

Pragmatic Halos

Peter Lasersohn

Language, Vol. 75, No. 3. (Sep., 1999), pp. 522-551.

Stable URL:

<http://links.jstor.org/sici?sici=0097-8507%28199909%2975%3A3%3C522%3APH%3E2.0.CO%3B2-D>

Language is currently published by Linguistic Society of America.



Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/lsa.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is an independent not-for-profit organization dedicated to creating and preserving a digital archive of scholarly journals. For more information regarding JSTOR, please contact support@jstor.org.

PRAGMATIC HALOS

PETER LASERSOHN

University of Illinois

It is a truism that people speak ‘loosely’—that is, that they often say things that we can recognize not to be true, but which come close enough to the truth for practical purposes. Certain expressions, such as those including *exactly*, *all* and *perfectly*, appear to serve as signals of the intended degree of approximation to the truth. This article presents a novel formalism for representing the notion of approximation to the truth, and analyzes the meanings of these expressions in terms of this formalism. Pragmatic looseness of this kind should be distinguished from authentic truth-conditional vagueness.*

1. **PRAGMATIC SLACK.** People speak with varying degrees of precision, and often speak quite loosely. Suppose, for example, I tell John that Mary arrived at three o’clock. In certain relatively unusual circumstances, the exact second of her arrival might be important, but most of the time this level of precision is not required. So if John finds out later that Mary didn’t arrive at three but at fifteen seconds after three, it would be unreasonable of him to complain ‘You said she came at three!’

But whether or not John is acting unreasonably in this situation, I think we have to concede that he is, strictly speaking, **RIGHT**: when I told him that Mary arrived at three, I said something that was literally false, not true. As people often do, I was just speaking a little loosely. My defense is not that I was telling the truth, but that what I said was ‘close enough’ to the truth for practical purposes.

Taking this position sets up an interesting semantic puzzle. If sentence 1 is true—literally true, not just close enough to the truth for practical purposes—only if Mary arrived precisely at 3:00, then it seems to be truth-conditionally equivalent to 2.

- (1) Mary arrived at three o’clock.
- (2) Mary arrived at exactly three o’clock.

But of course these sentences don’t mean the same thing. How can we express the difference between them?

The difference seems to be just that 2 does not permit one to speak as loosely as 1. Put differently, 1 allows greater ‘slack’ than 2 in determining just how close to the truth is close enough for practical purposes.

Now, one might have supposed that the degree of slack or looseness with which one may speak would be purely a matter of the pragmatic situation of utterance. But since 1 and 2 seem to differ precisely in how much slack they allow, and since this difference is tied directly to the presence or absence of the word *exactly*, it appears that we must build some means of dealing with slack into the lexical entries for particular words—and perhaps even into the compositional semantics.

Here is a related example: sentence 3 seems to involve some kind of quantification over the townspeople.

- (3) The townspeople are asleep.

* Versions of this article were presented at Tel Aviv University, Bar Ilan University, Hebrew University of Jerusalem, the University of Illinois, and Indiana University. Thanks to the audiences at those talks, and especially to Fred Landman, Susan Rothstein, Nirit Kadmon, Chris Barker, Christine Brisson, and the editors and referees for comments. Thanks to the University of Rochester for supporting a leave of absence, during which the initial research for this paper was carried out, and to Tel Aviv University for hosting me during this leave.

What kind of quantification is this? One might first suppose that it is universal quantification, so that 3 is analyzed as meaning something like 4.¹

- (4) $\forall x[\text{townsperson}(x) \rightarrow \text{asleep}(x)]$

But then one realizes that a speaker might very naturally use 3 even if some small percentage of the townspeople are, exceptionally, awake. Perhaps the quantification isn't universal after all, but only near-universal, in which case 3 should be accurately paraphrasable as something like 5.

- (5) More-or-less all of the townspeople are asleep.

But this leads to a problem, pointed out by Kroch (1974:190–91). Sentence 6 sounds contradictory, while 7 is clearly not contradictory.

- (6) Although the townspeople are asleep, some of them are awake.

- (7) Although more-or-less all the townspeople are asleep, some of them are awake.

I think Kroch is right in claiming that examples like 6 are contradictory—though this is not quite the right explanation for why they sound strange, as we shall see. But if 6 is contradictory, this suggests the quantification in 3 may have been universal rather than near-universal after all. But then how can we analyze 3 as allowing any awake townspeople at all, and how can we account for the difference in meaning between 3 and 8?

- (8) All the townspeople are asleep.

It seems to me that 3 allows only as many townspeople to be awake as can be, for practical purposes, ignored. That is, it allows exceptions not because its truth conditions explicitly allow for exceptions—the contradictoriness of 6 shows us that they don't—but because in most situations, one may speak a little bit loosely. One or two odd townspeople burning the midnight oil probably won't matter that much, so even if they are enough to make 3 literally false, the sentence is still close enough to the truth to serve its purpose.

This makes 3 truth-conditionally equivalent to 8. The difference in meaning, I suggest, is not truth-conditional, but simply in how much pragmatic slack they allow, in how much deviation from the truth they permit the pragmatic situation to license.

In further support of this idea, note that in pragmatic situations that make every exception relevant, sentences parallel to 3 do not allow any exceptions, even though the word *all* is not present. Suppose we are conducting an experiment on the nature of sleep. We have several people serving as experimental subjects there in our lab, lying on beds, dozing off one by one. In order for the experiment to proceed, we need to make sure that all of them are completely asleep; otherwise the experiment is ruined. In this sort of situation, if you assert 9, every last one of the subjects had better be asleep; exceptions are not tolerated.

- (9) The subjects are asleep.

The sentence itself may allow some slack, but this does not mean that every pragmatic situation will exploit the slack that the sentence makes available. If I am worried about the experiment being ruined, I may well hold the speaker to a stricter standard, and prohibit loose talk of the sort that comes close to the truth but doesn't quite hit it.

SCALAR ADJECTIVES are a third example where specific lexical items seem to play a

¹ This does not mean that we must take the definite determiner *the* as expressing universal quantification. The quantificational effect could come from a hidden verb phrase operator of the sort developed in Link 1987a, Roberts 1987, or from the lexical semantics of the predicate, as in Scha 1981, Lasnik 1995, and §4.2 below.

role in controlling how much deviation from the truth is pragmatically permissible. Scalar adjectives admit modification by degree adverbials such as *very*.² But some adjectives—even some close synonyms to scalar adjectives—do not permit such modification.

- (10) This ball is very round.
- (11) ?This ball is very spherical.
- (12) The surface was etched with very straight grooves.
- (13) ?The surface was etched with very linear grooves.

A reasonable way to account for this is to analyze scalar adjectives semantically as matching objects with levels on a scale, while nonscalar adjectives simply denote sets of objects, just like other predicates. Then, we formulate the semantics of *very* so that it makes direct reference to the scale introduced by the adjective it modifies, for example by raising the 'cutoff' point between those objects which qualify as round or those that don't, or by lowering the values onto which each object is mapped.

Since the semantics of *very* makes direct reference to scales, it should make no sense to combine *very* with a nonscalar adjective. Now we say that *spherical* is nonscalar, and simply denotes the set of spherical objects. Since *spherical* makes no reference to a scale, it should not combine with *very*, and we predict the contrast in 10 and 11. The same analysis can be applied to *straight* and *linear*, to give us the contrast in 12 and 13.

But now we have a problem: Which objects are spherical? Surely it is not unusual to talk about an object as being spherical even if its shape deviates slightly from that of a true sphere; indeed, outside of mathematical examples, it is unlikely that ANYTHING fails to deviate from this shape at least slightly; the same point holds for *linear*. We even have words whose function is to indicate that little or no such deviation is present, and these words can combine with words like *spherical* or *linear* perfectly naturally, no less than with *round* or *straight*.

- (14) This ball is perfectly round.
- (15) This ball is perfectly spherical.
- (16) The surface was etched with perfectly straight grooves.
- (17) The surface was etched with perfectly linear grooves.

Here is a puzzle: if *spherical* and *linear* are not scalar, how can we account for the intensifying effect of *perfectly* in examples like 15 and 17? With no scale to work on, *perfectly*'s semantics can't be analogous to *very*; indeed, it is hard to see how an operation on ordinary sets, with no reference to scales, could have an intensifying effect at all. We could try to fix things by claiming that *spherical* and *linear* are scalar after all, but if this is true, then why don't they combine with *very*, like *round* and *straight*?

I suggest that when we describe an object as spherical even though its shape deviates slightly from that of a true sphere, we are employing pragmatic slack—saying something which is literally false, but close enough to the truth for practical purposes. The purpose of *perfectly*, in examples like 15 and 17, is to take up some of this slack, to reduce the acceptable level of deviance from the truth allowed by the pragmatics. In this sense, *perfectly* is similar to *exactly* or *all* in our previous examples.

² I intend this simply as a statement about which adjectives combine with *very*, and not as a definition of scalarity. A referee suggests that we might want to regard *warm*, for example, as 'more of a scalar' than *round*; at least in the referee's dialect, *round* does not combine with *extremely*, *extraordinarily*, or *excessively*, although *warm* does (a judgment I do not share). See Bolinger 1972 for general discussion of degree words and scalarity.

To sum up: When extreme precision is not required, people accept utterances that deviate in minor ways from the truth. But the degree of deviation allowed is not determined solely by the pragmatics of the situation of utterance, but in part by the appearance of particular words within the utterance itself. We might call these words *SLACK REGULATORS*. A semantic theory ought to give some account of how slack regulators work. This article will make some proposals.

2. CLOSENESS TO THE TRUTH. I have spoken repeatedly of certain utterances ‘coming close’ to the truth without actually being true, and of slack regulators helping determine just how close to the truth is ‘close enough’. To give an analysis of slack regulators, therefore, we will need to make some sense of the notion of closeness to the truth. But, unfortunately, there seems little hope of coming up with a general metric by which we can assign any sentence its degree of closeness to the truth. A sentence may deviate from the truth in a myriad of ways, not all measurable in the same terms, and some perhaps not measurable at all.

If we attack the problem piecemeal, rather than aiming for full generality right from the start, it is not so hard to make sense of the notion. Even if we can’t give a general procedure for determining how close to the truth a given sentence is, we can certainly say how close to three o’clock Mary arrived, and this will give us some measure of the accuracy of *Mary arrived at three o’clock* (ex. 1). Likewise we can obtain some measure of accuracy for *The townspeople are asleep* (ex. 3) by determining what percentage of the townspeople are asleep, and so on. So, rather than appealing to a single, general notion of closeness to the truth, let us consider several ‘dimensions of closeness’: How close does the sentence come to getting the time right? How close does it come to picking out the right set of individuals? How close does it come to getting the shape of the ball right? And so on. There are undoubtedly many more dimensions of closeness besides these, and it is reasonable to suspect that different slack regulators make reference to different dimensions of closeness.

Of course we are concerned not just with measuring closeness to the truth, but also with determining just how close is close enough. This immediately raises the question, Close enough for what? Here again, a simple, very general answer is not likely to be forthcoming, but we may begin to get a handle on the problem by noting that it is not always degree of closeness per se that is important. What really matters is whether or not pragmatically relevant details and distinctions are represented. Suppose for example, we are waiting for the townspeople to fall asleep so we can attack the town. We may be perfectly satisfied to count the townspeople as asleep even if a few of them are awake—but only if the ones awake are not standing guard. In this situation, the distinction between the set of townspeople as a whole and the set of all the townspeople but a few insomniacs puttering around their houses may be irrelevant; the distinction between the set of townspeople as a whole and the set of all townspeople but a few vigilant lookouts is not. The absolute number of awake townspeople matters much less than the effect they may have on our attack.

Let us suppose, then, that close enough to the truth will normally mean ‘close enough not to obscure pragmatically relevant details or distinctions’. All sorts of factors play into which information is relevant and which is not. I must of necessity leave this part of the analysis vague, but assume that some elaboration of Grice’s (1975) theory of conversation will explain the contribution of the pragmatic situation to determining which details and distinctions may be ignored, even at the cost of not telling the truth.³

³ To sketch briefly: violations of Grice’s maxim of quality (calling for truthfulness) can be explained as clashes with the second maxim of quantity, which prohibits giving more information than is required. For

In the formal presentation below, I will simply speak of the pragmatic context *C* associating various sets with linguistic expressions, but it is understood that Gricean mechanisms are at work in determining this association.⁴

3. PRAGMATIC HALOS. There are any number of ways to model slack regulation and distance from the truth. This and the following sections offer one possible technique. The method I informally outline here (with a more formal discussion in the appendix) is intended to capture fairly directly the intuition that deviation from the truth is licensed because discourse participants find it convenient to ignore irrelevant distinctions and details.⁵ Slack regulators signal how much detail should be ignored.

The general idea is this: Each expression in the language, as usual, is assigned a denotation (relative to a model), which is used in calculating truth conditions. In addition, I assume that the pragmatic context associates this denotation with a set of objects of the same logical type as the denotation itself. Each object in this set is understood to differ from the denotation only in some respect that is pragmatically ignorable in context. The denotation itself is naturally included in this set. The phrase *three o'clock* for example, will be taken to denote a particular instant of time *i*; but the pragmatic context will also associate with *i* a set of times, each of which differs from *i* only in ways that may pragmatically be ignored in that context. These could be times that are just a few seconds before or after *i*, for example. Or the phrase *the townspeople* will denote a set of individuals. The context will also associate with this set a set of sets of individuals—sets that differ from the set of townspeople only in irrelevant ways, such as leaving out a few odd members whose quirky properties are pragmatically not a concern.

The set associated with the denotation of an expression may be ordered, either totally or partially, in such a way that the denotation forms a natural endpoint for the ordering. Let us assume that the denotation of *three o'clock* is some instant *i*; associated with this will be a set of times which differ from *i* only in pragmatically ignorable ways. This set may be ordered by the relation of 'closeness to *i*', such that *x* is closer to *i* than *y* is iff *x* falls between *y* and *i* on the time line.

Or, the relevant ordering on the set associated with the phrase *the townspeople* will be one of more closely approximating the set of actual townspeople, and so on. The relative position of the elements of a set according to such an ordering relation gives

example, a false claim that Mary arrived at three o'clock might be justified because a true claim would require giving superfluous information such as the number of seconds after 3:00 that she arrived; a false claim that the townspeople are asleep might be justified because a true claim would necessitate giving superfluous information such as a list of insomniac townspeople. But if Mary arrived several hours after three, or half the population of the town is insomniac, this information is no longer superfluous, and the second maxim of quantity no longer provides justification for violating quality.

⁴ As a referee points out, Grice (1975:45) allowed that the first maxim of quality might have a special status, and that 'other maxims come into operation only on the assumption that this maxim is satisfied'; the most familiar examples of quality violations are figurative, or, as the referee puts it, 'deviant': metaphor, irony, etc. In contrast, the analysis proposed here requires quality violations of a much less dramatic and more routine sort. If admitting this sort of maxim violation is a departure from Grice, it is not, I think, a very major one. As Grice continued, 'so far as the generation of implicatures is concerned [the first maxim of quality] seems to play a role not totally different from the other maxims.'

⁵ Of course in some cases, details are ignored of necessity rather than convenience, because they fall below the threshold of perceptibility. We might call a ball spherical for example, not because we are deliberately ignoring its oblong shape, but because its difference from a sphere is so slight that we literally cannot perceive it. Even in such cases, however, I think it is strictly false to say that the ball is spherical.

us a way of judging closeness to the truth—I will make use of these orderings in formulating the semantics of slack regulators.

Given an expression α denoting some object x , I like to think of the set the context associates with x as arrayed around x in a sort of circular cluster, so I will call this set, together with its ordering relation, the PRAGMATIC HALO of x , or, extending the terminology, as the pragmatic halo of α .

Complex expressions, and not just lexical items, may have pragmatic halos. Just how ‘wide’ the halo of a complex expression is depends in part on whether it contains a slack regulator as a constituent, such as *exactly*, *all*, or *perfectly*. This sensitivity to the appearance of specific lexical items suggests that some sort of compositional procedure is involved in halo assignment for complex expressions. A number of issues arise about the specific principles involved. Some of these issues are discussed in the appendix, but for here let us take a fairly simple procedure as the starting point for the discussion: assume that in the default case, the halo of a complex expression α is derived by applying normal semantic rules to all possible combinations of elements drawn from the halos of the immediate parts of α .

For example, suppose α has immediate parts β and γ , that β denotes some function f , γ denotes some object a in the domain of f , and the denotation of α is fixed by rule as the value which the function denoted by β returns when applied to the denotation of γ ; that is, $\llbracket \alpha \rrbracket = \llbracket \beta \rrbracket(\llbracket \gamma \rrbracket) = f(a)$.

Now, since β denotes a function f , its halo will be a set of functions, say $\{f, g, h\}$. And since γ denotes some object a its halo will be a set of similar objects, say $\{a, b, c\}$. To derive the halo of α , we simply take all possible combinations of a function in the halo of β applying to an object from the halo of γ : $\{f(a), f(b), f(c), g(a), g(b), g(c), h(a), h(b), h(c)\}$. (The technique is formally similar to that used for WH-questions in Hamblin 1973 or p-sets in Rooth 1985 and is easily generalized to a wider range of cases; see the Appendix.)

How is this set ordered? At a minimum, we may assume that the orderings of the halos of the constituent expressions are preserved in the ordering of the complex expression. So if b is closer to a than c is, then $f(b)$ will be closer to $f(a)$ than $f(c)$ is; likewise if g is closer to f than h is, then $g(a)$ will be closer to $f(a)$ than $h(a)$ is.

4. THE SEMANTICS OF SLACK REGULATION. Slack regulators are expressions that serve to readjust the pragmatic halo of the expressions they combine with. In fact, in all of my examples, this readjustment is a tightening—an elimination from the halo of those elements ordered furthest away from the core formed by the expression’s truth-theoretic denotation.

I assume a semantics in which verbs and other predicates have a hidden argument place for eventualities (that is, events, states and processes), roughly in the style of Davidson 1967; I will also assume that sentences denote sets of eventualities rather than truth values, so that *Mary arrived at three o’clock*, for example, denotes the set of Mary’s arrivals at three o’clock. A sentence is TRUE if its denotation is nonempty, FALSE if it is empty. So far as I can tell, nothing in the semantics of slack regulation requires such a system; I adopt it here simply as a convenient framework in which to discuss sentences with time reference. It should be possible to adapt the same basic techniques to systems of temporal semantics constructed along quite different lines, if desired.

4.1. EXACTLY THREE O’CLOCK. I turn now to example sentences, starting with the case of Mary’s not-quite-at-three-o’clock arrival. Let us assume that a phrase such as

three o'clock denotes some particular time i , and treat the preposition *at* as expressing a relation between eventualities and times, so that *at three o'clock* denotes the set of eventualities that occur at time i .

Let intransitive verbs denote relations between individuals and eventualities, so that *arrive*, for example, matches each individual x with events of x arriving (for the sake of simplicity, ignore tense); let *Mary* denote an individual as usual, and let *Mary arrived* denote the set of all those eventualities which this individual is paired with by the relation denoted by *arrive*—that is, the set of Mary's arrivals.

To derive the denotation of *Mary arrived at three o'clock*, we take the intersection of the set of eventualities denoted by *Mary arrived* with the set denoted by *at three o'clock*, yielding the set of Mary's arrivals which occur at three o'clock. The sentence is true iff this set is nonempty, false if it is empty.

What about the pragmatic halos? Suppose, for illustration, that there are two points in time close enough to i that the difference between them and i is ignored in context, so that the halo of *three o'clock* is the set $\{i, j, k\}$, ordered according the relation of closeness to i .⁶ To keep things simple, let us suppose that this is the only lexical expression in the sentence with a nontrivial halo; for the rest, simply assume that the halo contains the expression's denotation and nothing more.

From the principles set forth so far, it should be evident that the sentence *Mary arrived at three o'clock* will have as its halo a set with three sets of events as members, which we may regard as the set of Mary's arrivals at i , the set of her arrivals at j , and the set of her arrivals at k .

We will count a sentence as 'close enough to true for context C ' iff its halo relative to C contains at least one nonempty element. For example, suppose that Mary arrived only once, at k . With denotations and halos as given above, *Mary arrived at three o'clock* will have the empty set as its denotation (so the sentence is false), but its halo will be the set $\{\{x|x \text{ is an arrival by Mary at } i\}, \{x|x \text{ is an arrival by Mary at } j\}, \{x|x \text{ is an arrival by Mary at } k\}\}$. Where e is the event of Mary's arrival, this set is equal to $\{\emptyset, \emptyset, \{e\}\}$, or more simply $\{\emptyset, \{e\}\}$. Since this set contains at least one nonempty member, the sentence is close enough to true for its context, even though it is false.

Now to slack regulation. We take *exactly* (as applied to times)⁷ to be trivial, from a truth-conditional point of view: it simply denotes the identity function on times, so that *exactly three o'clock* has the same denotation as *three o'clock*.

The real effect of *exactly* is on pragmatic halos: we want the pragmatic halo of *exactly three o'clock* to include those elements of the halo of *three o'clock* which are closest to i (that is, to the actual time of three o'clock), eliminating outlying elements.

Shouldn't we actually eliminate ALL elements of the halo of *three o'clock* other than three o'clock, to assure that *Mary arrived at exactly three o'clock* is only close enough

⁶ The use of a finite set here is a simplification; assuming time is dense, the halos of time phrases will normally be infinite.

⁷ *Exactly* seems to behave differently when it applies to times than when it applies to cardinal numbers, where it does appear to have a truth-conditional effect. Sentence (i) seems true even if John has more than seven chairs, while (ii) does not.

(i) John has seven chairs.

(ii) John has exactly seven chairs.

It would be theoretically attractive to claim that (i) and (ii) are truth-conditionally equivalent, since this would allow a unified treatment of the temporal and cardinal uses *exactly*. But, I find that this idea is so gross a violation of my intuitions about the meaning of (i) that I cannot advocate it, despite its theoretical attractions. See Horn 1992, Geurts 1998 and the references cited there for discussion.

to the truth if she arrives right on the button? Well, maybe in some contexts, but even the word *exactly* sometimes leaves a little pragmatic slack. How often do we worry about milliseconds?⁸

Exactly, like any other expression, will have a pragmatic halo, provided by the context, but we shall have to place some constraints on it. Since the denotation of *exactly* is an identity function on times, the halo should be a set of functions on times. Each of these should differ from an identity function only in pragmatically ignorable ways. What would this mean? Presumably, a function that differs from an identity function only in pragmatically ignorable ways would be a function that maps each time t in its domain onto a time differing from t only in ignorable ways.

In effect, these functions together map t onto a set of very close times, which we may regard as a kind of 'extra-tight' halo surrounding t . That is, we can obtain an appropriate halo for *exactly three o'clock* by applying these functions to three o'clock itself, and gather the resulting values into a set. I adopt this strategy, though it means that the halo of this phrase will be derived by a somewhat different procedure than what is given by the default rule sketched above.

Now, to maintain the claim that the halo of a complex expression is derivable from the halos of its parts, we will have to recover the time three o'clock from the halo of the phrase *three o'clock* so that it can serve as the argument to the functions in the halo of *exactly*. Fortunately, this is possible, since *three o'clock* will always be the centerpoint of the halo. To calculate the halo of a phrase of the form *exactly T*, we apply the functions in the halo of *exactly* to the centerpoint of the halo of T , and take the union of the resulting values.

To guarantee that this procedure yields a reasonable halo and that it at least contracts, rather than expands, the halo of three o'clock, we should stipulate that the functions in the halo of *exactly* together should map a given time t onto a subhalo of the halo of t . That is, every time in the halo of *exactly T* should be in the halo of T , and if a time y is in the halo of *exactly T*, so should every time between y and T , as well as T itself.

As an example, assume again that *three o'clock* = i and that the halo of i is $\{i, j, k\}$, ordered according to closeness to i . Suppose that j is closer to i than k is. Then the halo of *exactly three o'clock* will be some subset of $\{i, j, k\}$ which at least contains i , and which may not contain k unless it also contains j .

Suppose that the halo of *exactly three o'clock* contains i and j but not k . If we assume as before that Mary arrived only once and that this was at k , we find that the halo of

⁸ An anonymous referee suggests that the halo should contract as far as possible, 'given contextually accessible measuring devices and practices'. This does not seem right, however, for two reasons. First, when fine discriminations of time are pragmatically relevant, but the available means of measurement are coarse, the use of *exactly* is not licensed. Suppose, for example, we are worried whether Mary made it to the train station on time or not, but our only contextually available time-measuring practice is to regard the position of the sun in the sky. Unless we are unusually expert at time-reckoning, we would not say that Mary arrived at exactly three o'clock in such a situation, even if we know she arrived as close to three o'clock as we can gauge, given the means at hand.

Second, in contexts where we do have devices available for measuring fine discriminations, but where these fine discriminations are pragmatically irrelevant, the word *exactly* can be felicitously used even if we are not committed to as close an approximation as our devices are capable of. Lots of people habitually wear watches with second hands, but we might reasonably describe Mary's arrival at a party, for example, using the sentence *Mary arrived at exactly three o'clock* even if she arrived more than a second afterwards. The most we can say, I think, is that the pragmatic context determines how tight a halo contraction *exactly* gives us, and that this varies from context to context.

the sentence *Mary arrived at exactly three o'clock* = $\{\emptyset, \emptyset\} = \{\emptyset\}$; since this set does not have any nonempty members, the sentence is not close enough to true for our context—although the sentence *Mary arrived at three o'clock* was close enough, even relative to the same denotations and halos.

4.2. ALL THE TOWNSPEOPLE. In treating the difference between *the townspeople are asleep* and *all the townspeople are asleep*, we will have to consider several issues in the semantics of plurality, which are largely independent of the semantics of slack regulation.

Sentences containing plural noun phrases with *the* or *all the* are often ambiguous, and may be interpreted either collectively or distributively. On the collective reading, the sentences 18 and 19 (from Link 1983:311) attribute raft building to the children as a group.

(18) The children built the raft.

(19) All the children built the raft.

On this reading, these sentences are not equivalent to those in 20 and 21, which attribute raft-building to each of the children individually.

(20) Every child built the raft.

(21) Each child built the raft.

To capture collective readings, we should let *the children* and *all the children* denote the group of children directly, with no quantificational force of their own. I model this 'group' simply as the set of children, and allow the predicate *built the raft* to take this set directly as an argument.

How then to account for the distributive readings of 18 and 19, and for the fact that 22 and 23 do entail that the individual townspeople are asleep?

(22) The townspeople are asleep.

(23) All the townspeople are asleep.

A broad array of arguments suggests that this effect is not due to any ambiguity in the noun phrases, but rather to the semantics of the predicates involved. (See Roberts 1987, Lasersohn 1995, and the references given there.) Accordingly, even in sentences like 22 and 23, I will treat the noun phrases as denoting the set of townspeople collectively, and let this set serve as an argument to the predicate *are asleep*. A quantificational effect is obtained by placing constraints on the semantics of the predicate.

I have been treating predicates as having a hidden argument place for eventualities, but if we allow not just individuals, but sets of individuals, to serve as arguments to a predicate, we should also consider letting not just eventualities, but also sets of eventualities, to serve as arguments. I will in fact assume that predicates can take sets of eventualities as arguments, and stipulate that a set of individuals *X* will stand in the *sleep* relation to a set of eventualities *Y* iff every member of *X* stands in the sleep relation to a member of *Y* and vice versa (and similarly for other lexically distributive predicates). This gives the effect that a group of individuals cannot be asleep without the members of the group being asleep. In the case of predicates that show an ambiguity between collective and distributive readings, I assume the distributive reading results from a hidden operator attached to the verb phrase, roughly as in Link 1987a and Roberts 1987.

Now, the halo of the noun phrase *the townspeople* should be a set of sets of individuals, which differ from the set of townspeople only in ways that are pragmatically irrelevant in context, ordered according to closeness to the actual set of townspeople.

For our purposes, we may assume that this ordering is just what the subset relation gives us, so that, where T is the set of townspeople, A counts as being at least as close to T as B is iff either $T \subseteq A \subseteq B$ or $B \subseteq A \subseteq T$. (Ultimately we may want to consider more sophisticated closeness relations.)

I treat *all*, like *exactly*, as denoting an identity mapping; its halo, like that of *exactly*, should be a set of functions approximating an identity function—that is, mapping each set onto a set that differs from it only in pragmatically ignorable ways. As before, we derive the halo of *all the N* by gathering into a set the values which these functions return for the centerpoint of the halo of *the N*—in the case of *the townspeople*, this centerpoint will be the set of townspeople. And as with times, we need to stipulate that these functions map any set A onto a subhalo of the halo of A (so every set in the halo of *all the N* will also be in the halo of *the N*, and if a set X is in the halo of *all the N*, so will any set Y which is at least as close to the denotation of *the N* as X is).

It should be possible to see now that *the townspeople are asleep* will be close enough to true in some contexts where *all the townspeople are asleep* is not, even though the sentences are truth-conditionally equivalent.

Note that if we analyze 24 as true iff at least one townspeople is awake (and if we assume that no one can be asleep and awake at the same time), then 24 and 25 will never both be true, and 26 will be contradictory—even though 25 pragmatically allows exceptions.

(24) Some of the townspeople are awake.

(25) The townspeople are asleep.

(26) Although the townspeople are asleep, some of them are awake.

Ex. 26 is not just contradictory, though, it is also quite odd pragmatically. But given the way I have set up the theory, there is no guarantee that contradictions are pragmatically odd. After all, a contradiction is just a sentence that is false in all circumstances; but I am explicitly allowing some false sentences to function pragmatically as though they were true. Isn't there a risk, in the sort of system I have set up, of erroneously labeling this sentence as 'close enough to true' for some contexts? What we want is for it never even to be close enough to the truth pragmatically, which is something stronger than mere truth-conditional contradictoriness.

I think we can understand the oddity of this sort of example if we bear in mind that utterances are licensed only when pragmatically relevant, and that pragmatic relevance is precisely what determines halo membership. One could felicitously assert the second clause in 26 only if there were some pragmatic relevance to the fact that some of the townspeople were awake—that is, if there were some pragmatic relevance to the distinction between the set of townspeople as a whole and the set of all the townspeople but the awake ones. But in that case, this distinction would not be ignorable in context, so the latter set would not be in the halo of *the townspeople*. Hence the first clause of 26 will never be close enough to true for its context if the context is one in which the second clause is assertable.⁹

⁹ Several people have suggested to me that the pragmatic oddity of this example might be tied in some way to the anaphoric relation between *them* in the second clause and *the townspeople* in the first clause. We might, for example, claim that the halo of a noun phrase represents different possibilities for who or what the noun phrase is being used to 'refer' to, in some pragmatic sense of reference (perhaps to be identified with Kripke's (1977) *SPEAKER'S REFERENCE*). Then we could claim that whichever element of the halo a noun phrase is used to refer to, any pronoun that is anaphoric to that noun phrase must refer to that same element. In this way, the second clause of 26 would always be interpreted as meaning that some of the group referred to in the first clause were awake, and we would get the pragmatic effect of a contradiction even when a great deal of leeway is allowed as to the referent of *the townspeople*.

4.3. PERFECTLY SPHERICAL. This example is of special interest because we want to capture the contribution of *perfectly* not just for examples like 27 but also for examples like 28.

(27) The ball is perfectly spherical.

(28) The ball is perfectly round.

We saw above that *round*, unlike *spherical*, is scalar; it accepts modification by *very* and similar expressions of degree. But it may also appear in sentences like 29, which does not contain a degree modifier.

(29) The ball is round.

There are a number of ways to treat this sort of sentence; for the sake of argument, let us assume that the pragmatic context provides a cutoff point the ball must reach if 29 is to be true (relative to that context). Note that the concern here is authentic truth, not pragmatic closeness to truth; scalarity gives rise to real truth-conditional vagueness, and not just pragmatic slack.

Now we face an interesting problem. If the job of *perfectly* in 27 is just to contract the pragmatic halo around the denotation of *spherical* shouldn't it do the same thing in 28, contracting the halo around the denotation of *round*? But suppose we are in a context that sets the cutoff point for *round* quite low on the scale, so that 29 is true in that context even though the ball is quite oblong. That is, in such a context, the denotation of *round* will contain the ball, even though it is oblong. But contracting the halo around the denotation of *round* does not remove anything from the denotation itself; the oblong ball will still be in the denotation even after any halo-contraction. Thus, contracting the halo will not give us the set of perfectly round objects at all, and we do not obtain the correct semantics for 28.¹⁰ We need a semantics for *perfectly* that manipulates scales for those predicates that provide them, in addition to manipulating pragmatic halos.

To do this, we will have to be a little more explicit about scales and scalar predicates. Suppose that a scalar adjective such as *round* receives a different denotation relative to different contexts, but that relative to each context, it simply denotes a set of individuals. Assume also that contexts may be ordered according to strictness, and that if context C_1 is stricter than context C_2 in the ordering, then the denotation of *round* relative to C_1 is a subset of the denotation of *round* relative to C_2 .

The notion of a scale may be easily defined in terms of this ordering, if desired; intuitively, position in the ordering corresponds directly to the position of the scalar cutoff point imposed by the context. At one end of the ordering, we find the denotation of *round* in the strictest possible contexts, which presumably contain only perfect circles, spheres, disks, and so on.

The truth-conditional contribution of *perfectly* involves an implicit quantification over contexts: an object a is perfectly round iff it falls into the denotation of *round* in all contexts (hence in the strictest contexts). Extending this idea to adjectives in general, note that if an adjective is not scalar (and hence does not show linearly ordered context-

But, Christine Brisson points out to me that pragmatic reference is not necessarily maintained in anaphora, as this analysis requires. Sentence (i), for example, seems acceptable even if not exactly the same boys went swimming as went fishing, provided both groups are good approximations of the group of boys as a whole.

(i) The boys went fishing and then they went swimming.

¹⁰ Thanks to Fred Landman for raising this issue to me.

tual variation in denotation), modification with *perfectly* will be semantically trivial, with no truth-conditional effect.¹¹

So far so good, but now consider adjectives such as *spherical*. Any contribution *perfectly* makes here should come from its effect on pragmatic halos, since *spherical* is not scalar. I assume that *spherical* denotes the set of perfect spheres. What will its halo be like? This should be a set of sets that differ from the set of perfect spheres only in pragmatically ignorable ways, for example by containing some objects whose shape deviates from that of a sphere in minor ways.

As with the other examples, we will want to impose a partial ordering on this halo, to represent closeness to the set of spheres. This raises several issues, but I will simply assume that (where X and Y are sets) X approximates the set of spheres at least as closely as Y iff the least round member of X is rounder than the least round member of Y . (More generally, X approximates Z at least as closely as Y iff the member of X that requires the greatest deformation to be converted to a member of Z requires less such deformation than the member of Y that requires the greatest deformation to be converted to a member of Z .) Note that the endpoint of this ordering will be the set of perfect spheres.

Now let the halo of *perfectly* be a set of functions approximating its denotation—that is, a set of functions, each of which maps a given set X onto a set that differs from what the denotation of *perfectly* maps X onto only in pragmatically ignorable ways. As with previous examples, the halo of *perfectly spherical* is derived by gathering into a set the values given by these functions for the endpoint of the halo of *spherical*.

As with *exactly* and *all*, some constraints must be placed on the functions in the halo of *perfectly*, so that any set in the halo of *perfectly spherical* will also be in the halo of *spherical*, and so that if Y is in the halo of *perfectly spherical* and X approximates the denotation of *spherical* at least as closely as Y does, then X must be in the halo of *perfectly spherical* as well. (A more precise consideration of these conditions is delayed until the Appendix.)

5. COMPARISON TO ALTERNATIVES.

5.1. STANDARDS OF CLOSENESS VS. STANDARDS OF TRUTH. The technique of halo manipulation contrasts with an alternative, and perhaps more familiar, way of approaching the problem of linguistic precision. In this alternative approach, one does not treat standards of precision as determining how close one must come to the truth; rather, they determine what is true and what is not. That is, truth itself is relativized to standards of precision.

In fact, we have just applied such a theory in analyzing the scalarity and truth-conditional vagueness of predicates like *round*. But one of my claims in this article is that pragmatic slack is not the same thing as scalarity or truth-conditional vagueness, but a separate phenomenon over and on top of it. To see the differences more clearly, let us consider in more detail the idea of relativizing truth to standards of precision, and how it could be applied to the examples here.

In this sort of analysis, one might claim that *Mary arrived at three o'clock* is true relative to one standard of precision even if she arrived at fifteen seconds after three, but false relative to a stricter standard. Similarly, one could claim that *the townspeople*

¹¹ Note that quantifying over contexts is not equivalent to modal quantification over possible states of affairs. To evaluate a sentence such as *John is perfectly happy*, for example, we do not imagine whether John would be happy no matter what context we place him in; rather we consider whether John's actual state would qualify as happiness no matter what context provides our standard of evaluation.

are asleep is true relative to a fairly liberal standard of precision even if some of the townspeople are awake, but false relative to a stricter standard. The same sort of claims could be made about *this ball is spherical*, and so on.

In this sort of system, we would assume that the pragmatic situation of utterance determines which standard of precision is in force on a particular occasion. Slack regulators are analyzed not as affecting the choice of standards, but as helping determine the truth value relative to a given standard. For example, *exactly* might be analyzed in such a way that 30 is true relative to a given standard *s* iff 31 is true relative to EVERY standard *s'* (including even the strictest standards).

(30) Mary arrived at exactly three o'clock.

(31) Mary arrived at three o'clock.

All could be given a similar analysis, so that 32 is true relative to *s* iff 33 is true relative to every standard *s'*, assuring that 32 allows no exceptions even if *s* is liberal enough to allow exceptions for 33.

(32) All the townspeople are asleep.

(33) The townspeople are asleep.

Perfectly, and no doubt many other slack regulators, could be given similar analyses.

This kind of analysis contrasts with the view advocated above in that it gives slack regulators a truth-conditional effect. Thus, 30 and 31 will no longer be logically equivalent, since 31 can be true while 30 is false; likewise 33 can be true while 32 is false.

Some might find this a welcome effect, since there does seem to be a difference in meaning in these examples, but these truth-conditional differences also produce some unpleasant side effects. For one thing, we need to deal with examples like 34:

(34) Absolutely all the townspeople are asleep.

Absolutely appears to be a slack regulator here, but do we really want to claim that 34 is truth-conditionally distinct from 32? How?—by letting 32 be true even if some of the townspeople are awake? This seems unpalatable. And if 33 can be true while 32 is false, it then becomes difficult to solve one of the original problems: how do we explain that 35 is contradictory?

(35) Although the townspeople are asleep, some of them are awake.

Some of the townspeople are awake should be true whenever *all the townspeople are asleep* is false. If we also claim that *the townspeople are asleep* can be true even while *all the townspeople are asleep* is false, there should be no contradiction.

Perhaps the analysis could be saved by fiddling with the semantics of *some*, for example by claiming that *some of the townspeople are awake* simply doesn't receive a truth value relative to those standards of precision that are liberal enough to make 33 true and 32 false. But then it no longer follows from *all the townspeople are asleep* that it is false that *some of the townspeople are awake*, which seems very wrong. Moreover, it will be difficult to get truth-value gaps in just the right places, especially since *some* occurs perfectly naturally in discourses where 33 is interpreted liberally enough to allow exceptions. We would not want to claim that the second sentence of 36, for example, lacks a truth value relative to any standard of precision liberal enough to allow exceptions to the first sentence.

(36) The townspeople are asleep. In fact, some of them are snoring.

These problems seem to me great enough to make the analysis quite unattractive. We may well want to claim that sentences like 33 allow exceptions relative to some stan-

dards of precision but not others. I had better give up on the idea that the choice of a standard affects its truth value.

These same considerations apply, I think, to other examples, since I want to account for the contradictoriness not only of 35, but also of examples like 37 and 38.

(37) Although Mary arrived at three o'clock, she didn't arrive until slightly after three o'clock.

(38) Although the ball is spherical, it is oblong.¹²

Suppose the ball is close to a sphere but slightly oblong. If *s* is a liberal enough standard of precision, it will make 39 true, but if it is this liberal, it will surely also make 40 true.

(39) The ball is spherical.

(40) The ball is oblong.

Examples like this suggest that something different is going on here than just relativization of truth to degrees of precision.

None of this prevents us from retaining a relativization of truth values to standards of precision in the semantics of predicates like *bald* and *heap*, which gives rise to the sorites paradox. Here, in fact, the relevant examples are not contradictory:

(41) Although John is bald, he has some hair.

We should distinguish, I think, between authentic semantic vagueness, in which the extension of a predicate really does not have well-defined borders, and mere pragmatic looseness of speech, in which we allow speakers a certain leeway even in the use of predicates whose extensions are quite sharply defined.

5.2. SCOREKEEPING AND SUPERVALUATIONS. An analysis somewhat closer to the one I advocate here was suggested by Lewis 1979, but that too seems to miss an important point. Lewis assumes a theory of vagueness based on van Fraassen's (1966) technique of SUPERVALUATIONS (see also Kamp 1975, Fine 1975). That is, if John is a borderline case of baldness, then a sentence such as *John is bald* is analyzed as simply not having a determinate truth value. However, a complex sentence like *John is bald or not bald* will have a determinate truth value, assigned by quantifying over ways of assigning a precise border between the bald and the nonbald. A sentence is 'simply true' if it is true relative to all ways of drawing the border, 'simply false' if it is false relative to all ways of drawing the border, and lacks a determinate truth value if it is true relative to some ways but false relative to others.

Lewis notes, though, that even if a sentence is false relative to some ways of drawing the relevant border, one might still treat the sentence 'more or less as if it is simply true', provided it is true for ENOUGH ways of drawing the border. In this case, Lewis calls the sentence 'true enough'. Exactly how many ways of drawing the border count as 'enough' varies according to context, and depends on which standard of precision is in force. A sentence may be true enough relative to a context employing a loose

¹² A referee suggested that this example did not really sound so bad, and noted that (i) below was perfect.

(i) The yard is triangular, but its corners are rounded.

My informants unanimously agreed that a yard could be triangular but with rounded corners; but they rejected by a margin of approximately ten to one the claim that a ball could be spherical yet oblong. (A similar margin rejected the claim that the townspeople could be asleep even while some of them are awake.) I can only conclude that not all geometric terms are the same, and that, at least in some contexts, *triangular* can be truthfully applied to objects that are not perfect geometric triangles. One informant explained his judgments by suggesting that *spherical* was 'more of a technical term' than *triangular*.

standard, even while failing to be true enough relative to contexts employing stricter standards.

Lewis suggests such an analysis not only for *bald*, but also for sentences such as *France is hexagonal* and *Italy is boot-shaped* and for the adjective *flat*. It seems reasonable to consider whether the same technique might also be adopted for the main examples in this article. We might, for example, treat slack regulators as manipulating how many ways of drawing a border count as enough in a given context, in which case they would affect the status of a sentence as true enough, but would not affect its status as simply true or simply false.

From my perspective, Lewis's term TRUE ENOUGH is objectionable, since it seems to imply that the issue is how true the sentence is, rather than how close to the truth it has to come. But this is probably not an important issue; Lewis appears to intend his term merely as a convenient shorthand label, and not as representing a serious commitment to degrees of truth.

A stronger ground for objection is that Lewis's system allows for contextual variation in whether a sentence is true enough ONLY when the sentence receives an indeterminate simple truth value—that is, only when the sentence is truth-conditionally vague. Sentences that are simply true will always be true enough, and those that are simply false will NEVER be true enough, because the whole notion of TRUE ENOUGH can come into play in a nontrivial way only for those sentences in the indeterminate border area between truth and falsity. Lewis allows a sentence to be true enough even while failing to be true—but in this case the sentence must also fail to be false; there is no way to capture the intuition that a sentence like *Mary arrived at three o'clock* might be false even while we treat it pragmatically as though it were true. For Lewis, contextual variation in whether or not a sentence counts as true enough arises only because there is an undefined border area around the extension of some predicates, not because we disregard certain types of pragmatically irrelevant falsehood.

5.3. MODELS AND DOMAIN MANIPULATION. Quantificational sentences can be judged true or false only relative to a particular domain of quantification, which is normally fixed pragmatically. For example, if, at the beginning of class, a teacher says 'Everyone is here, let's begin', the first clause will normally be deemed true if everyone enrolled in the class is present and false otherwise; the sentence does not claim that everyone in the entire universe is present.

Classically, we regard each model for the interpretation of a language as coming with its own domain of quantification, which serves also as the domain from which denotations for constants are built up. Of course some of these models will have quite inappropriate domains, but pragmatic factors will constrain the domain of the intended model. (A sentence is regarded as true if it is true relative to the intended model, and false otherwise. The other models are used primarily in defining logical relations like entailment and equivalence.) Specifically, the domain of the intended model should contain only those individuals that are pragmatically relevant to the discourse at hand. Irrelevant individuals are simply not represented in the model.

It is increasingly popular to assume that each quantificational noun phrase may introduce its own domain of quantification, or CONTEXT SET, rather than requiring all quantification to range over the entire domain for the model. (See especially Westerståhl 1985, von Stechow 1994.) But this does not affect the general point that the domains over which we quantify and from which we build up denotations for constants contain only pragmatically relevant individuals.

The same considerations apply not only to individuals, but also to any entities that we might quantify over or refer to, for example, times. In doing temporal semantics, one might consider a system in which the intended model for a discourse of English contained only those times that are pragmatically relevant to that discourse, and not the 'whole' structure of time, for example.¹³ (Or one might employ a whole hierarchy of increasingly fine-grained representations within a single model, as suggested in Link 1987b.)

In analyzing linguistic imprecision, we might also try exploiting the fact that irrelevant entities are not represented in the intended model as a means of representing pragmatic slack. Slack regulation could then be modeled as domain manipulation.¹⁴

All, for example, might be a signal that the domain of quantification is being expanded, essentially as in Kadmon and Landman's (1993) analysis of *any*. Just as Kadmon and Landman regard noun phrases of the form *any N* as semantically just like noun phrases of the form *a N*, but with the denotation of *N* 'widened' to include individuals that might otherwise have been regarded as pragmatically irrelevant, noun phrases of the form *all the N* can be regarded as semantically like those of the form *the N*, but again, with the denotation of *N* widened to include individuals that might otherwise be regarded as irrelevant. Hence *the townspeople* might leave out of its denotation a few odd exceptional townspeople, but the denotation of *all the townspeople* will be expanded to include even them.

Kadmon and Landman do not make the formal representation of widening completely clear, but there seem to be two options: (1) in the intended model, the domain of discourse may include some townspeople who are not in the denotation of the noun *townspeople* (unless it is widened by the presence of *all* or *any*); or (2) the model does not contain any townspeople other than those in the denotation of *townspeople* (even before widening), but widening involves changing to a different model, with more townspeople in it. In this second strategy, *all* or *any* would signal which model is the intended model, rather than having an effect on compositional denotation assignment in the usual sense.

The first strategy does considerable violence to the notion of denotation, since now the denotation of a noun (prior to widening) can no longer be identified with the set of things of which the noun may be truthfully predicated. Moreover, neither approach seems capable of explaining examples like 42 (mentioned above in n. 9).

(42) The boys went fishing and then they went swimming.

This sentence could be used felicitously in a situation where not exactly the same boys went swimming as went fishing, provided both groups are good approximations of the group of boys as a whole. But which group serves as the unwidened denotation of *boys*? Neither group seems appropriate. However, a pragmatic halos account can handle this sort of example with no problems, since both clauses of the conjunction may be close enough to true for their context even if the two groups are not identical.

A similar problem arises in examples involving conditional sentences. Suppose I am

¹³ Such an analysis might provide a means to account for examples like (i), from Partee (1973), even while treating tense morphemes as quantificational operators.

(i) I didn't turn off the stove.

While (i) may not mean that there was no time whatsoever at which I turned off the stove, it could plausibly be analyzed as meaning that there was no pragmatically relevant time when I did so. See, however, the problems for this general technique discussed below.

¹⁴ Thanks to a referee for suggesting that I consider this sort of analysis.

about to enter a casino, and say to you, 'If I lose the money, I'll meet you back here at 2:00.' In making this statement, I surely intend to include a situation in which I lose everything but a few odd pennies. In a domain-widening account, this means that the unwidened denotation of *the money* must not include those few odd pennies. But which pennies, specifically, are excluded? Choose any one of them; if it is the only one left, it hardly matters, so it should be excluded from the unwidened denotation. But what if the money consists of nothing but pennies? By this logic, the unwidened denotation of *the money* should be empty. In a pragmatic halos account, the speaker is pragmatically committing to the consequent clause as long as the antecedent clause is close enough to true for its context, which will be the case no matter which few pennies are left. A pragmatic halos analysis has the advantage of allowing us to consider the full set of reasonable approximations to the denotation at once, while a domain-widening analysis requires us to choose just one of these as the 'unwidened' denotation.

If we try to extend the technique to a wider range of examples, we encounter even more problems. Suppose, for example, we try to account for imprecise time reference by claiming that in the intended model, only pragmatically relevant times are represented, and in particular, that overly fine distinctions of time are not represented. In that case, if we are in a context where it is irrelevant to measure time down to the second, or even down to the minute, the smallest times in the model might be five-minute intervals, for example. Now, suppose that in the real world, Mary arrives at precisely fifteen seconds after three o'clock, and the intended model is like the real world, except that it contains no intervals shorter than five minutes. How can we assign a truth value to a sentence such as *Mary arrived at three o'clock* relative to this model? There are several options.

First, since three o'clock, being an instant, is simply not included in the model, we might claim that the sentence is false, or an example of presupposition failure. I think in fact the sentence is false in this situation, but just saying this much will not explain why the sentence is felicitous, which is what we were after in moving to a coarse-grained model to begin with.

Second, the phrase *three o'clock* could denote some time other than three o'clock itself. But again, this does direct violence to the notion of denotation; surely, what it means for *three o'clock* to denote a time *t* in the intended model is precisely that *t* is three o'clock. Even if we accept this idea, we are faced with the question of which other time will serve as the denotation. If we want the sentence to come out true as long as Mary arrived within five minutes of three o'clock, we somehow need all the five-minute intervals which include three o'clock to count as though they were the denotation. But of course there are an infinite number of such intervals, and we cannot limit ourselves to any one of them in particular. We might say that it denotes the whole set of these intervals, but this means that the logical type of the phrase will vary from model to model; it denotes a time in some models, and a set of times in others; with this sort of variation in type from model to model, the statement of compositional rules also becomes much more complicated.

A related problem comes up with the denotation of the verb. I have been assuming that verbs denote relations between individuals and events, Davidson-style; each event is associated with a time, understood as the time at which it occurs. Suppose *e* is an event that occurs at precisely fifteen seconds after three in the real world: which time will it be associated with in the model? Here again, we want to associate it with all the five-minute intervals that include three o'clock.

But to formally associate a short event with a longer interval or intervals in this way

is to suppress the distinction between the time *AT* which an event occurs and the times *IN* which it occurs. It is precisely this distinction that we are likely to need if we are to model the difference between the various aspectual classes. How are we to distinguish a punctual event from one that has an actual duration of five minutes? Yet it seems indisputable by now that punctuality and durativity are important semantic categories that must be distinguished from one another.

Moving away from an event-based semantics does nothing to solve the problem. Suppose for example, that intransitive verbs are taken as denoting sets of individuals, but index or relativize these denotations to times, as in the usual formalization of tensed predicate logic. Which individuals are included in the denotation of *arrive* relative to, say, the five minute interval centered on three o'clock? If we include all those individuals who arrive *DURING* that time, we once again lose the distinction between what happens in a time and what happens at a time.

Finally, and perhaps most importantly, this sort of analysis suffers from the same defect as any analysis that treats slack regulation as part of the assignment of truth values: it allows a sentence like *Mary arrived at three o'clock* to be assigned a value of 'true' even if she arrived only at some other time—that is, even if the sentence is false. There is no appeal here to tolerance of falsehood or approximation to the truth, yet these notions seem intuitively to be involved in the interpretation of sentences of this kind.

5.4. THE ANALYSIS OF *all*. The meaning of *all* and in particular the differences among *all*, *every*, *each*, and *the* have been very widely discussed, and it is beyond the scope of this article to give a detailed comparison of my analysis to the previous literature. But it may be useful to consider one prominent proposal, that of Dowty (1986).

Dowty presents a series of puzzles about *all*. In some examples, it simply seems to force a distributive reading, much like *every* or *each*.

(43) The students voted for the proposal.

(44) All the students voted for the proposal.

Sentence 43 may be interpreted as meaning merely that the proposal passed. That is, a majority of the students (or whatever number is required), voted for it, but it is entirely compatible with this sentence that some individual students voted against the proposal rather than for it. In contrast, 44 means that every student voted for the proposal.

Unlike *each* or *every*, however, *all* combines naturally with collective predicates.

(45) All the students gathered in the hall.

(46) ?Every student gathered in the hall.

In these cases, we do not obtain a distributive reading. Sentence 45 does not mean that each student gathered in the hall, which would hardly make sense anyway, since a single individual cannot gather.

The situation is made even more complicated by the fact that some collective predicates do not allow *all*, any more than *every* or *each*.

(47) *All the students are numerous.

(48) *Every student is numerous.

Dowty suggests that these facts may be accounted for by claiming that *all* always does have a kind of distributive effect; however, it is not necessarily the predicate itself that distributes to the members of the group denoted by the noun phrase modified by *all*, but rather the 'subentailments' of this predicate.

Dowty does not give a formal definition of subentailments, but the idea is that some

predicates, in order to apply collectively to an entire group, impose certain requirements on the members of the group. For example, *gathered in the hall* is a predicate that properly applies only to whole groups, not to individuals. But if *gathered in the hall* holds truthfully of a group *G*, then certain facts must hold about the individual members of the group: they must come into the hall, and remain long enough that they are all there at a common time. This requirement is a subentailment of the predicate *gathered in the hall*. Predicates that impose subentailments thus exhibit what Lasersohn 1995 calls PARTICIPATORY DISTRIBUTIVITY.

Once subentailments are recognized, ordinary distributivity emerges as just a special case. A predicate like *be asleep* applies to a group only if its individual members are asleep; this requirement is a subentailment of *be asleep*, and may be treated in parallel fashion to the subentailments of *gather in the hall*. In some cases, a predicate may impose subentailments on some but not all members of a group. For example, 43 certainly requires that SOME of the individual students voted for the proposal—enough in fact, for the proposal to pass. This again seems like essentially the same phenomenon as *gather in the hall* or *be asleep*; the difference is merely that the requirement is not imposed on quite so large a proportion of the group's members.

Now, Dowty suggests, we may analyze *all* as imposing a requirement that the subentailments of the predicate must be satisfied by each individual member of the relevant group. Thus 44 means that each individual student voted for the proposal, for example. Likewise, *all the children are asleep* will require each individual child to be asleep, and 45 will require each individual student to be in the hall.

Dowty suggests that the anomaly of examples like 47 can be explained by the fact that *be numerous* has no subentailments. All that is necessary for a group to qualify as numerous is that it be of sufficiently large cardinality; no requirements are imposed on the individual members. Since the purpose of *all* is to indicate distribution of subentailments, it cannot be felicitously used with predicates that lack them.

Dowty's analysis is problematic in several respects, most of which he points out himself.¹⁵ In some cases, a predicate imposes a requirement on some members of a group, but *all* does not distribute this requirement to the rest of the group as the analysis requires. If, for example, a group of integers sum to thirteen, at least one member of the group must be odd, but 49 does not require that all the integers be odd.

(49) All these integers sum to thirteen.

Another problematic case involves group actions in which each member performs a different task, as in 50.

(50) All the students built that cabin.

This sentence might describe a situation in which one student laid the foundation, a different student put up the frame, a different one put on the roof, and so on. There does not seem to be any action or property that EVERY member of the group has to have—yet the analysis requires some subentailment that is imposed on every individual member.

Dowty suggests that the relevant subentailment in this case is simply that every student be a member of the group of students that collectively built the cabin. But this is really a requirement on the group as a whole and not on the individual members. If we allow requirements this weak to qualify as subentailments, there no longer seems to be anything to rule out examples like 47. Why isn't *all* licensed here, since every

¹⁵ See Taub 1989 for another critique of Dowty's analysis.

student has the property of being a member of a group of students with sufficiently great cardinality?

Yet another problem is that the analysis predicts that *all* will be semantically vacuous in cases where the predicate is authentically distributive (that is, where the predicate holds true of a group if and only if it holds true of the group's members). In that case, the subentailments of the predicate are required to hold of all the group's members anyway, so *all* adds nothing. Yet there certainly seems to be some difference between, for example, 51 and 52.

(51) The children are asleep.

(52) All the children are asleep.

Dowty's response is to claim that even predicates that are normally considered distributive can hold of a group without holding of every member of the group; hence 51 can be true even if some of the children are awake. The exact proportion of members which must satisfy a predicate in order for the group as a whole to satisfy the predicate varies from predicate to predicate and may also be subject to pragmatic considerations; Dowty notes a parallel here to analogous issues in the semantics of generic statements, as discussed at length by Carlson 1977. But, to accept that 51 can be true even if some of the children are awake is to claim that the quantification here is near-universal rather than universal, with all the problems pointed out above for such a claim.¹⁶ One of my purposes here is to develop a solution to these problems, so we cannot accept, with Dowty, that sentences attributing a distributive predicate to a group can be 'not false or merely "approximately" true, but literally true' (103), even if not all members of the group satisfy the predicate.

It is perhaps worth noting that this last problem is not intrinsic to the idea of analyzing *all* in terms of subentailments; rather it is a problem in the analysis of distributive sentences without *all*. If one wanted, one could easily claim that *all* simultaneously tightens the pragmatic halo of a noun phrase and distributes the subentailments of the predicate to all members of the relevant group; there is no conflict between these ideas. In this case, *be asleep* would hold true of a group only if it holds of all members of the group; the appeal to subentailments would be vacuous in this example, and the maximizing effect would be due to the halo-tightening effect of *all*. For an example like *all the students voted for the bill*, one might claim that a group can vote for the bill even if only some of its members do; adding *all* would enforce this requirement on every member of the group, as in Dowty's analysis. Simultaneously, the halo for *the students* would tighten, assuring that the group of which every member is required to vote for the bill really is the group of all students, and not just some rough approximation of this group.

To keep the presentation simple, and because of the other problems with Dowty's analysis, I will not formalize a dual halos-and-subentailments analysis here. This leaves examples like the *vote for* case unaccounted for. Further research is needed to determine

¹⁶ A referee suggests that we might claim instead that, semantically speaking, 51 merely says that the predicate *be asleep* holds of the children as a group, relegating any distributive effects entirely to the pragmatics. In this case, the sentence would not involve any quantification at all, even in the lexical semantics of the predicate. This option is not open if we wish to maintain Dowty's analysis, however; if 51 does not actually entail that any individual children are asleep, then *be asleep* has no subentailments, and 52 should be anomalous, with a status similar to 47. Since it is not anomalous, we cannot treat the distributive effects here as purely pragmatic.

whether such examples should be explained using Dowty's subentailment technique or through some other mechanism.

5.5. RELEVANCE THEORY AND RESEMBLANCE AMONG PROPOSITIONS. Sperber and Wilson (1986) offer an analysis of 'loose talk' which, like that presented here, treats the looseness as essentially a pragmatic, non-truth-conditional phenomenon.¹⁷ Sperber and Wilson suggest that natural language discourse is not subject to a convention of literal truthfulness (such as Grice's maxim of quality), as is often assumed. Rather, there is a general expectation of 'an interpretive resemblance between the proposition expressed by the utterance and the thought that the speaker intends to convey' (170). Fully precise, literally truthful interpretations are just the limiting case of such resemblance.

Propositions are held to resemble one another to the extent that they share implications. The relevant notion of implication (hence also the notion of resemblance) is a contextual one. The 'context' of an utterance is here understood to be the set of background premises used in understanding the utterance. A proposition P resembles another proposition Q in context C to the extent that $C \cup \{P\}$ shares implications with $C \cup \{Q\}$.¹⁸ For example: Let C_1 be *{If it is winter, then it is cold, If it is cold, we should stay at home}*. Relative to C_1 , propositions 53 and 54 resemble one another closely, since they share the implications *it is cold*, and *we should stay at home*.

(53) It is winter.

(54) It is freezing cold.

In contrast, let C_2 be *{If it is winter, there are no flowers in the garden, If it is freezing cold, we should heat the greenhouse}*. Relative to this context, 53 and 54 do not share these implications, and so resemble each other less.

Now, if two propositions sufficiently resemble one another in a given context, they will share a large number of implications, and hence will license many of the same inferences. Yet it is possible that one of the propositions may be true and the other one false. In addition, an utterance expressing one of these propositions may require more processing effort than the other. Sperber and Wilson suggest that considerations of processing effort may in fact lead a speaker to utter a sentence expressing a false proposition, provided it sufficiently resembles the proposition the speaker intends to convey.

For example, suppose Marie lives in a suburb of Paris, one block from the Paris city limits. In the context of a social party in London, she might answer a question about where she lives by saying 'I live in Paris'. This utterance would be literally false, but not misleading. Marie may intend only to convey that she spends most of her time in the Paris area, knows Paris, lives an urban life, or the like. These implications are shared by the proposition expressed by her utterance, despite its literal falsehood. Of course she could have been literally truthful and said 'I live near Paris', but Sperber and Wilson suggest that this sentence requires more processing effort, and hence might trigger inferences to the effect that there is some relevant respect in which her residence is unlike one in Paris, which would be misleading. In contrast, in the context of an electoral meeting for a Paris local election, the issue of whether Marie lives inside the city limits would probably be relevant, and so her response would be different.

Although Sperber and Wilson do not directly address slack regulation, their general account of loose talk is the closest to the one suggested here of any of the proposals

¹⁷ Thanks to an anonymous referee for pointing out this article to me.

¹⁸ I have made some adjustments here to Sperber and Wilson's notation, in the hopes of increasing clarity.

I am aware of. There are some differences, including minor ones which it would be a mistake to make too much of: I have been interpreting sentences as denoting sets of events, and associating each sentence in context with a halo of similar sets of events; Sperber and Wilson instead speak of sentences as expressing propositions, and of a resemblance relation among propositions in context. But the appeal to events was made mainly for convenience. The general technique of halos is in no way dependent on it, and in any case, one could model propositions in terms of events if desired. Similarly I appeal to Gricean mechanisms in calculating the halos of lexical expressions, while Sperber and Wilson naturally appeal to their own framework of relevance theory. Without intending to discount the significant differences between Gricean and relevance-theoretic approaches to pragmatics, let me suggest that the choice between them has relatively little to do with the main concerns of this article. Far more important is the central point, shared with Sperber and Wilson, that much of the looseness in our talk is due not to vague truth conditions, but to pragmatic tolerance of falsehood.

In fact, Sperber and Wilson go further in this than I do here. They treat ordinary vagueness, of the sort that gives rise to the sorites paradox, as also involving pragmatic tolerance of falsehood rather than authentically vague truth conditions. They suggest that a person is truly bald, for example, only if the person has no hairs at all; even a single hair is enough to render a sentence like *John is bald* false. (Of course if John does have a few hairs, one might still reasonably use the sentence, despite its falsehood.)

It is doubtful whether this kind of approach can offer a general solution to the sorites paradox. One need only consider the other classic example, *heap*: how many grains of sand do we need before we obtain something analogous to a man with absolutely no hairs? Perhaps there is some exact number and we simply have no means of discovering it, as suggested by Williamson (1994), but this seems quite counterintuitive.

I should point out another important difference between my approach and Sperber and Wilson's. For Sperber and Wilson, one may utter a false sentence provided it expresses a proposition that sufficiently resembles the one one wants to convey. Resemblance is defined in terms of implication—and of course implication is a relation among propositions. Hence one cannot calculate resemblance until one has reached the level of propositions.

In contrast, in the system I advocate here, even individual lexical items are associated with pragmatic halos. Halos of course are sets of items resembling the denotation of the expression in question, so resemblance is calculated even at the level of individual words and phrases, and not just whole propositions. Indeed, halos for complex expressions are derived from the halos of their parts, so we may regard the resemblance relation as playing a role in the compositional semantics.

Do we need to build reference to pragmatic resemblance into the compositional semantic rules for a language, or can we keep it to the level of post-compositional pragmatics, as Sperber and Wilson seem to assume? I suggest that slack regulators such as *all*, *exactly* and *perfectly* provide evidence that we do need a compositional procedure for calculating the required degree of resemblance. Indeed, the lexical semantics of these words gives specific information on this matter.

Advocates of a purely pragmatic, post-compositional analysis might try to account for the halo-tightening effect of slack regulators by appealing to the fact that including them in a sentence involves some effort on the part of both speaker and listener. This effort would only be justified if there was some relevance to the fact that the word was included; hence a sentence containing one of these words must be interpreted differently than a corresponding sentence without it.

But while this kind of explanation successfully accounts for the fact that slack regulators have SOME kind of effect on interpretation, there is nothing in it which predicts what that effect will be. Why is the effect a tightening of slack, rather than, say, a loosening, or some completely different change? An analysis in which the effect is stipulated lexically at least allows us to say which specific effects are produced by which specific words.

6. FURTHER ISSUES. Several issues arise in the kind of analysis suggested here. I will not attempt to resolve all these issues here, and the discussion of some of them will be tentative and brief.

6.1. WILL THERE BE ANY TRUE SENTENCES LEFT AT ALL? Yes, of course there will. But I raise the question because some readers will undoubtedly worry that the analysis sketched above sets the standard of truth impossibly high. This reaction may be strongest for the example of Mary's arrival. Must we really say that 55 is false if she arrived just a few milliseconds after three? How about a few microseconds?

(55) Mary arrived at three o'clock.

Once we reach this scale we are sure to encounter difficulties: Just what instant is the instant of her arrival? The moment the tip of her nose crosses the threshold? The moment the outermost atom of the tip of her nose crosses the threshold? We will have to start worrying about quantum indeterminacy before long—but surely this cannot be a serious concern in the analysis of sentences about macroscopic objects like Mary!

Although we often speak of events like Mary's arrival as though they occurred at some particular instant, they really extend over a longer interval, and even this may not have well-defined endpoints. But if we speak of events as occurring at an instant when really they don't, then what we say is actually false. Sentences like 55 are never true in the real world. Mary's arrival could never take place at one single instant, as the sentence suggests; to speak this way is to idealize. But idealization is normal, and we treat sentences like 55 as close enough to the truth for practical purposes. This effect could be obtained by letting the pragmatic halo of *three o'clock* contain not just instants, but also longer intervals, which include 3:00 but extend beyond it slightly in either direction.

Even this view does not condemn us always to speak falsely, however. We need not go so far as Unger (1975), who, starting from the position that nothing is truly flat, eventually concludes that nothing is authentically true. If desired, the speaker can always preserve truth by making explicit allowances in the sentence, as in 56, for example.

(56) Mary arrived between 2:59 and 3:01.

Of course it would be quite burdensome to speak this way, and many people react to sentences like 56 as being pedantically overprecise—even though this sentence actually makes a LESS precise claim than 55, which sounds perfectly natural and not overprecise at all. How can this be? In the framework suggested here, this is explained by the way the pragmatic situation fixes halos: sentence 56 can be used felicitously only in a context where the distinction between times just a minute apart is relevant, hence its halo is relatively tight. In contrast, 55 can be used even in contexts where the distinction between such times may be pragmatically ignored; its halo may be considerably looser. The theory provides a way to distinguish between the level of precision encoded in the truth-conditional content of a sentence, and the level of precision assumed pragmatically in uttering the sentence, which may be something quite different.

6.2. HEDGES. A natural application of the formalism of pragmatic halos is in the analysis of expressions such as *roughly* or *loosely speaking* known since Lakoff 1973 as HEDGES. It should be stressed, however, that such expressions have a clear effect on the truth-conditional denotation of the expressions they combine with; they should not be analyzed merely as manipulating pragmatic slack. Consider sentence 57, for example.

(57) Loosely speaking, John is king.

This sentence might be used to mean that John functions as king in all practical respects, even though he doesn't technically have the title king. (He might officially be president-for-life, for example.)

The fact that John is not technically king does not render 57 false. On the contrary, I'm not sure it can really be true if John is king. Therefore we do not want to claim that *loosely speaking* is merely the signal of an extra-wide pragmatic halo.

I think it is more reasonable to suppose that *loosely speaking* serves to replace a sentence's denotation with a set whose members are drawn from the elements of the sentence's halo, excluding the denotation. (In the notation developed in the appendix, $\llbracket \textit{loosely speaking } \Phi \rrbracket^{M.C.} = \cup H_C(\Phi) - \llbracket \Phi \rrbracket^{M.C.}$.) In 57, for example, we might say that *king* denotes the set of kings, and has as its halo a set of sets, some of which include presidents-for-life, etc.; the sentence *John is king* will denote the set of eventualities of John being king, but its halo will include sets of eventualities some of which may be eventualities of John being president-for-life, or something else suitably kinglike. Then 57 will denote the set of these eventualities of John holding some kinglike title without actually being king.

If this approach to hedges is correct, they must be seen not as expanding the halo of the expression they modify, but instead as expanding the denotation into the halo. From one perspective, this may seem surprising; if expressions like *exactly* contract halos, isn't it natural to expect expressions like *loosely speaking* to expand them?

Not necessarily. What is common to slack regulators like *exactly* and to hedges like *loosely speaking* is that they both force the denotation and the halo to coincide more closely. *Exactly* does this by shrinking the halo so that it more closely matches the denotation; *loosely speaking* does it by expanding the denotation so that it more closely matches the halo, in effect converting pragmatic slack into semantic content. From this perspective, the two cases may be seen as symmetric. It is unclear to me exactly why these two possibilities appear to be the ones allowed; perhaps further research will illuminate this issue.

6.3. GRADING AND FUZZINESS OF HALOS. We have modeled pragmatic halos as structured sets—structured in the sense that they have a distinguished central member (the denotation), and in that the other members may be ordered with respect to their degree of closeness to this central member. Despite this structure, halos are SETS, in which even the outlying elements have full membership; our formalism does not provide for a gradual, continuous 'fading' from members to nonmembers. Is this right?

Certainly, even with respect to a given context, some distinctions may be more difficult to ignore pragmatically than others, and there may be no clear boundary between those distinctions that are ignorable and those that are not. Should we, then, treat halos as fuzzy sets, assigning each element of the halo a degree of membership between 0 and 1? In effect, this would treat the logic of the predicate *close enough to true for the context* as a version of fuzzy logic, even while allowing us to claim that the logic of the simple predicate *true* is not. Such a move brings with it most of the same advantages and disadvantages that attend fuzzy logic generally. I will not attempt a

general review of these here but refer the reader to Chierchia & McConnell-Ginet 1990: 391–95 for some discussion.

Even aside from general concerns about fuzzy logic, however, I think there is not really much to be gained by treating halos as fuzzy sets. First, note that halos already encode information about the relative ease or difficulty of ignoring various distinctions, in terms of the notion of relative closeness to the denotation. And although there is no clear boundary between ignorable and nonignorable distinctions, in actually choosing whether or not to ignore a particular distinction, there is no middle ground. If Mary arrived at 3:02, we may either ignore the distinction between 3:00 and 3:02 and say she arrived at 3:00, or we can maintain this distinction and say she arrived at 3:02; but there does not seem to be any way to halfway ignore the distinction and halfway maintain it. We must simply decide whether to include 3:02 in the halo for 3:00 or not. For this reason, there seems to be no harm in treating halos as ordinary, nonfuzzy sets, with halo membership reflecting an actual decision in context about which distinctions to ignore and which to maintain.

6.4. ATTENUATION OF HALOS IN THE SEMANTIC COMPOSITION. Our default rule for calculating halos has a kind of cumulative effect. If a complex expression consists of several constituents, each with its own wide halo (and no slack regulators), all those halos will ‘add up’ to produce an extra-wide halo for the complex expression.

For example, suppose in some context *C* the halo of *three o’clock* includes all times that are within fifteen minutes of 3:00. Now suppose that in this same context *C* the halo for *home* includes, say, all locations on the same city block as a given individual’s home. In this situation, if Mary arrives at 3:14, the sentence *Mary arrives at three* will be close enough to true for *C*; and if Mary is on her block, the sentence *Mary is home* will be close enough to true for *C*. But suppose she arrives on her block at 3:14, but does not actually enter her house until much later. Is the sentence *Mary arrives home at 3:00* close enough to true for *C*? As it stands, the analysis predicts yes, but perhaps we should consider the possibility that distinctions that are ignorable on their own may turn out to be too much to ignore when combined.¹⁹

My own intuitions about this example are not completely clear, but tend to support the view that examples like this should be considered close enough to true for the context described. If we want to account for the opposite intuition, the rules should be modified to include an ‘attenuation’ effect. I will not give the reformulation here, but the idea would be to eliminate outlying elements from the halos as composition proceeds up the tree. Since we may be dealing with several dimensions of closeness at once, a number of interesting and difficult issues arise as to what counts as outlying and how much of the halo to eliminate, but I must leave these questions to further research.

6.5. THE EFFECT OF LINGUISTIC CONTEXT. I have been assuming that halos are assigned to basic expressions by the pragmatic context. The rules assign the halo of a complex expression in a given context based on the halos of its constituent expressions IN THAT SAME CONTEXT. That is, there is no provision for context change within the sentence. It seems natural to expect, however, that one part of a sentence may contribute to the context of other parts, and that the appearance of certain expressions in a sentence might affect the initial assignment of halos to other basic expressions.

Suppose for example, that John is notorious for his lack of punctuality. Any statement about when John will arrive must be taken with a big grain of salt. In contrast, Mary

¹⁹ Thanks to Susan Rothstein for raising this issue.

is known to be extremely punctual. With these facts as background knowledge in the context, how wide should the halo of a phrase like *at three o'clock* be? One suspects that a wider halo will be assigned to a sentence like *John will arrive at three o'clock* than to a sentence like *Mary will arrive at three o'clock*, even if these sentences are uttered in the same context.

It is not clear that this suspicion is correct. In the situation just described, one may be more skeptical of a claim that John will arrive at three o'clock, but skepticism per se is no reason to assign a wide halo. But if we actually treat *John arrived at three o'clock* as though it were true, even when he arrives so late that we would not treat *Mary arrived at three o'clock* as true in analogous circumstances, then the assignment of a halo to *three o'clock* must be sensitive to the choice of subject noun phrase, *John* or *Mary*. Put differently, the subject noun phrase must form part of the context in which *three o'clock* is uttered and therefore affects its halo assignment.

To deal with such cases, we must move to a dynamic semantics, in which expressions (even subsentential expressions) map from contexts to contexts. In such a theory the context is recalculated with each new expression, and effects of this type are expected.

Since Stalnaker 1979 semantic theories along these lines have been developed in some detail, especially in the analysis of anaphora (see, e.g. Heim 1982, Groenendijk & Stokhof 1990, 1991). The kinds of context change needed to deal with anaphora are relatively systematic and amenable to formalization. Unfortunately, it is much less clear that the changes necessary to deal with halo assignment can be similarly formalized. (No doubt we must appeal to some Gricean or similar mechanism here, just as we did in the more static treatment of context assumed up to this point.)²⁰

But even if we cannot give a rigorous theory of what specific effect an utterance of the name *John* for example, will have on the context, we can still give a formal theory which represents this name as having an effect, and which treats the halo of an expression as assigned by a context obtained in part through processing the preceding discourse. The basic technique is familiar: we assign expression denotations relative to pairs of contexts, rather than just one, with the first context regarded as the context in which the expression is uttered, and the second one as the context produced by the utterance. The input context of an expression must match the output of the preceding expression; mappings for complex expressions are obtained by composing the mappings of the constituent expressions.

There is already considerable evidence that natural language semantics makes use of this sort of dynamic mapping from context to context, so it is hardly surprising to find that the assignment of pragmatic halos may also involve dynamic interpretation. The task of formalizing a dynamic treatment of halos will have to await another occasion. My purpose here is simply to give an initial presentation and justification for the

²⁰ A referee worries that the difficulty of formalizing this sort of context change calls into question any theory that makes use of a compositional procedure in halo assignment, as opposed to a purely pragmatic account. This worry is misplaced. The problem is in formalizing the effect of linguistic context on the initial assignment of halos to basic expressions, not in the compositional procedure for determining the halos of complex expressions. This problem comes up in any theory that admits an effect of linguistic context on halo assignment, whether that theory views halos as assigned on a purely pragmatic basis or on a mixed pragmatic/compositional basis as suggested here.

There is of course reason for concern whenever a theory cannot be stated with enough precision to be formalized, but the problem here is no worse than in Gricean pragmatic theories generally. Despite recent advances, we are quite some way from being able to formalize all the effects that context, including linguistic context, may have on interpretation.

idea of halos, and to sketch a technique for representing them. This seems most effectively done without the formal complications a dynamic logic inevitably involves.

7. CONCLUSION. I have argued that speakers make use of a certain kind of pragmatically licensed deviation from the truth, here called *PRAGMATIC SLACK*; and that natural languages are actually structured to exploit this fact, containing expressions whose function is to signal how much pragmatic slack should be allowed in the interpretation of an utterance. I termed such expressions *SLACK REGULATORS*, and gave an analysis of them. This analysis employed the device of *PRAGMATIC HALOS*; sets, approximating the denotation of an expression, which are fixed in part pragmatically, by ‘ignoring’ irrelevant details and distinctions, and in part structurally, through the use of slack regulators.

It is, I think, not too controversial to claim that speakers frequently say things that are not, strictly speaking, true, and that many such utterances are neither lies nor mistakes, but a completely normal part of honest, error-free discourse. It should not be surprising, therefore, to find that languages contain expressions whose semantics makes direct appeal to this phenomenon. I have argued that *exactly*, *all*, and *perfectly*, and perhaps also ‘hedges’, are such expressions. No doubt there are many others, in English and in other languages.

APPENDIX: SOME FORMAL DETAILS.

This appendix fleshes out some of the formal details of the system for the composition of pragmatic halos discussed in §§3 and 4. A fully formalized fragment of English will not be given, but I will try to make the basic technique reasonably explicit.

I assume an algebraic theory of compositionality, essentially as in Montague (1970). Syntactic structures are generated from basic expressions by recursive application of syntactic operations. Each syntactic operation F_i has a corresponding semantic operation G_i (though I will often leave reference to these operations implicit in the actual statement of the rules). Where α is any expression, I write $\llbracket \alpha \rrbracket^{M,C}$ for the denotation of α relative to model M and context C . Denotations of complex expressions are assigned compositionally, so that for any $\alpha_1, \dots, \alpha_n$, F_i , G_i : $\llbracket F_i(\alpha_1, \dots, \alpha_n) \rrbracket^{M,C} = G_i(\llbracket \alpha_1 \rrbracket^{M,C}, \dots, \llbracket \alpha_n \rrbracket^{M,C})$.

Relative to a given context C , each basic expression α is assigned a partially ordered set $\langle H_C(\alpha), \leq_{\alpha,C} \rangle$, the halo of α . $H_C(\alpha)$ is understood to be a set of objects which differ from $\llbracket \alpha \rrbracket^{M,C}$ only in ways which are pragmatically ignorable in C ; $\leq_{\alpha,C}$ is an ordering of $H_C(\alpha)$ according to similarity to $\llbracket \alpha \rrbracket^{M,C}$. We require that $\llbracket \alpha \rrbracket^{M,C} \in H_C(\alpha)$. (That is, the denotation of an expression is always included in its halo.) In addition, all elements of $H_C(\alpha)$ must be of the same logical type as $\llbracket \alpha \rrbracket^{M,C}$. Furthermore, we require that $\llbracket \alpha \rrbracket^{M,C}$ be the unique element y such that for all $x \in H_C(\alpha)$, $y \leq_{\alpha,C} x$. (The denotation of an expression is the centerpoint of the halo.)

As a default, I assume that halos for complex expressions are derived by applying normal semantic operations to all possible combinations of objects drawn from the halos of the immediate constituent expressions:

$$(58) \quad H_C(F_i(\alpha_1, \dots, \alpha_n)) = \{G_i(x_1, \dots, x_n) \mid x_1 \in H_C(\alpha_1), \dots, x_n \in H_C(\alpha_n)\}$$

The ordering relation is preserved in the composition:

$$(59) \quad \text{If } y, z \in H_C(\alpha_k) \text{ and } y \leq_{\alpha,C} z, \text{ then } G_i(x_1, \dots, y, \dots, x_n) \leq_{\alpha,C} G_i(x_1, \dots, z, \dots, x_n)$$

The two simpler cases of slack regulators, *exactly* and *all* can be treated as denoting identity functions: $\llbracket \text{exactly} \rrbracket^{M,C} = f: f(t) = t$, for all times t in M , and $\llbracket \text{all} \rrbracket^{M,C} = f: f(g) = g$, for all groups of individuals g in M . Where α is any time-denoting phrase, we naturally let $\llbracket \text{exactly } \alpha \rrbracket^{M,C} = \llbracket \text{exactly} \rrbracket^{M,C}(\llbracket \alpha \rrbracket^{M,C})$, that is, $\llbracket \alpha \rrbracket^{M,C}$; likewise where α is a group-denoting term, $\llbracket \text{all } \alpha \rrbracket^{M,C} = \llbracket \text{all} \rrbracket^{M,C}(\llbracket \alpha \rrbracket^{M,C})$, that is, $\llbracket \alpha \rrbracket^{M,C}$.

Since the elements of the halo of an expression are required to be of the same type as the denotation, $H_C(\text{exactly})$ will be a set of functions on times (approximating the identity function), and $H_C(\text{all})$ will be a set of functions on groups of individuals (again approximating the identity function). Since *all* and *exactly* regulate slack, we give separate rules to calculate the halos of expressions containing them, rather than appealing to the default rule in 58.

$$(60) \quad H_C(\text{exactly } \alpha) = \{x \mid \exists y[\forall z y \leq_{\alpha,C} z \ \& \ \exists f[f \in H_C(\text{exactly}) \ \& \ f(y) = x]]\}$$

$$(61) \quad H_C(\text{all } \alpha) = \{x \mid \exists y[\forall z y \leq_{\alpha,C} z \ \& \ \exists f[f \in H_C(\text{all}) \ \& \ f(y) = x]]\}$$

As discussed in §4, the idea here is to recover the centerpoint of the halo of α , then apply the functions in

the halo of *all* or *exactly* to this centerpoint, and then gather the resulting values into a set.²¹ Applied to the halo of *three o'clock*, for example, this should give the set of times which are values of three o'clock under some function which pragmatically approximates the identity function.

The halos of *exactly* and *all* should be constrained to assure that the halo of a complex phrase containing one of these words is at least as tight as the halo of the constituent phrase with which *exactly* or *all* combines, but not so tight as to exclude the centerpoint.

- (62) For all times t, t' :
 - a. If $t \in H_C(\text{exactly } \alpha)$ then $t \in H_C(\alpha)$
 - b. If $t \in H_C(\text{exactly } \alpha)$ and $t' \leq_{\alpha, C} t$, then $t' \in H_C(\text{exactly } \alpha)$
 - c. If for all t'' : $t \leq_{\alpha, C} t''$, then $t \in H_C(\text{exactly } \alpha)$
- (63) For all groups g, g' :
 - a. If $g \in H_C(\text{all } \alpha)$ then $g \in H_C(\alpha)$
 - b. If $g \in H_C(\text{all } \alpha)$ and $g' \leq_{\alpha, C} g$, then $g' \in H_C(\text{all } \alpha)$
 - c. If for all g'' : $g \leq_{\alpha, C} g''$, then $g \in H_C(\text{all } \alpha)$

The case of *perfectly* is somewhat more complex. Since it has a truth conditional effect when combined with scalar adjectives, as discussed in §4.3, it cannot be treated as denoting an identity function.

The general technique of pragmatic halos is consistent with a wide variety of approaches to scalar adjectives; for concreteness, I treat scalarity here in terms of contextual variation in denotation, as already informally discussed in §4.3. I assume that contexts may be ordered according to strictness with respect to the denotation of a scalar adjective α . ' $C \sqsubseteq_{\alpha} C'$ ' means that C is less strict than C' with respect to the denotation of α . We require that if $C \sqsubseteq_{\alpha} C'$, then $\llbracket \alpha \rrbracket^{M, C'} \subseteq \llbracket \alpha \rrbracket^{M, C}$. (So, for example, the denotation of *round* relative to a strict context will be a subset of the denotation relative to a less strict context.) A context C is maximally strict with respect to the denotation of α , notated ' $\max_{\alpha}(C)$ ', iff for all C' , $C' \sqsubseteq_{\alpha} C$.

For any expression α , let $\hat{\alpha}$ denote that function which maps any context C onto $\llbracket \alpha \rrbracket^{M, C}$. Abusing the terminology of Kaplan 1989, I call this the *character* of α .

Now we can treat *perfectly* as involving implicit reference to maximally strict contexts, by letting it combine semantically with the character of its complement: $\llbracket \text{perfectly } \alpha \rrbracket^{M, C} = \llbracket \text{perfectly} \rrbracket^{M, C}(\llbracket \hat{\alpha} \rrbracket^{M, C})$. The following definition gives the intended effect.

- (64) $\llbracket \text{perfectly} \rrbracket^{M, C}$ = that function f such that for all functions g from contexts to sets of individuals:
 $f(g) = \{x \mid \exists C[\max_{\alpha}(C) \ \& \ x \in g(C)]\}$

Thus, an object will be in the denotation of *perfectly round*, for example, only if it is in the denotation of *round* relative to a maximally strict context.

Where α is a nonscalar adjective such as *spherical*, we may assume that α denotes a constant function.²² Note that in this case, $\llbracket \text{perfectly } \alpha \rrbracket^{M, C} = \llbracket \alpha \rrbracket^{M, C}$.

$H_C(\text{perfectly})$ will be a set of functions of the same type as $\llbracket \text{perfectly} \rrbracket^{M, C}$, that is, functions that take characters of sets as arguments, and yield sets as values. Just as $\llbracket \text{perfectly} \rrbracket^{M, C}$ takes $\llbracket \hat{\alpha} \rrbracket^{M, C}$ rather than $\llbracket \alpha \rrbracket^{M, C}$ as its argument, we expect that the halo of *perfectly* α should be calculated from the halos of *perfectly* and $\hat{\alpha}$, rather than the halos of *perfectly* and α .

If $\llbracket \hat{\alpha} \rrbracket^{M, C}$ is a function from contexts to sets of individuals, then $H_C(\hat{\alpha})$ should be a set of such functions, and $\leq_{\alpha, C}$ should be a relation among such functions of closeness or similarity to $\llbracket \hat{\alpha} \rrbracket^{M, C}$. More specifically, let us assume that $f \in H_C(\hat{\alpha})$ iff for all C' , there is some $X \in H_C(\alpha)$ such that $f(C') = X$. For example, a function will be in the halo of $\hat{\text{round}}$ iff it maps each context onto a set that is in the halo of *round* in that context.

Since *perfectly* is a slack regulator, phrases containing it are not subject to the default rule for halo composition in 58, but have their halos assigned specially, just as with *all* and *exactly*.

- (65) $H_C(\text{perfectly } \alpha) = \{X \mid \exists g[\forall h \ g \leq_{\alpha, C} h \ \& \ \exists f[f \in H_C(\text{perfectly}) \ \& \ f(g) = X]]\}$

Here again, we recover the centerpoint from the halo of α ; the centerpoint will be identical with the denotation

²¹ Equivalently, we might apply the functions in the halo of *exactly* or *all* directly to the DENOTATION of α . This would make a somewhat simpler rule, but I adhere to the present format in order to derive the halo of the complex expression from the halos of the parts.

²² Of course nonscalar adjectives may be polysemous, with the context disambiguating among the various readings, and in that sense show 'contextual variation in denotation'. We need not take this to mean that $\hat{\alpha}$ denotes a nonconstant function, however. In my notation, $\hat{\alpha}$ must be understood as already having been disambiguated. Scalar adjectives do, as nonscalar adjectives do not, show contextual variation in denotation over and above the resolution of polysemy; and it is this variation that 'characters' are intended to model.

of $\hat{\alpha}$. I apply the functions in the halo of *perfectly* to this centerpoint, and gather the results into a set, to serve as the halo of *perfectly* α .

Of course, these functions need to be constrained in order to obtain a reasonable halo, just as with *all* and *exactly*, though the reference to characters and quantification over contexts makes the case of *perfectly* slightly more complicated.

- (66) a. If $X \in H_C(\textit{perfectly } \alpha)$ then there exists some $f \in H_C(\hat{\alpha})$ such that $\exists C[\max_\alpha(C) \ \& \ f(C) = X]$
- b. If $X \in H_C(\textit{perfectly } \alpha)$, and $\exists C[\max_\alpha(C) \ \& \ f(C) = X \ \& \ f'(C) = X']$, and $f' \leq_{\alpha, C} f$, then $X' \in H_C(\textit{perfectly } \alpha)$
- c. If for all f' : $f \leq_{\alpha, C} f'$, and $\exists C[\max_\alpha(C) \ \& \ f(C) = X]$, then $X \in H_C(\textit{perfectly } \alpha)$

The first clause assures that any set in the halo of *perfectly round*, for example, will be in the halo of *round* for some maximally strict context. The second clause assures that if one function more closely approximates that character of *round* than another, then the halo of *perfectly round* will not contain sets assigned by the less closely approximating function unless it also contains the sets assigned by the more closely approximating function. The third clause recovers the centerpoint of the halo of *round*—this should be identical with the denotation of *round*—and stipulates that the halo of *perfectly round* will contain the set this function yields in a maximally strict context, that is, the set of perfectly round objects.

If desired, the rules and conditions for *all* and *exactly* could easily be reformulated to match the format given for *perfectly*; for these words, the quantification over contexts is trivial. This would give a unified analysis, but at the cost of a more complicated treatment of *all* and *exactly*. I am not convinced that this sort of generalization to the worst case should be methodologically preferred, and so forgo the reformulation here.

REFERENCES

- BOLINGER, DWIGHT. 1972. Degree words. The Hague: Mouton.
- CARLSON, GREG. 1977. Reference to kinds in English. Amherst, MA: University of Massachusetts dissertation.
- CHIERCHIA, GENARO, and SALLY McCONNELL-GINET. 1990. Grammar and meaning. Cambridge, MA: MIT Press.
- DAVIDSON, DONALD. 1967. The logical form of action sentences. The logic of decision and action, ed. by Nicholas Rescher, 81–120. Pittsburgh: Pittsburgh University Press.
- DOWTY, DAVID. 1986. Collective predicates, distributive predicates, and *all*. Eastern States Conference on Linguistics 3.97–115.
- FINE, KIT. 1975. Vagueness, truth and logic. *Synthese* 30.265–300.
- GEURTS, BART. 1998. The mechanisms of denial. *Language* 74.274–307.
- GRICE, H. P. 1975. Logic and conversation. *Syntax and semantics*, vol. 3: Speech acts, ed. by Peter Cole and Jerry Morgan, 41–58. New York: Academic Press.
- GROENENDIJK, JEROEN, and MARTIN STOKHOF. 1990. Dynamic Montague grammar. Papers from the second symposium on logic and language, ed. by L. Kalman and L. Polos, 3–48. Budapest: Akadémiai Kiadó.
- , ———. 1991. Dynamic predicate logic. *Linguistics and Philosophy* 14.39–100.
- HAMBLIN, C. L. 1973. Questions in Montague English. *Foundations of Language* 10.41–53.
- HEIM, IRENE. 1982. The semantics of definite and indefinite noun phrases. Amherst: University of Massachusetts dissertation.
- HORN, LAURENCE R. 1992. The said and the unsaid. *Semantics and Linguistic Theory* 2.163–92.
- KADMON, NIRIT, and FRED LANDMAN. 1993. *Any*. *Linguistics and Philosophy* 16.353–422.
- KAMP, HANS. 1975. Two theories about adjectives. *Formal semantics of natural language*, ed. by Edward Keenan, 123–55. Cambridge: Cambridge University Press.
- KAPLAN, DAVID. 1989. Demonstratives. *Themes from Kaplan*, ed. by Joseph Almog, John Perry, and Howard K. Wettstein, 481–563. Oxford: Oxford University Press.
- KRIPKE, SAUL. 1977. Speaker's reference and semantic reference. *Contemporary perspectives in the philosophy of language*, ed. by Peter A. French, Theodore E. Uehling, Jr., and Howard K. Wettstein. Minneapolis: University of Minnesota Press.
- KROCH, ANTHONY. 1974. The semantics of scope in English. Cambridge, MA: MIT dissertation.
- LAKOFF, GEORGE. 1973. Hedges. *Journal of Philosophical Logic* 2.458–508.
- LASERSOHN, PETER. 1995. Plurality, conjunction and events. Dordrecht: Kluwer.

- LEWIS, DAVID. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8.339–59.
- LINK, GODEHARD. 1983. The logical analysis of plurals and mass terms: A lattice-theoretical approach. *Meaning, use and interpretation of language*, ed. by Rainer Bäuerle, Christoph Schwarze, and Arnim von Stechow, 302–23. Berlin: De Gruyter.
- . 1987a. Generalized quantifiers and plurals. *Generalized quantifiers: Logical and linguistic approaches*, ed. by Peter Gärdenfors, 151–80. Dordrecht: Kluwer.
- . 1987b. Algebraic semantics for event structures. *Proceedings of the eleventh Amsterdam colloquium*, ed. by Jeroen Groenendijk, Martin Stokhof, and Frank Veltman, 243–62. Amsterdam: ILLI.
- MONTAGUE, RICHARD. 1970. Universal grammar. *Theoria* 36.373–98.
- PARTEE, BARBARA. 1973. Some structural analogies between tenses and pronouns in English. *Journal of Philosophy* 70.601–10.
- ROBERTS, CRAIGE. 1987. *Modal subordination, anaphora and distributivity*. Amherst, MA: University of Massachusetts dissertation.
- ROOTH, MATS. 1985. *Association with focus*. Amherst, MA: University of Massachusetts dissertation.
- SCHA, REMKO. 1981. Distributive, collective and cumulative quantification. *Formal methods in the study of language*, ed. by Jeroen Groenendijk, Theo Janssen, and Martin Stokhof, 483–512. Amsterdam: Mathematical Center.
- SPERBER, DAN, and DEIRDRE WILSON. 1986. Loose talk. *Proceedings of the Aristotelian Society* 87.153–71.
- STALNAKER, ROBERT. 1979. Assertion. *Syntax and semantics*, vol. 9: *Pragmatics*, ed. by Peter Cole, 315–32. New York: Academic Press.
- TAUB, ALISON. 1989. Collective predicates, aktionsarten and *all*. *Papers on quantification*, ed. by Emmon Bach, Angelika Kratzer, and Barbara Partee, 337–68. Amherst, MA: Graduate Linguistic Student Association.
- UNGER, PETER. 1975. *Ignorance*. Oxford: Oxford University Press.
- VAN FRAASSEN, BAS. 1966. Singular terms, truth-value gaps, and free logics. *Journal of Philosophy* 63.481–95.
- VON FINTEL, KAI. 1994. *Restrictions on quantifier domains*. Amherst, MA: University of Massachusetts dissertation.
- WESTERSTÄHL, DAG. 1985. Determiners and context sets. *Generalized quantifiers in natural language*, ed. by Johan van Benthem and Alice ter Meulen, 45–71. Dordrecht: Foris.
- WILLIAMSON, TIMOTHY. 1994. *Vagueness*. London: Routledge.

Department of Linguistics
University of Illinois
Urbana, IL 61801
[laser@cogsci.uiuc.edu]

[Received 13 June 1996;
revision received 10 April 1998;
accepted 19 October 1998]