Linguist and Philos (2015) 38:1–29
DOI 10.1007/s10988-014-9164-6

RESEARCH ARTICLE

# Whole-for-part metonymy, classification, and grounding

**Alexandra Arapinis**

Published online: 17 February 2015


**Abstract** Since the early 1980s, metonymy has progressively gained central stage in linguistic investigations. The advent of cognitive linguistics marked a new turn in the study of this trope conceived, not as a deviation from semantic conventions (contra classical rhetorical theories), but as a phenomenon rooted in non-language-specific mechanisms of conceptualization of the world. Acknowledging that metonymy is ultimately cognitive in nature, this paper proposes to consider metonymy from its multiple levels of manifestation, integrating cognitive, pragmatic, semantic, but also ontological angles of approach. Taking whole-for-part (WP) metonymies as a case study, I aim to show how recent developments within these respective disciplines can enrich our understanding of such metonymic mechanisms, sometimes without even identifying them as such. This paper proposes to establish a dialog between these disciplines on the topic of WP-metonymy. So, after a presentation of the most standard cognitive and pragmatic approaches to WP-metonymy, I will argue for the relevance of recent semantic investigations on quantity gradability, and for the theoretical importance of keeping these two kinds of part-reference clearly apart. I will show that the literature on gradability provides strong semantic arguments for doing so. Finally, connecting the debate on WP-metonymy with the ontological debate on property inherence will open the way for a formal treatment of WP-metonymy within ground logic.

**Keywords** Metonymy · Whole–part · Classification · Pragmatics · Cognitive linguistics · Lexical semantics · Ontology · Grounding

A. Arapinis (✉)
Laboratory for Applied Ontology, Institute of Cognitive Science and Technologies, CNR, via alla Cascata 56/C, Povo, 38123 Trento, Italy
e-mail: aarapinis@gmail.com

A. Arapinis
Institut d'Histoire et de Philosophie des Sciences et des Techniques, IHPST (CNRS/Paris 1/ENS), 13 rue du Four, 75006 Paris, France

## 1 Introduction

Since the early 1980s, metonymy progressively gained central stage in linguistic investigations. The advent of cognitive linguistics marked a new turn in the study of what was long thought to be a trope or figure of speech; a turn that was coupled with a radical distancing from predominant structural semantic approaches to natural language. Instead of treating metonymy as a deviation from semantic conventions, the way traditional semantics did, cognitive linguistics opted for a more integrative view of language as part of the wider cognitive system, rooting metonymy (alongside metaphor) in non-language-specific mechanisms of conceptualization of the world. In a chapter dedicated to metonymy, Lakoff and Johnson (1980) were the first to point to the fact that metonymic concepts, indirectly pointing to an entity by referring to a related one, were in fact systematic in that they "structure not just our language but our thoughts, attitudes and actions […] [and are] grounded in our experience" (p. 39). In the same spirit, Panther and Thornburg (2004) more recently characterized "conceptual metonymy as a contingent, i.e. non-necessary, relation within one conceptual domain between a source meaning and a target meaning, in which the source meaning provides mental access to the target meaning. […] [S]uch metonymic relations [are regarded] as multipurpose conceptual devices not restricted to language but used in other semiotic systems and thinking as well. Furthermore, […] in a prototypical metonymy the source meaning is conceptually more prominent, i.e. more in the focus of attention, than the target meaning" (p. 91).

Starting from this very idea that metonymy is ultimately cognitive in nature, this paper proposes to consider metonymy from its multiple levels of manifestation, integrating cognitive, linguistic, but also ontological angles of approach. Taking whole-for-part (WP for short) metonymies as a case study, e.g.

*The windmill* [*the vanes*] *is turning*
*The pen is red* [*the ink, the surface*]
*My car* [*the engine*] *broke down*

I aim to show how recent developments within these respective disciplines can enrich our understanding of such metonymic mechanisms, sometimes without even identifying them as such. This paper proposes to establish a dialog between these disciplines on the topic of WP-metonymy. Through this confrontation of traditions and approaches, I will in particular establish that what is referred to as WP-metonymy in the cognitive and pragmatic literature, presented in Sect. 1, has also been broached under the heading of "classificatory properties" in recent lexical semantic work on gradable adjectives. Section 2 will thus focus on lexical semantic constraints on WP-metonymy, and in so doing, will dissolve a confusion that is often found in the cognitive literature between quantity gradability and WP-metonymy. This will prove all the more important given that the mereological relations that underpin these two kinds of conceptual/linguistic phenomena differ in nature (corresponding respectively to purely extensional and functional parthood relations). Therefore, clearly distinguishing the two will allow a sharper characterization of the conceptual relations underlying WP-metonymy in general.

Finally, Sect. 3 will complement our semantic analysis by introducing a related ontological debate concerning particularized qualities and the principle of bearer uniqueness (e.g. if particularized properties are individuated by their unique bearer, is the sharpness of the knife the same particularized quality as the sharpness of its blade?). Again, I will argue that cases that seem to defy the principle of bearer uniqueness for particularized properties, i.e. cases of so called "partitive predication", are precisely cases of WP-metonymy, where the (particularized) property attributed to the whole is a (particularized) property manifested by one of its parts. Bringing this last level of analysis into the picture will in particular allow us to use formal ontological tools for the semantic characterization of WP-metonymy. This will further bring out a parallel between the central, though admittedly imprecise, notion of noteworthiness used in cognitive and pragmatic approaches to metonymy, and the explanatory relation of grounding formally worked out by ontologists. Most importantly, we will see that combining part–whole relations with the notion of grounding results in a kind of functional parthood. This is in accordance with the general cognitive principle specifying that metonymic patterns do not rest on objective relations among entities (such as extensional part–whole relations), but are relative to the way we conceptualize the world in different contexts. The proposed analysis will thus dissipate the felt impression that WP-metonymy builds on mind-independent, purely objectual part–whole relations, which would be at odds with the assumed relativity of metonymy to conceptual domains.

As a whole, the paper follows the above-mentioned levels of expression of WP-metonymy along an increasing degree of constraints imposed on its realization. The different terminologies used to refer to the same underlying conceptual mechanism will be discussed at each level: "WP-metonymy" at the cognitive/pragmatic level, "classificatory predication" at the lexical semantic level, "partitive predication" at the ontological level. Once introduced, these terminological variations will be used interchangeably.

## 2 Previous accounts of WP-metonymy

### 2.1 Conceptual contiguity: the cognitive approach

While primarily concerned with metaphor, Lakoff and Johnson's seminal work *Metaphors We Live By* (1980), set the agenda of many of the cognitive linguistic investigations that followed, and renewed the interest for metonymy as a conceptual mechanism of primary cognitive importance. From there, cognitive linguistics progressively pulled metonymy at the central stage of linguistics, and out of its long dedicated rhetoric discipline. Broaching language from the perspective of the more general (non-specifically linguistic) human capacity to conceptualize the experienced world, metonymy is henceforth undisputedly acknowledged by cognitive linguists as a fundamental cognitive mechanism that originates in perception and bodily experience, and is as such constitutive of our everyday way of thinking.

Though summarizing *the* cognitive semantic approach to metonymy is a perilous task, given the lack of consensus in the field, some general principles do seem to be

indisputably shared. In particular, contrary to the traditional semantic view according to which words express well delimited meaning units that specify satisfaction conditions, cognitive linguists start from the hypothesis that linguistic symbols always evoke whole domains of knowledge. Concepts never come in isolation, and expressions do not merely refer to entities in the world. They do so against different frames of encyclopedic knowledge corresponding to different perspectives one might have on a given object or situation. So, different expressions can denote the same entity, but through concepts that relate to different domains of knowledge. As Croft and Cruse put it:

> Certain concepts 'belong together' because they are associated in experience […] They are not related […] by structural semantic relations; they are related […] by ordinary human experience. […] The need for another means for organizing concepts has been felt by researchers in cognitive psychology and artificial intelligence as well as in various branches of linguistics, and has led to a variety of similar proposals, each typically with its own name. Among these names are: frame, schema, script, global pattern, pseudo-text, cognitive model, experiential gestalt, base, scene (Croft and Cruse 2004, pp. 7–8).

Such domains or frames form coherent regions of knowledge, and function as the necessary basis against which (linguistically expressed) concepts are understood, and from which they cannot be alienated. This is most clearly evidenced by comparing coextensive expressions such as *meat* vs. *flesh*. While the two concepts share denotation, *flesh* is profiled against the domain of body anatomy, whereas *meat* evokes the domain of food (Langacker 1984, 1987, 1999).

In such a conceptual setting, metaphor and metonymy are conceived in terms of our general capacity to map information across and within conceptual domains. It is thus admitted that while metaphor builds on mappings across distinct domains (e.g. time *as* movement: *time flies by*), metonymy involves intra-domain mappings (e.g. author *for* book: *Proust is tough to read*). Though this might be slightly more controversial, one way to put the difference is this: metaphor draws an *is like* comparison relation between structurally similar source and target domains, while metonymy draws a *stand for* connection between entities in the same domain. Different kinds of metonymy (and metaphor) are thus individuated through the conceptual mappings they rely on, such as PRODUCER FOR PRODUCT (e.g. *I've got a Ford* for 'car'), POSSESOR FOR POSSESSED (e.g. *that's me* for 'my bus'), INSTRUMENT FOR PRODUCT (e.g. *did you hear the whistle?* for 'the latter's sound'), and many more, among which are WHOLE FOR PART metonymies which will retain our exclusive attention in this paper.

(1) a. The windmill [*the vanes*] is turning
    b. My car [*the engine*] broke down
    c. A sharp knife [*the blade*]
    d. A red pen [*the ink, the surface*]
    e. Dirty trousers [*some part*]

As with every other kind of metonymy, of which the list given above is but a very small sample, WP-metonymy operates an indirect focalization of a part *B,* by reference to its whole *A,* in contexts where *A* is more salient or easy to grasp than *B*. Following this line, the pervasiveness of WHOLE FOR PART metonymies, and perhaps the fact that they often go unnoticed, is naturally explained by that fact that, for humans, ordinary objects form natural units of perception and experience, whose salience is much greater than that of their parts. Furthermore, part–whole relations are cognitively so fundamental that they span through a great variety of conceptual domains, giving rise to less domain specific kinds of metonymies (compared for instance to the widely debated ORDER FOR CUSTOMER metonymy, e.g. *the ham sandwich left without paying*, which seems limited to the RESTAURANT frame and is therefore much more noticeable).

Besides this very general characterization of metonymy, on which probably most cognitive semanticists would agree, significant divergences arise when looking at the more detailed analysis of the data and the specific patterns they instantiate. In the case of WP-metonymy in particular, scholars have assumed part–whole relations to be more or less restrictive. Indeed, the examples given in (1) clearly involve proper subparts of the literal denotation. But what about (2)?

(2) a. She buttoned herself up [*a piece of clothing*] (Kleiber 1999)
b. The kettle is boiling [*the water*] (Langacker 1999)
c. The book is highly instructive [*the content*] (Barcelona 2008)
d. Exxon has raised its prices again [*a person responsible*] (Lakoff and Johnson 1980)
e. I have a temperature [*fever*] (Seto 1999)

All of these examples are found in the literature under the heading of WP-metonymy. However, it can be argued that they do not involve proper meronymic relations, at least following the linguistic tests for meronymy given in Cruse (1986) and Winston et al. (1987), probably the two most insightful and influential studies in the lexical semantics of part–whole relations. Example (a) involves an extension of part–whole relations to mere attachments, while (b) extends meronymic relations to topological inclusion. Examples (c) and (d) further involve part–whole relations on the whole conceptual domain rather than the denoted entity. It is thus argued that the informational content is a subdomain of the matrix–domain of books, which also includes a subdomain of physical objects. In the same way, the conceptual domain of institutions can be argued to include an abstract institutional subdomain along with an agentive subdomain. Finally, (e) involves category inclusion, viz. GENUS FOR SPECIE synecdoche (the whole scale of temperatures for those approximately above 37°C), also often assimilated to a kind of WP-metonymy.

Depending on the level of analysis, one may not wish to dwell on the specificities of each metonymic mapping, and may opt for a general unifying description of the above-mentioned phenomena. It is expected for instance that the linguistic and cognitive levels of analysis, though sharing an important number of explanatory principles, will constrain the exploitation of conceptual mappings in different ways. Deciding whether to describe particular metonymies as instances of the WHOLE

FOR PART pattern therefore leads to more fundamental issues like: How does our linguistic system constrain the great variety of conceptual metonymic patterns? Are all the conceptually relevant metonymic patterns actually expressed in language? And do they render similar linguistic effects? As pointed out by few cognitive linguists, language does not actually realize all the metonymic connections that could be relevant from a purely cognitive point of view. Thus, the conceptual analysis of metonymic mappings cannot dispense with a properly linguistic research.

For instance, there are significant differences among the examples in (2) concerning the availability of the metonymic target for subsequent predication (anaphoric or co-predicative). Thus, compare the following anaphoras built on (1.b), (2.b) and (2.c) respectively

(1.b′) My car broke down [*the engine*] though I changed it [**engine*] last week
(2.b′) The kettle [*the water*] is boiling. It [?*the water*] is hot
(2.c′) I didn't find this book [*the content*] very interesting so I returned it [*the physical copy*] to the library

At least from a purely linguistic point of view, the semantic distinction among different kinds of part–whole relations and other close parents does seem to make a difference. And even if one were to insist that they share a common conceptual core, viz. that they all qualify as WP-metonymies in some broad sense, this core would certainly have to be further qualified to account for the attested semantic divergences.

As we will now see, coordination tests have in particular been at the focus of Nunberg's (1995, 2008) and Stallard's (1993) respective work on the linguistic distinction between referential and predicative metonymy. According to Nunberg, "the various mechanisms of transfer can't be distinguished simply by pointing at the type of correspondence they exploit. And for this reason, the description of these mechanisms is fundamentally a linguistic problem, rather than a problem of conceptual analysis" (1995, p. 112). So let us turn to the linguistic counterpart of metonymic mappings and the linguistic issue of the locus of WP-metonymy (viz. referential vs. predicative). I will then look into lexical semantics and see how the predicative content actually triggers WP-metonymies, involving what I will call partitive[1] predicate uses. In so doing I will build on recent work on the semantics of color adjectives and part-structure modifiers. Throughout this semantic analysis, I aim to contribute to current lexical semantic debates, but I also hope to shed new light on the broader conceptual issues raised by WP-metonymy: (i) the terminological heterogeneity found in classifications of WP-metonymies and (ii) the formalization of "noteworthiness" or "salience" conditions on WP-metonymies.

## 2.2 WP-metonymy in language: predicative or referential?

According to Nunberg's pragmatic theory (also in: Recanati 2004), metonymy is a linguistic mechanism of meaning transfer that makes it possible to use the same

[1] I borrow this terminology from Schnieder (2004), to which I will return in Sect. 3.

expression to refer to disjoint and functionally related sorts of entities. Relying mainly on tests like anaphora and co-predication to distinguish between different kinds of linguistic metonymy, Nunberg identifies two sorts of meaning transfer: reference transfer (or deferred reference) and predicate transfer. Typical cases of deferred reference are:

(3) a. The omelette is at table 7
b. This (pointing to car keys) is parked outback

While the following examples illustrate predicate transfer:

(4) a. I am parked outback
b. Billy's shoes were neatly tied

In all these examples, there is indeed an intuitive functional relation between entities that allows for the metonymic readings. In (3a) the relevant functional relation is that between a restaurant client and the food he ordered; (3b) involves the relation between a car and the keys that open it; (4a) relies on the ownership relation between a car and a person; and finally (4b) involves the relation between shoes and their laces.

Nunberg further notes that, based on the sole nature of these transfer functions, there are two possible ways to proceed with the linguistic analysis: either (i) these functions operate on objects, allowing the subject term to refer, not to its literal denotation, but to a functionally related entity, or (ii) they can operate on the predicate, allowing it to express a different though related property. (3a) for example could be analyzed in two ways. Either the nominal phrase (NP) 'the omelette' has a transferred reference and denotes the person that made the order instead of the food ordered; or the predicate 'is at table 7' has a transferred reading and expresses the property of having been ordered by someone sitting at table 7.

According to Nunberg the choice between these two alternatives has to be guided by purely linguistic criteria, viz. using standard tests of anaphora and co-predication. Thus, while it is acceptable to extend (3a–b) respectively by adding

(5) a. He is getting impatient
b. It's a big Mercedes

it would be very odd to pursue the same sentences by

(6) a. * It is overcooked
b. * And fits only the front door

This is taken to show that, in (3a–b), there is in fact a referential shift, that is, 'the omelette' metonymically refers to the person who ordered it, and the demonstrative 'this', pointing to the car keys, refers to the car they open.

The same test is taken to establish that in (4a–b) the referent of the argument is left unmodified. What is transferred, Nunberg argues, is the predicate. Indeed, he points to the fact that predicate transfer naturally accounts for anaphora and

co-predications where the target is obviously the original literal reference of the subject NP.

(7) a. I am parked outback but I'll walk back home
b. My shoes are tied and dirty

Note that in (7b), the predicate 'dirty' couldn't be understood as applying to the laces, as opposed to the predicate 'tied'.

Following this criterion, WP-metonymy would be a kind of predicative metonymy. For example, with respect to (4a) (viz. *Billy's shoes were neatly tied*), Nunberg (1995) notes that "either *shoes* has a transferred reading where it refers to the shoe laces, or *tied* has a transferred reading, where it denotes the property that shoes acquire when their laces have been tied". But while both alternatives are equally plausible conceptually, anaphora and co-predication show that the metonymy operates on the predicate rather than the subject NP.

(8) Billy's shoes were neatly tied. They were blue.

(9) a. Billy's shoes were neatly tied but dirty.
b. ?? Billy's shoes were neatly tied but frayed.

Example (8) shows that the un-shifted nominal denotation is available for anaphoric binding—*blue* can only be interpreted as applying to the whole shoe and not the laces alone—while (9) shows that co-predication is only possible with predicates applying to the whole shoe, and is infelicitous with predicates that can only apply to the laces, like *frayed*.

However, if one follows Nunberg in using anaphora and co-predication tests to identify reference transfer, it seems that verbal ellipsis should play the same role in assessing the presence of predicate transfer. For example, consider the small dialog below:

A: Your shoes are too loosely tied
B: So is your bowtie.

If the predicate in A's utterance is shifted, roughly expressing the transferred property of *having laces that are too loosely tied*, then this analysis should seemingly further predict that B's utterance is equivalent to something like *Your bowtie has laces that are too loosely tied*. But this is certainly not a possible interpretation. Of course, even if Nunberg's theory doesn't address this issue, it is not theoretically insurmountable. In fact, there would be different ways to go to account for divergences in ellipsis. One possibility would be to defend that the transferred meaning is underspecified (a proposition along these lines is worked out in Egg 2003). An alternative would be to suppose that the transfer actually takes place once the elided material has been recovered. This will however not be discussed here. First, because Nunberg never actually developed such arguments. Second, because any theory that would account for divergences between the explicit

verbal phrase and the elided one would expectedly extend to anaphora, thus jeopardizing Nunberg's own test for reference transfer.

The ellipsis test in fact corroborates the felt impression that, in WP-metonymies, the predicate retains its original meaning. When I say for instance that a shirt is dirty, I am really attributing the property of being dirty. The predicate is thus not strictly speaking shifted, but exploits the inner-structure of the object it applies to in order to highlight a piece of it. The fact is that, even under Nunberg's analysis, predicative metonymy is not built on mappings between properties or relations in themselves. Rather, the idea is that the metonymy is always built on relations between entities, but that this relation is added to the content of the linguistic predicate in cases of predicative metonymy.[2] In this respect, Stallard's account of predicative metonymy is somewhat more appealing. As Stallard suggests, predicative metonymy should better be analyzed along the lines of predicative coercion, where the argument place is modified, but the argument NP itself remains unchanged. Under this account, predicative metonymy would thus involve no shift (as opposed to referential metonymy) in the sense that both the original meaning of the NP and the predicate remain unchanged (see Pustejovsky 1995, on coercion). In predicative metonymies, among which are WP-metonymies, the predicative content operates to highlight an entity metonymically related to its direct object (in our case a part). But following Stallard's definition of predicative metonymy, this involves no shift at all, neither in the meaning of the verb, nor of its expressed direct object.

It thus seems that the traditional account of metonymy as strictly involving a referential shift along some conceptual connection among entities could be preserved after all by sharply distinguishing between metonymy and coercion. One could argue that what is called predicative metonymy by Stallard is in fact not a metonymy, but rather a form of coercion. (Prandi 2004) pushes in this direction and restricts metonymy to mechanisms that operate on saturated entities (viz. object or process) and maintains referential shift as a definitional feature of metonymy. However, since the aim of the present paper isn't to settle such issues on metonymy itself, and knowing that coercion is generally presented as a *logical* form of metonymy (Pustejovsky 1995), I believe that breaking coercion apart from metonymy risks to confuse the reader while not crucially affecting the core of the argument. Indeed, like Stallard, Pustejovsky also defines coercion as an operation "[that] create[s] an *extension* of the NP meaning, [by] metonymic reconstruction." (Pustejovsky 1995, p. 200).

To conclude, though there is most probably no predicative meaning shift involved in WP-metonymy, and more generally in cases of predicative metonymy, it seems that one should look at the predicative content to see how it exercises pressure to highlight parts of its direct object. In other words, we now have to turn to a more detailed lexical analysis to identify the coercive effects of the predicates involved in WP-metonymy.

[2] Interestingly, a vast majority of cases of predicative metonymy discussed in the literature are cases where the predicate includes an NP that is interpreted as having a shifted reference (see Sweep 2009).

## 3 The systematic ambiguity between partitive and quantity-gradable readings

Interestingly, many of the examples of WP-metonymy found in the cognitive linguistic literature have been independently discussed in the domain of scalar and degree semantics. Examples like (10) have indeed received great attention from lexical semanticists studying gradable adjectives and comparatives.

(10) a. A red pen [*the ink, on the outside*]
b. A car that's damaged [*the engine*]
c. Dirty trousers [*some part*]

However, the existing gap between structural semantics and conceptual/cognitive approaches to language has resulted in a total absence of dialogue between these two traditions. Taking note of this common (though underexplored) ground, I propose to bridge this gap and see to what extent gradable adjectives supporting a quantity interpretation can trigger WP-metonymic readings. A closer look at the class of "quantity gradable adjectives" will show that they in fact support two quite distinct interpretations exploiting different mereological structures: a partitive use and a properly quantity-gradable use. As we will see, the latter of these interpretations rests on purely extensional mereological relations (division into portions), while the former involves functional mereological relations and functional parts. These have often been conflated in the cognitive linguistic literature, thus blurring the mereological constraints on WP-metonymy.

### 3.1 Gradable adjectives with quantity-based scales

In the discussion to follow, I will make use of the terminology of degree-based approaches to gradability, which is the most refined and most widely used among linguists.[3] Note however that, although it borrows the degree-terminology at some points, the main argument does not really take stance on the semantics of gradability and does not in any way commit to this particular approach. For this reason, this section will also remain informal and present the relevant data without entering into the details of their formal treatment.

Looking at the data, the usual way of arguing for the gradable nature of a predicate is by looking at comparative structures, measure and proportional modification. The fact that the examples in (10) can naturally be embedded in comparative constructions (11.a), and combined with measure and proportional modifiers (11.b) is generally taken as evidence of the gradability of the adjectives they modify (viz. *red*, *broken*, *dirty*).

[3] I will in particular heavily rely on Kennedy and McNally (2005) because they offer a good overview of the rich literature on gradability. Further significant references in this field can be found in their paper.

(11) a. This pen is redder than that one
This car is more damaged than that one
These trousers are dirtier than those ones
b. This pen is partly (half, 50%, etc.) red—This pen is completely (entirely, 100%, etc.) red
This car is partly (half, 50%, etc.) damaged—This car is completely (entirely, 100%, etc.) damaged
These trousers are partly (half, 50%, etc.) dirty—These trousers are completely (entirely, 100%, etc.) dirty

In fact, there are more than one possible readings for these sentences, corresponding to different scales along which the gradable properties are compared or measured. In particular, there is an obvious quantitative reading and a qualitative reading. An object can be more or less red, respectively damaged or dirty, in extent or intensity. In the present context, it is of course the quantitative gradable reading that is of primary importance, a reading that is available when adjectives can apply both to the whole and to its parts. "In general, any adjective that can apply not only to the whole of an entity but also to its parts could also be associated with a scale whose structure is based on the part structure of that entity" (Kennedy and McNally 2005, p. 365).

Furthermore, following Kennedy and McNally, gradable adjectives can be either absolute or relative, depending on certain features of the associated scale structure. For the present purpose, it is not needed to enter into the detailed description of scale structures as developed by the authors. The following informal characterization suffices to grasp the basic idea driving the proposed classification: relative gradable predicates are those predicates whose standard of comparison for the applicability of the predicate varies with the context, while absolute gradable predicates have a fixed standard of comparison for their applicability. Furthermore, absolute gradable adjectives come in two kinds; those with an absolute minimal standard and those with an absolute maximal standard. Considering quantitative scales only, this means that some gradable predicates, under their normal unmodified usage, require that the property be satisfied to a minimal extent, while others require that the property be satisfied by the entire entity.

To give flesh to these distinctions, consider the predicates *visible*, *flat* and *cooked*, which exhibit respectively an absolute minimal, an absolute maximal and a relative standard of comparison (Kennedy and McNally 2005). Being visible imposes a minimal comparison standard, given that the truth of a sentence like "*The spot is visible*" minimally requires a small part of the spot to be visible. On the contrary, by saying that "*The road is flat*" one can only mean that the entire road is flat, not just any minimal part. These two examples further contrast with gradable adjectives like *cooked*. According to the standards of French cuisine, a red meat qualifies as being cooked even if it is not cooked throughout (many would even say that it is overcooked when it is cooked throughout). But this is certainly not an absolute standard, as many other culinary cultures would have it that a piece of meat is raw if it is merely seared. On the quantitative reading, an absolute minimal standard is minimally greater than the total lack of the property, an absolute maximal standard

requires the presence of the property throughout, while a relative standard is contextually fixed somewhere along the way.[4]

As I said, Kennedy and McNally correlate their typology of standards of comparison with structural features of the scale involved, depending on whether the scale is totally or partially open or closed. But the important point here is that quantitatively gradable predicates with absolute minimal and relative standards of comparison can be true of an object when only part of it (more or less important) exhibit the particular quality. Under this definition, all such predicates then give rise to something similar to WP-metonymies. I will return to this point in a moment. For the time being, let us just mention that, the predicates in (10) are all gradable predicates with a minimal or a relative standard.

Focusing on color adjectives, Kennedy and McNally (2010) thus show that they all exhibit a qualitative and a quantitative gradable reading, as corroborated by (11). Furthermore, the extent to which an object has to be of color *X* (analogously, how intensely *X* it has to be) to qualify as *X* is argued to depend on contextually determined standards. Tomatoes for instance are said to be red, but it doesn't mean that absolutely all their parts are red. This is confirmed by the following tests.

(12) a. The tomato is red but it could be redder
b. Both of these tomatoes are red, but this one is redder than that one

This conforms to the criterion for relative adjectives, according to which "the truth conditions for a relative adjective entail only that its argument falls above a contextually determined standard of comparison" (Kennedy and McNally 2005, p. 359).

The predicates *dirty* and *damaged*, on the other hand, arguably impose a minimum-standard. They only require the minimal satisfaction of the corresponding property for the predication to be true. Looking at their negation provides evidence for this. Indeed, the negation of an absolute minimal adjective should entail the total absence of the corresponding property, as shown in (13) for *dirty* and *damaged*.

(13) a. ?? These pants are not dirty, but there is some dirt on them
b. ?? This car is not damaged, but there is a part of it that is damaged

In fact, these adjectives, together with their antonyms, belong to the particular class of total/partial (or universal/existential) pairs of adjectives (Yoon 1996; Rotstein and Winter 2004). These pairs have been shown to share the following common features. Each pair is associated to a base condition such that the existential/partial predicate only requires its argument to satisfy the base property to a certain extent, while the universal/total adjective requires its argument to be entirely free of the

[4] Opposed to the view I have just presented, a number of voices have been raised against the possibility of marking such a clear-cut distinction between absolute and relative adjectives. These have pointed to the fact that the notion of minimal departure from the lower degree on a scale was subject to contextual variations, and that a certain distance (minor exceptions) with respect to the maximum satisfaction of a property could be admissible in different contexts. As this is somewhat orthogonal to the present discussion, I will just refer the reader to McNally (2011) and Toledo and Sassoon (2011).

base condition.[5] Among such pairs are the antonyms *wet/dry*, *sick/healthy*, *flawed/flawless*, and many more (see Yoon 1996 for a partial list). A wet towel, for example, isn't necessarily an entirely wet towel, and it doesn't need to be entirely covered with dirt in order to be considered dirty. On the other hand, to be dry or clean means to be totally free of humidity and dirt respectively.

Again, the behavior of total and partial adjectives and their respective characteristic standards are best seen in light of modifier distributions and the inferences they support. Typically, partial adjectives allow inferences to the positive whereas total adjectives do not necessarily.

(14) a. X is more (Partial-Adj) than Y ⇒ X is (Partial-Adj)
b. X is more (Total-Adj) than Y ⇒ Y is not (Total-Adj)

(15) a. X is dirtier than Y ⇒ X is dirty ⇒ *but they are both clean* is unacceptable
b. X is cleaner than Y ⇒ Y is not clean ⇒ *but they are both dirty* is acceptable

Furthermore, there seems to be a clear distribution of modifiers such as *almost*, *perfectly*, *slightly* (Rotstein and Winter 2004).

(16) a. It's almost/perfectly clean
b. The towel is almost/perfectly dry
c. * It's almost/perfectly dirty
d. * The towel is almost/perfectly wet

(17) a. It's slightly dirty
b. The towel is slightly wet
c. * It's slightly clean
d. * The towel is slightly dry

Contrary to what much of the cognitive linguistic literature has suggested in dealing with these and similar examples, what is involved is not a WP metonymic mapping, but the predication of quantitatively gradable adjectives with non-maximum standards of application. As Kennedy and McNally put it, "[t]he fact that degree modifiers [and entailments] are sensitive to these features [viz. to the different default standards associated to gradable adjectives] argues for encoding them in the lexical semantics of gradable expressions" (2005, p. 348). The question then is: Does WP-metonymy reduce to gradability? If so, the analysis of metonymy as requiring a cognitively significant relation between the metonymic source (the whole) and the metonymic target (the part) would be seriously compromised. I will now endeavor to show that this is not the case. More precisely, I will argue that we have to distinguish between quantity gradable and partitive readings of predicates, a

[5] The different terminologies that have been used in the literature to describe this phenomenon reflect the distinct semantic approaches of the authors. More precisely, Yoon (1996) introduced the existential vs. universal terminology as she conceived of such pairs of antonyms in terms of quantification over parts. On the other hand Rotstein and Winter (2004) opted for a degree-based approach, the relevant standards associated to such pairs of adjectives being qualified as partial vs. absolute.

distinction that is semantically relevant and is important for the more general understanding of the conceptual ground of WP-metonymy.

### 3.2 WP-metonymy with gradable adjectives: partitive reading

The main argument that militates for distinguishing partitive readings from quantity readings of gradable adjectives draws on a small but significant point that was recently made by scholars working on color predicates. They noticed that it was perfectly correct to say of an apple

(18) This apple is completely red,

meaning that it is completely red *on the outside*. Kennedy and McNally (2010) briefly mention this data, suggesting that this points to a distinct non-gradable reading of color predicates.[6] Indeed, (18) appears to involve a use of the color predicate that does not really depend on how much of the object has a given color, but *where* it is of that color. The color property is relativized to a given part, and it is with respect to the relevant part that there is a gradation as to how much of it is of the given color. Now, such uses of color adjectives aren't isolated cases, and one could easily multiply the examples

(19) a. This orange is completely/almost red [*the flesh of a blood orange*]
b. This car is entirely/half/part blue [*the outside*]

Note also that similar examples are easy to come up with using partial predicates

(20) a. This table is completely/partly dirty [*the top surface*]
b. The bench is completely/partly wet [*the seating surface*]

Most importantly, in the case of partial/total pairs of adjectives, it appears that such partitive uses are completely independent from the partial nature of the predicate (contrary to what one could be tempted to think). This is supported by the fact that analogous examples can be given with their total antonyms, involving the very same parts.

(21) a. The table is (completely) clean [*the top surface*]
c. The bench is (completely) dry [*the seating surface*]

A further argument for distinguishing purely gradable uses of partial predicates from partitive ones comes from considering their application to plural arguments.

[6] In a version anterior to the one published, Hansen (2011) further developed the idea of part-variability within the degree-based approach. As it was cut from the final version, I will not enter into the details of his view. Though this paper follows a different modeling path, it is however important for me to mention this work as it brought the phenomenon to my attention, and reading this non-final version helped me shape my own view.

(22) a. The tables are dirty
b. The red pens are in the top drawer

In attributing dirtiness to tables, it seems quite reasonable to consider that while for eating purposes the relevant part is the top surface, in a surgical context, it is the whole table that is relevant. But it is impossible to use (22a) to describe a plurality of tables, some of which are dirty "in the restaurant sense" and others "in the surgery sense". The partitive use of the partial predicate *dirty* to a plurality imposes a uniform conceptualization of all the elements of the plurality.

Once the property *P* is (distributively) applied to the elements of the plurality with respect to a certain part, the partiality of *P* accounts for the fact that the part in question needs not be wholly *P*. So each table can be more or less, completely or partly dirty with respect to the relevant part. Note also that while the focalized part must be the same for all the elements of the plurality, the portion of that part that will satisfy the property, and license the dirtiness attribution, need not be the same for all the elements. For example, dirty tables "in the restaurant sense" will all be dirty with respect to their top surface, but each table can be dirty on different portions of its surface. The same point can be made with respect to (22b) and color predicates in general. One has to identify the salient part of the object to measure how much of it is of a certain color, and in dealing with pluralities, the relevant part has to be the same for each member.

Finally, Kennedy and McNally (2010) draw attention on the fact that it is impossible to compare two objects to determine which one is quantitatively more of a given color, when the two share the same color but with respect to different dimensions or parts. For example, a pen that is blue on the outside cannot be compared to a pen with blue ink to determine which one is bluer. The question doesn't make any sense. In this respect too, partial predicates and color predicates can be shown to behave analogously under their partitive use. To see this, consider the following scenario:

> Bob proposes to take Mary back to her house with his car. Mary looks inside Bob's car and answers: "Your car is too dirty. I refuse to get in!". She then adds: "I'll go with John. His car is cleaner."

It is reasonable to take Mary's first utterance to mean that Bob's car cabin is dirty. This is why she refuses to get in. It is also clear that her second utterance is suitable if John's car is cleaner than Bob's *with respect to the same part*. Taking into account the dirt on the engine of John's car to compare it with the state of dirt of Bob's car cabin would simply make no sense in this context.

One way of understanding the difference between gradable uses of predicates with non-maximal standards and their partitive uses is by appealing to Cruse's distinction between functional parts and pieces (1986). Intuitively, a piece is an arbitrarily delimited portion, characterized by mere topological inclusion, and therefore lacks the following three characteristics of functional parts: autonomy, non-arbitrariness of boundaries and determinate function with respect to the whole. To illustrate these features, Cruse (1986, pp. 157–159) uses the examples of human anatomy and typewriter parts vs. pieces. Pieces of a human body, respectively of a

typewriter, would typically result from cutting the body or typewriter with a hacksaw. By contrast, the separation of functional parts from the typewriter would require "undoing of screws and other means of attachment", and the separation of functional parts from a human body would have to be operated at their "natural" joints (e.g. the hips, the knees, the elbows, the wrists, etc.). The latter parts, but not the pieces, are autonomous in the sense that they would qualify as parts even if they had never actually been integrated to the whole. Keyboard buttons are typewriter-parts even if they have never composed any actual typewriter. Though perhaps more debatable, Cruse further defends that it is a matter of contingency that body parts are integrated to the body, and therefore equally qualify as autonomous parts (it may help to consider the limiting case of organs that are now grown in medical labs detached from any actual body). More striking are the non-arbitrariness of boundaries and the related functionality of parts, as opposed to pieces. Roughly, keyboard buttons have a writing function (hands have a grabbing function), which contributes to delimiting their boundaries within the whole typewriter (body respectively).

Using the piece vs. functional part terminology, one can describe quantity gradable uses of adjectives with non-maximal standards as involving the satisfaction of the property by some more or less extended pieces of an object. Partitive uses, on the other hand, involve the satisfaction of the property by some functional part. The fact that pieces are mere topological portions with no proper function within the whole makes it possible to map them onto quantity scales, measuring the extent to which they overlap with the whole. If the extent of the piece satisfying the property is above a certain standard, the property truthfully applies to the whole. On the partitive reading however, it isn't the measurable extent that is at stake, but rather, whether a given part contributes to some contextually determined function of the whole, and whether the property satisfied by the part is relevant for the part's contribution to the overall property.

One compelling argument for conceptualizing the distinction between gradable vs. partitive uses in terms of non-functional portions or pieces vs. functional parts, is that it provides an explanation of the following perceived contrasts.

(23) a. Your car is dirty [*the frame*/*the passenger cabin*/**the engine*]
b. My computer is broken [*the processor*/*??the screen*]
c. My car is the red one [*the frame*/**the car's machinery*]

One wouldn't normally consider a car as dirty just because the engine is dirty. In fact, when washing a car, one never washes the engine. Similarly, imagine a car whose machinery is entirely made of some red material or painted in red, but whose exterior is of a different color. It wouldn't normally qualify as red. Only the parts that are visible in the everyday use of the car are relevant. It also seems that a computer screen with a deep and large scratch would be described as broken, but the computer wouldn't necessarily. Adapting an example from Croft and Cruse (2004, p. 156), this is upheld by the fact that one could perfectly well say "The screen is broken, but the computer itself isn't". But it would certainly be stranger to say something like "The processor is broken, but the computer itself isn't".

### 3.3 A unified characterization of WP-metonymy: functional classification and partition

To sum up, I have shown that, even predicates that are sensitive to the *extent* to which objects manifest the designated property (i.e. involving mere portions), have a distinct use that is metonymic as far as it involves identifying which part bears the relevant property (again, to a lesser or greater extent). As already mentioned, Kennedy and McNally (2010) had pinned down this use without really working it out, and of course hadn't made the link with metonymy. Nevertheless, the connection they draw with classification presents significant and interesting analogies with existing definitions of metonymy, and in particular, Nunberg's definition of "noteworthiness" as a cognitive ground for metonymy. Here is how they characterize the classificatory (non-gradable) use of color predicates:

> What all these examples share is that having the property denoted by the color adjective is correlated with having some other property or properties which are relevant for some purpose or other. In the case of the green leaves, this might be the property of growing from a plant which will reproduce more leaves of the same color (vs. one that will not); in the case of the traffic signal, the property of indicating that one must stop or may go; and in the case of the pen, the property of producing blue ink vs. ink of some other color. These correlations constitute the basis for classifying objects […] arguably, we use the color term because color is an economical and easily observable identifier of this property (pp. 88–89).

Now compare this passage with the following one from Nunberg (1995, Sect. 3.2.) in which Nunberg presents noteworthiness as a condition for metonymic predicate transfer.

> Let me describe this condition by saying that predicate transfer is only possible when the property contributed by the new predicate is "noteworthy", which can be reckoned in either of two ways. The first type is exemplified in an utterance like "I'm parked out back," or "I'm in the phone book." The purposes of these utterances is to classify the members of one set of things (e.g., garage customers) relative to the immediate conversational purposes (of discharging the attendant's responsibilities, say), in virtue of their relation to the identifying properties of some other group of things (the cars). In cases like these we will say that a property is noteworthy if it offers a useful way of classifying its bearer relative to the immediate conversational interests.

The similarity of the two analyses in terms of classification is striking. Applied to WP-metonymy, this provides a very intuitive and general characterization: WP-metonymy involves a classificatory use of the predicate. This means that entities are contextually classified relative to some property or purpose, which in turn determines a functional partitioning of the entities, viz. a division into parts with properties that contribute to the classificatory property of their whole. Without entering into the details of functional parthood just yet, let me simply say that I here understand it along the lines of Cummins (1975), viz. a part of an object has the

function $P$ relative to the object's function $Q$, if the object's capacity to $Q$ is analyzable, in part, in terms of the part's capacity to $P$. It is this very notion of functional parthood that I will later seek to formalize within the logic of ground (Sect. 3.2). And though issues of functional integrity and partitioning are totally absent from Nunberg (1995) and Kennedy and McNally (2010), it is now clear that this is crucially what WP-metonymy really boils down to, viz. classification through part-properties.

The way Kennedy and McNally formally capture this correlation between the literally designated property (call it the "ground property") and classification is by introducing a free predicate variable $P_i$ to which the context assigns a value (the classificatory property), and a relational predicate $cor(P_i, Q)$ which expresses the correlation between $P_i$ and the ground property $Q$. The color predicate *green* for example receives the following value[7]

(24) $T(green_A) = \lambda x.P_i(x) \wedge \boldsymbol{cor}(P_i, \boldsymbol{green})$

It is worth noting that while Kennedy and McNally do acknowledge the relevance of part structure in the evaluation of what they call classificatory uses of adjectives, this is not reflected in (24). As it stands, (24) does not even require that the object denoted by the argument actually manifest the ground property for the predication to be true. What it requires is merely that the object manifests some classificatory property $P_i$ and that the latter be correlated with the ground property. An object might be truthfully described as having a certain color, under this analysis, even if it doesn't manifest the given color at all.

According to Kennedy and McNally, this is a desirable effect, though they are prepared to admit that such cases are in fact very rare. The only example they provide is that of the Catalan expression "black wine" (i.e. red wine in English). But what about the vast majority of cases where the object does in fact *partially* manifest the ground property, as in the above discussed examples? Looking precisely at these kinds of examples, the authors grant that there is some appeal to the hypothesis of an unexpressed part-referring variable in the semantic representation of the adjective. However, they finally choose to leave the connection between classification and part structure as a factor of indeterminacy that is left open by the semantic content of the adjective itself. They "assume that the correlation relation generally requires the object to also manifest the relevant color *somehow*, *somewhere*, at some point in time, but this variability, together with the underspecification of the […] correlated property, introduces a certain indeterminacy into the semantics […] that [is] consistent with the facts" (*ibid*: 94).

Such a justification for (24) is however unsatisfactory in many respects. First, the existence of marginal cases, like that of "black wine", cannot in and of itself motivate or guide a general semantic analysis of adjectives, even when restricted to colors. Saying that the indeterminacy introduced into the semantics is consistent

[7] Kennedy and McNally use the unmarked ***green*** notation for the value of the positive nominal predicate $green_N$, (which I here refer to as the ground property), while $green_A$ corresponds to the adjectival use of the predicate. ***T(a)*** stands for the translation of ***a*** into the higher order logical calculus that represents denotations.

with the facts therefore seems to be an overstatement. There is an overwhelming majority of cases that require the existence of some part manifesting the ground property expressed by the adjective, a property that affords mental access to some correlated classificatory property of the whole. As I have argued, this is indeed a much more general phenomenon than what Kennedy and McNally had thought while focusing on color predicates: a phenomenon that has been under discussion for quite some time, though not in structural formal linguistics, under the heading of WP-metonymy. It could further be argued that the Catalan use of the expression "black wine" (just like the expression "black coffee" in English, or "white wine") is idiomatic, and thus should not receive a standard compositional analysis based on the semantic content of its components.

The fact that the correlation between ground and classificatory property requires that the object "manifest the relevant [property] somehow, somewhere" therefore deserves to be further qualified. We saw, in (23), that the existence of a correlation between two properties, e.g. the presence of dirt and some more general perceptual classificatory property, wasn't enough to grasp the satisfaction conditions of the classificatory (metonymic or partitive) uses of the predicate *dirty*. The correlation between the presence of dirt and the more general visual aspect of a car, let's say, licenses the predication of dirt to the whole car only when it is manifest on the outer surface or the inside cabin, but not when it is present on the engine. And this extends to color predicates. There is for example an obvious correlation between the visual aspect of a pen and the presence of color (on the outside), and between the writing features of a pen and the presence of color (in the ink). But these correlations do not license the attribution of color to the whole pen when an invisible part other than the ink has the relevant color. (24) thus misses the connection between the fact that the ground property is correlated with some further classificatory property, and the fact that such a correlation strongly constrains the part where the ground property should be manifest.

## 4 The semantics and ontology of WP-metonymies

This last section will finally show how Kennedy and McNally's analysis of classificatory predicates can be amended to integrate both the idea of classification and that of part-reference. I will further argue that something like Nunberg's notion of noteworthiness has to be added in order to grasp the relation between classification of the whole and reference to parts satisfying the ground property. I will show that such a pragmatic notion can receive a rigorous formal treatment within the logic of "ground" (Fine 2012a), providing the right constraints for metonymic reference to parts. While not making explicit mention to functional parthood, we will see that the semantics proposed does in fact formalize this very idea. Roughly, functional parthood will be rendered in terms of the grounding relation between a part having a property and the whole having another one.

### 4.1 Predicative and argumental co-specification of WP-metonymies

So far, our discussion of proportional modifiers and comparative constructions established the contribution of the predicative content to WP-metonymic readings. We further saw that partitive uses of predicates involve the classification of the subject entity under some contextual property, the latter being correlated to the satisfaction of the base property by some relevant part. This last point is most clearly seen in light of co-predication data, where the predicates involved trigger metonymic readings that bring different parts into focus.

(25) The car is dirty [*the cabin or outer surface*] and damaged [*some mechanical part*]

But as we will now see, the content of the argument NP also contributes to such metonymic readings. Indeed the predicate *together* with the NP *co-determine* the parts that can be accessed through WP-metonymy. One way of assessing the contribution of the argument NP is to consider cases of ellipsis, where the explicit and the elided predicate admittedly share the same content but render different metonymic readings. This can only be accounted for in terms of differences in the respective arguments of the predicates.

(26) The car is clean, but not the house!

Obviously, while the outer surface of the car is relevant for interpreting the first predication, it is totally irrelevant for the interpretation of the second. We do not normally clean the exterior surface when cleaning a house.

To account for elliptic and co-predication data we can again take inspiration of the way analogous issues have been dealt with in the case of gradable adjectives. Indeed, similar ellipsis phenomena have been addressed by Kennedy (1999) who noticed that, in a sentence like (27), the comparison class or property against which the standard of tallness is fixed has to depend on the argument of the predicate:

(27) Beck is tall, and his 6-year-old daughter is too (Kennedy 1999)

(27) can only mean that Beck is tall *for an adult man*, while his daughter is tall *for a 6-year-old girl*. The comparison property cannot be the same for both. Kennedy's suggested solution was to introduce a context-dependent function **p(x)** "that takes an individual as argument and returns a comparison property based on the value of its argument" (1999, p. 130). However, while such a contextual function grasps the one sided dependence of the comparison property (transposed to our case, of the classificatory property) on the argument, it leaves out its dependence on the base property expressed by the predicate. In other words, it accounts for the ellipsis data but is unable to account for co-predications. What we need is a context-dependent function that expresses a double correlation of the kind described in (24), viz. a correlation with the base predicate, and of the kind just mentioned, viz. a correlation with the argument NP.

Something along these lines has in fact been proposed by Toledo and Sassoon (2011) who define the comparison class of gradable predicates as the value of a binary contextual function C($A$, $x$), depending both on the adjective and its argument. Applied to WP-metonymy and using the notation $P_{A,x}$ to designate the value of C for the arguments $A$ and $x$, we can thus amend (24) in a very straightforward way:

(28) $T(\mathrm{A}) = \lambda x.P_{A,x}(x)$

The metonymic uses of the predicates in (29) would then receive the analysis given in (30).

(29) a. The car is damaged
b. The car is dirty

(30) a. $P_{Damaged.the\ car}$ (the car)
b. $P_{Dirty.the\ car}$ (the car)

Rendering the following interpretation for (25)

(31) $P_{Dirty.the\ car}$ (the car) $\wedge$ $P_{Damaged.the\ car}$ (the car)

The previously problematic elliptic construction (26) is also accounted for very easily, yielding the logical form (32).

(32) $P_{Dirty.the\ car}$ (the car) $\wedge$ $P_{Dirty.the\ house}$ (the house)

There is nothing in the above that prevents the selection function C from mapping the property *dirty* onto different classificatory properties ($P_{Dirty.the\ car}$ and $P_{Dirty.the\ house}$) when applied to different objects (cars and houses respectively). One further desirable feature of this proposal is that it makes sense of a minimal constraint on the interpretation of ellipsis, identified in Kennedy and McNally (2010, p. 84). Namely, the elliptic predication of a property to two arguments of the same kind does not allow for the two occurrences of the predicate (explicit and elided) to be correlated with distinct classificatory properties. So, while the correlation between properties may vary from one context to another, or from one argument type to another, this cannot be so in a unique context with two arguments of the same type. Here is the scenario considered by the authors. Imagine two traffic lights A and B, such that A's body is painted black and A's red light is glowing, while B's body is painted red and B's green light is glowing. It is argued, correctly I think, that (33.a) does not have a reading equivalent to (33.b).

(33) a. The traffic signal A is red, and so is traffic signal B
b. Traffic signal A is red in one of the senses that the term is used, and so is traffic signal B

The analysis I proposed is perfectly coherent with the data in this respect. If the metonymic predicate and the subject are of the same type (e.g. red elliptically

predicated of two traffic signals), it is expected that the contextual function C will return the same value at both places.

One final element is now missing from our analysis. We need to connect the classificatory property contextually determined by C to the requirement that there be some "non-arbitrary" part of the subject entity that satisfies the base property. In Nunberg's words, we need to express the fact that, if parts have certain properties, this must be *noteworthy* enough to classify the whole under a correlated property, and thereby license the metonymic attribution of the part-property directly to the whole. Remember, we saw that while the shoes being untied could directly cause a person to lose them while walking, the shoes being frayed had no such direct implication. Hence the acceptability of the metonymic reading of *My shoes [the laces] are tied* as opposed to *My shoes [*the laces] are frayed*. As I understand this notion of noteworthiness, it carries an explanatory force relating the fact that a part has a property $P$ and the fact that the whole has a correlated property $Q$.

What is interesting in this "explanatory" interpretation of noteworthiness is that it brings out a parallel between the linguistic issue of WP-metonymy and some of the ontological discussions concerning the "principle of bearer-uniqueness" as it arises from theories of particularized qualities (also called tropes, moments or accidents). This parallel most clearly emerges in Schnieder (2004). Indeed, the question addressed by Schnieder is the following. Theories of particularized qualities, e.g. the redness of this particular apple or the sharpness of this particular knife, often assume that "each particularized quality has a unique object to which it belongs". This is the principle of bearer-uniqueness, which stipulates, for example, that the redness of this apple is distinct from the redness of that one, even if they are instances of the exact same hue.

Despite its ontological plausibility, this principle exposes the theories that endorse it to the following objection: How are we to make sense of the strong intuition that the redness of this apple is the same particularized quality as the redness of its skin, or that the sharpness of this knife is the same as the sharpness of its blade? To accommodate both the uniqueness principle and the seeming counter examples, Schnieder argues that the expression "$x$ is $F$" is in fact ambiguous, and proposes to distinguish strict property inherence, which satisfies the debated principle, and a partitive relation that "holds between a particularized quality and an entity *because* it holds between the quality and a *part* of the entity" (p. 218). The satisfaction conditions of partitive predications therefore introduce an explanatory relation represented by the connective *because*.

(34) $x$ $\text{is}_{\text{partitive}}$ $F$, and it is so *because* there is a (prominent) part $y$ of $x$, such that $y$ is $F$

Ambiguity therefore dissolves the apparent puzzle. "Some apples are red, we are inclined to say. They are red despite of the fact that of course most parts of them (their whole insides) are not red at all. They are red, because their skin is so. Thus, "is red" as applied to standard red apples is used partitively in the way described. The same holds for "is sharp", when we are talking about [a] butter-knife. The knife

is sharp, true, although many of its parts are not. But it is nevertheless, because its blade is so." (Schnieder 2004, p. 225).

It is clear that Schnieder's notion of partitive predication actually points to phenomena of WP-metonymy. But while his proposal to analyze such predications in terms of an explanatory relation between part-properties and whole-properties is very appealing (with some interesting logical properties to which I will turn in the last section), it is also circular and misses the classificatory roles of such predicative uses. That is to say, the part having a property (say being red) is not an explanation for the whole having the same property, but some further property that classifies the whole. In the case of the butter-knife, the sharpness of the blade does not ground the sharpness of the knife as a whole. It grounds the fact that the knife is useful to cut food. Integrating the contextual characterization of classificatory properties proposed in (28) and the *because* connective introduced by Schnieder renders the following semiformal definition of WP-metonymic predications

(35) $T(A) = \lambda \mathrm{x} \ . \ \exists y \leq x \ (P_{A,x}(x)$ ***because*** $A(y))$

As it is most often the case in the philosophical literature, I assume that *because* is factive, viz. that it can only hold between true propositions. The idea is that if P is true because of Q, then both P and Q have to be true. To illustrate (35), let us apply it to Schnieder's own example of the sharp knife. Informally, we would say that the knife is sharp if and only if it is good for cutting things (e.g. food) and it is so because the blade is sharp. Let me stress again that *because* is not treated here as a relation but as a connective with logical properties. In the next and last section, I make a brief incursion into the logic of *because* and show its relevance for the analysis of negative WP-metonymic predications.

### 4.2 WP-metonymy, grounding, and negation

Following Schnieder, I have introduced a primitive unanalyzed *because* connective in the formal characterization of WP-metonymy. This way of representing the connective in question has an obvious intuitive appeal, directly pointing to the explanatory relation at play. It is however under the name of "grounding" that it is most often discussed in the ontological literature and logically studied. I will here focus on the work of Fine (2012a, b), which provides formal rules for the pure logic of ground, its interplay with the standard logical connectives (impure logic of ground), and a fact-based semantic interpretation.[8] I will henceforth rather speak of grounding and will pursue our analysis within (a segment of) Fine's formal system. In particular, I will distinguish between full and partial ground *simpliciter*, where Fine further discusses the orthogonal dimension of strict versus weak ground, depending on whether a truth is allowed to ground itself or not. Furthermore, diverging from Fine's actual notation, which confusingly resembles the standard

[8] I refer the reader to Fine (2012a) for a presentation and discussion of the full system.

notation for the part relation, I will use the symbol '$\rhd$' to designate full ground and '$\gg$' to designate partial ground. Informally, $\varphi \rhd \psi$ will be read '$\psi$ is true in virtue of, or because of, $\varphi$'s being true'. So for instance a car engine being damaged constitutes a full ground for the car not being usable for transportation. Analogously, $\varphi \gg \psi$ will be read '$\psi$ is true partly in virtue of, or because of, $\varphi$'s being true'. For example, the well functioning of the car engine is a partial ground of the car being usable for transportation, given that other parts must also be working properly for the car to be functioning.

From these very basic examples, it seems clear that the *because* in (35) should correspond to full ground. Indeed, the sentence *My car is well functioning* cannot be interpreted metonymically to mean that the engine is functioning well (e.g. the metonymy is not felicitous if the gearbox is damaged). Whereas *My car is damaged* can be used to mean that the engine is damaged, since it fully grounds the fact that the car cannot serve for transportation. So (35) formally translates into (36)

(36) $T(Q) = \lambda \mathrm{x} \,.\, \exists y \leq x \,(Q(y) \rhd P_{Q,x}(x))$

Let us then see how the notions of full and partial ground are semantically interpreted in Kit Fine's fact-based semantics for the logic of ground. The basic idea is that, at the semantic level, each true proposition A gets interpreted in terms of its corresponding "verification set" [A], which contains all the facts that "*wholly relevantly*" verify A. Facts are here to be understood as portions of the world, and the verifiers of a truth A as the *minimal* facts that make A true in the world. For example, the fact that it is 38°C outside is a minimal verifier of the proposition that it is warm outside. However, the bigger fact that it is 38°C outside and that I am thirsty is not included in the verification set. The fact that I am thirsty is irrelevant for establishing that it is warm. This bigger fact is in this sense not a minimal verifier. More generally, "it should not be supposed that if [a fact] *f* verifies a truth A then any 'larger' fact [including *f*] must also verify A" (p. 168). Without entering into all the formal details of fact-based model structures, the satisfaction conditions of full and partial ground are respectively defined in the following way (with '.' being the fusion operation on facts):

$M \models (H_1, H_2, \ldots) \rhd G$
iff $[H_1, H_2, \ldots]$ *is a full ground for* $[G]$ *in* $M$
iff $[H_1].\ [H_2].\ \ldots \subseteq [G]$

$M \models H \gg G$
iff $[H]$ *is a partial ground for a* $[G]$ *in* $M$
iff there are some $F_1, F_2, \ldots$ such that $[H].F_1.F_2.\ \ldots \subseteq [G]$

This is read informally as: $[H_1, H_2, \ldots]$ *fully grounds* $G$ iff the facts that verify $[H_1, H_2, \ldots]$ also verify $G$; and $H$ *partially grounds* $G$ iff the facts that verify $H$, composed with facts verifying some further truths, verify $G$ ($F_1, F_2, \ldots$ being potential verification sets for some truths).

Applying these definitions to our previous car examples we obtain the following:

(37) a. $M \models$ *Damaged*(*my car*$_i$ *engine*) $\rhd \neg$*Transport* use(*my car*$_i$)
b. The facts that verify that my car engine is damaged minimally verify that my car cannot serve for transportation,
Viz. [*Damaged*(*my car*$_i$ *engine*)] $\subseteq$ [$\neg$*Transport use*(*my car*$_i$)]
c. My car [*the engine*] is damaged

(38) a. $M \models$ *Function*(*my car*$_i$ *engine*) $\gg$ *Transport* use(*my car*$_i$)
b. The facts that verify that my car engine is well functioning, together with the facts that verify that the gearbox is functioning, and so on for all the relevant mechanical parts, verify that my car can serve for transportation.
Viz. [*Function*(*my car*$_i$ *engine*)].[*Function*(*my car*$_i$ *gearbox*)].
… $\subseteq$ [*Transport use*(*my car*$_i$)]
c. My car [**the engine*] is functioning

At this point, the reader will of course have noticed that the above examples bring negation into play. They illustrate how WP-metonymic readings can be analyzed in terms of partial and full grounding, but at the same time raise the question of the way negation operates in such contexts. If it is enough that the engine of a car be malfunctioning to qualify the whole car as being damaged, the fact that the engine functions well is on the contrary only a partial requirement for the whole car to properly function. But this doesn't generalize. The situation is for example different in the case of a blue pen (qua writing). It is enough that the ink of the pen be blue, respectively not blue, for the whole pen to qualify as blue, respectively not blue. As I will now show, the logic of ground provides an explanation for such metonymic discrepancies related to negation, though not by the direct application of the rules for negation as formulated by Kit Fine (2012a, b).

Indeed, in considering the extension of the pure logic of ground to the standard logical connectives and in particular negation, Fine notes that, "it is hard to see how one might state the grounds for $\neg A$ in terms of $A$, since if $\neg A$ is a truth then $A$ is a falsehood. [He therefore suggests that] What we might do instead is to take the case in which A is logically complex and then state grounds for $\neg A$ in terms of the components of $A$" (p. 34). But this isn't enough to deal with negation of partitive predications. The problem we are facing here is this. These rules specify the grounds for a negated logically complex sentence S in terms of the grounds for S, viz. in terms of its constituents. In our case, however, we are looking to determine the grounds for a negated *logically simple* sentence (a sentence of the form $Pa$). Intuitively, if the engine's malfunctioning is a sufficient explanation of the car's malfunctioning, this should mean that the engine being in order somehow contributes to the car being in order. If the state $P$ of the engine explains the state $Q$ of the car, then the engine being $\neg P$ cannot be explanatorily irrelevant with respect to the car being $\neg Q$. But the negation principles formulated by Fine do not tell us what this explanatory correlation should be.

One way to get around this difficulty is to treat negation in terms of a big disjunction. As I will now show, though logically simple, any grounded truth is ground theoretically equivalent to a logically complex sentence, namely, the disjunction of all its possible full grounds. So we have

(39) $G \rhd S$ iff G $\rhd \vee\Gamma$ with $\Gamma = \{A_i\colon \Diamond\,(A_i \rhd S)\}$

The steps of the proof are in fact easy to follow and only require a very small fragment of the pure logic of ground and the disjunctive fragment of the impure logic of ground.

Fragment of the pure logic of ground

**SUBSUMPTION**
$$\frac{\Delta, \varphi \rhd \psi}{\varphi \gg \psi}$$

**TRANSITIVITY**
$$\frac{\varphi \gg \psi \quad \psi \gg \vartheta}{\varphi \gg \vartheta}$$

**AMALGAMATION**
$$\frac{\Delta_0 \rhd \varphi \quad \Delta_0 \rhd \varphi \ldots}{\Delta_0, \Delta_1, \ldots \rhd \varphi}$$

Disjunctive fragment of the impure logic of ground

$\vee$**I-L.** $A \rhd A \vee B$ $\quad$ $\vee$**I-R.** $B \rhd A \vee B$

$\vee\mathbf{I}^{+}$. $A,B \rhd A \vee B$ $\quad$ $\vee$**E.** $\dfrac{\Delta \rhd A \vee B}{\Delta \rhd tA; \Delta \rhd B; \Delta \rhd [A, B]}$

The left to right side of (39) is really straightforward. If $G$ is a full ground for $S$ then, by definition, $G$ is one of the $A_i$'s, and following the rules of introduction for disjunction, $G$ is a full ground for $\vee\Gamma$. The right to left side of (39) isn't very difficult to show either. By the rule of elimination of disjunction, we have that if $G$ is a full ground for $\vee\Gamma$, then $G$ is either (i) a full ground for one of the $Ai \in \Gamma$; or (ii) a full ground for a subset $\Delta \subseteq \Gamma$. If (i) is the case, then by transitivity[9] $G$ is a full ground for $S$. If (ii) is the case, then by amalgamation $\Delta$ is a full ground for $S$, and by transitivity $G$ is a full ground for $S$.

From the semantic perspective, this result states that the facts that verify $S$ verify at least some of the full grounds for $S$, and vice versa, the facts that verify some of the full grounds for $S$ verify $S$ itself. The first side might be less obvious to see, but it is in fact very natural. If a truth has one or more explanations, then the facts that make it true should also make (some of) its explanations true.

With this result we can finally account for the grounds for $\neg S$ in terms of the grounds for $S$ in full generality, covering both cases where $S$ is logically complex and simple. Namely, we can establish that

$$\text{If } G \rhd S \text{ then } \neg G \gg \neg(\vee\Gamma), \text{ and hence } \neg G \gg \neg S.$$

If $G \rhd S$ then by definition $G$ is one of the $A_i$'s, and therefore $\neg G \gg \neg(\vee\Gamma)$, viz. $\neg G$ is a partial ground for $\neg(\vee\Gamma)$. Indeed, there is a set of sentences, i.e. $\{\neg A_i : A_i \in \Gamma\}\backslash\neg G$, such that together with $\neg G$ they constitute a full ground for $\neg(\vee\Gamma)$. Furthermore, the negation of a full ground for $S$ is itself a full ground for $\neg S$, only if it is the unique possible full ground for S, that is, only if $i = 1$ and $\vee A_i = A_1$.

[9] The axiom of transitivity bears on partial ground. But given the subsumption axiom, it follows that full ground is also a transitive relation.

Going back to our metonymic examples, the logic of ground accounts very nicely and directly for the fact that, while the predication *my car doesn't work* can be interpreted metonymically as meaning e.g. *my car's engine doesn't work*—because, as we saw, the engine's not working is a full ground for the whole car not being able to perform its transportation function—on the other hand, *my car is working* does not have an equivalent metonymic reading, e.g. *my car's engine is working*. For the car to perform its transportation function, there are other important parts that have to be working properly along with the engine. And the ill functioning of each of these parts is a full ground for the car's not working, viz. corresponds to a possible metonymic interpretation of *my car doesn't work*. To put it in yet a different way, a working car is a car that is not malfunctioning, viz. is such that none of the causes for malfunctioning are realized. This can be contrasted with the example of the traffic signal (or the colored pen). According to our analysis, *the signal is red* has the metonymic interpretation *the light of the signal that is glowing is red*. This means that the red color of the glowing light is a full explanation or ground for the fact that the whole signal indicates to stop. But note that *the signal isn't red* also has the corresponding metonymic interpretation according to which the light that is glowing isn't red. This is because the red color of the glowing light is *the only* full ground for the whole signal to function as an indication to stop. Hence, its not being red is a full ground for the signal to function as an indication not to stop.

## 5 Conclusion

Let us finally take stock and see how our semantics analysis contributes to the more general understanding of WP-metonymy. How can the particular semantic proposal contribute to the more general understanding of its pragmatic and cognitive foundations? We started from the very general cognitive characterization of WP-metonymy as a mechanism exploiting conceptual connections between an entity and some of its parts; connections that are drawn against a conceptual domain or frame of knowledge. Such connections allow us to use the whole entity as a vehicle providing mental access to some salient part that is thereby indirectly brought into focus. This very general idea of conceptual frame or domain through which an entity is conceptually grasped, and against which its parts acquire a relative salience, finds a direct echo in our analysis under the guise of functional parthood. Indeed, I proposed to analyze the notion of noteworthiness or salience in terms of the explanatory relation of grounding among facts. I argued that the noteworthiness of a part having property $P$ for the purpose of classifying the whole under $Q$ had to be understood in terms of the explanatory relevance of the former with respect to the latter. In other words an object is $Q$ *in virtue of*, or *because*, one of its parts is $P$. But note that this amounts to a relation of functional parthood in the sense of Cummins (1975). According to the systemic definition of functions, an object has a function only with respect to its place in a broader system, viz. in as much as it contributes to explaining some of the capacities of the containing system. So, within the logic of ground, we actually arrived at a formal and rigorous reformulation of the intuitive

idea (see Sect. 2.3) that WP-metonymy involves reference to functional parts of integrated wholes, as opposed to mere portions arbitrarily delimited.

In this respect, the close inspection of the most recent work in structuralist semantics is shown to complement more traditional cognitive approaches to metonymy. It contributes to sharpen the description of metonymic mechanisms viewed from their purely conceptual angle. In particular, some examples that could be seen as contradicting our functionality requirement have been shown to belong to gradability rather than metonymy. This conforms with, and makes more precise, the cognitive assumption according to which the contiguity among entities driving metonymic patterns are not objective relations in the world, but are subject dependent and of a conceptual nature. WP-metonymy does not exploit purely extensional part–whole relations, but rather relies on experiential knowledge concerning the classificatory functions we attribute to objects in various settings, and the contribution of some parts to such functions.

One last general, and much more programmatic point that seems to emerge from the above discussion, concerning the overall phenomenon of metonymy, is the proximity of the relations of noteworthiness and explanation or grounding in metonymic patterns. Though this would require investigating a wider range of metonymic patterns, the recent developments in the logic of grounding, and of dependence more generally, seem to offer promising avenues of research at the intersection of formal ontology and conceptual metonymy.

**Acknowledgments** I would like to thank Friederike Moltmann, Benjamin Schnieder and Nathaniel Hansen for their suggestions on previous versions of this paper, as well as the reviewers for their insightful comments. Part of the work has been financed by the ODASP FP7-people-2013-IEF project (331196).

## References

Barcelona, A. (2008). Metonymy is not just a lexical phenomenon: On the operation of metonymy in grammar and discourse. In J. Nils-Lennart, D. Minuch, & C. Alm-Arvius (Eds.), *Selected papers from the Stockholm 2008 Metaphor Festival* (pp. 1–40). Stockholm: Stockholm UP.

Croft, W., & Cruse, D. A. (2004). *Cognitive linguistics*. Cambridge: Cambridge University Press.

Cruse, D. A. (1986). *Lexical semantics*. Cambridge: Cambridge University Press.

Cummins, R. (1975). Functional analysis. *Journal of Philosophy, 72*, 741–765.

Egg, M. (2003). Beginning novels and finishing hamburgers: Remarks on the semantics of to begin. *Journal of Semantics, 20*, 163–191.

Fine, K. (2012a). The pure logic of ground. *Review of Symbolic Logic, 25*(1), 161–187.

Fine, K. (2012b). Guide to ground. In F. Correia & B. Schnieder (Eds.), *Metaphysical grounding. Understanding the structure of reality* (pp. 37–80). Cambridge: Cambridge University Press.

Hansen, N. (2011). Color adjectives and radical contextualism. *Linguistics and Philosophy, 34*(3), 201–221.

Kennedy, C. (1999). *Projecting the adjective: The syntax and semantics of gradability and comparison*. New York: Garland Publishing Inc.

Kennedy, C., & McNally, L. (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language, 81*(2), 345–381.

Kennedy, C., & McNally, L. (2010). Color, context and compositionality. *Synthese, 174*, 79–98.

Kleiber, G. (1999). *Problèmes de sémantique. La polysémie en question*. Villeneuve: Presses Universitaires du Septentrion.

Lakoff, G., & Johnson, M. (1980). *Metaphors we live by*. Chicago: University of Chicago Press.

Langacker, R. W. (1984). Active zones. *Proceedings of the annual meeting of the Berkeley Linguistics Society, 10*, 172–188.

Langacker, R. W. (1987). Foundations of cognitive grammar. In R. W. Langacker (Ed.), *Theoretical prerequisites* (Vol. I). Stanford, CA: Stanford University Press.

Langacker, R. W. (1999). Grammar and conceptualization. *Mouton de Gruyter, 116*, 75–111.

McNally, L. (2011). The relative role of property type and scale structure in explaining the behavior of gradable adjectives. In R. Nouwen, R. van Rooij, U. Sauerland, & H. Schmitz (Eds.), *ViC 2009 (Papers from the ESSLLI 2009 workshop on vagueness in communication)* (pp. 151–168). Berlin: Springer.

Nunberg, G. (1995). Transfer of meaning. *Journal of Semantics, 12*(2), 109–132.

Nunberg, G. (2008). The pragmatics of deferred interpretation. In L. R. Horn & G. Ward (Eds.), *The handbook of pragmatics*. Malden, MA: Blackwell Handbooks in Linguistics.

Panther, K.-U., & Thornburg, L. L. (2004). The role of conceptual metonymy in meaning construction. *Metaphorik.de, 06/2004*, 91–113.

Prandi, M. (2004). *The Building Blocks of Meaning*. Amsterdam / Philadelphia: John Benjamins Publishing Company.

Pustejovsky, J. (1995). *The Generative Lexicon*. Cambridge, MA: MIT Press.

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

Rotstein, C., & Winter, Y. (2004). Total adjectives vs. partial adjectives: Scale structure and higher-order modifiers. *Natural Language Semantics, 12*, 259–288.

Schnieder, B. (2004). A note on particularized qualities and bearer-uniqueness. *Ratio, 17*, 218–228.

Seto, K. -I. (1999). Distinguishing metonymy from synecdoche. In: K. -U. Panther & G. Radden (Eds.), *Metonomy in language and thought* (pp. 91-120). Amsterdam: John Benjamins Publishing Company.

Stallard, D. (1993). Two kinds of metonymy. In *Proceedings of the 32nd annual meeting of the Association for Computational Linguistics*, ACL-93 (pp. 87–94), Ohio State University, Columbus, OH.

Sweep, J. (2009). Metonymy without a referential shift. In B. Botma & J. van Kamp (Eds.), *Linguistics in the Netherlands 2009* (pp. 103–114). Amsterdam: John Benjamins Publishing Company.

Toledo, A., & Sassoon, G. W. (2011). Absolute vs. relative adjectives—Variances within vs. between individuals. *Proceedings of SALT, 21*, 135–154.

Winston, M. E., Chaffin, R., & Herrmann, D. (1987). A taxonomy of part–whole relations. *Cognitive Science, 11*, 417–444.

Yoon, Y. (1996). Total and partial predicates and the weak–strong interpretations. *Natural Language Semantics, 4*, 217–236.