

Compound nominals, context, and compositionality

1 Introduction

English and other natural languages are learnable, productive, and systematic, and the standard explanation for these facts is their compositionality. However, it has long been noticed that there are expression classes that do not obviously conform to compositionality. Some of these are idiomatic (e.g., ‘kick the bucket’, ‘get the ball rolling’), but so long as idioms are finite, they are little threat to the general compositionality claim, since finite idioms can be acquired one at a time. Truly challenging cases are productive: indefinitely many instances of them can be generated. One such class is compound nominals (CNs).

CNs present an interesting challenge for compositional semantic theories. In English and many other languages they appear to be productively generated.¹ The lexical constituents of CNs make a regular semantic contribution to the meaning of the whole, even if that meaning is not exhaustively determined by their contribution, and this contribution is substantially greater than that present in even those idioms that inherit some semantic features from their constituents. So we can formulate and interpret compounds like ‘apple pie plate tray accident’ (= *accident involving the tray for plates used to serve pie made from apples*), or, borrowing an example from Gleitman and Gleitman (1970, 67), ‘volume feeding management success formula award’. Somewhat easier to parse examples of nominal compounds are:²

¹ For discussion of the productivity of compounding, and issues about productivity in morphology generally, see Bauer (1983), Ch. 6, and Lieber (1992).

² These examples are taken from the television show *Seinfeld*, a rich source of amusement and linguistic data.

1. astronaut pen = *pen used by astronauts*
2. bubble boy = *boy who lives in a bubble*
3. soup Nazi = *person who behaves like a Nazi about soup (?)*
4. man hands = *hands like those of a man*
5. rat hat = *hat made from rat fur*
6. earth outfit = *outfit worn by those on earth*

What these cases have in common is that it is *prima facie* difficult to see how to account for their semantics in compositional terms. This stands in sharp contrast to the semantics for many common NPs. ‘Black cat’, for example, picks out the intersection of the black things and the cats. As a predicating combination, it can be paraphrased ‘cat that is black’.³ No such account works for CNs. ‘Dog house’ does not pick out the intersection of the dogs and the houses, since that intersection is empty. Neither does ‘jewel heist’ pick out the intersection of the jewels and the heists. ‘Dog house’ means *house for dogs*, and ‘jewel heist’ means *heist of jewels*. But neither of these meanings can be derived given the meanings of the constituent expressions plus their mode of combination, at least not if we assume that the default mode of semantic combination corresponding to syntactic or morphological concatenation is set intersection, as consideration of more straightforward cases suggests.⁴

The semantic difficulty posed by CNs can be seen by considering cases such as ‘axe murderer’ and ‘child murderer’. Here the surface structure of the CNs is identical

³ But as Lahav (1989), Szabó (2001), and Reimer (2002) have argued, this may be deceptive. Semantics for adjectives is not trivial. The semantics for CNs that I will propose has much in common with attempts such as Szabó’s (2001) to outline contextually sensitive compositional truth conditions for adjectives.

⁴ Although I won’t be discussing them here, privatives (e.g., ‘fake gun’, ‘phony prince’, ‘former senator’, ‘imitation leather’, ‘ex-lineman’) are another notable case in which set intersection is not the appropriate semantic operation.

except for their first word, and the semantic contribution of their constituents is the same as that contributed to other expressions. However, the former compound is usually read as *one who murders with an axe*, while the latter means, on a readily available reading, *one who murders children*. The latter also has an agentive reading on which it means *child who murders*. But the agentive reading *axe who(?)that murders* is not usually available for the former compound; neither is the instrumental reading *one who murders with a child* readily available for the latter compound. Since the surface form of the CNs in each case is apparently identical, and the semantic value of their constituents is constant, we will have to look to factors besides surface structure to determine their meanings.

I argue that the appropriate semantic analysis of CNs treats them as expressions that are fully compositional, but context-sensitive. Specifically, I claim that their context-sensitivity can be traced to a tacit or ‘hidden’ indexical element. I develop and present this analysis by contrast with two alternative kinds of analysis according to which CNs are (1) compositional but ‘unspecific’ in their meaning, or (2) compositional but semantically many-ways ambiguous. I then defend the hidden indexical analysis by showing how it can deal with problem cases for the alternative analyses, and with some recent objections to contextualism.⁵

⁵ Some limits on the current study: I will focus here only on the most productive kinds of compounds (in English), namely nominal compounds, specifically N-N combinations. I will restrict my analysis to endocentric compounds, which tend to be concatenations of restrictive modifiers to a head, which in English is the rightmost word in the compound. Discussion of exocentric and copulative compounds will be omitted, as will neo-classical (based on Greek and Latin borrowing) compounds. Copulative compounds are those in which the constituents are semantically ‘equal’ in some way. Examples are ‘girlfriend’, ‘warrior-priestess’, and ‘north-west’. Exocentric compounds are those that lack a semantic head; examples are ‘blockhead’ and ‘butterfingers’. I suspect that exocentric compounds could be treated along the lines sketched in this paper, but such a treatment will have to wait for a later occasion. For further descriptions of compounding in English, see Adams (1973), Bauer (1983), Marchand (1969), and Selkirk (1982).

2 Theoretical perspectives on compounding

2.1 Singular unspecific meanings

2.1.1 Sainsbury's analysis

Sainsbury (2001) offers a unified compositional account of a number of alleged counterexamples to compositionality, including CNs. He suggests that we treat these expressions as having context-invariant *unspecific* meaning. ‘On the unspecific account, the determinate meaning does not supply specific information about how the satisfiers of the nouns are related’ (Sainsbury, 2001, 400).⁶ So ‘axe murderer’ would mean *axe somehow related to a murderer*, ‘child murderer’ would mean *child somehow related to a murderer*, ‘whiskey bar’ would mean *bar somehow related to whiskey*, etc. Rather than assigning several possible meanings to such expressions, the strategy is to assign a single unspecific meaning that can be *satisfied* in many ways, by many different circumstances. The semantic mode of combination for CNs is thus a univocal one. We can express such meanings as, for instance:⁷

$$7. \llbracket \text{traffic light} \rrbracket = \lambda x[\mathbf{light}'(x) \ \& \ \exists R \exists y(\mathbf{traffic}'(y) \ \& \ R(x, y))]$$

showing that the expression is satisfied by things that are lights and such that there is some relation that they bear to something that is traffic. The unspecific compounding rule would then be:⁸

⁶ Bauer (1979) states (although doesn't endorse) a similar view: ‘If only one verb is to account for the range of semantic relations that exists between the two elements of compounds it will have to be very abstract and have a vague meaning: I suggest something like “there is a connection between”’ (p. 46).

⁷ On notation: I use double brackets ($\llbracket \ \rrbracket$) to express the interpretation function. Following standard practice in truth-conditional semantics, I assume that this function maps linguistic expressions onto model-theoretic structures that are their semantic values.

⁸ Sainsbury does not offer this formalism, but it seems to capture what he has in mind. I will later question whether a single rule can adequately capture CN-internal quantifications, but I defer this subject until Section 3.1.

$$8. \llbracket N_1 N_2 \rrbracket = \lambda x[\mathbf{N}_2'(x) \ \& \ \exists R \exists y(\mathbf{N}_1'(y) \ \& \ R(x, y))]$$

That is, CNs tacitly quantify over possible relations that might hold between head and modifier.

Sainsbury motivates this unspecific meaning analysis with the claim that advocates of contextualist strategies for preserving compositionality (e.g., Stanley, 2001) have confused two separate things: (1) a single construction's having two different readings or meanings; and (2) a single construction's being satisfiable by two different 'ways in which one and the same meaning could be true' (Sainsbury, 2001, 396). Where we might be tempted, in the face of the apparent semantic diversity of CNs (and genitives, adjectives, and other cases Sainsbury treats in a similar manner), to posit either massive ambiguity or covert indexicality, the appropriate response is to realize that we have been locating the diversity at the wrong level, so to speak. There is not a diversity of meanings (or contents), but instead a diversity of entities, relations, and properties that can satisfy one and the same univocal construction class.

A problem for the unspecific meaning view of CNs is that it gives overly general, and intuitively incorrect, truth conditions for sentences containing them. For it does not constrain in any significant way the circumstances that can satisfy those expressions. Suppose 'traffic light' means *light having something to do with traffic*, as the formalization above suggests. Hence it is true of anything that satisfies such a description. The unspecific meaning account does not say what relation R the lights and traffic must stand in so that the expression applies to such objects. But consider the following possible objects and situations:

- the lights at the front of my car illuminate the road and the surrounding traffic;

- the lights in the windows of houses that can be seen while in traffic;
- the light on my car's dashboard, which I use to read the map while in traffic;
- a light that regulates its brightness by the amount of traffic passing on the nearby freeway; and
- the light by my bed, which is 500 meters from the traffic on the freeway.

All of these are lights that stand in some relation or other to traffic. The relations may be fairly strange, but relations they are, even if we would not often choose to remark on them. Are they all things to which 'traffic light' applies in one and the same sense?

According to Sainsbury's analysis, it seems so. One way to decide the issue is by examining quantified CNs. I might assert: 'There are three traffic lights in my garage.' This is partially analyzed as:

$$9. \text{three}'(\lambda x[\text{light}'(x) \ \& \ \exists R \exists y(\text{traffic}'(y) \ \& \ R(x, y))])$$

Upon examination, the garage contains one device with three colored lights for guiding traffic, one lamp that is visible to the passing traffic, and one small light by which I read maps while in traffic. You might reasonably conclude that I spoke falsely, or at least misleadingly. Those are all in some sense traffic lights; but not in the *same* sense.

Notice that the appearance of falsity here is not dispelled by adding hedges such as 'strictly speaking' or 'literally'.

10. Strictly speaking, there are three traffic lights in my garage

11. There are literally three traffic lights in my garage

Both of these sentences seem false as well. This is significant because hedges often function to direct attention towards the highest standards of applying or assessing the truth conditions of the statements they modify (e.g., 'Strictly speaking, Richard Nixon is

a Quaker’; ‘There are literally a million ants in my pants’). If effects of tacit quantification over all possible relations between head and modifier appear anywhere, they should appear in hedged constructions. But they don’t seem to.

Sainsbury might accept these consequences, given his intention to distinguish between multiple meanings and multiple satisfiers. Further evidence against the unspecific view is the fact that generic statements involving CNs appear to come out false on this analysis. Note first that not every modifier-head construction can easily be read generically. Many Adj-N combinations, in particular, seem to resist this reading: ‘green bottle’ or ‘big trout’ are not easily read as the names of kinds, since they don’t seem to denote classes that have any interesting sort of unity.⁹ But there are still CNs that can be so interpreted, e.g. ‘traffic light’, ‘jewel thief’, ‘log cabin’, and ‘fire door’. Consider the following sentences, which plausibly have generic readings:

12. Traffic lights are colored green, yellow, and red.

13. The traffic light was invented in Detroit in 1920.

14. A traffic light can be found at most intersections.

It’s a familiar point that generic claims need not be true of every instance of a kind. Some traffic lights lack a yellow light, for example, but (12) is still true. ‘Invented’ is a predicate that is best used to describe kinds, not individuals or classes, so ‘traffic light’ in (13) must be interpreted as kind-expressing. How to give an adequate semantics for genericity is a complex matter, and not one on which I take a decisive stand here. But if ‘traffic lights’ means *lights having something to do with traffic*, these generic statements are false.

⁹ That isn’t to say that it’s impossible to read these expressions generically. But it does require some heavy lifting by the context. If green bottles were customarily used to warn of a hazard, for example, we might say that ‘The green bottle saved many children’s lives’ (Krifka et al., 1995).

To see this, suppose that some generic statements involve attaching an operator to an expression that restricts its extension to just the stereotypical or prototypical members of the category that it expresses. While this is an oversimplified analysis, and one that has been criticized, it is useful for my purposes since it captures the intuition that generic statements are true only of the best examples of the category itself. The statistically typical (or otherwise ‘normal’) member of the extension of ‘traffic light’, on the unspecific analysis, does not have three colored lights, if that CN includes the examples listed above in its extension. Given the number of ways that lights can be ‘involved’ with traffic, there is little reason to expect lights unspecifically related to traffic to have anything much in common. Quantifying over all possible relations between Ns only generates an arbitrary collection of objects, not a collection that constitutes a kind or has a coherent stereotypical member. Lights that regulate traffic in the way typical traffic lights do can be found at most intersections; but lights somehow unspecifically related to traffic cannot. The generic statements plainly seem true, but they would not be on an unspecific meaning analysis of CNs.

Does this occur because compounds such as ‘traffic light’ are highly familiar and thus lexicalized? One might argue that the reason ‘traffic light’ cannot receive the eccentric interpretations given above is that it is listed with an independent meaning in the lexicon. The correctness of the unspecific account would then be shown by considering truly nonce compounds like ‘tape cat’ or ‘rosemary book’, where the only way that we can interpret them is as ‘N₂ having something to do with N₁’. Absent a context, this quantified expression is the best that we can do to guess their meaning, and this is just what the unspecific analysis predicts.

It might be true that absent a context all we can do is say that ‘rosemary book’ means *book having something to do with rosemary*. But this datum does not decide between the unspecific account and the contextualist account that I offer in Section 3.1. Direct comparison of the two positions will have to wait. However, we should note that other context-sensitive expressions also display this pattern. All that can be said about the meaning of ‘He is tall’, considered out of context, is that it means that *some (male) person is tall*. Absent a context, we cannot say what male person the pronoun picks out (since, absent a context, it doesn’t pick out *any* person in particular). Analogously, absent a context, ‘rosemary book’ does not express any relation in particular. In a context, however, it does. So the fact that non-lexicalized compounds, considered outside of a particular context of use, only permit us to gloss their meaning as *N₂ having something to do with N₁*, doesn’t by itself argue in favor of the unspecific analysis, since this fact by itself is compatible with the contextualist analysis I will present.¹⁰

2.1.2 Dowty’s analysis

Dowty (1979) offers a sophisticated, but still unspecific, interpretation rule for CNs. Changing his formalism slightly, it states:

$$15. \llbracket N_1 N_2 \rrbracket = \lambda x \exists P[P\{x\} \ \& \ \exists R[\mathbf{appropriately-classificatory}'(R) \ \& \ \forall y[P\{x\} \ \rightarrow \ [N_2'(y) \ \& \ \mathbf{typically}'(\wedge \exists z[N_1'(z) \ \& \ \sim R(y, z)])]]]]]$$

Prescinding from some of the rule’s complexity, we can see that he adopts essentially the strategy I ascribe to Sainsbury, namely taking CNs to contain tacit quantification over

¹⁰ I would add that if the quantifier in the unspecific analysis is unrestricted, some of the other objections raised in this section still apply. Not every relation listed in the possible glosses of ‘traffic light’ is one that we can commonly use ‘traffic light’ (or any other compound) to express. So it may be that even the unrestricted analysis would need to appeal to elements of context to restrict the permissible relations that can be introduced by a CN. See Section 2.1.2 for more discussion.

possible relations. However, Dowty's rule includes two significant refinements. First, following Zimmer (1971, 1972), he restricts the possible relations to those that are 'appropriately classificatory'. The effect is to restrict the domain of quantification in the hope of ruling out some of the bizarre counterexamples to Sainsbury's analysis. Second, he adds a clause indicating that the relation between N_1 and N_2 must be one that 'typically' obtains. I will say more about whether we should always include such a clause in Section 3. For now I will focus on the notion of appropriateness.

What makes a relation appropriately classificatory is not easy to specify. The intent, though, is clear: to rule out some of the bizarre relations listed above as counterexamples. For instance, one doesn't in general find readings of 'migraine medicine' that make it mean *medicine sold by a person who has a migraine*. That isn't a relation between migraines and medicine that we usually view as appropriately classificatory of a kind of medicine. But I am not certain—and I am not certain how to demonstrate—that one could *not* get such a reading. Maybe it would just require some rather elaborate back story. One point that is clear about the relation of being appropriately classificatory, though, is that its extension has to be relativized to a speaker, an audience, a conversational and cultural background, and other contextual properties. It seems fruitless to look for a set of relations that are appropriate independent of such factors.

Dowty's rule, then, depends on a restriction that is sensible only when linked to elements of context. But there is no apparent place in the rule itself for context to play a role. As it happens, though, we already have ways of incorporating dependence on context into semantics, namely via indexical expressions. In Section 3.1, I'll show how

appropriateness can play a role in interpretation of CNs without including an unrelativized predicate such as **appropriately-classificatory'** into their logical form (LF). (Another point is that we need to acknowledge more possible quantificational relations than the Dowty rule allows. Again, the approach I recommend satisfies this desideratum.)

Dowty might refine the rule to address the first point about context-sensitivity by adding an explicit context parameter to the restrictive predicate, such as **appropriately-classificatory'(R, c)**, meaning that R is appropriately classificatory in context c. This would have the effect of making context-dependence explicit. The question of how to decide between this approach and the indexical approach would then turn on more indirect considerations, such as whether the relational material can be quantified by higher expressions. The indexical approach, on one interpretation, predicts that this material should be quantificationally bindable, while Dowty's rule predicts the opposite. This is not an easy question to decide, but I will say a little more about it in Section 3.

2.1.3 A pragmatic enrichment analysis

Another alternative view of CNs can be derived from the work of advocates of so-called 'primary pragmatic processes' (p-processes). Here, in order to have a concrete and familiar exemplar, I will focus on the analysis given of these processes by François Recanati (1993, 2004). I don't know that Recanati himself would endorse this view of CNs; my interest here is only in exploring the resources that p-processes can bring to CN interpretation. On Recanati's view, p-processes are local processes that can affect the content of subsentential elements prior to the composition of a complete proposition. They contribute novel semantic content to what is said. Examples of these processes are

saturation, enrichment, loosening, and transfer. Saturation is mandatory, because open variable places such as indexicals need to be filled in for a complete proposition to be expressed; the other processes are optional. These processes all operate below conscious awareness and are non-inferential, relying instead on associations among intrasentential, within-discourse, and extralinguistic contextual elements for their operation.

The view of CN interpretation to be defended in this paper posits tacit indexicals that need contextual completion; hence it is based on primary pragmatic processes (namely saturation) in Recanati's sense. The present issue, then, must be whether it is plausible to view CN interpretation as *non*-saturation-based but still involving p-processes. The most plausible analysis along these lines would posit optional free enrichment of the head N of a compound. So, e.g., in 'cat flap', 'flap' might be enriched to mean *flap for use by (a)*, which combines with *cat* to produce the content *flap for use by a cat*. In another context, it might be enriched differently, perhaps to *flap covered by a picture of a cat*. This process, though, is local in that it functions before the semantic composition of what is said takes place, and optional in that it is always possible simply to let 'cat flap' mean what it means without enrichment: *cat flap*.

This last point highlights the way in which this is an unspecific meaning analysis: it posits that CNs have a conventionally determined meaning consisting only of the juxtaposed meanings of their constituents. This conventional meaning need not be computed, though. It can be, and typically is, elaborated by p-processes in a context-sensitive way.¹¹ Because of its context-sensitivity, this analysis does not fall prey to the objections to Sainsbury's account. In different contexts, a CN can be specialized in a way

¹¹ Because the number of possible contextually specified enrichments of a given CN is very large, possibly unbounded, it is not a multiple specific meanings view such as the one to be considered in the next section.

that restricts its truth conditions to rule out some of the implausible examples I listed. This is a virtue. However, the account is still open to objection. The primary objection is that it isn't clear that we should think that *cat flap*, taken simply as the juxtaposition of *cat* and *flap*, expresses a complete semantic constituent. If it doesn't, then no expression containing it expresses a complete proposition, which in turn means that there cannot be any such thing as the complete conventional meaning of a CN-containing expression. The argument for thinking that 'cat flap' is semantically incomplete turns again on the fact, noted in Section 2.1, that independent of any context it is unclear what things it is true of. Profligate context-sensitivity often gives rise to intuitions of incompleteness (although it is obviously not decisive, as Recanati [2004] discusses at length).

The p-process theorist cannot allow that CNs are incomplete, since incomplete expressions trigger grammatically mandated saturation. The account would then collapse into an indexical analysis. So he might reply that the conventional meaning of 'cat flap' is straightforwardly complete, and is given by the conjunction $\lambda x[\mathbf{cat}'(x) \ \& \ \mathbf{flap}'(x)]$. The morphological juxtaposition of N_1 and N_2 is set intersection; semantically, all CNs are literally predicating combinations. This meaning is probably satisfied by virtually nothing, but it is at least complete, hence open to enrichment by optional p-processes.

This analysis, however, does not seem plausible. An objection comes from considering hedged statements like those described in Section 2.1.1. Hedged statements, including those containing 'literally', do not appear to have the alleged conventional truth conditions the p-process analysis posits. Saying that something is 'literally' (or 'strictly', or 'really') a jazz cat is not to say that it is an x that is both a cat and jazz; it's more likely to say that it's a cat that enjoys (or seems to enjoy) hearing jazz. The conventional

readings of CNs are accessible only under wildly counterfactual conditions. The account leaves unexplained why the language would provide a productive device the semantics of which, taken in its conventional meaning, is almost always used to say something absurd (e.g., to describe something that is both a cat and jazz).

To sum up, then, the challenge for the p-process theorist would be to say what the complete conventional meaning of a CN is, and to show that it is compatible with judgments made under hedging operators. I haven't shown that answering this challenge is impossible, but its difficulty has to be reckoned against the background of available alternatives such as the indexicalist analysis defended here.

2.2 *Multiple specific meanings*

Judith Levi, in a widely discussed book (Levi, 1978), argues that compound nominals, as well as non-predicating adjective-noun combinations and certain nominalizations ought to receive a uniform semantic treatment.¹² On her influential analysis, CNs are ambiguous, and their ambiguity is accounted for by the fact that they are potentially derived from twelve underlying logical structures, each of which contains a different specific relationship between head and modifier. Underlying CN formation are nine Recoverably Deletable Predicates (RDPs) that are present at deeper levels of syntactic analysis, but lexically instantiated and deleted during surface transformations: CAUSE, HAVE, MAKE, USE, BE, IN, FOR, FROM, and ABOUT. There are a total of twelve possible underlying forms because CAUSE, HAVE, and MAKE can occur in

¹² To keep terminological confusion to a minimum, I will continue to use 'CN' to abbreviate 'compound nominal', although Levi uses it to abbreviate 'complex nominal'. My interest in this paper is principally in the narrower class, although the analysis might be extended to cover the broader class of complex nominals as well. I focus on Levi's account specifically because it is by far the most sophisticated theoretical analysis of CNs, but similar remarks could apply to other transformational accounts such as Lees (1960).

either active or passive forms: CAUSES/CAUSED-BY, etc. So ‘traffic light’ in its dominant reading might be derived from the underlying structure:

16. $[_{NP} [_N \text{light}]] [_S [_V \text{FOR}]] [_{NP} [_N \text{traffic}]] [_{NP} [_N \text{light}]]]$

to which we might assign the LF:¹³

17. $\lambda x[\text{light}'(x) \ \& \ \exists y(\text{traffic}'(y) \ \& \ \text{FOR}'(x, y))]$

Levi’s analysis can be seen as involving a set of separate rules, each introducing such an abstract predicate at the level of LF. The surface expression type itself would then be ambiguous among the twelve possible readings. For any token utterance, it is a matter of pragmatic decision to select the conversationally appropriate one, as in other cases of lexical, syntactic, and semantic ambiguity.

There are two major difficulties for the multiple meanings strategy.¹⁴ Both are familiar, so I review them only briefly. One is that to keep the list of possible meanings small they must be relatively unspecific. This again raises problems akin to those faced by a maximally unspecific analysis such as Sainsbury’s. The second difficulty is that the list of possible semantic material filling in a compound is extensive, and subject to surprising new additions. Hence no small list of sanctioned relations seems likely to capture all uses of CNs (Downing, 1977; Selkirk, 1982; Zimmer, 1971, 1972).

¹³ There is some question about how to map the ‘logical structure’ representation proposed by Levi onto the truth-conditional representation I am employing. For example, should we treat the initial logical structure as receiving the truth-conditional interpretation, or some intermediate structure level? At one point in the derivation of ‘traffic light’, we have a representation like $[_{NP} [_{NP} [_N \text{light}]] [_{ADJ} [_{PREP} \text{for}]] [_N \text{traffic}]]]$, which might be more appropriate for model-theoretic interpretation. Since for my purposes it isn’t important which level is best given a truth-conditional interpretation, I won’t select any one in particular. Dowty (1979), however, notes correctly that it’s a challenge in interpreting these structures that they are largely inexplicit about quantification. The solution I propose is explicit about the variety of CN-internal quantificational relations, which gives it an advantage.

¹⁴ Of course, Levi is working within a Generative Semantics framework, which assumes some form of lexical decomposition. Those who are opposed to lexical decomposition approaches will therefore not find her position congenial. To keep things manageable, I will set to one side objections that are grounded primarily in opposition to lexical decomposition as such.

Take the unspecificity objection first. Consider ‘city folk’, ‘urban riots’, ‘terrestrial life’, and ‘polar climates’ (her examples). Each of these might have as its dominant reading an expression including an underlying IN: *folk IN city*, *riots IN urban (areas)*, *life IN terra*, *climates IN polar (areas)*. But IN here must correspond to something more general than the English predicate ‘in’, since it includes relationships such as being located on or at. For example, does IN mean the same in ‘Italian book’ (=book IN Italian) as it does in ‘pocket lint’ (=lint IN pocket)? It isn’t easy to see what this common sense might be. Italian books are written in the Italian language, while pocket lint is lint that has its origins in one’s pocket (but needn’t be there anymore). The same problem can be seen with CNs involving FOR. Take the analysis of ‘traffic light’: *light FOR traffic* might be general enough to include some of the examples cited in Section 2.1.1. Compare this to ‘migraine medicine’ (=medicine FOR migraine), where the sense of FOR seems to require getting rid of the thing that it’s for. ‘Bus money’ and ‘lunch money’, meanwhile, are for acquiring, not getting rid of, bus rides and lunch. Similar questions can be raised about other RDPs.

On this interpretation, Levi’s proposal is a more constrained sort of unspecific meaning account. It posits several possible meanings for CNs that rule out some highly exotic relations between their constituents. The particular meanings she does posit, though, must be sufficiently general to capture the diversity of semantic relations that can be expressed in CNs. But insofar as these relations are highly general, the account still seems to assign CNs insufficiently discriminating truth conditions.¹⁵

¹⁵ It should be said that Levi herself notes that the RDPs are vague (pp. 82, 98). She holds that this level of vagueness ensures that we can hold onto semantic generalizations concerning CNs. I suggest that these generalizations are not as useful as they might appear.

A maximally specific meaning strategy might pursue Levi's analysis to its natural conclusion by including in the semantics a rule for every possible relation that can hold among constituents of a CN. Levi argues against this possibility on the grounds that such an approach obscures theoretically useful generalizations. I've suggested that these generalizations may not be as useful as they might appear. Levi might reply along the following lines. The grammar delivers only such rather abstract or unspecific meanings for CNs, but we may narrow these meanings in various ways in order to express propositions containing the more specific relations I have sketched so far. We plausibly require these processes in order to account for the varying senses of nouns such as 'rabbit' in 'He eats rabbit (meat)', 'He wears rabbit (fur)', and 'After the accident, the road was covered in rabbit (stuff)'.¹⁶ Combining vague linguistic meanings with processes of semantic specification might address the unspecificity worry: pragmatics lets us sharpen general or unspecific senses as the context demands.

However, this suggestion naturally leads to the second objection to the proposal: the list of RDPs does not merely include highly general members, but is also crucially incomplete. Certain attested relations do not seem to be merely specifications or refinements of existing RDPs. How might we amend the RDP list?¹⁷

Adding TO seems necessary, given 'moon mission', 'Everest expedition', 'home delivery', 'stage door', etc. Another needed construction is DURING and related temporal prepositions, as in: 'halftime show', 'afternoon tryst', 'summer soldier', 'bedtime story', and 'honeymoon argument'. MADE-OF makes numerous appearances: 'alligator shoes', 'bear rug', 'leather upholstery', 'rat hat'. Note, too, that MADE-OF

¹⁶ See Recanati (2004) for extensive discussion of specification and other free enrichment processes.

¹⁷ An extensive summary of relations discovered by other compound researchers can be found in Ryder (1994).

tends to suggest extra semantic material denoting the part of the animal being used. A ‘rat hat’ isn’t just a hat made of rat, or rat parts in general, but, typically, a hat made of rat *fur*. This material is specifically dependent on the semantic properties of the head.¹⁸

A rather productive relation that appears nowhere in Levi’s scheme is LIKE, or a similarity-denoting predicate. ‘Man hands’, ‘chicken legs’, ‘shark lawyer’, ‘crocodile grin’, are examples. These aren’t just *hands OF (a) man* or *legs OF (a) chicken*, which is probably the nearest RDP reading. They’re hands that are somehow, in some respect, like a man’s; e.g., *hands that are large like those of a man*. Levi argues that these resemblance readings are derived pragmatically, not semantically, by a rule such as: ‘If no literal meaning of a CN makes sense (or is relevant, or applicable, or internally consistent) *in a given context*, interpret a CN of the form N_1N_2 on a reading comparable to “ N_2 which is like N_1N_2 ” (p. 115, emphasis in original).¹⁹ This raises the question, though, why we cannot appeal to similar pragmatic principles whenever we have some degree of mismatch with RDPs. How far must the use of a CN in context depart from its RDP before pragmatic departures are licensed? The question is especially pressing given the vagueness of the RDPs themselves. In the account I will defend, no such decisions are necessary.

One might argue that LIKE is infrequent enough that Levi’s pragmatic rule will not be needed often. But Wisniewski and Love (1998) draw attention to the existence of

¹⁸ Levi claims that this ‘extra’ semantic material is part of our association with such compounds, but should not be included as part of their meaning. Nevertheless, it seems reliably to be part of what is said with these CNs. Levi might appeal to semantic specialization again in order to account for the addition of this content; but the issue here is whether the relation itself can be seen as the specialization of grammatically contributed material.

¹⁹ There are other, more technical difficulties for introducing resemblance readings into Levi’s system, stemming from the place of the predicate in the transformational cycle (roughly, on her view, deletion of LIKE would *mandate* deletion of another, embedded RDP, a demand that no other RDP imposes). Because these arguments depend so heavily on the details of her own system, I omit discussion of them here.

