

- Ross, John R. 1986. *Infinite syntax!* Norwood, NJ: Ablex.
- Sekerina, Irina A. 1997. The syntax and processing of scrambling constructions in Russian. Doctoral dissertation, City University of New York.
- Slioussar, Natalia. 2007. Grammar and information structure: A study with reference to Russian. Doctoral dissertation, Utrecht University.
- Titov, Elena. 2013. Do contrastive topics exist? *Journal of Linguistics* 49:413–454.
- Warren, Tessa, and Edward Gibson. 2002. The influence of referential processing on sentence complexity. *Cognition* 85:79–112.
- Yadoff, Michael. 1999. Formal properties of functional categories: The Minimalist syntax of Russian nominal and prepositional expressions. Doctoral dissertation, Indiana University, Bloomington.
- Zemskaja, E. A., ed. 1973. *Russkaja razgovornaja reč* (Russian spoken language). Academy of Sciences of USSR. Moscow: Nauka.
- Zubатов, Gerhild, and Uwe Junghanns. 1995. Syntax and information structure of Russian clauses. In *Formal Approaches to Slavic Linguistics 4: The Cornell Meeting*, ed. by Wayles Browne, Ewa Dornisch, Natasha Kondrashova, and Draga Zec, 289–319. Ann Arbor: Michigan Slavic Publications.

AGAINST LEXICAL SELF-  
REFERENCE

Eric McCready

Aoyama Gakuin University

This squib discusses a problem that arises when a standard degree-based semantics for intensifiers is combined with a second-order contextualist semantics for the predicate *average* on its concrete reading. In a nutshell, the combination requires that the argument of *totally average* be simultaneously average in every respect and not average at all in one particular respect. This problem is claimed to arise from allowing (in a sense) the denotation of *average* to refer to itself; the problem is then solved by prohibiting (by a combination of semantic and pragmatic means) self-reference at the lexical level.

### 1 Why the Totally Average Is Special

The semantics of the term *average* has received substantial attention in recent years in linguistics and philosophy (Carlson and Pelletier 2002, Kennedy and Stanley 2009), partly for reasons of intrinsic linguistic interest and partly because of philosophical arguments that have been made on the basis of sentences including this predicate. More

Thanks to Daisuke Bekki, Yasutada Sudo, two anonymous reviewers for *Linguistic Inquiry*, and the editors for extremely useful discussion and comments.

recently, intensifiers such as *totally* have gotten even more attention; here, the reasons are largely linguistic, and concern proposals for the proper semantics of intensification as well as so-called “superpropositional” uses of intensifiers to indicate epistemic confidence or the like (see Beltrama and Bochnak to appear or Kaufmann and McCreedy 2013 for proposals and data). While there may be no fully uncontroversial analysis on the market for either term, they appear to be fairly well-understood; however, as yet there has been no examination of the result of putting the two together. The aim of this squib is to point out a puzzling consequence of doing so, and to propose a way of eliminating it.

What does it mean to be totally average? The term *average* is standardly taken to do one of two things. On what Carlson and Pelletier (2002) call its *concrete* use, it indicates that the individual denoted by its argument satisfies some number of properties taken to hold of typical individuals; this use is exemplified in (1). On its *abstract* use, it is used to construct an “average individual” abstracted from concrete instances, as in (2). This second use appears not to arise in verbal predication; (3) only makes sense on the concrete interpretation. I will therefore put the abstract use aside in what follows (though see Carlson and Pelletier 2002, Kennedy and Stanley 2009 for two possible analyses of it).<sup>1</sup>

- (1) The average Texan eats a lot of meat.
- (2) The average Texan owns 1.3 guns.
- (3) John is average.

Now consider the variant of (3) in (4).

- (4) John is totally average.

Here, a verbal predication of *average* is combined with *totally*, a very common construction. Still, a strange consequence arises from it. Suppose that John is typical in every way. He satisfies all the properties typical of the group under discussion: he is average in height, weight, and intelligence; his taste in food exemplifies that of the normal citizen of his nation; and so on. In short, he satisfies all of the properties taken to characterize the average. (4) then appears true. However, observe that in fact John is decidedly *not* average: the typical person is not typical in every respect, but only in some. John’s degree of averageness

<sup>1</sup> A reviewer observes that (3) is hard to interpret in a null context, probably because it is difficult to guess in what sense John is supposed to be average without knowing what is under discussion (is he an average Texan? an average American? an average linguist?). The second-order analysis of *average* provided later in the squib gives one way to understand this intuition: without a context, it is hard to determine what set of properties is supposed to be used to compute John’s averageness. The whole issue relates closely to the problem of determining intended interpretations; see McCreedy 2012 for one possible view.

is well outside the average range. If we understand being totally average as satisfying all the properties of average people, John is thus not totally average after all.

An intuitive difficulty therefore arises in assigning (4) a truth value. The aim of this squib is to examine the problem in more detail and, on the basis of the results, propose a means of eliminating it on the basis of the observation that being *totally average* does not require being *averagely average*. Generalizing, it turns out to be sufficient to simply prohibit self-referential application in lexical denotations. This prohibition may be thought of as a pragmatically based restriction on the construction of lexical meanings, as will be discussed further in section 3.

## 2 Clarifying the Issue

One might react to the above discussion by concluding that the proposed analysis of *totally average* is just wrong. This is a possible conclusion; however, the analysis follows quite directly from standard treatments of intensifiers and the meaning of *average*. The aim of this section is to show how, which in turn requires a somewhat more formal treatment.

Property-level intensification is something that arises only with scalar predicates. (5a) is fine; (5b) is not (when the intensification is taken to be of the property denoted by the verb phrase). The reason is that *open* is associated with a scale—the scale composed of degrees of openness—and (*is a*) *teacher* is not, at least in the absence of lexical coercion (which is often available, though not in this case; see, e.g., Morzycki 2009).


- (5) a. The door is totally open.  
b. \*John is totally a teacher.

The use of intensifiers thus presupposes that the predicate a given intensifier combines with is scalar (at least on a property-modifying reading; see, e.g., Kaufmann and McCready 2013 for some discussion). One standard analysis of scalar predicates goes as follows: each predicate is associated with a degree-based scale and is taken to map the argument to the predicate to a degree on that scale; the predicate is then to be judged true of the individual if the degree the individual is associated with exceeds a contextually specified point on the predicate's scale (Kennedy 2007). In (6a), for instance, the predicate *tall* maps the individual denoted by *John* to that individual's degree of height, and, given that that point exceeds the contextually specified standard for tallness, the sentence is judged true.

- (6) a. John is tall.  
b.  $\llbracket(6a)\rrbracket = 1$  iff  $\mathbf{tall}(j) >_{\mathbf{tall}} s$ ,  
for  $s$  the contextual standard for tallness

On a version of Kennedy's implementation of the degree-based view, this semantics is derived in somewhat more formal detail as in (7).

Gradable adjectives denote functions from individuals to degrees and come paired with a covert element *pos* (indicating the “positive” or noncomparative form of the adjective), which provides a standard of comparison and relates it to the degree yielded by application of the adjective denotation to the individual that is the object of predication.



Intensifiers of the *totally* class—which we might call *maximal intensifiers*—can be understood in this way of looking at adjectival predication as setting the contextual standard to the maximum point of the scale associated with the predicate they modify. Kennedy and McNally (2005) give the denotation in (8) for *completely*, which can be taken to exemplify the property-modifying use of this class of adverbial.<sup>2</sup>

$$(8) \llbracket \text{completely} \rrbracket = \lambda g \lambda x \exists d [d = \mathbf{max}(<_g) \wedge g(x) = d]$$

Thus, scalar predicates appearing with maximal intensifiers can be taken to be true of individual arguments if the degree of the property the individual has is the highest made available by the scale. If there is no such maximal point (i.e., the scale lacks an upper bound), the intensification will be ungrammatical, as in (9a), for there will be no scale-maximal point available (Kennedy and McNally 2005); conversely, (9b) is true just in case the door has the maximal possible degree of openness (made available by its physical structure).

- (9) a. \*John is totally tall.  
 b. The door is totally open.  
 c.  $\llbracket (9b) \rrbracket = 1$  iff  $\mathbf{open}(d) = \mathbf{max}(<_{\text{open}})$

From this discussion, it is clear that (the concrete use of) *average* must be scalar, for if it was not, intensification would be ungrammatical; further, it must have a maximal point, for if it did not, intensification with maximal intensifiers would lead to ungrammaticality as well. Thus, the scale associated with *average* has an upper bound.

How can a scale be derived for *averageness*? For the concrete use, Carlson and Pelletier (2002) take *average* to present a set of properties taken to be typical of some group under discussion (simplifying slightly). Given a set *T* of typical properties, the averageness or typicality of an individual can be thought of as the proportion of the properties in *T* that the individual satisfies; indeed, it is possible to

<sup>2</sup> I have modified Kennedy and McNally’s (2005) original definition—their (62)—slightly to fit it to the degree semantics I am using.

derive a closed scale directly from  $T$ , in that an individual satisfying none of the properties in  $T$  will have degree 0, an individual satisfying all of them will have degree 1, and individuals satisfying some intermediate number will be assigned degrees of typicality corresponding to the proportion of  $T$ -properties they satisfy. A predication of *average* will then be taken to be true if the degree of averageness or typicality the argument has exceeds the contextual standard for counting as “average,” just as above.<sup>3</sup>

Putting these pieces together, (4) can be taken to be true if the number of typical properties satisfied by John is the highest possible—that is, if John is average in every respect. This will be so if every property in  $T$  is true of John. This means that truth judgments about whether an individual counts as completely average depend on what properties are included in  $T$ ; this point is crucial to the later discussion.

$$(10) \llbracket(4)\rrbracket = 1 \text{ iff } \mathbf{average}(j) = \max(\langle \text{average} \rangle \text{ iff } \forall P[P \in T \rightarrow P(j)])$$

The problem we are concerned with here becomes clear when we consider how to derive averageness for gradable predicates in the first place. Consider for instance (11).

$$(11) \text{ John's height is average.}$$

Presumably, (11) must be true in order for John to count as completely average. But when is it true? To determine this, we must compare the degree of tallness that John has to the average degree of height of the relevant population. Suppose that the average height of the population is obtained by averaging over the heights of all contextually salient individuals, understood as those individuals in a set of relevant people  $C$ , as in (12a). Given some function  $h$  yielding a set of values surrounding a particular point (essentially the pragmatic halos of Lasersohn 1999, which are themselves presumably contextually dependent, as discussed in for example McCready 2008) and understood as determining a range of values describable as average, as in (12b), it is possible to derive a range of average-enough values for height by

<sup>3</sup> This semantics for concrete uses of *average* is not fully satisfactory, in at least two respects. First, it assumes that all properties in  $T$  are equally important, which is likely not the case. Certain properties will be more characteristic of certain kinds of average individuals than others, an issue easily solved: different weights can be assigned to the various properties, and if one property is taken to be more crucial for typicality than others in a particular context, it can be weighted more heavily. Second, not all instances of nontypicality are created equal. If the average American male is (say) between 170 and 180 cm tall, two individuals of heights 181 cm and 211 cm are both nonaverage, but the second is unlikely to be considered sufficiently average to count as genuinely average in many contexts in which the first still will. Thus, severity of deviation from the average should count into the global computation of averageness. Since these issues are orthogonal to my points here, I will put these complications aside for the purposes of this squib.

applying  $h$  to the average height of the population, as in (12c), which in turn can be used to decide whether an individual's height counts as average (enough). Thus, (11) will be true if John's height lies in the range of values determined by  $h$  when applied to  $average_{tall}(C)$ , as in (12d).

- (12) a.  $average_{tall} = \frac{\sum_{x \in C}(\mathbf{tall}(x))}{|C|}$   
 b.  $h(n) = \{i | dist(n, i) < k\}$  for a contextually specified distance  $k$   
 c.  $h(average_{tall}) = \{i | dist(average_{tall}, i) < k\}$   
 d.  $\llbracket(11)\rrbracket = 1$  iff  $\mathbf{tall}(j) \in h(average_{tall})$

Thus, averageness is determined in three steps. First, a range of average values for gradable predicates is derived, yielding predicates of "average  $P$ -ness." These predicates are put together with other properties of average individuals in the comparison set to form a set of "average properties"  $T$ . Finally, to determine whether a given individual counts as average, the number of predicates in  $T$  that that individual satisfies is counted and compared to a contextually determined standard. This seems more or less adequate.

Now we come to the crux of the matter. Applying the method described above to the property of averageness yields a range of average values of averageness (13b).

- (13) a.  $average_{average} = \frac{\sum_{x \in C}(\mathbf{average}(x))}{|C|}$   
 b.  $h(average_{average}) = \{i | dist(average_{average}, i) < j\}$

By assumption, John's averageness does not lie in this range, for his averageness is total (i.e., maximal, so  $\mathbf{average}(j) = \max(<_{average})$ ). So (4) will be true if and only if John's degree of averageness is identical to the maximal point of the scale of averageness, meaning that it should come out true given the assumption above. However, John's degree of averageness lies outside  $h(average_{average})$  in nearly any situation: it is very rare in any normal population for maximal degrees of some property to be sufficiently average to count as average. Then, since John is not averagely average, his averageness is not in fact maximal, as there is now a property in  $T$  that he does not satisfy, and so  $\mathbf{average}(j) < \max(<_{average})$ ; thus, John is not totally average after all, and (4) should come out false. This is paradoxical, and clearly an unwelcome result. Note that this issue does not arise for simple predications of averageness; there, since not all properties in  $T$  need be satisfied, the property of being average can always simply be ignored for purposes of determining whether the predicated individual satisfies a sufficient number of  $T$ -properties.

The problem thus can be stated most directly as follows: *average* determines a scale of averageness. This scale comes with a contextual standard for averageness; if an individual's degree of averageness exceeds the standard, she is average. To count as averagely average, an individual's degree of averageness must fall within normal parameters

for the population, which normally will exclude the upper and lower bounds of the averageness scale. However, to count as totally average, an individual's degree of averageness must be maximal. Thus, being averagely average is required for being totally average, but is also incompatible with it.

### 3 Lexical Self-Reference Disallowed

Section 2 showed in formal detail what goes wrong when the (seemingly individually reasonable) analyses of *totally* and *average* discussed above are put together. In the remainder of this squib, I will diagnose the problem and propose a solution.

The difficulties arise from the combination of *totally/completely*, which require a scale-maximal point, and *average*, which provides a set of properties and disallows its argument from mapping to a scale-maximal point in any of them. It is thus possible to address the problem on the basis of either of these elements. There are two obvious possibilities for doing so: (a) either we can relax the maximality requirement, or (b) we can leave the property of being average out of the computation of averageness. There is also a third option: (c) rejecting the analyses of maximizing adverbials or scalar adjectives used in the discussion. I will consider these three options, concluding that (b) is the best choice.

Let me begin with option (c), that of rejecting the analyses I have proposed. The question is whether doing so can yield any analytical possibility that removes the problem yet is empirically adequate and ultimately conceptually distinct from the current proposal. The analysis of *average* I have proposed strikes me as more or less conceptually minimal. According to it, an individual counts as average just in case she exhibits a sufficient number of the properties of average individuals. Ultimately, this by itself is not completely adequate, as discussed in footnote 3, but it seems as if any analysis of *average* must take it or something very similar as a baseline (e.g., Carlson and Pelletier 2002); the notion of a set of typical or prototypical properties seems to be a necessity. Alternatively, one might object to my use of a degree semantics, together with a contextual standard, pointing perhaps to alternative degree-free theories of gradable predication such as that proposed by Beltrama and Bochnak (to appear). Beltrama and Bochnak make use of contexts that indicate the current "strictness" of judgments about the truth of predicates; they analyze predications of gradable predicates  $P(x)$  as saying that the strictness of the current context supports the judgment that  $P$  holds of  $x$ . Intensifiers like *totally* then indicate that  $P$  can be judged true of  $x$  in all possible contexts, including the strictest. This kind of view may be a reasonable option, but does it yield different predictions in the present case? After all, for someone to be judged average in the strictest possible context, she ought to satisfy all properties that typically hold of average individuals, which corresponds quite directly to a maximal point on a scale of averageness. In sum, it is not obvious that switching frameworks will allow us to avoid the problems I have discussed.

Let us now consider a different use of contexts corresponding to option (a). The maximality requirement can be relaxed by allowing arguments of *totally Adj* to satisfy *Adj* to a nonmaximal degree that is nonetheless close to maximal. It is already known that certain absolute predicates involving upper-closed scales may not always require genuine maximality for truth, despite requiring it in theory. Consider the predicate *closed*: to be closed, a door must be not at all open, meaning that a zero degree of openness (i.e., a maximal degree of closedness) is required for the truth of a predication of *closed*. This is the standard picture (cf. Kennedy and McNally 2005). This analysis predicts that adding *totally* should not change the meaning at all: the door is already required to be fully closed, so *totally* doesn't do any work, as discussed by Kaufmann and McCreedy (2013). To solve this problem, one can assume that *closed* in fact admits a bit of fuzziness at the top of the scale, and *totally* works to eliminate it, something in the way in which predicates like *definitely* can be used in the analysis of vagueness (e.g., Fine 1975). Assuming that *totally* itself also allows for a bit of indeterminacy might be enough to eliminate the self-contradictory quality we have found in *totally average*: we can simply let the degree of averageness required for the argument be a little less than the maximum. This will let us evade the problem; we can always just disregard the degree of averageness possessed by the individual that is the object of predication, eliminating the inconsistency discussed in section 2.

There are two reasons to be unsatisfied with this analysis. The first is that it weakens the semantics of *totally* to an unacceptable degree. In the previous paragraph, we saw that *totally closed* requires a higher degree of closedness than *closed* alone, indeed a genuinely maximal one; but building some nonmaximality into the denotation means that a genuinely maximal degree will not be obtained. Causing problems for the semantics of intensifiers across the board to handle the special case under discussion is not a sensible analytic strategy. Worse, this solution fails to address the actual source of the problem: the issue is not that the degree of averageness required by *totally average* is too high, but that being totally average with respect to averageness yields self-contradictory results. It is necessary to leave the property of averageness out of the computation of averageness, not to simply weaken the semantics in order to allow it to leave out an unspecified property.

To see this point, consider a semantics for intensifiers that would simply introduce a ‘halo’ around the maximal point and let anything with a degree of *P* within the halo count as *completely/totally P*, as in (14).

$$(14) \llbracket \textit{completely} \rrbracket = \lambda g \lambda x \exists d [d \in h(\mathbf{max}(<_g)) \wedge g(x) = d]$$

This semantics, when combined with  $\llbracket \textit{average} \rrbracket$ , is compatible with a situation in which the average degree of averageness of the population is extremely high and John is average in all but a few respects (which nonetheless are highly striking, given the general homogeneity of the population). The reason is that his degree of averageness still falls



within the halo of  $\mathbf{max}(\langle \textit{average} \rangle)$ , since he is average in all but a few respects. Thus, (4) comes out true on this semantics. However, intuitively, John is *not* completely average: he is quite nonaverage with respect to the whole population. This is a first reason to think the proposed fix is on the wrong track.

Another reason comes from considering the following question: what properties should be “left out” of the computation? There are many ways to be nonmaximal with respect to satisfying second-order properties such as averageness. The key question here is whether averageness itself should be considered in computing averageness. There are two possibilities: to consider averageness, or not to. Suppose that averageness is considered. Then predicating *completely average* of  $x$  requires that  $x$  have a degree of averageness within the halo of normal averageness, and so be quite normal. In normal populations with a wide distributional range of properties, this will either mean (a) that  $\mathbf{average}(x) \notin h(\mathbf{max}(\langle \textit{average} \rangle))$  at all, or (b) that the average degree of averageness in the population is sufficiently close to  $\mathbf{max}(\langle \textit{average} \rangle)$  that  $\mathbf{average}(x) \in h(\mathbf{max}(\langle \textit{average} \rangle))$  after all. But, intuitively, judging (4) does not in general introduce domain conditions of this sort. It seems the problem here is the inclusion of averageness in the computation; leaving it out removes the issue. I conclude that the best option is to leave averageness out entirely, rather than to directly alter the semantics of intensifiers, which, as we have seen, runs the risk of improperly weakening the truth conditions of sentences like (5a) (and indeed (4) itself).

This corresponds to option (b) that I mentioned above; simply to leave out the problematic property. The analysis of *average* in section 2 involves a set of properties  $T$  taken to hold of the average individual of the sort under consideration; predication of *average* in turn involves checking whether the argument satisfies enough of those properties to count as average. Putting this together with *totally* then requires “local” nonmaximality, in that all the properties associated with *average* must hold of its argument, which means that the argument must not have too extreme a degree of any of these properties when they are gradable, and “global” maximality, in that every property in  $T$  must hold of the argument. As we saw above, these two requirements are not compatible when  $T$  contains the property of being average; if it does, the argument must be both nonmaximal and maximal with respect to averageness. Simply leaving out the property of averageness removes the incompatibility.

I suggest that the solution to the problem focused on here can be stated as the principle in (15).

(15) *No Self-Referential Lexical Predication (NSLP)*

No predication of a lexical item results in truth conditions that refer to the result of that predication.

According to NSLP, it is not just *average* that cannot check whether it is correctly predicated of its argument in order to check whether it is correctly predicated of its argument; no lexical item can do so. This

is sensible. The whole attempt has a paradoxical flavor. In general, it cannot be expected that such predications would even yield a determinate truth value, since the possibility is opened for predicates that “loop,” which makes them problematic from a pragmatic perspective.<sup>4</sup> I thus propose that it is this principle that allows us to avoid the problem discussed in sections 1 and 2. The particular case of *average* falls out as one case ruled out by NSLP.

Is there any more general justification for (15)? In fact, this constraint is closely related to observations that have been made for other paradoxical phenomena outside the lexical domain, which in turn justify the “lexical” part of NSLP itself.

*Totally average* is an example of a semantic paradox of the type called *impredicative*, which is a class of paradoxes resulting from “definitions . . . which [depend] on a set of entities, at least one of which is the entity being defined” (Bolander 2014).<sup>5</sup> Moss (2014) calls this kind of thing *object circularity*. Compare Whitehead and Russell motivating their introduction of the “vicious circle principle”:

An analysis of the paradoxes to be avoided shows that they all result from a kind of vicious circle. The vicious circles in question arise from supposing that a collection of objects may contain members which can only be defined by means of the collection as a whole. . . . The principle which enables us to avoid illegitimate totalities may be stated as follows: “Whatever involves all of a collection must not be one of the collection”; or, conversely: “If, provided a certain collection had a total, it would have members only definable in terms of that total, then the said collection has no total.” (1910:2nd ed. 37)

Clearly, this kind of move has close connections with NSLP. However, observe that many paradoxes constructible in natural language do involve self-referential predication, such as the famous Liar Paradox: (in one formulation) a statement of *This sentence is false*. Plainly, this sentence has no determinate truth value, but, just as plainly, it is both grammatical and interpretable. Because of this, any ban on self-referential predication to be enforced in natural language must be restricted; it cannot be taken as a general pragmatic principle, for if it were completely general, the Liar Paradox could not be generated. My proposal is therefore to restrict it to the case of predication of individual lexical items.

However, this does not mean that pragmatics should be ignored.<sup>6</sup> NSLP is satisfied for the case of *average* by leaving the property of averageness out of the set of properties that are used in judging whether

<sup>4</sup> Indeed, if one takes into account considerations from the evolution of language, it seems likely that predicates with this character would “die out” over some period of time, given the possibility of looping and the concomitant loss of utility. Some relevant discussion can be found in Jäger 2014.

<sup>5</sup> I thank the anonymous reviewers for suggesting an explicit comparison to other cases of paradox.

<sup>6</sup> Thanks to the editors and reviewers for pressing this point.

an individual counts as average or not. The question is whether this leaving-out is the result of a semantic or pragmatic process. If it is pragmatic, then language users leave averageness out when setting up a set of properties over which to compute averageness; if semantic, then the property is simply not available in the first place, in the sense that it is not a possible part of the denotation. Adopting the semantic story requires explaining the senses of *possible* and *available* at issue in the previous sentence, which seems impossible without some appeal to pragmatics, since it is in part context-dependent what properties one must satisfy in order to count as average. I conclude that NSLP, while a lexical principle, is one crucially founded on pragmatic processes.

As far as I can see, NSLP lacks negative empirical consequences. It does not appear to rule out anything we would like to allow. One might worry about examples like (16), discussed extensively in property-theoretic semantics (Chierchia and Turner 1988), but here *fun*, while it is in a sense predicating itself of itself, is nonetheless not doing so in a way that violates NSLP, which explicitly refers to cases of self-reference. In (16), the predication of *fun* does not refer to its own result.

(16) Fun is fun.

The function of NSLP is to restrict the construction of the denotations of lexical items that are computed on the basis of sets of properties. Several such predicates (and constructions) have been discussed in the literature. One case involves terms of the form *N-ly* like *manly* and their crosslinguistic analogues such as the Japanese *N-rashii*. Here (according to the analysis proposed in McCready and Ogata 2007), a set of properties taken to stereotypically hold of *Ns* is constructed, and truth conditions are computed much as with the analysis of *average* in the previous section; *x* counts as *N-ly* (*N-rashii*) if and only if it has a sufficient number of the relevant properties. Terms like *N-like* (or the Japanese *N-mitai*) are similar, barring the particular properties selected. Something similar might be thought to arise in cases of focus reduplication (Ghomeshi et al. 2004) like *salad salad*, where prototypical members of a class are picked out; to the extent that prototypicality involves satisfying some set of properties, the denotations here can be viewed as rather similar.

NSLP is also needed for these other predicates in the maximal intensification case. The reason is that in order to check whether all predicates in the adjective denotation are satisfied, it would be necessary to check for satisfaction of the predicate itself were it to be included, so an infinite regress would arise.<sup>7</sup> For example, in the case of *manly*, one would have to check whether all properties of manly

<sup>7</sup> Thanks to Yasutada Sudo for this point.

individuals were satisfied, including that of being manly; but, given that manliness is included in the relevant set of properties, this cannot be checked without first checking whether *manly* itself applies, leading to a loop.

The difference with *average* is that, for these cases, satisfying the “main” predicate in the maximal manner required by *totally* requires satisfying all the “subpredicates”: thus, for example, being *totally manly* requires having all the properties manly things have. In contrast, *totally average* simultaneously requires having all the properties of average things *and* not having too many, for *average* looks to the midpoint of a scale rather than to the extremes. This seems to be the reason that special problems arise for the combination of a scalar theory of *average* involving a second-order construction and a maximal-degree theory of *totally*: in the case of *average*, and only that case, we get a kind of Liar-style paradox in which maximal property satisfaction within the set of properties leads to a lack of satisfaction of the main predicate, given modification with a maximal intensifier. In this sense, *average* itself is special.

Considering other cases of second-order predication of this kind also rules out the strongest kind of pragmatic analysis, on which no lexical restrictions are placed on second-order predicates at all, and instead the constraint arises via pragmatics. Consider (17), which is true if John is taller than everybody but himself: his own height is not considered, for doing so would lead to a contradiction. The reason seems to involve the way in which comparison sets are selected: pragmatics dictates leaving out those predicates that would lead to a contradictory interpretation.<sup>8</sup>

(17) John is taller than everyone.

This kind of analysis would work for *average* as well, for *totally average* also gives an interpretation that cannot be true given that *average* is included in the set of properties that are required to have values in the average range. But it will not work for cases like *manly*: maximal manliness is quite consistent with being maximally manly. Still, given that looping arises, it will be impossible to determine whether the predicate *totally manly* is satisfied, just as with other cases of self-reference. A stronger pragmatic principle ruling out loops completely also cannot be maintained, given the existence of the Liar Paradox, which depends on the possibility of nonlexical looping in interpretation. NSLP therefore appears to be an independently needed semantic principle, though a pragmatically grounded one.

<sup>8</sup> This case, and the corresponding analysis, were suggested by an anonymous reviewer.

## References

- Beltrama, Andrea, and Ryan Bochnak. To appear. Intensification without degrees cross-linguistically. *Natural Language and Linguistic Theory*.
- Bolander, Thomas. 2014. Self-reference. In *Stanford encyclopedia of philosophy*, ed. by Edward N. Zalta. Spring 2014 ed. Available at <http://plato.stanford.edu>.
- Carlson, Greg, and Francis Jeffrey Pelletier. 2002. The average American has 2.3 children. *Journal of Semantics* 19:73–104.
- Chierchia, Gennaro, and Raymond Turner. 1988. Semantics and property theory. *Linguistics and Philosophy* 11:261–302.
- Fine, Kit. 1975. Vagueness, truth, and logic. *Synthese* 30:265–300.
- Ghomeshi, Jila, Ray Jackendoff, Nicole Rosen, and Kevin Russell. 2004. Contrastive focus reduplication in English (the salad-salad paper). *Natural Language and Linguistic Theory* 22: 307–357.
- Jäger, Gerhard. 2014. What is a universal? On the explanatory potential of evolutionary game theory in linguistics. In *Formal approaches to semantics and pragmatics: Japanese and beyond*, ed. by Eric McCready, Katsuhiko Yabushita, and Kei Yoshimoto, 85–103. Dordrecht: Springer.
- Kaufmann, Magdalena, and Eric McCready. 2013. Maximum intensity. Paper presented at Keio University.
- Kennedy, Christopher. 2007. Vagueness and grammar: The semantics of relative and absolute gradable adjectives. *Linguistics and Philosophy* 30:1–45.
- Kennedy, Christopher, and Louise McNally. 2005. Scale structure and the semantic typology of gradable predicates. *Language* 81: 1–37.
- Kennedy, Christopher, and Jason Stanley. 2009. On ‘average’. *Mind* 118:583–646.
- Lasersohn, Peter. 1999. Pragmatic halos. *Language* 75:522–551.
- McCready, Eric. 2008. What man does. *Linguistics and Philosophy* 31:671–724.
- McCready, Eric. 2012. Emotive equilibria. *Linguistics and Philosophy* 35:243–283.
- McCready, Eric, and Norry Ogata. 2007. Adjectives, comparison and stereotypicality. *Natural Language Semantics* 15:35–63.
- Morzycki, Marcin. 2009. Degree modification of gradable nouns: Size adjectives and adnominal degree morphemes. *Natural Language Semantics* 17:175–203.
- Moss, Lawrence S. 2014. Non-wellfounded set theory. In *Stanford encyclopedia of philosophy*, ed. by Edward N. Zalta. Spring 2014 ed. Available at <http://plato.stanford.edu>.
- Whitehead, Alfred North, and Bertrand Russell. 1910. *Principia mathematica*, vol. 1. (2nd ed., 1925.) Cambridge: Cambridge University Press.