

# A Metasemantic Account of Vagueness<sup>1</sup>

Agustín Rayo

## 0 The plan

In this paper I will defend what might be called a *metasemantic* account of vagueness.<sup>2</sup> The proposal divides naturally into three parts. The first will be developed in section 1, and is based on the following three theses:

### CONVENTIONALITY

The meaningfulness of our sentences depends on the prevalence of suitable linguistic conventions.<sup>3</sup>

### DEGREE

Whether or not a convention prevails in a given population needn't be an all-or-nothing matter; it can come in degrees.

### BORDERLINE CASES

Whereas conventions about how to use sentences involving only attributions of

---

<sup>1</sup>I wrote this paper in 2003, and have since have grown dissatisfied with its content. I am unhappy with a number of issues, but the main misgiving is this: conventionality-based conceptions of linguistic representation are not very plausible as accounts of linguistic competence (Laurence (1996)), and I have come to believe that an adequate account of vagueness must be closely linked to an explanation of what our ability to use language consists in. At the same time, I remain convinced of the basic insight of the Metasemantic Account: that the root of vagueness lies not in the type of semantic-value that is best associated with an expression, but in the type of linguistic practice that renders the expression meaningful. I have decided to publish this paper in spite of its many deficiencies because it seems to me that the Metasemantic Account is underrepresented in the literature, and that the paper presents a version of the account that is both intuitive and detailed enough to allow for serious assessment. I have deleted a section on mental content in order to satisfy a request from the editors for a shorter paper, but refrained from making any other modifications to the original version of the paper. (For my current position on vagueness, see Rayo (2008).)

<sup>2</sup>Related proposals are set forth in Black (1949*b*) and Dorr (2003). For criticism of Black, see Williamson (1994) §3.2.

<sup>3</sup>In claiming that language is conventional I do not wish to claim, for example, that speakers generally mean by their words whatever they intend to mean. CONVENTIONALITY amounts to little more than the truism that words lack 'natural' meanings.




Figure 1: Prevalence of conventions about how to use the  $M_i$  ( $1 \leq i \leq 100,000$ ), where  $M_n$  is the sentence ‘An event involving exactly  $n$  persons is multitudinous’.

vague predicates to clear cases are often prevalent to a very high degree, conventions about how to use sentences involving attributions of vague predicates to borderline cases are prevalent, if at all, to lesser degrees. (See figure 1.)

These theses leave open the question of what a *semantics* for vague terms should look like. The second part of the metasemantic account addresses this issue, and will be the focus of section 2. Once the semantic proposal is in place, a number of issues in pragmatics will need to be taken up. Doing so constitutes the third part of the metasemantic account, and will be the task of section 3.

## 1 Metasemantics

### 1.1 Convention

What is it for a convention to prevail amongst members of a given population? The metasemantic account of vagueness is to a large extent independent of the details of one’s answer to this question, but it will be useful to fix ideas by considering the account of convention set forth in Lewis (1969) and Lewis (1973).<sup>4</sup>

---

<sup>4</sup>Lewis’s work on convention is criticized in Schiffer (1972), Burge (1975), Jamieson (1975), Grandy (1977), Schiffer (1987) and Gilbert (1989). Rival account of convention—on the basis of which the metasemantic account of convention could also be defended—are set forth in Schiffer (1972) and Gilbert (1989).

According to Lewis, a convention is a regularity of a certain kind. Specifically, a regularity  $R$  is a convention in a population  $P$  just in case the following conditions obtain within  $P$ :

1. There is general (or almost general) conformance to  $R$ .
2. Everyone (or almost everyone) believes that there is general (or almost general) conformance to  $R$ .
3. The belief that there is general (or almost general) conformance to  $R$  gives everyone (or almost everyone) good and decisive reason to conform to  $R$  herself.
4. There is a general (or almost general) preference for general conformity to  $R$  rather than slightly-less-than-general conformity.
5.  $R$  is not the only possible regularity meeting conditions 3 and 4.
6. It is a matter of common knowledge (or almost a matter of common knowledge) that conditions 1–5 obtain.

Consider an example. The regularity of driving on the left-hand side of the road counts as a convention amongst drivers on the British Isles. This is because they meet the following conditions: (1) almost everyone drives on the left-hand-side of the road most of the time; (2) almost everyone believes that almost everyone drives on the left-hand side of the road most of the time; (3) the belief that almost everyone drives on the left-hand side of the road most of the time gives almost everyone good and decisive reason to drive herself on the left-hand-side of the road (since driving on the right-hand-side of the road is extremely dangerous when others are driving on the left-hand-side, and almost everyone is adverse to dangerous driving); (4) there is a near general preference for everyone's driving on the left-hand-side of the road rather than almost everyone's doing so (since even a slight non-conformance to the regularity of driving on the left-hand-side of the road makes driving

in Britain much more dangerous, and almost everyone is adverse to dangerous driving); (5) driving on the left-hand-side of the road is not the only regularity satisfying (3) and (4) (since the regularity of driving on the right-hand-side of the road would have done just as well); and (6) it is a matter of common knowledge (or almost a matter of common knowledge) that (1)–(5) obtain.

## 1.2 Convention as a matter of degree

On one of the hills surrounding the football stadium there is a small village. Whether a great many people attend the daily football match is important to villagers. It is important to village drivers because traffic becomes unbearable. It is important to village bar-owners because, after well-attended matches, football fans are drawn in great numbers to the village for drinks. It is important to villagers normally uninterested in football because they might be interested in running down to the stadium to share the excitement. It is important to villagers who like football but dislike crowds because they might consider watching the match on television rather than going to the stadium. And so forth.

As a result, an interesting practice has taken hold in the village. A single villager—a different one each day—phones the stadium box-office early in the morning. After learning how many of the available 10,000 tickets have been sold for the day’s match, he or she decides whether to walk to the center of town and raise the village flag.

The practice emerged gradually and informally, and no explicit rule about when to raise the flag was ever set forth. Nonetheless, there is a clear correlation between flag-raising and box-office sales. When box-office sales are high, the flag is nearly always raised; when box-office sales are low, the flag is nearly never raised; when box-office sales are middling, the flag is sometimes raised and sometimes not. More specifically, the correlation is as represented by figure 2.

Say that box-office sales are ‘high’ if between 9,000 and 10,000 tickets were sold. Then




Figure 2: Correlation between flag-raising and box-office sales.

regularity of raising the flag when box-office sales are high constitutes a convention amongst villagers, in Lewis’s sense. For: (1) almost every villager raises the flag almost every time box-office sales are high (villagers are sometimes absent minded, and sometimes engage in pranks); (2) almost every villager believes that almost every villager raises the flag almost every time box-office sales are high; (3) the belief that almost every villager raises the flag almost every time box-office sales are high gives nearly every villager good and decisive reason to herself raise the flag when box-office sales are high (unless they are engaged in pranks, villagers are eager to avoid misleading others); (4) there is a near general preference for every villager’s raising the flag when box-office sales are high rather than nearly everyone’s doing so (since non-conformance to the regularity is likely to mislead villagers); (5) raising the flag when box-office sales are high is not the only regularity satisfying (3) and (4) (ringing the village bell would have done just as well); and (6) it is a matter of common knowledge that (1)–(5) obtain. (Similarly, the regularity of *refraining* from raising the flag when box-office sales are low constitutes a convention amongst villagers.)

Now consider the following question: Is it a convention amongst villagers to raise the flag when box-office sales are between 8,000 and 9,000? What about sales between 7,000 and 8,000, or between 6,000 and 7,000, or between 5,000 and 6,000? As we consider lower and lower ranges of box-office sales, there is less and less of a flag-raising regularity in

place, and less of a widespread belief in the presence of a flag-raising regularity. When does it stop being a convention to raise the flag?

It is hard to give a principled answer to these questions. For, on the definition above, ‘convention’ inherits the vagueness of occurrences of ‘almost’ in clauses 1–6. Fortunately, as Lewis himself observes, it is possible substitute a parameter for each (extensional) occurrence of ‘almost’ in his definition:

- 1'. Conformance to  $R$  takes place  $x_1$  percent of the time.
- 2'.  $x_2$  percent of the population believes that there is general (or almost general) conformance to  $R$ .
- 3'. The belief that there is general (or almost general) conformity to  $R$  gives  $x_3$  percent of the population good and decisive reason to conform to  $R$  herself.
- 4'.  $x_4$  percent of the population prefers general conformity to  $R$  rather than slightly-less-than-general conformity.
- 5'.  $R$  is not the only possible regularity meeting conditions 3' and 4'.
- 6'. It is a matter of common knowledge (to degree  $x_6$ ) that conditions 1'–5' obtain.<sup>5</sup>

One could also speak of degrees of belief and preference, but I will avoid doing so here for the sake of brevity.

We can now say that a regularity  $R$  is a convention to degree  $\langle x_1, x_2, x_3, x_4, x_6 \rangle$  just in case conditions 1'–6' obtain within  $P$ . This puts us in a position to give principled answers to questions like ‘To what degree is there a convention amongst villagers to raise the flag when box-office sales are between 6,000 and 7,000?’. For, on the assumption that we have enough information about the village (and, in particular, enough information about the

---

<sup>5</sup>There are different ways in which one could assign a degree of common knowledge to a population, but the details matter little for present purposes.

beliefs and desires of villagers), we will be in a position to say that a convention prevails to degree  $\vec{x}$ , for some particular vector  $\vec{x}$ .

Of course, establishing that a convention prevails to degree  $\vec{x}$  won't always answer the question of whether the convention prevails *simpliciter*, in the original sense of 'convention'. But it is important to note the the two senses of 'convention' are not in competition. Each of them can fruitfully be used to address a different range of issues. Here we will be concerned with the notion of conventionality to a degree.

It should now be clear why I think that the second of the three theses I introduced in section 0 is true:

#### DEGREE

Whether or not a convention prevails in a given population needn't be an all-or-nothing matter; it can come in degrees.

In the case of the village and the stadium, for example, raising the flag when box-office sales are between 9,000 and 10,000 is a convention to a very high degree, raising the flag when box-office sales are between 6,000 and 7,000 is a convention to a lesser degree, and raising the flag when box-office sales are between 0 and 1,000 is a convention to a very low degree.<sup>6</sup>

### 1.3 Meaning-fixation

Let us now turn to the first of the three theses introduced in section 0:

#### CONVENTIONALITY

The meaningfulness of our sentences depends on the prevalence of suitable linguistic conventions.

---

<sup>6</sup>Since degrees of conventionality are  $n$ -tuples of real numbers rather than real numbers, something needs to be said about what makes one degree greater than another. The details matter little for present purposes, but to fix ideas one might say that  $\langle x_1, \dots, x_n \rangle \leq \langle y_1, \dots, y_n \rangle$  iff  $\sqrt{(\alpha_1 x_1)^2 + \dots + (\alpha_n x_n)^2} \leq \sqrt{(\alpha_1 y_1)^2 + \dots + (\alpha_n y_n)^2}$  (for suitable weightings  $\alpha_1, \dots, \alpha_n$ ).

What is a ‘suitable linguistic convention’? As before, the metasemantic account of vagueness is to a large extent independent of one’s answer to this question. But it will be useful to fix ideas by considering a slight modification of the account of meaning-fixation set forth in Lewis (1973). Assume, as Lewis does, that propositions are sets of possible worlds. The suggestion is that a sentence  $s$  is used by population  $P$  to *mean\** a proposition  $p$  just in case it is a convention in  $P$  to act (or believe) in accordance with the following principles:<sup>7</sup>

1. TRUTHFULNESS

Assent to an utterance of  $s$  only if you believe  $p$ ; dissent<sup>8</sup> from an utterance of  $s$  only if you believe  $\bar{p}$ .

2. TRUST

Acquire the belief that  $p$  when others assent to an utterance of  $s$ ; acquire the belief that  $\bar{p}$  when others dissent from an utterance of  $s$ .

(where  $\bar{p}$  is the set-theoretic complement of  $p$ ; I use ‘mean\*’ rather than ‘mean’ so it is clear that I do not wish to claim that a sentence  $s$  expresses a proposition  $p$ , as used by  $P$ , just in case  $s$  is used by  $P$  to mean\*  $p$ ).

Thus, ‘Snow is white’ is used to mean\* that snow is white if (a) there prevails a convention to assent to utterances of ‘Snow is white’ only if one believes that snow is white and dissent from utterances of ‘Snow is white’ only if one believes it is not the case that snow is white, and (b) to acquire the belief that snow is white when others assent to utterances of ‘Snow is white’ and acquire the belief that it is not the case that snow is white when others dissent from utterances of ‘Snow is white’. And, of course, one can say that  $s$  is used to

---

<sup>7</sup>I assume for the sake of simplicity that  $s$  is devoid of context-sensitivity and ambiguity.

<sup>8</sup>I take dissent into consideration in order to avoid the conclusion that, if  $s$  is used by  $P$  to mean\*  $p$ , then  $s$  is used by  $P$  to mean  $q$  whenever belief in  $p$  necessitates belief in  $q$ . A more complete characterization of meaning\* would presumably take into account speech-acts such as supposition in addition to assent and dissent.



degree  $\vec{x}$  by  $P$  to mean\*  $p$  just in case the relevant convention prevails in  $P$  to degree  $\vec{x}$ .<sup>9</sup>

The Lewisian account of meaning-fixation can be expected to deliver the following result: for a wide range of cases,  $s$  is used by English speakers to mean\*  $p$ , where  $p$  captures the truth-conditions which are intuitively associated with  $s$ .<sup>10</sup>

## 1.4 Vagueness

For each  $n$  between 1 and 100,000, let  $M_n$  be the sentence  $\lceil$ An event involving exactly  $n$  persons is multitudinous $\rceil$ , and assume that the space of possible worlds under consideration is restricted to worlds in which the actual linguistic conventions prevail. Then, to a very high degree,  $M_1$  is used amongst English speakers to mean\* the empty proposition.<sup>11</sup>

---

<sup>9</sup>It is worth emphasizing that the present account of meaning-fixation is compatible with semantic externalism. Let speakers on Earth and Twin-Earth display identical linguistic behavior. There will still be a difference in what ‘water is wet’ is used to mean\* on Earth and Twin-Earth, provided there are suitable differences in the contents of the beliefs of Earthlings and Twin-Earthlings. Similarly, there will be a difference between what sentences involving the term ‘mass’ are actually used to mean\* by English speakers, and what they would have been used to mean\* by English speakers had Newtonian physics been correct, provided there would have been suitable differences in the contents of our beliefs had Newtonian physics been correct.

Could there be meaning\* in a population consisting of a single member? There could, provided the speaker’s practice gives way to regularities over time which constitute appropriate conventions of truthfulness and trust.

<sup>10</sup>See, however, footnote 17.

<sup>11</sup>This is because the following conditions obtain:

1. *Truthfulness (assent)*

English speakers nearly always conform to the regularity of refraining from assenting to utterances of  $M_1$ . And their conformity is independent of what they believe, as long as they believe that actual linguistic conventions prevail. In addition, English speakers nearly always believe that there is general (or almost general) conformance to this regularity. [And so forth.]

2. *Truthfulness (dissent)*

It is trivially true that English speakers conform to the regularity of dissenting from utterances of  $M_1$  only if they believe the set-theoretic complement of the empty proposition. In addition, English speakers nearly always believe that there is general conformance to this regularity [And so forth.]

3. *Trust (assent)*

If someone assents to an utterance of  $M_1$ , it is to be expected that others will be led to revise their beliefs about whether the assent was serious, or about the utterer’s linguistic competence, or about their own linguistic competence. In other words, English speakers nearly always conform to the regularity of *opting out* when someone assents to an utterance of  $M_1$ : they opt out of the assumption that it is a case of genuine assent, or opt out of the assumption that the space of possible worlds under consideration is restricted to those in which the actual linguistic conventions




Figure 3: Degree to which English speakers use  $M_k$  to mean\* the empty proposition.

We get the same result for  $M_2$  and  $M_3$ : when attention is restricted to worlds in which the actual linguistic conventions prevail,  $M_2$  and  $M_3$  are, to a very high degree, used by English speakers to mean\* the empty proposition. As  $k$  gets larger, however,  $M_k$  is used amongst English speakers to mean\* the empty proposition to a lesser and lesser degree. For, e.g. the regularity to refrain from assenting to utterances of  $M_k$  prevails to a lesser and lesser degree, and the belief that there is general (or almost general) conformance to this regularity prevails to a lesser and lesser degree. By the time  $k$  is close to 100,000,  $M_k$  is used amongst English speakers to mean\* the empty proposition to an extremely low degree, if at all. The situation is represented by figure 3.

Similarly, when attention is restricted to worlds in which the actual linguistic conventions prevail,  $M_{100,000}$  is, to a very high degree, used amongst English speakers to mean\* the universal proposition. The same is true for  $M_{99,999}$  and  $M_{99,998}$ . As  $k$  gets smaller, however,  $M_k$  is used amongst English speakers to mean\* the universal proposition to a lesser and lesser degree. By the time  $k$  is close to 1,  $M_k$  is used amongst English speakers

---

prevail. In addition, English speakers nearly always believe that there is general (or almost general) conformance to this regularity [And so forth.]

#### 4. *Trust (dissent)*

It is trivially true that English speakers nearly always conform to the following regularity: if someone dissents from an utterance of  $M_1$ , then others will acquire the belief that the the set-theoretic complement of the empty proposition obtains (or, in this case, retain that belief). In addition, English speakers nearly always believe that there is general conformance to this regularity [And so forth.]




Figure 4: Degree to which English speakers use  $M_k$  to mean\* the universal proposition.

to mean\* the universal proposition to an extremely low degree, if at all. The situation is represented by figure 4.

It should now be clear why I think that the third of the three theses introduced in section 0 is true:

#### BORDERLINE CASES

Whereas conventions about how to use sentences involving only attributions of vague predicates to clear cases are often prevalent to a very high degree, conventions about how to use sentences involving attributions of vague predicates to borderline cases are prevalent, if at all, to lesser degrees.

For values of  $k$  such that events involving exactly  $k$  persons are clear cases of multitudinousness or non-multitudinousness,  $M_k$  is to a very high degree used by English speakers to mean\* a proposition. But for values of  $k$  such that events involving exactly  $k$  persons are borderline cases of multitudinousness,  $M_k$  is used to lesser degrees, if at all, by English speakers to mean\* a proposition. The situation is as represented by figure 5.<sup>12</sup>

This concludes my defense of the three theses introduced in section 0. The next task is to say something about what form a compositional semantics for vague terms might take.

<sup>12</sup>As before, I assume that attention is restricted to worlds in which the actual linguistic conventions prevail. I also assume uses amongst English speakers of one of the  $M_k$  to mean\* anything other than the empty proposition or the universal proposition are of low enough degree to be negligible.




Figure 5: Degree to which English speakers use  $M_k$  to mean\* a proposition.

## 2 Semantics

### 2.1 Degrees of truth?

One might be tempted to think that the preceding section leads to a degrees-of-truth account of vagueness. In particular, one might think that the ‘degree of truth’ of a sentence should be identified with the degree to which the sentence is used to mean\* a true proposition in the relevant linguistic community.

That would be a mistake. To begin with, the choice of terminology is unfortunate. For talk of a sentence’s ‘degree of truth’ fails to distinguish between the claim that it is a matter of degree whether a sentence is used to mean\* a given proposition (true or otherwise) and the claim that it is a matter of degree whether a proposition is true. The present proposal holds the former as a basic tenet, but has nothing to do the latter.

More importantly, it is no part of the present proposal that the degree to which a logically complex sentence is used to mean\* a true proposition is a function of the degrees to which its components are used to mean\* a true proposition. (Conspicuously, it is not the case that a sentence is used to mean\* a true proposition to degree  $\vec{x}$  just in case its negation is used to mean\* a true proposition to degree  $\vec{1} - \vec{x}$ .<sup>13</sup> Quite the contrary: it is crucial for BORDERLINE CASES that sentences involving attributions of vague predicates

---

<sup>13</sup>Let  $\vec{1}$  be the maximum degree,  $\vec{0}$  be the minimum degree, and  $\langle x_1, \dots, x_n \rangle - \langle y_1, \dots, y_n \rangle$  be  $\langle x_1 - y_1, \dots, x_n - y_n \rangle$ .

to borderline cases *not* have this feature.) But the view that the ‘degree of truth’ of a logically complex sentence is a function of the ‘degrees of truth’ of its components is a central assumption of standard degrees-of-truth accounts of vagueness.

Actually, the very idea of directly reading-off the truth-status of sentences from the degrees to which they are used to mean\* true propositions is hopeless. Among other things,<sup>14</sup> no sentence too complex to be parsed by humans can be expected to be used to mean\* a proposition to any significant degree. In order to properly assess the truth-status of such sentences, what one needs to do is ascribe a compositional semantics to the relevant linguistic community on the basis of information pertaining to sentences which are used to mean\* propositions to significant degrees.

## 2.2 The Absolute-Semantics Picture

Let me fix some terminology. To assign *truth-conditions* to a sentence is to specify what the world must be like in order for the sentence to be true.<sup>15</sup> A (compositional) *semantics* is a (compositional) assignment of truth-conditions to sentences. A *classical* semantics is a compositional semantics in the style of Lewis (1970).<sup>16</sup> If  $\mathcal{C}$  is a classical semantics,  $\mathcal{C}(s)$  is the set of possible worlds in which the truth conditions which  $\mathcal{C}$  assigns to  $s$  are satisfied;  $\mathcal{C}$  is *vindicated* in  $P$  with respect to  $s$  (to at least degree  $\vec{x}$ ) just in case  $s$  is used by  $P$  to mean\*  $\mathcal{C}(s)$  (to at least degree  $\vec{x}$ ).

According to the *Absolute-Semantics Picture*—as I shall call it—the ascription of a classical semantics  $\mathcal{C}$  to a given population is subject to the following constraint:

---

<sup>14</sup>See also footnote 17.

<sup>15</sup>I do not presuppose that speakers must somehow have cognitive access to such specification.

<sup>16</sup>For sentences of the form  $\lceil P(a) \rceil$  the idea is roughly as follows. The semantic value of the name  $\lceil a \rceil$  is a function taking each centered-world  $w$  to an individual in  $w$ ; the semantic value of the predicate  $\lceil P \rceil$  is a function taking each centered-world  $w$  to a set of individuals.  $\lceil P(a) \rceil$  is true relative to a centered-world  $w$  just in case the individual assigned to  $w$  by the semantic value of  $\lceil a \rceil$  is a member of the set assigned to  $w$  by the semantic value of  $\lceil P \rceil$ . An utterance  $u$  of  $\lceil P(a) \rceil$  is true *simpliciter* just in case  $\lceil P(a) \rceil$  is true relative to the actual world centered on  $u$ . Use of centered-worlds will be omitted in what follows for the sake of simplicity.

$\mathcal{C}$  should only be ascribed to  $P$  if it is vindicated by  $P$  to a sufficiently large degree with respect to an appropriate range of sentences.<sup>17</sup>

I will argue in the following section that, in spite of its initial plausibility, the Absolute-Semantics Picture must be given up.

## 2.3 A semantics for vague terms

For  $i$  a natural number, let the  $\mathcal{C}_i$  be a family of classical semantics. They all agree about the semantic lexicon from which sentences are to be built and about the particular semantic structure that is to be assigned to each sentence. They also agree about which semantic values should be assigned to each item in the semantic lexicon, with one important exception: for each  $n$ ,  $\mathcal{C}_n$  assigns to ‘multitudinous’ the function taking each world  $w$  to the set of events in  $w$  involving at least  $n$  persons. (Assume, in addition, that the meaning\* of the  $M_n$  in  $P$  is as in section 1.4, and that relevant lexical items are assigned ‘homophonic’

---

<sup>17</sup>I have ignored an important qualification. As Stalnaker (1979) has emphasized, the proposition communicated by an assertion is not always the proposition expressed by the asserted sentence. In general, the proposition communicated by an assertion of  $s$  is the *diagonal*—the set of (contextually relevant) worlds  $w$  such that ‘ $s$  is true’ is true-in- $w$ . One should therefore expect meaning\* to track diagonals, rather than tracking propositions expressed (as the notion of vindication defined above presupposes). One way of addressing the problem is by enriching the notion of meaning\*: instead of taking meaning\* to relate a sentence to a proposition, one could take it to relate a sentence to a function  $f$  from possible worlds to propositions. At a first approximation, meaning\* might then be characterized as follows: sentence  $s$  is used by population  $P$  to mean\* function  $f$  just in case (1) it is a convention in  $P$  to act (or believe) in accordance with the following two principles: (a) assent to an utterance of  $s$  only if you believe the proposition consisting of worlds  $w$  such that  $w \in f(w)$  (and correspondingly for dissent), and (b) acquire a belief in the proposition consisting of worlds  $w$  such that  $w \in f(w)$  when others assent to an utterance of  $s$  (and correspondingly for dissent); and (2) if  $s$  contains no counterfactual conditionals,  $c$  is a counterfactual conditional of the form  $\lceil s' \rightarrow s \rceil$  and  $s'$  is used to mean\*  $f'$ , then  $c$  is used to mean\*  $g$ , where  $g(w)$  is the set of worlds  $w'$  such that some world in the intersection of  $f'(w)$  and  $f(w)$  is ‘closer’ (in some appropriate sense) to  $w'$  than any world in the intersection of  $f'(w)$  and the complement of  $f(w)$ . (Here I assume that the language in question is rich enough to contain counterfactual conditionals, and that there is some satisfactory way of picking them out.) The notion of vindication can then be redefined as follows:  $\mathcal{C}$  is *vindicated* in  $P$  with respect to  $s$  (to at least degree  $\bar{x}$ ) just in case  $s$  is used by  $P$  to mean\*  $f$  (to at least degree  $\bar{x}$ ), where  $\mathcal{C}(s)$  is the result of applying  $f$  to the actual world.

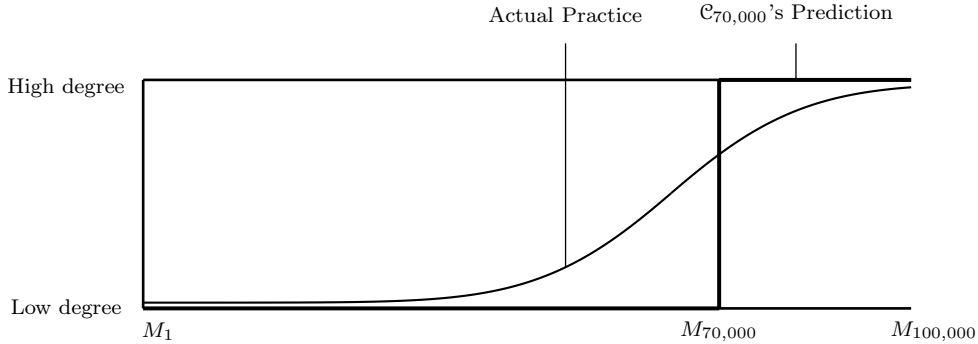


Figure 6: Degree to which  $M_k$  is actually used to mean\* the universal proposition in  $P$  vs. degree to which  $M_k$  should be used to mean\* the universal proposition in  $P$  according to  $\mathcal{C}_{70,000}$ .

semantic values by the  $\mathcal{C}_i$ .<sup>18</sup>)

As long as ascriptions of semantic theories are understood in accordance with the Absolute-Semantics Picture, no individual  $\mathcal{C}_i$  can provide us with a fully accurate description of  $P$ 's linguistic practice. Suppose, for example that  $\mathcal{C}_{70,000}$  is ascribed to  $P$ . Then, by the lights of the Absolute-Semantics Picture, one should expect  $M_0$  through  $M_{69,999}$  to be used by  $P$  to mean\* the universal proposition to a very low degree, and  $M_{70,000}$  through  $M_{100,000}$  to be used by  $P$  to mean\* the universal proposition to a very high degree; and this is only a rough approximation of the truth (see figure 6).

Ascribing a more complex semantics to  $P$  does not, by itself, constitute significant

---

<sup>18</sup>More specifically, we shall assume that the following is true for every  $i \in \mathbb{N}$ : (a)  $\mathcal{C}_i$  assigns to each of the  $M_n$  the semantic structure

$$\forall x[(\text{Event}(x) \wedge \exists!_n y(\text{Person}(y) \wedge \text{Involved}(x, y))) \rightarrow \text{Multitudinous}(x)]$$

(with standard definitional equivalences); (b) the semantic value of 'Event' according to  $\mathcal{C}_i$  is the function taking each world  $w$  to the set of events in  $w$ ; (c) the semantic value of 'Person' according to  $\mathcal{C}_i$  is the function taking each world  $w$  to the set of persons in  $w$ ; (d) the semantic value of 'Involved' according to  $\mathcal{C}_i$  is the function taking each world  $w$  to the set of ordered pairs  $\langle x, y \rangle$  for  $x$  an event in  $w$  and  $y$  involved in  $x$  in  $w$ ; (e) the semantic value of '=' according to  $\mathcal{C}_i$  is the function taking each world  $w$  to the set of ordered pairs  $\langle x, x \rangle$ ; (f) the semantic value of ' $\exists x$ ' according to  $\mathcal{C}_i$  is the function taking each world  $w$  to the set of non-empty sets of objects in  $w$ ; (g) the semantic value of ' $\neg$ ' according to  $\mathcal{C}_i$  is the function taking each set of worlds to its complement; and (h) the semantic value of ' $\wedge$ ' according to  $\mathcal{C}_i$  is the function taking each pair of sets of possible worlds to their intersection.

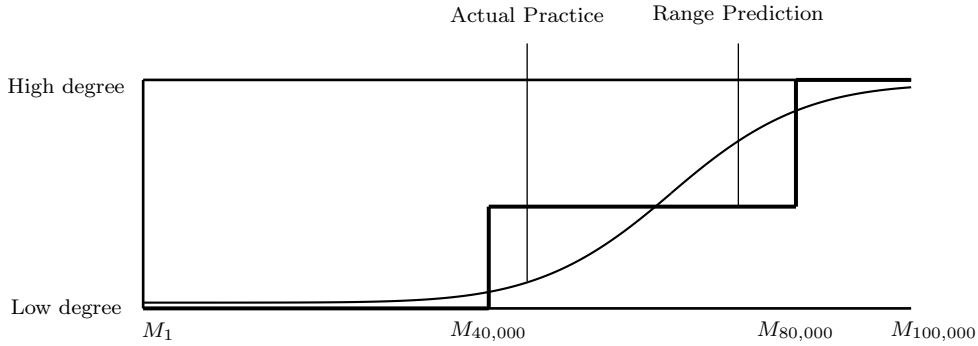


Figure 7: Degree to which  $M_k$  is actually used to mean\* the universal proposition in  $P$  vs. degree to which  $M_k$  should be used to mean\* the universal proposition in  $P$  according to a range ascription consisting of  $\mathcal{C}_{40,000}$  through  $\mathcal{C}_{80,000}$ .

progress. Little is gained, for example, by ascribing a *range* of  $\mathcal{C}_i$  to  $P$ , rather than an individual  $\mathcal{C}_i$ . Suppose, for example, that  $\mathcal{C}_{40,000}$  through  $\mathcal{C}_{80,000}$  are ascribed to  $P$ . Then, by the lights of the most natural way of extending the Absolute-Semantics Picture to range-ascriptions, one should expect  $M_0$  through  $M_{39,999}$  to be used by  $P$  to mean\* the universal proposition to a very low degree,  $M_{40,000}$  through  $M_{79,999}$  to be used by  $P$  to mean\* the universal proposition to a middling degree, and  $M_{80,000}$  through  $M_{100,000}$  to be used by  $P$  to mean\* the universal proposition to a very high degree. And, again, this is only a rough approximation of the truth (see figure 7).

In order for genuine progress to be made, the Absolute-Semantics Picture must be given up. Rather than attempting to describe  $P$ 's linguistic practice by ascribing a particular  $\mathcal{C}_i$  to  $P$  and insisting that the ascription satisfies the constraint imposed by the Absolute-Semantics Picture, one should set forth a method  $\mathcal{M}$  for gauging how apt a description of meaning\* in  $P$  each of the  $\mathcal{C}_i$  is, and describe  $P$ 's linguistic practice by saying of one or more of the  $\mathcal{C}_i$  how apt it is by the lights of  $\mathcal{M}$ . It is important that  $\mathcal{M}$  be chosen in such a way that it is clear what predictions about meaning\* one is entitled to make on the basis of an assignment of a particular degree of aptness to a classical semantics. But as long as that constraint is met, the choice of  $\mathcal{M}$  can be made on the basis of the particular needs



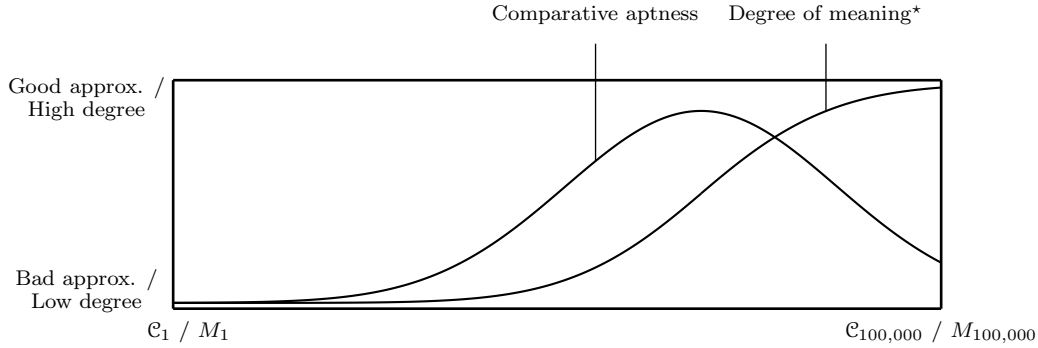


Figure 8: Comparative aptness of the  $\mathcal{C}_i$  (relative to ‘multitudinous’) / Degree to which  $M_k$  is used to mean\* the universal proposition in  $P$ .

of semanticists.

A method  $\mathcal{M}$  might measure aptness in two different ways. On the one hand,  $\mathcal{M}$  might measure the aptness of a semantics *relative to a given expression*. To fix ideas, think of the degrees of aptness of the  $\mathcal{C}_i$  relative to ‘multitudinous’ as given by the first derivative of the curve in figure 4 (see figure 8).<sup>19</sup> On the other hand,  $\mathcal{M}$  might deliver an overall degree of aptness for a given semantics. To fix ideas, think of the overall degree of aptness of  $\mathcal{C}_i$  as a weighted average of  $\mathcal{C}_i$ ’s degrees of aptness relative to a suitable class of expressions and the degrees to which  $\mathcal{C}_i$  is vindicated (in the sense of section 2.2) with respect to a suitable range of sentences.

Here is one example of a semantics for vague terms which is in accordance with the new approach. Say that a *weighted semantic space* is a pair  $\langle \mathcal{S}, \mathcal{W} \rangle$  such that: (a)  $\mathcal{S}$  is a set of classical semantics, all of which agree about a semantic lexicon from which sentences are to be built, and about the particular semantic structure that is to be assigned to each

<sup>19</sup>More precisely, assume that  $\vec{x}$  has  $m$  components, and let  $\vec{x}_j$  be the  $j$ th component of  $\vec{x}$ . For  $j \leq m$ , let  $f_j : \mathbb{N} \rightarrow [0, 1]$  be such that  $f_j(k) = \vec{x}_j$ , where  $\vec{x}$  is the degree to which  $M_k$  is used by  $P$  to mean\* the universal proposition, and let  $g_j : \mathbb{R} \rightarrow [0, 1]$  be the error function best approximating  $f_j$ . Say that the degree of approximation of  $\mathcal{C}_n$  is  $\langle g'_1(n), \dots, g'_m(n) \rangle$ . If degrees of approximation need be ordered, one can use the method described in footnote 6.

Relative to expressions such as ‘intelligent’, which have multiple dimensions of vagueness, one would need the degree of aptness of  $\mathcal{C}_i$  to be an  $n$ -tuple with different components corresponding to different dimensions.

sentence; and (b)  $\mathcal{W}$  is a function assigning a degree  $\delta_\alpha$  to each  $\mathcal{C}_\alpha$  and a degree  $\delta_\alpha^E$  to each pair  $\langle \mathcal{C}_\alpha, E \rangle$  (where  $\mathcal{C}_\alpha \in \mathcal{S}$  and  $E$  is an expression). Assume, moreover, that  $\mathcal{M}$  is an agreed method for gauging the degrees of aptness of the  $\mathcal{C}_\alpha$ . The proposal is then that a semantics for vague terms is a weighted semantic space, and that one is justified in ascribing a space  $\langle \mathcal{S}, \mathcal{W} \rangle$  to a population  $P$  only if the degrees assigned by  $\mathcal{W}$  are the degrees to which the members of  $\mathcal{S}$  approximate the truth about meaning\* in  $P$ , by the lights of  $\mathcal{M}$ .

Here is another example of a semantics for vague terms which is in accordance with the new approach. Assume, as before, that  $\mathcal{M}$  is an agreed method for gauging the degrees of aptness of the  $\mathcal{C}_\alpha$ . Let a semantics for vague terms be a classical semantics, and say that one is justified in ascribing a classical semantics  $\mathcal{C}_\alpha$  to the relevant population just in case  $\mathcal{C}_\alpha$  is an ‘optimal’ approximation of the truth about meaning\* in that population, by the lights of  $\mathcal{M}$ . (One could say, for example, that the semantics which is optimal amongst the  $\mathcal{C}_i$  with respect to ‘multitudinous’ is whichever comes closest to the point in which the second derivative of the curve in figure 4 is zero.) From the present perspective, the ascription of  $\mathcal{C}_\alpha$  to  $P$  should lead one to expect that  $\mathcal{C}_\alpha$  is an optimal description of meaning\* by the lights of  $\mathcal{M}$ , but *not* that  $\mathcal{C}_\alpha$  is a correct description of  $P$ ’s linguistic practice by the lights of the Absolute-Semantics Picture.

Different types of semantics for vague terms will capture different amounts of information about meaning\*, and will be more or less succinct. The first of the two proposals suggested above is capable of delivering more information about meaning\* than the second, but is less succinct. Fortunately, the proposals are not in competition. When properly understood, they both deliver accurate descriptions of the linguistic practice of the relevant population. One can therefore make use of whichever is best suited to one’s theoretical needs.

## 2.4 Supervenience

Any account of vagueness should allow for an explanation of how it is that the semantic facts supervene on the non-semantic facts. From our present perspective, the semantic facts pertaining to a population  $P$  are facts about meaning\* in  $P$  and facts about which semantics should be ascribed to  $P$ . The non-semantic facts include facts about the beliefs and desires of members of  $P$ , facts about the regularities that members of  $P$  conform to, facts about the environmental conditions under which members of  $P$  operate, and so forth.

According to the account of meaning-fixation set forth in section 1.3, a sentence  $s$  is used by  $P$  to mean\*  $p$  (to degree  $\vec{x}$ ) just in case an appropriate convention of truthfulness and trust prevails in  $P$  (to degree  $\vec{x}$ ). So, thanks to Lewis's account of convention, one has an explanation of how it is that the facts about meaning\* in  $P$  supervene on facts about the beliefs and desires of members of  $P$ , facts about the regularities that members of  $P$  conform to, and so forth.

On the present proposal, one should provide a semantics for vague terms by setting forth a method  $\mathcal{M}$  for gauging how apt a description of meaning\* in  $P$  a given classical semantics is, and go on to say of one or more classical semantics how apt it is by the lights of  $\mathcal{M}$ . So, for any choice of  $\mathcal{M}$  which makes clear how the degree of aptness of a classical semantics is to be determined on the basis of facts about meaning\*, one will have an explanation of how it is that facts about which semantics should be ascribed to  $P$  supervene on facts about meaning\* in  $P$ .

## 2.5 Characterizing vagueness

An answer to the question of what it is about vague terms that makes them vague flows naturally from the metasemantic account. Let  $\langle \mathcal{S}, \mathcal{W} \rangle$  be a weighted semantic space describing the linguistic practice of English speakers, and assume that  $\mathcal{S}$  consists of a large

and diverse family of classical semantics. One can then say that an expression  $E$  is vague to the extent that there is a range of semantics in  $\mathcal{S}$  which disagree about the semantic value of  $E$  along some appropriate parameter and to which  $\mathcal{W}$  assigns a significantly diverse but smoothly distributed range of degrees of aptness (relative to  $E$ ). In particular, one can say that ‘multitudinous’ is vague because there is a range of semantics in  $\mathcal{S}$  which disagree about the minimum number of persons an event must involve if it is to fall under ‘multitudinous’, and to which  $\mathcal{W}$  assigns a significantly diverse but smoothly distributed range of degrees of aptness (relative to ‘multitudinous’).

I would like to mention a virtue of the proposed characterization of vagueness. Because no appeal has been made to the notion of definiteness or its variants, there is no need to supplement the proposal with a characterization of *higher-order* vagueness.<sup>20</sup> The full depth of the notion of vagueness is captured from the start.

## 2.6 Linguistic competence

Suppose that  $\langle \mathcal{S}, \mathcal{W} \rangle$  is a weighted semantic space describing the linguistic practice of English speakers. Should one conclude that the linguistic *competence* of English speakers consists of (implicit or explicit) knowledge of  $\langle \mathcal{S}, \mathcal{W} \rangle$ ?

I would like to suggest not. It is sufficient for basic linguistic competence to know of some classical semantics that it is a rough approximation of the truth. For no more is required to be in a position to exercise reasonable conformance to conventions of truthfulness and trust prevailing in the relevant linguistic community.<sup>21</sup> To be minimally competent in the use of ‘multitudinous’, for example, it is sufficient to know of, e.g.  $\mathcal{C}_{70,000}$  that it is at least a rough approximation of the truth. If one knows, in addition, that none of the  $\mathcal{C}_i$  is a *perfect* approximation (and that the different degrees of aptness vary along the lines

---

<sup>20</sup>For a detailed characterization of higher-order vagueness, see Williamson (1999).

<sup>21</sup>Provided, of course, that one is competent with the workings of pragmatics.

described in section 2.5), then one also knows that ‘multitudinous’ is vague. So one has more than minimal competence in the use of the term.

## 2.7 Truth

Suppose  $\mathcal{M}$  is an agreed method for gauging the aptness of classical semantics, and imagine a perfect linguistic community: for any sentence  $s$  there is a proposition  $p$  such that speakers use  $s$  to mean\*  $p$  to degree  $\vec{1}$  (and use  $s$  to mean\* any other proposition to degree  $\vec{0}$ ). If all goes well, one can expect to find a classical semantics which is a *perfectly apt* description of the relevant linguistic practice by the lights of  $\mathcal{M}$ .

For better or worse, English speakers do not constitute a perfect linguistic community: single sentences with simple enough syntactic structures are used to mean\* different propositions to different degrees. If the description in section 1.4 is accurate, for example, then  $M_{80,000}$  is used both to mean\* the universal proposition to a high degree and to mean\* the empty proposition to a low degree. It is therefore unreasonable to expect to find a classical semantics which is perfectly apt by the lights of  $\mathcal{M}$ . Instead, one can expect to find a family of classical semantics to which  $\mathcal{M}$  assigns varying degrees of non-perfect aptness (and, perhaps, a classical semantics to which  $\mathcal{M}$  assigns ‘optimal’, though non-perfect, aptness).

When it comes to a perfect linguistic community, one can speak of the truth-status of any given sentence with no further ado: if  $\mathcal{C}^*$  is the classical semantics which describes the linguistic practice with perfect aptness (by the lights of  $\mathcal{M}$ ), one can say that  $s$  is true, as used by members of the perfect linguistic community, just in case  $\mathcal{C}^*(s)$  contains the actual world as a member. Not so for the imperfect community of English speakers. Even if  $\mathcal{C}^*$  turned out to be an ‘optimally apt’ description of the practice of English speakers (by the lights of  $\mathcal{M}$ ), it would be misleading to say that a sentence  $s$  is true, as used by English speakers, just in case  $\mathcal{C}^*(s)$  contains the actual world as a member. For by doing so one would ignore the differences between the community of English speakers and a perfect

linguistic community. It is better to say that a sentence  $s$  is true with respect to a classical semantics which is optimally apt just in case  $\mathcal{C}^*(s)$  contains the actual world as a member. More generally, if  $\langle \mathcal{S}, \mathcal{W} \rangle$  is a weighted semantic space describing the practice of English speakers, one can say that a sentence  $s$  is true with respect to a classical semantics which is apt to degree  $\delta$  just in case some  $\mathcal{C}_\alpha \in \mathcal{S}$  is such that  $\mathcal{W}(\mathcal{C}_\alpha) = \delta$  and  $\mathcal{C}_\alpha(s)$  contains the actual world as a member. (Similarly, one can say that  $s$  expresses proposition  $p$  with respect to a classical semantics which is apt to degree  $\delta$  just in case some  $\mathcal{C}_\alpha \in \mathcal{S}$  is such that  $\mathcal{W}(\mathcal{C}_\alpha) = \delta$ .)

I am not, however, suggesting that the notion of truth *simpliciter* be forgone. I will argue in section 3.2 that, even if an English *sentence* cannot be said to be true or false *simpliciter*, it can be used to make *assertions* which are true or false *simpliciter*.

Before bringing this section to a close, I would like to address some potential sources of concern:

1. Couldn't one give a supervaluational characterization of truth *simpliciter*? In particular, if  $\langle \mathcal{S}, \mathcal{W} \rangle$  is a weighted semantic space describing the practice of English speakers, couldn't one say that a sentence  $s$  is true *simpliciter* just in case  $\mathcal{C}_\alpha(s)$  contains the actual world as a member for every  $\mathcal{C}_\alpha \in \mathcal{S}$ ?

It is unclear that much would be gained by setting forth such a definition. For different sentences would be counted as true *simpliciter* on different choices of  $\mathcal{S}$ . But the question of which semantics to include in  $\mathcal{S}$  is to be answered on the basis of the theoretical needs of semanticists, not on the basis of information about the linguistic phenomenon under study. This is because by removing a semantics from  $\mathcal{S}$ , or adding a semantics to  $\mathcal{S}$  and extending  $\mathcal{W}$  in accordance with  $\mathcal{M}$ , one does not alter the accuracy with which meaning\* is described by  $\langle \mathcal{S}, \mathcal{W} \rangle$ ; what changes is the amount of information delivered. (And, of course, when  $\mathcal{S}$  contains *every* classical semantics which is apt to some positive degree, one should expect to get the unwelcome result

that a sentence like  $M_{100,000}$  is not counted as true *simpliciter*, since  $M_{100,000}$  is likely to be false relative to some classical semantics which is apt to a positive—albeit very low—degree.)

2. Don't we get a degrees-of-truth account of vagueness after all? For suppose  $\langle \mathcal{S}, \mathcal{W} \rangle$  is a weighted semantic space describing the practice of English speakers, and assume that the  $\mathcal{C}_\alpha \in \mathcal{S}$  are *normalized*, in the sense that the sum of the degrees of aptness assigned by  $\mathcal{W}$  to the  $\mathcal{C}_\alpha \in \mathcal{S}$  relative to any item in the shared semantic lexicon is  $\vec{1}$ . Couldn't one then say that the degree of truth of a sentence  $s$  is the sum of the degrees of aptness (relative to some lexical item occurring in  $s$ ) of the  $\mathcal{C}_\alpha \in \mathcal{S}$  such that  $\mathcal{C}_\alpha(s)$  contains the actual world?

One could certainly speak of such 'degrees of truth' if one wished, but it would be misleading to call the resulting proposal a degrees-of-truth account of vagueness. For, to return a theme from section 2.1, one will not generally get the result that the 'degree of truth' of a logically complex sentence is a function of the degrees of truth of its components, which is a central assumption of standard degrees-of-truth accounts of vagueness.

## 3 Pragmatics

### 3.1 The Stalnaker-Lewis Picture

Here is a rough outline of the Stalnaker-Lewis picture of conversation.<sup>22</sup> To each stage of a non-defective conversation corresponds a *context-set*: a set of possible worlds which are 'open-possibilities' as far as the presuppositions of participants at that stage in the

---

<sup>22</sup>See Stalnaker (1999) (especially chapters 1–5) and Lewis (1979).

conversation are concerned.<sup>23</sup> As the conversation evolves, the context-set changes. In particular, the context-set changes when an assertion is made. One way in which it changes is by registering the fact that an assertion has been made, i.e. by shifting so as to contain only possible worlds in which the relevant assertion takes place. But there is another, more important, way in which an assertion can alter the context-set: if participants accept the assertion, then the context-set shifts so as to include only possible worlds which are compatible with the content of the assertion. This is the assertion's *essential effect*.

On this picture, one may take the content of an assertion to be a subset of the context-set on which the assertion has its essential effect: the subset consisting of those worlds which are left as 'open possibilities' in light of what has been asserted. In this section I will try to explain, from the perspective of the metasemantic account of vagueness, how it is that the content of a given assertion may be determined.

The most straightforward proposal would be this: the content of an assertion of sentence  $s$  is the intersection of the context-set on which the assertion has its essential effect and the set of possible worlds in which the proposition expressed by  $s$  is true. Unfortunately, this suggestion is not available to proponents of the metasemantic account of vagueness. For, on the metasemantic account, it is not fully accurate to say of a sentence that it expresses a proposition *simpliciter*. One must say that it expresses a proposition *relative to a given semantics* (which is apt to a certain degree).

Fortunately, there is a different proposal available. Say that a semantics  $\mathcal{C}_\alpha$  *selects*  $A$  as the content of an assertion just in case  $A$  is the intersection of the context-set on which the assertion has its essential effect and the set of possible worlds in which the proposition expressed by the sentence asserted relative to  $\mathcal{C}_\alpha$  is true. The proposal is this: if amongst the sets of possible worlds selected as the content of an assertion by a sufficiently large

---

<sup>23</sup>For some purposes, it is best to think of 'open possibilities' not as (maximally specific) possible worlds, but as the elements of a partition defined on a set of (maximally specific) possible worlds. I shall ignore this complication for expositional purposes.



range of semantics there is one which is salient to participants in the conversation, then that proposition is the content of the assertion; otherwise, the assertion is *defective*, and has no content. Here are a couple of examples:

1. An assertion of ‘John is bald’ is to have its essential effect on a context-set consisting of worlds of the following two sorts:  $\alpha$ -worlds, in which John is a clear case of baldness, and  $\beta$ -worlds, in which John is a clear case of non-baldness. One can expect that any semantics which is apt to anything but the lowest of degrees (relative to ‘bald’) will select the set of  $\alpha$ -worlds as the content of the assertion. This should make the set of  $\alpha$ -worlds salient to participants in the conversation. Accordingly, the content of the assertion is the set of  $\alpha$ -worlds.
2. An assertion of ‘John is bald’ is to have its essential effect on a context-set consisting of worlds of following two sorts:  $\alpha$ -worlds, in which John has no hairs on his head, and  $\beta$ -worlds, in which John has  $n$  hairs on his head (appropriately distributed) for some  $n$  such that men with  $n$  hairs on their heads (appropriately distributed) are borderline cases of baldness. One can expect that any semantics which is apt to anything but the lowest of degrees (relative to ‘bald’) will select either of the following two propositions as the content of the assertion: (a) the set of  $\alpha$ -worlds, or (b) the set of all worlds in the context-set. This should make the set of  $\alpha$ -worlds salient to participants in the conversation, since accepting the set of all worlds in the context-set as the content of the assertion would make the assertion pointless. Accordingly, the content of the assertion is the set of  $\beta$ -worlds.

So far so good. But now consider a conversation in which nothing special is presupposed about the abundance of John’s hair. The context-set might start out containing a world in which John has  $n$  hairs on his head (appropriately distributed), for each  $n$  between 0 and 150,000. If that is the context-set on which an assertion of ‘John is bald’ is to have

its *essential* effect, then the assertion is likely to be defective. For one can expect that there will be a large number of rival propositions, each selected by a similarly broad range of semantics which are apt to similarly substantial degrees (relative to ‘bald’); so no one proposition will be salient to participants in the conversation. But how can this be right? Doesn’t everyday conversation supply examples of cases in which it is appropriate to assert that someone is bald even if one’s audience presupposes nothing about how much hair that person has?

The way out of the problem is to note that, regardless of what the original context-set consists in, an assertion of ‘John is bald’ will not usually have its *essential* effect on a context set containing too many worlds. For everyday conversation is governed by what Lewis (1979) calls ‘rules of accommodation’, according to which the context-set should evolve so as to make it the case that no assertion is defective. Thus, one can expect the assertion in our example to have its essential effect on a *modified* context-set: the result of excluding as many worlds as necessary from the original context-set for participants in the conversation to regard a selected proposition as sufficiently salient, and thereby forestall defectiveness. This will normally mean excluding worlds in which John is a borderline case of baldness, so that we are left with a situation such as the one described in example 1 above.<sup>24</sup> (Just which worlds are presupposed-away from the context-set is something it is up to each participant in the conversation to decide. If the modifications of different participants are close enough not to affect common goals, then the conversation may carry on without incident. Otherwise, there is a risk of defectiveness, and further adjustments may be necessary. See Stalnaker (1979).)

It is worth emphasizing that defectiveness-averting modifications of the context-set may not always be possible (as when a discussion of the phenomenon of vagueness is salient).

---

<sup>24</sup>It is worth noting that this sort of story predicts that, in contexts where nothing special is presupposed about the abundance of John’s hair, one will only assert ‘John is bald’ if one takes John to be a clear case of baldness. For, if John is a borderline case of baldness, the context set will exclude the actual world when it contracts to avert defectiveness, thereby falsifying the content of the assertion.

Our discussion of assertoric content has been developed on the basis of the Stalnaker-Lewis picture of conversation. But I hope that the plausibility of the following conclusions is relatively independent of some of the details of that picture:

1. Even if a sentence doesn't express propositions *simpliciter*, but only relative to some semantics or other, an *assertion* of that sentence can enjoy a determinate content because the situations which are regarded as 'open possibilities' for the purposes of the assertion can be fairly limited.
2. An assertion of a sentence involving vague terms may fail to enjoy determinate content if—despite the presence of rules of accommodation—enough scenarios are regarded as 'open possibilities'. If so, the assertion is pragmatically defective, and should be avoided.

### 3.2 Truth *simpliciter*

I am now in a position to say why I think that a sentence can be used to make an *assertion* which is true or false *simpliciter*, even if the sentence itself cannot be said to be true or false *simpliciter*, but only true or false relative to a semantics (which is apt to a certain degree). To wit: when the assertion enjoys a determinate content, its truth-value *simpliciter* may be identified with the truth-value of its content.

Since an argument's logical status does not depend on the meaning of its non-logical vocabulary, there is no need to give a pragmatic characterization of logical consequence: the standard model-theoretic characterization will do.<sup>25</sup> In particular, one gets the following welcome result. Let  $a_\phi$  and  $a_\psi$  be (non-defective) assertions of  $\phi$  and  $\psi$ , respectively, and

---

<sup>25</sup>All of this, *modulo* the sorts of considerations set forth in [Author's article]. Whether or not any English sentences turn out to be logically valid will depend on whether English particles such as 'and' and 'not' turn out to be precise terms expressing the logical connectives.

assume that there is no difference in the context-sets on which they have their essential effects. If  $a_\phi$  is true *simpliciter* and  $\phi \models \psi$ , then  $a_\psi$  must also be true *simpliciter*.

### 3.3 The Sorites Paradox

Here is a version of the Sorites Paradox as it arises with respect to the  $M_n$ :

Each of the following appears to be true:

1.  $M_{100,000}$
2.  $\neg M_1$
3.  $\forall k(M_{k+1} \rightarrow M_k)$

But, by iterated applications of universal instantiation and *modus ponens*, one can derive a contradiction from 1–3.<sup>26</sup>

In order to address the paradox, I will proceed in two stages. First I will argue that the metasemantic account can explain why 1–3 appear to be true. Then I will show that a contradiction is nonetheless averted from the perspective of the metasemantic account.

It is straightforward for a proponent of the metasemantic account to explain why  $M_{100,000}$  and  $\neg M_1$  appear to be true. For, according to the metasemantic account, typical assertions of  $M_{100,000}$  and  $\neg M_1$  are true (provided the conversation carries no false presuppositions). Indeed, in the absence of false presuppositions, the truth of such assertions is guaranteed by the relevant linguistic conventions, together with the norms governing the way in which an assertion acquires a content in a given context.<sup>27</sup> So all we need to explain

---

<sup>26</sup>I trust the reader will forgive use/mention infelicities in 2 and 3.

<sup>27</sup>This can easily be verified when propositions are assumed to be sets of possible worlds. For relative to all but a few compositional semantics (none of which is apt to anything but the lowest of degrees relative to ‘multitudinous’),  $M_{100,000}$  expresses the universal proposition. So, in all but the most devious of contexts, an assertion of  $M_{100,000}$  must have a true content—provided the context-set on which it is to have its essential effect is not marred by false presuppositions. (As Stalnaker has taught us, the

the appearance of truth of  $M_{100,000}$  and  $\neg M_1$  is the assumption that speakers' intuitions are suitably informed by their linguistic competence.

The real challenge is explaining the appearance of truth of  $\forall k(M_{k+1} \rightarrow M_k)$ . I will proceed by offering an explanation of why it seems so natural to think that there could be no counterexample to  $\forall k(M_{k+1} \rightarrow M_k)$ .

Here is the explanation. It is natural to suppose that, in order to produce a counterexample to  $\forall k(M_{k+1} \rightarrow M_k)$ , one must identify a true instance of the following:

(\*) Event  $e'$  is multitudinous, but event  $e$  is not;

(where it is presupposed that  $e'$  involves precisely one more person than  $e$ ). But it is a consequence of the metasemantic account that no such identification is possible. For, in any reasonable context, instances of (\*) will be unassertable as a matter of principle. The reason is this. If the discussion in section 3.1 is broadly along the right lines, an assertion can only escape defectiveness if the situations which are regarded as 'open possibilities' for the purposes of the assertion can be successfully classified by conversational participants into two distinct groups. But, when it comes to assertions involving 'multitudinous', our linguistic practice provides no more than rough guidance as to how such a classification is to be carried out (since 'multitudinous' is assigned no unique extension). This means that defectiveness can only be averted if some classification of the 'open possibilities' is salient enough to conversational participants that it can be singled it out merely on the basis of the rough guidance that linguistic practice supplies. The standard way for a classification to be rendered salient in such cases is for the range of 'open possibilities' to be fairly limited, as illustrated in section 3.1.<sup>28</sup> But it becomes impossible to exclude the right sorts of 'open

---

information conveyed by such an assertion is unlikely to coincide with its content. The content of the assertion is utterly uninformative because it consists of every world in the context-set. So the information conveyed must be given by the *diagonal*, which excludes worlds in the context-set where 'multitudinous' is used non-standardly. See Stalnaker (1979).)

<sup>28</sup>Given a sufficiently unusual context, there might be other ways. For instance, conversational participants might somehow presuppose that it is safe for 59,792 persons or less to get together, but unsafe for

possibilities’ when it comes to an assertion of (\*), for its conjuncts impose incompatible demands on the ‘open possibilities’ which are to be left out.

When propositions are taken to be sets of possible worlds, the argument can be spelled-out in further detail. Let  $\mathcal{C}_k$  be a semantics that places the cutoff point of ‘multitudinous’ at  $k$ . Then the content selected by  $\mathcal{C}_k$  for (\*) can only contain worlds in which  $e$  involves precisely  $k$  persons and  $e'$  involves precisely  $k+1$  persons, since these are the only worlds in which (\*) can be true, according to  $\mathcal{C}_k$ . But this means that no two semantics which differ in the extension they assign to ‘multitudinous’ can select the same non-empty content for (\*). So no non-empty content can be rendered sufficiently salient to participants in the conversation. The assertion must therefore be defective.<sup>29</sup>

So much for my explanation of why the Sorites premises appear to be true. How is it that a contradiction is nonetheless averted? From the perspective of the metasemantic account, the basic answer is straightforward: assertions of  $\forall k(M_{k+1} \rightarrow M_k)$  are *false* in any reasonable context because, for  $1 \leq k \leq 100,000$ ,  $\forall k(M_{k+1} \rightarrow M_k)$  expresses a necessarily false proposition according to all but a few compositional semantics (none of which is apt to anything but the lowest of degrees relative to ‘multitudinous’).

I would like to suggest, however, that the fact that typical assertions of  $\forall k(M_{k+1} \rightarrow M_k)$  are false is not an especially illuminating piece of information about the semantic status of  $\forall k(M_{k+1} \rightarrow M_k)$ . It is more a reflection of the theoretical machinery we chose to describe the practice of speakers—i.e. classical semantics—than of the practice itself.<sup>30</sup> In order to say something illuminating about the semantic status of  $\forall k(M_{k+1} \rightarrow M_k)$  we need a finer instrument than truth *simpliciter*; we need to say something such as the following:

---

59,793 persons to get together.

<sup>29</sup>Here I assume that the context-set on which (\*) is to have its essential effect contains no worlds that disagree about what our linguistic conventions are like. Otherwise, the assertion might be non-defective by having a determinate and non-empty diagonal content.

<sup>30</sup>It is worth emphasizing that, as far as the metasemantic account is concerned, one’s theoretical machinery needn’t be based on classical semantics. An approach based on non-classical semantics might do, provided one is able to set forth a suitable method for gauging how apt a description of meaning\* a given semantics of the relevant kind is, and supply a suitable pragmatics.

For any choice of  $\mathcal{C}_i$  meeting the constraints of section 2.3, the  $\mathcal{C}_i$  display a smooth and monotone gradation of degrees of aptness relative to ‘multitudinous’, which go from very low to very high as  $i$  increases.

### 3.4 Contextualism

A virtue of the present proposal is that it allows us to explain *why* it is that the sorts of contextual phenomena that motivate contextualist accounts of vagueness take place.<sup>31</sup>

Here are a couple of examples:

1. My only copy of *War and Peace* is colored a very light grayish-blue. In context A it is sitting amongst bright red books. When I assert ‘I would like you to bring me the blue book’, you bring me *War and Peace*. In context B my copy of *War and Peace* is sitting amongst bright cobalt-blue books. When I assert ‘I would like you to bring me the gray book’, you bring me *War and Peace*.

*Explanation:* In context A the context-set contains only open-possibilities according to which I would like you to bring me a book which is in fact bright red and open-possibilities according to which I would like you to bring me a book which is in fact *War and Peace*. In context B the context-set contains only open-possibilities according to which I would like you to bring me a book which is in fact bright cobalt-blue and open-possibilities according to which I would like you to bring me a book which is in fact *War and Peace*. The explanation can then proceed as in example 2 from section 3.1.

2. When it is contextually salient that two men have similar amounts and distributions of hair on their head, speakers will refrain from asserting that one of them is bald and the other is not.

---

<sup>31</sup>See Kamp (1981), Soames (1999) chapter 7, Raffman (1994), Raffman (1996) and Fara (2000). My examples follow Fara (2000).

*Explanation:* Parallel to the argument in section 3.3 for the conclusion that (\*) is unassertable.

## 4 Conclusions

I have argued for an account of vagueness according to which the root of vagueness lies not in the type of semantic-value that is best associated with an expression, but in the type of linguistic practice that renders the expression meaningful. I suggested, in particular, that conventions about how to use sentences involving attributions of vague predicates to borderline cases prevail to lesser degrees than conventions about how to use sentences involving attributions of vague predicates to clear cases.

This approach led to the conclusion that a sentence shouldn't generally be said to express a proposition *simpliciter*, but only relative to a given semantics (which is apt to a certain degree). I noted, however, that typical *assertions* will nonetheless turn out to be true or false *simpliciter*.

In arguing for the metasemantic account I sometimes presupposed particular theories of conventionality, meaning-fixation and conversational pragmatics. This was necessary to provide a detailed defense of the proposal. But, as far as I can tell, the proposal does not depend on the details of these presuppositions.<sup>32</sup>

---

<sup>32</sup>I am grateful to audiences at the University of Bologna, the University of California at Irvine, Glasgow University, the Jagiellonian University, MIT, the Ohio State University at Columbus, Princeton University, the University of St Andrews, Warsaw University and the University of Wisconsin at Madison. Special thanks are due to Jonathan Cohen, Juan Comesaña, Roy Cook, Richard Dietz, Patrick Greenough, Øystein Linnebo, Sebastiano Moruzzi, Graham Priest, Elliott Sober, Robbie Williams, Crispin Wright and an anonymous referee. I gratefully acknowledge that most of this paper was written during the tenure of a research fellowship at Arché, the AHRC Research Centre for the Philosophy of Logic, Language, Mathematics and Mind.



## References

- Black, M. (1949*a*), *Language and Philosophy*, Cornell University Press, Ithaca, NY.
- Black, M. (1949*b*), ‘Vagueness: An exercise in logical analysis’, *Philosophy of Science* **4**, 427–55. Reprinted in Black (1949*a*) and Keefe & Smith (1996).
- Bräuerle, R. et al., eds (1979), *Semantics from Different Points of View*, Springer-Verlag.
- Burge, T. (1975), ‘On knowledge and convention’, *The Philosophical Review* **84**, 249–55.
- Dorr, C. (2003), ‘Vagueness without ignorance’, *Philosophical Perspectives* **17**, 83–113.
- Fara, D. G. (2000), ‘Shifting sands: An interest-relative theory of vagueness’, *Philosophical Topics* **28**, 45–81. Originally published under the name ‘Delia Graff’.
- Gilbert, M. (1989), *On Social Facts*, Princeton University Press, Princeton, NJ.
- Grandy, R. (1977), ‘Review of lewis’s *Convention: A Philosophical Study*’, *The Journal of Philosophy* **74**, 129–39.
- Jamieson, D. (1975), ‘David lewis on convention’, *Canadian Journal of Philosophy* **5**, 73–81.
- Kamp, H. (1981), ‘The paradox of the heap’. in Mönnich (ed.) *Aspects of Philosophical Logic*.
- Keefe, R. & Smith, P., eds (1996), *Vagueness: A Reader*, MIT Press, Cambridge, MA.
- Laurence, S. (1996), ‘A chomskian alternative to convention based semantics’, *Mind* **105**, 269–301.
- Lewis, D. (1969), *Convention: A Philosophical Study*, Harvard University Press, Cambridge, MA.

- Lewis, D. (1970), ‘General semantics’, *Synthese* **22**, 18–67. Reprinted in Lewis (1983).
- Lewis, D. (1973), ‘Languages and language’, *Versus* **4**, 2–21. Appears in Italian translation as ‘Lingue e lingua’. Reprinted in English in Lewis (1983).
- Lewis, D. (1979), ‘Score-keeping in a language game’, *The Journal of Philosophical Logic* **8**, 339–59. Reprinted in Bräuerle et al. (1979); reprinted in Lewis (1983).
- Lewis, D. (1983), *Philosophical Papers, Volume I*, Oxford.
- Raffman, D. (1994), ‘Vagueness without paradox’, *Philosophical Review* **103**(1), 83–111.
- Raffman, D. (1996), ‘Vagueness and context-relativity’, *Philosophical Studies* **81**, 175–92.
- Rayo, A. (2008), ‘Vague representation’, *Mind* **117**, 329–373.
- Schiffer, S. R. (1972), *Meaning*, Clarendon Press, Oxford.
- Schiffer, S. R. (1987), *Remnants of Meaning*, MIT Press, Cambridge, MA.
- Soames, S. (1999), *Understanding Truth*, Oxford University Press, Oxford and New York.
- Stalnaker, R. C. (1979), ‘Assertion’, *Syntax and Semantics* **9**, 315–322. Reprinted in Stalnaker (1999), pp. 78–95.
- Stalnaker, R. C. (1999), *Context and Content*, Oxford University Press, Oxford.
- Williamson, T. (1994), *Vagueness*, Routledge, London and New York.
- Williamson, T. (1999), ‘On the structure of higher-order vagueness’, *Mind* **108**, 127–44.