

The Semantics and Ontology of *The Average American*

John Collins

Whether or not his newspaper and a set of senses reduced to five are the main sources of the so-called “real life” of the so-called average man, one thing is fortunately certain: namely, that the average man himself is but a piece of fiction, a tissue of statistics.

—Vladimir Nabokov, *Lectures on Don Quixote*

1: Introduction

The received wisdom is that truth-conditional semantics presupposes and entails certain externalist ontological commitments. The idea is simple enough: if a semantic theory specifies (compositional) truth conditions, then, for any given sentence that is true, the theory will specify the very conditions that hold such that the sentence is true. Presumably, such conditions are exactly ways the world might be that make our truths true and our falsehoods false. Semantics, on such a view, is world-involving merely by dint of being truth-involving. A general worry for this picture, no matter how otherwise tempting it might be, is the occurrence of *referentially defective* words and expressions, ones that appear properly to contribute to the truth conditions of sentences that may be variously true or false but which do not appear to have univocal referents we should otherwise sanction. Chomsky and others have claimed that, far from being outliers, a class of peculiar cases, referentially defective expressions in the shape of polysemous/copredicative constructions are utterly ubiquitous (Chomsky, 1975, 1977, 2000, 2003, 2013; cp., Hornstein, 1984; McGilvray, 1998; Pietroski, 2003, 2005; Collins, 2009, 2011, forthcoming a; Elbourne, 2011; Azzouni, 2013). The focus of the following will be the apparently referentially defective expression *the average American*. I shall

first spell out the problem it poses and upon that basis suggest a general two-tier approach to natural language semantics, under which the interpretation of syntax (the first tier) is not world-involving, but another level of interpretation is; this second tier of interpretation, however, is not specifically linguistic. So, if semantics is the interpretation of syntax conceived as structures specific to linguistic competence, then semantics is not world-involving. Secondly, the most thorough analysis of *the average+N* kind of construction, due to Kennedy and Stanley (2009), will be presented.¹ Their aim is to defend the orthodoxy; that is, they contend that the relevant phrases may be given a clear semantic analysis that renders them non-defective in the sense that the semantics delivers a ‘sensible ontology’. Thirdly, a number of problems for the analysis will be developed. Fourthly, an alternative account of the peculiar properties of the construction will be offered in line with the two-tier proposal previously suggested. The account preserves the defectiveness of *the average American*, but renders it unproblematic from a theoretical perspective that eschews ontological adequacy constraints upon a semantic theory.

One might wonder why *the average+N* construction merits such attention, especially given that the general anti-externalist line of reasoning Chomsky and others suggest does not rest upon the properties of relatively *recherché* constructions but on claims pertaining to the ubiquity of polysemy and copredication. The basic thought here is that lexical items have multiple construals that belie the idea that an item’s meaning can be given in terms of a univocal worldly referent; crucially, polysemy/copredication cannot be analysed away as a case of ambiguity. For example, *Germany* may be used to speak about, either separately or simultaneously, a geographical area, a population, a state, an idea, and perhaps other things. Similarly, *book* may be used to speak about, either

separately or simultaneously, a concrete particular, a type of object, a body of information/content, an idea, and perhaps other things. The same holds for virtually any nominal one cares to pick. Yet we are not ineluctably led to posit queer external objects that realise such collections of properties, which, perforce, have wildly divergent and apparently contradictory individuation conditions. The point, however, is not so much that such putative objects are impossible according to some or other metaphysical scheme, although an abstract object with mass and a geographical area with psychological states seem pretty contradictory; the point, rather, is that an understanding of the semantic properties of the nominals and so our competence with them in forming perfectly acceptable constructions, appears to be utterly unconstrained by any notion of an externalist ‘sensible ontology’, i.e., the queer objects neither explain the semantic properties nor, still less, reflect any otherwise plausible metaphysical presuppositions (Collins, forthcoming a).

Fascinating as the polysemy/copredication issue is, as indicated, it will not be my focus. Kennedy and Stanley are right to think that *the average*+N case is of especial interest precisely because it offers a linguistic example where settling on the syntax and compositional semantics *appears* to deliver a referentially defective expression that is thus a standing refutation of the ontological import of semantics in the absence of an account of why the appearances are misleading. Thus, the putative ontological import of semantics can be tested while bracketing the thorny issue of polysemy/copredication and broader issues of what there is to be referred to (one can bite various ontological bullets in order to keep semantics aligned with what there is).

2: Chomsky’s argument *sans* polysemy

Let's assume that an adequate semantic theory deals with truth conditions and so too the referential conditions of words and phrases. To know the meaning of a sentence is therefore to know something about the world, *viz.*, what it must be like for the sentence to be true or false. Thus, if we can say anything at all that is true, then our semantic theory will tell us something about the world, i.e., what it is like such that a sentence with such and such truth conditions is in fact true. Lots of bells and whistles might be added, concerning contextual parameters and circumstances of evaluation, but they don't alter the basic picture. Let's encapsulate this thought by saying that a semantic theory will entail (if we can say anything true at all) or presuppose whatever external ontology (externalia) makes our true sentences true. This is the standard view.

Famously, Lewis (1970, p. 190) rejected the so-called *Markerese* semantics of Katz and Fodor (1963), concluding that the putative semantics doesn't deal with 'the relations between symbols and the world of non-symbols—that is, with genuinely semantic relations'. The point to focus on here is not the virtues and vices of Markerese (assume that Lewis's negative judgement is correct), but the demarcation of semantics as being concerned with world-involving properties. This kind of complaint has become something of a shibboleth, where for a would-be semantic theory to be non-world-involving is tantamount to its being a case of false advertising (e.g., Evans and McDowell, 1976; Lepore and Loewer, 1981; Davidson, 1984; Higginbotham, 1985, 1989; Stanley, 2007). The attitude has acquired a textbook status, as it were. Dowty *et al.* (1981, p. 5) set their face against Markerese and other symbol-symbol semantic theories within linguistics: 'the proper business of semantics is to specify how language connects with the world – in other words, to explicate the inherent "aboutness" of language'.

Dowty *et al.* go on to endorse a correspondence theory of truth as the one fit for semantic duty. Chierchia and McConnell-Ginet (1990/2000, p. 53) appear to hold a similar view:

[T]he study of the relation of symbols to what they stand for [“refer to”, “denote”, “their informational content”] must indeed be a part of an account of meaning. For otherwise how could we understand the fundamental fact that configurations of symbols carry information about all the diverse aspects of our experience?

Jacobson (2014) echoes the view in her textbook on semantics. She

assumes that meaning is not just some string of symbols [generally, a representation], but rather some actual object out there in the world... [T]he point is that the grammar maps each linguistic expression into something beyond just a symbolic representation. Otherwise—as so aptly pointed out by Lewis (1970)—we are simply mapping one language (say, English) into another (what Lewis termed ‘Markerese’). Yet language is used to convey facts about the world; we draw inferences about the world from what we hear and we gain information about what is true and what is not. So semantics must be a system mapping a linguistic expression to something in the world (*ibid.*, pp. 27-8)

Of course, these are not just expressions of dogma. As Evans and McDowell (1976, pp. ix-x) rightly stress, it is a ‘fallacy’ to think that just because one needs language to state the truth conditions or meaning of a sentence, one is thereby merely relating one bit of language to another, rather than relating language to the world. Quite so! But it is another fallacy to think that just because one is using language (as opposed to mentioning it) that one is thereby managing to pick out something in the world out there anyway, or even that the theory that includes the relevant truth-conditional statements presupposes or

entails referential success. The point here is brought into sharp relief by *referentially defective* words and expressions. Let's characterise the class as follows:

(RD) A word or expression *e* is *referentially defective* iff (i) *e* apparently lacks a worldly referent even under the 'right' logical analysis, but (ii) *e* can occur as a proper constituent of a sentence that apparently can be true in a full-blown literal way.

Of course, at first blush, many phrases lack referents in an intuitive sense, such as quantifier phrases. I take it, however, that the right analyses of such phrases parcels the semantic properties in such a way as to relieve any such qualms. Thus, the hedges here indicate that one might end up claiming that no expression really is defective; still, referential defectiveness appears to be common. The RD-words that have hitherto invariably been discussed by philosophers are cases of fiction, myth, or some other form of inexistence (*Sherlock Holmes*, *Zeus*, *Father Christmas*, etc.). It is fair to say that perhaps all of the many positions considered for these cases remain live options. A more troubling case, however, that has attracted significantly less attention, is Hornstein's (1984) and Chomsky's (2000, 2003) case of *the average American/family/etc.*² There are three basic reasons to think that this expression is especially hard to handle while retaining the view that semantics is aligned with externalist ontology.

First, with fiction and myth, there is always the nuclear option, as it were, of simply denying truth to the relevant claims (Quine, 1953). The defectiveness is resolved merely by denying the second condition of RD. There is no temptation whatsoever to do likewise with the cases in (1):

- (1)a The average American has 2.3 children
- b The average American drives a Ford

Both look like straightforward empirical truths, which can, indeed, be verified by checking various facts about Americans.

Secondly, with fiction and myth, it doesn't seem too unpalatable to sanction the various objects and confer on them whatever status fits the claims we make about them. Thus, let there be Sherlock Holmes, *in some sense* (cp., Van Inwagen, 1977; Salmon, 1998; Kripke, 2013). Acknowledging his being is not yet to think of him, absurdly, as a concrete particular one could meet on Baker St., but it is to deny the first condition of RD. A reciprocal move, without further ado, is a dubious option for *the average American*. Again, considering the cases in (1), we see that positing an entity that is the average American over and above or somehow related to the particular Americans badly screws up the truth conditions, rendering the cases false, if not gibberish. The truth of (1a) appears not to commit us to, or even suggest, an American with 2.3 children additional to the other Americans. There is no such entity and nor can there be; rather, (1a) expresses something about all Americans in some abstract way, by way of a nominal that appears defective. Likewise for (1b). Only human beings drive Fords, and the average American is not amongst them, so we precisely don't want to account for the truth (1b) by positing an additional entity—the Americans there are suffice to fix the truth value of (1b). The average American as an additional entity appears to be supererogatory. Still, the semantic properties of the cases in (1) appear to render us with such an object, even though we can, as it were, see past the object to the truth of the claims. What makes the expression referentially defective is not that we cannot see how the host sentences can be true, but that such worldly truth conditions appear to swing free of what the expression

in fact encodes from a narrow linguistic perspective. My positive proposals about the construction will turn on this feature.

Thirdly, with fiction and myth, it is tempting to imagine that some form of pretence or figural stance or even an implicit intensional context is in play, which means that we ought not to treat the relevant sentences as expressing off-the-bat literal truth claims. There is little temptation, however, *pace* Yablo (1998), to render the kind of cases exhibited in (1) as similarly less than literal (or covertly intensional). The truth conditions they express, it seems, do not bear on anyone's conception of or presuppositions and attitude towards the relevant realm of facts. Of course, arguments may be had about the *the average+N* constructions in this regard, but a case has to be made, which has not been made (cp., Stanley, 2001).

All of that said, I do think that a general philosophical orientation is available that deals with *the average American* and fiction/myth uniformly; the key is to renounce the thought that referential defectiveness must somehow be explained away so that a sensible ontology issues from semantics. In this sense, one may reject *both* otherwise unmotivated analyses of the logical form of referential defective cases (*à la* Kennedy and Stanley) *and* the adoption of a figural stance towards what appear to be semantically 'normal' phenomena (*à la* Yablo) (see §5).

The real challenge Chomsky poses, however, is not merely the presentation of an especially troubling case of an RD-expression, after all, fictional cases are hard enough, and remain unsettled. Chomsky's thought, rather, is that referentially defective cases are simply a way of highlighting a gap between semantics and externalist ontology. That is to say, all else being equal, an adequate semantic theory on Chomsky's conception will treat

all cases on a par, regardless of whether they have an acceptable worldly referent or not. All else would not be equal if there were some particular linguistic/semantic reason to differentiate between expressions quite apart from one being *soi-disant* defective. Thus, Chomsky (1981, p. 324; 2000, p. 39) suggests that binding is an inquiry into referential properties, and so semantics, on one construal of the notion, but nothing in binding theory entails or presupposes a domain of language/mind-external entities that are otherwise coherent or merit a competent speaker-hearer's commitment. The reason for this is that the principles governing acceptable binding relations hold independently of any conception we have of the ontology of the referents in question. Put otherwise, we may say that Chomsky's point is that binding is indifferent to any criteria that would constitute a sensible ontology, a realm of entities speaker-hearers would otherwise sanction. The point here is intended to generalise from binding to the full range of semantico-syntactic phenomena linguists and philosophers typically investigate. In short, semantics as actually pursued is an inquiry into structural conditions on the kind of thoughts we can express with linguistic forms independent of how the world actually is that would make our thoughts true or false.

Consider the following cases:

(2)a [Barack Obama]_i has been damaged by the health care issue, but he_i remains likely to go down as a great president.

b. [The average American]_i is optimistic no matter the setbacks he_i faces.

As marked by the indexes, the coherence of intra-sentential referential dependence is invariant over the two cases even though, let us assume, none of us thinks that *the average American* designates a thing the way *Barack Obama* does; at any rate, the

understanding that (2b) is truth-apt does not obviously involve a commitment to such *recherché* metaphysics. (2b) is, of course, reminiscent of Geach's (1967) famous Hob-Nob kind of case, but the problem it poses is distinct. The problem Geach posed was how to make sense of the identity of intentional inexistents, such as witches, across anaphora. That problem arises precisely because there are no witches, although Hob and Nob believe there to be such. In the present case, the average American is not an intentional in-existent, but is still not a thing, and so not a thing to which we can refer back, even though (2b) is perfectly fine, in fact, true. The cases have this in common, though: the syntactic and semantic coherence of binding relations is apparently independent of the ontological status of the semantic values of the pronouns, which is just what one would expect, of course, if Chomsky is right.

It should be noted, too, that the *prima facie* problem (2b) and other binding cases pose for the alignment of semantic theory with an externalist ontology does not depend on construing the relevant pronouns as singular terms. For example, an E-type account of anaphora fails to ameliorate the problem.³ The basic idea behind the notion of E-type anaphora is that the bound pronoun has its reference fixed descriptively by way of the descriptive content of its antecedent (Evans, 1977; Neale, 1990). There are various ways of implementing the basic idea that pronouns have a descriptive reference-fixing content, and perhaps none of them are entirely happy, but a simple account will suffice for present purposes. For example, in *Bill bought a duck, but it ran away*, we can read *it* as not so much a bound variable, but as a term that has as its value whatever may value the definite description *the duck Bill bought* (or perhaps even just *the duck*). Crucially, however, *it* remains a pronominal rigid singular term, at least on Evans's initial account, so the

problem of what the pronoun refers to remains. Simply put, the problem with (2b) is not how the value of the pronoun may be fixed, but that it has a value at all! Besides which, as we shall see in §4.3, ‘*the average+N*’ can enter into ‘non-lazy’ binding relations (see note 3). Heim (1990) and Neale (1990) introduce the idea of D-type anaphora in response to certain problems with the E-type model. A D-type anaphora is a case where a description is a genuine semantic proxy for the bound pronoun, not a mere reference fixer, i.e., at the right level of analysis, the pronoun *is* a description. Again, though, even if the pronoun is treated descriptively, it still carries existential force in the sense of entailing an object that satisfies the description, but that is precisely what we don’t want in the present case.

Chomsky (1981, p. 324) offers the following general admonishment:

[I]nterpretation [of syntax] is not to be confused with what might be called “real semantics”, that is, the study of the relation between language or language use and the world. Rather it should be considered to be in effect an extension of syntax, the construction of another level of mental representation beyond LF [syntax], a level at which arguments at LF are paired with entities of mental representation, this further level then entering into “real semantic interpretation”.

So, ‘real semantics’ covers issues pertaining to how the cases in (1-2), say, can be true, or, indeed, how fictional discourse can be true, or how any words are used or otherwise understood to be true of a speaker’s environs. This is a far wider topic than is covered by linguistic semantics proper or the interpretation of syntax. That is a structural matter, mapping one system of internal representation onto another. If that is right, then semantics has no externalist ontology at all, for semantics does not relate symbols to

externalia. In effect, then, Lewis (1970) *et al.* are right about ‘real semantics’, but relations such as truth and reference are way down-stream of what can be narrowly fixed by the semantic interpretation of grammar/syntax, and such fixation involves numerous extra-linguistic capacities and factors pertaining to the occasion of speech or thought.

Such a position, of course, does not by itself amount to any positive account of syntax or the internal conceptual structure onto to which it maps, but nor is it meant to. For present purposes, adopting such a structuralist conception of natural language semantics amounts to the separation of ontological criteria from semantic conditions, an option Lewis *et al.* reject on nigh-on a priori grounds. That is, if the above reasoning is on the right lines, then a general structural orientation holds for the linguistic conditions that apply independently of any ontological criteria of what counts as a possible referent, such as binding, whereas ‘real semantics’, concerning truth and reference, applies beyond language narrowly construed, in the sense that an adequate semantics and syntax for language leaves worldly truth/reference conditions undetermined or potentially defective.

Lest this position be thought an odd quirk of Chomsky’s, the thoughts of other semanticists are worth noting. Krifka (1998, p. 198) writes

It was believed that model-theoretic semantics [as well as Davidson-style semantics] is forced to a ‘realistic’ view, in which natural-language expressions are interpreted by real entities, like objects and possible worlds, whereas cognitive semantics is concerned with cognitive models of reality. I don’t see that model-theoretic semantics has to be realistic in this sense. We can make use of the techniques developed in the model-theoretic tradition and assume that expressions are interpreted by elements of conceptual structures that in turn are related to ‘real’

entities by some extra-linguistic matching... [Semantics] attempts to capture certain properties of the way we see the world, not... to describe the world how it is.

In much the same vein as Chomsky, Krifka is here separating semantics as a narrow matter of the interpretation of syntax by other conceptual structure from ‘real semantics’ that hooks up what we say with bits of the environment, or at least what we think of as such.

Similarly, Asher (2011, pp. 44-8), without appeal to Krifka or Chomsky, endorses a ‘two-level’ account of linguistic content, a wholly internal level and a world-involving level. Crucially, the former does not determine the latter; the relation between the two levels is ‘at best a homomorphism’ (ibid., p. 48), by which is meant that the internal level mirrors the structure of our ontological commitments but does not encode such commitments.

Bach’s (1986, p. 575) slogan ‘No semantics without metaphysics!’ might appear to contradict the position Chomsky *et al.* are suggesting, but Bach’s further elaboration patently does not:

As a linguist, I feel perfectly justified in sidestepping such questions [of existence and metaphysics]. Consequently, I like to say that what I am doing is not metaphysics *per se* but *natural language metaphysics*... What we are doing is simply seeking linguistic evidence for the nature of the semantic structures that we seem to need to give a good account for the meanings of natural language expressions (Bach, 1989, p. 98)

For example, the semanticist may posit events, but the adequacy of a semantic theory that trades in such events does not *eo ipso* turn on the metaphysical status of events

independent of the requirements of the theory; indeed, the events posited might be incoherent from a metaphysical perspective (cp., Pietroski, 2015).

As indicated above, the general issues here raised about the foundations of semantic theorising do not devolve onto subtle consideration of *the average American*. Chomsky (2000, 2003, 2013) and others (Hornstein, 1984; Pietroski, 2003, 2005; Collins, 2009, 2011) have argued on the basis of polysemy/copredication considerations that no word at all carries externalist ontological import. Proper consideration of this issue would involve the assessment of many varied matters regarding both ontology and semantics, both in general and for particular constructions (see Ludlow, 2003, and Chomsky, 2003, and Borg, 2009, and Collins, 2009 for some critical debate, and Gross, 2006, 2015, for overviews of the issues). Still, the real interest *the average+N* construction has is exactly its offering a case of an apparently referentially defective complex expression which poses a challenge to the orthodoxy entirely independently of any general attitude towards the ubiquity of polysemy. In simple terms, even the most hard-headed, ontologically-minded, polysemy-rejecting conservative semanticist should wonder what on earth to say about *The average American has 2.3 children* and similar sentences.

It bears emphasis that the matter before us cannot be settled by any quick *a priori* move of the kind Lewis and the textbooks suggest. If Chomsky and the others are right, then no particular set of interpretive principles of language will, in fact, involve a presupposition or entailment of an externalist ontology, just as binding, in the shape of anaphoric interpretation, appears to be utterly indifferent between the status of Obama and the average American. For Chomsky to be wrong, it needs to be actually shown, not

merely insisted, that semantics is world-involving in the relevant sense. This point is not always properly registered.

It has seemed to many philosophers that weighty ontological commitments are simply built into truth-conditional semantics, or even the very idea of a semantic theory *per Lewis et al.* Kennedy and Stanley (2009) and Carlson and Pelletier (2002) certainly appear to think so, for they conflate Chomsky's rejection of externalism in semantic theory with some kind of unhinged rejection of truth-conditional semantics. Kennedy and Stanley (2009, pp. 585-6) write:

On the face of it, Chomsky's position sits oddly with his espousal of methodological naturalism. [Chomsky] 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 inquiry.

The accusation is that the kind of questions raised by Chomsky and others for the ontological involvement of truth-conditional semantics amounts to a kind of a priori harassment of successful science. Higginbotham (2009 p. 154) describes a position that endorses truth-conditional semantics without a concomitant commitment to the ontology that apparently thereby flows as being akin to 'bad faith' (cp., Burge, 2003, p. 466-7). As Higginbotham (2004, p. 575) says elsewhere: 'what turns up in the metaphysics of semantic investigation cannot be passed off as a mere manner of speaking, but constitutes

our best conception of the way the world is... [it] must be taken seriously'. Carlson and Pelletier (2002, p. 77) go further and associate Chomsky with, among others, Putnam, Cavell, Rorty, Derrida, and Eco, all of whom, in their different ways, reject truth-conditional semantics, or even semantics *tout court* as a systematic inquiry, on the basis of a priori considerations (well, I don't know anywhere where Derrida talks about such matters, but I guess we can guess). Carlson and Pelletier read Chomsky as suggesting that formal semantics is 'doomed to failure', 'bankrupt', an exercise in 'futility' (ibid., p. 77, 78, 81).

The thought seems to be that Chomsky and, by association, numerous others, mount a priori arguments against the very discipline of semantics *simply because* they offer considerations against the alignment of semantic theory with an externalist ontology. There are, however, no passages in Chomsky's texts that support such an extreme reading, which is, in fact, contradicted by Chomsky as quoted above, who further remarks that semantics consists of a 'great deal of illuminating work' (Chomsky, 2000, p. 174).⁴ What remains unsettled, let alone a priori obvious, is whether any of this illumination or 'extraordinary progress' is actually dependent upon or even suggestive of an externalist ontology, for the work might be best thought of, just as Chomsky and Kifka suggest, as illuminating the narrow interpretation of syntax independently of any commitments to how the world is or how our reference to it succeeds. Let us, then, consider *the average American* as a test case. Kennedy and Stanley (2009, p. 586) suggest that the behaviour of *the average American* amounts to Chomsky's 'strongest' argument for the referential defectiveness of natural language. 'Strongest' might be too

strong, but, for the reasons adduced, I think their focus is correctly targeted, so the little phrase deserves to be centre stage.

3: The semantics and ontology of *average*

Recall, the basic problem posed by *the average American* is that it appears to be syntactically and semantically on a par with kosher non-defective referential terms, which suggests that semantics does not reflect any ontological scruples, that, say, there is Obama but no average American. The standard riposte to this line of reasoning is that complex constructions such as *the average American* are not singular terms (Higginbotham, 1985; Stanley, 2001; Carlson and Pelletier, 2002; Ludlow, 2003; Kennedy and Stanley, 2009). After all, we cannot read off an expression's logical form from its surface form. As the matter stands, therefore, Chomsky's moral of the lack of contrast between (2a) and (2b) is hostage to the fortune of a naive semantic analysis of *the average+N* constructions. In effect, therefore, the common claim is that the relevant DPs are not referentially defective at all, for a 'right' semantics will reveal them to be suitably non-referential, and so not begetting of an ontological commitment towards an average American entity. This would be so if *the average+N* construction is not a unitary interpretable item at the right level of logical form.

The spirit of this rejoinder to the putative significance of the cases in (2) is perfectly sound. In particular, Stanley (2001) and Kennedy and Stanley (2009, p. 589) are correct to insist, *pace* Melia (1995) and Yablo (1998), that at least an argument is required if we are to accept Chomsky's (2000, p. 39) claim that the semantic contribution of *average* is

much like any other adjective, rendering the target construction much like any other definite [_{DP} the ADJ+N] construction.

This section will present Kennedy and Stanley's analysis. The following sections will raise some problems for it and offer an alternative analysis in line with Chomsky's animadversions against externalist ontology flowing from semantic considerations.

3.1: Kennedy and Stanley on 'abstract' *average*

Following Carlson and Pelletier (2002), Kennedy and Stanley focus on sentences such as *The average American has 2.3 children*, which have an 'abstract' rather than a 'typical/concrete' construal (this distinction will be explained shortly). I shall assume with all parties that the target sentence is acceptable, i.e., not semantically or grammatically anomalous. Kennedy and Stanley's chief stated goal is to establish that the abstractly construed *the average+N* construction is not a singular term or a generalised quantifier expression; the specification of the truth conditions of its host sentences, therefore, do not make reference to an entity we should otherwise not sanction. So, one might think of Kennedy and Stanley as offering a *paraphrase* strategy to secure sensible truth conditions for *the average+N* constructions. According to Carlson and Pelletier's (2002, p. 78) taxonomy of options, Kennedy and Stanley would indeed fit the description of a 'paraphrase away' account, albeit one of great sophistication (cp., Quine, 1960; Yablo, 2001). As Carlson and Pelletier (ibid., p. 80) remark, such accounts most often rest on 'metaphysical presumptions' about what should be ontologically sanctioned regardless of the apparent semantic properties of the language. In other words, the danger is that the metaphysical tail is wagging the semantic dog. Such an impugnation would be

unfair to Kennedy and Stanley. The main business of the analysis is to provide a general compositional account of *average* in terms of the independently accepted resources of semantic and syntactic theory, and one could accept the analysis without any concern for its advertised ontological consequences, which might be seen as mere motivational factors to consider the construction. We might think, therefore, of the proposal as a semantically legitimate analysis that has the kind of consequences a traditional paraphrase strategy seeks. Still, ultimately, I think Kennedy and Stanley are guilty of the kind of metaphysical presumption Carlson and Pelletier admonish against, even if the presumption does not affect the actual analysis offered. This is because, or so I shall contend, Kennedy and Stanley get the metaphysics of being abstractly average right, but fail to show that such metaphysics is encoded in the relevant semantico-syntactic structure.

I shall first sketch the account in outline; thereafter I shall provide more detail in discussion.

Kennedy and Stanley's principal claim comes in two parts. Firstly, *the average* is taken to be a semantically fused expression, with *the* being pleonastic, not the standard definite determiner. Still, *the average*+N is syntactically realised as a DP. Secondly, *average* in the 'abstract' sense (a notion to be explained shortly) expresses a *measure function* that maps from the division of a *degree relation* expressed by the scope of the DP (e.g., *number of children*) by its restriction (e.g., *Americans*) to an exact degree (e.g., 2.3). So, *The average American has 2.3 children* comes out as meaning that 2.3 is the degree of the measure of offspring divided by their American parents. There are, of course, issues here about who counts as American and an offspring, but decisions on such

matters are not covered by the analysis; after all, one could think, perhaps on error-theoretic grounds, that no statements involving abstract *average* are true precisely because of such potential vagaries and indeterminacies, but that would not condemn the analysis. The crucial factor is that the truth conditions of the relevant sentences do not involve an actual thing who is the average American. Furthermore, the analysis does not generalise to ‘typical’ or ‘concrete’ readings of *average*, where, say, *the average American* denotes any number of concrete Americans who are contextually assessed as typical according to the relevant dimension (e.g., eating habits, political views, etc.). The two construals of *average* appear to be quite different; in particular, the ‘typical/concrete’ reading entails concrete instances, which the ‘abstract’ reading does not. We can express this as an *existential test*:

(ET) If *The average N is F* entails that there is at least one N that is F, then *average* has a ‘typical/concrete’ construal.

An abstract construal of *average* fails ET, not so much because there might just happen not to be any N that is F, but because the predication of F to a particular N is semantically anomalous; that is, it is simply not an available thought that a particular American could have 2.3 children in the relevant sense (a chopped up child doesn’t count). Consider:

- (3)a The average American has 2.3 children
- b #The typical American has 2.3 children

No American can possibly have 2.3 children (in the intended sense), so no typical American can do so either. This phenomenon needs to be explained by any adequate semantics of *average*. It is a real virtue of Kennedy and Stanley’s analysis that it goes so far to capture this quite subtle difference. As we shall see, though, while the construals of

average are distinct, they are perhaps not as distinct as Stanley and Kennedy's analysis predicts (cp., Carlson and Pelletier, 2002).

Another nice consequence of the analysis is that it allows us to see *average* as a proper compositional constituent of its host structures. Higginbotham (1985) suggested that *average*, when construed abstractly, functions as an adverbial. This means that (3a) would correspond in meaning to (4):

(4) Americans, on average, have 2.3 children

The problem with the proposal now is that we are obliged somehow to reanalyse *have 2.3 children*, for, again, no American can possibly have 2.3 children, but the predicate in (4) appears precisely to attribute that property (Kennedy and Stanley, 2009, p. 609).⁵ Treating *average* as a measure function avoids this problem as *have 2.3 children* is not a direct predicate of the surface DP in (3a), but a degree relational argument of a function, whose numerical argument has moved into a higher scope position (see below).

The semantic account as sketched is intended to interpret a particular logic-syntactic form. Central to the analysis of ('abstract') *average* on offer is that the scope or predicate of the *average* DP expresses a degree relation. This poses a problem of how to separate the degree argument (the numeral expression, such as 2.3) and the degree relation (the scope of the DP, such as *x has n children*), which appear to be fused in the predicate, say, *has 2.3 children* (Kennedy and Stanley, 2009, pp. 625-7). The solution proposed takes the numeral expression to move over the DP *the average American*, by way of general syntactic operation that allows for items to move from a 'base' position in order to form scope relations (the operation standardly goes by the name QR). The movement means that *average* now takes as an argument the predicate formed by the movement of the

numeral. This gives us the following kind of logico-syntactic form, what Kennedy and Stanley refer to as ‘Logical Form’:

(5) $[_{XP} [2.3 [_{DP} [\text{the average}] [\text{American}]] [\lambda n \text{ has } n \text{ children}]]]$,

where the function expressed by *average* may now take its arguments as independent items. My query about this proposal concerns the formation of the scope $[\lambda n \text{ has } n \text{ children}]$ as a putative independent argument of the *average* function. Kennedy and Stanley (2009, p. 620) have the following to say about the general case:

The semantic effect of QR is [...] 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.

If this ‘mix’ were a mere matter of presentation, as Kennedy and Stanley suggest, there should be little complaint, but it does raise substantive questions. Kennedy and Stanley appear to be influenced by Heim and Kratzer’s (1998, pp. 184-8) semantic analysis of QR for the case where the raised quantifier DP leaves a trace in the object position of a VP, such as in *John offended every linguist*. Heim and Kratzer’s desideratum is to get *every* to take two $\langle e, t \rangle$ -type expressions (1-place predicates) concomitant with the verb *offend* taking an $\langle e \rangle$ -type expression as its object (a quantifier DP is of type $\langle \langle e, t \rangle, t \rangle$). Their solution is to take $[\alpha 1 [\text{John offended } t_1]]$ as a case of predicate abstraction, where ‘1’ is an index for the trace of the moved DP *every linguist*, such that the value of α is $[\lambda x. \text{John offended } x]$. This is not quite what Kennedy and Stanley propose, for Heim and Kratzer keep λ -abstraction properly speaking out of the syntax, proposing instead an extra projection that is functionally interpreted as abstraction, which turns the clause into an

<e, t>-type function. Still, it is difficult not to see the posited projection just as a ‘ λ ’—what else could it be? Thus, substantive issues arise pertaining to the syncategorematic ‘blemish’ of λ -abstraction as implemented by both proposals and whether, as Kennedy and Stanley suggest, the mix of meta-theoretical description and syntax proper can be disentangled. I think there is cause for scepticism here. One may appeal to λ -abstraction as part of one’s metatheory or as a notational convenience to mark an interpretation, but it is unclear how λ -abstraction is supposed to be implemented by QR *qua* syntactic operation, i.e., the new projection is headless, corresponding to no lexical item, and so appears to be parachuted into the syntax with no independent rationale.⁶ These are general issues, though, so I shall not pursue them further here.

Even from this sketch, I hope it is obvious that Kennedy and Stanley’s analysis is rich in consequence and interest. I shall, though, raise three empirical problems, which in turn will motivate an alternative proposal.

4: Three empirical problems

Three empirical issues complicate Kennedy and Stanley’s account; the first two relate to the separation of ‘abstract’ from ‘typical’ construals of *average*. The last bears upon binding. The proposal also faces the problem that it appears that abstract readings of DPs are available without *average* modifying the nominal, i.e., the relevant proposition can be expressed without use of *average*. I shall discuss this topic in relation to my positive proposal, for such a possibility is not in itself a fatal flaw: *average* may have an interpretation when it occurs, and it is hardly an empty expression, of course. Still, if the same content can be delivered without *average*, then the apparent ontological absurdity

of the full DP, *the average American*, cannot be avoided by decomposing *average* and the item it modifies, because there is no modifier!

4.1: Copredication

Carlson and Pelletier (2002), discussed by Kennedy and Stanley (2009, pp. 634-8), point to copredication cases, such as (6):

(6) The average American has 2.3 children and drives a domestic automobile.⁷

As we shall see below, there are lots of different ways of framing copredications, The basic feature, however, is that a single expression or phrase, a DP subject in the case of (6), can simultaneously serve as arguments of predicates or modifiers that appear to select for categorically distinct entities. Recall that an essential component of Kennedy and Stanley's view is that the numeral expression within the predicate of the 'abstract' *average* DP subject moves to scope over the DP. (6) poses a problem in this regard, for if the numeral moves, then it moves out of one constituent phrase of a phrase that is an island to movement, i.e., a conjunction of VPs as in (6).⁸ If the numeral is interpreted *in-situ*, then the wrong reading is produced: the conjunctive predicate would appear to attribute the property of having 2.3 children and driving a domestic automobile, but no American has such a property. To avoid this result, the first VP in (6)—*has 2.3 children*—must somehow be interpreted in an abstract way independently of the second VP that applies to typical Americans. Kennedy and Stanley suggest that there is a covert 'abstract' *average* adjoined to the first VP, which allows for the movement of the

numeral expression to scope over it within the VP. In effect, (6) is equivalent to (7), where the higher *average* is ‘typical’ and the VP internal *average* is ‘abstract’.

(7) The average_(t) American has an average_(a) of 2.3 children and drives a domestic automobile.

This trick works and constructions such as (7) are witnessed; indeed, if we are happy with a covert ‘abstract’ *average*, the semantic account Kennedy and Stanley propose can work smoothly for the kind of copredication cases Carlson and Pelletier consider (but perhaps not others—see below). As it is, though, a covert ‘abstract’ *average* cannot be so innocently posited. Kennedy and Stanley are all too aware that an overgeneration worry arises, i.e., if a covert ‘abstract’ *average* is available, then what conditions preclude it from rescuing apparently deviant sentences. This suggests a general question of why we should think that there is a covert *average* in the problematic copredication cases such as (6). That is, if no clear licensing conditions can be formulated for a covert *average* (i.e., specified conditions that generate just the right readings), the positing of a covert *average* for the copredication cases is but a hopeful speculation, even if its occurrence would genuinely solve the copredication problem.

Kennedy and Stanley (2009, p. 637) point to the following kind of example from the *New York Times*:

(8) One survey... 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.

The second sentence here is naturally interpreted as the men having on average 12.7 partners and the women having on average 6.5. This reading is said to provide ‘compelling evidence for the existence of a covert abstract ‘average’ (*op cit.*). The suggestion is that the antecedent occurrence of *median* primes an abstract construal of the following sentences, which otherwise would express metaphysically bizarre propositions (*ibid.*, p. 638). Hence, the presence of a covert modifier expressing the content of mean *average* may be hypothesised as an explanation of the acceptability of such discourse.

The priming effect, however, seems redundant, for the last sentence of (8) is fine by itself, i.e., if read on its own, it clearly would be construed as expressing a proposition about mean sexual partners. So, an antecedent average term is not required for a covert abstract reading of a subsequent clause. Moreover, priming generally does not discriminate lexical or syntactic content from general psychological association. So, priming is not required for the abstract construal of bare NP plurals, and, anyway, even if the abstract construal were licensed by priming alone, the construal might still be a broad psychological effect and so not be evidence for any hypothesis about the syntax or logical form of the relevant sentences. In sum, therefore, the discourse cited as ‘compelling evidence’ for the syntactic presence of a covert ‘abstract’ *average* is quite unlike (6) in precisely the relevant respect, *viz.*, (6) is supposed to contain a covert ‘abstract’ *average*

without a licence from an antecedent primer, so why should the fact that one can prime an abstract reading amount to a licence for the presence of a modifier at logical form?

Kennedy and Stanley do address the divergence between the cases. As mentioned above, the proposal as it stands faces an overgeneration problem: if a covert ‘abstract’ *average* is available, what conditions preclude its occurrence rescuing (9)?

(9) ?The typical American has <on average> 2.3 children

In other words, the unacceptability of (9) militates for the priming/licensing account of covert *average* insofar as a covert *average* (indicated by the angled brackets) is not properly licensed. It remains unclear why (9) is out, though; and what of the copredication cases?

Kennedy and Stanley (2009, p. 638) say the following:

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 [(8)], the first two sentences of the passage explicitly mention medians. And while the same cannot be said of [(6)], it may 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 going on here).

The first two sentences amount to the claim that the presence of covert ‘abstract’ *average* in (8) is a discourse effect, where priming is supposed to license items at logical form (a parse is not good enough, as the account on offer is of logical form/syntax, not merely a parse). If this claim is correct, then (8) does not provide *any* evidence for the presence of covert *average* in (6), because (6) lacks the requisite priming information. The thought expressed in the final sentence seems to finesse the initial suggestion: perhaps a ‘typical’

average might suffice to license a covert ‘abstract’ *average*. (9), then, is out because there would be a single predicate associated with two *averages*, whereas the licensing works by the covert *average* not being in competition with the typical *average*. All that said, (8) in particular provides no evidence for the proposal. On the contrary, it is not the occurrence of ‘typical’ *average* that licenses covert ‘abstract’ *average*, but the predicate containing a relevant numerical expression. That is, we could think of the ‘abstract’ reading as coerced from a ‘typical/concrete’ reading by way of a predicate that imposes an abstract reading on the subject. I shall say more on this once we get to the positive proposal about *average*. *Pro tem*, there are two clear reasons to think that it can only be the predicate that licenses the abstract construal of *the average+N*, not another occurrence of *average* or an intrinsic feature of the DP.

First, Kennedy and Stanley basically present *the average+N* to be ambiguous between a ‘typical/concrete’ and an ‘abstract’ reading. A case is disambiguated in terms of the predicate it takes, whether, for example, a given instance of the restriction may realise the property attributed by way of the predicate, e.g., whether a particular American could have 2.3 children or not. So, ‘typical/concrete’ *average* couldn’t by itself license (be sufficient for) a covert ‘abstract’ *average* in a subsequent clause; the relevant predicate would have to be in place to render the ‘abstract’ reading coherent. Furthermore, there is nothing the least bit incoherent in the average American having two children, from which it does not follow, on the abstract reading, that a particular American has exactly two children, but many might, of course. So, it appears to be just the semantic absurdity of the property expressed by the predicate being attributed to an instance of the restriction that triggers the ‘abstract’ reading. This indicates, in fact, that there is no ambiguity specific

to *average* at all, and certainly not one that requires a special syntax. Whether *average* is read typically or abstractly just depends on the kind of predicate employed, i.e., whether or not it can hold of instances of the complement inside the DP (more on this below).

Secondly, if the ‘typical/concrete’ *average* did alone license covert ‘abstract’ *average*, then some constraint must be in operation to stop every use of ‘typical/concrete’ *average* priming a covert *average*. This is just the overgeneration worry generalised. Again, what precludes such rampant ambiguity is surely the character of the predicate in the first place, not any priming an antecedent *average* might trigger.

A general problem now arises. If what has just been argued is correct—that ‘abstract’ *average* is licensed by at least the choice of the ‘right’ abstract predicate—then a straightforward compositional account of the constructions is surrendered, for how the DP is construed will depend on the predicate, not on any intrinsic property of the DP, let alone the occurrence of ‘typical’ *average* (more on this below).

As things stand, then, it appears that no coherent licensing conditions hold for the occurrence of a covert *average*. If that is so, then the copredication cases remain as outstanding problems for the proposal.

A further problem is the position and category of the would-be covert *average*. In the above examples, it appears as a VP adjunct, but it is not easy to see where it would appropriately occur in cases such as exemplified in (10):

(10)a The average American drinks too much, which is not good for his 2.3 children

b Discounting immigration, population growth will soon be reversed in the US.

Health care and tuition costs are rising and the average American can no longer

afford to raise his 2.3 children in the manner of twenty or thirty years ago

In these cases, the putative covert adjunct could not modify *2.3 children* without giving the wrong reading (the children are not average). The adjunct may modify the VPs (TPs) (*is not good for his 2.3 children/ can no longer afford...*), but this, again, gives the wrong reading, i.e., one where the drinking, on average, is not good for the children or where, on average, the raising of 2.3 children can't be afforded.⁹ (10) also indicates a binding problem for the Kennedy and Stanley proposal, to which I shall turn to in §4.3.

It is, of course, unsurprising that such cases should be problematic, for they instance the more general phenomenon of copredication or acceptable zeugmaticity that appears to resist any straightforward ontological construal (Chomsky, 2000; Pietroski, 2003; Collins, 2009). Moreover, copredicational cases naturally occur with generics, of which *the average American is/has F* is a case, I think, as I shall argue below.

4.2: Joe Shmoe/Joe Bloggs

The second empirical problem for the account is the phenomenon of apparently syntactically simple expressions that semantically behave much as *the average+N* does, viz., names such as *Joe Shmoe* and *Joe Bloggs*.¹⁰ Both names are naturally read with a 'typical' construal. Kennedy and Stanley do not venture an account of typical *average*, so we may leave that construal of the names as an independent problem. The names, however, also admit copredicational instances of typical and abstract construal:

(11)a Joe Shmoe drinks too much, which is not good for his 2.3 children

b Joe Bloggs eats take-away on the weekend, struggles with his mortgage, has 2.3

kids, and drives a second-hand car.

The problem Kennedy and Stanley's account faces as regards the above examples is here compounded, for one is now obliged to treat *Joe Shmoe* as licensing covert 'abstract' *average*, even though it is syntactically atomic, i.e., there is no token of 'typical' *average* to license abstract *average*, even if the item could be found a covert place in (11). Both names may seem ripe for a descriptive account, given acceptable paraphrases of them. The referential defectiveness of the names is not ameliorated, though, merely by our ability to paraphrase away the names, for such a paraphrase must be linguistically licensed if the defectiveness is to be explained while preserving the ontological import of semantics. Mere evasion of referential defectiveness does not militate for a linguistic account. Moreover, it is not obvious how the putative descriptive content would work in interaction with binding. At any rate, if the names are syntactically simple, they cannot be treated along the lines proposed for the compositional *the average+N* phrases.

4.3: Binding

Finally, it is curious that Kennedy and Stanley do not relate their analysis to binding phenomena, which pose a fundamental problem for the account, and which motivated Chomsky's (1981) anti-externalist construal of semantics in the first place. Consider the cases in (12):

- (12)a The average American has 2.3 children, but he only has 1.4 pets
- b The average American loves himself more than his 2.3 children
- c The average American and his 2.3 children live in a semi-detached property

(12a) appears to be a straightforward case of anaphora (not bound variable anaphora). In particular, the antecedent of the pronoun must be *the average American*, which means that the whole DP must be a c-commanding semantic unit relative to the pronoun. It is not clear how this squares with Kennedy and Stanley's account under which *average* takes the restriction of its host DP as an independent argument and is not semantically a unitary phrase. The answer here, however, might involve a covert *average* in the lower clause, assuming that is acceptable. (12b) poses the same kind of problem regarding both the reflexive and the possessive pronouns, with the additional problem that the case is similar to (10), where it is difficult to see how a covert 'abstract' *average* could take the appropriate argument while the right binding relations are preserved. The same problems and an additional one arise with (12c). Here, according to Kennedy and Stanley's account, the numeral must move out of its coordinated position, but this is in violation of the island constraint previously discussed that bans movement out of one of a pair of coordinated positions. A covert *average* within the coordinated DP might save the analysis, with the numeral not scoping over the whole coordinate phrase, but it remains unclear what the position of the covert *average* must be in such a case. The DP must be preserved as the argument of the predicate, but the DP itself contains the copredication.

The problems I have raised for Kennedy and Stanley's analysis are not decisive, but they are not trifling either: the problems (i) potentially undermine the very idea that there is a unique 'abstract' *average* that demands a distinct analysis and (ii) raise doubts about the analysis being able to handle binding phenomena, which were the very cases that concerned Chomsky.

5: *The average American*: A non-sensible proposal

There are, then, outstanding problems for Kennedy and Stanley's analysis of 'abstract' *average*. I suspect that a somewhat simpler analysis is available that is not beset by the problems enumerated. The analysis I shall offer presents the behaviour of 'abstract' *average* as being, in effect, not part of narrow semantics, determined by the interpretation of syntax, but an effect of *coercion* triggered by the reading of particular predicates in relation to the relevant DPs. This is in line with the kind of two-tier model suggested by Chomsky and Krifka as advertised above. The account is designed to evade the problems that arise for the Kennedy and Stanley account and be otherwise consistent with the general alternative model that dissociates externalist ontology from semantics.

As suggested above, I think that a profitable way of approaching *the average+N* is that it is subject to *coercion*; in particular, that the 'abstract' construal marks a departure from the default 'typical' reading of *average* triggered by way of the use of a predicate that cannot hold for instances of N, i.e., a predicate that pertains to abstracta rather than concreta. To say that *the average American* is default construed as typical means that speakers construe the expression as pertaining to Americans that are typical or characteristic along some relevant dimension, where, per (ET), at least one American will satisfy the predicate that combines with the DP (cp., Carlson and Pelletier, 2002). Given the choice of a kind-selecting predicate that cannot be satisfied by any given American, the DP is coerced so as to be understood as pertaining to some abstraction over the N. One may, of course, seek to express this basic idea in externalist terms, as a referential type shift from a count noun to a kind expression. My claim, as we shall see, though, is that whether one gets a typical or abstract reading is independent of the basic syntax of *the average+N* phrase, which remains invariant. The syntax gives rise to distinct

conceptual structures, neither of which by itself fixes reference, but which do record the conceptual distinction between typical and abstract. To render the proposal in externalist terms, though, is to operate at the level of post-linguistic processes, for, if I am right, a speaker is free to figure out how an abstract construal can be true without the world containing child-bearing abstract kinds, even if the conceptual structure the interpretation of the syntax delivers specifies exactly such a kind. In effect, then, abstract *the average+N* is rendered as referentially defective, for it is linguistically understood as picking out an abstract kind, even though there is no such entity with 2.3 children, and the speaker knows this.

It might be thought that if the abstract reading is coerced, then at least the typical reading must be fixed by language alone. This is not so. All that is meant here is that the abstract construal of *the average+N* depends upon a particular choice of predicate in a way the typical construal does not, because it remains available even with coercive predicates. As an adnominal adjective, *average* cannot modify an N to produce a kind construal independent of a particular abstraction-selecting predicate; otherwise put, we do not have overgenerated readings whereby, say, *The average American drives a Ford* may be read as false on the basis that abstracta don't drive. In distinction, *The average American has 2.3 children* does have an absurd reading, where the subject is read typically. So, *average* modifies count readings, and since the bare NP *sans* the modifier admits both typical and abstract readings, as we shall see, the abstract coercion depends just on the choice of predicate in relation to the count noun, not on the intrinsic semantics or syntax of the complex nominal *the average American*.

Assume, then, as Chomsky (2000) suggests, that *average* is not in any way syntactically special, and serves as an adjunct to a nominal projection within a DP much as any other adjective. Ignoring much complexity, we may depict the relevant structure as in (13):

(13) [DP The [NP [ADJ average] [NP American]]]₁₁

Again ignoring irrelevant detail, we can readily take (13) to serve as an agentive argument for any relevant verb/predicate:

(14) [TP [DP The [NP average [NP American]]] [T' [VP drives a Ford]]]

For (14) to have a ‘typical’ reading is for it to be mapped onto a conceptual representation that, in effect, renders the adjective as a quantificational relation over Americans and Ford drivers. We can assume that this follows from the lexical content of *average*. Some structure approximate to (15), therefore, will be the relevant conceptual representation:

(15) (*Gen x*)[American(*x*) ∧ P(*x*)](Drives-a-Ford(*x*))

It bears emphasis that I intend (15) to be a depiction of a conceptual structure, not logical form in any syntactic sense. (15), in effect, constrains the interpretations that can be had of (14) to be such that, typically, Americans who satisfy some cluster of properties P that unpacks the content of *average*, will drive a Ford. It is then a further process to arrive at ‘real semantics’, as Chomsky calls it, whereby a speaker can assess what would, in fact, make it true on a given occasion.

If we stick with the same syntax as (14), abstract *average* must be arrived at differently. Thus, (16) cannot be the right structure for *The average American has 2.3*

children, for a competent speaker-hearer knows that there can be no x that has 2.3 children, so denies the existential entailment (16) makes available:

(16) $(\exists x)[\text{American}(x) \wedge P(x)](\text{Has-2.3-children}(x))$

In effect, then, if a speaker treats *The average American has 2.3 children* as true, then she must at least cancel any instance of the generalisation, which are not merely empirically false, but incoherent, and so the relevant syntax is mapped to an abstraction that matches the abstraction of the predicate. In other words, the DP is coerced out of its normal mapping in order to support the predicate that fails to hold under the normal mapping:

(17) $(\exists k)(\forall x)[(x \in k \rightarrow \text{American}(x) \wedge P(x)) \wedge \text{Has-2.3-children}(k)]$

(17) is a conceptual representation that constrains interpretations of the relevant sentence to be thoughts concerning a kind whose members are only Americans who satisfy some relevant properties (being adults, etc.), and that kind has the property of having 2.3 children. I don't want to make any firm commitments here about the nature of kinds in general, save that kinds have members, but predications to kinds qua kinds do not entail that any individual members have the property predicated, and that the property holding of the kind is not reducible to some shared property of the members (see note 15). So, the kind that I propose is the average American in the abstract sense is a statistical object that satisfies the features just described. What renders the object statistical is precisely that it has a property determined by a calculation over some class of objects rather than some shared property of all the objects. It might now seem, though, as if no real progress has been made, for surely abstract kinds do not have 2.3 children any more than individuals do. I shall have more to say on this matter shortly. For the moment, my idea is that the kind has been coerced into existence, if you will, simply because an apparent truth is at

hand that is not supported by the default mapping given the predicate, i.e., it is not conceptually available to think of a particular American as having 2.3 children, so some other concept must be mapped onto the DP that rules out that possibility but which somehow supports the predication. It remains to be seen how a speaker-hearer might render (17) true. That is my basic story.

As an analogy for the coercion effect I have in mind, consider how a proper name may be coerced into a mass term by way of an appropriate predicate (so-called *grinding*):

(18) After the steam roller had finished with him, Bill smeared all over the road

Assume that *Bill* is default mapped onto an individual concept, i.e., if you want another reading, some functional items must be merged, such as plurality feature or a determiner (*Most Bills*, etc.). No individual, however, can satisfy a mass-selecting predicate such as *was smeared*. Hence, *Bill* is coerced from an individual to a mass construal in the sense that *Bill* gets mapped onto a mass concept. The individual and mass construals are related, of course, because *Bill* is now read as pertaining to the mass that constitutes Bill the individual, not some or other mass.

As it is, Kennedy and Stanley (2009, p. 10) reject a generic account of ‘abstract’ *average* on the basis that ‘characterising’ generics quantify over individuals, and so do not pattern with the abstract construal of *the average+N*, i.e., characterising generics can have instances such that at least one instance of the subject can satisfy the predicate (even if none happens to), whereas ‘abstract’ *average+N* does not have any coherent instances (this is an expression of the existential test (ET) discussed above). Kennedy and Stanley are right about the semantic properties of ‘abstract’ *average* in this regard, but wrong about generics, I think, so the latter might offer a model of the former after all.

It is certainly true that there is something of a consensus that a large class of generic statements have a standard tripartite quantificational structure: in the simplest case, a *generally/typically* covert operator (*Gen x*), a restriction, formed by the subject nominal, and a scope, formed by the predicate. Clearly, on such an account, individuals satisfying the restriction and scope are quantified over; hence it is that the coherence of instances of the predication are licensed, even if not entailed.¹² If we bracket the complex issues pertaining to the syntactic plausibility of *Gen* as a covert operator (see Collins, forthcoming b), such an understanding of genericity is just another way of saying that standard generics are brought about by a mapping to a structure akin to (16) above. We also have kind-selecting predicates, however, where the DP is not construed as a collection or a plurality, and so does not involve generic quantification over the individuals of the kind. I am suggesting that we should construe ‘abstract’ *average+N* in precisely this way as (17) depicts.¹³ This approach does not involve burdening the world with abstract kinds as well as individuals; abstract kinds of the sort (17) trades in, I propose, have no categorical fidelity to the way the world is anyway (whether the world contains kinds at all is an issue for general metaphysics and science, not semantics). Rather, as indicated, the relevant kind is an aspect of a conceptual structure that is mapped onto the syntax because of a failure of the predication to have any possible instances. The subject is construed as expressing a statistical entity dependent on the predicate of which it is the subject. Still, if the sentence is treated as expressing a truth, then a speaker must find a way of relating the conceptual structure to the world, but neither the syntax nor the conceptual representation determines how the world must be such that the relevant utterance might be true.

In general, the semantics of kinds and their relation to ‘typical’ or ‘characterising’ generics is complicated and controversial, but it is not unreasonable to view the kind construal as often coerced from a non-kind generic reading given the choice of certain predicates, especially where the DP is not a bare kind term, such as a Latin name of a species. This is not to claim that all generic DPs are systematically ambiguous, no more than that DPs generally are systematically ambiguous. The claim, rather, is that the construal of a DP is relative to its predicate. Hence it is that non-kind DPs can express kinds given the choice of a predicate that cannot hold of instances of the nominal contained in the DP.¹⁴ Consider:

(18)a Lions evolved into ferocious predators

b The dodo is extinct

c White rhinos are widespread

None of these admit a coherent instance of a particular token of the kind possessing the respective predicated property: no particular lion could have evolved (in the relevant sense), just as no particular dodo can be extinct, and no particular white rhino can be widespread.¹⁵ Whether we construe the DP of a generic statement as expressing an abstract kind or a collection of individuals depends upon the predicate, just as whether we construe *the average American* to be typical or abstract depends upon the predicate, the former if we are talking about the car s/he owns, the latter if we are talking about her/his 2.3 children. *The average American* accepts the predicates exemplified in (18):

(19)a The average American evolved into a couch-potato

b The average American is extinct/widespread/rare in this part of New York

I take it that these examples exhibit abstract construals of *the average American* insofar as they do not support any existential instances, much as kind generics do not, i.e., no particular American evolved into a couch-potato (perhaps the current generation were couch-potatoes from birth) just as no particular American can be extinct or widespread in a given part of New York. *The average American* in such cases appears to pick out a statistical kind, much as it does when subject to the predicate *has 2.3 children*.

This reading is further supported by the fact that (i) construed either typically or abstractly, *the average+N*, just like a generic, only accepts stative predicates, and that (ii) bare plurals, definites, and indefinites can be coerced as abstract, just as the same range of DPs can be generic:

(20) The/an American(s) has (/have) 2.3 children₁₆

On an out-of-the-blue reading, the acceptability of the cases in (21) is clearly degraded given the absence of *average*, and it is hardly as if *average* is pleonastic in some sense. Still, an abstract construal via a relevant predicate is readily available where, say, one is comparing birth rates between populations:

(21) Due to government intervention, Chinese birth-rates remain stable in the face of immigration, whereas in the US, immigration has affected the rates. The figures for China remain constant, with a Chinese family having 1.9 children, whereas in the US a family's children has risen to 2.3.

So, the abstract construal appears to be possible just as if it were a species of kind construal, and the availability of the construal turns on the character of the predicate, just like a generic, rather than an intrinsic semantic property of the nominal (it doesn't follow that a kind construal is always available via coercion; see below). I am assuming, note,

that *the average*+N is default construed as typical because of the presence of *average*, and can be coerced into an abstract construal. The cases in (20) and (21), however, show that the same content can be expressed without the presence of *average*. So, the abstract kind reading I am proposing is not dependent on *average*, but is dependent on the predicate being one that coerces an abstract reading. If this is so, then, the abstract reading of *the average*+N does not depend on the content of *average* in the way Kennedy and Stanley (2009) propose. Were it to, the availability of (20) and (21) would be anomalous.¹⁷

The proposal on offer also avoids the problems I raised for Kennedy and Stanley's account. First, since the account involves no QR operation, it does not turn on problematic issues pertaining to the proper interpretation of QR and the unconstrained positing of extra projections.

Secondly, copredication is straightforwardly accommodated too on the current proposal. Just as *book* may simultaneously be construed as referring to a type, a token, and content, depending on the predicate, so *the average American* may simultaneously be construed as typical and abstract, which is, in fact, characteristic of generics. It is the relevant predicates that coerce the readings, given the openness of the nominal to the copredication. The nominals by themselves as lexical items do not refer to one thing or another. Thus, assume that *The average American drives a Ford and has 2.3 children* has a syntactic structure approximate to that in (22) (irrelevant details elided):

(22) [TP [DP the average American] [&P [_{vP} <the average American> drives a Ford] [&' and
[_{vP} PRO has 2.3 children]]]]¹⁸

We may take the subject of *drives a Ford* to contribute to a generic-like conceptual structure, and the PRO subject of *has 2.3 children* to map onto a kind structure via coercion by the predicate. This issues in a conceptual structure that amounts to a conjunction of (16) and (17):

(23) $(Gen\ x)[American(x) \wedge P(x)](Drives-a-Ford(x)) \wedge (\exists k)(\forall x)[(x \in k \rightarrow American(x) \wedge P(x)) \wedge Has-2.3-children(k)]$

So, the one overt DP is interpreted in distinct ways given the different subjects of the relevant ν Ps. PRO is cointerpreted with the overt DP, and so is default construed as the same, but the distinct predicate of the PRO subject coerces a distinct construal at the level of conceptual structure. It will be noted that (23) does *not* entail that there is something that both drives a Ford and has 2.3 children, just as (22) does not.

This proposal runs counter to Liebesman’s (2011) ‘simple’ account of generics. Liebesman objects to the standard tripartite account of generics precisely on the basis of the copredication of individual-level and kind-level predication:

(24)a Mosquitoes are widespread and irritating

b Lions, which evolved into ferocious predators, are in my garden¹⁹

Liebesman suggests that the right account is that all generic nominals are kind-referring, but are semantically underspecified as regards how the predications can hold of the kinds as inherited from the members. Ironically, this proposal precisely fails for kind-selecting predicates where no inheritance relation can be specified, such as *evolve*, where for a kind to evolve does not involve a change in its members, but a redistribution of properties over members of the kind. In this sense, a subject of *evolve* is a statistical entity (see note 14). Equally, Leslie’s (2015, pp. 45-6) riposte to Liebesman in defence of the standard

account of genericity is problematic. Leslie argues that copredication cases such as in (24a) are to be semantically analysed as involving a conjunction of two clauses, one kind-predicating, the other type-lowered from a kind predication to an individual-level predication. For example:

(25) $\lambda x [\text{widespread}(x)](\text{mosquitoes}) \wedge \text{GEN}y [(\Downarrow(\text{mosquitoes}))(y)](\text{irritating}(y))$

The structure in (25) accurately depicts the intended meaning of (24a), and is akin to the structure offered in (23) in virtue of both of them distinguishing between the subjects of conjoined predicates (ν Ps). The problem is that there is no obvious linguistic licence to treat the predication of (24a) as clausal conjunction featuring two tokens of *mosquitoes*, which may thus be differentially construed via a selective type-shifting to just one of the tokens. Perhaps, as with (22), the hypothesis of a covert argument for *irritating* is not so implausible, although the case is much less clear between a predicative participle and a full ν P. The proposal, however, *pace* Leslie's apparent intentions, doesn't offer a general model of copredication, for copredication occurs both within a DP, in the shape of relative clauses, as (24b) exemplifies, and via modifiers and across anaphora.

(26)a Widespread irritating mosquitoes prove to be a pain to the tourists

b Widespread mosquitoes are a nuisance, but a good repellent works on them

The dual modification in (26a) is not reducible to distinct individual-level and kind-level predications, for the irritating mosquitoes are the widespread ones; besides, again, there is no clear syntactic rationale for the relevant reanalysis. The problem with (26b) is analogous. The anaphoric *them* is dependent on *widespread mosquitoes*, so no type lowering is available for just *mosquitoes* as a construal of *them*. (See Collins, forthcoming b for greater detail on these issues.)

Thirdly, for the same reason, binding does not create any problems: the pronoun is simply referentially dependent on the kind expressed by the *average* DP, just as one would expect for normal cases. Fourthly, *Joe Schmoe/Bloggs* cases are easily accounted for. They are simply names that pick out the statistical entity expressed by *the average American/British* person abstractly construed; hence it is that the terms do not admit non-stative predicates.

Let us now return to the issue of the truth of (18), which appears to commit the speaker-hearer to abstract kinds having 2.3 children, which might seem as unacceptable as individual Americans having 2.3 children. For instance, Koslicki (1999, p. 449) raises the thought that *the average man* denotes a kind only to dismiss the idea, for ‘there is no such thing as the average man (who has all kinds of strange properties, such as having 2.2 children)’. This thought is correct insofar as the average anything is not a natural kind out there anyway awaiting our discernment of it. On the current proposal, the abstract kind *k* specified in (17) is specified as part of a conceptual representation that is mapped onto the relevant syntactic structure. The kind, therefore, is precisely not an independent ontological commitment of the speaker-hearer, no more than a speaker employing the concept of *Donald Duck* is therefore committed to Donald Duck, more about whom shortly. My proposal is that the average American is a statistical kind, whose existence we come to acknowledge only by way of taking as true that something has 2.3 children (or similar properties). That is, only a statistical kind could have such a number of children, for such a property is statistical. So, *The average American* abstractly construed is referentially defective insofar as no *thing* can have 2.3 children in any normal sense, but we equally understand that the sentence is made true by the distribution of offspring

over relevant Americans, which is just our acknowledgment of the statistical nature of the claim at hand and so our treatment of the average American as a statistical kind just for the purposes of understanding how the relevant sentence could be true, i.e., *pace* Koslicki, the kind construal does not suggest that an abstract entity gave birth to or adopted 2.3 children. We do not treat the sentence as obviously false or some kind of category mistake, because it is easy to work out what the coercion involves, much like the example of mass coercion above. Indeed, as intimated, *evolve* works much the same way, where if one talks of tigers evolving, one is not, unless confused, talking about a change of the properties of individuals or a change in some abstract object; instead, one is making a statistical claim about a change in the distribution of properties over a population, precisely where no individual does evolve. Thus, the property of having 2.3 children can hold for some class of Americans, not understood as a collection of individuals (hence the coercion), but as a statistical entity ('a tissue of statistics', as Nabokov puts it) derived from the class of Americans relative to the predicate chosen. So, *the average American* expresses a statistical object precisely because it has a statistical property, such as having 2.3 children.

The basic point here is to separate a story of truth making from bare truth conditions a competent speaker-hearer understands. The former can be highly complex and depart in various ways from what a speaker may have an understanding of, whereas the latter come with basic linguistic competence (cp., Azzouni, 2010; Asay, 2013; Crane, 2014). It doesn't follow, of course, that the truth-making conditions will be obscure to a subject who understands the relevant sentence, although in many cases they might be. Thus, in the present case, if I am right, a subject's semantic competence will deliver up a typical

construal of the DP, which is then coerced into making reference to a statistical entity, which, in turn, may be eschewed in one's recognising that the truth of the matter devolves upon how things are for the American individuals. The crucial point is to distinguish between these three steps. Note, too, that the occurrence of *average* is not essential for the abstract construal, for one can have exactly the same abstract kind reading in its absence—witness (20) and (21). In effect, Kennedy and Stanley are more or less right about what would make the relevant *abstract average* sentence true, but it doesn't follow that the syntax or logical form of the sentence reflects the world in that way.

A critic might respond as follows.²⁰ 'A simpler account that is suggested by your proposal is that *the average American* does refer to an abstract kind; so, *strictu dictu*, the truth of the relevant claims does involve ontological commitment to such kinds. However, there are no such entities, or at least none that undergo childbirth. So, a speaker of the relevant sentence would have said something false, but still somehow manages to communicate a coherent thought concerning the distribution of off-spring over Americans'.

The complaint is well taken, but offers a more problematic account than the one on offer. First, it is precisely a virtue of the present account that *the average+N* does not have any determinate reference as a matter of language alone. What creates the abstract kind reading is the presence of a predicate that coerces the syntax to be mapped to a structure that involves an abstract statistical kind for a speaker who wants the sentence to be coherent, indeed, true. Such a speaker is able, in the normal run of things, to work out what would make the sentence true, even if they demur when it comes to ontologically

committing to statistical entities, in the manner indicated. A speaker needn't in any sense think the sentence is false because there are no abstract kinds with 2.3 children. Rather, she may reason that what *makes* the sentence true is simply not specified in her linguistic understanding of the sentence.

In fact, as intimated in §2, the same approach works for the broader class of referential defective expressions, and here I follow the lead of Azzouni (2010) on fiction and myth. Consider:

(27) Donald Duck made millions for Walt Disney

Suppose a speaker reckons (27) to be an outright truth, even knowing that there is no Donald Duck. Suppose further that the facts of natural language are such that the syntactico-logical form of (27) is insensitive to the inexistence of Donald Duck, so that that syntax of (27) is mapped onto a conceptual representation featuring a Donald Duck notion just as if the sentence were *John Wayne made millions for John Ford*. This is not a conundrum, for a speaker can reckon (27) to be true and deploy a representation featuring a Donald Duck concept, and yet appreciate that what makes the sentence true will be a set of complex factors involving film and merchandising rights, and so on, but no Donald Duck, for there is no such entity. Analogously, a speaker can figure out what makes true *The average American has 2.3 children* by way of deploying a concept of a statistical abstract kind *k* without being committed to the world containing any such independent kind. Space precludes a thorough assessment of the treatment of the panoply of referential defectiveness along these lines, but I think it is clear that the basic approach is both coherent and fundamentally divorces externalist ontology from semantics.

A final concern is meta-theoretical: ‘What is on offer is not so much a theory, but merely a packaging of the phenomena, which makes appeal to unanalysed pragmatic processes of construal. We are asked to trade in a theory for a description.’²¹ The complaint here can be expressed as an overgeneration worry: if no linguistic constraints are placed upon the interpretive role of the posited mechanism of coercion, then, contrary to fact, nigh-on any argument should be subject to coercion by the relevant predicate. As it is, the various relevant linguistic forms bear witness to systematic patterns of interpretation. Consider the following:

(28) The/a happy American has 2 children

(28) has a kind construal, even though the predicate clearly cannot coerce this construal in the way *has 2.3 children* does, for individual Americans may have 2 children (lots do).

On the other hand, the cases in (29) don’t have kind construals:

(29)a. Bill has 2 children

b Some happy American has 2 children

c Few happy Americans have 2 children

The question is why there is such a pattern, if the availability of a kind reading or not is at the behest of an extra-linguistic mechanism of coercion.

First, in general, the complaint being entertained presupposes that if coercion is an extra-linguistic mechanism, then it operates free of any linguistic constraints. This claim, though, should be rejected. Everything I have said so far is perfectly consistent with there being constraints from both syntax and lexical semantics on when coercion may apply. Equally, nothing I have said entails that the relevant kind construals can only be triggered by the sort of semantic anomaly exhibited by taking the predicate *has 2.3 children* to

apply to individuals. That is just one case where coercion does apply to deliver a particular kind construal. In short, the account on offer is highly constrained; it is just not constrained by way of fixing the difference between typical and abstract readings on the basis of a difference of syntax or linguistic logical form. With these points in mind, let us briefly turn to the phenomena in (28) and (29), starting with the latter first.

It is true that the cases in (29) do not have generic or kind readings, and so coercion does not and cannot induce such readings. Yet why should anyone imagine it could? I take it to be common ground that names (in singular use) and quantifier DPs resist any generic or kind reading because the DPs express precise cardinality values per the standard generalised quantifier treatment. The meaning of the DPs in this respect, therefore, is fixed and so excludes any possible coercion to a non-precise cardinality as a matter of lexical semantics. Besides, coercion should not apply anyway because no anomaly arises between the predicate and subject DP. In sum, coercion is not a magical process that may turn any DP into a kind term, but a process that may apply only where the language itself (syntax + lexical meaning) does not fix an interpretation. The account is, therefore, constrained precisely by what we independently know about the possible interpretations of linguistic structure.

Turning to the (28) cases, we reach a similar conclusion. All of the cases are ambiguous independently of any issue arising with coercion. In particular, DPs headed by articles (definite and indefinite) can be construed as singular (existential) or generic or kind in part depending on the predicate, where I take none of the options to be default, in distinction to how *the average American* is default typical for the reasons given above. So, here, it is not so much that an anomaly triggers kind coercion, but that a kind reading

simply remains an option given a choice of predicate such as *has 2 children*. On the other hand, a kind reading is excluded by a stage-level predicate, such as *is in my garden*. Again, then, the account on offer is perfectly happy with the pattern of constraints to which the cases in (28) bear witness.

Of course, nothing in my account explains why some DPs admit kind readings whereas others do not; why, that is, the articles are less constraining than the quantifiers in this respect, but otherwise cross-classify *vis-à-vis* other properties (for example, weak determiners, including *some* and *a/an* can wide-scope out of islands, unlike, *every*, *most*, etc.). The complaint being rebutted is only that the account is unconstrained, not, absurdly, that it doesn't explain phenomena that falls outside of its brief. I take the divergence in available readings for quantifier DPs and article-headed and bare plural DPs to be an open question; for present purposes, it suffices that the account on offer is perfectly well constrained by whatever factors explain such differences in the ways just rehearsed.

The balance sheet of the proposal as it stands looks pretty healthy. It deals with all of the phenomena smoothly, including copredication cases; it patterns *the average+N* with generics; it suggests how the 'typical' and 'abstract' construals are related; and it does not require any peculiar syntax or semantics just for the one construction. On the debit side, as the orthodoxy would have it, the analysis does not deliver a sensible ontology as a reflex of semantico-syntactic structure, but, if Chomsky and others are right, that is exactly as one should expect once the linguistic business of semantics is separated from the metaphysical business of ontology.

6: Conclusion

The status of ontology in relation to semantic theory cannot be settled by getting straight the semantics of *the average American*. Indeed, if one is minded to buy a general conception of lexical items as being polysemous, much as Chomsky (2000) and others commend, then the status of our little DP can seem trifling. Kennedy and Stanley are right, though, that the fate of the phrase is significant, for it offers a *prima facie* case of referential defectiveness that cannot be resolved, as I argued above, by the kind of standard moves employed towards fiction and myths that seek to preserve the alignment of semantics with acceptable ontology. Moreover, the DP is semantically and syntactically complex, and so whatever the phrase does mean should be compositionally resolvable. Thus, the case poses a severe challenge quite independently of whatever attitude one may have towards the general polysemy of lexical items. If, then, Kennedy and Stanley can show that, in fact, the best semantico-syntactic account of the DP renders it non-defective, not only is a significant case for semantic anti-externalism knocked out, but a general method seems to be supported of furnishing sophisticated analyses of would-be referential defectiveness that reveal the cases to be illusory. If I am right, the dialectic decidedly goes in the other direction in favour of Chomsky. Notwithstanding its ingenuity, the analysis Kennedy and Stanley offer suffers from a range of problems due to the very devices that are designed to reveal the underlying logical form of the DP not to be a unitary referential term. Additional and questionable resources are required to account for copredication and binding in the form of a covert *average*, but it is exactly copredication and binding considerations (*inter alia*) and lead Chomsky and others to a principled separation of ontology from semantics independently of any issue pertaining to polysemy. Still, Kennedy and Stanley are right: it is not enough for any party to point to

this or that phenomenon and imagine that the apparent properties of the relevant constructions will settle such a high-level issue as the ontological import of semantics. In this spirit, I offered an alternative account of *the average+N* DPs that avoids the problems that beset Kennedy and Stanley's proposal while crucially retaining the referential defectiveness of the DP. *The average American has 2.3 children* can be perfectly true (I assume it is true), while the DP fails to refer to anything that might be deemed to be an external entity that has 2.3 children. What makes the sentence true is a complex statistical fact that the sentence does not encode. Rather, a speaker of a token of the relevant sentence coerces the DP into expressing a statistical abstraction that has no instances that can satisfy the predicate, just like a kind statement can be coerced given a suitable kind-selecting predicate. The speaker does not, therefore, have a crazy metaphysics, for she understands, we may suppose, how to spell out what makes the sentence true, which involves a systematic departure from what she knows merely by dint of being competent with the sentence, an aspect of which is to treat the DP as a unitary semantico-syntactic phrase. In effect, Kennedy and Stanley get the metaphysics of being abstractly average right, but Chomsky's challenge was to show how the right metaphysics, or any sensible metaphysics, is fixed by the relevant linguistic properties, a challenge that remains standing.²²

Notes

¹ Stanley (2001) presented a quite different analysis of *the average+N* construction, which, I take it, is to be read as superseded by the later analysis. For one thing, the analysis posits a covert measure adjective, whose role is now taken up by *average* itself.

² The *the average*+N construction is also discussed by Melia (1995) and after him by Yablo (1998), who both take it to be referentially defective. See Stanley (2001) for effective criticism of their wider aims.

³ In response to Ludlow's (1999, p. 174) reporting on Chomsky's binding-based reasoning, Stanley (2001, p. 58, n. 36) rejects the significance of binding due to the cases being 'lazy', and so effectively succumbing to an e-type account. It bears emphasis, however, that one of Ludlow's examples from Chomsky is not lazy, but genitive:

(i) That his income is falling bothers John Doe

Ludlow (2003) appeals to the potential for an E-type account to at least complicate Chomsky-style arguments as regards anaphoric reference to flaws and such like.

⁴ Unfortunately, Chomsky's 'autonomy/independence thesis' has often been read as an anti-semantic doctrine. In fact, from Chomsky's earliest work onwards, the doctrine has consisted of two methodological precepts, neither of which is anti-semantic: (i) syntax is not reducible to semantics and (ii) progress in semantics can be made by discerning the syntactic constraints on what phrases/sentences can mean. For discussion of the history, see Collins (2008, pp. 39-45).

⁵ For detailed discussion of other problems with the adverbial analysis, see Carlson and Pelletier (2002, pp. 82-4).

⁶ Heim and Kratzer (1998, p. 188) do wonder what empirical evidence there could be to rule out their extra projection, and we may suppose that Kennedy and Stanley would likewise wonder what evidence would tell against their extra projection as a reflex of QR. Interestingly, Heim (1997) presents ellipsis phenomena that militate for a uniform formula analysis of VPs and TPs including QR-extraction sites, which does not involve

λ -abstraction either in the syntax or in interpretation. However the facts pan out, the extra projection is *conceptually* problematic precisely because it lacks any syntactic rationale. For example, the relevant constructions do not feature relative pronouns, which are readily treated as λ -terms, or tough-movement, which is often analysed as involving an operator. As it stands, therefore, Kennedy and Stanley's proposal is syntactically *sui generis* or ad hoc, as is Heim and Kratzer's proposal, for the posited projection is not *really* a projection at all, for it lacks a head and any syntactic life; its sole *raison d'être* is to turn syncategorematically a clause into a monadic predicate.

⁷ Carlson and Pelletier (2002) argue that *average*+Ns constitute a distinctive class of expression that is 'ecumenical', i.e., neither a quantifier nor a singular term. *Average*+Ns, according to Carlson and Pelletier (2002, p. 87) 'do not refer—either to an entity or to a set', where 'refer' picks out our ordinary, non-technical notion. Still, *average*+Ns do have semantic values that contribute to the semantic values of their host structures. Such values are sets of properties of a kind that admit the calculation of an average (individuals cannot be averaged). Which properties these are remains obscure, but Carlson and Pelletier (ibid., pp. 91-2) assume that *average* introduces a partition function that selects those properties of the value of the modified noun that may be averaged (weight of Americans, kind of car owned by women, etc.). The account thus offers a unitary account of 'abstract' and 'concrete' construals insofar as the properties that may be averaged divide between the scalar (e.g., weight) and non-scalar (e.g., kind of car owned); so, the account is partly designed to account for copredication cases. For present purposes, I am happy to leave such an account unmolested, for while Carlson and Pelletier, like Kennedy and Stanley, seek to save *the average*+Ns for a truth conditional semantics, the result, so

far as I can see, is neither animated by nor satisfies a constraint to deliver a ‘sensible ontology’. That said, as we shall see with Kennedy and Stanley’s account, it is a mistake to put any load upon the adjective *average*, for an abstract construal of a DP is available without it.

8 The standard example of such a ban in operation is *wh*-movement out of a conjunctive object:

- (i) Bill loves Mary and Jane
- (ii) Bill loves Mary and WHO
- (iii) * Who does Bill love Mary and.

9 A reviewer suggested that (10a) is outright unacceptable. If that is right, then its deviance in fact supports Kennedy and Stanley’s proposal insofar as no position for a covert *average* is available. Informal surveys of mine indicate that speakers are generally OK with (10a), and I think (10b) offers no problem whatsoever. It is also worth noting that judgements of the relevant cases are somewhat sensitive to familiarity. For example,

(i) is odd:

- (i) The average American has 0.017 pet iguanas

There is nothing deviant about (i), however, relative to *The average American has 2.3 children*. Any problem with (10a), I suspect, is simply due to uses of nominals such as *2.3 children* being fairly restricted to certain specialised discourses where statistics are at issue; hence it is that (10b) is perfectly OK (thanks to Zoltan Szabó for discussion on these cases).

10 Expressions for enemy combatants fit the same pattern (*Tommy, Fritz*, etc.), as well as some slurs.

¹¹ The complexity elided comes in at least three flavours. Firstly, the nominal projection within the DP appears to be hierarchical, especially with regard to the placement of adjectives (see, for example, Scott, 2002; Cinque, 2010). Secondly, it is now common to split the NP projection into a covert *nP* projection that selects for an NP. Thirdly, adjectives are patently not uniform in their semantic and syntactic behaviour. For present purposes, though, we may take Chomsky's claim and my elaboration of it to be the modest one that *average* is not bespoke, but syntactically patterns with other non-gradable abstract adjectives.

¹² The consensus is not universal. Koslicki (1999) argues that generic DPs denote plurals (second-order concepts); Liebesman (2011) argues that generic subjects uniformly denote kinds—no covert *Gen* operator exists or is required (cp., Collins, forthcoming b). See Mari *et al.* (2012) for an overview of recent research on genericity.

¹³ Carlson (1977) initially argued for a uniform kind model of genericity, and Chierchia (1998) later developed the view with type-shifting against a strict dual view under which a tripartite analysis is offered for the 'typical' generics and a kind-referring analysis is offered for the remainder (Krifka *et al.*, 1995).

¹⁴ I assume that there are at least some kind-selecting predicates that obligatorily take kind subjects (terms that can only be so construed) as opposed to plural terms. It might be that some apparent kind-selecting predicates, such as *rare*, *common*, etc., are quantificational predicates, i.e., they express a second-order property pertaining to the distribution of instances of the first-order property (Krifka, *et al.* 1995, pp. 95-8). It is certainly the case that the relevant predicates can take non-kind subjects (*Elephants with*

one leg are rare). For my purposes, it will suffice for there to be uncontroversial kind-selecting predicates (e.g., *evolve*, *extinct*) that take *the average American* as subject.

15 It is tempting to think that kind predication can be reduced to a complex predication to instances of the kind or a plurality, not to the kind itself. Some such account is tempting for those theorists who think of generics as uniform. Koslicki (1999), for instance, proposes a uniform plural model for all generics as a ‘first-level’ analysis of their logical form (i.e., generic subjects denote pluralities). She seeks to deal with kind-level or second-order predicates that do not license instances of the predication by way of a separate ‘second-level’ analysis that relates the predicate holding of a plural term to a first-order predication that is intended to express something more informative about the truth conditions of the generic statements. In effect, the predicates are treated as plural predicates with an associated story of what makes the predicate true of the plurality in terms of the state of the members of the plurality (cp., Liebesman, 2011). Thus, she writes: ‘any characteristic of the species involves ultimately *something* happening to its individual members; but to express exactly what this something is will involve ‘unpacking’ the higher-order plural predicate in question’ (ibid., p. 457). So, as a ‘first-level’ analysis, *all* generic DPs express a concept or property that holds for all and only those things that are N, and the predicate of the DP expresses some property of what falls under the concept expressed by N. This form of account might work well for so-called *quantificational predicates*, such as *rare*, *common*, *widespread*, which appear to pattern much like collective predicates taking plural arguments (*surround*, *gather*, *meet*, etc.), which block particular instances and existential inferences too. The account, however, performs poorly at best with explicit kind-selecting predicates. Thus, for the kind-

selecting predicate *extinct*, Koslicki offers the ‘second-level’ analysis just mooted above that expresses only a first-order property (e.g., ‘*Extinct* holds of a plurality N iff there used to be members of N, but now there aren’t any’). Even if the ‘unpacking’ works for *extinct* (which appears not to be the case, since not any old collection can be extinct), no first-order property will serve to ‘unpack’ *evolve*. For lions to have evolved does not involve ‘ultimately *something* happening’ to individual lions. This pattern is exhibited by the following:

(i)a Tigers (with one leg) are rare/widespread/common

b Tigers (#with one leg) are extinct/evolved into ferocious predators

More generally, it is a truism that *evolve* and *extinct* do not express a property of a mere collection. Take (ii):

(ii) The large carnivores that roamed the Highlands are now extinct

This is not a claim about the fate of the individual animals, for, trivially, all animals die, but a claim about the kind to which the relevant animals belong, *viz.*, it has no extant members. Thus, I submit that *no* first-order unpacking will eliminate reference to the kind in such cases.

¹⁶ I do not here suggest that all of these forms equally submit to a generic construal in general. Since Lawler (1973), it has often been noted that indefinite DPs admit a generic construal only when the predicate expresses some kind of essential property or principled connection between the property expressed and the subject. Thus, it is pretty close to impossible to hear (i) as generic, and so it is semantically anomalous; (ii) is also difficult on such a reading, although much less so, I think. On the other hand, (iii) and (iv) are naturally read as generic:

- (i) #A dodo is extinct
- (ii) ?A madrigal is popular [on a generic reading]
- (iii) Madrigals are popular
- (iv) A madrigal is polyphonic

Although there is certainly something to this restriction, it is more of a default condition for non-kind-selecting predicates. Thus, the correct claim hereabouts is that an indefinite is always open to a generic construal with a predicate that *is* optionally generic, i.e., not stage-level (*in the garden*) or kind-selecting (/quantificational) predicates. So, perhaps no setting will induce a generic reading for (i), because of its kind-selecting predicate, but (ii) succumbs to such a reading easily:

- (v) Among the elderly, a madrigal is popular.

This issue, however, does not affect the present points, for the relevant abstract kind-selecting predicates do not fall under the potential restriction, i.e., they are ones that would admit a generic reading with an indefinite DP subject (cp., Greenberg, 2003).

¹⁷ This is predictable, if we take the abstract predicate to coerce a kind-reading of the subject, for kinds go unmodified and resist adverbial quantification (Cohen, 2001):

- (i)a #The average dinosaur is extinct
- b #Dinosaurs are typically/usually extinct

One explanation for this is that such modifiers by themselves entail the availability of an existential inference, which the predicate blocks. In our target case, therefore, *average* is rendered superfluous with the whole DP coerced into a kind reading, a fact borne out by (20) and (21).

18 For present purposes, it doesn't matter if PRO is renounced in favour of a copy of *the average American* or an ellipsis; all that is crucial is that the lower vP has a subject cointerpreted with the higher subject.

19 Conjunctive copredication of individual-level and stage-level predicates is not entirely happy:

(i) ?Lions evolved into ferocious predators and are in my garden

The existential paraphrase is clearly as bad for (i) as it is for any individual-level predication:

(ii) * There are lions evolved into ferocious predators and in my garden

20 A reviewer raised the following as an objection rather than as a positive proposal.

21 An anonymous referee raised this objection.

22 My thanks go to two anonymous referees, Zoltán Szabó, and Nathan Klinedinst for insightful and constructive criticism, and Paul Pietroski, Gennaro Chierchia, Steve Gross, and David Liebesman.

References

Asay, J. 2013: *The Primitivist Theory of Truth*. Cambridge: Cambridge University Press.

Asher, N. 2011: *Lexical Meaning in Context: A Web of Words*. Cambridge: Cambridge University Press.

Azzouni, J. 2013: *Semantic Perception: How the Illusion of a Common Language Arises and Persists*. Oxford: Oxford University Press.

Bach, E. 1986: Natural language metaphysics. In R. B. Marcus, G. Dorn, and P.

Weingartner (eds.), *Logic, Methodology and Philosophy of Science VII* (pp. 573-595). Amsterdam: Elsevier.

- Bach, E. 1989: *Informal Lectures on Formal Semantics*. New York: SUNY. Press.
- Borg, E. 2009: Must a minimalist be a semantic internalist? *Proceedings of the Aristotelian Society (Supplementary Volume)*, 83: 31-51.
- Burge, T. 2003: Psychology and the environment: reply to Chomsky. In M. Hahn and B. Ramberg (eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (pp. 451-470). Cambridge, MA: MIT Press.
- Carlson, G. 1977: *Reference to Kinds in English*. Ph.D. diss., University of Massachusetts at Amherst.
- Carlson, G. and Pelletier, F. J. 2002: The average American has 2.3 children. *Journal of Semantics*, 19: 73-104.
- Chierchia, G. 1998: Reference to kinds across languages. *Natural Language Semantics*, 6: 339-405.
- Chierchia, G. and McConnell-Ginet, S. 1990/2000: *Meaning and Grammar: An Introduction to Semantics* (Second Edition). Cambridge: MIT Press.
- Chomsky, 1975: *Reflections on Language*. New York: Random House.
- Chomsky, N. 1977: *Essays on Form and Interpretation*. Amsterdam: North-Holland.
- Chomsky, N. 1981: *Lectures on Government and Binding: The Pisa Lectures*. Dordrecht: Foris.
- Chomsky, N. 2000: *New Horizons in the Study of Language and Mind*. Cambridge: Cambridge University Press.
- Chomsky, N. 2003: Reply to Ludlow. In L. M. Antony and N. Hornstein (eds.), *Chomsky and His Critics* (pp. 287-295). Oxford: Blackwell.

- Chomsky, N. 2013: Notes on denotation and denoting. In I. Caponigro and C. Cecchetto (eds.), *From Grammar to Meaning: The Spontaneous Logicality of Language* (pp. 38-45). Cambridge: Cambridge University Press.
- Cinque, G. 2010: *The Syntax of Adjectives: A Comparative Study*. Cambridge, MA: MIT Press.
- Cohen, A. 2001: On the generic use of indefinite singulars. *Journal of Semantics*, 18: 183-209.
- Collins, J. 2008: *Chomsky: A Guide for the Perplexed*. London: Continuum.
- Collins, J. 2009: Methodology, not metaphysics: against externalism. *Proceedings of the Aristotelian Society (Supplementary Volume)*, 83: 53-69.
- Collins, J. 2011: *The Unity of Linguistic Meaning*. Oxford: Oxford University Press.
- Collins, J. forthcoming a: The copredication argument. *Inquiry*
- Collins, J. forthcoming b: Genericity sans Gen. *Mind and Language*.
- Crane, T. 2014: *The Objects of Thought*. Oxford: Oxford University Press.
- Dowty, D., Wall, R., and Peters, S. 1981: *Introduction to Montague Semantics*. Dordrecht: Kluwer.
- Elbourne, P. 2011: *Meaning: A Slim Guide to Semantics*. Oxford: Oxford University Press.
- Evans, G. 1977: Pronouns, quantifiers and relative clauses (I). *Canadian Journal of Philosophy*, 7: 467-536.

- Evans, G. and McDowell, J. 1976: Introduction. In G. Evans and J. McDowell (eds.), *Truth and Meaning: Essays in Semantics* (pp. vii-xxiii). Oxford: Oxford University Press.
- Geach, P. 1967: Intentional identity. *Journal of Philosophy*, 64: 627-632.
- Greenberg, Y. 2003: *Manifestations of Genericity*. London: Routledge.
- Gross, S. 2006: Can empirical theories of semantic competence really help limn the structure of reality? *Noûs*, 40: 43-81.
- Gross, S. 2015: Descriptive semantic externalism. In N. Riemer (ed.), *Routledge Handbook of Semantics* (pp. 13-29). London: Routledge.
- Heim, I. 1990: E-type pronouns and donkey anaphora. *Linguistics and Philosophy*, 13: 137-177.
- Heim, I. 1997: Predicates or formulas? Evidence from ellipsis. In A. Lawson (ed.), *Proceedings from SALT 7* (pp. 197-221). Ithaca: CLC.
- Heim, I. and Kratzer, A. 1998: *Semantics in Generative Grammar*. Oxford: Blackwell.
- Higginbotham, J. 1985: On semantics. *Linguistic Inquiry*, 16: 547-593
- Higginbotham, J. 1989: Elucidations of meaning. *Linguistics and Philosophy*, 12: 465-517.
- Higginbotham, J. 2004: On linguistics in philosophy, and philosophy in linguistics. *Linguistics and Philosophy* 25: 573–584.
- Higginbotham, J. 2009: Two interfaces. In M. Piatelli-Palmarini, J. Uriagereka, and P. Salaburu (eds.), *Of Minds and Language: a Dialogue with Noam Chomsky* (pp. 142-154). Oxford: Oxford University Press.
- Hornstein, N. 1984: *Logic as Grammar: An Approach to Meaning in Natural Language*.

Cambridge, MA: MIT Press.

Jacobson, P. 2014: *Compositional Semantics: An Introduction to the Syntax/Semantics Interface*. Oxford: Oxford University Press.

Katz, J. and Fodor, J. 1963: The structure of a semantic theory. *Language*, 39: 170-210.

Kennedy, C. and Stanley, J. 2009: On 'average'. *Mind*, 118: 583-646.

Koslicki, K. 1999: Genericity and logical form. *Mind and Language*, 14: 441-467.

Krifka, M. 1998: The origins of telicity. In S. Rothstein (ed.), *Events and Grammar* (pp. 197-236). Dordrecht: Kluwer.

Krifka, M, Pelletier, J. F., Carlson, G. N., ter Meulen, A., Chierchia, and Link, G. 1995: Genericity: an introduction. In G. Carlson and F. J. Pelletier (eds.), *The Generic Book* (pp. 1-124). Chicago: Chicago University Press.

Kripke, S. 2013: *Reference and Existence: The John Locke Lectures*. Oxford: Oxford University Press.

Lawler, J. 1973: *Studies in English Generics*. Ph.D. diss., University of Michigan, Ann Arbor.

Liebman, D. 2011: Simple generics. *Noûs*, 45: 409-442.

Lepore, E. and Loewer, B. 1981: Translational semantics. *Synthese*, 48: 121-133.

Leslie, S-J. 2015: Generics oversimplified. *Noûs*, 49: 28-54.

Lewis, D. 1970: General semantics. *Synthese*, 22: 18-67. References to the reprint in D. Lewis (1983), *Philosophical Papers. Volume I* (pp. 189-229). Oxford: Oxford University Press.

Ludlow, P. 1999: *Semantics, Tense, and Time: An Essay in the Metaphysics of Natural Language*. Cambridge, MA: MIT Press.

- Ludlow, P. 2003: Referential semantics for I-languages? In L. M. Antony and N. Hornstein (eds.), *Chomsky and His Critics* (pp. 140-161). Oxford: Blackwell.
- Melia, J. 1995: On what there's not. *Analysis*, 55: 223-239.
- Mari, A., Beyssade, C. and Del Prete, F. 2013: Introduction. In A. Mari, C. Beyssade, and F. D. Prete (eds.), *Genericity* (pp. 1-92). Oxford: Oxford University Press.
- McGilvray, J. 1998: Meanings are syntactically individuated and found in the head. *Mind and Language*, 13: 225-280.
- Neale, S. 1990: *Descriptions*. Cambridge, MA: MIT Press.
- Pietroski, P. 2003: The character of natural language semantics. In A. Barber (ed.), *Epistemology of Language* (pp. 217-256). Oxford: Oxford University Press.
- Pietroski, P. 2005: Meaning before truth. In G. Preyer and G. Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (pp. 255-302). Oxford: Oxford University Press.
- Pietroski, P. 2015: Framing the event variable. *Erkenntnis*, 80: 31-60.
- Quine, W. V. O. 1953: On what there is. In W. V. O. Quine, *From a Logical Point of View* (pp.). Cambridge, MA: Harvard University Press.
- Quine, W. V. O. 1960: *Word and Object*. Cambridge, MA: MIT Press.
- Salmon, N. 1998: Nonexistence. *Noûs*, 32: 277-319.
- Scott, G-J. 2002: Stacked adjectival modification and the structure of nominal phrases. In G. Cinque (ed.), *The Functional Structure in DP and IP: The cartography of Syntactic Structure, Volume 1* (pp. 91-121). Oxford: Oxford University Press.
- Stanley, J. 2001: Hermeneutic fictionalism. *Midwest Studies in Philosophy*, 25: 36-71.
- Stanley, J. 2007: *Language in Context: Selected Essays*. Oxford: Oxford University

Press.

Van Inwagen, P. 1977: Creatures of fiction. *American Philosophical Quarterly*, 14: 299-

308.

Yablo, S. 2001: Go figure: a path through fictionalism. *Midwest Studies in Philosophy*,

25: 72-93

Yablo, S. 1998: Does ontology rest on a mistake? *Proceedings of the Aristotelian Society*

(*Supplementary Volume*), 72: 229-262