Are mental representations underdeterminacy-free?

Claudia Picazo (UB)

(draft)

Abstract: According to some views (Carston, Fodor), natural language suffers from underdeterminacy, but thought doesn't. According to the underdeterminacy claim, sentence types underdetermine the truth-conditions of sentence tokens. In particular, the semantics of a predicate type seems to underdetermine the satisfaction conditions of its tokens. By contrast, mental representation-types are supposed to determine the truth-conditions of its tokens. In this paper I critically examine these mixed views. First, I argue that the arguments supporting the indispensability of including in one's theory mental representations that are free of the underdeterminacy exhibited by natural language are not sound. As a result, the possibility that mental representation-types are as underdetermined as natural language sentence-types has not been ruled out. Second, I argue that Carston's *ad hoc* concept-types are as underdetermined as word-types. I finish by arguing that mental representations are also underdetermined in a second sense—mental representation-tokens only determine a partial function from possible worlds to truth-values.

1. Underdeterminacy

There are reasons to think that the truth-conditional content expressed by an utterance of a well-formed declarative sentence S is not determined by the semantics of S. Here I will focus on the underdeterminacy that can be traced to predicates. The contribution a large class of predicates make to the truth-conditions of an utterance seems to be affected by contextual factors that are not part of the predicate's meaning[1]. Consider the following examples, discussed by Carston (2002):

(1) The kettle is black.

[1] Charles Travis has made this point by presenting a variety of examples in which the satisfaction conditions of some predicate shifts across contexts. His examples involve colour predicates ('is green'), artefact terms ('is a desk'), magnitudes ('weighs 80kg'), etc.

(2) Anne is happy.

(3) I want to meet some bachelors.

Imagine an old aluminium kettle that, after years of use, has turned black on most of its surface. As has been repeatedly noted, sentence (1) can be used to describe the original colour of the kettle, or the observable colour of most of its surface. The truth-conditions of an utterance of (1) will consequently vary, and two utterances of (1) used to describe the old aluminium kettle (in the same state) can have different truth-values. Similarly, 'happy' can be used to describe a range of positive emotions. Someone can be rightly described in an occasion as 'happy' because of being momentarily in a state of intense joy, whereas in other occasions a more stable positive feeling is required in order to count as 'happy'. As to the last sentence, Carston notes that 'bachelor' can be used to refer to unmarried men or in a more restrictive sense excluding men that have committed themselves to celibacy. Thus, if a woman tells her friend that she is going through a divorce and wants to meet some bachelors, 'bachelor' would most likely be restricted to unmarried heterosexual men. Moreover, who exactly counts as an unmarried man can be decided differently on different occasions. The extension of 'bachelor' in an occasion of use could also include men that are currently going through a divorce and exclude men that are involved in long-term relationships (men who have legal partners but are not legally married), or exclude men married in a religious ceremony but not legally (or the other way round).

As a result, sentences (1)-(3) can express different truth-conditions in different contexts, even when the referent of indexicals and definite descriptions is kept fixed. In this sense, their linguistic meaning can be said to underdetermine their truth-conditions in an occasion of use—the satisfaction conditions of some expressions in those sentences are sensitive to facts about the conversation in which they are used and, probably, to the interlocutors' interests and communicative intentions. If that is so, then we need truth-conditional pragmatics[2].

I will not offer here any new argument in support of this underdeterminacy claim. Given the variety of examples contextualists have put forward, I think the claim is already well justified. Moreover, the number of examples suggests that the phenomenon of shifts in satisfaction conditions is ubiquitous. This motivates a radical underdeterminacy claim according to which most, or perhaps all, well-formed sentences in natural languages are such that their linguistic meaning underdetermines their satisfaction conditions. I will take this radical underdeterminacy claim as a starting point. If this claim is correct, then the truth-conditions of an utterance are partly a pragmatic business.

[2] Truth-conditional pragmatics is the view that the meaning of a sentence does not determine the truth-conditional content of an utterance of it (even after fixing the referent of indexicals), for truth-conditions need to be pragmatically supplemented.

The question now is: What sorts of representations exhibit truth-conditional underdeterminacy? I think that, besides natural language sentences, underdeterminacy could be a feature of Mentalese sentences (mental representations). In particular, both sorts of representations might exhibit what I will call Type-Underdeterminacy[3]:

> **Type-Underdeterminacy:** A non-indexical structured representational item[4] S is type-underdetermined if and only if there are tokens of S that have distinct truth-values.

That tokens of S have distinct truth-values means that the type does not determine a unique truth-value (given the state of the world). The previous considerations about sentences (1)-(3) suggest that *linguistic* representational items, i.e., sentences, exhibit Type-Underdeterminacy. Once it is accepted that natural language underdetermines truth-conditions in this sense, it might be disputed whether other representational items are also type-underdetermined. We can take Type-Underdeterminacy to be a feature of representations in general, and generalize the underdeterminacy claim to other structured representational systems—saliently, to mental representations. Or, we can take it to be restricted to language and try to find another representational system that is, in this sense, underdeterminacy-free. Again, the candidates are mental representations. Following this second option, the underdeterminacy detected in natural language can be explained as concerning the relation between natural language sentences and other representational items.

Fodor (2001) and Carston (2002) put forward an approach in which language suffers from Type-Underdeterminacy, but thought does not. Given that they endorse the Language of Thought (LOT) hypothesis, their claim is to be understood as saying that natural language sentences are type-underdetermined but LOT sentences, i.e., mental representations, are not. The underdeterminacy detected in sentences (1)-(3) concerns, according to this view, the relation between utterances and the mental representations expressed by them. What is underdetermined is what mental representation corresponds to an utterance of a natural language sentence S, S's meaning being compatible with, let's say, two different mental representations. Mental representations are seen as free of Type-Underdeterminacy. Pragmatics is supposed to bridge the gap between sentences and mental representations: our

[3] In section 5 I introduce a different notion of underdeterminacy, namely Token-Underdeterminacy. I am indebted to an anonymous referee for helping me distinguish the two notions.

[4] By 'structured representational item' I mean a representational item that is identified by its structure, such as a sentence. The definition is intended to apply both to sentences in natural language and to sentences in Mentalese.

mind reading abilities allow us to infer which thought (i.e., which mental representation) a given utterance expresses. I will call approaches that fit this second option 'mixed views'[5].

A feature of the mixed view that is open to criticism has to do with its reliance on the existence of items (mental representations) that are identified by a structure and have fully determined truth-conditions independently of the context of use[6]. The first challenge for these approaches arises from the fact that we haven't been shown that other representational system behave differently from natural language when it comes to truth-conditions. Natural language sentences are representational items (i.e., items that we typically use for representing) that are identified by structure, i.e., syntax. If two tokens have the same syntactic form, then they are tokens of the same sentence. According to the Type-Underdeterminacy claim for natural language, there is a gap between these structural items and the truth-conditions they express in an occasion of use. In the case of natural language sentences, structure doesn't automatically get us truth-conditions. The problem is that, once we have seen that this is what happens with sentences, the assumption that there are other representational items identified by structure with context-independent truth-conditions needs to be justified. Why is it not the case that mental representations, like sentences of natural language, also express different truth-conditions at different tokening contexts? This line of reasoning can be used in order to motivate a generalization of Type-Underdeterminacy to mental representations[7].

The point I will address in this paper concerns the relation between representations in the mental realm and truth-conditions. The plan is the following. In section 2, I will present the mixed view in more detail. In section 3, I will assess what I call the indispensability arguments. Proponents of the mixed view argue that there are strong reasons for including in one's theory these structured representational items that are (allegedly) free of Type-Underdeterminacy. The main idea here is that they play a role that only a representation that does not exhibit Type-Underdeterminacy could play. In this sense, they are indispensable. If their arguments work, then we have reasons for positing mental representations that are free of Type-Underdeterminacy. Against these views, I will argue that their arguments fail to establish that representations that are free of Type-Underdeterminacy are indispensable. As a consequence, alternative approaches are, at least, tenable. In section 4, I will argue, contra the mixed views, that there are reasons to doubt that Mentalese sentences are unlike natural language sentences regarding Type-Underdeterminacy. In particular, I will argue that Carston's description of the process of *ad hoc* concept creation suggests that Mentalese sentences are as type-underdetermined as natural language sentences. In section 5, I will introduce a second sense of

[5] I call these approaches 'mixed' because they include two sorts of representations—type-underdetermined and non-type-underdetermined.
[6] Travis (2000) calls this the Janus-faced picture of thoughts.
[7] This generalization has been pursued by Travis (2000) and Searle (1983).

'underdeterminacy'—Token-Underdeterminacy—and argue that mental representations are also underdetermined in this second sense.

2. The mixed view

The mixed view makes two claims. First, it claims that linguistic representations, i.e., sentences with their linguistic meanings, or sentences plus the contextual information determined by their linguistic meanings, are type-underdetermined. Second, it claims that, besides sentences, there are other representational items—mental representations—that are not type-underdetermined. Which of this second kind of item is expressed by a use of a sentence is to be determined via pragmatic interpretation.

Carston's theory is explicitly a version of the mixed view. Carston (2002) writes:

> [T]he position I've been arguing for is that there are no eternal sentences in natural languages (that is, no sentences which encode a proposition or thought which is constant across all contexts), from which it follows that the linguistic underdeterminacy of the proposition expressed by an utterance is an essential feature of natural language (2002, p. 42).

> It is fully propositional conceptual representations, rather than sentences, or even utterances of sentences, that are the primary bearers of truth conditions (2002, p. 60).

According to Carston, there are fully propositional representations to be had, they are just not encoded by natural language sentences. Rather, they correspond to mental representations. Strings of concepts are supposed to be free of Type-Underdeterminacy. The arguments that Carston presents in support of the underdeterminacy claim include the examples presented at the beginning of this paper. Concerning predication, she relies on Travis's examples[8].

Given that linguistic meaning underdetermines the thought expressed by a sentence in an occasion of use, there must be some mechanism that enables the interpreter to grasp the thought that the speaker intends to communicate by means of a sentence, i.e., a mechanism that bridges the gap between sentences and propositional thoughts. In Carston's theory, the mechanism consists in the creation of new concepts, slightly different from the concepts encoded in language. Concepts of this second kind are called *ad hoc* concepts. The idea is that

[8] Besides rejecting eternal predication, Carston also rejects eternal reference. However, here I will only consider the underdeterminacy that can be traced to predication.

the interpreter uses his ability for pragmatic interpretation in order to recover the thought that the speaker intends to communicate by using as input the encoded content of a sentence and the available contextual information. Often the recovery of the intended thought is achieved via the adjustment of the encoded concepts. In those cases, a new concept is 'constructed on-line (on the fly) in response to specific expectations of relevance raised in specific contexts' (2002, p. 322).

Carston follows relevance theory in her understanding of what a concept is. Atomic concepts consist on three kinds of information: logical content (a set of inference rules capturing analytic implications of the concept), encyclopaedic knowledge (scientific information, general knowledge about the object, personal observations) and lexical properties (phonological and syntactic properties). Complex concepts are structured strings of atomic concepts. Language codifies concepts such as CAT[9], with certain logical content (if something is a cat, then it is an animal), encyclopaedic knowledge (cats are domestic animals, visual images of cats), and lexical properties. As I said, in a conversation, in the process of utterance interpretation, interlocutors construct *ad hoc* concepts by adjusting the information of the lexically encoded concepts to the specifics of the context. As is common, Carston distinguishes two pragmatic processes of adjustment: narrowing and broadening. In cases of narrowing, the concept is made more specific. Let us imagine a use of (3) by a woman who is chatting with a friend about her desire to meet a man, get married and have children (Carston 2002, p. 326). The encoded concept BACHELOR makes reference to non-married men. The encyclopaedic entry might contain information about different types of bachelors: irresponsible and forever-young-and-free, capable of long-term commitment, etc. However, in this conversation the concept expressed by the word 'bachelor' is more specific: the speaker wants to meet some men eligible for marriage (capable of long-term commitment, heterosexual). Given what the hearer knows about the speaker (marital interest), during the process of interpretation a new *ad hoc* concept will be constructed excluding from the encyclopaedic entry features standardly associated with bachelors as irresponsible and forever-young-and-free. The account for broadening is symmetrical. Consider the sentence 'France is hexagonal'. France is not a geometrical hexagon. However, loosely speaking, its shape can be considered hexagonal. What is going on here is that whereas the concept of HEXAGON includes only strict hexagons, the *ad hoc* HEXAGON* includes shapes that deviate to some degree[10].

[9] I use capital letters for encoded concepts. *Ad hoc* concepts are marked with an asterisk.

[10] Let me note that Carston's proposal has two problematic aspects. The first has to do with the relation between encoded concepts and *ad hoc* concepts. It is not at all clear whether lexically encoded concepts as HEXAGON are being conceived here as determining an extension, or whether it is only *ad hoc* concepts as HEXAGON* that do so. In her (2002), Carston holds both that there is no eternal predication (one of the reasons why natural languages are underdetermined), and that lexically encoded concepts have extensions. For example, she writes, about narrowing, that 'the extension of the concept pragmatically constructed is a subset of the extension of the lexical concept from which it has been

Fodor (2001) also holds a version of the mixed view. He takes compositionality to be non-negotiable and argues that, between language and thought, whichever is compositional is the one that has content in the first place. He further takes language to be not compositional. Language being non-compositional seems to mean that linguistic meaning does not determine the (truth-conditional) content of complex expressions (of declarative sentences)—i.e., that language is type-underdetermined. By contrast, the contents of the simple constituents of thought, together with a mode of composition, are supposed to determine the (truth-conditional) content of complex thoughts[11].

Both Fodor and Carston hold a representational theory of mental content according to which mental representations are sentences in LOT. Whereas natural language sentences are taken to exhibit Type-Underdeterminacy, Mentalese sentences are supposed to be free of it[12]. As a consequence, natural language sentences are not apt to encode Mentalese sentences. In this sense, thoughts are ineffable—although we can often express them, we cannot find a sentence whose linguistic meaning corresponds to them.

Now, the problem with ineffability is that it has the unpalatable consequence that those mental representations that are free of Type-Underdeterminacy can neither be encoded in a natural language sentence nor consciously entertained. Strictly speaking, we cannot encode those representations in natural language. We can express our thoughts via natural language, i.e., communicate them by using a sentence, but this will involve *ad hoc* concepts. Moreover, it seems that we cannot consciously entertain them. If conscious thought at least sometimes occurs in natural language, as introspection suggests, then conscious thought might also fail to

---

derived.' (2002: 325). This presupposes that the lexical concept has an extension. However, Carston's radical underdeterminacy claim is incompatible with lexically encoded concepts having extensions: if the lexically encoded concept BACHELOR has an extension, then 'is a bachelor' should have a constant extension. And if so, then there is eternal predication. Nonetheless, simply dropping the assumption that lexically encoded concepts determine extensions is also problematic, for encoded and *ad hoc* concepts have the same structure. As long as encoded and *ad hoc* concepts are described in analogous terms it is mysterious why they should behave differently. The second problematic aspect has to do with the supposed atomic character of *ad hoc* concepts. *Ad hoc* concepts, just like lexically encoded concepts, are supposed to be atomic. Carston writes: 'This term [*ad hoc*] is used to refer to concepts that are constructed pragmatically by a hearer in the process of utterance comprehension. The idea is that speakers can use a lexically encoded concept to communicate a distinct non-lexicalized *(atomic)* concept, which resembles the encoded one in that it shares elements of its logical and encyclopaedic entries, and that hearers can pragmatically infer the intended concept on the basis of the encoded one.' (Carston 2002, p. 322. Emphasis added.) However, it is not clear what is meant by 'atomic' here, given that they are created by adding or subtracting information from already existing atomic concepts, and so they seem to be complex. I will argue (sections 3 and 4) that both options are problematic. If *ad hoc* concepts are atomic, then Carston cannot avail herself of the productivity argument (section 3.1)—something she can do if *ad hoc* concepts are complex. However, if they are complex, as the description of the creation of BACHELOR* suggests, i.e., if they are compositions of lexically encodable concepts, then one should expect that they be as underdetermined as combinations of words are (section 4.1).

[11] Recanati (2004) could be read as an instance of the mixed view. The process of meaning modulation could be understood as a form of sense or concept creation although, as Carston (2015) notes, *ad hoc* concept creation and modulation are described in different terms.

[12] Carston (2002) admits some context-dependence in thought, namely the presence of indexicals. However, she doesn't admit context-sensitive predicates at the level of thought.

encode determinate mental representations-types[13]. What are the grounds, then, for assuming that there is a level of representations that are similar to natural language sentences yet not type-underdetermined?

3. The indispensability arguments

In this section, I will review and reject two arguments that have been offered in support of the claim that mental representations (Mentalese sentences) cannot exhibit Type-Underdeterminacy (Fodor 2001, 2003; Carston 2002). Fodor presents the arguments as concerning compositionality and does not distinguish them. However, in order to appreciate the dialectics of the discussion it is better to assess them separately.

There are different principles of compositionality. If we focus on natural language we can distinguish, at least, compositionality of meaning from truth-conditional compositionality[14]:

Meaning Compositionality: The meaning of a well-formed declarative sentence S is determined by the meaning of the expressions in S and the syntactic structure of S.

Truth-conditional Compositionality: The truth-conditions of a well-formed declarative sentence S are determined by the semantics (or the meaning) of the expressions in S and the syntactic structure of S.

It is this second principle that is at stake here[15]. The idea is that the truth-conditional content of a representation is exhausted by the semantic content of its simple constituents and their arrangement. If natural language is type-underdetermined, then it is not truth-condition compositional[16]. However, it can still be meaning compositional. If mental representations are also type-underdetermined, then they are not truth-condition compositional either, which means that the truth-conditional content of its tokens is not exhausted by the concepts that form them and their arrangement. Their truth-conditions would depend on something else, as Travis sometimes puts it.

Why do mental representations have to be truth-condition compositional? Fodor (2001) addresses the question whether it is language or thought that has content in first instance.

[13] As Martínez-Manrique and Vicente (2004) argue, ineffability can be problematic for views that hold an availability principle, as Recanati (2004).

[14] See Searle (1980) and Clapp (2012) for similar distinctions. Clapp notes that this distinction undermines the systematicity and productivity arguments. He focuses on systematicity.

[15] Fodor (2001) does not state any principle of compositionality. However, it is in this sense of 'compositional' that language fails to be compositional.

[16] Following Carston (2002), I have framed the discussion in terms of underdeterminacy. However, instead of talking about sentences of natural language and Mentalese being Type-Underdetermined one could talk of them failing Truth-conditional Compositionality.

He claims that compositionality is non-negotiable and takes it that between language or thought the one which is compositional is the one which has content in first instance. Given the linguistic evidence (including the underdeterminacy arguments), it seems that language is not compositional. Because of this, Fodor assumes that thought is. However, one must ask why compositionality is supposed to be non-negotiable. Fodor mentions productivity and systematicity: 'Nobody knows exactly what compositionality demands, but everybody knows why its demands have to be satisfied. Here too the arguments are familiar; and, in my view, they're decisive. Both human thought and human language are, invariably, productive and systematic; and the only way that they could be is by being compositional'[17] (Fodor 2001, p. 6). I will focus on productivity. So the first argument is that mental representations have to be compositional because they are productive, and compositionality is the best (or the only) explanation we have for this.

This argument, as I will show, is off-target. Given Fodor's notion of productivity, meaning compositionality is sufficient for a system of representations to be productive. Moreover, if it is creativity we are interested in, or our ability to think new thoughts, then the best explanation given the scenarios of underdeterminacy involves the creation of *ad hoc* concepts, as Carston defends. But then productivity is not the key to our ability to think new thoughts—the creation of new concepts is. It follows that a system that is not truth-condition compositional can still be productive, and that some forms of creativity are unrelated to productivity. So Fodor hasn't in fact provided any conclusive reason in support of the claim that mental representations must be free of Type-Underdeterminacy.

There is, however, a second indispensability argument. It has to do with the individuation of content. The idea is that if mental representations were not truth-condition compositional, then thought would be ambiguous or equivocal: the same representation could express different truth-conditions. The outcome would be that we wouldn't be able to tell some thoughts apart. I will argue that the conditional is false. Mental representations (types) could be equivocal, yet the tokening context could resolve the equivocation.

3.1. Productivity

This is Fodor's notion of productivity: 'Productivity is the property that a system of representations has if it includes infinitely many syntactically and semantically distinct symbols' (Fodor 2001, p. 6). Natural languages are productive in this sense: there is an infinite number of well-formed, meaningful sentences. Given that the number of simple expressions is finite, compositionality is regarded as explaining the productivity of language: there can be an infinite

[17] It is already odd that Fodor is here taking language to be compositional, when he explicitly rejects it. I will not try to solve this apparent contradiction.

number of meaningful complex expressions (sentences) because their semantics is determined by the semantics of simple expressions plus their syntax.

The relation between productivity and compositionality is often used to explain our ability to understand new sentences. As it has been repeatedly noted, we are able to understand utterances of sentences we have never heard before. The best explanation for this ability is that we understand these new sentences because we know the meaning of the simple expressions that form them and the syntactic rules. A finite mind can thus be reconciled with the capacity to interpret an infinite number of sentences.

Despite this argument, it has been questioned that knowledge of meaning and syntax suffices for working out the truth-conditions of an utterance, where this questions the idea that language is truth-condition compositional. Searle (1980) argues that knowing the meaning of 'cut', 'the', and 'sun' doesn't automatically enable us to understand an utterance of 'Cut the sun!'—we might fail to see what are the satisfaction conditions of this order, what exactly we are supposed to do. We lack some background that enables us to see what action would fit the order. This suggests that language is not truth-condition compositional, but only meaning compositional. There certainly is something we understand when we first hear 'Cut the sun!', but we don't automatically grasp satisfaction-conditions.

What about thought? As Fodor (2003) notes, we are able to entertain new thoughts—potentially, an infinite number of them. However, being finite creatures, we only possess a finite number of simple concepts. Mental representations being productive would explain our infinite ability. Again, the reason why we are able to entertain an infinite number of thoughts could be that simple concepts can be arranged so as to form an infinite number of complex mental representations. Now, it is important to note two things. First, as Fodor defines productivity, something equivalent to meaning compositionality would be sufficient for mental representations to be productive. Being able to entertain an infinite number of mental representations only requires that simple concepts can be arranged so as to form an infinite number of complex concepts. However, Meaning Compositionality for mental representations is compatible with mental representations exhibiting Type-Underdeterminacy (i.e., with them failing to be truth-condition compositional).

Second, it is one thing to have the capacity to form or entertain indefinitely many mental representations, it is another to have the capacity to entertain indefinitely many truth-conditional contents. One could have the latter without having the former. Imagine a group of people who only possess two concepts, HUNGRY and THIRSTY[18], and no rule of composition (so no complex mental representation). Are these people only capable of entertaining two truth-conditional contents? No. Imagine that one of them tokens the concepts HUNGRY at 10am. And then, he tokens it again at 5pm. Even if the mental representation

[18] This example is inspired on Perry (1986). Clapp (2012) uses a similar example.

does not include the time, it is possible that the tokening context adds it and that the truth-conditional contents he entertains are different—the first is true if and only if he is hungry at 10am, whereas the second is true if and only if he is hungry at 5pm. So he can entertain an indefinite number of truth-conditional contents by tokening only one simple concept. As a consequence, our ability to think new thoughts does not inescapably go via compositionality.

What about Carston's view? Carston's concern is directly related to the underdeterminacy scenarios. In her explanation of these cases, speaker and hearer token a concept that they create on-line, as a response to the specifics of the occasion. Now, if that on-line created concept is, as Carston labels it, atomic, then she is not in a position to use the productivity argument is support of her view. Recall that productivity is supposed to reconcile our ability to think an infinite number of thoughts with our having a finite mind. The traditional answer is that we can create an infinite number of complex concepts. However, Carston's explanation is a different one: in the underdeterminacy scenarios we think new thoughts because we create new simple concepts. Creativity, then, has to do with an ability to create indefinitely many concepts that fit the indefinitely many situations we encounter.

As a conclusion, Fodor hasn't offered any conclusive argument to the effect that mental representations are truth-condition compositional, and not merely meaning compositional—Meaning Compositionality being compatible with Type-Underdeterminacy. On the other hand, if *ad hoc* concepts are really atomic, and not a combination of pre-existing concepts[19], then Carston cannot use the productivity argument.

3.2. Equivocation

It has been argued (Fodor 2001, 2003; Carston 2002, 2008) that ambiguity and equivocation cannot occur at the mental level, since it is thoughts that disambiguate sentences. According to this line of reasoning, it makes no sense to take thoughts to be ambiguous themselves: if they were, there would be nothing that could disambiguate them. The argument can be reconstructed as a *reductio*. Some English words, as for example 'bank', have distinct meanings (financial institution, side of the river). Suppose that the same happens at the mental level, that is, suppose that English speakers only have one concept BANK. Now, if that is so, then thoughts about financial institutions and thoughts about the side of a river are indistinguishable. But, clearly, thoughts about financial institutions and thoughts about the side of a river are easily distinguishable. So it is not the case that English speakers have one ambiguous concept BANK. Rather, they must have two different concepts, corresponding to the two different meanings of

[19] As I noted in sec. 2, Carston's description of the process of creation of *ad hoc* concepts raises some doubts that they are not a combination of pre-existent representations. If they are, in this sense, complex, then she can avail herself of the productivity argument. However, a different worry would arise, as I argue in sec. 4.

the word 'bank'. Moreover, these two concepts are the key to resolving the ambiguity with the English word 'bank'.

If this argument works for ambiguity, then it also works for other forms of equivocation such as the ones involved in (1)-(3). Again, if speakers only have one concept corresponding to the different senses of 'black', 'bachelor' and 'happy', then thoughts about a kettle being superficially black would be indistinguishable from thoughts about a kettle being originally black; thoughts about bachelors capable of long-term commitment would be indistinguishable from thoughts about not legally married men; and we could not discriminate thoughts about the different degrees and kinds of positive emotions that 'happy' seems to cover. The idea, in short, is that whereas language can, and often does, equivocate, thought cannot. Thought resolves equivocation. As Fodor puts it:

> [W]hereas the content of a sentence may be inexplicit with respect to the content of the thought it expresses, a thought can't be inexplicit with respect to its own content; there can't be more—or less—to a thought than there is to its content because a thought just is its content. (Fodor 2001, p. 14; quoted by Carston 2008, p. 339).

Fodor and Carston endorse a representational theory of mental content according to which 'having a thought with a particular content *P* involves the occurrence (the mental 'tokening', as it is often put) of a sentence of the language of thought (Mentalese) that means that *P*.' (Carston 2002, p. 74). In this framework, Mentalese sentences (and in particular, their types) disambiguate or resolve the equivocations of natural language sentences. Whenever there are two distinguishable thoughts, there must be two different Mentalese sentence-types being tokened.

Now, as Clapp (2012) has argued, the thesis that a thought just is its content does not entail the thesis that mental representations, conceived as Mentalese sentences (and specifically, sentence-types), have context-independent truth-conditions. 'Thought' can mean thought-content or thought-vehicle (see also Recanati, 2007). In the content sense, and assuming content is conceived as truth-conditions, it is a conceptual truth that a thought just is its truth-conditional content. However, Mentalese sentences are thought-vehicles. Thought-vehicles are not truth-conditional content themselves—at most, they determine truth-conditional content, or are the bearers of truth-conditional content. Thus seen, the argument from equivocation relies on the premise that the only things that can disambiguate between two uses of an ambiguous or equivocal sentence—the only thing that can distinguish their truth-conditions—are thought-vehicles, and more specifically, types of thought-vehicles. Hence, the argument concludes that, whenever uses of a sentence involving the word 'bank' ('black', 'bachelor' and 'happy') differ in truth-conditions, these uses must be expressions of different Mentalese sentence-types.

Making the premise explicit and applied to equivocation, the argument is the following:

1. Two uses of an equivocal sentence can (and often do) differ in truth-conditions.

2. The only thing that can resolve the equivocation (=account for the difference in truth-conditions) is a mental representation-type which is conceived as non-equivocal.

C. Therefore, there must be non-equivocal mental representation-types.

However, premise 2 is false, because mental representation-types can be equivocal, yet the tokening context can resolve the equivocation. To show this, I will turn to various examples in which the same mental representation-type gives rise to different truth-conditions in different contexts of tokening. Let me start with an argument by Recanati (2007) concerning demonstratives and progressively move to general terms.

Recanati (2007, p. 13-14) asks whether it is possible to find thought-vehicles with different truth-conditions at different contexts. His affirmative answer is supported by the following example. Imagine a subject who entertains the thought corresponding to 'This man is happy' while looking at a certain man—Bob. His thought is true if and only if Bob is happy. Now, had the context been different, the truth-conditions of his thought-vehicle could have been different. Suppose that it is Bill, not Bob, the man that the subject is looking at. In this case, the truth-conditions of his thought-vehicle would depend on Bill's properties, not Bob's. It seems that the same thought-vehicle can have different truth-conditional content. If the fact that the visual information that the subject has is different is considered a problem, we can imagine that Bob and Bill are visually indistinguishable for the subject (they are twins).

This example shows that the tokening context can affect the truth-conditional content of a mental representation-type. Let's consider another example. Suppose that some mental representations, those concerning the present, are time neutral. Take the thought corresponding to 'Diana is happy'. To say that this mental representation is time neutral is to say that it contains no element encoding the time at which Diana is supposed to be happy. Would tokens of this mental representation equivocate as to when Diana is supposed to be happy? Not necessarily. The time can be provided by the tokening context. A tokening of the mental representation corresponding to Diana is happy at time t will be true if and only if Diana is happy at t, whereas a tokening of the mental representation corresponding to Diana is happy at time t* will be true if and only if Diana is happy at t*. Not everything that affects truth-conditional content needs to be encoded in the mental representation[20].

Can the tokening context resolve the equivocation of predicates and general terms? A potential problem comes from the fact that, whereas time and place are automatically given by the tokening context, the same is not true of the satisfaction conditions of a predicate or a

[20] Perry (1986) argues that something similar happens with the location needed to get the truth-value of an utterance of 'It's raining'. Although the location is not, as he puts it, articulated in the sentence, it is necessary in order to get a truth-value. His proposal is that it is provided by the context.

general term. In the previous example, time can be said to only take one value—something like now, let's say. When we move to the cases of underdeterminacy discussed at the beginning of this paper things get more complicated. The predicate 'is black' can get different satisfaction conditions—roughly, black in some specific relevant part, originally black, black all over, black in such-and-such observation conditions, etc. The act of tokening the mental representation corresponding to 'The kettle is black', of itself, is not going to decide what counts as 'black'. I think, however, that there is something else.

Typically, utterances take place in the context of a conversation. Following Lewis (1979), we can think about the context as the conversational score. The conversational score can be seen a specification of the state of the conversation. What is interesting now is that the interpretation of a given utterance can depend on the conversational score. One of the examples Lewis discusses involves the verbs 'coming' and 'going'. These expressions require a point of reference. Now, this point of reference can be part of the conversational score, for instance because of being given by a previous utterance:

> One way to fix the point of reference at the beginning of a narrative, or to shift it later, is by means of a sentence that describes the direction of some movement both with respect to the point of reference and in some other way. "The beggars are coming to town" requires for its acceptability, and perhaps even for its truth, that the point of reference be in town. Else the beggars' townwards movement is not properly called "coming". This sentence can be used to fix or to shift the point of reference. When it is said, straightaway the point of reference is in town where it is required to be. Thereafter, unless something is done to shift it elsewhere, coming is movement toward town and going is movement away. If we are told that when the soldiers came the beggars went, we know who ended up in town and who did not. (Lewis 1979, p. 351).

Something similar could be going on in a conversation in which 'That kettle is black' is uttered. Imagine a pair of friends who have just moved together. In their new flat, there is no kettle and so one of them says 'We need a kettle, let's go buy one'. On their way to the shop, they talk about what kind of kettle they will buy. They say things like 'I don't like aluminium kettles, they are so ugly', and 'Let's buy something modern and elegant'. When they arrive to the shop and one of them says 'Look, that kettle is black' the previous conversation has made it clear that the colour predicate applies to the apparent surface of the kettle. The dialogue is enough for us to see what kind of kettles count as black (nor burnt ones, for example, not ones that look black on black and white photographs but not in normal conditions). Had the conversation been different, the satisfaction-conditions of the predicate could have been different.

Thought episodes also take place in context. Instead of a conversational score, we have a mental score. Whether they are chatting or they are independently thinking about the kettle they need, it is plausible to think that the two friends consider the need to buy a kettle, think about the kinds of kettles they like, etc. Suppose, now, that the mental representation they token at the shop includes the lexically encoded concept BLACK and not an *ad hoc* concept. Is their thought distinguishable from one in which only originally black things count as black? It is, for the mental score determines that what is at stake is the aspect of the kettle. The same mental representation type, tokened in two different thought contexts, can express different truth-conditions.

In conclusion, the indispensability arguments fail to secure the claim that mental representations cannot be type-underdetermined. Alternative approaches are tenable.

4. Persistent underdeterminacy

In the previous section I have considered and rejected two arguments in support of the claim that mental representations do not suffer from Type-Underdeterminacy. I will now argue that mental representations involving *ad hoc* concepts should be expected to be type-underdetermined.

The reason has to do with the source of Type-Underdeterminacy. Let us go back to linguistic Type-Underdeterminacy. The linguistic meaning of sentences (1)-(3) type-underdetermines the truth-conditions expressed by its tokens. It might seem that the reason these sentences suffer from Type-Underdeterminacy lies in the fact that they do not make everything explicit. For example, (1) doesn't explicitly tell us where, in what spots, the kettle is supposed to be black; (2) doesn't specify the sense of 'happy' in which Anne is happy; and (3) doesn't make explicit what kind of bachelor the speaker wants to meet or who counts as a bachelor. However, as soon as we try to find other sentences that explicitly state where the kettle is supposed to be 'black', etc., we realize that, again, the new sentences do not make everything explicit either[21]. We can replace (1)-(3) with more complete versions:

(1*) The kettle is superficially black.
(2*) Anne has a stable feeling of happiness.
(3*) I want to meet some bachelors interested in long-term relations with women.

These sentences have solved some doubts, but not others. (1*) doesn't specify how much of the surface needs to be black for the utterance to be true, or what exactly counts as the

[21] See Searle (1978, 1980), Travis (1997), Cappelen and Lepore (2005) and Recanati (2010) for similar arguments.

surface (what about black plastic glued to the original surface?). (2*) doesn't specify how long a feeling needs to last in order for it to count as 'stable'. And the same goes for (3*), which doesn't give any clue as to what the threshold is for a relation to count as a 'long-term' one. So they are, in this sense, equivalent to (1)-(3): they also exhibit Type-Underdeterminacy.

The question now is whether it is possible to come up with a more complex sentence-type that leaves no doubts open. We, normal speakers, might be unable, for our cognitive capacities are limited. The question, however, is whether, were our capacities to be improved, we could come up with a sentence that encodes the content that utterances of (1)-(3) convey. There is a reason why we should give a negative answer. If we go to (1*)-(3*), we will see that the reason why these sentences also exhibit Type-Underdeterminacy is that the new expressions in the sentence can be understood differently on different occasions. For example, the threshold for something to count as 'stable' can vary across contexts, depending on what is at stake. To be sure, a speaker could use (2*) and solve some doubts concerning a use of (2) ('What do you mean 'Anne is happy'? Has she won the lottery?'). But this doesn't mean that he has thereby found a sentence that is free of Type-Underdeterminacy. The same happens with 'superficially'. For example, now it is open how much of the surface needs to be black in order for an object to be 'superficially black', or what counts as the object's surface. The added material comes with different possible interpretations, with different satisfaction conditions in different contexts. Complexifying a sentence is not the route to overcome Type-Underdeterminacy.

Let us suppose that the mental representation-types corresponding to sentences (1)-(3) are composed of the concepts BLACK, HAPPY, BACHELORS, etc. Since adding linguistic expressions is not a way to eliminate the Type-Underdeterminacy, adding concepts to mental representation like SURFACE, STABLE or LONG-TERM will not be a way of getting mental representations that are free of Type-Underdeterminacy. Combinations of concept-types corresponding to words are as type-underdetermined as the combinations of words themselves are. The problem is that it hasn't been shown that there are ways of conceiving the mental representation-types corresponding to utterances of (1)-(3) other than combining the concepts we already have—those corresponding to 'bachelor', 'black', 'happy', 'surface', 'long-term', 'stable'. As they are described, *ad hoc* concepts are no exception.

As I said, *ad hoc* concepts are supposed to be non-linguistically encoded atomic concepts. They don't correspond to words, and they are not combinations of concepts. So they might seem to escape this problem. However, they are created by activating some information belonging to the encyclopaedic entry of linguistically encoded concepts, and this information sounds very similar to the kind of information that we can encode in natural language and has different satisfaction conditions in different occasions. Concerning (2), Carston writes:

> Suppose, as above, that the context is one in which the addressee knows that the speaker, who is a woman, wants to get married and have children. Having accessed the lexical concept BACHELOR, which makes available its associated logical and encyclopaedic information, he uses a subset of this information to construct a more specific concept BACHELOR*, which is relevant (that is, gives rise to cognitive effects) in the context. The encyclopaedic entry might well contain information about certain sorts of bachelor, the irresponsible, fun-loving, forever-young-and-free sort, the elderly, solitary, misogynous sort, and those who are youngish, heterosexual, and capable of long-term commitment, i.e. eligible for marriage. Given the hearer's alertness to the speaker's marital interest, it is probable that information about this third sort of bachelor will be more highly activated than that about either of the others, so it will be accessed first, together with the logical entry, and used to construct the *ad hoc* concept. Provided this gives rise to a satisfactory range of cognitive effects, it is retained as the intended interpretation. In different contexts, other narrowings might be effected, yielding concepts which denote different subsets of the category of unmarried adult males. (Carston 2002, p. 326)

In order to create the *ad hoc* concept, some information included in the encyclopaedic entry of BACHELOR is activated. Now, this information doesn't seem to be essentially different from the information we can encode in natural language. As I mentioned, 'being capable of long-term commitment' can be interpreted in different ways in different occasions, and the same happens with 'solitary', 'irresponsible', and so on. This information, per se, is no different from the information contained in 'bachelor'. Just as the latter can be understood in different ways, so can the former. Even if *ad hoc* concepts are supposed to be atomic, they are created by selecting features, so, in the end, they are not very different from a combination of concepts. And combinations of concepts encodable in natural language (like LONG-TERM COMMITMENT, SOLITARY, IRRESPONSIBLE) are still type-underdetermined.

The problem lies on the assumption that some structured representational items beyond language are free of Type-Underdeterminacy. We have seen that, in language, given a general expression *F* and a referential expression *a*, it is not automatically decided, by the semantics of *F* alone, whether *a* satisfies *F*. In general, some doubts concerning how to apply *F* must be solved. Concepts, as common nouns or adjectives, also have a general content. So it is plausible to think that the mental context is doing some work, just like the conversational context does. This is specially pressing if, as Fodor and Carston do, we conceive mental representations as sentences of a language[22].

[22] Wittgenstein ([1953] 2009, 139-141) considers whether understanding a word (in his example, 'cube') can be explained as having a picture before one's mind (an image of a cube) and asks the question whether it is correct to apply the word 'cube' to a prism. He argues that the picture, the mental image, can be both made to fit and not to fit this particular use of the word 'cube', depending on the method of projection one applies. Now, if the method of projection is made part of the image that comes before one's mind, then the same problem arises again: how am I to know how to apply the image of the

At this point, a proponent of the mixed view might have recourse to an externalist theory of mental content for *ad hoc* concepts[23]. In order to counter the previous argument, he could claim that the content of the *ad hoc* concept-type (a concept-type that is created on-line) is determined by the world itself[24]. If I say 'The kettle is black' referring to a very dirty kettle I am looking at, the content of the ad hoc BLACK* (the content of this new concept-type) would be, according to this externalist reply, determined by the world itself—by how the kettle is. Or imagine Peter, a friend of Anne who, while chatting with her, tokens a mental representation that could be expressed by saying 'Anne is happy'. According to this externalist position, the content of his representation will depend on how Anne happens to be. The *ad hoc* HAPPY* is not created by adding or subtracting information from the lexically encoded HAPPY. Rather, it would be a new concept-type whose content is determined by the world itself. If, at the time of Peter's tokening, Anne experiences a stable positive emotion, then that will be the content of Peter's thought[25]. Thus, an externalist could claim that mental representations have determinate content, even though these contents might not be fully transparent to us and despite the fact that we cannot encode them in natural language.

I think that this externalist position faces two important worries. First, there are cases where the world itself does not help resolve the equivocation. Belleri (2014) imagines such a case. I adapt from her example. The predicate 'is black' can be used to describe, among others, the original aspect of an object or its apparent surface. In this second sense, a burnt or painted kettle can be rightly described as 'black'. Now, we can imagine a kettle that is both originally black and painted black, or burnt, or very dirty. Imagine, further, someone thinking, of that kettle, what could be expressed by the sentence 'The kettle is black'. Is his mental representation about the original colour or about the apparent surface? Here, the world doesn't help decide. If the content of the *ad hoc* BLACK* is determined by the world itself, then it doesn't resolve the equivocation between the possible interpretations of 'black'. However, in many cases, the context will resolve the equivocation—for example, in cases where the person who tokens the mental representation is checking whether his flatmate has properly washed the kettle after using it.

---

method of projection? Even if one abandons the idea that understanding a word has to do with having a mental image, the meaning of a general word seems to involve some kind of descriptive content. The method of projection problem arises for this descriptive content as well.

[23] I owe this objection to an anonymous referee.

[24] The externalist approach would need to be one according to which the world determines the content of the concept-type—a concept type that is created during the conversation. If the world were supposed to determine the content of the concept-token, the approach would be compatible with Type-Underdeterminacy.

[25] It is not at all clear that this kind of externalism will even be applicable to all cases. Although it might sound appealing for cases of singular representation, it is not easy to see how the world could resolve the equivocation involved in (3), given that a mental representation corresponding to this sentence will not be about any particular individual.

More importantly, in the underdeterminacy scenarios the relevant understanding of the predicate doesn't seem to be determined by the world. Rather, it depends on what is at stake—on what is relevant for the purposes of the conversation or the thought episode. Imagine someone who wants to buy a modern, elegant kettle. Looking at a kettle, he tokens a mental representation that could be expressed by saying 'The kettle is black'. Let us imagine now that the relevant kettle is only black on its interior—a part of the kettle the subject cannot perceive—and that he was misperceiving the object because, let's say, the kettle is placed behind a dark glass. Should we say that his mental representation is such that the concept he tokens is one that is true of black-in-the-interior objects? I think this would be an odd description of the situation. Rather, given his purposes, it would be more reasonable to take his thought to be true if and only if the relevant kettle is superficially black. Typically, it is mental score, rather than the world itself, that resolves equivocation. But if it is conceptual material that is meant to fix the ensuing underdeterminacy, then the problems described in this section arise.

5. Partiality

In this section I will introduce a second notion of underdeterminacy and argue that mental representations are also underdetermined in this second sense[26].

So far underdeterminacy has been understood as concerning types. However, it can also be understood as concerning tokens of structured representations:

> **Token-Underdeterminacy:** A token of a structured representational item S is token-underdetermined if and only if for some possible states of affairs its truth-value is indeterminate (i.e., if and only if it determines a partial function from possible worlds to truth-values).

Token-Underdeterminacy can be motivated by appealing to what Waismann called the 'open texture' of our words. Here is a case Wittgenstein discusses ([1953] 2009, 80). Imagine that I invite a friend over and say "There is a chair over there". When he approaches it, it disappears, but seconds later it appears again. Is my utterance true in the scenario described? The conversation might have fixed a relevant understanding of the word 'chair' (for example, one that excludes doll chairs and baby chairs) but still fail to determine whether the thing that appears and disappears counts as a 'chair'.

[26] In the previous sections I have addressed the problem of underdeterminacy as it arises in the literature concerning language and thought (mainly, in Carston's work). This section addresses a separate problem concerning a different notion of underdeterminacy.

Token-Underdeterminacy becomes interesting when we realize that something similar to the case described by Wittgenstein can happen in more mundane scenarios. Imagine that instead of a chair that disappears I have some boxes that I usually use as chairs, or some pieces of modern furniture. There can be conversations about chairs that would leave it open whether these objects count as 'chairs'—conversations taking place at buildings where there are no such objects. Nonetheless, in other conversations it will be determinate whether these objects count as 'chairs' (saliently, in conversations at my place). Given a state of affairs (an apartment with some boxes and pieces of modern furniture), it can be the case that a token T_1 of 'There's a chair over there' is true at this state of affairs, whereas a second token T_2 is false and the truth-value of a third token T_3 is indeterminate (where T_1, T_2 and T_3 occur in different conversations but the state of the relevant objects and the value of indexicals remain constant). Different tokens of the same sentence can leave the truth-value of different state of affairs indeterminate.

As I mentioned in the previous section, natural language sentences do not make everything explicit. In this sense, they can be thought of as leaving some questions open[27], such as: where does a kettle need to be black in order to count as 'black'? Or, for how long does one have to have a positive emotion in order to count as 'happy'? Usually, the context of use resolves these questions. For example, if everything goes well, previous conversation will fix whether 'black' is supposed to apply to the whole surface or a specific part of it, or whether 'happy' refers to a rather stable emotion or not. Depending on what is at stake, the question will be answered in one way or another. However, there can be questions that are left unanswered in a conversation. For example, all that is said in a conversation might leave unanswered questions about abnormal illumination conditions or about for what period of time the superficial colour is supposed to last (does a kettle that has been painted black with a painting that will only last for a month—because it progressively disappears as it is washed, let's say—count as 'black' in a conversation about superficially black kettles?). Whenever some questions remain unanswered there might be possible states of affairs for which the utterance does not determine a truth-value. As I will argue, the same goes for mental representations.

Here is an example. Imagine someone asking for a red pen to mark some essays. He utters the sentence 'I need a red pen. I've some essays to mark'. Now, 'red pen' can be used to refer to pens with red ink or to superficially red pens, etc. To what does it refer in this case? The second sentence solves this doubt. Because of it, it might seem plausible to assume that the mental representation that the speaker was tokening contained information about the ink. However, there are other things that the sentence 'I need a red pen' does not make explicit. For example, it does not settle whether the ink must look red on recycled paper, or on yellow paper. Imagine that there are no recycled paper sheets available in the building where the conversation takes place and that, because of this, the speaker never prints the essays that his

[27] See Travis (1989) for this kind of approach to underdeterminacy.

students send him on recycled paper. His mental representation (token) will very likely simply leave that doubt open. As a result, it will not determine a truth-value for any possible state of affairs. In what follows I develop this point in some more detail.

Here is a well-known example by Travis, often used as an argument for the Type-Underdeterminacy of natural language (it is similar to the black kettle example Carston mentions):

> Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood. (Travis 1997, p. 111)

The example discloses some features that might matter to the application conditions of the colour predicate, such as the origin of the colour and the part of the object that exhibits it. It might seem plausible to say, about this example, that the mental representation that Pia tokens is (at least in the first occasion) one that specifies that the leaves are superficially, or visibly, green. Now, there are other examples disclosing other possible doubts concerning the application of the colour term. Let us call this kind of example a Travis case. There are other possible Travis cases. Here is one (based on some remarks by Travis 1997): Jean's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, he paints them. He has no green paint, and he enjoys experimenting, so he uses blue and yellow paint. He paints small yellow and blue dots. The paints don't get mixed. However, at a certain distance (0.5m, let's say), the leaves look green. Returning, he reports, 'That's better. The leaves are green now.' He speaks truth. A photographer friend then phones, seeking green leaves for photography of a single leaf, taken at a close distance. 'The leaves (on my tree) are green,' Jean says. 'You can have those.' But now Jean speaks falsehood.

This example discloses something else that the sentence 'The leaves are green' doesn't make explicit: the distance at which the object is observed might matter. If Pia's mental representation is free of Token-Underdeterminacy, it must determine whether Jean's leaves would satisfy her use of 'are green'. However, this might very well be indeterminate. Pia need not have considered the possibility that at a certain distance objects might look different (as it typically happens with printed colours). In Jean's scenario, this, however, is relevant, and this doubt can be considered as solved by the context. But this does not imply that the doubt is solved in every context in which the predicate 'is green' is used.

Let us go back to 'bachelor'. The meaning of this word gives rise to the following questions: is someone married by some religious ritual but not legally a bachelor? Is someone

in the middle of a divorce a bachelor? Is someone not married but who has a legal partner a bachelor? Is a man who married two women simultaneously in a ceremony taking place in his home country but that now lives in some European country (where polygamy is not permitted) a bachelor? These questions will be answered differently in different cases, and in some cases, some of them will not be answered at all. Imagine the following scenarios.

Case 1: Anne is interested on the relation between marital status and the amount of taxes paid. She suspects that being married has not entailed any economic advantage for her. Thinking that it would be interesting to ask for personal experiences, she tells a friend: 'I want to meet some bachelors'. In this setting, it is plausible to take Anne's mental representation to token a narrower concept than BACHELOR, roughly equivalent to non-legally-married-men. However, it is not plausible to take her mental representation to solve, for example, the question about men going through a process of divorce or about men who got married in countries with marital laws radically different from the ones in the country where she lives, even though some of these men happen to live in the same country as she does. These questions simply didn't occur to her.

Case 2: Tom wants to meet someone special and begin a long-term relationship, but most of the men he knows are married or have a partner. He says to a friend: 'I want to meet some bachelors'. Again, it is plausible to take his mental representation to solve some questions, but not all of them. For example, his mental representation will very likely exclude men who are married by some ritual regardless of whether they are legally married, and it will also exclude men married in other countries even though they don't count as legally married in the country where he lives. The point is that the concept he tokens need not decide whether a man in the middle of a divorce satisfies his use of 'bachelor', or what happens with men who got married in order to get a European nationality, or with men that got divorced but are still involved in a relationship with their ex-husbands.

As a conclusion, mental representations can be Token-Underdetermined. It is plausible to think that tokens of mental representations determine only a partial function from possible states of affairs to truth-values[28].

5. Conclusions

I have argued against the claim that mental representations are underdeterminacy-free. I have distinguished two notions of underdeterminacy—Type-Underdeterminacy and Token-Underdeterminacy. In sections 3 and 4 I have considered and rejected two arguments in support of the claim that mental representations must be free of Type-Underdeterminacy: the

[28] Belleri (2014) defends a related view according to which the notion of determinacy for the contents of our thoughts is context-relative.

productivity argument and the argument from equivocation. Against them I have argued, first, that the productivity argument only establishes that natural language and mental representations must be meaning-compositional, and this is compatible with Type-Underdeterminacy. Moreover, I have noted that views that have recourse to online concept creation cannot avail themselves of the productivity argument. Second, I have argued that the argument from equivocation relies on a false premise—namely, the premise that mental representation-types are non-equivocal and, as such, they are suitable to resolve the equivocations involved in an utterance of a natural language sentence. The tokening context can also resolve the equivocation. As a result, proponents of the mixed view have failed to provide good reasons in support of their view. After this, I have argued that *ad hoc* concepts should not be expected to behave differently than lexically encoded concepts. Hence, mental representations, even when they involve *ad hoc* non-lexically encoded concepts should be expected to be Type-Underdetermined.

In section 5 I have argued that there are reasons to take mental representations to be token-underdetermined. I have done so by presenting some cases where it seems reasonable to conclude that the truth-value of the mental representation-token is indeterminate at some possible states of affairs. The resulting picture—that mental representations are token-underdetermined—is admittedly a radical one. In particular, it is more radical than approaches that take mental representations to be type-underdetermined. The thesis that mental representations are type-underdetermined motivates approaches according to which the situation where a representation is tokened plays a crucial role in resolving equivocations and, consequently, in determining the truth-conditional content of the token[29]. This is compatible with the tokens being free of Token-Underdeterminacy. By contrast, the partiality argument motivates an approach according to which tokens exhibit this second kind of underdeterminacy. Although at the tokening context the questions concerning application conditions that are relevant for whatever is at stake will typically be answered, other questions that are not relevant can remain unanswered.

References:

Belleri, D. (2014). You can say what you think: vindicating the effability of our thoughts. *Synthese* 191 (18):4431-4450.

Cappelen, H., & Lepore, E. (2005). *Insensitive Semantics: a Defence of Semantic Minimalism and Speech Pluralism.* Wiley-Blackwell.

Carston, R. (2002). *Thoughts and Utterances.* Blackwell.

[29] As an example of an account along these lines, see Corazza and Dokic's situated minimalism (2012).

Carston, R. (2008). Linguistic communication and the semantics/pragmatics distinction. *Synthese* 165 (3):321-345.

Carston, R. (2015). Contextual adjustment of meaning. In Riemer, N. (Ed.). *The Routledge Handbook of Semantics* (pp. 195-210). Routledge.

Clapp, L. (2012). Is even thought compositional? *Philosophical Studies* 157 (2):299-322.

Corazza, E., & Dokic, J. (2012). Situated minimalism versus free enrichment. *Synthese*, *184*(2), 179-198.

Fodor, J. (2001). Language, thought and compositionality. *Mind and Language* 16 (1):1-15.

Fodor, J. (2003). *Hume Variations*. Oxford University Press.

Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic* 8 (1):339--359.

Martínez-Manrique, F., & Vicente, A. (2004). Overhearing a sentence: Recanati and the cognitive view of language. *Pragmatics & Cognition*, 12(2), 219-251.

Perry, J. (1986). Thought without representation. *Proceedings of the Aristotelian Society* 60 (1):137-151.

Recanati, F. (2004). *Literal meaning*. Cambridge University Press.

Recanati, F. (2007). *Perspectival Thought: A Plea for Moderate Relativism*. Oxford University Press.

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

Searle, J. (1978). Literal meaning. *Erkenntnis*, *13*(1), 207-224.

Searle, J. (1980). The background of meaning. In *Speech act theory and pragmatics* (pp. 221-232). Springer Netherlands.

Searle, J. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge university press.

Travis, C. (1989). *The uses of sense: Wittgenstein's philosophy of language*. Oxford University Press.

Travis, C. (1997). Pragmatics. In Hale, B. & Wright, C. (Eds.), *A Companion to the Philosophy of Language* (pp. 87-107). Blackwell.

Travis, C. (2000). *Unshadowed Thought: Representation in Thought and Language*. Harvard University Press.

Wittgenstein, L. (2009). *Philosophical Investigations, 4th edition (trans. Hacker and Schulte)*. Wiley-Blackwell.