

The Average American has 2.3 Children

GREG CARLSON
University of Rochester

FRANCIS JEFFRY PELLETIER
University of Alberta

Abstract

Average-NPs, such as the one in the title of this paper, have been claimed to be ‘linguistically identical’ to any other definite-NPs but at the same time to be ‘semantically inconsistent’ with these other definite-NPs. To some this is an ironclad proof of the irrelevance of semantics to linguistics. We argue that both of the initial claims are wrong: *average*-NPs are *not* ‘linguistically identical’ to other definite-NPs but instead show a number of interesting divergences, and we provide a plausible semantic account for them that is *not* ‘semantically inconsistent’ with the account afforded other definite-NPs but in fact blends quite nicely with one standard account of the semantics for NPs.

1 INTRODUCTION

We often hear sentences like (1) on the TV and read them in the newspaper:

(1) The average American has 2.3 children.

Other than questions like ‘Are they counting Americans living abroad?’, ‘Are they counting non-citizens living in America?’, ‘Are they counting adoptions?’, ‘Are they counting children as eligible to have children?’, and the like, sentences like these seem to elicit no particular puzzles of interpretation. We know that they do something like: discover the number n such that there are n Americans (citizens? living in the United States? of a certain age?) and discover the number m of children (who weren’t adopted?) and the ratio of m to n is 2.3. It might seem to go without saying that there are no Americans who have 2.3 children, although facts about every American (of the right age, domicile, etc.) go into determining the truth of (1).¹

Another meaning of ‘average American’ is illustrated in (2):

¹ So in reply to the arguments we cite below from Chomsky and Hornstein, one might point out that sentences that employ *average*-NPs are, in fact, ‘about the world’, as remarked to us by James

- (2) a. John is an average American when it comes to barbecuing.
b. Average Americans cook hamburgers on the barbecue.

Here we mean something like: on some set of features that we deem relevant, such as preparing food on a BBQ, John acts as most (or, typical) Americans do. An average American of this type is one who has typical properties. In this meaning there can be many average Americans; and not all Americans go into determining the truth of sentences like (2b). For example, wealthy Alphonse, who would never cook a hamburger on the barbecue, might be claimed not to be a counterexample to (2b) on the grounds that he is not an average American. Therefore, ‘average’ in this one of its meanings, is an adjective much as ‘tall’. If ‘American’ is true of all Americans (designates the set of Americans), then ‘average American’ is true of all the average Americans (designates the set of average Americans—a subset of the Americans) much as ‘tall American’ is true of all the tall Americans (also a subset of the Americans). That is, in this meaning ‘average’ is a function with sets as arguments that yields as a value some subset of the argument, and the resulting phrase is a CN-type meaning, in Montague Grammar terms. We can differentiate this ‘concrete reading’ (so-called because the CN is true of concrete individuals) from the ‘abstract reading’ (where it is only true of something abstract, if it is true of anything at all) in a variety of ways. For example, the abstract sense of ‘average’ may not readily appear in a noun phrase with any other adjectives to its left, whereas the concrete sense appears naturally in such contexts. Further, modification of the adjective with words like ‘very’ or ‘highly’ or ‘somewhat’ is much more natural with the concrete sense than the abstract one:

- (3) a. John is an intelligent average American
b. ?The intelligent average American family has 2.3 children.
- (4) a. Highly/somewhat/very average Americans cook hamburgers on the barbecue.
b. ?The highly/somewhat/very average American family has 2.3 children

Since in general, for any set X, there are numerous exemplars of being an average X, in the concrete sense just identified, it is rare for the

Higginbotham. If, for instance, you know that the average family has 2.3 children, and you know the number of families, then you can calculate the number of children there are, which is a fact determined by the world. Saying this, however, does not give an actual analysis that sheds light on the semantics of such constructions... a task we take as central to a full defence of semantics against the argument (see Section 2).

definite ‘*the average X*’ to employ ‘average’ in this particular concrete sense. Instead, such definite *average*-NPs usually employ ‘average’ with the meaning used in sentence (1), the abstract sense. However, such examples can be constructed, as in

- (5) a. Of those people barbecuing next door, I want to talk to the average American, not the wealthy one.
 b. Of those people barbecuing next door, I want to talk to the average American, not the German.

in which the definite NP employed picks out a concrete individual. More often, however, the concrete sense is use quantificationally, while the abstract sense is nearly always expressed with a definite determiner in the singular.

We call NPs of the sorts in (1) and (2) *average*-NPs; and we subdivide them into the abstract and concrete types corresponding to (1) and (2) respectively.² Our focus here is entirely on the abstract sense. *Average*-NPs have not been discussed much in the linguistic or formal semantic literature, nor is the distinction between abstract and concrete *average*-NPs widely acknowledged. (But see Higginbotham 1985: 586–587; Stanley 2002: Appendix).

2 AN OFT-USED ARGUMENT

A certain argument has divided that group of researchers which contains theoretical and philosophically-interested linguists and semantically-oriented philosophers. This group can be cut into two by their reaction to the argument: there are those who think it is decisive, while the others dismiss it out of hand as embodying a fundamental confusion. The particular form of this argument in which we are interested is often associated with Noam Chomsky, but it is not only Chomsky who has put it forward nor is it only those who are directly influenced by Chomsky who have put it forward; the general argument can be found in classic philosophical writings. Here is a particularly clear statement of the basic argument.

“No one wishes to claim that there are objects that are average men in any meaningful sense. . . If this is the case, it would appear that several currently influential approaches

² *Average*-NPs also seem semantically related to generic-NPs and other similar constructions (see Krifka *et al.* 1995). Such NPs as *The average American*, *The typical dog*, *The normal student*, *The frequent flyer*, and the like have something in common about how their meaning is related to instances. But consideration of this would take us too far afield.

to the semantics of natural language are misconceived. For example, programs of the Davidsonian variety... will inevitably construe quantification in natural language objectively, and reference as a relation between the term and an object... If the above observations are correct, then for many cases, how meaning is determined has nothing at all to do with notions like 'object', 'truth', or a 'model'. The key notions are not semantic, but syntactic or pragmatic. The semantic notions turn out to have no role in generalizations that explain how meaning functions in natural language." (Hornstein 1984: 58)

As we remarked, (Chomsky 1995: 29) has given this general argument form using various examples in addition to 'average man' in making the point.

"If I say one of the things that concerns me is the average man and his foibles or Joe Sixpack's priorities or the inner track that Raytheon has on the latest missile contract, does it follow that I believe that the actual world, or some mental model of mind, is constituted of such entities as the average man, Joe Sixpack, priorities, and inner tracks?"

In 1981 (324) he puts the point

"If I say the flaw in the argument is obvious, but it escaped John's attention, I am not committed to the absurd view that among the things in the world are flaws, one of them in the argument at question. Nevertheless, the NP *the flaw in the argument* behaves in all relevant respects in the manner of the truly referential expression *the coat in the closet*..."

And in 1986 (44–45) the argument goes thus:

"I think that much of this work [that Chomsky is recommending—GNC/FJP] is not semantics at all, if by semantics we mean the study of the relation between language and the world—in particular, the study of truth and reference. Rather, this work deals with ... levels of mental representation, including representations of syntactic and lexical form and others called 'models' or 'pictures' or 'discourse representation' or 'situations', or the like... If I say 'John thinks that he is intelligent', *he* may refer to John, but not if I say 'He thinks John is intelligent'. We can account

for such facts by a theory of the structural configurations in which a pronoun can acquire its reference. . . The same principles, however, apply to such sentences as ‘The average man thinks he is intelligent’, ‘He thinks that the average man is intelligent’ (or ‘John Doe thinks that he is intelligent’, where ‘John Doe’ was introduced as a designation for the average man). But no one assumes that there is an entity, the average man (or John Doe), to which the pronoun is permitted to refer in one but not the other case. If I say ‘John took a look at him, but it was too brief to permit a positive identification’, *it* can refer to the look that John took; . . . although no one believes that there are looks that a person can take, to one of which the pronoun *it* . . . refers.

Even if we restrict ourselves to the context ‘there are. . .’ we can hardly assume that there are entities in the world, or in the world as we believe it to be, that correspond to the terms that appear (‘there are chances that are too risky to take’, ‘there are opportunities that it should not be passed up’, etc.).

One can speak of ‘reference’ or ‘co-reference’ with some intelligibility if one postulates a domain of mental objects associated with formal entities of language by a relation with many of the properties of reference; but all of this is internal to the theory of mental representation; it is a form of syntax.”

While the citations here have been to Chomsky and like-minded colleagues, essentially the same argument has been used by a wide range of linguists who agree with Chomsky on little else (e.g. Fauconnier, Jackendoff, Lakoff, Langacker, Talmy), as well as by numerous philosophers (e.g. Cavell, Putnam, Rorty) and literary theorists (e.g. Derrida, Eco). The general issue as proposed in these arguments it is that there are many expressions which we use in language that clearly do not have any correlate in ‘reality’. And so therefore any theory of language which has a component that requires all expressions to ‘hook up with the reality’ is doomed to failure. (It is doomed to failure under even weaker assumptions: if it just requires *some* expressions to ‘hook up with reality’, requires syntactically identical expressions to be treated identically in this regard, and believes that the expressions we have been considering are syntactically identical to these ones that ‘hook up with reality’). Since, formal semantics, of whatever stripe, is

(or contains) such a component, and therefore should be abandoned as fundamentally misguided. The nub of the matter is that in ordinary language there are some expressions which for all the world function in language exactly like so-called referring expressions, and yet they blatantly do not refer; 'the average American' is a paradigm case of this. Since these expressions are linguistically indistinguishable from the referring expressions, they should be treated identically. But since the one group clearly does not 'hook up with the world', then the others must not either. And thus (formal) semantics is bankrupt. This is related to familiar puzzles in philosophy, where much has been made of examples like *for John's sake*, *the flaw in the argument*, *the golden mountain*, *the round square*, *the possible fat man in the doorway*, and the like.

In the literature, responses to this argument have taken four forms:

- A. Semantics is all in the head
- B. Semantics is all in our social practices
- C. Bite the bullet
- D. Paraphrase away

Chomsky's view of this argument is that it justifies the 'all in the head' option, with some fancy footwork concerning how 'semantics' is really a sort of relation between a 'domain of mental objects' on the one hand and a (mental) set of 'formal entities of language' on the other. As Chomsky has put the point (rather more bluntly than others who have taken this route) this is not semantics in any standard sense of the word, and is more properly called 'syntax'. Chomsky apparently sees no need for any semantics other than this 'syntax'.³

Chomsky's position in this matter has not met with approval amongst formal semanticists, who instead take one of the other three attitudes when faced with these examples. We do not here wish to engage those who opt for (B), the 'Wittgensteinian', 'embodied', 'radical interpretation' view that semantics is all a matter of social practices. Instead we will briefly discuss reactions (C) and (D), the 'bite the bullet' option and the 'paraphrase away' option. Overall, our

³ In his discussion of Chomsky's attitude toward semantics, Ludlow (forthcoming) distinguishes three 'conceptions' of semantics. R_0 is model-theoretic semantics ('especially that grounded in Montague'), seen as a relation between a syntactically specified language and a mental model. (Stated this way R_0 also includes the old Katz & Fodor (1963) stuff that Lewis (1972) called 'mentalese' as well as Fodor's more recent conceptions of a Language of Thought.) Ludlow cites LePore (1983) and Higginbotham (1990) for reasons that 'no one adopts the R_0 style of semantics any more'. An R_1 semantics is one where 'a term itself refers to the world' and an R_2 semantics is one according to which 'a speaker uses a term to refer to certain aspects of the world under certain contextual conditions'. Ludlow thinks that R_2 is 'the view that most philosophers hold', and he takes Chomsky however to be arguing against R_1 -style semantic theories.

view is that the ‘bite the bullet’ option, according to which terms like ‘the average American’ ‘really’ do refer to something—perhaps to something ‘abstract’, perhaps to something ‘fictional’, perhaps to something ‘non-existent’—can be made a bit more palatable than a first glance might suggest, and we will mention some ways that such theorists might wish to explore. However, we think that option (C) is probably as problematic as the Chomskian mentalistic option even with the ameliorations that we will mention. So, while it perhaps is not so outrageous a theory as the Chomsky argument might have us believe (the theory need not say that ‘the average American’ ‘really’ refers to something, and that the world ‘really’ contains the average American. . . but that it is invisible and does not enter into the causal order of things because it is abstract or fictional or non-existent, even though it ‘subsists’), it nevertheless remains a theory that in the present context is too weak a response to the charges cited above. Or at least, it should not be considered as long as there is another reasonable contender; and according to Chomsky, the mentalistic theory *is* on the table.

The final type of response is what we call the ‘paraphrase away’ option. In this type of analysis, the syntactic structure of a sentence is suggested to be something other than what it might appear to be, and that if one appropriately readjusts the syntactic structure one finds that the ontological objections are no longer valid. Consider, for instance, the suggestion by (Quine 1960§44) that belief-reports such as (6a) do not have the ‘logical form’ of a relation between a subject and a proposition, but instead are to be analysed as a simple predication of a monadic predicate being true of the subject, as indicated in (6b) with the indivisible monadic predicate ‘believes-Cicero-denounced-Cateline’. Or consider the ‘adverbial theory of perception’ as advocated by Chisholm (1957, 1963) and Scheffler (1964), in which sentences like (6c,d) are analysed as (6e,f):

- (6) a. Tom believes Cicero denounced Cateline.
 b. Tom believes-Cicero-denounced-Cateline.
 c. John is seeing green.
 d. John is seeing a green tomato.
 e. John is seeing green-ly/ in a green manner.
 f. John is green-ly seeing a tomato/ is seeing a tomato as green.

The apparent embedded sentence in (6a) is analysed as a (semantically) unanalysable part of the ‘absolute general term’ (as Quine puts it, p. 216). One way to understand this recommendation is to take the embedded sentence as expressing a *modification of the method of believing*

that Tom is manifesting. That is, to take the embedded sentence as an adverbial modification of 'believes'. Philosophers who share Quine's ontological prejudices take this adverbial modification as a way to avoid any commitment to propositions. In a similar manner, although 'green' is a noun in (6c), it gets paraphrased adverbially; and the adjectival 'green' in (6d) is also analysed as a kind of adverbial modification of seeing. According to this view, taken by philosophers who have certain sorts of ontological scruples, we do not 'really' see colours (for example) but we see coloured-ly. ('Colours don't really exist; what exists are objects that are coloured'). Similar comments hold for other perceptual features of vision, as well as for other senses. Only with items that are ontologically independent can we be perceptually acquainted, according to this view. Taken as a piece of linguistic advice, this view holds that certain nouns and adjectives 'really' are not what they appear, and should be assimilated to some other grammatical category. So, in the present case we replace an NP adverbially, as a type of perception; we replace a modified noun by an adverbial modification of perception. And in so doing we rearrange the syntax in such a way as to avoid committing ourselves to (possibly) unwanted ontological categories.

The 'paraphrase away' option is one, we wish to emphasize, that has a great deal of historical precedent, and has contributed greatly to the development of linguistic semantics and the philosophy of language. For example, Davidson's (1967) highly influential analysis of event sentences in which adverbial modifiers are treated as predicates applying to events has had unquestionably contributed greatly to the study of language; similarly, his (1968) paratactic analysis of belief sentences illuminated the topic of propositional attitudes. Russell's theory of descriptions likewise falls into this camp. In linguistics, many (though certainly not all) of the analyses of word meanings and sentence meanings in the Generative Semantics tradition were based on paraphrase, and there seems little doubt that this has likewise contributed much to the subsequent study of semantics and syntax.

Despite the attractiveness of the paraphrase option generally and of these results in particular, this methodology nonetheless has the drawback of not being able to say exactly *when* one should paraphrase in one way versus when one should paraphrase in the other way. And in the present example this seems especially significant, because the decision seems to rest exactly on what one's metaphysical presumptions are, rather than the dependency being in the opposite direction. It also has the drawback of not faithfully honouring the grammatical structure of the sentences under consideration, and of constructing the 'logical

form' as a matter of non-syntactically based preconception on the part of the theorist.

Suppose we therefore shun the option of invoking the paraphrase technique and continue to treat *average* noun phrases as ordinary NPs. What then is the alternative to viewing these NPs as referential? One obvious alternative is that perhaps they are quantified NPs rather than referential NPs. An example of this sort of view would be Russell's treatment of definite descriptions, where the apparently referential phrase 'the such-and-so' is taken to contain a hidden quantifier. (At least according to some accounts, e.g. Neale (1990), where this is taken to be a quantified NP. Others deny that there is *any* NP at all in the result since these are 'incomplete symbols' that are to be eliminated. We will not try to adjudicate this matter.) A different possibility would be to question the view that every NP is either a quantificational NP or a referential NP. That is, perhaps we should question the idea that there are only two types of NPs. This will be the route we suggest. We will claim that at least for some *average*-NPs the proper analysis is that their semantic value is non-quantificational and non-referential. And so we claim that there is a category of non-quantificational, non-referential NPs—which we call 'encuneral NPs'.⁴

Our purpose here is to consider the types of analyses of the abstract *average*-NPs that could be made by someone who is committed to the usefulness of semantic theory. It is certainly true that one who, like Chomsky and Hornstein, is antecedently convinced of the futility of formal semantics will not find in these analyses an absolutely convincing reason to re-instate the field. But certainly there needs to be *some* semantically plausible analysis available before any progress at all can be made in overcoming their sceptical position. To this end we will consider, and reject, the one suggestion that has been made in the literature—in the 'paraphrase away' tradition—questioning its adequacy as an analysis of the phenomena surrounding *average*-NPs. We follow this with a quite different proposal, an 'encuneral' NP analysis, which we think sets out a firmer foundation for disposing of the anti-semantics position that *average*-NPs are enlisted in aid of. At the end of the paper, we return to the question of whether, in fact, the 'bite the bullet' approach might be not so objectionable after all.

⁴ 'Non-quantificational, non-referential NPs' is abbreviated NQNR NPs, and it is pronounced as if it were spelled 'encuneral NPs'. Further exploration into this notion is in Hoffweber and Pelletier (forthcoming). For present purposes, we are taking the analysis of definites and indefinites in the 'DRT' tradition (Kamp 1981; Heim 1984) as quantified, and not as a third category. However, this assumption is made for merely expository purposes.

3 THE ADVERBIAL APPROACH

One analysis that has been given to meet this ontological scepticism, and which accommodates the abstract sense of *average*-NPs, is due to Higginbotham (1985). According to Higginbotham (1985: 587) the appropriate analysis of the abstract sense of sentences employing *average*-NPs should be adverbial: (7a) is to be paraphrased as (7b), (7c), or (7d):

- (7) a. The average American has 2.3 children
 b. Americans on average have 2.3 children
 c. Americans have on average 2.3 children
 d. Americans have 2.3 children, on average

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. Nor does Higginbotham provide any further description of how this analysis is to proceed; his interest in this article was instead merely to show that there was *some* analysis that might overcome the Chomsky/Hornstein objections, and for this purpose the analysis is very suggestive. While in the "paraphrase away" tradition, Higginbotham's analysis has the merit of being syntactically plausible enough (Larson 1998), so the analysis does not require different interpretations for similar structures based purely upon prior ontological considerations.

One basic problem with the adverbial suggestion concerns sentences with multiple *average*-NPs. (8a) should be analysed as something like (8b) or (8c), if an adverbial analysis is correct:

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

Yet the proposed analysis seems instead to give us something like:

- (8) d. ??Americans, on average, know little about Mexicans, on average

which does not seem to be a correct analysis. And in some multiple *average*-NP sentences, not all occurrences can be eliminated no matter how clever the paraphrase. Consider (9a) and its plausible paraphrase (9b), and contrast it with the implausible paraphrase (9c). Note furthermore that (9d) says something different from (9a): (9d) says that women do not admire *men*, whereas (9a) only says that, on average, women do not admire the *average* man.

-
- (9) a. The average woman does not admire the average man
 b. On average, women do not admire the average man
 c. ?? The average woman does not admire men, on average
 d. On average, women do not admire men

A further problem with the adverbial suggestion has to do with *average*-NPs that do not seem at all equivalent to the adverbial paraphrase; indeed, the adverbial paraphrase often seems not to make sense.

- (10) a. The average tiger hunts at night.
 b. ??Tigers on average hunt at night
 c. ??Tigers hunt on average at night
 d. ??Tigers hunt at night, on average
- (11) a. The average Russian man wears a hat
 b. ??Russian men on average wear a hat
 c. ??Russian men wear on average a hat
 d. ??Russian men wear a hat, on average
- (12) a. The average American owns a car
 b. ?? Americans on average own a car
 c. ?? Americans own on average a car
 d. ?? Americans own a car, on average
- (13) a. The average 50-year-old American male is worried about his waistline
 b. ?? 50-year-old American males on average are worried about their waistlines
 c. ?? 50-year-old American males are worried on average about their waistlines
 d. ?? 50-year-old American males are worried about their waistlines, on average

One other possible problem with Higginbotham's particular formulation of the adverbial suggestion is the appearance of a bare plural rather than an indefinite singular in the proposed paraphrase. Plurals of (non-collective) nouns, such as 'gangster', can appear with collective predicates (such as 'scatter'), but the singular version may not. However, abstract *average*-NPs do not allow collective predicates:

- (14) a. *The average gangster scatters when the police arrive
 b. *A gangster scatters when the police arrive

- c. Gangsters scatter when the police arrive

Average-NPs can take collective predicates, but only with an appropriate collective noun such as ‘army’. when the singular form itself otherwise allows collective predicates:

- d. The average army disperses when defeated.
e. The army (\star gangster) dispersed after the defeat.

So abstract *average*-NPs seem to function as genuine singulars, and not as plurals. This may be an objection easily circumvented in the spirit of Higginbotham’s analysis, but it does at least show that the analysis requires, at the very least, some adjustment of detail.

To sum up, at this point we have expressed some reservations about the adequacy of the adverbial analysis, at least enough to justify exploring other possibilities. Other objections will be put forward in the context of our actual proposal in Section 5. In the next section we make some brief comments on peculiarities in the semantics of ‘average’. In the following section we give an ‘encuneral NP’ analysis, inspired by comments of Bacon (1973), which we present as pointing in the direction of an adequate response. This analysis takes the denotation of *average*-NPs to be just a set of (non-quantificationally-defined) properties.

4 SOME OBSERVATIONS AND QUESTIONS ABOUT AVERAGING

What examples like (10)–(13) show us is that we need a better understanding of just what properties can be ‘averaged’, and how that ‘averaging’ is done. Consider the truth,

- (15) a. The average American is right-handed

In what sense are we ‘averaging’ here? After all, about 15% are left-handed, so why do not we have

- b. The average American is 85% right-handed

instead? It is not correct to say that it is because no one can be 85% right-handed, or that ‘it does not make any sense’ to think of the property 85% right-handed, for the very same thing could have been said about (1), (some variant of) which is paradigmatically true. Why is it weird to say

- (16) a. ??The average couch has 5.2 legs
 b. ??How far from the sun is the average planet
 c. ??How large is the average planet
 d. ??What colour is the average bird
 e. ??Is the average bird multicoloured

Another feature worth noticing is a certain ambiguity in averaging:

- (17) a. The average trouser size for men is 36.
 b. The average trouser size for men is 40.

(17a) could be true from the trouser manufacturers' point of view: take the average of the sizes of all the individual pairs of trousers they produce and sell, and it will be 36 inches; yet (17b) may be true from the individuals' point of view: take all the men and average their trouser sizes, and it comes out to be 40. (This might occur if people with large waists buy fewer pairs of trouser.) Any analysis of 'average' will need some sort of variable or parameter whose value can be altered so as to yield distinct meanings, in this and perhaps many other instances.

Another consideration is anaphora with *average*-NPs, where a definite pronoun can apparently refer to the denotation of the *average*-NP. This formed part of the initial sceptical argument by showing that *average*-NPs were indistinguishable from 'ordinary referential NPs'. The adverbial analysis seems unable to account for these facts.

- (18) a. The average tiger is 10 feet long. *It* hunts at night, seeking small mammals.
 b. The average working man punches a time clock. *He* does manual labour for an employer.

Note that on the adverbial analysis, since the second sentence contains no instance of 'average', one should find no 'averaged' reading for the sentence containing the pronoun. However, such a reading does appear.

Another problem, this one more 'logical' than linguistic, is that analyses of *average*-NPs have to take care to avoid a version of the 'third man argument'. Most of us think that (19a) is true and that its oddness, if any, is due to its obviousness; but even those who think that (19a) is somehow deficient in meaning agree that (19b) is true.

- (19) a. The average American is an American
 b. The average American is a human.

But now consider the census that was recently taken. Whatever the difficulties with counting large populations, no one argues that one more American, the average American, should be added to the total because of the truth of (19a)! If we were instead counting all humanity, no one would argue that another person should be added, the average American (nor one for the average Canadian, the average Mexican, . . .) because of the truth of (19b). The point is that, although the average American is an American, and the census is intended to enumerate all Americans, it should not count the average American. This logical puzzle needs to be dealt with by any account of the abstract notion of ‘average’.

We have already remarked on the fact that the average American is right handed, and not 85% right handed. So sometimes when the average of a property is computed for some *average*-NPs, we do not end up at the statistical norm. It should also be mentioned that sometimes there simply is no average computed and ascribed to ‘the average X’, even though each individual X has a value for that property. Neither (20a) nor (20b) are true, even though each American is either male or female;⁵ on the other hand (20c) is (plausibly) true.

- (20) a. The average American is female
- b. The average American is male
- c. The average elementary school teacher in America is female

Just how peculiar this can seem is illustrated by the following pair:

- (21) a. The average couch is 2.1 meters long
- b. ?? The average couch has 5.2 legs

Note once again that it is not appropriate to account for the difference between (21a) and (21b) on the grounds that there could be a 2.1 meter long couch, but not a couch with 5.2 legs, for there likewise cannot be an American family with 2.3 children.

Many of these puzzles could receive at least a first attempt at a solution if only we had a better understanding of the process of ‘averaging’. Why is it, for example, that certain properties do form natural averages but others do not? And can this be captured in any systematic way? Why is it that when some properties form averages, the result is not the intuitive average (e.g. ‘handedness’)? In pursuit of understanding these phenomena, we could investigate properties X, of which it makes sense to average the property X (e.g. size), but such that

⁵ If it is true that there are three or four or five genders for people, then the example can be appropriately altered to accommodate this.

‘the average X’ doesn’t make sense along some dimension (e.g. size of the average planet). And we want to know the conditions under which ‘the average X’ means something different from ‘Xs, on average’. We have no completely clear answers to such questions at the moment, but note them as facts that any fully adequate analysis would need to give some account of. At the moment, however, we see no reason to think that any of the facts above favour one line of analysis over another, and to these considerations we now return.⁶

5 AN ENCUNERAL ANALYSIS OF *AVERAGE*-NPs

We have looked at one possibility for a ‘paraphrase away’ analysis of *average*-NPs, the adverbial analysis. In this section we wish to explore an analysis that treats the semantic value of *average*-NPs as being just sets of properties.

Saying that their semantic *value* is a set is different from saying that *average*-NPs *refer* to such a set. There may be some technical notion, ‘denotation’ perhaps, according to which the semantic value of a linguistic item is that item’s denotation. But we are here using ‘reference’ in a normal, ordinary, and non-technical sense—and it is in this sense that we are claiming that *average*-NPs do not refer—either to an entity or to a set.⁷ Although many noun phrases, chiefly quantified noun phrases, are commonly analysed as sets of properties (that is, as generalized quantifiers), we claim that *average*-NPs are not quantificational NPs either, not in the normal, ordinary, and non-technical sense where a quantified NP picks out a domain of individuals and introduces into the evaluation of a sentence some (sub)part of that domain (the all, most, some, none, etc. subpart).

As we shall see, on our analysis the definite article contributes nothing to the meaning of the noun phrase (it is not a definite description) because the output of combining of ‘average’ with the common noun it modifies will be identical to the meaning of the whole noun phrase with the determiner ‘the’ added. It shares this feature with

⁶ Another matter that deserves further consideration is the relation(s) between *average*-NPs and other constructions such as:


- (a) The occasional sailor drops in.
- (b) The exceptional singer will have a four-octave vocal range.
- (c) The ideal woman, to the Greeks, was strong.
- (d) The typical untrained dog comes only for food.
- (e) The usual people met at the usual places

(See Larson 1985, for further examples and discussion.)

⁷ Perhaps the phrase ‘the set of properties expressed by *the average American*’ is a referring expression, referring to a certain set of properties. But *the average American* itself does not refer to this set, at least not on our analysis (nor on most analyses).

Higginbotham's adverbial analysis. This may raise some questions about the handling of anaphora, which we address below, but if successful it becomes as ontologically unobjectionable as the adverbial analysis.⁸ This idea for treating *average*-NPs was first noted, to our knowledge, by Bacon (1973), but a detailed analysis was not presented there.

Let us start with a straightforward example and use it as a scaffolding on which to build when dealing with less straightforward cases. I have a class of 75 students, and I have given them an exam and graded it, each grade being an integer. I then seek to find the average test score (i.e. to determine the denotation of 'the average test score'): I total the test scores and divide by the number of tests. Suppose the average exam score comes out to be 83.5 by my calculations. Essentially what we want to do now is to reproduce this calculation in the semantics of the phrase 'the average test score' (treating 'test score' as a compound noun). We are going to assume that the meaning is calculated according to the following syntactic arrangement.



Let us start by considering the phrase 'test score' itself. The most straightforward thing to try is to treat it as a set of numbers:

$$\{x : \llbracket \text{test-score} \rrbracket(x)\}$$

This immediately fails, of course, since the set $\{a, a\}$ is identical to the set $\{a\}$, so we cannot represent the fact that several students might have attained the same score. Backing off slightly, then, we might resort to representing the numbers as an ordered n -tuple (there are other possibilities, but for simplicity's sake we're taking test scores to be numbers). We indicate this by using the convention of angle brackets. In addition, the set of exams administered needs to be 'contextually determined' so the set does not include test scores from exams given

⁸ One anonymous referee remarks 'an account that hypostasizes (intensional) properties, properties of properties, and so on, is beyond good and evil already.' We are not sure whether this is or is not an endorsement of our proposal.

other by other people or to other people. Let's fudge a bit at this point and include a contextualizing parameter, c , to make this clear:

$$\{\langle x : \llbracket \text{test-score}^c \rrbracket(x) \rangle\}$$

What happens then when we add 'average'? It seems we want a formula which reproduces the calculation. That is, we need access to the sum of the exam scores and to the cardinality of the set (indicated by vertical bars in our notation) – information that we get from the context, somehow:

$$\begin{aligned} \text{average}(\{\langle x : \llbracket \text{test-score}^c \rrbracket(x) \rangle\}) &= \left\{ \sum(\{\langle x : \llbracket \text{test-score}^c \rrbracket(x) \rangle\}) \right. \\ &\quad \left. \div \left| \{\langle \llbracket \text{test-score}^c \rrbracket(x) \rangle\} \right| \right\} \end{aligned}$$

Since this calculation produces a unique number, in context, the definite article ι is added yielding a denotation, in this example, of 83.5:

$$\iota\gamma[y = \sum(\{\langle x : \llbracket \text{test score}^c \rrbracket(x) \rangle\}) \div |\{\langle x : \llbracket \text{test score}^c \rrbracket(x) \rangle\}|]$$

So long as we can deliver on the promise of c , which gives us the relevant exam, the relevant scores, and the relevant number of people who took it (something any analysis, it would appear, must also do), we have given a possible semantics here of 'the average test score'. Note that this analysis, which treats numbers as entities, yields an entity for its denotation. When we turn to NPs with other head nouns and attempt the same thing, matters do not go so smoothly, which helps to motivate an encuneral, set of properties, analysis.

Take, for instance, 'the average lion' (which is tawny, eats meat, lives about 14 years, weighs some 250 pounds, and so forth). Since the set (whether ordered or not) defined by $\{x : \llbracket \text{lion}^c \rrbracket(x)\}$ in a given context is going to be a set of lions and not a set of numbers, we need to ask, in what sense is it possible to average individuals such as lions themselves? Here, we are at something of a loss. What, for instance, is the average of Jill, Alice, Bob, Jack, and Mandy? Not only do we have no idea how to answer such a question, but it seems clear that there simply *is* no answer to it. We could give their average height, or weight, or times they have visited Paris, for example, but not what their average is *simpliciter*. We conclude that we cannot average individuals as such, but only (certain) properties of those individuals. That is, in order to deal with such examples, we need to accord 'average' some means of access to properties of individuals, at the very least.

But it cannot be simply the full set of properties had by each member of the set that is ‘averaged’: the properties need to be appropriately partitioned so we do not end up averaging, for instance, heights mixed with weights, or times the dentist is seen annually with points scored in high school basketball games. It would appear that a good start on this partitioning would be to define functions associated with properties which return numerical values—a measure. Consider the property of weight, for example. The function which would return a numerical value would be something like ‘weighs [x units]’; that function for trips to the dentist annually would be ‘sees the dentist [x times] annually’. We might push ahead with this strategy, but even at this early point we ought to pause and ask if it is sufficiently general. Essentially, what this says is that we can average only those properties which have some standard measure along some dimension expressible in numbers, or, somewhat more weakly, an adjective with a determiner position marking extent. However, it seems we can quite sensibly say things like the following:⁹

- (22) a. The average student at this university is from the Northeast.
b. The average biology major expects to attend medical school.
c. The average lion is a carnivore.
d. The average automobile is powered by petrol.

In these and many other cases the properties ascribed have no natural expression in numbered units—one is either from the Northeast, or not, expects to attend medical school, or not, and so forth. What could ‘average’ mean with respect to properties like these? It does not appear to be a simple matter of counting up and determining either majority or plurality. For instance, while we believe that the plurality of Americans are brown-eyed, and the majority of Americans are female, it is false (and maybe not even sensible) to say that the average American has brown eyes, or that the average American is female. Even completely reliable statistics on these matters (as opposed to our impressionistic beliefs) would not suddenly sway our opinion in such cases.

For want of clear intuitions on such matters, we are going to leave it at the level of basic observation that there are properties denoted by *average*-NPs which are not naturally reducible to numbers along a scale of measurement, and there are dimensions (such as eye-colour or gender, as suggested above) along which no secure judgments are available (though such dimensions might be properties of average

⁹ This shows a weakness in the account given in Stanley (2002), where it is claimed that all averaging functions (or at least ones like these) generate numerical scales.

individuals in other populations—it makes sense to say that the average Russian physician is female, for instance—a claim we have already made in discussing (20)).

It remains clear, however, that one does average the comparable properties across individuals, and does not average different properties (such as height and weight) together. We say ‘comparable’ rather than the ‘same’ properties, because we wish to deal with cases such as region of the country one is from or expected major (as in (22) above), where, for instance ‘being from the Northeast’ and ‘being from the South’, though different properties, are comparable. So we require some partition-of-properties-of-individuals function, which we will call **Part**^l. This **Part**^l function operates on the Common Noun Phrase (i.e. the noun plus modifiers) that is being ‘averaged’ and 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.¹⁰ For example, if we are computing a semantic value for ‘the average American’, then **Part**^l(*American*) will first determine what the appropriate partitions of properties for Americans 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. It will *not* construct a partition corresponding to gender or eye colour, nor to any other features for which there are not properties of the average American. Intuitively, partitions are like very ‘general properties’ that can be ‘particularized’ in many different ways. So, the ‘number of children’ partition can be particularized in all sorts of ways—‘less than a dozen’, ‘more than her mother’, ‘an odd number’, ‘exactly three’, and the like. Our idea is that the second job of **Part**^l(*American*) is this: for each partition thus computed, to pair up each instance of ‘American’ with *the most specific way* of particularizing that partition which the instance exemplifies. Molly, let us suppose, is an American, and on the ‘number of children’ partition she exemplifies all of ‘less than a dozen’, ‘more than her mother’, ‘an odd number’, and ‘exactly three’. But ‘exactly three’ is the most specific particularization that she instantiates for this partition. And so, one of the ordered pairs within the ‘number of children’ partition of **Part**^l(*American*) will be ⟨Molly, exactly 3⟩. **Part**^l(*American*) is thus a set, each member of which is itself

¹⁰ We use ordered pairs for the representation of monadic properties: $\langle a, P \rangle$ means that a exemplifies the property P . We will not here discuss how this might be extended to relational properties in general.

a set (a partition of *American*) of ordered pairs, where each ordered pair has an instance of ‘American’ as its first member and its most specific instantiation of the possible values of the partition as its second member.¹¹ We return to the issue of how these properties give rise to inferences involving *average*-NPs later in this section.

For its part, ‘average’ introduces a function we call **Ave’**, which can be seen as performing two tasks. First, it deploys the appropriate **Part’** function as described just above (appropriate for the CN argument of ‘average’). Its second task is to do some computation on the result constructed by **Part’**, to figure out the average value corresponding to each partition. Recall that this is not a ‘uniform process’: some of these partitions (weight and height, perhaps) generate their value by a simple process of summing the values found in the second position of the ordered pairs in this partition and dividing by the cardinality of that partition (alternatively, dividing by the cardinality of the CN argument to ‘average’). But others of these partitions have their average computed differently: handedness, for example, makes a binary choice of right or left, presumably based on the fact that a large majority exhibit one value. (If the split were 55%–45%, then the average American probably would not be handed . . . in the same way that the average American is not gendered.¹²) And there are many other types of ‘averaging’ that we have discussed earlier, such as those in (10)–(13) and (22). Further, as we said in discussing example (17), there is often an ambiguity as to how the ‘averaging’ is going to be done. We presume that there are different **Ave’** functions, or possibly different modes of individuation in certain cases, available for this purpose. However it is done, the point of the **Ave’** function is to generate a property in some way for each partition that was created by **Part’**. The semantics for ‘average (*N*)’, where *N* is a common noun meaning and *c* is a context variable is therefore

$$(23) \text{ a. } \text{average} (N) = \{ Q : Q \in \mathbf{Ave}'(\mathbf{Part}'\lambda P\{\langle x, P \rangle : N^c(x) \& P(x)\}) \}$$

or more simply

$$\text{b. } \text{average} (N) = \{ Q : Q \in \mathbf{Ave}'(\mathbf{Part}'(N^c)) \}$$

¹¹ It is not clear to us whether the notion of ‘a partition of a CN’ is definable in purely semantic terms (dependent on only the semantic value of the CN), or is instead a matter of ‘being a partition of that CN in a context’. If it is a purely semantic notion, it is then also not clear to us whether ‘being the most specific instantiation of a partition’ is definable in purely semantic terms, or whether it is a matter of ‘most specific in a context’.

¹² An example that might be relevant is ‘The average baseball pitcher is right-handed’, which strikes us as false given disproportionate numbers of pitchers who are left-handed.

i.e. the result is a set of properties Q such that for each partition of N , Q is the ‘average’ value on that partition.¹³ So the output of applying ‘average’ to a noun is not a set of individuals, but rather a set of properties.

The next question in line is to ask what work the definite article has to do when added. The analogical strategy based on ‘the’ picking out a unique individual from a set of individuals won’t do here, as we want all the average properties that have been calculated to remain: we do not need to identify a single ‘most relevant’ property as the denotation of the NP in context. In fact, things would work quite nicely if the definite article were simply not there—we already have a noun phrase meaning which insofar as we can see could do the work itself without the need for an additional definite article. So on this analysis, the cleanest way to treat the definite determiner would be to accord it the identity mapping, i.e. it makes no semantic contribution to the whole. Is this an unhappy outcome? There is precedent for positing expletive determiners in the syntax (e.g. Vergnaud & Zubizarreta (1992); Longobardi (1994)), and this could well be the case here. In any event, it would seem a consequence of the analysis presented.

There is a positive side to things, as well. Since the adjective ‘average’ returns not a set of individuals, but a set of properties, it would seem to be unlike other more normal adjectives which operate on and return just sets of individuals. Thus, other adjectives should not be able to operate on the output of this abstract sense of ‘average’; and this seems correct. If we assume, as is highly plausible, that the leftmost adjective is the last semantically added, we find, as noted above, that ‘average’ must be leftmost. If an adjective is to the left of ‘average’, it must be read appositively, as in (24b) (where it is the average American who is fat, unlike (24a)), or else one must interpret ‘average’ as picking out a set of average-type individuals. In contrasting (24a,b), this seems correct.

¹³ Technically speaking, this is an ‘intensional’ construction because it treats the common noun N (or, the N^c) as a CN-meaning, rather than as a syntactic unit. The **Part**’ function in (23a) works on the meaning of N^c , and therefore it is only if two different phrases have the same meaning (in context), that **Part**’ would yield the same partitioning (and then **Ave**’ would yield the same set of average properties). We could make it ‘fully extensional’ if we defined **Part**’ to operate on the set of instances of N^c , constructing the partitions by considering all the properties that every one of the instances of N^c had. In this case, if two common nouns were co-extensive (in context), then they would generate the same partitions. It seems to us that the extensional version is wrong; it sometimes also seems to us that even the intensional version is wrong, and that it might be possible for two different terms X and Y to have the same meaning (in context) and nonetheless for ‘the average X ’ and ‘the average Y ’ to be different because of a difference in the types of partitions that are relevant to X and Y . But we won’t follow up this line of thought now.

- (24) a. The average fat American eats too much sugar.
b. The fat average American eats too much sugar.

Adjectives like ‘former’ that are purely attributive, and cannot be read appositively, simply cannot appear to the left of ‘average’ in the abstract sense:

- (25) a. The average former U.S. Senator has 6.8 grandchildren.
b. ??The former average U.S. Senator has 6.8 grandchildren.

In (25b) ‘average’ must be read as picking out a set of individuals, and so some particular individual must have 6.8 grandchildren, which is not possible. This seems a correct and salutary prediction of this analysis.

A further possibly quite good result is that it offers an understanding of how we might analyse such apparent names as ‘Joe Sixpack’, ‘John Doe’, ‘John Q. Public’, and the like. These, too, would be sets of properties without referential value, there being no particular individual serving as the name’s referent. Note that the ‘adverbial analysis’ could not in any straightforward sense be invoked to give a similar account.

With these thoughts in mind, let us return to our initial motivating example (1):

- (1) The average American has 2.3 children

The construction just mentioned proceeds by finding, first, the partitions of *American*. As we mentioned, these will be such ‘general properties’ as ‘height’, ‘weight’, ‘number of children’, ‘handedness’, ‘food preference’, and so forth. Exactly which properties form the partitions is a mystery not solved here; but somehow we speakers, relying on much world-knowledge, are able to determine them. In the computation of the truth-value of (1), each partition is filled up with all the ordered pairs such that: the first member is chosen from the set of Americans while the second member is the ‘most specific particularization’ exemplified by that American for that partition. For each of these partitions we construct the appropriate ‘average’, giving us a set of (average) properties; and the set thus constructed is the semantic value of ‘the average American’. If ‘has 2.3 children’ is one of these (average) properties, then (1) is true; if some other value incompatible with ‘has 2.3 children’ is one of the properties, then (1) is false. If there is no value for the ‘has *n* children’ partition—that is, if there is no such partition—then (1) is semantically anomalous in the way we have discussed for various examples. (1) is true just in case ‘has 2.3 children’ is in the set of properties that is the semantic value of ‘the average American’. Speaking more generally (and ignoring niceties of use/mention):

(26) ‘The average X is P ’ is true (in a context c) iff $\llbracket P \rrbracket \in \mathbf{Ave}'(\mathbf{Part}(\llbracket X^c \rrbracket))$

One notices a similarity here with treatments in which proper nouns have as their semantic value the ‘property set’ of the item denoted by the proper noun (i.e. the semantic value of a proper noun is the set of all the properties of the item denoted). The semantic value of ‘Hillary Clinton’, for example, is the set of all of Hillary Clinton’s properties. The sentence ‘Hillary Clinton lives in New York’ is true just in case ‘lives in New York’ is one of the properties in the set corresponding to ‘Hillary Clinton’. Generally, where PN is a proper noun,

(27) ‘PN is P ’ is true (in a context c) iff $\llbracket P \rrbracket \in \llbracket PN \rrbracket$

The only difference between (26) and (27) is the set indicated by the subject term of the sentence in question: in (27) it is the set of all the PNs properties, in (26) it is the set of the properties of the average X . The difference between proper nouns (in this analysis) and our analysis of *average*-NPs is merely this: proper nouns have an ‘individual sublimation’ (to use a Montague Grammar term, see Dowty *et al.* 1981: 220ff) whereas *average*-NPs will not, due to the way the properties are constructed.¹⁴ The way the properties are constructed also appears to be very uncharacteristic of quantificationally-determined property sets (supposing one attempts to treat ‘average’ as expressing a determiner-type meaning). Intuitively, it simply does not matter how many or what portion of the individuals in the domain have the property of the predicate ascribed to the average. For instance, if every Harvard undergraduate has an SAT score of 1400, or half of them, or 3%, or 75% percent, or none at all, this alone makes no decisive difference whatsoever regarding whether the *average* Harvard undergrad has a score of 1400. Such facts, however, are *decisive* in the case of determiners such as ‘most’, ‘a few’, ‘all’, ‘no’, and so forth. The instances where a property has an empty extension also work differently for *average*-NPs than for quantified NPs. Suppose, as seems reasonable, that the two verb phrases ‘has 2.3 children’ and ‘has 8.7 children’ are both (necessarily) extensionally empty; that is, $\llbracket \text{has 2.3 children} \rrbracket = \llbracket \text{has 8.7 children} \rrbracket = \emptyset$ under all conditions. So, ‘all/no/some Americans have 2.3 children’ has the same truth-value as ‘all/no/some Americans have 8.7 children’. *Average*-NPs obviously do not work similarly. If they

¹⁴ Is it possible for there to be an individual X who just happens to have precisely the properties constructed for the average X ? We doubt it. We think there will always be some cases of constructed properties that cannot be had by any individual... as the property of ‘number of children’ is for Americans. But it is logically possible.

did, then *any* predicate with the denotation of the empty set should be uniformly true or be uniformly false. Given that there is some particular value of n which makes the sentence ‘The average American has n children’ true (we have been using the figure of 2.3), it would then be the case that the average American also has 2.4 children, 8.7 children, π children, is a married bachelor, is a round square, etc.

Since, however, we are not making use of known quantificational structures, we do owe some account of how entailments among the properties of an *average*-NP are calculated. For instance, if the average Californian owns a pet fish, then it follows that the average Californian owns a fish and owns a pet. And we certainly would want to claim that ‘the average American is an American’ is tautologous. So, how do we avoid, on this account, the possibility that an *average*-NP is just a completely unstructured ‘jumble’ of properties? The answer, we expect, lies in the precise details of how ‘averaging’ occurs, but we have only presented an outline which leaves this question only partially addressed. What we can observe, however, is that the properties in the set behave (upward) monotonically: that if the extension of a property P in the set defined by some *average*-NP is a subset of the extension of another property P' (to speak of properties as sets, for the time being), then P' is also in that set of properties. Without going into the how and why, we note this as a fact about the overall structure of the set of properties defined by *average*-NPs, and any more particular analysis must achieve this result.

The observation that the semantics of *average*-NPs conforms with the general analysis of NPs in Montague Semantics, where they are generally sets of properties, leads to the following interesting conclusions. Implications like

The average American has a pet dog \models The average American has
a pet
The average American drinks whisky in the evening \models The
average American drinks alcoholic beverages in the evening
The average American has a pet dog. The average American
drinks whisky in the evening. \models The average American has
a pet dog and drinks whisky in the evening.

suggest that *average*-NPs might be filters.¹⁵ Although we have suggested above that only the most specific instances of properties go into determining the values of *average*-NPs, it may be that the full

¹⁵ A set of sets P is a filter iff (i) $A \in P$ and $A \subseteq A'$ implies $A' \in P$, and (ii) $A \in P$ and $B \in P$ implies $(A \cap B) \in P$.

range of true sentences about ‘the average *N*’ might contain some less informative statements that arise systematically from the most informative ones. With regards to the above inferences, it might be a matter of conversational informativeness that we attribute ‘owning pet dogs’ to the average American rather than ‘having a pet’. The analysis of *average*-NPs proposed in this paper opens up this sort of further analysis as well as possible attempts to define the set of concrete average Americans in terms of properties of the (abstract) average American. All these sorts of connections form avenues for future research.¹⁶

The encuneral analysis, where some expressions are taken to describe a set of properties, which we have proposed here for *average*-NPs (and, potentially, other types of NPs) questions a common assumption in semantics that all noun phrases are analysable either as quantificational or as referential, there being no third alternative. (See Dever 2001, for a recent discussion and a partial defence of this common assumption; see also Neale 1990, 1993.) Our analysis questions this simple bifurcation of NPs, and if ultimately successful opens the way for thinking about what other types of noun phrase may fall into this third category of noun phrase meanings.

Finally, let us address the issue of anaphora that we mentioned above, as illustrated in (18), repeated here.¹⁷

- (18) a. The average tiger is 10 feet long. *It* hunts at night, seeking small mammals.
 b. The average working man punches a time clock. *He* does manual labour for an employer.

Note that such examples present a real difficulty for the ‘adverbial’ analysis discussed earlier, as the pronoun cannot plausibly incorporate the meaning of the adverb ‘on average’ into it. But these examples also pose an apparent difficulty for our encuneral analysis as well, since the canonical uses of pronouns refer back to an entity introduced in prior discourse, or deictically present in the context of the utterance. But on the encuneral analysis, there is no entity to refer to at all, and so there is an apparent problem. However, it is also generally agreed that there are other uses of pronouns that call for a different type of analysis, in which the pronouns are (often) regarded as standing not for an individual but rather a set of properties (these would be ‘E-type’

¹⁶ Our thanks to an anonymous referee for the comments that have found their way into this paragraph.

¹⁷ There is an additional reading available for the pronouns in these instances, where the pronoun refers back to the kind tigers, or working men in general, and not to the average tiger or working man. We set aside that reading.

pronouns (Evans 1977, 1980; Cooper 1979) and especially ‘pronouns of laziness’ (Karttunen 1969; Geach 1962; Jacobson 2000)). Consider an example such as (28).

(28) John wants to buy a new car. *It* must have a sun roof and leather seats.

Here, we consider the intensional reading of the first clause, in which there is no particular new car that John has in mind. In this instance, what entity does the pronoun in the next clause refer back to? It would seem to have to be a car, but of the set of all cars, none can be identified as the car referred to. Rather, a paraphrase of the pronoun—itsself in an intensional context—would be something like: ‘the car/any car that John eventually purchases (if any)’, a description that is analysable not as referring to an individual but rather as a set of properties. Given this analysis, it is plausible to think that such uses of pronouns will also give an account of such instances as we find in (18).

This line of analysis, to the extent that it is successful, maintains a ‘hookup between language and the world’ while at the same time honouring the views that *average*-NPs do not refer to anything in the world, but that they function in language exactly like terms which do refer to things in the world. As we said, the fact that we have found a possible account that does all this will not change the minds of anyone who is antecedently convinced of the futility of traditional semantics. But we believe that it should show that this argument is not conclusive and that it should stop being employed as a justification for this position.

6 COULD *AVERAGE* -NPs REFER TO ENTITIES?

The previous section was dedicated to showing that there is a line of analysis that can successfully meet the objections of those who express a scepticism about the possibility of there being any semantic analysis of *average*-NP’s that will be ontologically ‘reasonable’. But, as we mentioned earlier, not everyone shares these concerns. So, let us consider the possibility that *average*-NPs might involve a reference to an entity. This analysis, of course, walks straight into the teeth of the ontological objections, but we do find it worthy of consideration nonetheless. The positing of non-existent entities, we note, is not unheard of even in fairly recent philosophical discussion (e.g. Parsons 1980; Meile 1995).

One possibility, suggested to us in conversation with Jason Stanley (see also Stanley 2002), is that *average*-NPs do, in fact, denote some actual individual or other (though, in sentences such as (4)

above, another analysis is exceptionally called for, but not in our other examples). This suggestion would have the merits of dealing with anaphora straightforwardly, be ontologically unobjectionable, and would automatically (like the encuneral analysis above) avoid the ‘third man’ problem, at least for the non-exceptional sentences.

We do not find this particular suggestion initially promising for the following reasons. If, say,

- (29) The average Harvard undergraduate has a combined SAT score of 1400

then the denotation of ‘the average Harvard undergraduate’ would have to be one of those students who did score exactly that on the exams. Suppose there are by chance no such students. Then this sentence could not be true, which does not seem quite right to us. Further, since there is no student fitting that description, there is no average undergraduate at Harvard at all (for, if there were one, s/he would have an SAT score of 1400). And since there is no average undergraduate at Harvard, this eliminates the possibility of constructing any true sentences at all about the average Harvard undergraduate: s/he has no average weight, no average height, no average grade point, etc. But suppose then, quite contingently, that there just happened to be some Harvard student who received a 1400 SAT score; then suddenly one *could* construct such all these other sentences meaningfully! This does not seem correct.

But even if we could somehow convince ourselves that such puzzles could be overcome, we might ask ourselves what the sentences would entail under such an analysis. Suppose there happened to be just one such student with that test score. Then that person would be some particular person, e.g. Sally Jones. It should be possible to assert of ‘the average Harvard undergraduate’ any property whatsoever that Sally Jones may have. So, if Sally is a freshman, then the average Harvard undergrad (being Sally Jones) would be a freshman; if she watches no television, then the average Harvard undergrad watches no television; if she is 27 years old, then the average Harvard undergrad is 27 years old, and so forth. But it seems that the average undergrad might have received a test score of 1400, watch 2.7 hours of television per day, *and* be 20 years old, all at the same time.

Then, of course, there is the more likely situation where there is more than one Harvard student who scores 1400 on the SAT exams. Since ‘the’ presupposes uniqueness, under these circumstances, one would expect some form of presupposition failure in ‘the average Harvard undergraduate has 1400 combined SAT scores’, but we detect none. The overall problem here is that contingent facts about how

many, if any, Harvard undergraduates actually score exactly 1400, has *nothing* to do with our intuitions about whether (29) is true, or false, or odd, or entirely acceptable.

There may be sensible rejoinders to this line of thought, but we do not envision them at the moment. So, if *average*-NPs denote an entity, and it cannot be any normal, concrete entity, what sort of entity might that be? One way of softening, but perhaps not fully escaping, the ontological objections might be to employ the following line of reasoning. There are, it may appear, entities we refer to all the time that we also, by general agreement, do not believe to actually populate the world. We talk about Superman, Santa Claus, Sherlock Holmes, and other fictional entities in exactly the same way we talk about real individuals. Suppose ‘the average Harvard undergraduate’ is one of these sorts of entities. This escapes the variant of the ‘third man’ problem mentioned above (when taking a census of London in the late 1800s we would not also count Holmes and Watson, even if they are Londoners), and it allows for an entity to be anaphorically referred to by definite pronouns in subsequent discourse. It also provides a natural type of denotation for proper names like ‘Joe Sixpack’, ‘John Doe’, or ‘John Q. Public’. Note we can also say things about these entities which seem to be true or false (e.g. Sherlock lived on Baker Street, and Superman wears a cape, but Sherlock did not live in Thailand, nor does Superman wear the same get-up that Batman wears). So, the reasoning would go like this: You tell us the ontological status of Superman, something that needs to be contended with in any case, and then we’ll tell you the ontological status of the average French voter. In neither case would we be positing objectionable entities in the real world . . . at least, not unless your analysis of fictional entities is objectionable.

There is one striking similarity between fictional entities and *average*-NPs, and one striking difference. The similarity is that both types denote entities that have ‘property gaps’ in them. For instance, any individual, let us assume, is either male or female, yet the average American would seem to be neither, as we have said above. Similarly, fictional entities also exhibit property gaps: Sherlock Holmes did not have a liver weighing exactly 6 pounds, nor any other particular weight, but he did have a liver. Santa Claus is neither right-handed, nor left-handed (nor ambidextrous), but manipulates things with his hands nonetheless.

The striking difference between fictional entities and the denotations of *average*-NPs is that, while the fictional entities’ properties are simply stipulated (e.g. if Doyle had decided to make Holmes live on Downing Street instead, Holmes thereby would have), *average*-NPs

have properties which are not so stipulated but are instead determined by the character of the world. This difference may be enough for many to regard it as crucial, showing the fictional entities approach to be unpromising. But there are a few things that can be said in its defence. For one thing, *some* properties of fictional individuals are also calculated from knowledge of the real world, much as properties of *average*-NPs. For instance, we may not know exactly how tall Holmes was, but we do not believe he was about one inch tall. However, if the usual height for a man in the real world hovered around one inch, then we would ascribe this property to him. Or, if knowing that Santa lives at the North Pole, we calculate that he lives in a cold, uninhabited place, this calculation is founded on the truth in the actual world that the North Pole is cold and uninhabited; were the North Pole warm and heavily populated, then Santa would live in a warm, heavily populated region. The calculations behind the ascribing of properties to the average French voter are different in character, but perhaps not in kind. Furthermore, exactly how we determine the properties of fictional entities should not be confused with their ontological status *per se*. I establish the existence of my dog by direct observation, and the existence of Nelson Mandela by listening to news reports, but this does not result in each having a different ontological status: they are both individuals in the world.

Such an analysis depends crucially on the treatment of names of fictional characters as names of entities. We cannot go into an evaluation of this possibility here, and, even if this were an entirely defensible position, the further step still needs to be taken of defending treating *average*-NPs in the same way. On the other hand, perhaps fictional entities are, like *average*-NPs on the encuneral analysis, sets of properties.

7 CONCLUDING COMMENT

We started this paper with a consideration of an argument that has been put forward by Chomsky and some associates, involving *average*-NPs. The thrust of this argument is that *average*-NPs and ordinary proper nouns (and other referring expressions) behave linguistically in ways that are indistinguishable from one another, and yet it is patently obvious that *average*-NPs do not refer to anything concrete in the world. It is concluded that, regardless of whether proper nouns refer to things in the world, this must be irrelevant to their linguistic understanding . . . since that understanding is exactly the same as with *average*-NPs. But traditional semantics is fundamentally based on there

being some relation of reference, which will ‘hook up language and the world’. Therefore traditional semantics is, at its conceptual roots, a misguided enterprise. We took this argument seriously, and searched for a way to maintain this ‘hookup between language and the world’, while at the same time honoring the view that *average*-NP’s do not refer to anything in the world but nonetheless function just like other NP meanings do—as sets of properties. In the end we reconsidered the ‘bite the bullet’ point of view, and suggested an avenue of approach for those who, unlike Chomsky and others, might not find the ontological objections quite so convincing after all.

Neither analysis gives a complete account of all the questions and data we noted, in particular they are silent on the details of how and when ‘averaging’ takes place. These questions remain for future work.

ACKNOWLEDGMENTS

This paper is related to Hofweber and Pelletier (forthcoming). The present paper argues for an ‘encuneral NP’ analysis of *average*-NPs, while Hofweber and Pelletier argue for the plausibility of encuneral NPs generally. Talks based on the two papers have been given at the Amsterdam Colloquium (1999), the Rutgers Workshop on Philosophy of Language (2000), Michigan State University (2001), and the University of Calgary (2001). Our thanks go to those audiences for comments and criticisms. In particular, we would single out Theo Janssen (University of Amsterdam), Peter Pagin (Stockholm University), Jason Stanley (University of Michigan), and Ali Kazmi (University of Calgary) for their extremely helpful remarks, some of which we have shamelessly incorporated into this paper. Two anonymous referees also greatly contributed to its final form.

GREG CARLSON
Department of Linguistics
University of Rochester
Rochester, NY
USA 14627
e-mail: Carlson@ling.rochester.edu

Received: 19.10.2001
Final version
received: 10.04.2002

FRANCIS JEFFRY PELLETIER
Department of Philosophy
University of Alberta
Edmonton, Alberta
Canada T6G 2E5
e-mail: jeffp@cs.ualberta.ca

REFERENCES

- Bacon, J. (1973) 'Semantics of generic THE'. *Journal of Philosophical Logic* 2:323–339.
- Chisholm, R. (1957) *Perceiving: A Philosophical Study*. Cornell University Press. Ithaca.
- Chisholm, R. (1963) 'The theory of appearing'. In M. Black (ed.), *Philosophical Analysis*. Prentice-Hall. Englewood Cliffs, NJ.
- Chomsky, N. (1981) *Lectures on Government and Binding*. Foris. Dordrecht.
- Chomsky, N. (1986) *Knowledge of Language: Its Nature, Origin and Use*. Praeger Publishers. NYC.
- Chomsky, N. (1992) 'Language and interpretation: Philosophical reflections and empirical inquiry'. In J. Earman (ed.), *Inference, Explanation and other Philosophical Frustrations*. University of California Press. Berkeley, 99–128.
- Chomsky, N. (1993) 'Explaining language use'. *Philosophical Topics* 20:205–231.
- Chomsky, N. (1995) 'Language and nature'. *Mind* 104:1–61.
- Cooper, R. (1979) 'The interpretation of pronouns'. *Syntax and Semantics* 10:61–92.
- Davidson, D. (1967) 'The Logical Form of Action Sentences'. In N. Rescher (ed.), *The Logic of Decision and Action*. University of Pittsburgh Press. Pittsburgh, 81–95.
- Davidson, D. (1968) 'On saying that'. *Synthese* 19:130–146.
- Dever, J. (2001) 'Complex demonstratives'. *Linguistics and Philosophy* 24:271–330.
- Dowty, D., Wall, R., & Stanley, P. (1981) *Introduction to Montague Semantics*. Reidel. Dordrecht.
- Evans, G. (1977) 'Pronouns, quantifiers, and relative clauses, I and II'. *Canadian Journal of Philosophy* 7:467–536 and 777–798.
- Evans, G. (1980) 'Pronouns'. *Linguistic Inquiry* 11:337–362.
- Geach, P. (1962) *Reference and Generality*. Cornell University Press. Ithaca, NY.
- Heim, I. (1984) '*The Semantics of Definite and in Definite Noun Phrases*', Ph.D. dissertation, University of Massachusetts/Amherst.
- Higginbotham, J. (1985) 'On semantics'. *Linguistic Inquiry* 16:547–593.
- Higginbotham, J. (1990) 'Contexts, models, and meanings: A note on the data of semantics'. In R. Kempson (ed.), *Mental Representations: The Interface Between Language and Reality*. Cambridge University Press. Cambridge, UK, 29–48.
- Hofweber, T. & Pelletier, F.J. (forthcoming) 'Encuneral NPs'.
- Hornstein, N. (1984) *Logic as Grammar*. MIT Press. Cambridge, MA.
- Jacobson, P. (2000) 'Paycheck pronouns, Bach-Peters sentences, and variable-free semantics'. *Natural Language Semantics* 8:77–155.
- Kamp, H. (1981) 'A Theory of Truth and Representation'. In J. Groenendijk, T. Janssen & M. Stokhof (eds), *Formal Methods in the Study of Language*. Mathematisch Centrum. Amsterdam, 277–322.
- Karttunen, L. (1969) 'Pronouns and variables'. In *Papers from the Fifth Regional Meeting of the Chicago Linguistic Society*. CLS. Chicago, 108–116.
- Katz, J. & Fodor, J. (1963) 'The structure of a semantic theory'. *Language* 39:170–210.
- Krifka, M., Pelletier, F.J., Gregory, C., Meulen, A.ter, Gennaro, C., & Link, G. (1995) 'Genericity: An introduction'. In Gregory Carlson & Francis Jeffrey Pelletier (eds), *The Generic Book*. University of Chicago Press. Chicago, 1–124.

- Larson, R. (1985) 'Bare-NP adverbs'. *Linguistic Inquiry* **16**:595–621.
- Larson, R. (1998) 'Events and modification in nominals'. In D. Strolovitch & A. Lawson (eds), *Proceedings from Semantics and Linguistic Theory (SALT) VIII*. Cornell University Press. Ithaca, NY, 145–168.
- LePore, E. (1983) 'What model-theoretic semantics cannot do'. *Synthese* **54**:167–187.
- Lewis, D. (1972) 'General semantics'. In D. Davidson & G. Harman (eds), *Semantics of Natural Language*. Reidel. Dordrecht, 169–218.
- Longobardi, G. (1994) 'Reference and proper names'. *Linguistic Inquiry* **25**:609–666.
- Ludlow, P. (forthcoming) 'Referential Semantics for I-Languages?' In N. Hornstein & L. Antony (eds), *Chomsky and his Critics*. Blackwells. Oxford.
- Meile, J. (1995) 'On what there's not'. *Analysis* **55**:223–229.
- Neale, S. (1990) *Descriptions*. MIT Press. Cambridge, MA.
- Neale, S. (1993) 'Term limits'. *Philosophical Perspectives* **7**:89–123.
- Parsons, T. (1980) *Nonexistent Objects*. Yale University Press. New Haven.
- Quine, W. Van O. (1960) *Word and Object*. MIT Press. Cambridge, MA.
- Scheffler, I. (1964) *The Anatomy of Inquiry*. Routledge & Kegan Paul. London.
- Stanley, J. (2002) 'Hermeneutic fictionalism'. *Midwest Studies in Philosophy* **25**: (forthcoming)
- Vergnaud, J.-R. & Zubizarreta, M.-L. (1992) 'The definite determiner and the inalienable possession construction in French and English'. *Linguistic Inquiry* **23**:595–652.