

On “average”*

CHRISTOPHER KENNEDY AND JASON STANLEY

March 9, 2009

This paper investigates the semantics of sentences that express numerical averages, focusing initially on cases such as “The average American has 2.3 children.” Such sentences have been used by both linguists and philosophers to argue for a disjuncture between semantics and ontology. For example, Noam Chomsky and Norbert Hornstein have used them to provide evidence against the hypothesis that natural language semantics includes a reference relation holding between words and objects in the world, whereas metaphysicians such as Joseph Melia and Stephen Yablo have used them to provide evidence that apparent singular reference need not be taken as ontologically committing. We develop a fully general and independently justified compositional semantics in which such constructions are assigned truth conditions that are not ontologically problematic, and show that our analysis is superior to all extant rivals. Our analysis provides evidence that a good semantics yields a sensible ontology. It also reveals that natural language involves genuine singular term reference to numbers.

1. Introduction

According to the standard conception of natural language semantics, its purpose is to give an account of the relation between a sentence, on the one hand, and the information about the world communicated by an utterance of it, on the other. In constructing a semantic theory for a language L , we gather intuitions from native speakers of L about the truth or falsity of sentences of L with respect to various possible situations. We use their reactions to form hypotheses about the meanings of the words in L , and the ways in which, together with information from the context of use, their meanings compose to yield the truth-conditions of different utterances of sentences of L . The theories we construct from such data map words (perhaps relative to contexts) onto objects, events, situations, functions, or properties, and construct from these mappings assignments of truth-conditions to sentences of L , relative to contexts. So, the standard conception of natural language semantics coheres well both with the sort of data semanticists use in forming their theories, as well as the theories thereby constructed.

Quite obviously, natural language semantics does not tell us which kinds of things are in the world. A semantic theory for a language containing adverbs that exploits quantification over events does not tell us that there are events. Rather, what it tells us is that if there are no events, then numerous utterances of sentences containing adverbs are false. So, semantic theory can play a role, albeit a limited one, in the project of telling us what kinds of things are in the world. It can tell

*To appear in *Mind*.

us what the costs would be of denying the existence of certain kinds of entities. If, for example, Donald Davidson is correct that a sentence such as “John kissed Bill” is true only if there was an event of kissing, then it follows that if there are no events, then no utterance of this sentence could be true. But because semantic theories involve assignments of truth-conditions to sentences of a language (relative to contexts of use), they do tell us something about the costs of various metaphysical views. If a straightforward semantic theory for arithmetic is true, then a sentence such as “There is a prime number between two and five” entails the existence of numbers. As a result, a nominalist who rejects the existence of numbers is committed either to rejecting the simple semantics, or to rejecting the truth of “There is a prime number between two and five”. Finally, it is clear that linguistic semanticists are aware of these commitments, and use them in evaluating the plausibility of semantic theories.¹

The reason that natural language semantics can play a role, albeit a limited one, in the project of telling us which kinds of things are in the world is because the central notions of natural language semantics are semantic ones, namely reference and truth. The bulk of the empirical data that semantic theories are designed to capture consists of speakers’ judgements about the truth and falsity of various sentences relative to different possible situations. However, one might have thought that semantic notions such as reference and truth were too metaphysical to be scientifically respectable. There are, for example, famous sceptical arguments that seem to show that semantic notions such as reference cannot be reduced to physically acceptable ones (see Quine 1960, ch. 2; Kripke 1982, ch. 2). Many philosophers have on such a basis concluded that reference and truth are not sufficiently naturalistic notions to be explanatory planks in a scientific theory. One problem with such arguments, as (Chomsky 2000, ch. 4) has emphasised, is that no clear meaning has been given to the term “physical” in such discussions. Another problem is that it is far from clear that sceptical arguments and a priori metaphysical claims should impinge on naturalistic enquiry. As Chomsky writes (2000, p. 77) writes:

Let us also understand the term “naturalism” without metaphysical connotations: a “naturalistic approach” to the mind investigates mental aspects of the world as we do any others, seeking to construct intelligible explanatory theories, with the hope of eventual integration with the “core” natural sciences.... There are interesting questions as to how naturalistic enquiry should proceed, but they can be put aside here, unless some reason is offered to show that they have a unique relevance to this particular enquiry [the study of language and the mind]. That has not

¹Examples of this abound. To take just one, from the literature on plurals, Godehard Link rejects set-theoretic accounts of plural reference on the grounds that they involve a “mysterious transition from the concrete to the abstract” (Link 1998, p. 2). In short, Link finds it ontologically objectionable to take singular reference to be to concrete entities, but plural reference to be to abstract entities, such as sets.

been done, to my knowledge. Specifically, sceptical arguments can be dismissed in this context. We may simply adopt the standard outlook of modern science....

So a methodologically naturalist attitude towards the theory of meaning involves bracketing sceptical arguments and a priori metaphysical worries about semantic notions.

Ironically, Chomsky himself has argued for decades that a truly scientific semantics should not appeal to a relation of reference between words and things. As he writes (2000, p. 17), "In general, a word, even of the simplest kind, does not pick out an entity of the world, or of our 'belief space.' Conventional assumptions about these matters seem to me very dubious." To adopt a denotational semantics is to "go beyond the bounds of a naturalistic approach," even presumably that of Chomsky's methodological naturalist, who eschews metaphysical constraints on scientific enquiry. According to Chomsky (2000, p. 132):

As for semantics, insofar as we understand language use, the argument for a reference based semantics (apart from an internalist syntactic version) seems to me to be weak. It is possible that natural language has only syntax and pragmatics; it has a "semantics" only in the sense of "the study of how this instrument, whose formal structure and potentialities of expression are the subject of syntactic investigation, is actually put to use in a speech community," to quote the earliest formulation in generative grammar 40 years ago, influenced by Wittgenstein, Austin, and others (Chomsky 1955, 1957: 102-3).

On the face of it, Chomsky's position sits oddly with his espousal of methodological naturalism. As the quote makes clear, he has not changed his position on the naturalistic acceptability of semantic notions in fifty years, despite the extraordinary progress that has occurred in that time-period with the use of such notions. Furthermore, the philosophers whose influence he acknowledges, such as Wittgenstein, were clearly influenced in their rejection of the semantic project by the very sceptical arguments whose force Chomsky rejects in genuinely naturalistic enquiry.

Chomsky has a number of different kinds of reasons for his scepticism about the semantic project. Some of them seem to us to be inconsistent with methodological naturalism. But other reasons he has given are in the naturalist spirit; specific arguments concerning various natural language constructions, the referential analysis of which is flawed. Our purpose in this paper is to investigate in detail what we take to be his strongest such argument against the thesis that the study of natural language exploits a genuine reference relation.

One sort of reason Chomsky gives for thinking that the relation between terms and their semantic values is not the relation of reference is that he thinks that, in gathering data about meaning, we are not actually eliciting speaker intuitions about the truth and falsity of the sentences of the language. Suppose "a" is a

singular term whose semantic value is an object a , and “ P ” is a predicate term that denotes a property P of such objects. If we were eliciting speaker intuitions about truth and falsity, then speakers should tell us that a sentence of the form “ a is P ” is true only if a exists, and has the property P . But according to Chomsky, speakers often tell us that sentences of the form “ a is P ” are true, when it is obvious by a little reflection, that the singular terms in them do not refer to anything. That is, according to Chomsky, there are certain apparently singular terms that we do not think of as referring to anything, which sometimes even appear as the arguments of predicates that denote properties that are not plausibly true of anything, and yet the result is a perfectly coherent and informative sentence. If so, then we are not gathering information about the genuine truth conditions of sentences when we are eliciting speaker intuitions.

One class of example Chomsky gives concerns sentences like “London is a city in England.” According to Chomsky, native speakers will tell us that this sentence is actually true. But Chomsky thinks it is quite clear to all that the city of London, the standard semantic value of the noun phrase “London”, does not exist (Chomsky 2000, p. 37). We certainly do not accept his reasons for so thinking. Nevertheless, even if we did, this would not give us a reason to reject semantic theories that assign to the sentence “London is a city in England” truth conditions that require there to be a genuine entity in the world that is actually called “London”. It would just give us a reason to conclude that no non-negated sentences containing the word “London” are true. That is, if someone believed that London did not exist, we expect that they would also report that the sentence “London is a city in England” is not actually true. So the putative non-existence of London is not a good ground for thinking that we are not eliciting speaker intuitions about truth and falsity in gathering data about meaning.²

Chomsky does, however, have considerably more persuasive examples to provide of the phenomenon in question. The most compelling examples are sentences containing definite descriptions based on the adjective “average” such as the following:

- (1) a. The average American has 2.3 children.
- b. The average Freddie Voter belongs to 3.2 airline programs.
 (www.freddieawards.com/events/17/trivia.htm)

Sentences such as (1a-b) can certainly express truths. But presumably there is no one in the world who has 2.3 children or belongs to 3.2 airline programs. According to Chomsky, if an otherwise successful semantic theory assigns a semantic value to

²Chomsky presents a number of other sentence types as instances of this schema, as well. For example, he seems to think that an utterance of “That is a flaw in his argument” can be true, even if there are no flaws in the world. This example is not ideal, however, because many of us believe that there are flaws in the world, and even ones that can be demonstrated. We therefore consider this example as on a par with Chomsky’s scepticism about London.

“the average NP” which predicts that (1a-b) are true just in case there are entities that have 2.3 children and belong to 3.2 airline programs, then the relation between this term and its semantic value is certainly not the relation of reference.³ If it were, then the truth of (1a-b) would require the existence of objects that manifest these impossible properties, but this does not accord with our intuitions.


(2a-b) make a similar point, though in a somewhat different fashion:

- (2) a. Although we did not measure this in our study, I can say from other work the average German sees his doctor 13 times a year, the average Swiss sees his doctor 7.5 times a year and the average Briton 3.5 times. (*British National Corpus*)
- b. The average member of a national committee has served 8 years in that capacity. The average Republican member has served 9.2 years, while the average Democrat has seen 7.7 years in his office. [...] In addition, the average Republican has held public office 9.6 years and the average Democrat 10.5 years. (*Wallace S. Sayre (1932), ‘Personnel of Republican and Democratic National Committees’, The American Political Science Review, Vol. 26, No. 2, pp. 360-362.*)

While the property of seeing one’s doctor 7.5 times per year is not incoherent in the same way as the property of having 2.3 children (we might, for example, count examinations as full visits and follow-ups as half visits), the truth of (2a) does not commit us to the existence of any individuals who actually have this property. Likewise, the truth of (2b) does not commit us to the existence of actual Republicans and Democrats who have served 9.2 years or 7.7 years, respectively, as shown by the fact that the second sentence of this example could be followed by an utterance along the lines of “...though no single individual has served for those precise amounts of time.”

A natural reaction to examples like (1a) is to deny that phrases like “the average American”, in such constructions, function semantically like proper names or normal definite descriptions. This is the path we will pursue in this paper. However, one must tread carefully here. It would be question-begging to conclude from

³We say “the average NP” rather than “the average N” because we assume that the structure of the extended nominal projection is as in (i), where the determiner is of category D and projects a DP (Abney 1987).



the fact “the average American has 2.3 children” can be true even though it is obvious that there is no average American thing, without detailed argument or analysis, that “the average American” is not an expression in the same semantic category as phrases like “Bill Clinton” or “the young man”, particularly if we find that “the average American” behaves, in linguistically relevant respects, in the same way as expressions such as “the young man”, or “the red car on the corner”.

This is the core of Chomsky’s argument. He agrees that there is some relation — call it R — between expressions and their semantic values; his challenge is to the thesis that R is a relation between words and things (properties, events) in the world. According to (Chomsky 2000, p. 35) some linguistically significant properties that invoke this relation are the ones that determine whether the various forms of the third person singular pronoun in the following examples can be understood as co-valued with the definite description “the young man:” in (3a-b) this is possible; in (3c) it is not (see also Hornstein 1984, pp. 58-59).

- (3) a. The young man thinks he is a genius.
 b. His mother thinks the young man is a genius.
 c. He thinks the young man is a genius.

As Chomsky (*Ibid.*, p. 39) continues:

Explanation of the phenomena of example [(3)] ... is commonly expressed in terms of the relation R. The same theories of binding and anaphora carry over without essential change if we replace “young” in example [(3)] by “average”, “typical”, or replace “the young man” by “John Doe”, stipulated to be the average man for the purposes of a particular discourse.... In terms of the relation R, stipulated to hold between “the average man”, “John Doe”, “good health”, “flaw”, and entities drawn from D, we can account for the differential behaviour of the pronoun exactly as we would with “the young man”, “Peter”, “fly” (“there is a fly in the coffee”).... It would seem perverse to seek a relation between entities in D and things in the world — real, imagined, or whatever — at least, one of any generality.

In short, Chomsky argues that in linguistically significant respects, “the average man” is just like “the young man”. In both cases, we need to speak of the relation that regulates co-valuation of descriptions and pronouns, but only with “the young man” do we think that this relation is a word-world relation. This is a mistake, according to Chomsky: the fact that the same relation applies to “the average man” shows that it is not be a word-world relation after all.

Chomsky is not the only one who has exploited these constructions for theoretical gain. We shall use the phrase “genuine singular term” to refer to an expression whose semantic function is to designate or pick out a single object. The meaning of a genuine singular term therefore involves a commitment to the existence of

a unique individual. The category of a genuine singular term crosscuts syntactic and semantic categories — both ordinary proper names and ordinary definite descriptions, such as "the young man across the room", are genuine singular terms in our sense, despite their other differences. Philosophers have typically used phrases like "the average so-and-so" to argue that many apparent genuine singular terms are not genuine singular terms.

For example, Gilbert Ryle (1949, p. 18) argues that the person who takes the mind to be an independently existing thing distinct from operations of the body is making the same mistake as someone who takes "the Average Taxpayer" to be a genuine singular term in our sense, designating a fellow citizen. Following Melia (1995), Stephen Yablo takes phrases such as "the average so-and-so" to be examples of *representationally essential metaphors*, and argues on this basis that the phenomenon of representationally essential metaphors is pervasive. A philosopher who seeks, in Yablo's apt phrase (1998, footnote 22), "ontology-free semantic productivity" will find appeal to these constructions hard to resist. For example, philosophers of mathematics who are suspicious of the existence of mathematical entities such as numbers appeal to sentences containing definite descriptions such as "the average American" in support of their positions. According to this line of thinking, the truth of " $2+2=4$ " is consistent with the non-existence of numbers, for the same reason that the truth of "the average American has 2.3 children" is consistent with the non-existence of the average American. In the case of " $2+2 = 4$ ", number terms appear to be genuine singular terms. But sentences like "the average American has 2.3 children" demonstrate that apparent genuine singular reference is not always genuine singular reference. Just as "the average American" is not a case of genuine singular reference in "the average American has 2.3 children", so the nominalist maintains that the number terms in arithmetic are not genuine singular referring terms. On this line of thinking, the truth of sentences like "the average American has 2.3 children" shows that we can accept non-negated sentences containing apparent genuine singular reference to objects, without accepting the existence of those objects. As Joseph Melia (1995, p. 229) writes "[t]he mere fact that we quantify over or refer to a particular kind of object in our best theories does not, *pace* Quine, necessarily mean that we ought to accept such kinds of objects into our ontology. We should not always believe in the entities our best theory quantifies over."

Of course, some occurrences of number terms do not appear to be genuine singular terms at all. In constructions such as "two men are at the bar", the number term "two" is standardly taken as a quantificational determiner use of "two", rather than a singular term whose function is to designate an object (this is what Dummett (1991) calls the "adjectival" use of number terms). This raises the possibility that number terms do not occur as genuine singular terms in natural language. If number terms in natural language are uniformly quantificational determiners, then, as many philosophers have argued, we should not take their apparent

genuine singular term uses in arithmetic at face-value. Perhaps, as Harold Hodes (1984, p. 140) has urged, the “person on the street” thinks of “ $5 + 7 = 12$ ” as saying that “if one takes five objects and then another seven distinct objects then one has twelve objects in all.” If so, then number terms are not genuine singular terms, even in arithmetic.⁴

In this paper, we provide a compositional semantics for “average” as it occurs in sentences like (1a) and (1b). The significance of our analysis is as follows. First, our account of the semantics of instances of “the average NP” does not involve any dubious ontology of ghostly average things. Our work therefore provides evidence that adhering to commonsense ontology about semantic values results in better semantic theory. Secondly, we show that sentences like “the average American has 2.3 children” do not provide any support for nominalist programs in ontology. Phrases such as “the average American” do not even purport to be genuine singular terms. Third, our analysis reveals that the number terms that occur in such sentences are in fact genuine singular terms, and along the way we will raise further problems for the project of explaining away apparent genuine singular term uses of number terms. “Average” sentences are therefore important for our understanding of the relation between semantics and ontology, but for essentially the opposite reasons most philosophers have assumed.

The paper is organised as follows. We begin by reviewing and rejecting previous accounts. We then explain and motivate a semantics in which sentences like these can be true without entailing the existence of objects that correspond to the average American, the average Freddie Voter, etc., or satisfy properties like having 2.3 children. We will argue that the occurrence of “average” in these examples is one instance of a more general, cross-categorical term that is used to express numerical averages. We will show how the various instantiations of this term are related, and how sentences like those in (1a) and (1b) can be assigned fully compositional interpretations using mechanisms that are independently necessary to account for a set of other constructions.

2. Abstract and concrete uses of “average”

It is standard, in the small literature on this subject, to distinguish between two uses of “average”. On the one hand, there is the use that we see in sentences such as (1) and (2). On the other hand, there is the use in sentences such as (4a-b).

⁴See Hofweber (2005) for a very recent defence of this position. A problem for this strategy is if it turns out, as Frege suggests in *The Foundations of Arithmetic*, that in the statement of the meaning of a quantifier such as “four moons”, we have to appeal to a genuine singular term use of “four”. As he writes (1980, section 57): “For example ‘The proposition ‘Jupiter has four moons’ can be converted into ‘the number of Jupiter’s moons is four’. Here the word ‘is’ should not be taken as a mere copula as in the proposition ‘The sky is blue’. This is shown by the fact that we can say: ‘The number of Jupiter’s moons is the number four, or 4.’” In fact, in the standard representation of a generalised quantifier meaning, number terms are used as singular terms in exactly the way that Frege indicates.

- (4) a. The average New Yorker is stressed out.
 b. The average philosopher is absent minded.

As Carlson and Pelletier (2002, p. 74) write, "Here we mean something like: on some set of features that we deem relevant... An average American of this type is one who has typical properties. In this meaning, there can be many average Americans." And indeed, this use is synonymous with "typical", which can replace it without a significant change in meaning:

- (5) a. The typical New Yorker is stressed out.
 b. The typical philosopher is absent minded.

Carlson and Pelletier call the use of "average" as it occurs in sentences like (4a-b) CONCRETE, since it is true of concrete individuals, and uses of "average" as it occurs in the examples in (1) and (2), ABSTRACT. Though we disagree with Carlson and Pelletier on the semantic mechanisms underlying these two uses of "average", we will adopt their vocabulary in this paper.

A concrete use of "average" expresses different properties relative to different contexts of use. Relative to one context, it can express (say) the property of being typical in terms of wealth. Relative to another context, it can express (say) the property of being typical in terms of how one cooks one's meals. A concrete occurrence of "average" is therefore context-dependent relative to its 'dimension of typicality'. It is also gradable: one object can be more or less typical (relative to a dimension of typicality) than another. As such, it accepts various kinds of degree morphology, as shown by the naturally occurring examples in (6).⁵

- (6) a. Cruise lines have also decided to target a more average american. (to fill all those cabins in the water) So, you see a lot more people who do not go to the symphony, or theater, who work 60 hours a week, and who are generally tighter with there money.
 (www.cruisemates.com/forum/viewtopic.php?t=540988)
 b. The most average American is Bob Burns, a 53-year-old building maintenance supervisor in Windham, Conn. [...] He is the one perfectly average American.
 (www.cbsnews.com/stories/2005/12/13/eveningnews/main1124183.shtml)
 c. As an outsider Bush looks like a perfectly average American to me. He loves nascar and knows no history or geography...
 (www.lioncity.net/buddhism/lofiversion/index.php/t25936.html)

Abstract uses of "average" differ in several ways from concrete uses. In sentences of the sort we are focusing on, in which "average" is part of a definite description in subject position, the truth conditions of examples with concrete "av-

⁵(6a) and (6c) highlight the flexibility in determining what properties can be taken into consideration in particular contexts — in these cases, political blogs — when measuring typicality.

verage" can typically be paraphrased as generic statements. For example, (4a) is true just in case it generally holds that individuals who are New Yorkers and are average relative to the class of New Yorkers (a kind of relativity that may or may not be similar to the familiar comparison class relativity of gradable predicates) are also stressed out. This strategy fails utterly for abstract "average", however. It would be wrong to say that (1a) is true just in case it generally holds that if an individual is American and is average (relative to other Americans), then that individual has 2.3 children: no American has 2.3 children!

Given these considerations, it is not surprising that replacing "average" with "typical" in sentences like (1a) and (1b) results in anomaly: (7a) and (7b) are not paraphrases of (1a) and (1b), but are rather understood as generic statements, and are odd precisely because they entail the existence of individuals who have fractional children and belong to fractional airline programs.

- (7) a. #The typical American has 2.3 children.
b. #The typical Freddie Voter belongs to 3.2 airline programs.

The contrast between (8a) and (8b) makes a similar point.

- (8) a. The average French woman today is 137.6 pounds, compared to 133.6 pounds in 1970. (*www.msnbc.msn.com/id/11149568/*)
b. ??The typical French woman today is 137.6 pounds, compared to 133.6 pounds in 1970.

(8b) does not involve a commitment to impossible individuals (such as people with fractional children); it is odd because it describes a highly unlikely scenario: one in which it is generally the case that contemporary French women have a very specific weight of 137.6 pounds. The use of a specific measurement introduces a high standard of precision, but this clashes with the inherent imprecision of a generic statement. The fact that no comparable anomaly arises in (8a) suggests that the semantics of abstract "average" does not involve generic quantification over individuals, but rather some kind of reference to actual averages, i.e., to numbers or amounts which may (or may not) be precise.

Finally, unlike concrete "average", abstract "average" is not gradable, as shown by the anomaly of (9a) and by the fact that (9b) entails that there is a Republican member of Congress who has served for 9.2 years (namely, the most average one).

- (9) a. #The most average American has 2.3 children.
b. The most average Republican member of Congress has served 9.2 years.

In what follows, we will take it to be diagnostic of the distinction between concrete and abstract "average" that the former can be replaced without significant change in meaning or acceptability by the term "typical", disprefers precise measurements,

and can be modified by degree morphology, while the opposite holds of the latter.

3. Previous proposals

We have a number of distinct complaints about each previous approach to the problem of "average" sentences, which fall into three categories: analysis-specific empirical or conceptual shortcomings, non-compositionality, and lack of generality. We will discuss the first two points in the following subsections as they apply to specific accounts; the third point is that any account that is specifically designed to deal with "the average American" will fail to explain the fact that the abstract interpretation of "average" appears in a variety of different constructions. Several additional (and arguably more colloquial) uses of abstract "average" are illustrated in (10).

- (10) NYU has reported that the 53 teens have lost an average of half of their excess weight over the past year, and that's truly excellent, considering that their average weight was 297 pounds at the beginning! So, assuming that they should weigh an average of, oh, 125 pounds, they were an average of 175 pounds overweight, which means they'd lost an average of 87 pounds over the year - spectacular weight loss, IMO, even though we are talking about averages here. (From a posting on www.fitnessblogonline.com)

The examples in (11a-e) illustrate the different ways of expressing the content of the underlined part of (10) (with word order variations given in parentheses).

- (11) a. The average weight of the teens in the study was 297 lbs.
(The teens' average weight was 297 lbs.)
b. The teens in the study averaged 297 lbs in weight.
c. The teens in the study weighed an average of 297 lbs.
d. The teens in the study weighed on average 297 lbs.
(The teens in the study weighed 297 lbs on average.)
(On average, the teens in the study weighed 297 lbs.)
e. The average teen in the study weighed 297 lbs.

That these are all instances of abstract "average" is shown by the fact that the numeral in each example can be felicitously modified by "exactly" and by the fact that each of these examples could be true even if no individual student among the group of 53 weighed 297 lbs. These examples make it quite clear both that we are dealing with numerical averages here, and that such meanings are a part of everyday, colloquial English. The semantic analysis of abstract "average" that we develop in section 4 is unique in that it not only accounts for (11e), it also generalises to the entire array of "average" constructions in (11a-d). On that count alone, then, it achieves a higher level of explanatory adequacy than the alternatives we

discuss below.

3.1. *The pretence account*

Perhaps the most straightforward account of abstract “average” in definite descriptions is the pretence account. According to this analysis, there is no special abstract meaning of “average”. Though there is no individual that satisfies the description “average American”, we pretend that there is one when we utter sentences such as (1a), and we allow for the possibility that this pretend individual has (otherwise impossible) properties such as having 2.3 children. Whether it is true in the pretence that the average American has 2.3 children depends on the relevant distribution of facts in the real world. On this view, abstract readings are not due to a special semantic content for certain uses of “average”; they arise because we can pretend that certain ordinary semantic contents are true. The pretence account of “average” has come in for significant criticism in Stanley 2001; here we reiterate some of those criticisms, and add some additional ones.

According to the pretence account, the definite description “the average American” is a genuine singular term, like “the young American on the corner”, or “the nice boy next door”. It is just that when we utter “The nice boy next door is going to college”, we are not pretending that there is a nice boy next door (we are instead presupposing that there is one), whereas when we utter (1a), we are pretending there is an average American. The pretence account accords with Chomsky’s view that “the average American”, in its abstract use, is no different than other definite descriptions.

But there are a host of differences between descriptions containing abstract uses of “average” and ordinary descriptive phrases. First, abstract uses of “average” can only occur with the determiner “the”. The following sentences quite clearly involve concrete “average”, in that they have meanings that remain the same if “average” is replaced by “typical”, and they commit the person who asserts them to the existence of individuals with impossible numbers of children.

- (12) a. Every average American has 2.3 children.
 b. Most average Americans have 2.3 children.
 c. Some average American has 2.3 children.

The impossibility of quantification over pretend individuals — which is what would be required to maintain abstract interpretations in (12a-c) — is mysterious if the difference between e.g. “the average *NP*” and “the young *NP*” has nothing to do with the syntactic or semantic behaviour of these phrases, but rather only with whether or not they are being evaluated literally or under a pretence.

It is worth emphasising how serious a problem it is for the pretence account of abstract uses of “average” that it only can co-occur with the determiner “the”. For example, even the Russellian translation of “the average American” is infelici-

tous:

(13) #There is one and only one average American, and he has 2.3 children.

The occurrence of "average" in (13) does not allow an abstract use. This is deeply mysterious, if the correct account of an abstract use of "the average American" involves pretence, rather than something to do with the semantic content of "average".

Perhaps (13a-c) do not allow abstract uses of "average", because a sufficiently clear context has not been set up. Let us suppose that the following are all true:

- (14) a. The average Swede has 1.3 children.
 b. The average Norwegian has 1.2 children.
 c. The average Dane has 1.4 children.

According to the pretence account, we pretend that there is an average Swede with 1.3 children, and an average Norwegian with 1.2 children, and an average Dane with 1.4 children. If so, (15) should be both felicitous and true, but it is neither.

(15) #Every average Scandinavian has between 1 and 1.5 children.

In particular, it does not allow a reading where it simply states the conjunction of (14a-c), as it should if the pretence account were correct. It is possible to convey this information, but only if we replace "every" with "the" in (15), further illustrating the importance of the definite determiner in licensing the abstract interpretation of "average".

A further problem for the pretence account of abstract uses of "average", also emphasised in Stanley 2001, is that unlike other adjectives, one cannot place adjectives between "the" and the adverbial use of "average"

- (16) a. The old fancy car is parked outside.
 b. The fancy old car is parked outside.
 (17) a. The average conservative American has 1.2 guns.
 b. #The conservative average American has 1.2 guns.
 (18) a. The average red car gets 2.3 tickets per year.
 b. #The red average car gets 2.3 tickets per year.


If the abstract use of "average" simply had to do with a pretence governing the relevant instance of "the average NP", rather than any fact about the compositional semantics of the phrase, then it would be mysterious why one could not place adjectives between the abstract use of "average" and the determiner "the".

In sum, there are a host of distributional facts about "average" that are rendered completely mysterious by the pretence account. These distributional facts

strongly suggest that the abstract use of “average” emerges from facts about the meaning and compositional structure of the relevant constructions, rather than an attitude of pretence we have towards ordinary contents. More generally, this type of account provides absolutely no explanation of the similarity in meaning between sentences like those in (11), in which “average” appears in different syntactic contexts.

3.2. *Stanley 2001*

Stanley (2001) proposes a very different kind of theory, according to which instances of “the average NP”, when “average” has an abstract use, denote degrees on a contextually salient scale. According to Stanley, the syntactic structure of an instance of “the average NP”, when it has an abstract use, is as shown in (19), where O denotes a function from properties to measure functions (functions from objects to degrees on a contextually salient scale), whose domains are restricted to the extension of that property.



So, relative to a context in which height is salient, $\llbracket O \rrbracket^c(\llbracket American \rrbracket^c)$ yields a function from Americans to their heights. The denotation of “average” then operates on the resulting measure function in an appropriate way, returning a property that is true of degrees that correspond to the average of the values obtained by applying the measure function to its domain. Since such a property is true of only one degree (the actual average), it is just the sort of thing that can combine with the definite article.

There are several advantages to this theory. First, it exploits resources familiar from other domains, in particular the semantics of gradable expressions. Second, it explains why constructions in which adjectives occur between the definite description and the adverbial occurrence of “average” are semantically deviant: in e.g. (17) and (18), we are trying to compose properties of individuals (the denotations of “conservative” and “red”) with a property of degrees. Finally, it predicts that an abstract use of “average” is only licensed when there is a contextually salient scale. This explains why the only reading of a sentence such as (20) is one in which “average” can be paraphrased by “typical”:

(20) The average American worker votes Democratic.

There are, however, a number of significant disadvantages of Stanley’s ac-

count. First, the postulation of the *O* operator is somewhat ad hoc. Secondly, the empirical claim that "the average NP" denotes a degree is questionable. Stanley provides examples like (21) as support for this point, claiming that this sentence just expresses an ordering between degrees.

(21) The average Norwegian male is taller than the average Italian male.

But if we replace the definite noun phrases in this example with clear degree denoting expressions, as in (22), the result is odd, precisely because the adjective "tall(er)" (like "conservative" and "red") expects an individual-denoting argument.

(22) ??179.9 cm is taller than 176.9 cm.

Conversely, when we modify to the example to make the measure phrases acceptable, as in (23a), the definite descriptions become infelicitous (23b):

(23) a. 179.9 cm is a greater height than 176.9 cm.
 b. ??The average Norwegian male is a greater height than the average Italian male.

For similar reasons, Stanley's theory also has trouble with identity statements. Suppose the average height of the students in class 101 is the same as the average height of the students in class 201. Then the theory predicts that (24) is true, which is clearly not the case.

(24) The average student in class 101 is the average student in class 201.

Finally and most significantly, although Stanley's analysis provides us with a response to the referential challenge of "the average NP" (by denying that such constituents denote individuals), it does not help us with the second part of the challenge: explaining how the predicates with which these constituents compose end up having the meanings they have. Stanley expresses the important intuition that the sentence in (25a) has the truth conditions in (25b).

(25) a. The average American has 2.3 children.
 b. The average number of children that an American has is 2.3.

But merely assuming that the definite description in (25a) denotes a degree does not help us understand how the rest of the pieces of the sentence come together to give us the truth conditions paraphrased in (25b). In particular, we have no compositional account of how "have 2.3 children", which looks like a property of individuals, gets turned into a property of degrees that is true of a degree just in case it is equal to 2.3. Without this piece of the puzzle, we do not have a real explanation of how (25a) comes to have the truth conditions paraphrased in (25b).

3.3. *Carlson and Pelletier 2002*

Like Stanley, Carlson and Pelletier (2002) claim that definite descriptions like “the average American” do not involve reference to individuals in the first place; they differ in analysing them as sets of properties, rather than as degrees. An important feature of Carlson and Pelletier’s analysis is that it unifies abstract and concrete uses of “average” under a single denotation. Their analysis is thus intended to capture the natural readings of sentences like those in (26), as well as incontestably abstract uses of “average” such as the occurrences of “average” in the sentences in (1).

- (26) a. The average tiger hunts at night.
 b. The average Russian wears a hat.
 c. The average American owns a car.

As we shall see, this move is both a principal virtue and a principal vice of their theory.

In Carlson and Pelletier’s analysis, the set of properties introduced by an “average” DP is derived by subjecting the denotation of its nominal argument to what they call a partition function **part** and a special kind of averaging function **ave**, as spelled out in (27) (where f_c is a contextually restricted variant of f , the property contributed by the noun; see Carlson and Pelletier 2002, p. 92).

$$(27) \quad \llbracket \text{average} \rrbracket^c = \lambda f. \{Q \mid Q \in \mathbf{ave}(\mathbf{part}(f_c))\}$$

The two crucial features of their semantics are the functions **part** and **ave**; we discuss each in turn.

The job of the partition function is to take a NP denotation (“CN” in the quote below) and yield an object over which the averaging function can operate. In particular, **part**:

...has the dual jobs of (a) finding the appropriate partitions of the properties indicated by the CN it is operating on, and (b) for each partition thus constructed, building the set of ordered pairs made up of individual CNs and value-on-that-partition. For example, if we are computing a semantic value for “the average American”, then $\mathbf{part}(\text{american}'_c)$ will first determine what the appropriate partitions of properties for American are — for instance, it will pick out “height”, “weight”, “number of children”, “food preferences”, etc., for all those types of properties we are used to seeing in reports of the features of average Americans. (p. 91)

For each dimension that is relevant for partitioning the set corresponding to the NP denotation, the **part** function produces a set of ordered pairs $\langle x, v \rangle$ where x is an individual in the partition, and v is that individual’s “most specific value” along

the relevant dimension. So, if Kim has (exactly) 3 children, then $\mathbf{part}(american'_c)$ will include only $\langle Kim, 3 \rangle$ along the *number-of-children* partition, and no other Kim-containing ordered pairs in that partition.

The job of the averaging function **ave** is then to range over all of these sets of pairs and “do some computation ... to figure out the average value corresponding to each partition” (p. 92). Let us assume for the moment that this is simply a matter of summing up the values of the second members of each pair and dividing by the cardinality of the partition set. Assuming that “the” is semantically vacuous in such constructions, “the average American” ends up denoting a set of properties, as shown in (28).

$$(28) \quad \llbracket average \rrbracket(\llbracket American \rrbracket) = \mathbf{ave}(\mathbf{part}(american'_c)) = \\ \{ \lambda x.x \text{ has } 2.3 \text{ children, } \lambda x.x \text{ weighs } 150.25 \text{ lbs, } \lambda x.x \text{ is } 64 \text{ in tall, } \lambda x.x \text{ is con-} \\ \text{cerned about the economy, } \lambda x.x \text{ eats too much fast food, ...} \}$$

Thus, on Carlson and Pelletier’s analysis, *the average American* denotes a set of properties, and has the same semantic type as a generalised quantifier (type $\langle \langle e, t \rangle, t \rangle$). “The average American has 2.3 children” is true just in case the property of having 2.3 children is an element of the set of properties in the denotation of “the average American”. Crucially, since this constituent is not a referring term, Chomsky’s argument is defused.

There are some definite virtues to this analysis. First, the fact that the “average NP” constituent denotes a set of properties ensures that no adjectives may intervene between “average” and the (vacuous) definite determiner, assuming that other adjectives expect to combine with simple properties (or functions from properties to properties; Carlson and Pelletier 2002, p. 94). Second, in providing a uniform treatment of abstract and concrete “average”, this analysis appears to be well-equipped to handle examples like (29a-b), in which the two elements of the conjoined VP appear to require different senses of “average”: the abstract one for “have 2.3 children” and “belongs to 3.3 frequent flyer programs”, and the concrete one for “drives a domestic automobile” and “prefers to fly nonstop”.

- (29) a. The average American has 2.3 children and drives a domestic automobile.
b. The average traveller belongs to 3.3 frequent flyer programs and prefers to fly nonstop.

We defer a detailed discussion of such cases until section 5.2, but it should be clear how Carlson and Pelletier’s account can handle them. As long as we have a suitable account of property coordination, (29a), for example, will work out to be true as long as “has 2.3 children” and “drives a domestic automobile” denote properties that are in the set of properties introduced by “the average American”.

Carlson and Pelletier’s analysis also has a number of serious shortcomings, however, which ultimately weaken it as a semantic analysis of “average” construc-

tions, and thereby undermine its strength as a response to Chomsky's challenge. Some of these problems stem from the attempt to unify the concrete and abstract meanings. First, this analysis fails to explain why only the concrete version of "average" is gradable. Gradability is typically analysed in terms of semantic type, such that gradable predicates introduce degrees while non-gradable ones do not (see Kennedy 1999 for discussion). If there is no semantic difference between abstract and concrete "average", then it is unclear how their differing behaviour with respect to degree modification and other tests for gradability, discussed in section 2, can be accounted for.

Second, this analysis offers no explanation for why most of the other forms of "average" in (11) have only the abstract interpretation:

- (30) a. ??The average time of a tiger's hunting is at night.
 (cf. The typical time of a tiger's hunting is at night.)
 b. ??A hat averages a Russian's dressing style.
 (cf. A hat typifies a Russian's dressing style.)
 c. ??The average of a 50-year old man's worries is his waistline.
 (cf. The focus of a 50-year old man's worries is his waistline.)

If the abstract and concrete uses of "average" involve the same lexical item, then the same range of meanings ought to be available to all of its grammatical forms, contrary to the data in (29) (cf. "typical"/"typically"/"typify"). Instead, the fact that the abstract and concrete meanings have distinct distributions can be taken as further evidence that they are associated with distinct lexical items.⁶

Third, it remains unclear to us exactly how this analysis actually works for examples that we (and they) classify as concrete uses of "average". For example, the interpretation of a sentence such as (31) presumably involves a partitioning of the domain introduced by the nominal "Russian" that contains a set of ordered pairs of Russians and some maximally specific way of particularising that Russian's dressing habits.

- (31) The average Russian wears a hat.

But what in the world could this be, what counts as such a particularisation, and most importantly, how do we compute averages over such values?⁷

⁶The only form other than adjectival "average" with a concrete meaning is "on average" (see footnote 8 for additional discussion):

- (i) On average, tigers hunt at night.

This is unsurprising if concrete "average" is underlyingly a gradable predicate, as we argued above: most terms in this semantic category have both adjectival and adverbial forms, with the latter often (though not always) morphologically marked by a preposition or affix.

⁷These questions do not arise for the analysis that we develop in section 4, which presumes that abstract and concrete "average" correspond to different lexical items. In our analysis, the former is a

Finally, even in the case of partitionings in which it is clear what the value of the second element of each partition pair is supposed to be and how it can be used to compute an average, as in the *number-of-children* example discussed above (where this value is a number), trouble looms when we reflect on how exactly the set of properties in the denotation of “average American” is supposed to be generated. It is not enough just “to do some computation .. to figure out the average value corresponding to each partition” (Carlson and Pelletier 2002, p. 92), we also need to say how we move from the right averages to the right properties. That is, if the *number-of-children* partition contains the pairs $\{\langle Kim, 2 \rangle, \langle Hannah, 8 \rangle, \langle Bill, 3 \rangle, \dots\}$ and the average we compute based on the second elements of these pairs is 2.3, we need to make sure that the corresponding property is the property of having 2.3 children. But the average is simply a numerical value, so there is no way to ensure that it is used to generate this particular property as opposed to some other one, such as the property of owning 2.3 cars or the property of being a member of 2.3 frequent flyer programs.

Intuitively, what we want is for the dimension on which the partitioning is initially constructed to determine the property derived by the averaging component of “average”. Carlson and Pelletier do not tell us how precisely this should be accomplished, however, or even whether it can be done compositionally, a shortcoming that weakens the strength of their analysis as a response to Chomsky’s criticisms. In contrast, the analysis that we will present in section 4 avoids this shortcoming by dispensing with Carlson and Pelletier’s partitioning machinery entirely, and instead deriving the truth conditions of “average” sentences strictly in terms of the meanings of the other constituents of the sentences in which they appear.

3.4. Higginbotham 1985

We conclude with a look at Higginbotham’s (1985) discussion of “average” NPs, which does not quite qualify as an analysis, as we will show below, but which, together with ideas from Stanley’s account and our modified version of Carlson and Pelletier’s analysis, forms the starting point for our own proposals. (It also provides a new set of problems to be explained, as we will see.) Specifically, Higginbotham suggests that prenominal “average” can function as an adverb as well

functional expression that computes numerical averages. The latter, on the other hand, is a gradable predicate with a meaning similar to that of “typical”, as we showed in section 2, which can be used in generic statements about individuals; this is what we assume to be the case in examples like (31). While there are interesting and complicated questions about how exactly to characterise the lexical semantics of gradable predicates and of typicality predicates in particular, we contend that such questions are independent of the semantic analysis of abstract “average”.

Of course, since our analysis makes a semantic distinction between abstract and concrete “average”, new questions will arise as to how it deals with examples such as those in (29). We will answer these questions in section 5.2.

as an adjective, where the former corresponds to its abstract interpretation and the latter to its concrete one. On this view, (1a) should be understood as equivalent to (32).⁸

(32) Americans, on average, have 2.3 children.

Since (32) does not contain a definite description that has to be analysed as making reference to odd entities, an analysis of (1a) in terms of this structure would bypass Chomsky's metaphysical worries.

One might think that Higginbotham's proposal is ad hoc, because he is suggesting that what appear to be adjectives really are adverbs. But there is a class of adjectives that are semantically parallel to adverbs, in that they quantify over the events introduced by a VP rather than the individuals introduced by the nominal that they are in construction with, such as so-called "frequency adjectives" like "occasional". Despite appearances, "occasional" in (33a) does not modify its noun complement.

(33) a. The occasional sailor strolled by.
b. Occasionally, a sailor strolled by.

⁸Actually, it is actually not entirely clear what Higginbotham would say about (1a), though our assessment below (which Carlson and Pelletier (2002) evidently share) seems likely. The reason is that Higginbotham does not actually discuss examples involving number terms, such as (1a). Instead, he is concerned specifically with the example in (i), from Hornstein (1984).

(i) The average man is worried that his income is falling.

(i) has a concrete meaning, which can be paraphrased in the usual way (*the typical man...*). This reading is true even if no man in the upper 5% of the income bracket is concerned about his income, for example, because such men are not average in the sense relevant to the concrete reading (in this context). Higginbotham proposes the adverbial analysis to account for a second reading of (i) that is falsified in such a situation, and is in fact the only possible interpretation of (ii).

(ii) On average, men are worried about their falling incomes.

However, it is unlikely that this is a true abstract reading, since there is no actual averaging going on. Instead, this is more likely a concrete reading in which "on average" is simply modifying events or situations, as is certainly the case with *in general* in (iii).

(iii) In general, men are worried about their falling incomes.

That *in general* involves a concrete reading is shown by the entailments it generates in examples involving number terms. Both of (iva-b) have readings that require the existence of men who lost exactly 12.7% of their incomes; in (ivb), this corresponds to the concrete interpretation of "on average".

(iv) a. In general, men lost exactly 12.7% of their incomes last year.
b. On average, men lost exactly 12.7% of their incomes last year.

Only (ivb), however, also has a reading that does not require the existence of such men, but is true as long as the the average of all the losses is exactly 12.7%. This is the abstract interpretation of "on average".

(33a) does not involve a commitment to the position that a unique sailor who strolled by has the property of being occasional; instead, “occasional” functions semantically like an adverb of quantification, here, so that (33a) is semantically equivalent to (33b). Larson (1998) has shown how prenominal “occasional” can be given a compositional analysis in which it takes sentential scope as an event quantifier (see also Stump 1981); Higginbotham’s idea is that whatever mechanisms are at work in examples like (33a) should apply equally to prenominal “average”, so that as far as the semantics is concerned, we are always dealing with meanings like (32).

Carlson and Pelletier (2002, pp. 82-4) provide several criticisms of Higginbotham’s idea. The first is that the approach does not successfully generalise to uses of “average” such as those found in (34):

- (34)
- a. The average tiger hunts at night.
 - b. The average Russian man wears a hat.
 - c. The average American owns a car.
 - d. The average 50-year-old American man is worried about his waistline.

As Carlson and Pelletier rightly point out, adverbial paraphrases do not accurately reproduce the relevant readings of the sentences in (34). For example, (34a) asserts of the typical tiger that it hunts at night, and leaves open the activities of atypical tigers. (35), on the other hand, asserts of all tigers that their typical hunting is nocturnal, leaving open when the atypical hunting takes place.

- (35) Tigers, on average, hunt at night.

This criticism is not entirely persuasive, however, because all the occurrences of “average” in (34) involve concrete uses, rather than the abstract uses. For example, they can all be adequately paraphrased with the use of the term “typical”:

- (36)
- a. The typical tiger hunts at night.
 - b. The typical Russian man wears a hat.
 - c. The typical American owns a car.
 - d. The typical 50-year old American man is worried about his waistline.

Since “typical” is synonymous with the concrete use of *average*, and is not synonymous with the abstract use of *average*, the lack of a full paraphrase between adjectival and adverbial variants of the examples in (34) is not necessarily a problem for Higginbotham’s account of abstract *average*.⁹

⁹In fact, we think that (35) also involves concrete “average”, but instead of functioning as a gradable property true of appropriately typical individuals as in (34a), in (35) it is functioning as a gradable property true of appropriately typical events. While there is surely some correlation between typical individuals and typical events, the fact that we have modification of different semantic objects in these two cases is enough to explain the lack of synonymy.

Carlson and Pelletier's second criticism of Higginbotham's approach is that it does not account "for sentences with multiple 'average' NPs" (p. 82). As they point out, (37a) is not synonymous with (37b):

- (37) a. The average American knows little about the average Mexican.
 b. Americans, on average, know little about Mexicans, on average.

But this criticism also misses its mark. The use of "average" in (37a), again is not the ontologically worrisome abstract reading. Rather, it too expresses the same meaning as "typical". (37a) means:

- (38) The typical American knows little about the typical Mexican.

Carlson and Pelletier also raise some worries about the syntactic processes involved in Higginbotham's account. As they write:

There are a couple of syntactic manipulations involved here about which Higginbotham does not give details: (a) the definite singular NP has become a plural NP, (b) it is not specified whether the adverbial phrase is to be attached to the VP or to the S (or elsewhere) and (c) no information is given concerning what variables (if any) the adverbial may bind. (ibid.)

This is a more serious objection, and we agree that any proposal that attempts to explain the abstract interpretation of prenominal "average" in terms of adverbial constructions like (32) must be accompanied by a compositional analysis that relies on a minimum number of construction-specific assumptions (ideally, zero). It should also explain the relation between abstract "average" and the definite determiner ("occasional" is not so picky, occurring both with "the" and *a*), and the fact that no additional adjectives can appear to the left of "average".

Most importantly, such an analysis must explain why (32) is not itself subject to a variant of Chomsky's challenge, one that focuses on the compositional contributions of the predicate term to the truth conditions rather than the referential properties of definite descriptions. The problem becomes clear when we take a closer look at the verb phrase in this example, "have 2.3 children". Outside of "average" sentences, this constituent clearly introduces the property of having 2.3 children, as shown by the following examples, all of which have truth conditions that make reference to such a property.

- (39) a. Kim has 2.3 children.
 b. Every/no/that American has 2.3 children.
 c. Most/many/few Americans have 2.3 children.
 d. Americans always/usually/rarely/never have 2.3 children.

This is not surprising: if the meaning of “have 2.3 children” is compositionally derived from the meanings of its parts, and if noun phrases of the form “ n children” introduce existence entailments about quantities of children of size n , then we expect “have 2.3 children” to denote the property that it does in fact denote in these examples.¹⁰

The puzzle is why it evidently *fails* to denote such a property in “average” sentences. This question applies equally to sentences like (1a), in which “average” is a prenominal modifier, and to Higginbotham’s (32), in which it is part of an adverbial expression. In short, merely assuming that prenominal “average” can be interpreted adverbially tells us nothing about the semantic contribution of “have 2.3 children”, and in particular, it does not explain why this constituent fails to introduce the property of having 2.3 children, as it does in the various examples in (39). Thus while we accept Higginbotham’s important insight that there is a relation between prenominal and adverbial “average”, in the absence of an explicit compositional semantics for constructions like (1a) and (32), this remains merely an observation of a correlation rather than an analysis, and certainly not an explanation of anything.

4. An “average” semantics

4.1. Preliminaries

Let us begin our own analysis by returning to our earlier observation that the abstract interpretation of “average” can be found not just in definite descriptions or adverbs, but in fact appears in a number of distinct construction types, which we repeat below in (40).

- (40) a. The average weight of the teens in the study was 297 lbs.
(The teens’ average weight was 297 lbs.)
b. The teens in the study averaged 297 lbs in weight.
c. The teens in the study weighed an average of 297 lbs.
d. The teens in the study weighed on average 297 lbs.
(The teens in the study weighed 297 lbs on average.)
(On average, the teens in the study weighed 297 lbs.)

¹⁰There is a more general problem here about the status of fractional number terms such as “2.3”, “two and a third”, etc. In standard generalised quantifier theory, number terms denote relations between sets: “two As are B” is true just in case the cardinality of the intersection of the As with the Bs is 2. In a sentence like “There are 2.5 oranges on the table”, however, “2.5” cannot plausibly be analysed in this way, since cardinalities correspond only to whole numbers. Alternative formulations of number terms as existential quantifiers run into a different problem: their restrictions do not appear to provide the right set of objects for them to quantify over. As Nathan Salmon writes (1997, p. 4): “The orange-half on the table is not an orange, and hence is not in the class of oranges on the table.” It is therefore unclear how “2.5 oranges” ends up having the meaning it does. We return to this point in note 16 below.

- e. The average teen in the study weighed 297 lbs.

The semantic generalisation that we can draw from these examples is that, independent of its grammatical category and syntactic position, abstract “average” requires three semantic arguments: a measure function (here based on the meaning of “weight”/“weighed”), a domain (provided by the nominal “the teens in the study”), and an average, the result of dividing the sum of the values derived by applying the measure function to each object in the domain by the set’s cardinality (“297 lbs.”). In other words, all of the examples in (40a-e) convey the information in (41) (possibly along with other, construction-specific aspects of meaning that we abstract away from here), where **weight** is a function from objects to their weights, T is the set of teens in this particular study, and *297 lbs* is a degree of weight.

$$(41) \quad \frac{\sum_{x \in T} \mathbf{weight}(x)}{|T|} = 297 \text{ lbs}$$

In prose: the sum of the weights we get by applying the **weight** function to all of the objects in T , divided by the number of elements in T is *297 lbs*.

Our challenge is to show that we can get from each of the different syntactic forms in (40) to truth conditions equivalent to (41) — and in particular, that we can get from (40d-e) to (41) — without doing violence to generally accepted assumptions about the nature of semantic composition. Our strategy will be to assume that one of the forms in (40) is basic, provide it with a denotation that derives the truth conditions in (41), and show how this basic denotation, together independently justified assumptions about possible Logical Forms and compositional operations on them, can be used to derive appropriate truth conditions for all forms of abstract “average”.

Before proceeding, we want to make explicit some assumptions about the syntax-semantics interface that we will be working with, as well as some notational conventions that we will adopt for expository purposes. Regarding the former, we will adopt the general framework of semantic interpretation outlined in Heim and Kratzer 1998 (and assumed in some form or other by a wide range of work in generative grammar), in which semantic interpretations are based on a syntactic level of Logical Form, which may differ in configuration in certain well-defined ways from the syntactic level that feeds pronunciation. In particular, Logical Forms may differ from surface representations in the positions of argument terms and quantifiers, in ways that we will make explicit in section 4.2.2, with the result that scopal relations between different terms are transparently represented in the syntax. Nothing in our analysis hinges on this particular assumption about the syntax-semantics interface, and all of our proposals could be recast in terms of a framework that directly interprets surface representations (we discuss specific alternatives in section 4.2.3); our decision to adopt a framework that posits a level of Logical Form

is based primarily on the assumption that it will provide the most transparent and accessible means of explaining our proposals.

Logical Forms are directly mapped to truth conditions by the composition rules of the language, which we assume to be sensitive to the semantic types of the denotations of (atomic and complex) syntactic objects, and to include at least rules of function application and function composition, defined in the usual ways. We assume at least three atomic semantic types: individuals (type e), truth values (type t), and degrees (type d); the latter corresponds to measures of quantity, degree or number.¹¹ In specifying denotations, we will follow standard convention in using angle-bracket notation to indicate complex types, so that an expression of type $\langle a, b \rangle$ denotes a function from things of type a to things of type b . For simplicity, we will omit type specifications from statements of denotations when possible, either specifying in the text what type a particular argument is, or following a general convention of using variables from the set $\{x, y, z\}$ for things of type e and variables from the set $\{n, m, d\}$ for things of type d . Finally, we will use predicate logic as a metalanguage for representing truth conditions (rather than English as in Heim and Kratzer 1998), defining any new symbols that we introduce (such as **weight** above).

We make two further assumptions in order to maximise the clarity of the exposition. The first is a simplifying assumption: we will treat the domain argument of “average” in all cases as denoting a set, ignoring the fact that the linguistic expression that provides this argument may take different forms (a bare plural, a definite plural, a conjunction structure, a bare noun, etc.) and also ignoring the potentially important contribution of verbal particles like “each”, “per year” and so forth. A fully comprehensive linguistic analysis will likely need to assume that the denotations of the different forms of “average” specify mappings from different kinds of expressions to sets (or to more structured objects, such as pluralities), but since it is straightforward to define such mappings and since our more general proposals are consistent with different analytical options, we will talk in terms of sets in what follows.

Second, although we assume that one of the crucial semantic components of averaging is a measure function (type $\langle e, d \rangle$), as described above, in all of the constructions we examine the actual linguistic terms that provide this component denote degree relations, either type $\langle e, \langle d, t \rangle \rangle$ (such as the noun “weight”) or type $\langle d, \langle e, t \rangle \rangle$ (such as the verb “weigh”):

- (42) a. $\llbracket \text{weight}_N \rrbracket = \lambda x \lambda d. \mathbf{weight}(x) = d$
 b. $\llbracket \text{weigh}_V \rrbracket = \lambda d \lambda x. \mathbf{weight}(x) = d$

¹¹A complete interpretive framework will also need to include at least the type of events, and possibly times and possible worlds as well, but since none of the constructions we discuss require reference to such types we will set them aside in what follows.

Degree relations (either lexical or derived) can easily be converted into measure functions, however, so we will use the following abbreviatory conventions in our semantic representations to simplify the notation:

- (43) a. If $f \in D_{\langle e, \langle d, t \rangle \rangle}$, then $f_{meas} = \lambda x.max\{d \mid f(x)(d)\}$
 b. If $f \in D_{\langle d, \langle e, t \rangle \rangle}$, then $f_{meas} = \lambda x.max\{d \mid f(d)(x)\}$


See e.g., Cresswell 1977; Heim 1985; Klein 1991; Carpenter 1997 and Kennedy to appear for the use of such conversions in the semantic analysis of comparatives.

4.2. Analysis

4.2.1. Basic cases

We begin with the assumption that the form of “average” in (40a), which combines directly with a measure noun, reflects the basic meaning of the term. This assumption, while arbitrary, is based on an informal search of the British National Corpus for collocations of “the average”, “an average” and “on average”, which suggests that the measure noun-modifying form in (40a) is by far the most frequent. Nothing hinges on this particular assumption, however, and our analysis is completely consistent with another (or a more abstract, category-neutral) form being basic.

The structure of a noun phrase containing this form of “average” is as shown in (44) for “the average weight of the teens”.



Assuming that the noun “weight” denotes the degree relation in (42a) and that the plural DP *the teens* introduces a set as discussed above (which we will abbreviate throughout as *teens'*), this structure indicates that the core meaning of *average* is the function **average** in (45) (where f_{meas} is the measure function based on f , as defined in (43)).¹²

¹²The assumption that “average” forms a constituent with the measure noun independent of the PP in (44) may appear unjustified, given that adjectives typically modify full NPs (nouns and their arguments; in this case *weight of the teens*) rather than nouns. There is some reason to believe that this structure is correct, however, and may even be a case of compounding rather than adjectival modification. First, unlike the form of “average” in “the average American”, this form may and typically must be rightmost:

- (i) a. The unexpected average weight of the teens

$$(45) \quad \mathbf{average} = \lambda f \lambda S \lambda d. \frac{\sum_{x \in S} f_{meas}(x)}{|S|} = d$$

Composition of the nominal portion of (44) gives us (46a), which spells out as the property of degrees in (46b) after lexical insertion and λ -conversion.¹³

$$(46) \quad \begin{array}{l} \text{a. } \mathbf{average}(\llbracket \mathbf{weight} \rrbracket)(\llbracket \mathbf{the teens} \rrbracket) \\ \text{b. } \lambda d \left[\frac{\sum_{x \in teens'} \mathbf{weight}(x)}{|teens'|} = d \right] \end{array}$$

This property is true of a degree if it equals the average weight of the teens, and further composition with the definite article will result in a definite description that picks out the unique degree that satisfies this property. The net result is that (40a) is predicted to be true just in case the average weight of the teens equals the degree denoted by *297 lbs.*, which is exactly what we want.

The verbal form of “average” in (40b) can be analysed in much the same terms, the only difference being the order of argument composition. Taking the surface syntax as a guide, the verbal form differs from the basic form in selecting the degree argument first, then the measure argument (which can also be implicit if the context is rich enough, as in as in “The teens averaged 297 lbs.”), and finally the domain argument, as shown in (47).

b. ??the average unexpected weight of the teens


Second, in some languages, this form quite clearly involves compounding. This is illustrated by the Norwegian data in (ii),

- (ii) a. Den norske gjennomsnittslønnen er 500,000 kroner.
 the Norwegian average.salary.DEF is 500,000 kroner
 b. Den naaverende gjennomsnittsalderen paa studentene er 24 aar.
 the current average.age.DEF on students.DEF is 24 years

¹³Our use of **weight** in (46) to represent the result of applying the conversion operation to the degree relation denoted by “weight” reflects the fact that all of (ia-c) are equivalent.

- (i) a. $[\lambda x \lambda d. \mathbf{weight}(x) = d]_{meas}$
 b. $\lambda z. \mathit{max}\{d \mid \mathbf{weight}(z) = d\}$
 c. **weight**

This equivalence holds for any lexical degree relation. In the case of the derived degree relations we will introduce shortly, we will spell out the result of f_{meas} conversion using λ -terms like (ib).



An appropriate meaning for verbal “average” can be defined in terms of **average** as in (48).


(48) $[[_{\text{V}} \text{average}]] = \lambda d \lambda f \lambda S. \mathbf{average}(f)(S)(d)$

Composition of the various constituents in (47) gives (49a), which maps onto (49b) after lexical insertion and λ -conversion, which is in turn equivalent to (49c).

- (49)
- a. $[[\text{average}_{\text{V}}]]([\text{297 lbs}])([\text{weight}])([\text{the teens}])$
 - b. $\mathbf{average}([\text{weight}])([\text{the teens}])([\text{297 lbs}])$
 - c. $\frac{\sum_{x \in \text{teens}'} \mathbf{weight}(x)}{|\text{teens}'|} = 297 \text{ lbs}$

4.2.2. *Derived degree relations*

We now turn to the nominal form of “average” in (40c), which can be analysed semantically in exactly the same way as verbal “average”, even though its syntactic properties are different. Assuming the structure of (40c) is as shown in (50), the denotation we want is the one in (51).



(51) $[[_{\text{N}} \text{average}]] = \lambda d \lambda f \lambda S. \mathbf{average}(f)(S)(d)$

Composition is then straightforward: “average” combines first with the measure phrase “297 lbs”, then with the measure verb “weigh”, and finally with the subject, resulting in (52a). (We assume for simplicity here that “an” and “of” are semanti-

cally vacuous.)

- (52) a. $\llbracket \text{average}_N \rrbracket(\llbracket 297 \text{ lbs} \rrbracket)(\llbracket \text{weigh} \rrbracket)(\llbracket \text{the teens} \rrbracket)$
 b. $\mathbf{average}(\llbracket \text{weigh} \rrbracket)(\llbracket \text{the teens} \rrbracket)(\llbracket 297 \text{ lbs} \rrbracket)$
 c. $\frac{\sum_{y \in \text{teens}'} \mathbf{weight}(y)}{|\text{teens}'|} = 297 \text{ lbs}$

Given the denotation in (51), (52a) is equivalent to (52b), which spells out as (52c) after lexical insertion and λ -conversion. (40c) is thus correctly predicted to be truth-conditionally equivalent to (40a) and (40b).

In examples like (40c), the degree relation that **average** converts into a measure function is lexical, provided directly by the verb “weigh”. However, in many other constructions involving “an average of”, the degree relation is not lexical but instead must be derived in the syntax. (53) is an example of such a construction.

- (53) The teens ate an average of 17.5 hamburgers each.

The degree relation we want in order to get the right truth conditions for this example is the relation between quantities n and individuals x that is true just in case the number of hamburgers that x ate equals n , which we represent informally in (54).

- (54) $\lambda n \lambda x. x \text{ ate } n \text{ hamburgers}$

If (54) is supplied as the second argument of nominal “average”, and the plural subject as the third argument, the truth conditions we will ultimately end up with are those represented in (55).

- (55) $\frac{\sum_{y \in \text{teens}'} \lambda x. \max\{n \mid x \text{ ate } n \text{ hamburgers}\}(y)}{|\text{teens}'|} = 17.5$

Given that (55) correctly characterises the meaning of (53), the question is how we get from the verb phrase “eat an average of 17.5 hamburgers” to the degree relation in (54). In fact, this is exactly the question that we kept running up against in our discussion of previous approaches to “the average NP” in section 3. Recall from that discussion that the problem we confronted was how to avoid interpreting a verb phrase like “have 2.3 children” in a way that did not entail of any entity that it has 2.3 children, an entailment made by any approach that assumes that numerals are quantificational determiners that combine with nominals to yield generalised quantifiers. In order to derive a degree relation like (54), and avoid these problems, we need to give up this assumption. Instead, we need to recognise that number terms lead a dual life. In addition to their use as quantificational determiners and corresponding relational meanings, they also occur as

singular terms. As such, they can saturate a degree/quantity position inside the noun phrase, and take scope independently of the rest of the noun phrase in which they occur.

That number terms can occur both as quantificational determiners and as singular terms is an already familiar point. As Gottlob Frege (1980, section 57) writes in a famous passage:

I have already drawn attention above to the fact that we speak of “the number 1”, where the definite article serves to class it as an object. In arithmetic this self-subsistence comes out at every turn, as for example in the identity $1 + 1 = 2$. Now our concern here is to arrive at a concept of number usable for the purposes of science; we should not, therefore, be deterred by the fact that in the language of everyday life number appears also in attributive constructions. That can always be got round. For example, the proposition “Jupiter has four moons” can be converted into “the number of Jupiter’s moons is four”.

Similarly, Michael Dummett (1991, p. 99) notes:

Number-words occur in two forms: as adjectives, as in ascriptions of number, and as nouns, as in most number-theoretic propositions.

As both Frege and Dummett emphasise, the paradigmatic use of number terms in arithmetical contexts is as singular terms, rather than as quantificational determiners. Since abstract uses of *average* involve arithmetical contexts, it is no surprise that number terms occurring in sentences containing them are the former rather than the latter.¹⁴

Philosophers of mathematics since Frege have been aware that number terms often behave both syntactically and semantically as singular terms. However, in sentences containing abstract uses of “average” (such as (1a) and (1b)) the number term superficially appears in quantificational determiner position. But here surface syntax is not always a good guide to semantic type. Consider, for example, sentences like (56a-b).

- (56) a. John ran 2.3 miles.
b. Bill weighs 70 kilograms.

¹⁴Indeed, much philosophy of mathematics consists of attempts to reduce one of these uses to the other. Platonists such as Frege consider the use of number terms as singular terms as central and the quantificational determiner use to be misleading. In contrast, those hostile to arithmetical platonism tend to view the use of number terms as quantificational determiners as central, and the use of number terms as singular terms as peripheral. Such authors think of generalised quantifiers such as *two men* as having, for the semantic value, a set of properties; the number term *two* contributes, not a number, but a function from properties to the characteristic function of such a set (see e.g., Barwise and Cooper 1981; Keenan and Stavi 1986).

Here, the terms “2.3” and “70”, like the uses of “2.3” and “3.2” in (1a) and (1b), superficially appear to be in quantificational determiner position. However, semantically they are clearly not quantificational determiners. The term “2.3” in (56a) does not express a relation between the set of miles (whatever that would mean) and the set of things (distances?) that John ran. Instead, it provides the value of a certain kind of measurement, namely a measurement of the distance that John ran (in miles). Our discussion in what follows is neutral between the Fregean position that the singular term use of number terms is fundamental and the alternative position that there are genuinely two distinct uses of number-terms, as quantificational determiners (the “adjectival use”) and singular terms. What is crucial for our analysis that we need to recognise a genuine singular term use of number terms, not that it is the only such use.

The hypothesis that number terms (can) genuinely denote numbers even in constructions in which they superficially appear to function as quantificational determiners is developed in detail in the work of Manfred Krifka (1989; 1990) (see also Cresswell 1977), and used to account for a range of facts involving aspectual composition and the relation between nominal and verbal reference. In Krifka’s analysis, plural count nouns do not denote simple properties of (plural) individuals, but rather relations between individuals and numbers (or degrees/amounts; as noted above, we treat all of these things as instances of the same semantic type) of the sort shown in (57) for the noun “hamburgers”. (Here the variable x ranges over plural rather than atomic individuals; see Link 1983.)¹⁵

¹⁵A slightly different approach to number terms, but one that is conceptually related to the version we adopt in this paper, is advocated by Martin Hackl (2001). Hackl provides arguments from the syntax and semantics of comparative quantifiers like “more than three” that the determiner “many” introduces a number argument, as specified in the denotation in (i).

$$(i) \quad \llbracket \text{many} \rrbracket = \lambda n \lambda P \lambda Q. \exists x [|x| = n \ \& \ P(x) \ \& \ Q(x)]$$

Composition of “many” with a number term returns an expression with the semantic type of a quantificational determiner (type $\langle\langle e, t \rangle, \langle\langle e, t \rangle, t \rangle\rangle$) and a denotation that is identical that of the number term on the corresponding generalised quantifier analysis.

$$(ii) \quad \llbracket \text{many} \rrbracket (\llbracket \text{three} \rrbracket) = \lambda P \lambda Q. \exists x [|x| = 3 \ \& \ P(x) \ \& \ Q(x)]$$

If the surface string like “three hamburgers” is (or at least can be) of the form shown in (iii), where “many” is deleted from the surface representation, then we allow for the possibility that the number term can take scope independently of the rest of the phrase, leaving behind a variable over degrees/amounts.

$$(iii) \quad [_{\text{DP}} \text{ [three } \cancel{\text{many}} \text{] hamburgers}]$$

Since nothing in our proposals hinges on a choice between a Krifka-style or a Hackl-style analysis, we adopt the former because it allows us to keep the syntactic representations as simple as possible. It is interesting to note, however, that Hackl’s system provides a simple account of the (apparent) dual life of number terms as determiners and singular terms: the bit of phonology pronounced /θri/ is the pronunciation both of the number term “three” (which denotes a quantity) and the quantificational

$$(57) \quad \llbracket \text{hamburgers} \rrbracket = \lambda n \lambda x. \text{hamburgers}'(x) \ \& \ |x| = n$$

Composition with a number term returns a property that is true of pluralities of hamburgers whose cardinality is equal to the number denoted by the number term. The compositional analysis of “three hamburgers”, for example, is shown in (58).

$$(58) \quad \begin{aligned} \llbracket \text{three hamburgers} \rrbracket &= \\ \llbracket \text{hamburgers} \rrbracket(\llbracket \text{three} \rrbracket) &= \\ \lambda x. \text{hamburgers}'(x) \ \& \ |x| = 3 \end{aligned}$$

This property may then compose with a verb meaning, saturating an open argument, and the variable corresponding to this argument will ultimately be bound by a default existential quantifier, deriving truth conditions that are equivalent to what we get on a standard generalised quantifier semantics.¹⁶

What is important for our purposes is that on this analysis, a number term or other amount term saturates the degree argument of a plural noun, and so can in principle take scope independently of the rest of the noun phrase, leaving a degree variable in its place. This provides us with a straightforward means of deriving the degree relation in (54) and providing a compositional analysis of (53). The analysis runs as follows.

First, we assume that “an average of 17.5” is a constituent in this example that occupies the same syntactic position as a simple number term. As such, it may

determiner “three ~~many~~” (which denotes a relation between sets). In other words, in Hackl’s system, number terms are not ambiguous — they are always singular terms — but the surface form hides an underlying structural ambiguity.

¹⁶We are simplifying Krifka’s proposal somewhat, bringing the semantics of plural count nouns closer to Cresswell’s analysis (1977, p. 277). Krifka actually analyses count nouns as in (i), where HAMBURGER is a property of hamburger-stuff, and NU(HAMBURGER) is a measure function that returns the degree to which x constitutes a “natural unit” relative to the HAMBURGER sort.

$$(i) \quad \llbracket \text{hamburgers} \rrbracket = \lambda n \lambda x. \text{HAMBURGER}(x) \ \& \ \text{NU}(\text{HAMBURGER})(x) = n$$


This proposal is very similar to one articulated by Nathan Salmon (1997, p. 10), who suggests that “...numbers are not merely properties of pluralities simpliciter, but relativized properties. They are properties of pluralities relative to some sort or counting property.” An advantage of this analysis is that it provides a semantics for nominals with fractional number terms (see note 10): assuming that the (sorted) NU function is not constrained to return whole numbers as values, a phrase like “2.5 oranges” denotes a property that is true of orange stuff whose measure equals 2.5 orange-units. However, the analysis also involves a commitment to the position that count nouns are in some fundamental sense semantically the same as mass nouns, in that they denote properties of quantities of stuff, rather than properties of atomic objects. This hypothesis raises a number of significant linguistic and philosophical questions, but since our analysis of “average” does not require us to adopt Krifka’s full proposal, we will leave their investigation for another occasion, and work with plural noun denotations of the sort shown in (57). As we will see in the next section, such denotations are sufficient for our purposes because the fractional number terms in “average” sentences do not saturate the degree argument of a plural noun, but rather the degree argument of “average”. This is why the verb phrase in “The average American has 2.3 children” does not denote the property of having 2.3 children.

take scope independently of the rest of the noun phrase. There are many different ways of accounting for the scopal properties of various expressions, which differ primarily in their assumptions about the relation between (surface) syntax and the truth conditional interpretation. As noted above, we will state our analysis here in terms of the framework developed in Heim and Kratzer 1998, in which scope relations are encoded in a syntactic representation of Logical Form that is derived from a surface representation by a transformational operation of quantifier raising (QR). Our proposals, however, are entirely consistent with alternative interpretive frameworks in which scope relations are derived from surface syntactic representations through type- and category-shifting rules.


Quantifier raising has two crucial consequences for the syntactic representation, stated in (59).

- (59) QR of a constituent α to some other constituent β :
- i. leaves a variable-denoting expression indexed i (which we represent as t_i) in the base position of α , and
 - ii. affixes an occurrence of i to β .


This higher occurrence of i interacts with Heim and Kratzer’s composition rule of Predicate Abstraction, which dictates that the $[i \beta]$ constituent is interpreted as a function of type $\langle a, b \rangle$, where a is the semantic type of the variable left behind by α and b is the type of β (see Heim and Kratzer 1998, p. 186). The semantic effect of QR is thus that of λ -abstracting over the base position of the raised constituent. To reflect this fact, we will mix syntactic and semantic representations a bit in our Logical Forms to make the semantic consequences of QR clear, and represent structures in which α raises to adjoin to β as in (60).



The motivation for QR is typically assumed to be type-mismatch: QR provides a kind of “repair strategy” that allows for interpretability without type-shifting. For example, assuming that generalised quantifiers are type $\langle \langle e, t \rangle, t \rangle$ (properties of properties) and transitive verbs are type $\langle e, \langle e, t \rangle \rangle$ (relations between individuals), a sentence such as (61), with a quantified noun phrase in direct object position, is uninterpretable because the verb and the object cannot compose.



QR repairs the type mismatch by raising the quantifier to adjoin to a node of type t (here the sentence node), creating an expression of type $\langle e, t \rangle$ (the function derived by abstracting over the base position of the quantifier), as shown in (62).



Assuming the quantificational determiner denotation for “every” in (63a), this Logical Form is fully interpretable via function application, with the nominal complement of “every” providing its restriction and functional constituent derived by QR providing its scope.


$$(63) \quad \llbracket \text{every} \rrbracket = \lambda f_{\langle e, t \rangle} \lambda g_{\langle e, t \rangle} \cdot \{x \mid f(x)\} \subseteq \{y \mid g(y)\}$$

Composition will assign truth conditions to (62) that render it true if the (contextually restricted) set of hamburgers is a subset of the set of things that Kim ate, and false otherwise.

As far as the semantics is concerned, QR could also apply even when there is no type-mismatch, targeting e.g. a type e argument of a transitive verb, or in the cases we are interested in, the type d number argument of a plural noun. Such a move is typically semantically vacuous, however: all other things being equal, application of QR to a term that is interpretable in its base position will derive a Logical Form whose truth conditions are equivalent to those of the corresponding structure without movement, since application of the function created by QR to the moved expression will have the effect of “putting it back” in the semantics (λ -conversion). It is therefore sometimes assumed that whenever there is reason to believe that e.g. a type e like the name “Jones” undergoes QR (assuming that names do denote individuals), it does so because it has taken on a generalised quantifier denotation (as in Montague 1974): it denotes the set of properties that Jones has ($\lambda f_{\langle e, t \rangle} \cdot f(\text{jones}')$) rather than the individual Jones (jones'). To keep things as simple as possible, we will not make the corresponding assumption about number terms here (that e.g. “3” denotes the generalised quantifier $\lambda f_{\langle d, t \rangle} \cdot f(3)$ whenever the number term undergoes QR), though our proposals are entirely consistent with

such a move.¹⁷ Instead, we will assume that QR of a number term is always an option, even when it is interpretable in its base position.¹⁸

Returning to (53), a Logical Form that with the desired truth conditions can be derived by raising “an average of 17” to adjoin to the verb phrase, as shown in (64).



Assuming existential closure over the variable introduced by the object, the denotation of the sister of “an average of 17.5” is (65), which is a more precise characterisation of the degree relation that we posited earlier in (54).

$$(65) \quad \lambda n \lambda x. \exists y [ate'(y)(x) \ \& \ hamburgers'(y) \ \& \ |y| = n]$$

Composition may then proceed as described above, deriving the truth conditions in (55). In effect, the LF we are positing for (53) is a variant of the synonymous sentence in (66), which uses verbal “average”.

(66) The teens averaged 17.5 in number of hamburgers eaten.

Before moving to the next section, we should say a few words about our assumption that “an average of 17” — and by extension, number terms in general — can undergo quantifier raising. While our assumptions about semantic type certainly allow for this option, one might object that the syntax of English does not allow for such structures, pointing to the impossibility of overt extraction of

¹⁷Note that a semantic analysis of number terms as generalised quantifiers is crucially distinct from that of number terms as quantificational *determiners* (which then combine with NPs to form generalised quantifiers), in that it is rooted in a more basic treatment of number terms as singular terms, in a manner completely parallel to Montague’s treatment of names.

¹⁸In point of fact, QR of the number terms in the examples we discuss below *can* be motivated by considerations of uninterpretability of a slightly different sort. Although the number terms could remain in their base positions with no problem, this would result in Logical Forms that would not satisfy the interpretive requirements of the different forms of “average”. Only through QR of the number terms (or equivalent operations in a framework without movement and LF) can we generate LFs that provide “average” with all of the arguments that it needs.

number terms in examples like (67a-b).¹⁹

- (67) a. *How many did they eat *t* hamburgers?
 b. *It was 17 that they ate *t* hamburgers.

However, there are other kinds of examples which suggest that English syntax does allow for such structures. One case involves quantity comparisons like (68).

- (68) Miller has hit more big shots in playoff games than O'Neal has hit free throws. (*Chicago Tribune, June 3, 2000*)

There is ample syntactic evidence that the comparative clause in examples like this (the complement of "than") involves *wh*-movement, and in particular *wh*-movement of the amount term associated with the nominal "free throws", as shown in (69a) (see e.g., Bresnan 1973; Chomsky 1977; Heim 1985; Hackl 2001; Kennedy 2002; and many others).²⁰

- (69) a. [*wh* O'Neal has hit [*t* free throws]]
 b. *max*{*n* | O'Neal has hit *n* free throws}

The structure in (69a) can be straightforwardly mapped onto an interpretation along the lines of (69b), which involves quantifying over the amount/degree position inside the noun phrase, and the resulting degree description then provides one of the arguments to the comparative relation. (We will have more to say about this relation in the next section.)

Another piece of evidence that the syntax-semantics interface allows an amount term to scope independently of the rest of the NP comes from so-called "reconstruction effects" in quantity questions like (70) (Heycock 1995).

- (70) How many people did Jones decide to hire?

(70) can be interpreted either as a question about the number of people who were actually hired, presupposing the existence of such individuals, as paraphrased in (71a); or as a question about the amount that was decided on, independent of whether anyone was actually hired, as paraphrased in (71b).

¹⁹In fact, Krifka himself takes facts like these as problems for a syntactic implementation of his account of the number-of-events reading of a sentence like "Four thousand ships passed through the lock" that involves scoping the number term (see Krifka 1990, p. 502).

²⁰A simple illustration of this is the fact that this position cannot be filled by an overt amount term: the various options in (i) are completely ungrammatical, even though they are in principle coherent things to say (with the amount terms giving the actual number of free throws that O'Neal has hit, and the rest of the sentence saying that Miller has hit more shots than that).

- (i) *Miller has hit more big shots in playoff games than O'Neal has hit few/17/not many free throws.

- (71) a. What is the number of people such that Jones decided to hire them?
 b. What is the number such that Jones decided to hire that many people?

Different syntactic and semantic mechanisms have been proposed to derive this ambiguity (see Fox 1999 for discussion of alternatives); what is crucial for us is that the reading in (71b) involves scoping only the amount quantifier above the intensional verb *decide* and interpreting the rest of the nominal in its base position in the embedded clause (hence the label ‘reconstruction’).

We take facts like these to support the conclusion that the mapping between syntax and semantics in English is such that an amount/number term amount term may take scope independently of the nominal that appears as its sister in the surface form. For the purposes of this paper, we will assume that this relation is mediated by a syntactic level of Logical Form, and that whatever constraints rule out overt movement of a number term in examples like those in (67) do not apply to covert movement. However, our proposals are perfectly consistent with alternative analyses that achieve the same results through type-shifting or some other mechanism.

4.2.3. Parasitic scope and “the average American”


We are now ready to tackle the final two cases of averaging: the adjectival “average” construction in (40e) and the adverbial form in (40d). For expository purposes, we will frame the discussion in terms of the examples we began the paper with in (72a-b).

- (72) a. The average American has 2.3 children.
 b. Americans have 2.3 children on average.


Recall from our discussion in sections 2 and 3 that one of our questions about such examples is why the verb phrases do not denote the property of having 2.3 children. (Or, to put it more generally, why these examples do not entail the existence of individuals with fractional children.) We now have an answer to this question. Since number terms may take scope independently of the noun phrases in which they occur, these examples can be associated with Logical Forms in which the number term has been raised out of the VP, leaving behind a constituent of the form “have *n* children”. We have already seen that this type of constituent can be used to build the measure function argument of “average”. Given our semantics for plurals, the measure function in this case will be one that maps individuals to the number of whole children that they have, which is exactly the measure function that we need in order to compute the correct truth conditions for (72a-b). The number term “2.3” may then be supplied as the degree argument of “average”, rather than the degree argument of the noun “children”, thereby avoiding problematic entailments about fractional children.

There is a complication, however, which we illustrate with a discussion of (72a). (The same considerations apply to (72b).) Initially, things appear straightforward. First, we assume with Carlson and Pelletier (2002) that the definite article in these constructions is vacuous, which we will indicate by referring to the adjective as “th’average.” (We will have more to say about this issue in section 5.3.) The surface syntax of (72a) suggests that “th’average” combines first with the domain term (“American”), so we just need to determine the order of composition of the degree relation and the average. When we actually try to construct candidate Logical Forms, however, we run into a problem. Applying QR to the number term derives either (73a) or (73b) (depending on whether the number term raises above or below the subject), but neither of these representations are adequate.

(73) a.



b.




The problem with (73a) is that the constituent that should be provided as the degree argument of “th’average” (the number term) and the one that should provide the degree relation (the scope of the number term, marked by λn) form a syntactic constituent. Compositionality dictates that these two elements also form a semantic constituent, and indeed, this is the normal result of interpreting representations involving quantifier raising (either the raised expression is the function and the λ -term is its argument, or vice-versa). But this means that neither term can be supplied as arguments to “th’average”. The problem with (73b) is worse, since “th’average” is itself contained inside the scope of the number term, so there is no way that the latter could provide one of the arguments of the former. The potential Logical Forms for (72b) will be identical in the relevant respects, since the number term will have to raise to a position above or below the adverb.

In short, we appear to have no way of separating the constituent that provides the average (the raised number term) from the constituent that provides

the degree relation (its scope) and supplying them as independent arguments to “th’average” or “on average”. To be more precise, we have no way of doing this and simultaneously maintaining first, that semantic composition is local (defined over immediate subconstituents), and second, that natural language makes use of a very limited set of composition rules: function application, function composition, predicate abstraction, and perhaps a small number of lexical type-shifting principles. If we were to give up these assumptions, we would have various options for dealing with the structures in (73a-b). For example, we could handle (73a) by invoking a non-local composition rule such as (74a).

- (74) a. If α has the form in (74b), where $\llbracket A \rrbracket$ is a function of type $\langle b, \langle c, \dots \rangle \rangle$, $\llbracket B \rrbracket$ is type b and $\llbracket C \rrbracket$ is type c , then $\llbracket \alpha \rrbracket = \llbracket A \rrbracket(\llbracket B \rrbracket)(\llbracket C \rrbracket)$.



But this rule results in a system that is not fully compositional, as β is assigned no denotation.

To avoid this problem, we might instead posit the composition rules in (75a-b).

- (75) a. If β is a constituent with daughters B, C , then $\llbracket \beta \rrbracket = \langle \llbracket B \rrbracket, \llbracket C \rrbracket \rangle$.
 b. If α is a constituent with daughters A, β , $\llbracket A \rrbracket$ is a relation whose domain consists of pairs of items of type b and c , $\llbracket \beta \rrbracket = \langle \llbracket B \rrbracket, \llbracket C \rrbracket \rangle$ such that $\llbracket B \rrbracket$ is type b and $\llbracket C \rrbracket$ is type c , then $\llbracket \alpha \rrbracket = \llbracket A \rrbracket(\llbracket \beta \rrbracket)$.

But these rules are clearly ad hoc, designed to handle “the average American” but not obviously relevant to other constructions. If it turned out that the only way to provide a fully compositional account of (72a) and related examples involved invoking ad hoc or non-local composition principles such as (75a-b) or (74), then it would be fair to say that our analysis does no better in responding to Chomsky’s challenge than the analyses we criticised in section 3.

Fortunately for us, it turns out that the assumptions that we need to make in order to provide a compositional account of the truth conditions of (72a) and (72b) are ones that already have a substantial amount of independent support, and that have been invoked in order to account for the interpretation of comparative constructions (Heim 1985; Bhatt and Takahashi 2007; Kennedy to appear), distributive interpretations of plural NPs (Sauerland 1998), and noun-modifying uses of “same” and “different” (Barker 2007). Specifically, we need to allow for the possibility that a third constituent can intervene between a scope-taking constituent and the function-denoting constituent that normally serves as its scope, a configuration that Barker (2007) dubs *parasitic scope*.

Comparatives like (76) provide a good illustration of parasitic scope.²¹

(76) More people live in New York than Chicago.

Most work on comparatives assumes that “more” and the phrase introduced by “than” form a constituent in Logical Form. There is also a substantial amount of syntactic evidence that the comparative standard in an example like (76) (the complement of “than”) can be a simple noun phrase, rather than an underlyingly clausal structure (Hankamer 1973).²² Such a structure requires the denotation for “more” in (77), which is looking for two individual arguments — a standard of comparison y and a target of comparison x — and a degree relation (Heim 1985; Hoeksema 1983; Bhatt and Takahashi 2007; Kennedy 1999, to appear).

(77) $\llbracket more \rrbracket = \lambda y \lambda f_{\langle d, et \rangle} \lambda x. f_{meas}(x) \succ f_{meas}(y)$

The degree relation is converted into a measure function and applied to the target and standard, so that the truth conditions of a comparative are ultimately stated in terms of an asymmetric ordering between two degrees.

In (76), the standard argument is directly provided by “Chicago” (“than” is typically assumed to be vacuous). In order to derive the right truth conditions, the target argument should be “New York”, and the degree relation should be the one in (78).

(78) $\lambda n \lambda x. n \text{ people live in } x$


²¹By far the most complete semantic characterisation of parasitic scope is the one developed by Barker (2007) to account for the sentence-internal reading of “same” in examples like (i) (where sameness is calculated with respect to the books read by the entities falling under “everyone”, rather than a book previously mentioned in the discourse).

(i) Everyone read the same book.

Barker shows that the interpretation of examples like (i) crucially require “same” to intervene between the quantifier and its nuclear scope, much like “the average American” intervenes between a number term and its scope, as illustrated below. Barker ultimately states his analysis in terms of a type-logical grammar with continuations, rather than in terms of Heim and Kratzer-style LFs (though he also shows how the latter approach would work). This type of approach may ultimately prove to have certain empirical and theoretical advantages over one stated in terms of Logical Forms (see Kennedy and Stanley 2008 for discussion), but it has the disadvantage of being highly technical and difficult to understand for those not familiar with the logic. Since our goal in this paper is not to choose between different compositional approaches to the syntax-semantics interface, but rather to show that abstract “average” can be accounted for in terms of independently motivated assumptions about the properties of this interface, we will refer readers who wish to see a fully worked-out logic of parasitic scope to Barker 2007 and will formulate our own analysis in terms of a syntactic representation of Logical Form.

²²In the latter case, the problems described below do not arise, since the meaning of the standard works out to be something like *the number n such that n people live in Chicago*, and there is no need to recover a degree relation from the rest of the clause.


We can derive such a relation by raising both “more than Chicago”, which saturates the degree argument of the plural NP (the same slot occupied by a number term), and “New York”, but only if QR of the former can target the functional constituent created by QR of the latter; this is the sense in which the former is “parasitic” on the scope of the latter. The resulting LF is shown in (79).



QR of “New York” derives the constituent whose left-hand daughter is λx . Subsequent QR of “more than Chicago” to adjoin to this constituent derives the constituent whose left-hand daughter is λn , which denotes precisely the degree relation in (78). This relation is supplied as the second argument to the comparative morpheme, which gives us exactly the truth conditions we want:

$$(80) \quad \lambda x.max\{n \mid n \text{ people live in } x\}(NY') \succ \lambda x.max\{n \mid n \text{ people live in } x\}(Chicago')$$

Returning now to “the average American”, the Logical Form we need in order to derive the correct truth conditions for (72a) is one that is parallel in the relevant respects to (79). Specifically, we need a representation in which “the average American” is parasitic on the scope of the number term, as shown in (81).




Assuming the denotation of “th’average” in (82), this LF can be straightforwardly interpreted using only function application; there is no need to invoke arbitrary interpretive principles of the sort shown in (74) and (75) above. The full derivation of the truth conditions is shown in (83).

- (82) $\llbracket [A \text{ th'average}] \rrbracket = \lambda S \lambda f \lambda d. \mathbf{average}(f)(S)(d)$
- (83) a. $\llbracket [A \text{ th'average}] \rrbracket(\llbracket [American] \rrbracket)(\llbracket [\lambda n \lambda x.x \text{ has } n \text{ children}] \rrbracket)(\llbracket [2.3] \rrbracket)$
 b. $\mathbf{average}(\llbracket [\lambda n \lambda x.x \text{ has } n \text{ children}] \rrbracket)(\llbracket [American] \rrbracket)(\llbracket [2.3] \rrbracket)$
 c. $\frac{\sum_{y \in \mathit{American}'} \lambda z. \mathit{max}\{n \mid z \text{ has } n \text{ children}\}(y)}{|\mathit{American}'|} = 2.3$

Constructions with adverbial “on average”, such as (72b), are analysed in exactly the same way. Assuming that “on average” attaches to VP (it can be preposed and included in VP-ellipsis), we can analyse its meaning as in (84) and posit the Logical Form in (85) for (72b).

- (84) $\llbracket [PP \text{ on average}] \rrbracket = \lambda f \lambda d \lambda S. \mathbf{average}(f)(S)(d)$



(86) shows the derivation of the truth conditions of this LF; once again, we end up with a meaning that is equivalent what we get in the other “average” constructions.

- (86) a. $[[\text{PP on average}]](\lambda n \lambda x. x \text{ has } n \text{ children})(2.3)([[\text{Americans}]])$
 b. $\mathbf{average}(\lambda n \lambda x. x \text{ has } n \text{ children})([[\text{Americans}]])(2.3)$
 $\sum_{y \in \text{Americans}'}$ $\lambda z. \max\{n \mid z \text{ has } n \text{ children}\}(y)$
 c. $\frac{\sum_{y \in \text{Americans}'}}{|\text{Americans}'|} = 2.3$

We have thus accounted for our two most difficult cases — “the average NP” and “on average” — without positing any special interpretive mechanisms beyond those that are independently necessary to account for other constructions. The final piece of the analysis was the assumption that natural language allows for the possibility that some expressions can take another expression’s nuclear scope as an argument — in Barker’s terms, they may be parasitic on the scope of another term. The literature on comparatives, plurals and “same”/“different” that we have cited indicates that such an option must be available to the interpretive system; “average” can be viewed as further evidence for this conclusion.²³

4.3. Summary

In this section, we have provided a semantics of averaging according to which (morphosyntactically) definite noun phrases of the form “the average NP”, with an abstract interpretation of “average”, are semantically not referring expressions, but rather what we might call *averaging expressions*. As such, they do not involve reference to bizarre individuals or predications of odd properties (such as the property of having 2.3 children). Crucially, our analysis is fully compositional, and accounts for the interpretation of “the average NP” and its adverbial cousin “on average” strictly in terms of independently justified assumptions about the syntax-semantics interface. Finally, our analysis provides an empirical advantage over all previous analyses in extending beyond these two forms of “average” and explaining how the various ways of expressing averages illustrated in (40) give rise to the same core truth conditions.

²³Important questions remain about how exactly parasitic scope should be accounted for in the grammar, and what its implications are for the syntax-semantics interface. For example, in the system developed in Barker 2007, parasitic scope is a consequence of the overall logic of quantification. In contrast, Sauerland (1998) and Bhatt and Takahashi (2007) derive parasitic scope from the syntax of quantifier raising: given the statement of QR in (59), parasitic scope arises when a constituent B raises to adjoin to the λ -term created by QR of another constituent A. The “average” data may indicate that this derivational approach is not general enough, however. Adverbs are typically assumed to occupy fixed positions in the syntactic representation, so there would be no way to derive the representation in (85) through the operation of QR alone. See Kennedy and Stanley 2008 for a more detailed discussion of the implications of “average” for the grammatical characterisation of parasitic scope and for the syntax-semantics interface.

5. Extending the analysis

5.1. Comparison

Extending our view beyond simple examples like “The average American has 2.3 children,” an important question is how our analysis handles comparative constructions such as (87a), which can be interpreted as in (87b). (This example can of course also have a concrete interpretation, whereby it is claimed that the typical Norwegian male is taller than the typical Italian male.)

- (87) a. The average Norwegian male is taller than the average Italian male.
 b. The average height of Norwegian males exceeds the average height of Italian males.

Recall that a problem for Stanley’s (2001) analysis, in which “average” DPs denote degrees, is that true degree-denoting expressions cannot appear as arguments to *taller than* (see (22) in section 3.2). Our analysis has no such problem, and more importantly, it straightforwardly maps (87a) into truth conditions parallel to (87b).


To see how, we must say a few more words about comparative constructions. As we illustrated in the previous section, comparatives in which “than” is followed by a DP in the surface form may be interpreted ‘directly’, sometimes using parasitic scope. However, it is generally accepted that comparatives which have this structure on the surface are syntactically ambiguous between an underlying form that mirrors the surface structure (as in the previous section) and one in which the standard constituent has an underlying clausal structure. Specifically, the ‘comparative clause’ has a syntactic analysis as a *wh*-movement structure in which a null operator binds a degree variable inside a copy of the gradable predicate that appears in the matrix, which is elided in the surface form together with other material that is identical to material in the matrix sentence. Such structures are interpreted as properties of degrees, and directly provide the standard argument for a clausal variant of “more”, whose denotation is shown in (88) (see e.g., Hankamer 1973; Bresnan 1973; Chomsky 1977; Heim 1985; Kennedy 1999; Lechner 2001 and many others). (The *max* operator in (88) returns the maximal degree that satisfies its type $\langle d, t \rangle$ argument.)

$$(88) \quad \llbracket more_{clausal} \rrbracket = \lambda f_{\langle d, t \rangle} \lambda g_{\langle d, t \rangle} . max(g) \succ max(f)$$

As we saw above, the comparative morpheme and comparative clause form a constituent which undergoes QR at LF and normally directly binds the degree argument of a gradable predicate. To handle (87a), all we need to do is assume that “the average Norwegian male” and “the average Italian male” can parasitically take the scope arguments of the two degree operators — the entire comparative constituent in the matrix clause and the null degree operator in the comparative clause — as their arguments, as shown in (89) (where “Norwegian” and “Italian”

abbreviate “Norwegian male” and “Italian male”, respectively).

(89)



Given the semantics for the comparative in (88), (89) is true just in case the relation in (90) holds.

$$(90) \quad \max([\alpha]) \succ \max([\beta])$$

Assuming that “than” and the null operator in the comparative clause are both semantically vacuous (as is standardly done; movement of the null operator creates a degree property in line with the semantics of quantifier raising discussed in section 4.2.2), the denotations assigned to the constituents marked α and β in (89) are as shown in (91).

$$(91) \quad [\alpha] = \lambda d. \frac{\sum_{y \in Norwegian'} \mathbf{height}(y)}{|Norwegian'|} = d \quad [\beta] = \lambda d'. \frac{\sum_{y \in Italian'} \mathbf{height}(y)}{|Italian'|} = d'$$

Putting everything together gives us (92) as the denotation of (87a), which is exactly what we want. (Here we use the ι operator rather than max to reflect the fact that the sets of degrees that satisfy the properties in (91) are singletons.)

$$(92) \quad \iota d \left[\frac{\sum_{y \in Norwegian'} \mathbf{height}(y)}{|Norwegian'|} = d \right] \succ \iota d' \left[\frac{\sum_{y \in Italian'} \mathbf{height}(y)}{|Italian'|} = d' \right]$$

The same kind of analysis will extend to examples like (93a), assuming the pronoun in the comparative clause can be analysed as a covert definite description,

so that the sentence's Logical Form looks like (93b) (see Evans 1977, Neale 1990, and especially Elbourne 2005).

- (93) a. The average American has more cars than he has children.
 b. more [than th'average American λm [has m children]] [th'average American λn [has n cars]]

More generally, the analysis we have outlined here should in principle be applicable to any construction whose compositional interpretation involves degree relation, either derived (via movement of a number term or degree quantifier, as in the comparative) or lexical, if the syntax of the construction allows for parasitic scope. Space prohibits us from fully exploring this prediction here, but we know of no counterexamples to this prediction.

5.2. *Anaphora and conjunction*

Recall that one of the virtues of Carlson and Pelletier's (2002) assumption that there is no distinction between concrete and abstract "average" is that they can account for examples like (94a-b), in which the "average" NP combines with two verb phrases, each of which appear to require a different sense of "average".

- (94) a. The average American has 2.3 children and drives a domestic automobile.
 b. The average traveller belongs to 3.3 frequent flyer programs and prefers to fly nonstop.

On their analysis, this is expected, since they explicitly deny a semantic distinction between concrete and abstract "average". These examples appear to raise a significant challenge for our proposals, however.

First, recall that our analysis assumes that "the average NP" (on the abstract interpretation) must combine with a degree relation, which is created by raising a number term out of the VP and invoking parasitic scope to bind off the base position of the number term. In order to construct the right sort of Logical Form in examples like (94a-b), however, the number term would have to raise out of one sub-part of a conjunction structure, in violation of the Coordinate Structures Constraint (Ross 1967). If Quantifier Raising obeys syntactic constraints, then such movement should be impossible, and the number term would instead have to remain in its base position. But this would in turn mean that the only option for interpreting the conjoined VP would be as a complex property, rather than a degree relation. The conjoined VP in (94a), for example, would denote property of having 2.3 children and driving a domestic automobile. This VP would then be predicated of the subject — which would necessarily involve concrete "average" (to avoid a type mismatch) — generating an entailment that some American (the average one) has 2.3 children. This is clearly the wrong prediction, as (94a) lacks

such an entailment, and indeed the first part of (94a) appears to have the usual abstract meaning.

We could avoid this problem by instead hypothesising that (for whatever reason) the number term can raise out of the left-hand part of the conjoined VP. This would result in an interpretable Logical Form in which the degree relation that the “average” NP combines with is based on the two conjuncts. In (94a), for example, movement of the number term to a position above the subject, plus parasitic scope, will derive the degree relation in (95).

(95) $\lambda n \lambda x. x \text{ has } n \text{ children and } x \text{ drives a domestic automobile}$

This relation is of the appropriate type to combine with *the average American*, and plugging it in as the measure argument to **average** will ultimately derive the truth conditions in (96).

$$(96) \quad \frac{\sum_{y \in \text{American}'} \lambda x. \max\{n \mid x \text{ has } n \text{ children and } x \text{ drives a domestic auto}\}(y)}{|\text{American}'|} = 2.3$$

This is not what we want, however: the value returned by the measure function for any American who does not drive a domestic automobile will be zero, which means that as long as a sizable portion of Americans do not drive domestic cars, the ‘average’ in (94a) should be understood as being much lower than the actual average number of children that objects in the domain have. This is not how we understand this sentence, however; instead we understand it as in (97), where the noun-modifying “average” is the concrete one and the one that combines directly with the number term is abstract.

(97) The average American has an average of 2.3 children and drives a domestic automobile.

Our semantics handles (97) with no problem, because the job of doing the averaging is taken over by “an average of”, as outlined in section 4.2.2. If it were possible to show that (94a) and similar examples could somehow be analysed as including a covert occurrence of “an average of”, the problem that they present for our proposals would disappear.

In fact, there is reason to believe that such an option is possible. First, examples like the following show concrete and abstract “average” can be combined in the same sentence ((98a) is from the same source as (1b)):

- (98) a. The Average Freddie voter took an average of 14.9 domestic trips in the past year. (*www.freddieawards.com/events/17/trivia.htm*)
 b. The average Democratic senator from a red state enjoyed a +26.7 average net approval rating (which equals roughly a 63% approval rating), whereas the average Republican senator from a red state had just a

+17.2 average net approval rating.
 (politicalinsider.com/2007/06/who_is_the_most_popular_group.html)

These examples all have multiple occurrences of “average”, where the ones associated with the number terms/measure nouns are presumably abstract, and the ones contained in the subjects (“the average Freddie voter”, “the average Democratic senator from a red state”, etc.) are concrete. And indeed, the latter can be replaced with *typical* with no change in meaning, indicating no inherent problem with having both concrete and abstract “average” in the same sentence.

Evidence that abstract “average” may sometimes be covert comes from examples like the following, which appeared in the *New York Times*:

- (99) One survey, recently reported by the federal government, concluded that men had a median of seven female sex partners. Women had a median of four male sex partners. Another study, by British researchers, stated that men had 12.7 heterosexual partners in their lifetimes and women had 6.5.
 (www.nytimes.com/2007/08/12/weekinreview/12kolata.html)

The crucial part is the third sentence: “Another study ... stated that men had 12.7 heterosexual partners in their lifetimes and women had 6.5.” While it is in principle possible for people to have fractional sexual partners, such an interpretation is not the most natural one for this sentence; instead it is understood as in (100), providing compelling evidence for the existence of a covert abstract “average”.

- (100) Another study stated that men had an average of 12.7 heterosexual partners in their lifetimes and women had an average of 6.5.

Whether this covert element is in fact the nominal form or some other form (e.g., a covert occurrence of “on average”), and whether it is a true null expression, something derived through ellipsis, or the result of syntactic reanalysis are not questions that we can answer at this time. What is important is that the facts indicate that at least in certain contexts, it is possible to parse sentences like the last one in (99) and, we claim, those in (94), as containing a covert occurrence of abstract “average”.²⁴

²⁴This explanation of the facts in (94) can be extended to variants like (ia), in which the second conjunct has a pronominal subject, or (ib), based on an example from Chomsky discussed in Ludlow 1999, p. 174.

- (i) a. The average businesswoman belongs to 3.3 frequent flyer programs and she prefers to fly nonstop.
 b. In your report on the average businesswoman, you failed to note that she belongs to 3.3 frequent flyer programs.

We assume the overt occurrences of “average” in these examples are concrete ones, that the pronouns are covert definite descriptions anaphoric to the “average” NPs (Evans 1977; Neale 1990; Elbourne 2005), and that like the cases discussed above, the correct truth conditions are derived by inserting a

Of course, if this explanation is correct, then it is appropriate to ask why sentences like (101a-b) are anomalous: should not world knowledge force a parse with covert “average” to avoid nonsensical truth conditions?

- (101) a. ??The typical American has 2.3 children.
 b. ??The typical traveller belongs to 3.2 frequent flyer programs.

We suspect that for any sentence of the relevant type, a parse involving covert “average” is dependent on a certain amount of contextual priming. In (99), the first two sentences of the passage explicitly mention medians. And while the same cannot be said of the examples in (94), it may very well be the case that the use of concrete “average” is itself enough to license a covert occurrence of abstract “average” (or reanalysis, if that is what is actually going on here).

Another question that arises if our explanation of the examples in (94) is correct is whether the adjectival form of “average” in the variants of (101a-b) that we began this paper with is ever abstract. If concrete “average” can license a covert instance of “an average of”, then in principle all sentences of “the average American” type could be handled in this way. While this is in principle possible, given the fact that abstract “average” clearly has instantiations in all other categories, and the fact that independently attested principles of semantic composition can derive the correct truth conditions for an adjectival form of abstract “average”, as we showed in section 4.2.3, it seems unlikely that the learner would fail to posit a such lexical item. We will therefore continue to assume that abstract “average” is what we normally see in sentences like “The average American has 2.3 children”, and that reanalysis in terms of a covert abstract average occurs only as a last resort in cases like those in (94).

5.3. Why “the” average American?

We conclude with a few thoughts on the status of the definite article in “the average NP”: why is abstract “average” restricted to occur with “the”, and why is “the” semantically vacuous? Although our stipulation in section 4.2.3 that “the” is vacuous does not put our proposal in any worse shape than other proposals (Carlson and Pelletier, for example, make the same assumption), we actually think there may be a plausible morphosyntactic and semantic answer to this question.

Our answer builds on earlier proposals by Svenonius (1992) and Larson (1998), who argue that certain adjectives incorporate into the determiner position. Svenonius provides evidence for this claim from Norwegian. As shown in (102a-b), Norwegian definite nouns must appear without a determiner when they are bare, but with a determiner when preceded by an adjective.²⁵

covert abstract “average”.

²⁵In the following examples, (**x*) means that the form is ungrammatical when *x* is present and grammatical when *x* is omitted; *(*x*) means the opposite.

- (102) a. (*den) møtet
 (*the) meeting.DEF
 b. *(den) viktige møtet
 *(the) important meeting.DEF

Certain adjectives, including *samme* “same” and *første* “first” can license definite marking on the noun in the absence of a determiner, however:

- (103) a. samme trøtte maten
 same boring food.DEF
 b. første viktige møtet
 first important meeting.DEF

This option is available only when the adjectives are leftmost, however; when they are themselves preceded by another adjective, the determiner is again obligatory:

- (104) a. *(den) trøtte samme maten
 *(the) boring same food.DEF
 b. *(den) viktige første møtet
 *(the) important first meeting.DEF

Svenonius takes these facts to indicate that *samme* and *første* are ‘determining adjectives’, which bear a definiteness feature that allows them to incorporate into the determiner position in (103), licensing the morphology on the noun. When another adjective intervenes, such incorporation is blocked, and an overt determiner must be inserted. Larson (1998) makes use of Svenonius’ ideas in his analysis of the prenominal “occasional” in examples like (105), arguing that it incorporates into the determiner, thereby forming a quantifier that ranges over both individuals and events.

- (105) The occasional sailor passed by.

We would like to suggest that something similar is going on with prenominal “average”, whereby it incorporates into the determiner position at LF, and the surface realisation of “the” is a kind of expletive element inserted to satisfy morphophonological requirements (e.g., definiteness features on D) at PF. (Alternatively, we might take our notation of “th’average” in (81) as closer to the truth and posit a complex determiner analogous to “another”.) The fact that no other (non-appositive) modifiers may intervene between “average” and “the” then falls out from conditions on locality of movement, which prohibit a head from crossing intervening landing sites. On the semantic side, the fact that we get definiteness features on the determiner can be justified based on the semantics of the entire construction, which ends up picking out a unique degree.

6. Conclusion

Semanticists generally assume that, as Richard Larson and Gabriel Segal write (Larson and Segal 1995, p. 5), "part of the pretheoretical domain of semantics concerns the relation between language and the world." As Gennaro Chierchia and Sally McConnell-Ginet note (1990, p. 55), from this "denotational point of view, symbols stand for objects. Consequently, configurations of symbols can be used to encode how objects are arranged and related to one another." The constructions we have been discussing pose a *prima facie* worry for the denotational perspective. If the denotational perspective requires the truth of "The average American has 2.3 children" to entail the existence of an average American who has an impossible number of children, then the denotational perspective is incorrect. What we have shown in this paper is that sentences such as "The average American has 2.3 children" do not in fact entail the existence of Americans with impossible numbers of children. The occurrence of "the average American" in such constructions does not stand for an object; it is not a singular term occurrence. But this is a conclusion fully consistent with the denotational, truth-conditional perspective in semantics, rather than in opposition to it.

Our specific analysis also has other foundational consequences. As we have seen, many authors suspicious of the claim that numbers are objects have taken the quantificational determiner use of number terms to be central, and the use of number terms as singular terms to be peripheral. The work of Thomas Hofweber provides one very recent example. Recall Frege's point that "Jupiter has four moons" can be converted into the apparent identity statement "the number of Jupiters moons is four". Hofweber (2007) argues that the occurrence of "four" in the latter sentence is not in fact a genuine singular term use. The reason "four" occurs where it does in this construction is that it has been moved for conversational purposes to a focus position in the sentence. (Though as Moltmann (2008) points out, such a syntactic transformation is "very implausible".) Semantically, it functions as a determiner. If so, Frege's attempt to provide natural-language examples of singular term uses of number terms fail. In other work, Hofweber argues that the apparent singular term use of number terms in arithmetic is also not a genuine singular term use (Hofweber 2005) Like Hodes (1984), he regards a sentence such as " $5 + 7 = 12$ " as "really" telling us that when we take five objects and seven distinct objects, we have twelve objects. So Hofweber takes himself to have explained away all apparent singular term uses of number terms. If philosophers such as Hodes and Hofweber are correct, numbers are not needed as the referent of any genuine singular terms either in arithmetic or in natural language.

One consequence of our analysis is that programs of the sort envisaged by Hodes and Hofweber are, at the very least, seriously incomplete. As we have made clear in our critical discussion of other proposals, in a construction such as "The average American has 2.3 children", the number term cannot be plausibly analysed as

a quantificational determiner. Instead, in our final semantic analysis, “2.3” occurs as a genuine singular term, one that flanks an identity sign. In fact, as we have also suggested, one can make a similar point with other uses of number terms, such as the use of “2.3” in a construction such as “John lives 2.3 miles away”, and possibly also in constructions like Salmon’s “There are 2.3 oranges on the table”, which (unlike “average” sentences) generate actual entailments about fractional objects. Though Frege did not draw our attention to these constructions, these uses of number terms to provide measures are perhaps better examples of uses of number terms as singular terms than the ones he provided. It is an interesting consequence of our analysis that in a construction many have appealed to in arguments against taking singular reference to numbers at face value, we find perhaps the most compelling examples of such reference.²⁶

Department of Linguistics
University of Chicago
 1010 E. 59th St.
 Chicago, IL 60637
 USA
 ck@uchicago.edu

CHRISTOPHER KENNEDY

Department of Philosophy
Rutgers University
 26 Nichol Avenue
 New Brunswick, NJ 08901
 USA
 jasoncs@rucss.rutgers.edu

JASON STANLEY

References

- Abney, S. P. 1987. *The English Noun Phrase in its Sentential Aspect*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, Mass.
- Barker, C. 2007. ‘Parasitic scope’. *Linguistics and Philosophy*, 30:407–444.
- Barwise, J. and Cooper, R. 1981. ‘Generalized quantifiers and natural language’. *Linguistics and Philosophy*, 4:159–219.
- Bhatt, R. and Takahashi, S. 2007. ‘Direct comparisons: Resurrecting the direct analysis of phrasal comparatives’. In EDITORS (ed.), *Proceedings of SALT 17*. CLC Publications, PAGES.

²⁶We are grateful to two anonymous reviewers and the editor of *Mind* for thoughtful comments on an earlier draft, and to Peter Alrenga, Chris Barker, Cian Dorr, Richard Heck, Harold Hodes, Thomas Hofweber, Peter Ludlow, Roger Schwarzschild, Ted Sider, and audiences at St. Andrews University, the University of Chicago, Michigan State University, the University of Maryland, Semantics and Linguistic Theory 18, and the XII World Congress of Philosophy for helpful feedback. This paper is based in part upon work supported by the National Science Foundation under Grant No. BCS-0620247.

- Bresnan, J. 1973. 'Syntax of the comparative clause construction in English'. *Linguistic Inquiry*, 4:275–343.
- Carlson, G. and Pelletier, F. J. 2002. 'The average American has 2.3 children'. *Journal of Semantics*, 19:73–104.
- Carpenter, B. 1997. *Type-Logical Semantics*. Cambridge, MA: MIT Press.
- Chierchia, G. and McConnell-Ginet, S. 1990. *Meaning and Grammar: An Introduction to Semantics*. Cambridge, Mass.: MIT Press.
- Chomsky, N. 1955. *The Logical Structure of Linguistic Theory*. New York: Plenum. (published in 1975).
- 1957. *Syntactic Structures*. The Hague: Mouton.
- 1977. 'On wh-movement'. In P. Culicover, T. Wasow, and A. Akmajian (eds.), *Formal Syntax*. New York: Academic Press, 71–132.
- 2000. *New Horizons in the Study of Language*. Cambridge: Cambridge University Press.
- Cresswell, M. J. 1977. 'The semantics of degree'. In B. Partee (ed.), *Montague Grammar*. New York: Academic Press, 261–292.
- Dummett, M. 1991. *Frege: Philosophy of Mathematics*. London: Duckworth.
- Elbourne, P. 2005. *Situations and Individuals*. Cambridge, MA: MIT Press.
- Evans, G. 1977. 'Pronouns, quantifiers, and relative clauses, I & II'. *The Canadian Journal of Philosophy*, 7:467–536, 777–797.
- Fox, D. 1999. 'Reconstruction, variable binding and the interpretation of chains'. *Linguistic Inquiry*, 30:157–196.
- Frege, G. 1980. *The Foundations of Arithmetic*. Evanston, IL: Northwestern University Press.
- Hackl, M. 2001. 'Comparative quantifiers and plural predication'. In K. Megerdooian and L. Bar-el (eds.), *Proceedings of WCCFL XX*. Somerville, MA: Cascadilla Press, 234–247.
- Hankamer, J. 1973. 'Why there are two 'than's in English'. In C. Corum, T. C. Smith-Stark, and A. Weiser (eds.), *Proceedings of the 9th Annual Meeting of the Chicago Linguistics Society*. Chicago, IL: Chicago Linguistics Society.
- Heim, I. 1985. 'Notes on comparatives and related matters'. Ms., University of Texas.
- Heim, I. and Kratzer, A. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.
- Heycock, C. 1995. 'Asymmetries in reconstruction'. *Linguistic Inquiry*, 26:547–570.
- Higginbotham, J. 1985. 'On semantics'. *Linguistic Inquiry*, 16:547–593.
- Hodes, H. 1984. 'Logicism and the ontological commitments of arithmetic'. *Journal of Philosophy*, 81:123–149.
- Hoeksema, J. 1983. 'Negative polarity and the comparative'. *Natural Language & Linguistic Theory*, 1:403–434.
- Hofweber, T. 2005. 'Number determiners, numbers, and arithmetic'. *The Philosophical Review*, 114:179–225.

- 2007. 'Innocent statements and their metaphysically loaded counterparts'. *The Philosopher's Imprint*, 7.
- Hornstein, N. 1984. *Logic as Grammar*. Cambridge, Mass.: MIT Press.
- Keenan, E. and Stavi, J. 1986. 'A semantic characterization of natural language determiners'. *Linguistics and Philosophy*, 9:253–326.
- Kennedy, C. 1999. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. New York: Garland. (1997 UCSC Ph.D thesis).
- 2002. 'Comparative deletion and optimality in syntax'. *Natural Language & Linguistic Theory*, 20.3:553–621.
- to appear. 'Modes of comparison'. To appear in the Proceedings of the Chicago Linguistic Society 43.
- Kennedy, C. and Stanley, J. 2008. 'What an *average* semantics needs'. Paper presented at SALT 18, University of Massachusetts, Amherst, 21 March, 2008.
- Klein, E. 1991. 'Comparatives'. In A. von Stechow and D. Wunderlich (eds.), *Semantik: Ein internationales Handbuch der zeitgenössischen Forschung* (Semantics: An International Handbook of Contemporary Research), chap. 32. Berlin: de Gruyter, 673–691.
- Krifka, M. 1989. 'Nominal reference, temporal constitution and quantification in event semantics'. In R. Bartsch, J. van Benthem, and P. van Emde Boas (eds.), *Semantics and Contextual Expression*. Stanford, CA: CSLI Publications, 75–115.
- 1990. 'Four thousand ships passed through the lock: Object induced measure functions on events'. *Linguistics and Philosophy*, 13:487–520.
- Kripke, S. A. 1982. *Wittgenstein on Rules and Private Language*. Cambridge, Mass.: Harvard University Press.
- Larson, R. 1998. 'Events and modification in nominals'. In D. Strolovitch and A. Lawson (eds.), *Proceedings from SALT VIII*. Ithaca, NY: CLC Publications, 145–168.
- Larson, R. K. and Segal, G. 1995. *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, Mass.: MIT Press.
- Lechner, W. 2001. 'Reduced and phrasal comparatives'. *Natural Language and Linguistic Theory*, 19:683–735. Unpublished ms., Universität Tübingen.
- Link, G. 1983. 'The logical analysis of plurals and mass terms: A lattice theoretical approach'. In R. Bäuerle, C. Schwarze, and A. von Stechow (eds.), *Meaning, Use, and the Interpretation of Language*. Berlin: de Gruyter, 302–323.
- 1998. *Algebraic Semantics in Language and Philosophy*. Stanford, CA: CSLI Publications.
- Ludlow, P. 1999. *Semantics, Tense and Time: An Essay in the Metaphysics of Natural Language*. Cambridge, MA: MIT Press.
- Melia, J. 1995. 'On what there's not'. *Analysis*, 55:223–229.
- Moltmann, F. 2008. 'Reference to numbers in natural language'. Unpublished ms.

- Montague, R. 1974. 'The proper treatment of quantification in ordinary English'. In R. Thomason (ed.), *Richard Montague: Selected Papers*. New Haven, Connecticut: Yale University Press, 247–270.
- Neale, S. 1990. *Descriptions*. Cambridge, MA: MIT Press.
- Quine, W. v. O. 1960. *Word and Object*. Cambridge, Mass.: MIT Press.
- Ross, J. R. 1967. *Constraints on Variables in Syntax*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, Mass.
- Salmon, N. 1997. 'Wholes, parts, and numbers'. *Noûs Supplement: Philosophical Perspectives, 11, Mind, Causation, and World*, 31:1–15.
- Sauerland, U. 1998. 'Plurals, derived predicates and reciprocals'. In U. Sauerland and O. Percus (eds.), *The Interpretive Tract, MITWPL 25*. Cambridge, Mass.: MIT Working Papers in Linguistics, 177–204.
- Stanley, J. 2001. 'Hermeneutic fictionalism'. *Midwest Studies in Philosophy*, 25:36–71.
- Stump, G. 1981. 'The interpretation of frequency adjectives'. *Linguistics and Philosophy*, 4:221–257.
- Svenonius, P. 1992. 'The structural location of the attributive adjective'. In E. Duncan, D. Farkas, and P. Spaelti (eds.), *Proceedings of the 11th West Coast Conference on Formal Linguistics*. Stanford, Calif.: CSLI Publications, 438–454.
- Yablo, S. 1998. 'Does ontology rest upon a mistake?' *Proceedings of the Aristotelian Society Supplement*, 72:229–261.