

Utterance content, speaker's intentions and linguistic liability*

Claudia PICAZO JAQUE

Received: 13/07/2016

Final version: 17/05/2017

BIBLID 0495-4548(2017)32:3p.329-345

DOI: 10.1387/theoria.16720

ABSTRACT: According to contextualists, communication has to do with pragmatically adjusted content, not with conventional meaning. This pragmatic content is sometimes identified with speaker meaning or with the thought the speaker intends to express. I will argue that given the sociolinguistic role of utterance content—the fact it provides reasons for action, liabilities and entitlements—locutionary content should not be modelled as a variety of speaker meaning.

Keywords: what is said, utterance content, accountability, linguistic liability, speaker meaning.

RESUMEN: De acuerdo con los contextualistas, la comunicación tiene que ver con contenidos pragmáticamente ajustados al contexto. Este contenido pragmático es a veces identificado con el significado del hablante o con el pensamiento que el hablante tiene intención de expresar. Argumentaré que, dada la función sociolingüística de la noción de contenido de la preferencia—el hecho de que proporciona razones para la acción, responsabilidades y permisos—el contenido locucionario no debe ser modelado como una variedad de significado del hablante.

Palabras clave: lo que se dice, contenido locucionario, responsabilidad lingüística, significado del hablante.

1. *Speech act content and communication*

In talking we typically acquire commitments. If I promise that I will ϕ , then I commit myself to ϕ -ing. If I assert that p , I commit myself to the truth of p and of having reasons supporting p . I am accountable for what I have asserted. Thus, there is a connection between speech act content and the acquisition of commitments and liabilities. When we engage in a conversation, we keep track of those commitments and liabilities. In this sense, conversations can be seen, as Brandom (1994) puts it, as involving deontic scorekeeping.

* I am indebted to Inés Crespo, Andreas Heise, Palle Leth, and Manuel Pérez Otero for helpful discussions and comments on earlier versions of this paper. I have presented versions of this material at the III Congreso de graduados de la Sociedad Española de Lógica, Metodología y Filosofía de la Ciencia and the VIII Congreso de la Sociedad Española de Filosofía Analítica. I am grateful to these audiences, especially to Jordi Valor. I also thank two anonymous reviewers for this journal for helpful comments and suggestions. Financial support: Ministerio de Economía y Competitividad (Government of Spain) and European Union, Grant FFI2015-63892-P (MINECO, AEI/FEDER, EU).



This is not the only way in which we can look at conversations. Language is an efficient tool for conveying thoughts. Although there certainly is non-linguistic communication, engaging in a conversation is one of the main ways in which we share thoughts. Deontic scorekeeping and ‘belief’ scorekeeping go hand in hand. Typically, when I assert that *p* I not only acquire some commitment and accountability but also convey my belief that *p*. If everything goes well, the hearer will recover the content of my belief.

Despite the fact that these two dimensions are intertwined, I think it is useful to distinguish two projects. On the one hand we have what could be called a theory of linguistic commitments and liabilities—a branch of speech act theory. This theory deals with the responsibilities that speakers acquire in performing a speech act. On the other hand, we have the theory of communication or of sharing thoughts. The two projects aim at making sense of different practices. The first aims at making sense of our linguistic commitments. The second aims at explaining the transmission of thought via language. Given that they are two different projects, it can be the case that a notion that is useful for one project is not central to the other, and vice versa. This, I think, is what happens with the notion of utterance content.

Communication succeeds whenever the hearer recovers the thought intended by the speaker. Davidson (1990) writes:

What matters to successful linguistic communication is the intention of the speaker to be interpreted in a certain way, on the one hand, and the actual interpretation of the speaker’s words along the intended lines through the interpreter’s recognition of the speaker’s intentions, on the other. (Davidson, 1990, 311)

What matters to communication is that there is convergence among interlocutors, that the speaker manages to get his thought across. In order for this to be the case, the hearer must recover what the speaker means (the thought he intends to communicate). The notion of speaker meaning will arguably be a key notion in any theory of communication: what matters to successful communication, we can say, is that the hearer recovers what the speaker means. But it is not clear that we will need, on top of that, a notion of utterance content (the truth-conditions of a sentence-in-use, where these can transcend conventional meaning). Plausibly, conventional meaning will figure prominently in any explanation of linguistic communication. The easiest way to convey a thought is to find a sentence whose conventional meaning coincides with the content of the thought. However, there are cases of successful communication where the thought intended by the speaker and recognised by the hearer does not even coincide with the conventional meaning. Malapropisms and slips of the tongue provide good examples, but we can think also about speakers speaking a language their hearers do not fully understand but making themselves understood by the use of gestures. It seems that what is relevant to convergence is speaker meaning, not utterance content—not what the utterance *really* says.

In spite of the notion of utterance content not being crucial for explaining convergence, we need a notion of utterance content for other purposes. We sometimes care not only about what others mean but also about what they say and what they are accountable for. For example, we care about the information others put forward and hold them responsible for its accuracy. In these cases we might be interested not only on what people mean, but also on what they actually say—for we obtain knowledge and coordinate action on the

basis of what they say, and consequently attribute liabilities. Hence, the notion of locutionary content is a sociolinguistic one: it has its home in the public sphere, in the information or otherwise the contents we put forward in a speech act, and is, in this sense, inseparable from the acquisition of commitments on the part of the speaker. In contrast with communication as convergence, where the question about what an utterance actually says might be of limited interest, in holding speakers accountable and, in general, in keeping track of the deontic scorekeeping of a conversation, what an utterance actually says becomes an important issue. Locutionary content is content speakers are accountable for.

Camp (2006) has defended a similar approach. According to Camp, we can develop a useful notion of 'what is said' by focusing on the commitments that speakers undertake. In particular, Camp works towards an identification of what a speaker says with an utterance of a sentence and the commitments hearers are entitled to attribute to the speaker by virtue of the utterance.¹ Interestingly, she supports her approach with some observations about ordinary uses of 'say'. As she points out, indirect reports are appropriate only when the speaker is reported as having said something to which he has explicitly committed himself—as opposed to, for example, contents that are merely suggested or insinuated. If, for example, as an answer to the question how good as a philosopher John is I utter the sentence 'John's command of English is excellent, and his attendance at tutorials has been regular' it would be inappropriate to report me as having said that John is a bad philosopher, for I have not openly committed myself to that claim. By contrast, it would be appropriate to report me as having said that John's command of English is excellent—a claim I have explicitly committed myself to. This speaks in favour of modelling what is said as being related to the commitments one acquires in speaking.

Now, the notion of 'what is said', or locutionary content, has lately been the topic of much debate. In particular, contextualists have argued that it should not be identified with conventional meaning,² for what is said with an utterance of a sentence *S* can go beyond what is lexically encoded in the conventional meaning of *S*. In order to show this, some philosophers have used indirect reports,³ since sometimes the correct report of an utterance is not a homophonic one, but instead one that adds more linguistic material. Imagine I say 'He's tall'. What exactly is the commitment that I have undertaken? What is the appropriate report? Probably, the appropriate report will make reference to some comparison class ('He's tall for a 5-year-old child'). The commitment I have acquired will also be relative to this more precise content, and not a generic claim ('He's tall, no matter with whom you compare him'). This, together with other considerations, has motivated the claim that what is said is often the outcome of a process of completion or modulation.

The previous example motivates a question: how is the comparison class determined? Does it depend on the speaker's intentions? In what follows I will argue that, in defective contexts, that is, contexts where the speaker meaning of a given utterance is not aligned with the interpretation a normal interpreter is most likely to arrive at, liabilities and commitments tend to be aligned with the salient interpretation (that of a normal interpreter,

¹ Given that she wants to isolate 'what is said' from commitments related to contents that can be derived from it in context, her definition is more sophisticated than this.

² See for example Carston (2002), and Recanati (2004).

³ See Recanati (2006) and Wieland (2010).

i.e., one who has mastered the language and is not cognitively impeded), not with the speaker's intentions. Given the relation between locutionary content and commitments, I conclude that we should not model 'what is said' as being fixed by speaker's intentions.⁴ The plan is the following. In section 2 I will introduce the contextualist debate. In section 3, I will argue that accountability doesn't tend to align with speaker meaning. After this, I will consider some objections (section 4). I will finish by addressing Perry's view on what he calls 'the forensic', i.e., our ordinary notion of saying's sensitivity to responsibilities (section 5).

2. *The contextualist debate: modulated locutionary content*

It is nowadays common to hold the view that the content of an utterance of a sentence *S* need not be (or cannot be, according to radical contextualism) equivalent to the linguistic meaning of *S*, even in absence of indexicals.⁵ The reason is that a non-indexical sentence can express different sets of truth-conditions when used in different context. Take as an example 'Tipper is ready'. In a conversation about an interview Tipper is about to have, an utterance of it will very likely say that Tipper is ready for the interview, and will be true if and only if Tipper is ready for the interview. By contrast, if the sentence is used in a conversation about going out for dinner, an utterance of it will say that Tipper is ready to go out for dinner. Similar examples can be created for a variety of predicates, including 'weighs 79 kilos' ('Hugo weighs 79 kilos' can be used to say that Hugo weighs 79 kilos naked, before breakfast or that Hugo weighs 79 kilos with his clothes on), 'needs a red pen' ('Paul needs a red pen' can mean that Paul needs a pen with red ink or that Paul needs a superficially red pen), and so on.

There is a hot debate about the semantics/pragmatics. We can distinguish two issues. First, it is open to discussion whether predicates express, of themselves, properties, and, consequently, whether non-indexical sentences express truth-evaluable content or whether, by contrast, a context of use is needed in order to get something truth-evaluable. According to minimalists (Borg 2004, 2012; Cappelen and Lepore 2005), the semantic content of well-formed declarative sentences is truth-evaluable. According to radical contextualists (Searle 1978; Travis 1989, 2008), there is no truth-evaluable content outside a context of use. Between these two options, there is a range of intermediate positions, usually considered versions of contextualism.⁶ Some philosophers take it that the semantic content of some well-formed declarative sentences is not truth-evaluable, but instead that the proposition these sentences express must be somehow completed (Bach 1994). Others agree that semantics delivers a truth-evaluable proposition, but claim that this proposition is often pragmatically adjusted, where the result of such an enrichment must be considered as the literal meaning of the utterance (Recanati 2010). Here is where the second issue enters the picture. One can concede that predicates semantically express properties, but argue that

⁴ Perry claims that 'for the philosophy of language we want a concept that is closely tied to the actual structure of intentions, and immune to forensic issues' (2009, 194). He focuses on demonstratives, but also mentions predicates. Korta and Perry (2011) generalize this view. I will focus on predicates.

⁵ As an example of the first option, see Recanati (2004, 2010). More radical views are held by Carston (2002), and Travis (2008).

⁶ See Borg (2012) for a nice presentation of the different positions.

these properties are not relevant for communication, broadly understood. On this view, the notion of content relevant for communication will often, if not always, be the result of pragmatic adjustment. Despite the disagreement among answers to the first question, there is a certain consensus that there is a legitimate notion of content (or 'what is said') that often goes beyond the semantically determined meaning and that is the one relevant to our communicative practices,⁷ because of being the one available to speakers, the one needed to derive implicatures, the one relevant for practical reasoning, etc. I will refer to this kind of content as the (locutionary) content of the utterance.

In what follows, I leave aside the question about whether sentence-types (as opposed to utterances) are also bearers of truth-evaluable content and focus on how to model utterance content.⁸ The question I will address here concerns the determination of the property expressed, or otherwise the satisfaction conditions of the predicate. I will take as common the ground the claim that the relevant truth-conditions of many utterances are a pragmatic affair. What I am interested in is what 'pragmatic' means here. To some, pragmatic content means speaker-meant content. By contrast, my aim is to argue that the content of an utterance of a sentence, what is usually called 'what is said', the modulated or adjusted meaning or the explicature,⁹ when underdetermined by semantics, should not be modelled as being determined by the speaker's intentions.

3. *What is said and locutionary commitments*

Many of the examples used in order to motivate the need for modulated meanings show that talk is intertwined with extralinguistic actions. Take as a sample the oft-discussed ex-

⁷ On this point, as I mentioned earlier, minimalists and contextualists can agree. Thus, Carston (on the contextualist side) approvingly discusses Borg's claim (on the minimalist side) that semantics does not deliver the content relevant for communication: '[Borg] discusses and puts to rest (hopefully forever) the idea that a semantic theory is required to deliver for any sentence *s* the content *p*, where for a speaker *S* who utters *s*, '*S* said that *p*' is a correct indirect report. As she argues, both intuitive pre-theoretic notions of 'what is said' and the technical Gricean notion of 'what is said' are a matter of communication (or speaker meaning), hence involve consideration of speaker intentions, and so fall outside the domain of a formal semantic theory, while any more restrictive notion of 'what is said' turns out either to be dependent on the very semantic judgements it is meant to constrain or to be redundant in determining semantic content since the real work is being done by syntactic features of the sentence uttered. As suggested above, once we recognise that what the contextualist is focussing on is this pragmatic notion of what is said (that is, the truth-conditional content of an utterance or speech act) as distinct from a formal 'pure' level of linguistic semantics, much of the fire goes out of the minimalist/contextualist debate' (Carston 2008, 361). Note that Carston identifies utterance content and speaker meaning.

⁸ By 'utterance content' I mean the truth-conditions of a use of a sentence, where these can transcend the truth-conditions determined by the linguistic meaning of the sentence.

⁹ Different authors use different terms to refer to the truth-conditional content expressed by an utterance, with the locution 'what is said' probably being the more standard one. Recanati (2004, 2010) also usually refers to modulated meanings, i.e., meanings adjusted in context (enriched or loosened). Relevance theorists as Carston (2002) use the expression 'explicature'. It is beyond the scope of this paper to assess whether there are some differences among these notions.

ample involving the expression ‘red apple’. The situation is the following:¹⁰ Carmen arrives home and says ‘I’m very hungry, is there something I can eat?’ María, pointing to the kitchen, replies ‘There are some red apples on the table’. This assertion will very likely motivate (and justify) certain actions (going to the demonstrated place, grabbing an apple, eating it). If it turns out that it is false, i.e., that there are no red apples on the table, the addressee is entitled to complain. Communication involves, among other, transmission of knowledge, reasons for action and the acquisition of justifications and liabilities. The notion of ‘what is said’ should make sense of this—for example, it should explain how is it that, if Carmen finds out that the apples were rotten (red in the inside) she is entitled to complain, and that nonetheless her action was justified.

Consider another example. Hugo and Laura are waiting for Paul. They are going out for dinner, and Paul is getting dressed. Hugo says ‘Paul is ready’. Laura grabs her coat and opens the door. As it turns out, Paul is not ready to go out for dinner, but he is ready for many other things (for instance, to take a shower). We need a notion of ‘what is said’ that, at least, (i) explains Laura’s actions (as being rational) and (ii) according to which Laura is justified in asking Hugo for an explanation. Why? Because we, normal speakers, plan our behaviour on the basis of other people’s utterances that, as Hugo’s or María’s, need to be completed or adjusted in context. Moreover, not only do we plan our behaviour, but it also seems rational or justified to do so. When analysing how content depends on context we should not ignore this fact. On the contrary, the discussion about how to model content should make sense of this practice.

With this in mind, let us address the following question: how is the property expressed by a predicate in an occasion of use determined? A particular line of thought, sometimes more assumed than scrutinised, takes the content of the utterance to be determined by speaker meaning. An example:

[Discussing an utterance of the sentence ‘That is empty’, where ‘that’ refers to a can] Did I mean “contains no Sprite” or “contains nothing at all”? The word “empty” can be used to get at a number of properties, depending on what one intends to say with it, and it can also be used less definitely. If I meant, in my first remark, about the can on the table, that it contains no Sprite, what I said was true. If I meant *completely* empty, in the state required by the city of Palo Alto if I am to toss it in the recycling bin, what I said was false. If I had no intention either way, then I didn’t say anything clear enough to have a definite truth value one way or the other. (Perry 2009, 197)

For Perry, the truth-conditions of ‘That is empty’ are contingent upon what the speaker intends to say. Similar approaches are held by Carston (2002), and Bach (1994). These views suggest that the content of the utterance is identical to the thought the speaker intends to communicate, or otherwise determined by what the speaker means. Let’s call this view Speaker Centred Meaning (SCM).

A couple of qualifications are needed. First, content is, on this view, constrained by linguistic meaning, so it’s not completely up to the speaker. Moreover, what the speaker means is supposed to be interpretable. It is part of the technical notion of speaker meaning that a

¹⁰ I adapt from Wieland (2010). She uses the example to argue that homophonic indirect reports are not always correct.

speaker who means that *p* by utterance *U* intends his audience to recognise his intention and to interpret that *p* on the basis of that recognition (Grice 1957, 1969). Or, as Davidson (1986) pointed out, given that one cannot intend what one knows to be impossible, a (rational) speaker cannot mean something he knows to be non-interpretable. Because of this, SCM is not necessarily the view that if a speaker 'means' that today is sunny by uttering a random sentence ('Green is my favourite colour', for example), then the content of his utterance is that today is sunny, regardless of the linguistic meaning of the sentence used.¹¹

Second, SCM is a view about what Recanati (2004) calls 'primary pragmatic processes' such as enrichment or loosening, not secondary processes as implicatures. Recanati (2004), as other contextualists do, distinguishes three levels of content: (i) linguistic or conventional meaning, (ii) what is said, the literal content of an utterance, and (iii) implicatures. SCM is a view about the second level.

In contrast to SCM, other philosophers take content to be determined by the available interpretation (Recanati 2004), the reasonable interpretation (Travis 1989) or some feature of the context of use, such as the purposes of the conversation (Travis 2008; Kölbel 2008) or the Question Under Discussion (Schoubye and Stokke 2016). I will not try to establish an alternative model to SCM here. My aim is to compare it to a broad approach that would model locutionary content as being close to the salient interpretation. By 'salient interpretation' I mean the interpretation that a normal, informed, rational interpreter is most likely to arrive at, given the previous discourse and the linguistic and extralinguistic surroundings of the utterance. This broad approach is, I think, compatible with two kinds of theories. According to the first, locutionary content is determined (in the metaphysical sense) by the salient interpretation, or by the interpretation that normal speakers would arrive at (perhaps under some idealization). According to the second, locutionary content is determined by some objective parameter, such as the Question Under Discussion, with this parameter being available to normal speakers. In this second view, salient interpretation merely tracks locutionary content.

Now, why shouldn't we model the content of an utterance as being determined by speaker's intentions? Because a notion of content determined by speaker meaning makes the wrong predictions when it comes to linguistic liability. Speakers can generally be held accountable for different levels of content. In the simplest case, one is committed to the con-

¹¹ There can be contexts where a speaker means that today is sunny by an utterance of 'Green is my favourite colour'. Zarraluki's novel *La historia del silencio* depicts such a context for the sentence 'I've forgotten the spaghetti': 'I had just come back from the supermarket and I emptied the shopping bags out on the kitchen table. Irene was already boiling water for the pasta in a pan. Then I gave her a heartbroken look before announcing that I had forgotten the spaghetti. From that moment on, in our private language 'I've forgotten the spaghetti' came to mean giving up on something out of a sense of overwhelming fatigue. And so, when Irene once managed to go four days without smoking, she said I've forgotten the spaghetti and lit a cigarette. And I said it while lying in bed, the moment I woke up, the day I decided to abandon my stubborn efforts to go the gym every morning...' (electronic version consulted). My point is that SCM is not committed to the claim that, regardless of the context, a speaker can say that today is sunny by uttering no matter which sentence. Unless there is some kind of code shared by speaker and hearer, a rational speaker will not take his addressee to be able to interpret no matter what. Because of that he will not mean, i.e., intend to say, that today is sunny by uttering a random sentence. According to Davidson, one simply cannot form such an intention.

ventional meaning of the words he utters. If, in normal circumstances, I assert ‘France is a European country’, then I commit myself to the claim that France is a European country. We also hold speakers accountable for the locutionary, or modulated, content of their utterances. Imagine that Hugo says ‘Paul is not ready’, in a scenario where Paul is not ready to go out for dinner, but he is ready for other things. If, in this case, we held Hugo responsible for the conventional meaning of his words, then we could complain that he had said something false, for there is something Hugo is ready for. However, it would be more natural to hold him accountable for the truth of the enriched content *that Paul is not ready for dinner*.¹²

Let us go now to the previous examples. According to SCM, whether María and Hugo said something true or false depends, inter alia, on what they meant. If María meant that there were some red-in-the-inside apples, then what she said is true. If Hugo meant that Paul was ready for a shower, then what he said was true. Or, to put the point as concerning the information transmitted, what info they conveyed by their utterances depends on what they meant. The information they put forward in their assertions is, according to SCM, contingent upon their communicative intentions. Let us stipulate that María meant that there were some red-in-the-inside apples and that Hugo meant that Paul was ready for a shower. If that is so, according to SCM, the content of their utterances would be that there were some red-in-the-inside apples and that Hugo meant that Paul was ready for a shower. By contrast, normal speakers would take their utterances to convey different information and would be justified in doing so and entitled to complain. More importantly, what María and Paul mean, when not made explicit, does not go into the deontic scorekeeping of the conversation—which shows that the notion of content at stake here is not speaker centred.

I will consider a last example, this time a more detailed scenario. Consider the following conversation: François, Marie and Julien work together. They are proofreading a book and writing down some corrections, each of them working on a different copy. François is using a red pen to mark typos. Marie underlines in green grammatical mistakes. Julien’s job is to improve the style of the writing, and writes his suggestions in blue on the margins of the document. From time to time, Julien marks a typo (in red), or underlines a grammatical mistake (in green), and so do the others. François’ pen has run out of ink. He says: ‘I can’t finish the book, I need a new pen’. Marie replies: ‘There are some pens in that box’. François checks and replies: ‘None of them is red. I can’t use them for my corrections. I need to buy a new red pen. Does someone else need a new pen?’ Julien answers: ‘I need a blue pen, could you buy one for me? I will give you the money later.’ François buys two pens, one with red ink, one with blue ink, none of which is superficially blue, and gives the blue one to Julien. Each of them has cost 1€.

Does Julien owe 1€ to François? I think the answer is ‘yes’. Julien asked for a certain kind of pen and committed himself to paying for it. François fulfilled his request and, as a consequence, is now entitled to ask for the money. However, SCM might predict otherwise. According to this view, what Julien asked for, and consequently the commitment he acquired, is determined by his intentions, in particular, it depends on which property

¹² In some cases speakers might even be held accountable for contents that are merely insinuated or implicated, as the Liberace case shows (more on this later). This might suggest that implicated content should also be modelled as not being determined by speaker meaning. Nonetheless, I think that as we depart from explicitly articulated content things get more complicated, and I will not discuss implications here.

he intended to communicate when he uttered the words 'I need a blue pen'. If, despite the surroundings of the conversation, he had in mind a superficially blue pen (because he was thinking about wearing a pen in his pocket, and blue fits his necktie specially well), then he asked for a superficially blue pen. Imagine that was precisely the case. Julien was not paying full attention to the conversation. Instead, he was thinking about his new necktie. When he heard François asking if someone needed a pen, he was confused about what kind of pen François was talking about. Nonetheless, he answered his question, with the intention to ask for a superficially blue pen. As a consequence, according to SCM, François did not fulfil the request, for he didn't buy a superficially blue pen.

SCM runs counter our practices of attributing responsibilities. We hold Julien as being committed to pay François the money back because, given the circumstances of the conversation, the salient interpretation of his words was that he needed a pen that writes in blue. Speakers are, of course, authoritative about what they mean, and Julien may complain that that's not what he meant. This non-intentionality might work as an excuse. This, however, doesn't prevent François from being entitled to ask for the money. At most, Julien could retreat to the conventional meaning of his words (perhaps 'blue in some sense or other') but not claim that he is only committed to what he meant.

Given that we generally don't hold speakers accountable for the locutionary contents they mean, I conclude that we should not model what is said as a variety of speaker meaning. In the next section I will consider some objections that could be raised in defence of SCM.

4. *Objections and replies*

OBJECTION 1: JULIEN HAS NOT ACQUIRED ANY RESPONSIBILITY, BECAUSE HE DIDN'T HAVE THE INTENTION TO ASK FOR THAT PARTICULAR KIND OF PEN

One can try to invalidate the argument by claiming that there are no responsibilities involved in the example, for one is responsible only for what one intentionally does. However, this is not true in general. If I am playing with a ball near a window, and I accidentally hit the glass with the ball and as a result it breaks, I am responsible for it despite the fact that I didn't have the intention to break it. I did break the window, despite the fact that I didn't mean to. Speech acts are similar to other acts. One thing is what someone intends to do, a different thing is what one does. If everything goes well, the two will coincide. But that is no reason to conclude that what one does is determined by what one intends to do. Speakers sometimes acquire commitments that go beyond what they have considered. For example, we are, in general, committed not only to the content of our utterances but also to its logical consequences. Moreover, speakers are sometimes mistaken about the content of the words they utter. The phenomenon is quite trivial: there are slips of the tongue, malapropisms, confusion about whom a proper name refers to, and so on.¹³ Why shouldn't the same happen in cases where the content of a predicate goes beyond its linguistic meaning?

¹³ Very often, the fact that something was said unintentionally, by mistake, serves as an excuse. Following Austin (1956) we can distinguish excuses from justifications. Imagine someone being accused of doing something wrong (for instance, asserting something false). According to Austin, providing a justification amounts to admitting that one has done what he is being accused of but claim that it was not

OBJECTION 2: FRANÇOIS IS ONLY SECONDARILY ENTITLED TO ASK FOR THE MONEY

A second possibility is to focus on François' entitlement to ask for the money and to claim that it is of a special kind. One can do this by distinguishing primary and secondary entitlement.¹⁴ Imagine two individuals, A and B. A says to B: 'You will find some snacks at the office. They're for free'. A goes to the office and eats some of the snacks without paying for them. This action is appropriate according to B's assertion. Let's say that an action is primarily entitled by a speech act if and only if it is in accordance with the content of the speech act. By contrast, an action is secondarily entitled by a speech act if and only if it is in accordance with what the hearer takes to be the content of the speech act (and he is being rational, we can add). Secondary entitlement has to do with what is correct to do given the available evidence. Imagine that B was not talking to A about the office they both work in, but to someone else, C, about the office C works in, but that A couldn't see C. A misinterprets B, goes to the office and eats some of the snacks without paying for them. There is some sense in which A is justified in such an action, since, as far as he knows the snacks are for free. Something similar could be going on with François. Given the available evidence, he takes it that Julien wants a pen with blue ink. Since he could not have rationally interpreted the utterance otherwise, he is justified in having bought that particular pen, and now he is secondarily entitled to ask for the money.

I grant that there might be something like secondary entitlement, and this could maybe explain the impression that François has not made any mistake within a speaker centred approach. However, this is insufficient to explain how is it that Julien has incurred in some kind of mistake. This is what needs to be explained.

OBJECTION 3: JULIEN'S RESPONSIBILITY HAS TO DO WITH THE FACT THAT HIS WORDS WERE VERY LIKELY TO BE MISINTERPRETED, AND HE SHOULD HAVE CONSIDERED THIS

Another objection might agree that Julien has acquired some responsibility, but disagree on its origin. The reason why Julien is financially responsible is that he should have paid attention to the conversation and be aware that his words would be misunderstood. The idea here is that our intuitions about who is responsible in this case do not track content, but something else—our obligation to be good communicative partners and avoid using words that can easily be misinterpreted, let's say.

This objection targets the claim that responsibilities show that locutionary content aligns with salient interpretation and not with speaker meaning. It does not constitute a novel defence of SCM, for it doesn't provide any new reason to the effect that we should model locutionary content as a variety of speaker meaning—we don't have any reason for holding Julien accountable for what he meant instead of holding his responsible for the sa-

something wrong. By contrast, one might accept that the action was wrong but don't accept responsibility because, for example, it was done unintentionally. Austin reserves the term 'excuse' for this second case. Malapropisms and slips of the tongue are generally excused, for they are unintentional. However, it would be odd to conclude that the speaker who commits a malapropism didn't say what the sentence he uttered meant—it would be more natural to say that he didn't mean it.

¹⁴ This is inspired on DeRose's (2002) distinction of primary and secondary propriety.

lient interpretation of his words. I think, however, that there are two reasons why the objection is problematic.

The first is that it is more difficult to run this objection when it comes to other examples. It is probably true that some of the responsibilities we attribute have to do with speakers using misleading words. There might be cases where speakers are responsible (and blameworthy) because of saying something misleading. It can be discussed whether Julien's case is one of those. But there are others where this objection would be odd. Imagine that, in the situation depicted earlier, Laura says: 'We're going to be late for dinner. We should leave now. Is Paul ready?' and that Hugo answers: 'He's ready', meaning that he is ready to have a shower. Or, in a more radical example, that Laura asks 'Is Paul ready for dinner?' and Hugo, who was not paying full attention, answers 'He's ready', meaning that he is ready to have a shower. Here, the speaker clearly didn't choose the right words in order to communicate his thought. Nonetheless, when it comes to the truth-conditions of his utterance, it would be odd not to model them as depending on the question the utterance is an answer to.

The second reason is that the objection suggests that speakers are only accountable for the explicit content, or articulated content, of their utterances. In other cases where we attribute responsibility, we would do so because the speaker shouldn't have used words that would very likely be misinterpreted. However, we generally recognise commitments beyond explicit content that are content-based. Borg (2016, 352) alludes to a 1950s case where a newspaper was sued for libel and found guilty because of the implicated content of the following paragraph: '[Liberace is] the summit of sex —the pinnacle of masculine, feminine, and neuter. Everything that he, she, and it can ever want ... a deadly, winking, sniggering, snuggling, chromium-plated, scent-impregnated, luminous, quivering, giggling, fruit-flavoured, mincing, ice-covered heap of mother love'. Liberace claimed that the newspaper has described him as a homosexual. Despite the fact that the newspaper argued that the paragraph was not meant to describe Liberace as a homosexual, it lost the trial. The judge considered that the paragraph was a libel, presumably because it depicted Liberace as a homosexual. In this case, the newspaper was found accountable for a content that was not explicitly stated. Moreover, this case suggests that the intentions of the speaker (or writer) are not decisive in establishing accountability.

OBJECTION 4: THE RELEVANT INTENTIONS HERE SHOULD NOT BE IDENTIFIED WITH THE SPEAKER'S INTENTION TO CONVEY A CERTAIN THOUGHT BUT WITH HIS INTENTION TO ANSWER A QUESTION OR CONTRIBUTE TO A CONVERSATION

One might have recourse to the intention to defer the adjusted meaning of a given sentence or expression. In the previous example, the speaker plausibly has the intention to answer the question 'Is Paul ready for dinner?', and so he has fixed the topic of his utterance deferentially. With this correction, SCM can explain the previous example.

In order to work, this view needs to be a bit more sophisticated. Conversations don't happen in the vacuum. They are often accompanied by extralinguistic actions, and interlocutors obtain a lot of information from this extralinguistic context. So the topic of the conversation can be fixed implicitly—for example, by making some gestures, like pointing to the watch, and so on. It can be the first person to talk the one that is wrong about what the conversation is about. However, even if it is true that very often the topic of one's ut-

terance is fixed deferentially, and that speaker's intentions to answer a certain question are very relevant in taking their utterances to actually be answers to those questions, SCM goes one step forward. It doesn't merely say that speaker's intention to participate in a conversation are relevant—it claims that we should identify the content of an utterance with what the speaker means, that is, with the thought he intends to communicate. But these two can easily come apart: a speaker might have the intention to contribute to the on-going conversation but, because of being wrong about what the conversation is about, intend to say that Paul is ready for a shower by uttering the sentence 'Paul is ready'.

OBJECTION 5: IF JULIEN WAS THINKING ABOUT SOMETHING ELSE, THEN IT IS NOT DETERMINED WHAT THE CONVERSATION WAS ABOUT AND AS A CONSEQUENCE IT IS NOT DETERMINED WHAT HE ASKED FOR

This objection calls into question the possibility of having a situation of the kind described, in which what the conversation is about is fixed but at the same time one of the participants is unaware of it, or simply mistaken about it. I grant that the content of some of our utterances is simply underdetermined: there can be two possible interpretations and no way to decide between them. This happens, for instance, whenever we have two interlocutors talking past each other. However, there are also cases where one of the participants might disagree about the topic of the conversation and yet this is clearly fixed. Imagine that we have four interlocutors; three of them agree that they are talking about an interview, and there's even a CV is on the table that they check from time to time. The fourth interlocutor takes it that they are talking about going out for dinner—because that's what they were talking about ten minutes earlier, let's say. Despite this interlocutor not being in line with the others, this is a case in which the topic is fixed. The fourth interlocutor is simply mistaken. The blue pen example depicts one of these cases.

One might still object that what is said might be construed as being fixed by the shared beliefs of the interlocutors. Instead of the intentions of the speaker, we could focus on the intentions the speaker shares with the hearer. This, however, is too demanding, for it rules out the possibility of having a content in the previous cases (there were no shared beliefs about what the conversation was about, for instance).

A variant of this objection might complain that the example is too much of an idealisation. In real life, Julien would complain that he was misunderstood and make explicit what he meant. I grant that. Moreover, in many cases we only care about what the speaker meant. However, this is not of use for the proponent of SCM. An avowal of what a speaker meant is not a statement of the content of a previous utterance. It simply is a way of making explicit what someone intended. So, if Julien was to say 'I meant a superficially blue pen' that might explain why communication was unsuccessful, but it doesn't settle the issue about who should pay for the pen. Making explicit what one means does not amount to fixing the content of a previous utterance.

OBJECTION 6: OUR ATTRIBUTIONS OF RESPONSIBILITIES ARE WRONG

Another possibility is to insist on SCM and to argue that we are conflating epistemic and metaphysic determination. In the metaphysic sense, the argument would go, what is said is determined by the speaker's intentions. Hearers, on the other side, use the context in order

to determine (epistemic sense) what the speaker means. If so, given that the responsibilities a speaker acquires correspond to the content of his speech acts (and not to what others take to be the contents), our attributions of responsibilities are wrong in all these cases. We hold speakers responsible for available interpretations, and not, strictly speaking, for the content of their utterances. It seems to us that Julien owes 1€ to François, but this is simply a mistake on our part.

This strategy faces an important problem. People have responsibilities, commitments and entitlements because we take them to be responsible, committed or entitled. These deontic statuses, as Brandom (1994) calls them, are not independent of our practices: they do not pre-exist them. Therefore, it is not very plausible to claim that we are that wrong. If we find out that the attribution of responsibilities is, as a general rule, sensitive to the salient interpretation and not to what the speaker means, then insisting that we are wrong in our attributions does not seem like a good option¹⁵.

OBJECTION 7: LIABILITY TRACKS SALIENT INTERPRETATION, NOT WHAT IS SAID

According to this line of thought, responsibilities go hand in hand with the salient interpretation but, nonetheless, the correct interpretation is the one the speaker intended. Our attribution of responsibility is sensitive to the surroundings of the conversation (to the activity in place, previous discourse, and so on) in a way that our notion of content should not be. The blue pen example doesn't force the proponent of SCM to revise his approach, but only to note that responsibilities have nothing to do with the domain of semantics.

It can be argued that there are here two notions of content at stake. On the one hand, we hold speakers responsible for the salient interpretation, and, on the other, because of independent reasons, we have an intention-based notion. On this reading, if I am responsible for a certain speech act, it is because I uttered some words whose salient interpretation (first notion of content) was such and such, not because I asserted (second notion of content) that such and such. Julien owes 1€ because he uttered a sentence whose salient interpretation in the context (again, first notion of content) was that he wanted a pen with blue ink—not because he asked for (again, second notion of content) a pen with blue ink. John Perry, focusing on demonstratives, has defended an approach along these lines. I examine it in the next session. Before that, let me note that we already have two notions: what is said and speaker meaning. The question I will discuss concerns whether what is said can in turn be understood in two different ways. Perry's view is that we have an ordinary notion of what is said and a semantic one (or, better, that the notion needed in semantics is not the ordinary one). Against this, I will argue that the notion we need in semantics should be faithful to the ordinary practice and, therefore, that we should stick to the pair utterance content/speaker meaning.

5. *A plea for the forensic*

I have argued that linguistic responsibilities are sensitive to the salient interpretation. As I pointed out in the last objection, one could accept that our ordinary notion of content,

¹⁵ This does not apply, of course, to cases where we have reasons to think that the attributor is wrong.

the one we hold speakers responsible for, is sensitive to what is salient or other contextual factors or features of the conversation but still claim that this is not the notion of content needed in semantics. Perry (2009) defends such an approach for demonstratives.¹⁶

Perry discusses the following example. In the context of a conversation with some students, Professor Z says: ‘That is a picture of Carnap. He was the greatest philosopher of the twentieth century’. As it happens, in the room where the conversation takes place there are two paintings, one of Carnap (C), one of Spiro Agnew (A). Professor Z’s intention is to refer to C, but he extends his finger in a way that an uninformed interpreter (one not knowing how Carnap looks like) would take to point to A. After that, one of the students buys a portrait of Spiro Agnew and hangs it with the words ‘Rudolf Carnap’. In this example, there is a mismatch between the speaker’s intentions and the salient interpretation. Which of these, if any, determines the referent of the demonstrative? Perry writes: ‘I’m inclined to agree that Professor Z did say something that entailed that Agnew was the greatest philosopher of the twentieth century, even though he didn’t mean to. Because he was careless, he said something he didn’t mean to say’ (Perry 2009, 191).

In spite of this, Perry takes it that our ordinary notion of what is said —the one according to which Professor Z said something that entailed that Agnew was the greatest philosopher of the twentieth century— is not the one needed in semantics, precisely because “saying something is a social act, which has effects on others in virtue of the words used, their meanings, and other publicly observable indications of the speaker’s intentions” (2009, 191). For these effects and the intertwined liabilities Perry reserves the term ‘the forensic’. Our ordinary notion of saying, the thought goes, contains a forensic element. What one says, on the ordinary sense of saying, is connected to what he is responsible for. The idea is that the semantical notion of content, by contrast, should be tied to speaker’s intentions: “The careless gesture misled the alumni in a way he was responsible for, and our ordinary forensic concept of what it said is sensitive to this. But I think our semantical concept of locutionary content should not be. A prima facie plausible principle is that if one is sincere, and semantically competent, and makes no verbal slips, what one says will have the same content as the belief one intends to express by saying it” (2009 193). Let us call this principle Perry is putting forward the Speaker principle:

Speaker principle: If a speaker is sincere, and semantically competent, and makes no verbal slips, what he says will have the same content as the belief he intends to express by saying it.

I agree that, if everything goes well, what the speaker means will be identical to what he says (assuming he is speaking literally). However, this also holds for the interpreter: if everything goes well, what a competent interpreter interprets will be identical to the content of the utterance. Let us call it the Interpreter principle:

Interpreter principle: If an interpreter is semantically competent, and he doesn’t mishear the utterance, what the utterance says will be identical to the output of his interpretation.

¹⁶ Perry’s main claim is that the relevant intentions in the determination of the referent of a demonstrative are what he calls ‘directing intentions’. A similar approach to the forensic and utterance content is put forward in Korta and Perry (2011).

The question now is: in cases of mismatch, which principle is to have more weight? In particular, why should our semantical concept of utterance content privilege the Speaker principle?

According to Perry, we need to explicate the notion of what is said and purge it of the forensic element. However, the problem of substituting our ordinary notion of content by an intentionalist one is that at some point semantics needs to connect with use. The output of semantics is the input to pragmatics. For instance, semantics establishes the truth-conditions of a variety of sentences relativized to contexts, and pragmatics might tell us that one must assert only what is true (or what one knows, one what one has reasons to take as true), or that speakers aim at saying true things. This aim at truth plays a crucial role in our practice of gaining information through speech. And it is this practice that we, at the end of the day, aim at explaining.

As an illustration of the relation between semantics and pragmatics, let me quote Stalnaker's discussion of Lewis:

An utterance event takes place; certain sounds are produced. The sounds constitute the utterance, in a particular situation, of a sentence of a language with a certain semantics. The job of the semantics is to define a function that delivers a certain output, where the input is the sentence that was uttered, plus whatever features of the situation may be relevant to determining the output.

On the Lewis picture, the output is a truth-value, and Lewis's background story of what is going on when one engages in the activity of assertive speech is something like this: It is presumed that speakers aim to say things that are true, and so the addressee uses her knowledge of the semantics plus other facts on which the truth-value of the utterance depends to figure out what the world must be like in order for that utterance to be true, and then infers that the world is that way. (Stalnaker 2014, 20)

Thus, the output of semantics is of pragmatic interest.¹⁷ The reason why it is interesting to have a semantic theory delivering truth-values is that, thus construed, knowledge of semantics enables the speakers to figure out how the world is on the basis of a speaker's utterance—on the assumption that the speaker speaks truth. As I said before, one important aim in doing semantics is to understand our linguistic practices. In the previous example, semantics delivers an output that plays a role in accounting for our practice of transmitting information. Framing content in terms of possible worlds, semantics might tell us that the content of an utterance is the set of possible worlds that such and such. Pragmatics, then, tells us that, in virtue of this, an assertion reduces the set of worlds to which the real world belongs, thus allowing the interlocutors to gain information.

Attributing responsibilities is also a linguistic practice. For instance, we hold speakers responsible for what they assert, and are entitled to ask for a justification. However, if we replace the ordinary notion of content by a speaker centred one, we can no longer explain this practice, for it would turn out that semantic content goes one way, whereas the attribution of responsibilities goes another, together with the ordinary notion of content. Thus, the proper input to our linguistic practices of acquiring linguistic liabilities or motivating actions is not an intention-based content.

¹⁷ Similar views are held by Brandom (1994, 143-145), who also quotes Dummett, and MacFarlane (2014, 54-55).

As a result, explaining some of our central communicative practices requires us to have a notion that is faithful to our ordinary notion of content, and, in particular, that contains a forensic element. A possible objection is that I am focusing on cases that are not instances of successful communication—no thought has been transmitted from speaker to hearer. I grant that one can reserve the term ‘communication’ for cases that fit the following model:¹⁸ a speaker in context C has a thought p he wants to share with a hearer, he finds a sentence S that encodes p in C or, in the present discussion about modulated contents, that somehow captures p, utters it, the hearer decodes S and infers (or somehow grasps) p. However, there are two reasons for wanting a notion of content that reaches beyond this model. First, it is not at all clear that communication, as the Lockean model captures it, calls for a notion of utterance content. Communication, in this model, succeeds whenever the thought the hearer recovers is identical, or at least relevantly similar, to the one the speaker intended to convey. What the actual content of an utterance was is a matter of limited interest once we have convergence among interlocutors—assuming all we wanted to explain was communication in the Lockean sense.

And second, we still need a notion of utterance content that grounds social interaction involving language, and this notion should be in play in cases in which the speaker fails to get the thought across. As I argued in my reply to objection 5, there are conversations in which not all of the participants are in line with each other. However, in some of these, utterances have a determinate content. A context being defective doesn’t automatically prevent an utterance from having a content.

REFERENCES

- Austin, John L. 1956. A plea for excuses: The presidential address. *Proceedings of the Aristotelian Society* 57: 1-30.
- Bach, Kent. 1994. Conversational Implicature. *Mind & Language* 9/2: 124-162.
- Borg, Emma. 2004. *Minimal semantics*. Oxford: Oxford University Press.
- . 2012. *Pursuing Meaning*. Oxford: Oxford University Press.
- . 2016. Exploding explicatures. *Mind & Language* 31/3: 335-355.
- Brandom, Robert. 1994. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge: Harvard University Press.
- Lepore, Ernest and Cappelen, Herman. 2005. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.
- Camp, Elisabeth. 2006. Contextualism, metaphor, and what is said. *Mind & Language* 21/3: 280-309.
- Carston, Robyn. 2002. *Thoughts and Utterances*. Oxford: Blackwell.
- . 2008. Minimal semantics - by Emma Borg, *Mind and Language* 23/3: 359-367.
- Davidson, Donald. 1986. A Nice Derangement of Epitaphs. In Ernest Lepore, ed., *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*, 433-446. Oxford: Blackwell.
- . 1990. The structure and content of truth. *Journal of Philosophy* 87/6: 279-328.
- DeRose, Keith. 2002. Assertion, knowledge, and context. *Philosophical Review* 111/2: 167-203.
- Gauker, Christopher. 1992. The Lockean theory of communication. *Noûs* 26/3: 303-324.
- Grice, H. P. 1957. Meaning. *Philosophical Review* 66/3: 377-388.
- . 1969. Utterer’s meaning and intention. *Philosophical Review* 78/2: 147-177.

¹⁸ Gauker (1992) calls this the Lockean theory of communication.

- Korta, Kepa and Perry, John. 2011. *Critical Pragmatics: An Inquiry Into Reference and Communication*. Cambridge: Cambridge University Press.
- Kölbel, Max. 2008. Introduction: "Motivations for Relativism". In Manuel García-Carpintero and Max Kölbel, eds., *Relative Truth*. Oxford: Oxford University Press.
- MacFarlane, John. 2014. *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
- Perry, John. 2009. Directing intentions. In Joseph Almog and Paolo Leonardi, eds., *The Philosophy of David Kaplan*, 187-201. Oxford: Oxford University Press.
- Recanati, François. 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- . 2006. Crazy Minimalism. *Mind and Language* 21/1: 21-34.
- . 2010. *Truth-Conditional Pragmatics*. Oxford: Oxford University Press.
- Schoubye, Anders J. and Stokke, Andreas. 2016. What is Said? *Noûs* 50/4: 759-793.
- Searle, John R. 1978. Literal meaning. *Erkenntnis* 13/1: 207-224.
- Stalnaker, Robert. 2014. *Context*. Oxford: Oxford University Press.
- Travis, Charles. 1989. *The Uses of Sense: Wittgenstein's Philosophy of Language*. Oxford: Oxford University Press.
- Travis, Charles. 2008. *Occasion-Sensitivity: Selected Essays*. Oxford: Oxford University Press.
- Wieland, Nellie. 2010. Context Sensitivity and Indirect Reports. *Philosophy and Phenomenological Research* 81/1: 40-48.
- Zarraluki, Pedro. 2006. *La historia del silencio*. Barcelona: Editorial Anagrama. (English edition: 2014. *The History of Silence*. Translated by Nick Caistor and Lorenza García. Madrid: Hispabooks Publishing.)

CLAUDIA PICAZO JAQUE is PhD candidate at the department of Philosophy of the University of Barcelona and a member of the LOGOS Group. She works on issues in the Philosophy of Language, especially on contextualism.

ADDRESS: Departamento de Filosofía. Univ. de Barcelona. C/ Montalegre 6. 08001 Barcelona. Email: claudia.picazo@gmail.com