

Examining the nature of referential metonymy

Josephine Bowerman

Abstract

Working within the framework of Relevance Theory, I argue that referential metonymy can be analysed in properly inferential terms, and can be treated as a distinct type of lexical innovation, alongside lexical ‘modulation’ (narrowing and broadening) and motivated neologism. I address such issues as motivations for referring metonymically and factors constraining the choice of metonymic referring term. In addition, I consider how the interpretation of referential metonymy may routinize, and how frequency of use and routinization may affect the storage and processing of a metonymic interpretation, examining empirical data on metonymy processing from Frisson and Pickering (1999).

Keywords: Referential metonymy; Relevance Theory, lexical innovation; underspecification

1 Introduction

Metonymy is traditionally treated as merely a rhetorical device or figure of speech in which, for the name of a thing, we substitute the name of an attribute of that thing, or of something closely related to it (e.g. *hands* for *workers*, or *skirt* for *woman*). In order to further our knowledge of the phenomenon, I set out to elucidate the nature of referential metonymy, exemplified in (1)-(3), wherein the referring terms ‘*ham sandwich*’, ‘*Shakespeare*’ and ‘*muscle*’ are being used metonymically because they express attributes of/things relating to the actual target referent:

- (1) The *ham sandwich* has left without paying.
- (2) Put *Shakespeare* on the top shelf.
- (3) If they won’t pay, send in the *muscle*.

In (1) the referent is a restaurant customer but the referring term is the word for his food order; in (2) the referent is a book but the referring term is the name of the author of the text; and in (3) the referents are hired tough guys but the referring term is the word for their most notable property, i.e. their sizeable muscles.

I restrict my focus to referential metonymy because it is the prototypical example of metonymy and as such has been studied in most depth. Working within the framework of Relevance Theory (hereafter, RT) I aim to discover whether referential metonymy can be treated as a type of lexical innovation in its own right, and whether it is amenable to a fully inferential analysis.

I will show that referential metonymy is a motivated phenomenon in which the linguistic properties of the metonymically-used word provide evidence of the speaker’s intended referent, such that the target interpretation can be recovered inferentially. The comprehension of referential metonymy is therefore compatible with the RT view of utterance processing, which sees interpretation as fully inferential (Sperber & Wilson 1986/1995, 1998, 2002; Wilson & Sperber 2002). I will also argue, by comparison with other lexical pragmatic phenomena, namely ‘modulation’ (narrowing and broadening) and motivated neologism, that referential metonymy is indeed a distinct type of lexical innovation which involves ‘repurposing’ an existing expression to pick out a target entity. Finally, I will consider empirical evidence relating to my arguments.

2 Motivations for referring metonymically

Before I examine the interpretation of referential metonymy, and assess whether or not the phenomenon may be classed as an independent type of innovative word-use, I will begin by investigating why and how speakers refer metonymically. I take as my starting point RT accounts such as Papafragou (1996), Jiang (2013) and Rebollar (2015), in which referential metonymy is seen as a speaker's intentional use of an expression to refer to an entity that does not fall under the literal denotation of the expression. These accounts draw on Kaplan's (1989, p. 559-561) claim that a speaker's intention to use an expression to refer to a given entity irrespective of any prior meanings associated with, or pre-established uses of, the expression (the 'referential' intention) can dominate over the intention to use an expression with the meaning given it by the person from whom the expression was learnt (the 'attributive' intention). For example, although the established usage of the expression '*ham sandwich*' is to pick out the bread-and-meat snack, in (1) above, uttered in a restaurant context, the speaker's critical intention in using this expression is to refer to a particular customer who ordered a ham sandwich. Also pertinent is Donnellan's (1966) slogan: expressions denote, people refer. Consider (4):

(4) I spent the summer reading *Dickens*.

While the encoded meaning of the expression '*Dickens*' denotes the flesh-and-blood author Charles Dickens, the person who utters (4) uses this expression to refer to the works of Charles Dickens. This highlights a distinction between the linguistically-specified referent of an expression and the entity/entities which a speaker can use the expression to pick out. It also raises the question of why this distinction should arise: what reasons might a speaker have to fix a novel referent for an existing expression?

Speakers may refer metonymically in situations where referring literally would impose unnecessary extra processing costs on the audience. For example, suppose that a waitress in a crowded restaurant full of businessmen on their lunch break intends to refer to a specific diner. However, using an expression which literally refers to her target entity (e.g. '*businessman*') would, in the communicative context, lead to referential ambiguity because it would not identify a unique individual, and would thus require of the audience extra cognitive effort to resolve the ambiguity and work out the intended referent. Furthermore, in a hectic restaurant environment the audience, i.e. other waitresses, will have to deal with many other cognitive tasks in addition to utterance interpretation, and therefore may not have the time or processing resources to interpret a literal referring expression such as '*the man who ordered the ham sandwich*' which, although it provides disambiguating information, is long and syntactically complex.

Given the RT assumptions that all acts of ostensive communication, including utterances, communicate a presumption of their own optimal relevance (i.e. that the ostensive stimulus is both sufficiently relevant to be worth the audience's processing effort, and also the most relevant stimulus the communicator is willing and able to produce), and that the overall relevance of an utterance is a function of cognitive effects against processing effort (Sperber & Wilson, 1986/1995), the extra effort required for ambiguity resolution or processing a complex expression may result in the utterance failing to meet expectations of relevance, such that the audience abandon utterance processing without having recovered the speaker's intended referent. It is therefore in the best interests of the speaker to choose a referring expression that makes reference resolution easy— quick and unambiguous. Minimising the effort required for interpretation increases the likelihood that the audience will maintain

attention and fully process the speaker's utterance, thus increasing the likelihood that the intended referent will be recovered (Allott, 2013, p. 81).

Further, as Papafragou (1996, p. 186) notes, a speaker might refer metonymically to provide access to more contextual implications than would be available if reference were made literally, or to express attitudinal/affective information. By way of illustration, consider an alternative scenario. A guest at an office Christmas party is talking to her friend and, wishing to refer to their colleague Dave, who is wearing a distinctive outfit that night, utters (5):

(5) Look at the *green trousers* dance!

While it would be quicker and easier to refer to Dave by simply using his name, which is known to both speaker and addressee, picking him out by using the expression '*green trousers*' creates a humorous effect through reducing a human being to his garments. It also encourages the addressee to explore more carefully the implications of the referring expression '*green trousers*' and thereby derive extra conclusions, e.g. that Dave is, like his outfit, loud and tasteless. Here, the extra effects derived presumably outweigh any extra processing costs that their derivation incurs. Just as minimising the effort required for utterance processing is in the speaker's best interests, so too is maximising the benefits of the enterprise for the audience: providing extra, useful information will also increase the chances that the audience will pay attention, fully process the utterance, and remember the message that the speaker wants to convey (Allott, 2013, p. 81).

Finally, Jiang (2013, p. 11-13) and Rebollar (2015, p. 196-7) claim that speakers use metonymy to fill gaps in the lexicon if there is no existing word to refer to the target entity. For example, if the two friends at the office party do not know the name of the enthusiastically jitterbugging man in emerald flares, they must use an alternative referring expression which allows for economical-efficient referent identification, and possibly also implies additional conclusions. Summarising these points, we could thus say that the overarching motivation for referential metonymy is the speaker's realisation that, with respect to the entity she has in mind and intends that her audience should pick out, the available means for referring to that entity are not fully adequate.

3 Selecting a referring expression

When using the available linguistic resources as they are ordinarily used would prevent the speaker from fulfilling her communicative intention, as regards getting her audience to come to pick out a specific entity, the speaker's challenge is then to find a better way of referring. The most economical-efficient way to do so is to 'repurpose' an existing expression for the task. Yet this leads us to question how the existing term which is to be 'repurposed' in referential metonymy is chosen. For example, why does the restaurant waitress who wishes to refer to a specific customer use the expression '*ham sandwich*' to pick out her target entity?

If metonymy is indeed to optimise reference-making in the absence of an adequate literal way to refer, the speaker must ensure that although the referring term she uses does not literally identify her target entity, it is nevertheless the best 'evidence' she is able to provide for her referential intention (and, in some cases, for her intention to achieve certain effects), given the RT view that linguistic meaning, rather than fully determining communicated content, acts as a clue or pointer to the speaker's intended meaning. The 'evidence' must allow the audience to work out the intended referent in a way that is as low in processing costs and/or as rich in beneficial effects as possible, such that relevance is maximised. This can be achieved by choosing an existing expression which refers to a cognitively salient feature/property of the

target entity, i.e. a feature/property within the mental representation of the target entity that is more highly activated or accessible than other properties. Entities can be identified through their salient properties (Rebollar, 2015, p. 195), thus by getting the audience to focus on a particular, especially attention-commanding aspect of the target entity (e.g. for the restaurant customer to whom the waitress wishes to refer, the property of ‘being a ham-sandwich orderer’), it is made easier for the audience to pick out the entity itself.

When we speak of a ‘salient’ feature/property of an entity, we must bear in mind that cognitive salience is a relative notion that depends on the demands of the situation and thus cannot be independently defined, as Jiang (2013) emphasises. While there may be properties of an entity that are ‘inherently’ attention-commanding, e.g. concrete, functional or interactional properties (see Langacker (1987, p. 385-6) and Radden & Kovecses (1999) on the general cognitive determinants of salience), ultimately, the most salient property of an entity is the property that in the context at hand is the most relevant (yielding the most cognitive effects relative to processing effort) and therefore the most worthwhile for an agent to attend to. Importantly, the inherent salience of a property does not guarantee its usefulness in every situation. For example, if the situation demands economical-efficient identification of a target entity, the most relevant property of the entity will be one which, regardless of whether or not it is inherently salient, is individuating, i.e. in the context, uniquely picks out the entity without introducing any referential ambiguity which would require extra processing effort to resolve. Alternatively, if the situation requires us to derive certain implicit conclusions concerning a target entity, it may be the case that the assumptions associated with an inherently salient property of the entity in question would not allow us to draw the intended conclusions, and thus the inherently salient property would not be the most relevant property of the entity.

A speaker whose communicative goal is to achieve economical-efficient reference resolution and/or to convey contextual implications about a particular entity must therefore, in making reference to her target entity, get her audience to focus on the property of the entity that is most relevant with respect to her intentions. To do this, she must ‘repurpose’ as her referring expression the existing expression that, in the communicative context, introduces a property of the target entity which is uniquely identifying, thereby making the entity maximally accessible and minimising the effort required to pick it out, and/or which achieves additional effects.

We are now able to account for why the waitress chooses the referring expression ‘*ham sandwich*’ to identify her target customer. In the restaurant context, ‘food ordered’ is already a highly relevant property of customers, as attending to this property enables waitresses to plan and act successfully (taking the correct order to the correct customer). Moreover, for any given customer, this property is likely to be one of the only properties to which a waitress attends. This makes ‘food ordered’ an effective individuating property of customers. Therefore, by referring to a specific target customer using a term e.g. ‘*ham sandwich*’ that expresses the customer’s food order, the waitress is likely to have provided sufficient evidence of her referential intention to enable an addressee to identify the intended customer without incurring unnecessary processing costs. In addition, given sufficient time and processing resources, addressees can consider the assumptions associated with the referring term ‘*ham sandwich*’ and derive extra conclusions, e.g. that the intended referent—the customer who ordered the ham sandwich—is, like his choice of snack, traditional and boring. Although on this occasion drawing such conclusions is not central to the success of the metonymic use of ‘*ham sandwich*’ because the speaker’s primary goal is to ensure economical-efficient identification of her intended referent, any further implications that addressees can derive may increase the overall relevance of the utterance.

4 Interpreting referential metonymy

Having considered motivations for and factors constraining the production of referential metonymy, we can now explore how addressees interpret metonymically-used referring expressions. First, I will briefly present the RT approach to utterance interpretation. RT claims that in comprehension, our goal is to construct a hypothesis about the speaker's intended meaning that satisfies the expectations of relevance raised by the utterance. This task comprises three sub-tasks: (i) constructing hypotheses about the explicit content of the utterance, (ii) constructing hypotheses about intended contextual assumptions, and (iii) constructing hypotheses about intended contextual implications (implicated conclusions). These hypotheses are developed on-line and in parallel (Sperber & Wilson, 2002, p. 261). Considerations of relevance may lead us to expect a particular conclusion, or type of conclusion; these expectations may contribute via backwards inference to the identification of explicit content and contextual assumptions which act as explicit and implicit premises that warrant, via forwards inference, the anticipated conclusion (Carston, 2004). Utterance interpretation is therefore fully inferential: each sub-task in the overall process of constructing a plausible interpretation for the speaker's utterance involves a non-demonstrative inference process (Sperber & Wilson, 2002, p. 262). The mutual parallel adjustment of explicit content, context and cognitive effects is constrained by the RT comprehension strategy, which states that interpreters should follow a path of least effort and test interpretive hypotheses in order of accessibility, stopping when expectations of relevance are satisfied (or abandoned) (Sperber & Wilson, 2002, p. 259).

Let us apply the RT approach to the following examples of referential metonymy to see that, although these cases differ in terms of 'creativity', and in terms of how addressees make use of the metonymic referring expression in reference resolution, they are nevertheless both interpreted using the same mechanism of mutual parallel adjustment.

We will begin with (4), repeated here as (6), which is a less creative instance of referential metonymy:

(6) I spent the summer reading *Dickens*.

The speaker of (6) does not intend '*Dickens*' to be understood literally, as referring to Charles Dickens. This would yield an ungrammatical, nonsensical interpretation of (6) that would fail to satisfy our expectations of relevance. Rather, her intended referent is the works of Dickens. However, she has chosen '*Dickens*' as her referring expression because the property of her target entity that this word expresses (the name of the author of the works) is individuating, thereby enabling us to pick out that entity with minimum effort; and it allows us to derive her intended conclusions.

In order to grasp the intended referent of the metonymic use of '*Dickens*' in (6), we 'zoom in' on especially accessible encyclopaedic assumptions of the encoded meaning of the word, e.g. the information that Dickens was an author. Further, given that online processing proceeds incrementally (e.g. Just & Carpenter, 1980; Tyler & Marslen-Wilson, 1977), incoming words may influence predictions concerning the intended interpretation of the utterance. In (6), when the verb '*read*' is encountered, its selectional properties will make certain interpretive hypotheses about the explicit content of the utterance highly accessible, for instance that if the verb is followed by a noun-phrase complement, that noun phrase will refer to something readable. These assumptions and hypotheses combine in mutual parallel adjustment. Inference runs forwards from the hypothesis that the noun phrase will refer to a readable object, and from the assumption that Dickens was an author: authors write readable works, therefore the most plausible referent for '*Dickens*' is the works of Dickens. There may also be backwards

inference from expected cognitive effects (for example the anticipatory hypothesis that the speaker's utterance will achieve relevance by describing an activity that took up her whole summer), which contributes to the identification of explicit content and contextual assumptions that together will warrant the derivation of the contextual implications of the utterance. We therefore derive an interpretation for (6) where the term '*Dickens*' is understood as picking out the works of Dickens. This allows for sound inference to our expected conclusions: given that the works of Dickens are numerous and lengthy, we can infer that the speaker of (6) was occupied with her reading for the entirety of the summer. As this derivation shows, the comprehension of (6) is a fully inferential process, with both explicit and implicit content worked out by abductive reasoning.

Now consider a more novel example, (7), uttered in a crowded bar to refer to a hirsute man sitting nearby:

(7) The *beard* looks miserable.

Once again, the linguistic context will influence the accessibility of interpretive hypotheses about the explicit content of (7). For example, the verb phrase '*looks miserable*' may lead us to predict that its subject will be human. However in this instance, to identify the speaker's intended referent we do not need to focus on contextually relevant encyclopaedic assumptions about the referring expression '*beard*'. This is because a literal beard is physically present and directly perceivable, and the speaker's use of the word '*beard*' draws our attention to this entity. Crucially, when we attend to the beard, we simultaneously attend to an individual who could plausibly be the speaker's intended referent, by virtue of the fact that the beard is on his face. Thereby, through her choice of referring expression, the speaker leads us to focus on her target entity. Moreover, the speaker has presumably chosen the word '*beard*' as her referring expression because no other men present in the bar but her intended referent have facial hair, and the word '*beard*' therefore expresses an individuating property of her target individual. Thus, when we attend to the literal beard we can see nearby, we simultaneously attend to the only person who could be the speaker's intended referent.

Given that interpretation follows a path of least effort, the first interpretive hypotheses to be tested and developed by the mutual parallel adjustment process will be the highly accessible hypotheses that the subject of (7) is human and that the speaker's intended referent is most likely to be the only bearded man visible in the immediate environment. Inference runs forward from these hypotheses to the conclusion that the intended referent of '*beard*' is indeed the bearded man nearby. In addition, backwards inference from specific expectations about the cognitive effects that the utterance of (7) will achieve, for example the expectation that the speaker will provide us with useful information about other people in the bar, may also help us to recover an interpretation in which the referent of '*beard*' is taken to be a fellow drinker. Having derived the explicit content of the utterance, we can go on to infer further contextual implications which will contribute to the overall relevance of the utterance, e.g. that if the '*beard*', i.e. the bearded man nearby, is looking miserable, we should go over and attempt to cheer him up. Here, too, we can see that each of the sub-tasks contributing to the recovery of a plausible reading of (7) that satisfies our expectations of relevance involves non-demonstrative inference, such that the overall comprehension process is fully inferential.

5 Comparison with lexical modulation

As my analyses of the above examples of referential metonymy make clear, it is possible to provide an adequate explanation of the phenomenon in RT terms without claiming that in the

course of interpretation, an occasion-specific sense, or ad-hoc concept, is constructed from the linguistically-specified meaning of the metonymically-used referring term as the result of interaction between encoded information, contextual information and relevance-based expectations (see e.g. Wilson and Carston (2007) on ad-hoc concepts). That is to say, in order to deal with the phenomenon on a properly inferential approach, there is no need to treat it as a variety of lexical ‘modulation’ in which the linguistically-specified meaning of the metonymically-used referring term is narrowed, such that the communicated sense is more specific than the encoded sense; broadened, such that the communicated sense is more general than the encoded sense; or a combination of the two (Carston, 1997). Indeed, a modulation account appears to be wrong for referential metonymy.

Consider first narrowing, which involves using a word to pick out a proper subset of the items that fall under its linguistically-specified denotation, as in (8) and (9):

- (8) I’ve got a *temperature*.
 (9) Josie is avoiding *bachelors*.

(8) communicates the message that the speaker has a specific kind of temperature, i.e. a higher-than-normal one. (9), given a particular set of background assumptions, could be understood as conveying that Josie is steering clear of a certain subtype of unmarried men, for example the kind that still live with their mothers. The words ‘*temperature*’ and ‘*bachelors*’ come to express the ad-hoc concepts TEMPERATURE* and BACHELORS*, which have a restricted denotation compared to that of the encoded meaning from which they are derived. In order to construct such ad-hoc concepts, contextually relevant encyclopaedic properties are added to the encoded meaning, e.g. for ‘*bachelors*’, the property of living with their parents.

In referential metonymy, however, the overall effect is not to highlight a proper subpart of the linguistically-specified denotation of the metonymically-used referring term. It is plainly not the case that the intended referent in the ‘*ham sandwich*’ example—the customer who ordered the ham sandwich—could be thought of as a subtype of ham sandwiches. Likewise, in (10), there appears to be no way in which the entity that the speaker intends to refer to by her use of the word ‘*Mozart*’, i.e. a piece of music by that composer, is a specific subtype of the entity which is picked out by the encoded meaning of the word, i.e. the man Wolfgang Amadeus Mozart:

- (10) I’m playing *Mozart*.

Rather, the metonymic use serves to pick out an entity which lies outside of the linguistically-specified denotation of the referring term.

A further important difference between referential metonymy and narrowing is that in referential metonymy, the referring term itself continues to mean what it means on its normal usage. It does not come to express, through the addition of encyclopaedic properties, an ad-hoc concept which is a more specific version of the sense it is ordinarily used to convey. All that changes is the entity that it is used to pick out. It can thus be concluded that referential metonymy cannot be explained as a variety of narrowing.

We turn next to broadening. Broadening results in a widening of the linguistically-specified denotation of the broadened term, such that the term can apply to entities that lie outside of its linguistically-specified denotation. The denotation of the ad-hoc concept which results from broadening contains all the entities contained in the denotation of the original literal concept, plus more. In (11) for example, ‘*boiling*’ expresses the ad-hoc concept BOILING*, which includes in its denotation not only things that are boiling in the strict sense (having a

temperature of 100 C) but also further cases of things that are hot but not boiling, such as very warm bathwater.

(11) The water is *boiling*.

Yet compare the metonymic use of the expression ‘*ham sandwich*’, which denotes, or refers to, a person, not a person in addition to all the bread-and-ham snacks. The referring expression comes to pick out an entity that does not fall within its linguistically-specified denotation, which may appear similar to the way in which a broadened expression can apply to entities that lie outside of its linguistically-specified denotation. Importantly, however, the metonymic denotation does not stand in a superset relation with the denotation of the encoded literal concept, unlike the denotation of a broadened concept.

Broadening crucially involves dropping logical/definitional properties from the encoded meaning of the broadened expression, e.g. for ‘*boiling*’, the property of being at 100 C. However, this again seems very different to what occurs in the comprehension of referential metonymy. Consider some different examples: (12), a familiar case, where the word ‘*Vietnam*’ is used to refer to the Vietnam War; and (13), a more novel case, where the word ‘*sheep*’ is used to refer to someone born in the Year of the Sheep (originally from a television program; cited in Gerrig (1989) and Papafragou (1996)).

(12) *Vietnam* was a military disaster.

(13) You should avoid marrying a *sheep* at all costs.

In order to derive an interpretation of (12) and (13) that satisfies our expectations of relevance, we cannot take the literal referent of ‘*Vietnam*’ and ‘*sheep*’ to be the speaker’s target entity, and must find a novel referent for these expressions. However, in (12), we do not come to pick out an entity of which the property of being a military disaster can appropriately be predicated and which is therefore most plausibly the speaker’s intended entity by dropping logical properties of the encoded meaning of ‘*Vietnam*’ (i.e. the property of being an South-East Asian country), and possibly also adding encyclopaedic properties in the process, to end up with an ad-hoc concept, VIETNAM*, that includes the target entity in its denotation. Likewise, although (13) is a more creative instance of referential metonymy, in order to identify the speaker’s intended referent we do not need to drop from the encoded meaning of ‘*sheep*’ the logical property of literal sheephood, such that ‘*sheep*’ comes to convey an ad-hoc concept, SHEEP*, which can apply to human beings.

Indeed, we do not modulate the linguistically-specified meaning of the metonymically-used referring expression in any way. The referring expression conveys the same sense when used metonymically as it does when used literally. Although the referring expression comes to pick out a novel referent, this is not achieved by dropping logical properties from its encoded meaning to derive an ad-hoc concept whose denotation is sufficiently expanded compared to that of the input meaning as to include the target entity. Rather, the communicative context makes certain encyclopaedic assumptions of the encoded meaning of the referring expression especially relevant and accessible: in (12), that a disastrous war took place in Vietnam, and in (13), that the sheep is a Chinese zodiac symbol. These assumptions provide a uniquely identifying property of the speaker’s intended referent—the location of the target war in (12), and the Chinese zodiac sign of the target individual in (13)—thereby enabling the audience to economically-efficiently pick out that referent. Thus it is by virtue of the unmodulated linguistically-specified meaning of the metonymically-used referring expression that recovery of a novel referent is accomplished.

This shows that referential metonymy cannot be treated as a variety of broadening. It therefore appears that referential metonymy is not amenable to an analysis in terms of either type of lexical ‘modulation’, narrowing or broadening, as also noted by e.g. Recanati (2004, p. 26) and Wilson and Carston (2007, p. 254). Whereas in cases of modulation the linguistically-specified meaning of an expression is pragmatically adjusted to derive the speaker’s intended interpretation, in cases of referential metonymy, no such ad-hoc concept construction is involved in the recovery of the speaker’s target referent. Note, too, that not only is no new occasion-meaning expressed by the metonymically-used referring term, but my derivations in §4 also assume no change to the standing meaning of the metonymically-used referring expression.

6 Comparison with neologism

Given that it seems referential metonymy is not reducible to lexical ‘modulation’, we could argue that the phenomenon is a distinct type of innovative word-use, involving the fixing of a novel referent for the metonymically-used referring expression. However, it may be possible to treat referential metonymy as a variety of neologism, or word coinage, which would undermine the hypothesis that referential metonymy is an independent phenomenon of innovation. For example, Papafragou (1996, p. 182) claims that metonymically-used referring terms function as ‘newly coined names’ for the speaker’s intended referent, and Rebollar (2015, p. 196) speaks of ‘ad-hoc name creation’ for identification purposes. In addition, the need to fill gaps in the lexicon is identified as a motivation for metonymy by Jiang (2013) and Rebollar (2015), which implies that metonymic usages can be seen as supplying new words.

Wilson (in Wilson & Falkum, 2015) goes further, explicitly arguing that metonymic usages arise through spontaneous processes of word coinage and should be analysed as ‘motivated’ neologisms, i.e. neologisms in which the linguistic properties of the new word (here, the metonymically-used referring term) provide a clue to the speaker’s intended meaning. This argument seems to be driven by two main goals. Firstly, Wilson aims to provide an account of metonymy that is not formulated in terms of lexical modulation, but also does not resort to citing code-like ‘transfer of meaning’ rules (‘author for work’, ‘foodstuff for consumer’, etc.). While these mappings are central to the cognitive linguistics approach to metonymy (see e.g. Lakoff, 1987; Kövecses & Radden, 1998; Radden & Kövecses, 1999; Panther & Thornburg, 2003), wherein metonymy is seen as the natural language reflex of what is primarily a cognitive phenomenon involving a relationship of contiguity between two items in a single cognitive domain, such a code-based account seems only able to describe the input-output relation for certain highly frequent, familiar metonymic usages, and cannot account for how these relations arise nor how they are comprehended when they are first used, unless we adopt the unsatisfactory position that these rules are innate. Secondly, Wilson—while acknowledging the important role for associative relations, i.e. spreading activation patterns, in metonymy comprehension—seeks to accommodate these within a fully inferential approach. Wilson thus claims that metonymy interpretation involves the same mechanism as is used in the interpretation of other motivated neologisms such as denominal verbs or noun-noun compounds. That is to say, a hearer using the RT comprehension heuristic will treat the metonymically-used referring term as an ostensive stimulus and infer the intended meaning by mutual adjustment of explicit content, context and cognitive effects, following a path of least effort.

Wilson may also be influenced by evidence from acquisition which shows that even prelinguistic infants spontaneously produce motivated ostensive stimuli, including stimuli that seem to be based on the ‘metonymic’ principle of picking out a salient property/feature of the

target referent as a way to identify that referent, e.g. smacking lips for 'food' or making a 'vroom vroom' noise for 'car' (Falkum, Recasens & Clark, 2013). It is likely that infants 'innovate' in this way to compensate for a lack of established stimuli for conveying the intended message. The same applies to older infants' uses of novel denominal verbs, e.g. '*to gun*' for 'shoot', and noun-noun compounds, e.g. '*clown boy*' for 'boy who is a clown' (Falkum, Recasens & Clark, 2013): young children use language creatively to fill vocabulary gaps, when they want to pick out a specific referent but lack the words to refer to their target entity. This suggests that metonymic usages and phenomena which are traditionally taken as clear cases of neologism have a common motivation. Adults may also be driven to innovate for the same reason, but it is likely that their motives are more varied and less 'egocentric': adults may take into account the needs of their audience, and innovate in order to reduce processing costs and/or increase positive cognitive effects. Given that children's creative uses of language, including metonymic uses, fill vocabulary gaps, their innovations are clearly creating 'new words'. Adults typically have an established word in cases where a child might employ an innovative usage, therefore the novelty of the child's innovation is more apparent. This highlights the sense in which metonymic uses can be seen as new coinages.

Yet, although we can identify phenomena such as denominal verbs and noun-noun compounds as examples of neologism/new coinage, there nevertheless appears to be a lack of clarity as to how we should understand the notion of 'new coinage'. For instance, on the basis of the observation that children use language creatively when they are lacking a 'label' for, or way to pick out, some entity/concept they want to talk about, we could view a 'new coinage' as a 'new name/label'. This interpretation would also encompass adult's lexical innovations. For example, a metonymic usage motivated by the need to refer more efficiently or with more effects could be seen as the creation of a new name/label for the target referent, simply with the caveat that a new name/label is needed not because of a vocabulary gap, but because the existing way to pick out the target referent is inadequate in the communicative context. Accepting this understanding of 'new coinage' would mean that the notion applies to four distinct phenomena: metonymy, modulation, classical neologisms like denominal verbs, and out-and-out inventions e.g. *wilgrid* 'thorny, knotty problem'. These phenomena could be seen as falling on a continuum of creativity, with respect to how much is changed in order to coin a new name/label. Metonymy, at the lower end, only changes the referent of the metonymically-used word. At the upper end, invention is the most creative way to coin a new name/label: rather than repurposing an existing word, we produce an entirely novel form.

Invention aside, metonymy, modulation and classical neologisms all involve reusing an existing expression in some way; in Wilson's terms, these are the 'motivated' phenomena, in that the linguistic properties of the repurposed word provide evidence of the speaker's intended meaning (presumably, this is what influences the speaker's choice of word to be repurposed). These phenomena would therefore be amenable to a fully inferential analysis, and Wilson's goal of getting away from an unexplanatory analysis of metonymy which appeals to 'transfer of meaning' rules would thus be met. This seems to be a significant advantage of understanding 'new coinage' as 'new name/label'. However, classing both metonymy and modulation as new coinages may be an unwanted outcome for Wilson, who seems to want to keep metonymy firmly apart from modulation. Given that metonymy, modulation and classical neologisms would all receive an inferential analysis on this approach, the challenge would then be to find a way of clearly distinguishing the three phenomena. This would involve showing adequately that although the interpretive mechanism, i.e. the mutual parallel adjustment of explicit and implicit content which driven by the search for relevance and constrained by the RT comprehension procedure, is the same in all three cases, the outcomes of the operation of that mechanism are different (thereby meeting Wilson's goal of not treating metonymy as a case of

narrowing or broadening). It would also involve considering how the demands of communication and interpretation might lead to one phenomenon rather than another.

A further problem is that there is a strong sense in which the cases with no form-changes, i.e. metonymy and modulation, are not overtly novel, especially in contrast to e.g. a new compound or an outright invention, and even when compared to denominal verbs (e.g. ‘party_V’ = to attend lots of parties_N) and deverbal nouns (e.g. ‘scratch_N’ = slight mark/injury caused by scratching_V) which, although identical to the ‘parent’ noun/verb in bare form, nevertheless appear with verbal/nominal inflections (*he partied all night; three scratches*), such that the surface form of the new word is different to that of the existing word from which it is derived. Although in metonymy and modulation the reference (metonymy) or the meaning (modulation) of the innovatively-used word changes, on the surface it retains its old form and for this reason it could be argued that there is no true ‘coinage’ taking place.

Alternatively, we could view a ‘new coinage’ as a ‘new form/meaning pairing’. That is to say, a new coinage should involve, minimally, a change to the meaning of an existing expression. Note, though, that the term ‘meaning’ can be ambiguous between ‘encoded’ (or ‘standing’) meaning and ‘expressed’ (or ‘occasion’) meaning. We must therefore clarify whether a new coinage requires a change to the standing or the occasion meaning of an expression. Armstrong (2016, p. 101-2) draws a distinction between ‘lexical shifts’, in which the semantic properties of an existing word are revised, and ‘lexical expansions’, in which a new ‘character’ is introduced into the language, claiming that only with the latter type of lexical innovation is the set of linguistic expressions available for use between a speaker and her audience enriched by the addition of a new linguistic expression not already in circulation among the interlocutors. Presumably, ‘character’ is meant in the Kaplanian sense of ‘standing meaning’, a context-invariant constraint on the range of semantic contents that can be assigned to utterances of an expression. If so, this suggests that the coinage of a new word requires a new encoded meaning.

However, given that on my account a metonymically-used referring term comes to pick out a new referent but keeps its old standing meaning, it appears that referential metonymy does not involve a new form/meaning pairing and should therefore be excluded from the class of new coinages. This would be an undesirable result for Wilson. It also seems that ‘modulation’, in which a new ‘occasion’ meaning is derived for the modulated word on the basis of its constant standing meaning, cannot be seen as a variety of new coinage. This result is less problematic for Wilson, who holds cases of modulation—old words getting new ‘meanings’—apart from ‘true’ motivated neologisms.

Yet what of the ‘form’ part of the ‘form/meaning pairing’? We could focus on the factor of overt novelty, discussed above, and argue that for a lexical innovation to count as a ‘new coinage’, not only must we end up with a new standing meaning, we must also create a new surface form. On this approach, both classical neologisms (new meaning, new form created from an old word/old words by morphosyntactic operations such as category change or compounding) and inventions (new meaning, entirely novel form) would count as new coinages. However, a serious issue immediately becomes apparent. If we take true motivated neologisms to crucially involve the creation of a new linguistic expression, there again seems to be no way in which referential metonymy could be included in this class, given the claim that in referential metonymy, only the referent of the metonymically-used word changes. Wilson describes metonymies as ‘denominal nouns’, yet the creation of such expressions does not seem to correspond to any kind of morphosyntactic operation of the sort undergone by classical motivated neologisms: there is necessarily no category change, unlike with e.g. denominal verbs, and there is no compounding, affixing or other surface evidence of a form-change. It appears that if we accept the ‘new coinage as new form’ treatment, referential

metonymy is in fact excluded from the very class of phenomena to which Wilson argues it belongs.

While this outcome would clearly be problematic for Wilson, it does not rule out a pragmatic account of referential metonymy in which the phenomenon is not reduced to modulation. In §4 I have provided fully inferential analyses of some examples of referential metonymy which, if plausible, suggest there is no need to group metonymy with other phenomena of lexical innovation such as modulation or motivated neologism in order to deal with it on a properly inferential approach. Indeed, we have seen that metonymy is not modulation, and nor does it seem to be new coinage. Rather, metonymy is plausibly a distinct motivated phenomenon alongside modulation and neologism, which is interpreted by the same single inferential mechanism of mutual parallel adjustment. While it would be undesirable to posit multiple interpretive mechanisms (for example, to claim that certain linguistic phenomena are interpreted by an associative, code-based mechanism while others are interpreted inferentially), I can see no objection to proposing that there are multiple different kinds of lexical innovation, especially given that we already take modulation alone as subdividing into narrowing and broadening, and that RT sees broadening as covering a range of cases including hyperbole and metaphor (e.g. Wilson & Carston, 2007).

Due to the demands of communication (the message a speaker wants to convey, the ways in which she can use existing linguistic resources to satisfy her communicative intention) and of interpretation (the way an audience uses a speaker's linguistic 'evidence' to recover the intended meaning given the time available, their processing resources and their own informational needs/preferences), a metonymic usage of an expression does something specific: it fixes a new referent for the expression and provides us with a novel way of picking out the target entity (this is what makes it a case of 'innovation'). Yet no new sense comes to be conveyed by the expression, and its linguistically-specified meaning does not change. Further, the creation of a new form is not part of the innovation. Referential metonymy, as a phenomenon in its own right, can therefore be clearly distinguished from modulation (old word, new occasion meaning) and neologism (new word, new standing meaning). However, this distinction can be made without also needing to claim that metonymy requires a specific interpretive mechanism involving 'transfer of meaning' rules.

7 Metonymic patterns

Using an expression metonymically to pick out a novel referent may be a one-off event, or a metonymic usage may be limited to a specific context (e.g. the '*ham sandwich*' case, which is localised to its restaurant context) and/or a specific group of people: the context creates the conditions which motivate speakers to use words metonymically, as well as providing information and assumptions which are used to interpret the metonymy. Metonymic usages may also be more general. If a particular metonymy is highly relevant and useful, it will come to be used increasingly frequently as a way to refer to the target entity, and the interpretation process will routinize. This reduces processing costs, as it means that effortful, (more) conscious derivation of the intended interpretation is no longer required on every occasion the metonymy is encountered. In addition, a consequence of routinization is that interpreters may need to rely less on contextual support, and the metonymy can be used outside of its original context.

Interestingly, exploring how metonymy interpretation routinizes could shed new light on the 'conceptual metonymies' or associative mappings (e.g. 'producer for product', 'place for inhabitants', 'clothing for wearer', etc.) that are so central to the cognitive linguistics account of metonymy. As noted, metonymy interpretation is likely to routinize when the metonymy in

question is used very frequently—whether in a specific context or generally. A frequent metonymy becomes so on account of being an especially useful—that is to say, economical-efficient and/or effect-rich—way of picking out a certain referent. It picks out its target referent so successfully because the metonymically-used expression refers to a highly relevant, uniquely identifying property of that target referent. It is therefore plausible that, across the type or class that a given referent belongs to, the same properties will be useful identifying factors for all tokens of the type or all members of the class. For example, the referent in the ‘*ham sandwich*’ case belongs to the class of restaurant customers; for all members of that class, their food order is (in the restaurant context) a useful identifying property.

This suggests a generalisation: a ‘class-level’ identifying factor can be used to metonymically refer to any member of the class in question (e.g. ‘food ordered’ can be used to refer to any restaurant customer). Extracting, and storing in memory, patterns of this kind is useful to us in both metonymy production and comprehension. In production, it provides us with a set of ‘templates’ we can make use of on any occasion when we need to find a more efficient way of referring to a particular entity: if we can identify that entity as belonging to a specific class to which one of our stored patterns applies, all we need to do to extend the pattern to the case at hand is to find the contextually appropriate token of the class-level identifying factor and use as our referring expression the expression for which this token is the linguistically-specified referent. We are thereby likely to have met our goal of referring more efficiently, because the pattern we have applied has been extracted and stored due to its usefulness in reference-making. This could account for the productivity of metonymic patterns (see e.g. Pustejovsky (1995) on ‘metonymic polysemy’). In comprehension, metonymic patterns reduce processing effort by constraining our interpretive hypotheses. If we are able to identify the linguistically-specified referent of the metonymically-used referring term as belonging to a particular category of class-level identifying factors, we can access the relevant pattern and use it to identify the class to which the target referent belongs; we need then only home in on a specific class-member in order to recover the speaker’s intended referent. For instance, when a waitress utters (14), her colleague must identify the referring expression ‘*falafel wrap*’ as an instance of the identifying-factor type ‘food ordered’, but she can then retrieve from memory the pattern ‘food ordered for food orderer’, and this will constrain her search for the speaker’s intended referent to just the class of food orderers.

(14) The *falafel wrap* wants to know if the coffee is organic.

This is an example of a very localised pattern which is only useful, and therefore worthwhile extracting and storing, for a restricted group of language users, i.e. restaurant staff. However, there are also generalisations which are less-context specific and common to a greater number of language users, and indeed appear to be so broadly useful as to recur cross-linguistically, such as ‘container for contents’ as in (15), ‘producer for product’ as in (16) and ‘place for people’ as in (17) (Pustejovsky, 1995; Srinivasan & Rabagliati, 2015).

(15) The *kettle* is boiling. (= the water in the kettle)

(16) The guests admired his new *Picasso*. (= painting by Picasso)

(17) *France* has a healthier relationship with food than *America*. (= the French, the Americans)

Yet it is plausible that we can treat all these generalisations as arising from the routinization of metonymy interpretation, driven by the pressure to reduce processing costs in online interpretation which ultimately stems from the human tendency to maximise relevance when processing stimuli. Unlike on the cognitive linguistics approach, which sees such patterns as

somehow ‘inbuilt’, we are able to see them as emergent and, moreover, motivated. We can thus account for the crucial role of associative relations in metonymy interpretation in an explanatorily adequate manner, rather than by arbitrarily listing contiguity-based metonymic ‘mappings’, and can integrate our account into a properly inferential approach to referential metonymy.

8 Effects from frequency and routinization

Given that the frequency of a metonymic usage can lead to the routinization of its interpretation, we might also ask whether high frequency and routinization affect the storage and processing of a metonymic interpretation. This is an important question, as the answer could threaten the conclusions I have drawn concerning the nature of metonymy.

One type of answer may be that frequency and routinization create pressure to encode the metonymic interpretation of an expression as a separate lexical entry, meaning that at least certain metonymic interpretations are stored like the distinct meanings of a homonym (a case of two (or more) unrelated lexical items which happen to share a phonological form, e.g. ‘bank₁’ = side of a river, ‘bank₂’ = financial institution). For example, for the word ‘*Dickens*’, we may end up with two lexical entries, ‘*Dickens*₁’ = the author and ‘*Dickens*₂’ = the author’s works. This would threaten my claim that metonymy is not a variety of neologism, as it would suggest that in using an expression metonymically, a new standing meaning comes to be introduced into the language, albeit one which shares a linguistic form with the old (literal) standing meaning.

The hypothesis that frequent, routinized metonymic interpretations end up stored as distinct lexical entries makes several predictions, firstly that metonymic interpretations of this type will be processed differently to other metonymic interpretations. Instead of the metonymic interpretation being derived online, when the metonymic expression is encountered, representations of both its literal and its figurative referent would be initially activated before contextual information rapidly resolves the ambiguity, with the inappropriate interpretation being quickly suppressed, just as in homonymy processing (e.g. Swinney, 1979). Furthermore, we would expect to find relative frequency acting as a key determinant of the order in which the potential interpretations of the metonymic expression are accessed (Frisson & Pickering, 2007, p. 152).

However, empirical evidence suggests that these predictions do not hold. Frisson and Pickering (1999) conducted an eye-tracking experiment with a self-paced reading task, which focussed on the processing of sentences containing (i) expressions with a familiar ‘place for institution’ metonymic interpretation, e.g. ‘*convent*’ in (18); and (ii) expressions with a familiar ‘place for event’ metonymic interpretation, e.g. ‘*Vietnam*’ in (19):

- (18) That blasphemous woman had to answer to the *convent*.
 (19) A lot of Americans protested during *Vietnam*.

Their results revealed that people experience no difficulty with reading sentences like (18), where the context supports an interpretation of the sentence in which ‘*convent*’ receives its metonymic ‘institution’ reading, compared to sentences like (20), where the context supports an interpretation of the sentence in which ‘*convent*’ receives its literal ‘place’ reading:

- (20) These two businessmen tried to purchase the *convent*.

Processing times in the two conditions, literal and figurative, were equally fast, both during initial reading and later processing (as indicated by measures such as total fixation time for a word). Crucially, Frisson and Pickering (1999, p. 1376) discovered that the relative frequencies of the interpretations of a metonymic word did not affect processing times: participants were no quicker when reading a more frequent interpretation (whether literal or figurative) than when reading a less frequent interpretation. This suggests that the interpretations of a familiar metonymic item are not activated in order of frequency. By contrast, in homonymy processing relative frequency has a major effect on the order in which meanings are accessed. In addition, the ‘basic’ interpretation of a metonymic word, operationally defined by Frisson and Pickering (1999, p. 1373) as the interpretation listed first in dictionary definitions (for ‘*convent*’, the ‘place’ interpretation), was not found to be read faster than the derived interpretation (for ‘*convent*’, the ‘institution’ interpretation). This suggests that in processing frequent, routinized metonymies, language users do not initially access the ‘basic’ interpretation of the metonymic expression before subsequently accessing the derived interpretation. Frisson and Pickering (1999) thus concluded that there was no evidence for any ordering of the interpretations of a familiar metonymic item, whether in terms of frequency or of ‘basicness’. The data therefore indicate that metonymic expressions do not have the same processing profile as homonyms, which suggests that familiar, routinized metonymic interpretations are not stored in the same way as the distinct meanings of a homonym.

Frisson and Pickering (1999; see also Pickering & Frisson, 2001) argue that their results support an ‘underspecification’ model of processing for polysemous words, i.e. words that, like ‘*convent*’ and ‘*Vietnam*’, which have both a literal and a derived metonymic reading, are ambiguous between multiple different but related interpretations. The key claims of the model are that a polysemous word has only one encoded meaning which is an abstraction over the features of specific familiar interpretations, and that this single meaning is initially activated when the word is encountered. Note, however, that as Frisson (2009) highlights, there is a lack of clarity in the literature on ‘underspecified’ meanings as to whether the lexical semantic representation of a polysemous word should be seen as rich (as in Pustejovsky’s (1995) Generative Lexicon framework) or impoverished, possibly providing little more than pointers to conceptual space (e.g. Winograd, 1983; Carston, 2002; but see Carston, 2016, p. 157-8, and forthcoming, for a critique of the ‘impoverished’ position).

The underspecified meaning is taken to be compatible with all established interpretations of the polysemous word, making no distinction between primary and derived interpretations. For example, the meaning of ‘*Vietnam*’ which is initially accessed in (19) above would be underspecified with respect to whether it refers to the country or, metonymically, to the war that took place there. There is no immediate activation of any one fully specified referent. Contrast the processing of homonymous words, where all the different interpretations are initially activated. In addition, due to the fact that an underspecified meaning is the same for all related interpretations of a word, the model predicts that there should be no competition for activation between alternative interpretations, which again differs from many accounts of homonym disambiguation, where competition between meanings is taken to be the cause of processing difficulty (Rayner & Duffy, 1986). It thus appears from Frisson and Pickering’s (1999) evidence that, although a familiar, routinized metonymic interpretation of an expression is not encoded as a distinct entry in the mental lexicon, the standing meaning of the expression may nevertheless change in order to cover the established metonymic reading—yet at this stage, we cannot say more about the nature of such a change, e.g. whether the standing meaning becomes richer or leaner.

Problematically for my claim that referential metonymy is a distinct kind of lexical innovation, this change could be taken to indicate that at least certain, highly frequent and routinized metonymies are in fact examples of neologism as on Wilson and Falkum’s (2015)

account, involving the introduction of a new standing meaning, i.e. the underspecified meaning, into the language. However, the putative change can plausibly be seen as an 'adaptation' of the existing standing meaning of a metonymically-used word to cover its established interpretations—which seems fundamentally different in nature to the creation of an entirely novel standing meaning as in cases of neologism like compounds and denominal verbs, where the lexical entry for the new linguistic expression is presumably stored separately from the lexical entry for the word(s) from which it is derived. If there is indeed a change to the encoded meaning of a metonymically-used word, as Frisson and Pickering's (1999) evidence suggests, we could argue that while the resulting underspecified meaning would be undeniably novel, the change would not cause the lexical inventory to expand through the inclusion of a new standing meaning, and we could therefore conclude that no 'new coinage' takes place. The empirical data thus would not threaten my argument that referential metonymy is not reducible to neologism. Nevertheless, there is a certain tension between the notion of a change to the encoded meaning of a metonymically-used referring expression and my claim that a metonymically-used word has its old standing meaning and only its referent changes. In addition, given that the underspecified meaning for expressions with a familiar metonymic interpretation would be compatible with both the literal and metonymic readings of the expression, there is a sense in which such expressions could no longer be seen as 'metonymic', because they would no longer serve to pick out a new referent that is not covered by the linguistically-specified meaning of the expression.

However, let us consider why the postulated change to the encoded meaning of a referring expression with a highly frequent, routinized metonymic interpretation might occur. It may plausibly arise as a way to reduce processing costs, in response to usage of the expression. The resulting underspecified meaning would comprise features of, and thus be compatible with, the metonymic interpretation, which would facilitate the inferential move from encoded meaning to metonymic interpretation because on encounter with the expression, the metonymic interpretation would immediately be made highly accessible as a plausible hypothesis about the speaker's intended meaning and, according to the RT comprehension procedure which claims that interpretive hypotheses are tested in order of accessibility, may be the first hypothesis tested, should additional contextual information make this hypothesis more accessible than any other anticipated interpretations. The RT comprehension procedure further claims that our search for the speaker's intended meaning stops when we arrive at an interpretation which satisfies our expectations of relevance. Given that a frequent, routinized metonymic interpretation comes to be frequent and routinized on account of its relevance and usefulness, it is likely that on the occasion in question, the metonymic interpretation will indeed meet our expectations of relevance and will therefore be the interpretation we accept as the one most plausibly intended by the speaker. In this way, the underspecified nature of the encoded meaning of the referring expression would mean that when the metonymic interpretation is the interpretation required by considerations of relevance, we would be able to arrive at this reading quickly, with a minimum of processing effort. It is therefore possible not only to account for why the interpretation of referential metonymy may undergo routinization, as in §7, but also to propose a satisfactory explanation, which is compatible with the RT cognitive principle of relevance, of why the encoded meaning of a referring expression with a highly frequent, routinized metonymic interpretation may become underspecified.

Indeed, we could think of a diachronic progression for a metonymic reading: from an increase in frequency, depending on how relevant it is, to a routinization of its interpretation, and finally to a change to the encoded meaning of the metonymically-used expression. This may help to resolve the tension noted above, as it suggests that if the standing meaning of a metonymically-used expression changes, the change would be gradual, a function of the frequency with which the expression is used metonymically rather than an immediate and

necessary consequence of metonymic use. It also implies that such a change may not affect all metonymically-used words, only those whose metonymic usage attains an especially high frequency. We can thus acknowledge that repeated metonymic use of a referring expression may lead to diachronic changes to the standing meaning of that expression. However, we can also maintain that, for any given occasion when we use an expression to refer metonymically, although the expression comes to pick out a novel referent, its standing meaning does not change in the process of reference resolution.

9 Conclusion

Through close examination of examples of referential metonymy, comparison with other phenomena of lexical ‘flexibility’, namely modulation and neologism, and evaluation of empirical data, I have been able to elucidate the nature of referential metonymy and to show that referential metonymy can be accounted for in RT-compatible terms:

- Referential metonymy is a distinct phenomenon of lexical innovation involving the use of an existing expression to pick out a new referent.
- It is motivated by the communicative need to find a way of referring to a target entity which will enable the audience to recover the intended entity with minimum effort and/or maximum cognitive effects.
- The metonymically-used referring term is chosen because, in the communicative context, it introduces a property/feature of the target referent which is uniquely identifying and/or which conveys additional intended implications. It is therefore the best evidence of the speaker’s intended interpretation.
- The interpretation of referential metonymy is fully inferential.
- In interpretation, the metonymically-used referring expression does not come to express an ad-hoc concept. Referential metonymy therefore cannot be seen as a variety of lexical ‘modulation’.
- Further, referential metonymy does not involve ‘coining’ a new word, as in neologism.
- The interpretation of highly frequent metonymic usages may become routinized through the extraction of general metonymic patterns.
- Frequency of use and routinization of interpretation may lead to the encoded meaning of a metonymically-used referring expression changing to become underspecified as to whether the expression gets its literal or its metonymic interpretation.

Previous RT accounts of referential metonymy (e.g. Papafragou, 1996; Jiang, 2013; Rebollar, 2015) have argued for a fully inferential treatment of the phenomenon. The novel contributions here have been to closely analyse metonymy interpretation, thereby robustly demonstrating that referential metonymy is not reducible to ‘modulation’; to give an explanatorily adequate account of why we find metonymic patterns; and to offer a careful critique of the claim that metonymy can be treated as neologism, uniting theoretical arguments with experimental data from psycholinguistics in support of my position that referential metonymy is a distinct type of lexical innovation.

References

- Allott, N. (2013) Relevance Theory. In Capone, A., Lo Piparo, F., & Carapezza, M. (Eds.) *Perspectives on Linguistic Pragmatics*. New York: Springer. 57-98.
- Armstrong, J. (2016) The problem of lexical innovation. *Linguistics and Philosophy* 39 (2): 87-118.
- Carston, R. (1997) Enrichment and loosening: complementary processes in deriving the proposition expressed? *Linguistische Berichte* 8, Special Issue on Pragmatics: 103-127.
- Carston, R. (2002) *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Carston, R. (2004) Relevance Theory and the saying/implicating distinction. In Horn, L., & Ward, G. (Eds.) *The Handbook of Pragmatics*. Oxford: Blackwell. 633-656.
- Carston, R. (2016) The heterogeneity of procedural meaning. *Lingua* 175-176: 154-166
- Carston, R. (forthcoming) Polysemy, pragmatics and lexicon(s).
- Donnellan, K. S. (1966) Reference and definite descriptions. *Philosophical Review* 77: 281-304.
- Falkum, I., Recasens, M., & Clark, E. V. (2016) 'The moustache sits down first': On the acquisition of metonymy. *Journal of Child Language* 43. Doi: 10.1017/S0305000915000720.
- Frisson, S. (2009) Semantic underspecification in language processing. *Language and Linguistics Compass* 3 (1): 111-127.
- Frisson, S., & Pickering, M. J. (1999) The processing of metonymy: Evidence from eye movements. *Journal of Experimental Psychology: Learning, Memory and Cognition* 25: 1366-1383.
- Frisson, S., & Pickering, M. J. (2007) The processing of familiar and novel senses of a word: Why reading Dickens is easy but reading Needham can be hard. *Language and Cognitive Processes* 22(4): 595-613.
- Gerrig, R. J. (1989) The time course of sense creation. *Memory and Cognition* 17(2): 194-207.
- Jiang, X. (2013) Referential metonymy: Reference transfer and pragmatic motivations. *UCL Working Papers in Linguistics* 25: 1-19.
- Just, M. A., & Carpenter, P. A. (1980) A theory of reading: From eye fixations to comprehension. *Psychological Review* 87: 329-354.
- Kaplan, D. (1989) Demonstratives. In Almog, J., Perry, J., & Wettstein, H. K. (Eds.) *Themes from Kaplan*. New York: Oxford University Press. 481-563.
- Kövecses, Z., & Radden, G. (1998) Metonymy: developing a cognitive linguistic view. *Cognitive Linguistics* 9(1): 37-77.
- Lakoff, G. (1987) *Women, fire and dangerous things: what categories reveal about the mind*. Chicago: University of Chicago Press.
- Langacker, R. (1987) *Foundations of cognitive grammar*. Stanford: Stanford University Press.
- Panther, K-U., & Thornburg, L. (2003) *Metonymy and Pragmatic Inferencing*. Amsterdam/Philadelphia: John Benjamins.
- Papafragou, A. (1996) On metonymy. *Lingua* 99: 169-195.
- Pickering, M. J., & Frisson, S. (2001) The semantic processing of verbs: Evidence from eye movements. *Journal of Experimental Psychology: Learning, Memory and Cognition* 27: 556-573.
- Pustejovsky, J. (1995) *The generative lexicon*. Cambridge, MA.: MIT Press.
- Radden, G., & Kövecses, Z. (1999) Toward a theory of metonymy. In Panther, K-U., & Radden, G. (Eds.) *Metonymy in Language and Thought*. Amsterdam: John Benjamins. 17-60.
- Rayner, K., & Duffy, S.A. (1986) Lexical complexity and fixation times in reading: Effects of word frequency, verb complexity, and lexical ambiguity. *Memory & Cognition* 14: 191-201.
- Rebollar, B. E. (2015) A relevance-theoretic perspective on metonymy. *Social and Behavioral Sciences* 173: 191-198.
- Recanati, F. (2004) *Literal Meaning*. Cambridge: Cambridge University Press.
- Sperber, D., & Wilson, D. (1986/1995) *Relevance: Communication and cognition*. Oxford: Blackwell.
- Sperber, D., & Wilson, D. (1998) The mapping between the mental and the public lexicon. In Carruthers, P., & Boucher, J. (Eds.) *Language and Thought: Interdisciplinary Themes*. Cambridge: Cambridge University Press. 184-200.
- Sperber, D., & Wilson, D. (2002) Pragmatics, modularity and mind-reading. *Mind & Language* 17: 3-23.
- Srinivasan, M., & Rabagliati, H. (2015) How concepts and conventions structure the lexicon: Cross-linguistic evidence from polysemy. *Lingua* 157: 124-152.
- Swinney, D. A. (1979) Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior* 18: 645-659.
- Tyler, L. K., & Marslen-Wilson, W. D. (1977) The on-line effects of semantic context on syntactic processing. *Journal of Verbal Learning and Verbal Behavior* 16: 683-692.
- Wilson, D., & Carston, R. (2007) A unitary approach to lexical pragmatics: relevance, inference and ad hoc concepts." In Burton-Roberts, N. (Ed.) *Pragmatics*. London: Palgrave. 230-259.

-
- Wilson, D., & Falkum, I. (2015) *Explaining metonymy*. Talk presented at the conference on New Developments in Linguistic Pragmatics, Lodz, Poland, April 2015.
- Wilson, D., & Sperber, D. (2002) Truthfulness and relevance. *Mind* 111: 583-632.
- Winograd, T. (1983) *Language as a Cognitive Process. Volume i: Syntax*. Reading, MA.: Addison-Wesley