) performed by other speakers. In so doing, reporters are interested neither in systematicity nor in generality; they aim to convey something about a particular act in a particular context C to a particular audience situated in a different context C*. Reporters draw on information about the specific intentions of, knowledge about, and the history of a reported speaker in C and (maybe) similar features of an audience in C*. These are features one does *not* want to solicit when the aim is systematic and general.[1]

The present discussion both accepts and rejects Cappelen & Lepore's central claims in this passage. I will accept that the aims of the semanticist and the speaker differ, and I

---

[1] Cappelen & Lepore (1997b), pg. 289. Note that Cappelen & Lepore's aim is totally different from the aim of this paper. They say that it is *because* reporters are doing these things that we should not expect a semantic analysis to capture reports. Their goal is to protect semantics from what look like threatening examples from indirect reports. They acknowledge the semanticist's inability to explain indirect reports, but maintain that this is because reporting is something that speakers do for non-semantic reasons. They do not offer or entertain a full-fledged pragmatic theory to explain the relevant linguistic features of indirect reports.

will accept that the semanticist does not have a comprehensive analysis of 'to say' available to her. But I will reject that there are no general and systematic claims that can be made about the pragmatics of indirect reports. Although speakers are attempting to "convey something about a particular act in a particular context C to a particular audience situated in a different context C*" when they indirectly report, they still do so in systematic ways, and under identifiable constraints.

In what follows, I introduce the principal philosophical and pragmatic points of interest in the study of indirect reports. I will argue that indirect reports are more pragmatically complicated than direct reports. They lack explicit, conventional markings in written text (such as italics or quotation marks), and they permit a great deal of communicative deviation. Presumably, a felicitous indirect report needs to convey the content communicated in the original utterance—but there is a fair amount of reasonable disagreement as to what this might constitute. Felicitous indirect reports allow for the substitution of co-referential terms, context shifting, translation across languages, paraphrase, reporting by inference, and the use of demonstratives to stand in for utterances. Indirect reports also work as standard pragmatic vehicles for humor, irony, implicature, insult, praise, etc.

I will discuss several possible pragmatic approaches to the analysis of reports. In particular, I will discuss the ways in which reporting practices are similar to other forms of metarepresentation. I also assess the analyses available depending on whether language is conceived of as a verbal or written phenomenon. Philosophers have tended to analyze language in its written form, largely ignoring verbal communication. Linguists tend to do the opposite. Given that direct reports are explicitly marked only in written texts, we can see that a verbal orientation would suggest that analyses of direct and indirect reports should largely be the same. If this is the case, then much of the philosophical literature on quotation should be re-assessed.

2. Puzzles about indirect reports

2.1.Indirect reports and the use/mention distinction

It is standard in philosophical discussions of quotation to distinguish between the use and mention of a term or expression. The idea is straightforward on its face: some language is used meta-linguistically to talk about other language. A speaker can say (1) and (2) coherently and felicitously:

(1) Boston is a large American city.
(2) 'Boston' has six letters.

In (1) the term 'Boston' is being used; in (2) it is being mentioned. One is not saying in (2) that the city itself has six letters, but rather that the name of the city is comprised of six letters. (2) is a claim about language; (1) is a claim about the world.

Given this distinction it is natural to think of reports as special, meta-linguistic uses of language since they involve mentioning the speech of others.[2] Consider direct reports:

(3) The President said, "Osama bin Laden resisted."

In this case, the speaker of (3) is directly reporting the President's speech; the use of quotation marks indicates that the speaker of (3) is merely mentioning and not using the words "Osama bin Laden resisted." Now consider a case of an indirect report:

(4) The President said that Osama bin Laden resisted.

Direct written reports employ an arguably semantic tool[3] to mark them as such, i.e., quotation marks, italics, or angle brackets. Indirect reports have a less explicit tool: the phrase 'said that' or its equivalent. One possibility to consider is that the phrase 'said that' has a similar function to written quotation marks insofar as it indicates that the words to follow are being mentioned and not used. As we will we see in what follows, this is a

---

[2] It is not widely agreed upon that quoting is equivalent to mentioning. It is clear that there are forms of mentioning that do not involve quoting. This leaves open the possibility that quoting is a subspecies of mentioning. Some quotation theorists (see Saka (1998) and Washington (1992)) argue that the analyses of quotation and mentioning will be closely related, if not the same. Others (see Cappelen & Lepore (2009)) disagree that quotation and mentioning correspond in this way.

[3] In the quotation literature it is not universally accepted that quotation marks (and their corollaries) are semantic devices. I am glossing over these disagreements here.

tough case to make given the varieties of indirect reports. The use/mention distinction is typically directed at strings of words and not at linguistic content.

Could there be an alternative to the use/mention distinction for indirect reports? In this case, there would be a distinction between a speaker using content (e.g., the communication of linguistic content that takes place in the original utterance) and a speaker mentioning content (e.g., the reporting of linguistic content that takes place in the indirect report). Consider (5).

(5)   Osama bin Laden resisted.

When the President utters (5), he is using this expression in order to mean something or other. When a reporter utters (4), she is merely mentioning whatever it was that that the President meant in uttering (5), but is not attempting to mean whatever it was that the President himself meant with his utterance of (5). The reporter in (4) distances herself from the meaning conveyed by the President given that her speech act is one of *reporting* and not one of *describing* what unfolded in Osama bin Laden's compound.[4] This becomes clearer when we consider another case of indirectly reporting (5).

(6)   The President said that the leader of al-Qaeda resisted.

In this case, the speaker is using a co-referential substitution for the name 'Osama bin Laden' (see section 2.4). In doing so the speaker is mentioning—in a manner of speaking —what the President communicated, although not his exact locution. Admittedly, there is a bit of clumsiness in applying the standard use/mention distinction to indirect reports in this way. But this should highlight that there are interesting questions to be asked about what exactly is being reported in an indirect report as well as what is doing the reporting (see section 2.9).

2.2. Indirect reports and context sensitivity

---

[4] Indirect reports have thus been described as containing "two voices" and as being micro-narrations. In reporting, the speaker is using her own voice in order to convey the voice of another (or from another context); in doing so she is establishing a kind of narrative with the locution "said that" that presupposes both common knowledge between the reporter and audience of the reporter's goal, and common knowledge of relevant features of the context of the original utterance. See Capone (2010).

Indirect reports are linguistic objects that can convey content from one context to another. They need not always do this (e.g., consider the case of a speaker reporting on her immediately prior utterance, "I said that we're going to be late!"). When reports do convey content from context to context, it is essential that relevant features of the reported context are assumed in the reporting context. Presumptively, these features of the reported context would be constituted in shared knowledge.

Given this, we can say that indirect reports are themselves context sensitive; a complete understanding of the content of the report requires knowing *something* about the context in which it was uttered. It has also been argued that indirect reports themselves provide a test for context-sensitivity.[5] An aspect of the debate on the appropriate divide between pragmatics and semantics has been the extent to which the meanings of terms are sensitive to context. Another way of thinking about this is that if an expression is identified as 'context-sensitive', then its meaning is underdetermined without knowledge of context. As an example, consider a clearly context-sensitive expression such as 'today' embedded in an indirect report:

(7)   Marylee said that she will arrive today.

As reported, the content of Marylee's original utterance is underdetermined. It depends, at the very least, on the day of her original utterance. (7) is a felicitous report if Marylee's utterance of "I will arrive today" occurred on the same day as the utterance of (7). Likewise, (7) is a felicitous report if Marylee's utterance of "I will arrive tomorrow" occurred the day before the utterance of (7), etc. The way in which indirect reports can test for the context-sensitivity of an expression is by seeing if it survives being reported (without quotation marks but with the same actual expression) in a new context. We can see that obvious indexicals such as 'I', 'you', 'he', 'she', 'today', and 'tomorrow', and demonstratives such as 'this' and 'that' do not survive this test and so must be context-sensitive. On the other hand, other candidates for context-sensitivity such as gradable adjectives do pass this test—at least at first glance. For example, consider an expression

_____

[5] See Cappelen & Lepore (2004).

such as 'is tall'. Consider whether, when embedded in an indirect report, this expression compromises the felicity of said report:

(8)   She said that her niece is tall.

It's arguable that an utterance of "My niece is tall" by the original speaker contains the context-sensitive expression 'is tall'. After all, there are many ways of being tall, and many comparison classes against which height can be measured (e.g., tall for a seven year old, tall for a WNBA player). Given this context sensitivity, we might expect (8) to be underdetermined in the same way that (7) is. Yet, this is not obviously the case.[6] This tells us that indirect reports themselves might provide a tool for determining the extent to which language is context-sensitive, and thereby provide insight to the boundary between semantics and pragmatics.

2.3. Indirect reports and inter-substitutability of co-referential terms

Co-referential terms and expressions can be substituted for one another in indirect reports felicitously. For example:

(9)
    A: This morning I had pancakes, toast, and coffee.
    B: *A* said that she had breakfast this morning.

As we can see, inter-substitutable co-referential terms in felicitous indirect reports fail strict semantic analysis and yet seem to reflect our practices of indirect reporting. Consider this *co-referential substitution*:

(10)
    A: My favorite *tapa* is *patatas bravas*.
    B: *A* said that her favorite *tapa* is the third item on your menu.

In this case, the term '*patatas bravas*' is substituted with a definite description with a value that can only be determined in the reporting context. It would be implausible to suggest that the original speaker meant anything like *the third item on your menu* in the original context of utterance. Nevertheless, ordinary reporting practices take advantage of

---

[6] See Wieland (2010a) for an opposing view.

this sort of inter-substitution. In other cases, *co-referential substitution* will fail. Consider this case:

> (11)
>> *A*: The word 'mendacious' means *lying*.
>> *B*: *A* said that the word 'lying' means *lying*.

While 'mendacious' and 'lying' are semantically equivalent and thereby intersubstitutable, it would be infelicitous to report *A*'s utterance in this way. (The reason for this is partly because (11) involves the use of quotation marks, marking this as a mixed report. See section 2.6.) Co-referential substitution depends largely on maintaining the propositional content of the original report. But indirect reports do not rely on strict propositional preservation; in everyday felicitous reporting contexts, propositional content varies so much that it casts into doubt the idea that reports of any kind *require* such preservation.

2.4. Indirect reports and propositional content

Co-referential substitution is not the only way for an indirect report to felicitously modify an original utterance. An examination of the varieties of indirect reports suggests that indirect reports need not represent the propositional content of the original utterance in its entirety or without alteration. This, perhaps, is the primary reason why a semantic analysis of indirect reports is of limited utility in comparison to a pragmatic account. For example, consider what we can call *conjunction elimination*:[7]

> (12)
>> *A*: I went to the taco stand and bought a soda.
>> *B*: *A* said that she went to the taco stand.

The propositional content of *A*'s utterance minimally includes two conjuncts. *B*'s report of *A*'s utterance leaves out one of the conjuncts, so *A*'s utterance and *B*'s report do not contain the same propositional content. Consider another case. Call this *modifier elimination*:

> (13)
>> *A*: I had some delicious nachos for dinner.

---

[7] See Cappelen & Lepore, (1997b), pg. 282. They refer to these and some of the following examples as those with "partial semantic overlap" and suggest that such examples will require a pragmatic analysis.

> *B*: *A* said that she had nachos for dinner.

This is also a case where the propositional content does not survive from *A*'s utterance to *B*'s report. In cases such as these, the felicity of a report that alters a modifier cannot be determined along semantic lines. To see this, consider a related possibility, call it *modifier introduction*:

> (14)
>> *A*: I met a woman at the party.
>> **B*: *A* said that she met a beautiful woman at the party.

Most (but not all) of the time, *B*'s report will not be felicitous. Similarly, not all cases of *modifier elimination* will be pragmatically felicitous. Consider:

> (15)
>> *A*: I had some low-fat nachos for lunch.
>> *B*: Did *A* have anything healthy to eat today?
>> **C*: *A* said that she had nachos for lunch.

Some *modifier eliminations* and *modifier introductions* alter the original utterance in a pragmatically infelicitous way and some do not. These are governed by pragmatic constraints on relevance and not semantic rules (or so I argue in the final section of this paper).

Another example of how an indirect report can modify original propositional content is when the report involves an *inference* based on the content of the reported utterance and knowledge of context. Consider this example of an *inferential indirect report*:[8]

> (16)
>> *A*: I didn't fail any students.
>> *B*: Professor *A* said Maryanne passed her exam.

Just as long as *B* knows that Maryanne is one of *A*'s students, then *B* can felicitously report *A*'s utterance in this way. The fact that the intersubstitutability of co-referential terms and paraphrase on the basis of inference are not only possible but commonplace suggests that an indirect report does not function to replicate the original utterance, and it does not even function to convey content that is identical to the original utterance, but rather its pragmatic function is to convey whatever is relevant about the original utterance

---

[8] This example is from Cappelen & Lepore, (1997b), p. 285.

to the reporter and audience given new facts about the reporting context. Relevant features include the word choice of the original speaker, propositional content of the original utterance, intonation, stress, or accent of the original utterance, the broad structure or word order of the original utterance, etc. Similarly, facts about the reporting context are just as relevant to the nature of the report. For example, if the original utterance is in French, and the audience only speaks Italian, the report will be a translation.

Finally, one further category of report to consider is the implicature. Consider the following report:

(17)
    A: It's awfully stuffy in here.
    ?*B* [to *C*]: *A* said that she would like for you to open a window.

What is interesting about this report is that it affords a natural felicitous reading and an equally natural infelicitous reading. The implicature *that* A *would like for you to open a window* is natural and readily available to *B*. It's less clear that it's felicitous to say that *A* said such a thing. This case differs from the *inferential indirect report* because while speaker *A* did not utter the words attributed to her by *B* in either case, in an inferential indirect report, *B*'s report of what *A* said logically follows from what *A* said (and other facts known in context). It would be reasonable for *B* to object *that she did not say this* in the case of a reported implicature. For either of these cases, reading them as felicitous requires a robustly pragmatic interpretation of 'said that'—one that many would not be willing to accept.

2.5. Are indirect reports opaque?

Given the above discussion, we can address another corollary between indirect and direct reports. Direct reports are typically described in the philosophical literature as 'opaque'. This means that when synonymous expressions are substituted into the original utterance in quotation, the quotation can fail to preserve truth-value. Given the discussion in sections 2.3 and 2.4, we can see that indirect reports are opaque – *some* of the time. This

suggests that the locution 'said that' and quotation marks function similarly. However, opacity itself is tricky to identify. An expression is said to be opaque if exchanging synonymous terms alters the truth-value. A standard example looks like this:

(18) 'Bachelor' has eight letters.
(19) 'Unmarried man' has eight letters.

The substitution of 'unmarried man' for 'bachelor' results in a failure to preserve truth-value in the move from (18) to (19). But as we know from examples (9) and (10), not all substitutions result in a failure to preserve truth-value. (The preservation of truth-value in examples (9) and (10) is not merely coincidental but is, in part, a consequence of the flexibility inherent in indirect reports.) However, consider the following indirect report:

(20)
A: Superman can fly.
?*B*: *A* said that Clark Kent can fly.

Is this report felicitous? The answer here is not obvious. On one hand, it is plausible to say that *B*'s report in (20) is not felicitous because, at the very least, *A* would not accept *B*'s paraphrase of what she said.[9] Moreover, it would be false to ascribe the belief that *Clark Kent can fly* to *A* because, despite the co-reference of 'Superman' and 'Clark Kent', *A* does not believe this.[10] Although *B*'s report has the same propositional content as *A*'s original utterance, the report is nonetheless false because it does not reflect what *A* believes. Given this, it looks as if indirect reports are in fact opaque.

On the other hand, *B* is not reporting what *A* believes; *B* is reporting what *A* said. It seems plausible that in the context of *B*'s utterance, *B* and her audience know that Superman = Clark Kent. They might know about Clark Kent's secret identity *and* know that *A* does not possess such knowledge. Imagine for a moment that, in the reporting context, it is common knowledge that Superman = Clark Kent, but the extent of Superman's powers is not common knowledge. Under such a description, *B* might be

---

[9] Cf. Capone (2010).

[10] There is a substantial literature on belief ascription given co-referential substitution. For a good discussion please see Jennifer Saul (1998).

making a reasonably informative report. This suggests that indirect reports are *not* opaque.

The discrepancies in indirect reports prompt us to conclude that the opacity of indirect reports cannot be semantically analyzed in the same way as direct reports. Whether or not indirect reports are opaque depends on pragmatic principles and contextual features. In particular, it depends on the common knowledge that exists between reporter and audience and the communicative goal that the reporter has.

2.6.Mixed reports

It is difficult to get far in a survey of indirect reports without recognizing that a great deal of reporting takes the form of a mixed report, or a mixed quotation. The analysis of mixed reports borrows from analyses of quotation and analyses of indirect reports. Consider this standard example:

(21) Quine said that quotation "has a certain anomalous feature."

One might claim that a mixed report such as (21) does not have the same pragmatic latitude as a pure indirect report. Further, one might claim that (21) requires, at the very least, that Quine used the particular locution, "has a certain anomalous feature," and not just that Quine uttered an identical speech act type with translatable propositional content. Is a pragmatic account committed to denying this?[11] This is an interesting question, and the answer to it turns on *what* is being reported in any given report—direct or indirect (see Section 2.9). It seems arbitrary to presuppose that quotation marks are supposed to capture locutional accuracy rather than accuracy of content, motivation, tone, or purpose. What is being aimed for in any given report might depend on the conversational goals of the reporter and audience. This is not to deny that locutional accuracy is not a normative standard in many settings (e.g., journalism or academic papers). This is compatible with it being far less normative in everyday discourse,

---

[11] As suggested by Cappelen & Lepore (2009).

particularly when that discourse is verbal (see Section 3.1) and the quotation marks are implied, gestured, or marked in some other way.

While mixed reports might be particularly semantically troublesome, they are less so under a pragmatic analysis. The quoted phrase indicates that the reporter is establishing distance between herself and the original utterance (i.e., she is mentioning and not using this utterance), and that she is doing so in such a way as to establish that her report is accurate along some dimension relevant to the reporting context. If the reporting context is journalistic or academic, it is very likely that her report will be true only if what falls between the quotation marks is an identical locution to that which was uttered by the original speaker. Given that the starting point of a pragmatic analysis concerns the conversational goals of the speakers, and given that the principal pragmatic measure of reports is felicity (and not truth per se), the mixed report is analyzed as any report would be. The reading of quotation marks under this analysis of indirect reports is just that they draw special attention to (or distance from) some portion of the report. The reporter has already established that she is not using this language (by way of the phrase 'said that') and, further, that she is using quotation marks to re-affirm that point with respect to a particular passage.

This way of interpreting the mixed report makes less of the differences between the direct report, the mixed report, and the indirect report. Quotation in a mixed report is a means of establishing both accuracy and distance with respect to the original report. But quotation itself can operate on various levels, from locution to illocution and even to perlocution. A pragmatic analysis should be ready to recognize this. (See section 2.9.)

2.7. Indirect reports as social objects

The analyses offered in 2.4, 2.5, and 2.6 of propositional content in reports, opacity, and quotation marks are undoubtedly unconventional. Many of the claims are defiant of a straightforward semantic analysis of the phrase 'said that'. And yet they may be most appropriate for a pragmatic perspective on these phenomena. One of the reasons is that pragmatics, as a discipline, is disinclined (relative to semantics as a discipline) to concern

itself with the utterance in isolation. It is far more concerned with what has variously been called the *context* or the *situation* of utterance. And contexts and situations are usually social events. They involve a reporter, an audience, a place, goals, and past conversation. They also involve cognitive mechanisms, loads, limits, and predispositions. The pragmatist is more likely to situate the report in the reporting context with a simultaneous representation of the original situated speech act.

On the face of it, this seems like an appropriate approach for the analysis of indirect reports. After all, they explicitly involve the recreation and representation of what someone else said, and often what they meant, why they said it, and how they said it, all while maintaining the report's relevance to the current speech situation. As a speech act, it's explicitly intercontextual and metarepresentational. The challenge is to maintain accuracy of description of the speech act itself while maintaining cognitive, empirical plausibility. Talk of *context* and *situation* and multi-level representations is widespread.[12] Yet I do not think there is a decisive account of these terms that maintains empirical plausibility—for example, the account of metarepresentation needs to characterize the semantic and pragmatic function of a report without proposing a structure of representation that is implausibly complicated. (See section 3.2 for further discussion.)

2.8. As pragmatic vehicles: humor, irony, insult, praise

Indirect reports are used as a number of different kinds of pragmatic vehicles. They are also affected by the non-literal content of the original speech act. Irony presents a clear case for the need for a pragmatic account of indirect reports. Consider this case:[13]

> (22) After a really bad philosophy talk, *A* says:
> *A*: That was, like, *really* good.
> *B₁: A said that the talk was really good.
> *B₂*: A said that he didn't like the talk much.

---

[12] For an elaborate account of the semantic and pragmatic details of metarepresentation and indirect discourse, see Recanati (2000). For an incisive criticism of this attempt, see Ludwig (2003).

[13] Originally in Larson, R. & Segal, G. (1995). Quoted in Cappelen & Lepore, (1997b), pg. 284.

Even if we might say that B_1 literally reports what A said, we recognize that it is an infelicitous report of A's utterance since A was using irony. A reasonable constraint on the practice of reporting is that consistency of speech act type be maintained.[14] If an utterance is meant to be ironic, the report is infelicitous if it reports that utterance literally. For example, Romeo says that Juliet is the sun, and Mercutio reports: "Romeo said that Juliet is the sun. But she's not a giant ball of fire. So Romeo said something false." Mercutio's report is clearly infelicitous because it fails to treat Romeo's original utterance as figurative. Likewise, Mercutio's conclusion that Romeo has said something false is also clearly wrong. If a speaker uses an interrogative and is reported as having made an assertion, the report is, again, infelicitous. (Such reports could be made felicitous by substituting 'said that' with the appropriate term, 'asked whether', 'joked that', etc.)

However, some indirect reports use the act of reporting in order to insult, praise, or perform other speech acts. In doing so, the reporter does not maintain speech act consistency with the original utterance. It's arguable that if the reporter did so, the report would not serve its purpose—namely, to insult, praise, etc. Consider this use of humor, insult, and scorn in the report but not in the original speech act.

(23)

> *A*: I voted for Sarah Palin.
> *B*: *A* said that she voted for Caribou Barbie.

Now there's a clear sense in which A did *not* say that she voted for Caribou Barbie—after all, there is no such person. But there are also easily recognizable contexts in which B and her audience recognize that 'Sarah Palin' and 'Caribou Barbie' are co-referential even if 'Caribou Barbie' is not the conventional name for Sarah Palin (rather, it's a mildly insulting moniker). Is this just like any other case of co-referential substitution? A could readily object that she did not say that she voted for 'Caribou Barbie', especially since she likely finds that name insulting. This example suggests that there is a general constraint that consistency of speech act content be maintained; but it is the violation of

---

[14] See Capone (2010).

this constraint that, in part, makes *B*'s utterance humorous or insulting. It is the very recognition of constraints on such reporting practices that makes these kinds of non-literal reports possible through pragmatic manipulation.

There is a related phenomenon in the use and analysis of slurs. (Note that (23) is not an example of a slur, but rather an insult or expression of scorn and contempt.)[15] There is reasonable debate as to whether a report of a slurring expression is itself offensive—or *as* offensive, or offensive in the same way as the original use of the slur. If a speaker were to report the slurring speech of another, it's reasonable to think that the speech is being used mentioned and not used (see section 2.1). As such the reporter is not slurring in reporting the use of a slurring term. Speakers often resist this, however. In reporting on the surprising use of a slurring term, speakers might say, "*A* said that she is a —I can't bring myself to use the word," or "A said that he is an 'n-word'"—in each case the speaker is unwilling to even repeat the slur in question for fear of using the slur inadvertently. This may be explained by the possibility that some speakers regard slurring terms as non-referring. If they report the original utterance, using the slurring term, the speakers may regard this is an endorsement of the slurring term—at the very least as an endorsement of its referential status. This maintains the insight that reports are acts of mentioning and not using; but simultaneously explains why reporters are hesitant to even report on others' slurs.

2.9. What is an indirect report reporting?

As discussed in section 2.1, indirect reports differ from direct reports at least insofar as they lack conventional markings in written texts. (Some have suggested that implied quotation marks are essential features of verbal direct reports as well.) It is also clear that the expression 'said that' and its corollaries mark a report and establish pragmatic distance between the current speaker and the reported speech. This presents two related questions: what is doing the reporting and what is being reported?

---

[15] Anderson & Lepore, 'Slurring words'.

The literature on quotation might be informative here, although it is not going to answer our questions directly. In the analysis of quotation, some have claimed that whatever is inside of quotation marks is a name (or is being named).[16] As such, the quoted expression is unanalyzable. This improbable view of quotation has been widely rejected. Another view of quotation treats quotation as a description of an earlier utterance.[17] Both the name view and the description view (as we might call them) tell us more about the report (or quote) than they do about what the report (or quote) is operating on. Neither give us clear account of what is being reported, whether locution, propositional content, speech act, etc. Neither the name view nor the descriptive view are defended as plausible accounts of quotation anymore.

No clear winner has dominated the contemporary quotation literature, but there are three major contenders: the demonstrative theory, the disquotational theory, and the use/identity theory.[18] Although each of these views primarily provides a semantic analysis of quotation marks themselves, we can look to them to get an answer to the question of this section as well. The demonstrative theory tells us that quotation marks are demonstratives referring to whatever token expression is between them. Without getting into the details of this view let me just point out one pertinent feature. Under this view, the quotation marks demonstrate *something* about the expression token, but don't indicate what that *something* is. This is a semantic weakness of the view in that it underspecifies a general rule for the use of quotation marks. But it is a pragmatic strength of the view in that it ties the *something* to the goals of the speaker, the common knowledge between speaker and audience, the relevant features salient in the reporting context, and so on. The use/identity theory of quotation also has an answer to the question of this section. It

---

[16] See Tarski (1933) and Quine (1940, 1961).

[17] See Geach (1957).

[18] For the classic defense of the demonstrative theory (also called the paratactic theory), see Davidson (1979) and Seymour (1994). Accounts of the disquotational theory can be found in Richard (1986) and Ludwig and Ray (1998). Accounts of the use/identity theory can be found in Washington (1992), Saka (1998), Reimer (1996), and Recanati (2000, 2001). An overview and analysis of all of these views can be found in Cappelen & Lepore (2009).

claims that by using quotation marks (implicitly or explicitly), a speaker mentions the same token being uttered between the quotation marks. (As this view has developed, its defenders have claimed that it might also mention a type or a shape of an expression, or even something more loosely related.) The central claim of this theory is that quotation marks (and perhaps all reports) just serve to indicate that mentioning and not using is taking place.

In the case of indirect reports specifically, the problem is slightly different. Given that there is no claim to direct reporting, it is unnecessary to suppose that the report looks like or names the original utterance in some exact way, or that what follows the locution 'said that' is disquotationally equivalent to the original utterance. In fact, indirect reports don't function in that way at all. Indirect reports shift with context (see section 2.2), allow for paraphrase and inter-substitution of co-referential terms (see section 2.4), and are not bothered by translation. Given this, the direct report is only a "picture" of the original utterance if one squints.

So what is the locution 'said that' reporting? The literature on quotation has been focused on the locution itself—quotes somehow refer to a series of words in a particular order—because this is a way of thinking about quotation that strongly resonates. Work on indirect reports has been less concerned with the locution and more concerned with the content and the speech act that delivers that content. If we think of what is being reported as a contentful speech act, there are still a number of possibilities for how to think of the report. For example, the report might be of the type instantiated by the original token expression, a type otherwise related to the original token (through paraphrase or inference, perhaps), some other token related to the reported token, or the original token itself. It could be that an indirect report is less directed toward conveying a contentful speech act than in the concepts contained in the reported expression, or some other possibility. It's not altogether clear that we can choose only one of these options, since what is reported may be a representation of the original utterance either in form (broadly speaking) or content (broadly speaking), depending on the reporter's goals. The great diversity among reports leaves the question of this section far more open than it has been

treated in the literature on quotation. For example, indirect reports, in their pragmatic guise, can be used to report on the tone of voice, accent, or attitude of the original speaker. The goal in making the report may not be accurate representation, but humor, scorn, irony, etc. (see section 2.8). In these cases, propositional content and replication of the original locution may be insignificant.

Some have argued that indirect reports must always report the minimal propositional content of the original utterance.[19] This is to say that no matter what the report is *doing* in terms of the conversational goals of the reporter or audience expectations, the report is only a felicitous report if minimal propositional content is preserved in the reporting. Consider the following exchange:

(24)
　　A: The cat is on the mat.
　　B: *A* said that the cat is on the mat.

Whatever *B* is doing by reporting *A*'s utterance, she reports that *A* said *that the cat is on the mat* and *B*'s report is true iff *A* said *that the cat is on the mat*. This seems plausible, dully so, on its face. But in making this claim the defenders of this view also make the claim that any given utterance may communicate any number of other propositions in addition to the minimal one that can be linguistically decoded from the original utterance. Here's a further elaboration of this position, from Cappelen & Lepore:

> 'What an utterance of (1) says depends in part on the contextually salient comparison class, standards of measurement, and other such things.
> (1)　Serena is really smart.
> ...Suppose all we tell you is that Venus uttered (1). We predict the following:
> * There is a sense in which you can understand what Venus said, viz. that Serena is really smart.
> * You can repeat what Venus said, i.e. do what we're about to do right now, *viz*. say the same as Venus did: Serena is really smart.
> * You can indirectly report Venus by uttering (2):
> (2)　Venus said that Serena is really smart.'[20]

And another example of how they think content is shared between contexts:

> In some unspecified context, Herman utters: 'Napolean was an interesting guy'. Herman thinks it is obviously true that other people in other contexts have said, could have said, and will say

---

[19] See Cappelen & Lepore (2004) and Borg (2004).

[20] Cappelen & Lepore (2006), pp. 1020-1021.

exactly what Herman said with this speech act.[21]

We can imagine for these claims the extension from saying the same thing in different contexts to another form of content sharing—reporting. The reason they think the content-sharing that occurs when we report is so easy is because all instances of the same locution share minimal semantic content once basic ambiguities have been resolved and indexicals have been specified. And why not? The question of what a report is reporting can't be unanswerable—this must not be the case given the ease with which we navigate indirect reports in conversation.[22]

If the semantic minimalists' answer to the question of this section is just that an indirect report necessarily reports minimal propositional content and optionally reports a plurality of additional content, the pragmatic theorist on these issues should probably resist. It's easy to easy why philosophers and semanticists have tended to think about reporting in terms of semantic identity. But the work of interpreting real world indirect reports has less to do with options semantically triggered by locutions such as 'said that' and more to do with speakers embedded in a context. In a given context, reporters need to convey some understanding or representation of the original context; simultaneously, they need to convey their own reasons for reporting the original utterance, and the relevance of the report to the current context. Sometimes the latter will simply be to convey information, but, oftentimes, indirect reports will be vehicles for a multitude of other speech acts. Wilson (2000) uses the expression "the exploitation of resemblances" to explain the cognitive processes in play when interpreting an indirect report. I think this is the correct way of thinking about it given the variety of reports introduced in this section. The interpreter takes into consideration some collection of clues that are linguistically encoded and takes into consideration contextual information that would lead to the most salient, least taxing interpretation (more on this in section 3). I have tried to emphasize thus far that this available and salient interpretation may resemble the original utterance in a host of ways. Many of these ways may be paralinguistic (e.g., I

---

[21] With modification from Cappelen & Lepore (2004), pg. 127.

[22] See Wieland (2010b) for an opposing view.

might report on the utterance of *A* in order to practice my accent in *A*'s dialect), but they are nonetheless common to real world reporting practices.

3. Pragmatic accounts of indirect reports

3.1. Corollaries with pragmatic accounts of direct reports

How we distinguish between direct and indirect reports differs depending on what we think a language is. Saka (2011) distinguishes between a 'phonocentric' and a 'graphocentric' approach to the study of language. The former he attributes to linguists, arguing that linguists by and large identify language as a natural, verbal (and signed) phenomenon. Written language is, at best, a secondary form of language, meant only to represent verbal speech, and perhaps not a properly linguistic phenomenon at all. The 'graphocentric' approach to the study of language, on the other hand, is the approach most often adopted by philosophers of language, who conceive of language as an abstract (non-natural?) phenomenon. Philosophers of language are most likely to analyze linguistic reality as found in written text and to treat verbal speech as a secondary form of language riddled with 'performance' errors. Saka convincingly argues that settling debates about direct and indirect reports requires first settling questions about the proper approach to language itself. For example, consider these two examples from Recanati (2001, pg. 661):

(25) Alice said that life is difficult to understand.
(26) Alice said that life is 'difficult to understand'.

Recanati's somewhat surprising claim is that, when spoken, these two sentences are, in fact, the same. They do not differ in linguistic meaning, but rather differ only paralinguistically. From a phonocentric perspective this shouldn't be surprising. If we were to imagine a speaker reporting on Alice's assertion we would hear the same thing in both cases. There would be nothing linguistically marked as mentioned rather than used, or quoted directly rather than indirectly. Of course, for philosophers of language, the use/ mention distinction is so salient, and the intuition that quotation marks themselves have a

semantic role is so strong, that Recanati's claim seems preposterous. If Recanati is correct, or if, more generally speaking, the phonocentric approach to language is the correct approach, the distinction between direct and indirect reports does not dissolve, but it does not rely on what it has been thought to rely on—namely, the presence or absence of quotation marks.

On the other hand, when Recanati analyzes (25) and (26) graphocentrically, he concludes that, while there is no difference between the linguistic meanings of (25) and (26), the quotation marks in (26) indicate that the phrase they contain is being demonstrated, rather than represented in some way that differs from the meaning of (25).[23] This debate is not settled. The semantic role of quotation marks and phrases such as 'said that' are still in question, and the pragmatics of how speakers mark off, demonstrate, or report other speech or text is underexplored.

Finally, it's worth noting that speakers (as opposed to writers) mark the distinction between direct and indirect reports in a large variety of ways. Direct reports are marked in English with interjections such as, "and this is a direct quote," or modifiers such as "literally," or "verbatim," or "word-for-word." For example, the speaker in (26) may, in a natural speech setting say, "Alice said that life is—literally—difficult to understand." This would indicate to the audience that the reporter is telling us, directly, what it is that Alice said.[24] Reporters also use air-quotes or say "quote-unquote" in order to indicate that the report is direct rather than indirect, although it's possible that this borrowing from written text only occurs in post-literate settings. (Saka (2011) notes that the use of air-quotes and the phrase "quote-unquote" are often used in clumsy and incoherent ways. This seems right, and seems to suggest that insofar as speakers are borrowing quote-talk to set up a direct report they are doing it without relying heavily on the actual conventions of written text.) These variations in how speakers mark direct reports in English are evidence that they are committed to locutional accuracy, and don't draw a

---

[23] Note that he is arguing against Cappelen & Lepore (1997a), who reject this analysis.

[24] Here we can see how underdeveloped this analysis is. The interjection of 'literally' clearly sounds like it is marking off a direct quote. On the other hand, is it functioning as a demonstration? This is less clear.

strong distinction between direct and indirect reports. Speakers certainly distinguish between loose paraphrase and strong accuracy, but it's not obvious that they draw this distinction on the basis of locutional accuracy rather than on the basis of some other kind of representational accuracy. Getting at the heart of what is taking place in a report requires a more complete theory of what reporters and interpreters are each doing in the reporting context.

3.2. Resources for a complete pragmatics of indirect reports

One possible pragmatic framework for understanding indirect reports is to treat them as a kind of metarepresentation. The task of a reporter is to represent relevant features of the reported context to the audience and thereby convey something about that earlier context. The felicity of the report is constrained, in part, by the cognitive demands of interpreting the report given the current context and shared goals. It is plausible, to a certain extent, that the act of reporting involves having a theory of mind in order to understand what was said and meant in the original context, what the audience knows about that earlier context, and a representation of the two contexts and the respective agents' epistemic attitudes toward those contexts.

There are a few apparent problems with analyzing indirect reports in this way. The first is that understanding the locution 'said that' appears to be a relatively simple cognitive process—one not requiring robust metarepresentational abilities. After all, very young children have the ability to deploy and interpret 'said that' before they have a fully-formed theory of mind. The second, and related, problem is that analyzing indirect reports as metarepresentational suggests that speakers and hearers over-represent in a way that is empirically implausible. Grice originally proposed such a model as described here by Wilson:

> For a Gricean speaker's meaning to be conveyed the speaker's intentions must be not merely recognized but transparent, in a sense that seems to be definable only in terms of an infinite series of metarepresentations. The speaker must not only (1) intend to inform the hearer of something,

and (2) intend the hearer to recognize this informative intention, but (3) intend the hearer to recognize the higher-order intention in (2), and so *ad infinitum*.[25]

This was an assumption of Gricean semantics and early intention-based semantics, albeit recognized as flawed from early on. Stephen Schiffer correctly describes this approach as "superintellectualist":

> In *Meaning* I suggested that in an adequate account of declarative speaker-meaning the speaker must intend the audience to have what I called a "truth-supporting" (as opposed to a pragmatic or prudential) reason for his belief that *p*, but this now strikes me as an especially absurd instance of the superintellectualist conception of mind implicit in the whole IBS approach. I doubt that ordinary speakers can be relied on to have sophisticated beliefs about the sort of grounds they intend others to have for their beliefs. Also it now seems to me possible that there should be acceptable cases of meaning where the audience is not intended to have any reasons for the belief that the speaker intends to activate by his utterance.[26]

In this passage Schiffer is criticizing both the explicitly representational burden as well as the discursive and truth-oriented assumption in this philosophical account. Assuming that a speaker believes *such-and-such* and that his audience recognizes that he believes *such-and-such*, and so on, where each *such-and-such* is identical in linguistic meaning and logical form is implausible. In the case of reporting, this superintellectualist conception of mind is only intensified: the reporter must convey the original speaker's intentions alongside her own, with the relevant information from both contexts, all the while expecting her audience to recognize both the original intention, the intention in reporting, and the higher-order intentions *ad infinitum*. Even if we were to find a way of preventing an infinite iteration of representations, this superintellectualist metarepresentational account of reporting seems too cognitively demanding to be plausible.

While the standard view within pragmatics is that reports are best explained as a kind of metarepresentation, the strongest (but still underdeveloped) pragmatic account avoids the problems of *over*-representation and explicit discursivity. In the case of indirect reports these acts of metarepresentation are even more salient. For example, consider the following exchange:

(27)

---

[25]Wilson (2000), pg. 415.

[26]Schiffer (1987), pp. 246-7.

A: I'm hungry.
B₁ [to *C*]: A said that she's hungry.
B₂ [to *C*]: A said that she would like dinner now.

The report in *B₁* is more direct than the one in *B₂* and presumably requires minimal cognitive resources in order to execute it. The report in *B₂* is indirect and inferential. In this case, *B* understands *A*'s utterance, the meaning behind *A*'s utterance, and the relevance of it to *C*. In doing so, she infers that *A* would like to satisfy her hunger with dinner. But both *B₁* and *B₂* are candidates for felicitous reports. One possible explanation for this is that metarepresentation in general, and in indirect reports in particular, is largely non-discursively inferential and not committed to the recovery of literal meanings (the 'what is said' traditionally construed). Since it appears that young children use various forms of inferential reasoning as a means of language acquisition, the inferences themselves cannot antecedently require strong discursive abilities or the recovery of fully-formed literal meanings.[27] It should be clear that Grice was aware of these problems and thought it would be sufficient that speakers and hearers regard themselves and their audiences as representing in the requisite ways even if they don't actually do so. Sperber and Wilson (in various places) have defended a more plausible (if still cursory) account. Consider the following:

> Inferential comprehension starts from the recovery of a linguistically encoded sentence meaning, which is typically quite fragmentary and incomplete. The goal of pragmatic theory is to explain how the hearer, using available contextual information, develops this into a full-fledged speaker's meaning. The Communicative Principle of Relevance motivates the following comprehension procedure, which, according to relevance theory, is automatically applied to the on-line processing of attended verbal inputs. The hearer takes the linguistically decoded sentence meaning; following a path of least effort in the accessing of contextual information, he enriches it at the explicit level and complements it at the implicit level until the resulting interpretation meets his expectations of relevance; at which point, he stops.[28]

The virtue of this account is that it is general but unified enough to make sense of the puzzles described in section 2. It also—to its credit—makes it easy to see how communication can go wrong even when linguistically encoded sentence meaning is readily available to both speaker and interpreter and how the same process can have degrees of sophistication. When an audience misunderstands a speaker—as happens—it

---

[27] See Wilson (2000) and references therein.

[28] Wilson (2000), p. 420.

is possible that enrichment went wrong at either the explicit or implicit level or that the interpreter's expectations of relevance were inappropriate. After all, these basic pragmatic principles are supposed to explain why it is that complicated communication (in the form of reports, for instance) go well much of the time but also go partially or entirely wrong much of the time as well.

This account also explains the observation that I made above, namely that it seems as if young children are able to deploy and interpret 'said that' even prior to developing a robust theory of mind. In describing this comprehension procedure, Wilson notes that it is compatible with speakers of different abilities. She describes three such speakers who have interpretive strategies she labels *naive optimism*, *cautious optimism*, and *sophisticated understanding*.[29] Each interpretive strategy is associated with a different level of linguistic comprehension and/or theory of mind. For example, a typical language learner progresses through these strategies from infancy to adulthood when speakers and interpreters are typically able to represent complicated theories of mind with respect to one's audience and other speakers that one may be reporting. And as Wilson points out, these categories correspond, to a certain extent, with some of the interpretive obstacles encountered by autistic people. Of course, a complete pragmatic theory will likely propose a more finely-grained non-linear account of the possible interpretive strategies an interpreter might have. But it seems to me that this account, as I've only give the bare outlines of it here, is the best pragmatic (or even bolder: linguistic) framework for the explanation of reporting practices. However, there is still work to be done in philosophy and pragmatics on this meta-linguistic and meta-representational phenomenon.

Looking ahead, what will be the components of a complete theory of indirect reports?

 i. An account of the reporting context and the aims of that context. This would minimally require an understanding of the boundaries of context itself as well as an account of the common knowledge between reporter and audience. Of course,

---

[29] These are described in Wilson (2000), pp. 421-422. She borrows these three categories from Sperber (1994).

much work has been done by pragmatic theorists to account for context itself. It has been beyond the scope of this paper to describe this work. My suggestion, moving forward, is that the reporting context differs from ordinary speech contexts and, as such, will be described in different terms.

ii. A full account of the metarepresentational processes needed in order to represent the relevant features of the original context within the reporting context. This metarepresentational cognitive task requires a triple level of embedding of linguistic understanding, common knowledge, and conversational goals given the framework for comprehension described above. As indicated above, it should be scalar both respect to age and ability level but also context.

iii. An account of the socio-cultural/linguistic interpretive defaults that might be employed interpreting or representing the original utterance as well as the reported utterance. This will require a tandem account of the understanding of socio-cultural/linguistic defaults for the act of reporting itself.

iv. An account, given (i), of what counts as expressing the same content. What has not been accounted for in this paper, but what is certainly of interest to philosophers, is the way in which reporting is an act that has standards of evaluation. In this paper I have used the broad term 'felicity' to describe a successful report. Contexts vary with respect to the required strength of fidelity between the original utterance and the report of this utterance. In some, but not all, cases it is appropriate to assess whether the report is *true*. In some, but not all, cases it is appropriate for there to be *identity* between the original utterance and the report. In other cases, the two utterances, whether in content, form, or something else entirely paralinguistic, need to resemble each other in some other way altogether. In this paper, I have argued that what counts as *saying the same thing* in the original utterance and the report will depend on the norms of discourse for the relevant domain and the conversational goals of the reporter and her audience. This implies that there is no unified analysis of 'said that' or its corollaries that is either available or desirable. However, this does not excuse pragmatic theory from establishing the general principles that constrain speakers in particular cases.

The principles of cooperation proposed by Grice, the general principles of relevance theory, alongside the few additional pragmatic principles aimed at indirect reports that have been proposed here and in recent work on indirect reports go a long way toward explaining many of the linguistic puzzles raised by indirect reports. Indirect reports are proving to be one of the most illuminating linguistic phenomena for understanding the pragmatic/semantic divide, the proper analysis of 'what is said', metarepresentation, and quotation.

References

Anderson, Luvell and Ernest Lepore. Slurring words. In *Perspectives on pragmatics and philosophy*. *Noûs* 97: 25-48.

Borg, Emma. 2004. *Minimal Semantics*. Oxford: Oxford University Press.

Capone, Alessandro. 2010. On the social practice of indirect reports (further advances in the theory of pragmemes). *Journal of Pragmatics* 42: 377-391.

Cappelen, Herman, and Ernest Lepore. 1997a. Varieties of quotation. *Mind* 106: 429-450.

—. 1997b. On an alleged connection between indirect speech and the theory of meaning. *Mind & Language* 12: 278-296.

—. 2004. *Insensitive semantics*. Oxford: Basil Blackwell Publishers.

—. 2006. Shared content. In *The Oxford Handbook of Philosophy of Language*. New York: Oxford University Press.

—. 2009. Quotation. In *The Stanford Encyclopedia of Philosophy (Fall 2009 Edition)*, ed. N. Zalta .
<http://plato.stanford.edu/archives/fall2009/entries/quotation/>

Davidson, Donald. 1968. On saying that. *Synthese* 19: 130-146.

—. 1979. Quotation. In *Inquiries into truth and interpretation*, 79-92. Oxford: Oxford University Press.

Geach, Peter. 1957. *Mental acts.* London: Routledge Kegan Paul.

Larson, R. and Segal, G. 1995. *Knowledge and meaning*. Cambridge, MA: MIT Press.

Ludwig, Kurt. 2003. François Recanati's *Oratio oblique, oratio recta: an essay on metarepresentation*. *Philosophy and Phenomenological Research* 66: 481-488.

Ludwig, Kurt and Georges Ray. 1998. Semantics for opaque contexts. *Philosophical Perspectives* 12: 141-166.

Quine, W.v.O. 1940. *Mathematical logic*. Boston, MA: Harvard University Press.

—. 1961. *From a logical point of view*. Cambridge, MA: Harvard University Press.

Recanati, François. 2000. *Oratio obliqua, oratio recta: an essay on metarepresentation*. Cambridge, Mass: MIT Press.

—. 2001. Open quotation. *Mind* 110: 637-687.

Reimer, Marga. 1996. Quotation marks: demonstratives or demonstrations? *Analysis* 56: 131-142.

Richard, Mark. 1986. Quotation, grammar, and opacity. *Linguistics and Philosophy* 9: 383-403.

Saka, Paul. 1998. Quotation and the use-mention distinction. *Mind* 107: 114-136.

—. 2011. Quotation and conceptions of language. *Dialectica* 65: 205-220.

Saul, Jennifer. 1998. The pragmatics of attitude ascriptions. *Philosophical Studies* 92: 363-389.

Schiffer, Stephen. 1987. *Remnants of Meaning*. Cambridge: MIT Press.

Seymour, Michel. 1994. Indirect discourse and quotation. *Philosophical Studies* 74: 1-38.

Sperber, Dan. 1994. Understanding verbal understanding. In *What is intelligence?*, 179-198. Cambridge: Cambridge University Press.

Sperber, Dan and Dierdre Wilson. 1986. *Relevance: communication and cognition*. Oxford: Blackwell.

Tarski, Alfred. 1933. The concept of truth in formalized languages. In *Logic, semantics, metamathematics*, 2nd ed., 1983, 152-278. Indianapolis: Hackett.

Washington, Corey. 1992. The identity theory of quotation. *Journal of Philosophy* 89: 582-605.

Wieland, Nellie. 2010a. Context sensitivity and indirect reports. *Philosophy and Phenomenological Research* 81(1): 40-48.

—. 2010b. Minimal propositions and real world utterances. *Philosophical Studies* 148(3): 401-412.

Wilson, Dierdre. 2000. Metarepresentation in linguistic communication. In *Metarepresentations: a multidisciplinary perspective*, 411-448. New York: Oxford University Press.