

Vague Representation

AGUSTÍN RAYO

The goal of this paper is to develop a theory of content for vague language. My proposal is based on the following three theses: (1) language-mastery is not rule-based—it involves a certain kind of *decision-making*; (2) a theory of content is to be thought of instrumentally—it is a *tool* for making sense of our linguistic practice; and (3) linguistic contents are only *locally* defined—they are only defined relative to suitably constrained sets of possibilities.

The goal of this paper is to develop a theory of content for vague language. My proposal will be based on the following three theses:

- (1) Language-mastery is not a matter of gaining cognitive access to semantic rules corresponding to the expressions of one's language and learning how to deploy the rules in linguistic practice. Instead, it involves a form of *decision-making* which I call 'semi-principled'.
- (2) A theory of content is to be thought of instrumentally: it is a *tool* for making sense of our linguistic practice.
- (3) Linguistic contents are only *locally* defined: they are only defined relative to suitably constrained sets of possibilities.

Although these theses are independent of each other, we will see that the second and third gain plausibility in light of the first, and that there are potential worries about the third that are adequately addressed by appeal to the second.

In the course of the paper I hope to show that the three theses combine to form an organic picture of vague language, and that the picture is attractive enough to lend plausibility to the theses. The picture is divided into two main parts, which will be developed in sections 1 and 2, respectively. The first part is an account of linguistic usage; the second part is an account of linguistic content, based on the account of usage.

1. Usage

The purpose of this section is to argue that linguistic practice involves semi-principled decision-making. I will begin by setting forth an example, and using it to explain what semi-principled decision-making consists in. I will then consider a series of variations of the example, and argue that there is an important analogy between the last few members of the series and our actual linguistic practice.

1.1 The examples

Example 1: Indication

You are a participant in a game-show. After examining John's face, you are asked to name a positive integer. If the integer corresponds to John's age in days, you win a prize. If not, you receive a punishment. How will you proceed?

Not every integer will strike you as equally plausible in light of your examination. But the examination will not immediately put you in a position to specify an integer. The problem is that the information you are able to extract from the examination is in a form that is not perfectly suited to the task at hand. Whereas the information you are able to extract is in terms of skin-textures, fat-distributions, feature-proportions, etc., the task at hand calls for an integer. And there are limits in your ability to go from the one to the other. You are therefore not in a position to express the result of your examination as a statement of John's age in days. Your examination can be expected to be *tolerant* with respect to age-in-days-judgements, in the following sense: whereas some pairs of integers will strike you as plausible to very different degrees in light of the examination (e.g. 25×365 vs 75×365), pairs of neighbouring integers will strike you as similarly plausible in light of the examination (e.g. 25×365 vs $(25 \times 365) + 1$).

Three points are worth emphasizing. First, what is at issue here is not whether there is, in fact, a precise correlation between people's facial characteristics and their age in days; what is at issue is your *ability* to use information about facial characteristics to choose an integer. Second, your predicament is vagueness-independent in the following respect: it remains in place even if one assumes that John's age and features are perfectly precise, that the measurements you make in the course of your examination are perfectly precise, and that all the concepts you use in your deliberations are perfectly precise. Third, just as you are unable to express the result of your examination as a statement of John's age in days, you are unable to express the result of your exam-

ination as a probability distribution on the positive integers. Your examination is tolerant with respect to judgements about the adequacy of probability-distributions just as it is tolerant with respect to judgements about John's age in days. So even if you have resolved to act so as to maximize expected utility, your examination fails to yield a decision about how to proceed.

How then will you decide which integer to select? Since you are unable to single out any particular integer on the basis of your examination, your decision must be partly determined by considerations of a secondary sort: you might rely on a hunch, or make an arbitrary choice, or favour numbers that are easy to name in some preferred notation, or use one of the integers that audience-members call out.

I shall call decisions of the kind you face 'semi-principled'. In general, a semi-principled decision is a decision made under the following conditions: (1) although you have grounds for making your decision, your grounds are tolerant with respect to your range of options (i.e. whereas some options strike you as desirable to very different degrees on the basis of your grounds, options that are close to each other along some suitable parameter strike you as similarly desirable on the basis of your grounds); and (2) on account of such tolerance, you are unable to single out any one your options without going beyond your grounds: secondary considerations are necessary.

One would certainly like to know more about the psychological mechanisms underlying semi-principled decision-making. It would be helpful to know, in particular, how one succeeds in consolidating considerations of different sorts into a single decision. For it is not that your examination yields a definite range of 'leading candidates' from which a final choice is to be made on the basis of a secondary consideration. What happens is presumably that you use secondary considerations to choose an integer and go ahead with it as long as it is 'good enough' by the lights of your examination—and it is not entirely clear what it takes for you to judge an integer to be 'good enough'. However this may be, there are two points that it is important to be clear about. First, somehow or other we are able to come up with sensible choices when we face semi-principled decision-making. So the relevant psychological mechanisms must be in place. Second, to the extent that there is a mystery about how the mechanisms underlying semi-principled decision-making work, it is an *empirical* problem, and calls for empirical investigation.

Let me make a final observation about the case:

IMPROVEMENT FROM LIMITATION

Your decision becomes easier when there are suitable gaps between your options.

Suppose, for example, that the game-show host limits your answers to '1', ' 18×365 ', ' 45×365 ', and ' 100×365 '. Then your decision can be expected to be much easier than if no restrictions had been placed. Notice, however, that the transition between scenarios in which the gaps are big enough for you to be able to make your decision with confidence and scenarios in which the gaps are not big enough is a gradual one. Consider a series of scenarios σ_1 – σ_{20} . For each $n \leq 20$, scenario σ_n is such that the host limits your answers to ' 44×365 ' and ' $\lceil (44 + n) \times 365 \rceil$ '. In all scenarios John is 44×365 days old, and has features which are normal for his age. As n grows larger, you grow more and more confident that choosing ' 44×365 ' would be more sensible than choosing ' $\lceil (44 + n) \times 365 \rceil$ '. But you can be expected to lack a principled criterion for deciding when such confidence is confidence enough.

Example 2: Age-ranges

The game-show is as in the preceding section, except that this time you are asked to specify a *range* of positive integers after examining John's face. If John's age in days is within the specified range, you win a prize which is inversely proportional to the size of the range. If John's age is outside the specified range, there is a punishment instead of a prize. How will you proceed?

As before, you face a semi-principled decision. On account of its tolerance with respect to your options, you are unable to extract any particular range of integers from your examination. Whereas some age-ranges spanning the same number of years strike you as plausible to very different degrees, similar age-ranges spanning the same number of years strike you as similarly plausible.

As before, IMPROVEMENT FROM LIMITATION holds: your decision would become easier if there were suitable gaps between your options. Suppose, for example, that John has the features of a 30-year-old, and that the game-show limits your options to (29×365) – (31×365) and (79×365) – (81×365) . Then your decision can be expected to be much easier than if no restrictions had been placed. (Presumably ' (29×365) – (31×365) ' would make a sensible choice.)

In this second game-show scenario, however, there is an additional observation to be made:

FURTHER SECONDARY CONSIDERATIONS

Because you face a semi-principled decision, you will not be able to ground your decision solely on your examination. You might rely on hunches, or on some degree of arbitrary choice. But in the current scenario there are additional considerations that play a role. If you are risk-averse, for example, you will choose a more generous age-range than if you are not. And if you place great value on receiving a large prize but care comparatively little about being punished or receiving small prizes, you will choose less generous age-ranges. In all likelihood, you will find that there is no one size of age-range that stands out to you as best all things considered. So your decision about size may again rely on hunches, or involve some degree of arbitrary choice.

Example 3: Signalling

The game-show is as in the preceding section, except this time you are not allowed to examine John. Instead, you and I will be working as a team, and will be prized or punished together. I know John's exact age, and it is my job to convey as much of this information to you as I can. There are, however, limits to what I am allowed to do. I have a stack of photographs of people of different ages (none of which is John), and am allowed to hand you one of my photographs. But we are not allowed any other form of communication, and we are not allowed to come to prior agreements about how the photographs are to be interpreted. We have both been shown my stack in advance. All of this is common knowledge.

As before, you face a semi-principled decision. As before, IMPROVEMENT FROM LIMITATION holds. As before, FURTHER SECONDARY CONSIDERATIONS holds. But this time there are two additional observations to be made:

SENSITIVITY TO EXPRESSIVE RESOURCES

Your choice of age-ranges will be sensitive to what you know about my expressive resources.

If you know, for example, that my stack contains only two photographs, one of a young man and one of an old man, you can be expected to choose more generous year-ranges than if you know that my stack is long and varied.

INSTABILITY¹

The same signal can be used in different ways, depending on the context.

Consider two different scenarios. In Scenario A the game-show host limits your answers to $(20 \times 365) - (25 \times 365)$ and $(45 \times 365) - (50 \times 365)$. In Scenario B the host limits your answers to $(45 \times 365) - (50 \times 365)$ and $(70 \times 365) - (75 \times 365)$. In both scenarios I hand you a photograph of an 18-year-old. In Scenario A you choose $(20 \times 365) - (25 \times 365)$; in Scenario B you choose $(45 \times 365) - (50 \times 365)$, even though you ruled it out in Scenario A on the basis of the same photograph.

It is important to note that in both scenarios I am being cooperative. For I have chosen my photograph in such a way that it is clear to you how I propose that it be used to carry out the task at hand. Notice, relatedly, that it might be *preferable* for me to hand you a photograph of a man whose age is very different from John's. Suppose John is 54, and the game-show host limits your answers to $(53 \times 365) - (55 \times 365)$ and $(56 \times 365) - (59 \times 365)$. It will typically be more helpful of me to hand you a photograph of an 18-year-old, than to hand you a photograph of a 54-year-old.

Example 4: Public language

The game-show is as in the preceding section, except this time I am given a list of adjectives ('young', 'old', 'youngish', 'middle-aged', 'very old', etc.) instead of a stack of photographs. I am allowed to convey information about John's age by reading out one of the adjectives on my list, but we are not allowed any other form of communication. Nor are we allowed to come to prior agreements about special ways of understanding the adjectives on my list: we must rely on our competence as English speakers. We have both been shown my list in advance. All of this is common knowledge.

As before, you face a semi-principled decision. As before, IMPROVEMENT FROM LIMITATION holds. As before, FURTHER SECONDARY CONSIDERATIONS holds. As before, SENSITIVITY TO EXPRESSIVE RESOURCES holds. (If you know, for example, that my list contains only two adjectives, 'young' and 'old', you can be expected to choose more generous year-ranges than if you know that my list is long and varied.) And, as before, INSTABILITY holds.

¹ This sort of phenomenon is discussed in Graff 2000. Although a precursor of this point can be identified in the previous examples, the observation really only comes into its own in the case at hand.

To bring the point about INSTABILITY home, consider two different scenarios, analogous to those considered in the preceding section. In Scenario A the game-show host limits your answers to $(20 \times 365) - (25 \times 365)$ and $(45 \times 365) - (50 \times 365)$. In Scenario B the host limits your answers to $(45 \times 365) - (50 \times 365)$ and $(70 \times 365) - (75 \times 365)$. In both scenarios I say 'young'. In Scenario A you choose $(20 \times 365) - (25 \times 365)$; in Scenario B you choose $(45 \times 365) - (50 \times 365)$, even though you ruled it out in Scenario A on the basis of the same adjective. As before, it is important to note that in both scenarios I am being cooperative. For I have chosen my utterance in such a way that it is clear to you how I propose that the utterance be used to carry out the task at hand. Notice, moreover, that it might be *preferable* for me to choose an adjective that would not be used to describe John in ordinary circumstances. Suppose John is 54, and the game-show host limits your answers to $(53 \times 365) - (55 \times 365)$ and $(56 \times 365) - (59 \times 365)$. It will typically be helpful of me to say 'young'.

Example 5: Assertion

You are asked to help single out the murderer from a police lineup. The lineup consists of 60 men. For each n between 21 and 80, man n is n years old. Age aside, the men are very much alike. Your instructions are to select a range of men in the lineup, and inform the police that everyone outside that range can be ruled out as a suspect. All you have to go on is the final words of a dying witness: 'the murderer is a young man'.

As before, you face a semi-principled decision. This means that you will be unable to choose any particular range on the basis of the witness's assertion: secondary considerations will be necessary. But it does not mean that you are not in a position to help the police. You can still use the witness's assertion to rule out some of the men in the lineup as suspects—man 80, for example. (It is important to be clear that what is at issue here is not whether the *content* of the witness's assertion determines a particular range; what is at issue is your *ability* to use the assertion to choose a particular range, whatever the assertion's content.)

As before, IMPROVEMENT FROM LIMITATION holds. (If the lineup was limited to men 25 and 78, your decision would be much easier.) As before, FURTHER SECONDARY CONSIDERATIONS holds. (If you know that anyone you rule out will immediately be set free, you can be expected to be particularly cautious about ruling people out; if you know that people you fail to rule out will most likely go to jail, you can be expected to be particularly cautious about leaving people in.) As before, SENSITIVITY TO EXPRESSIVE RESOURCES holds. (If you know that the witness had

a very limited vocabulary, or very little time left to say her last words, you can be expected to be particularly cautious about ruling people out.)

Now take a case in which the witness is still alive, and is with you as you make your selection. Then INSTABILITY holds as well. To see this, consider two different scenarios. In Scenario A the police has narrowed down the lineup to men 21 and 37. In Scenario B the police has narrowed down the lineup to men 37 and 78. In both scenarios the witness asserts ‘the young man is the murderer’. In Scenario A you choose the man 21; in Scenario B you choose man 37, even though you ruled him out in Scenario A on the basis of an assertion of the same sentence. (As before, it is important to note that in both scenarios the witness is being cooperative. For she has chosen her assertion in such a way that it is clear to her audience how she proposes that the assertion be used to carry out the task at hand.)

1.1.1 Tolerance refined

The notion of tolerance will be important in what follows. So before bringing this section to an end I would like to supply a more precise characterization of tolerance. As used above, the notion of tolerance is defined in the context of a *decision*: when a subject is deciding amongst the options in a certain range, the subject’s grounds for the decision may be said to be tolerant or not relative to her options. From now on, I will focus on the special case of decision-making in which a *speech-act* has been performed, and a subject decides which of the possibilities in a given set to rule out on the basis of that speech-act.

I shall say that speech-act σ is *tolerant* (for subject S in context C , relative to set of possibilities P) just in case:

- (1) There is a parameter Π such that the possibilities in P differ with respect to Π .

(*Example:* Let P be the set of possibilities p_n ($18 \leq n \leq 100$), where p_n is the possibility that John be exactly n years old. Then a natural choice for Π is John’s age.)

- (2) There are possibilities p and p' in P that differ greatly with respect to Π and are such that, in C , S is able to discern much better reasons on the basis of σ for ruling out p than for ruling out p' .

(*Example:* Let P and Π be as above, and let σ consist of my handing you a photograph of a 21-year-old. In the context of

the game show, you are able to discern much better reasons on the basis of σ for ruling out the possibility that John is 74 than for ruling out the possibility that John is 22.)

- (3) If p and p' are possibilities in P that are similar with respect to Π , then, in C , any reason S is able to discern on the basis of σ for ruling out p is, to roughly the same extent, a reason for ruling out p' .

(*Example:* Let P , Π , and σ be as above. In the context of the game show, any reason you are able to discern on the basis of σ for ruling out the possibility that John is 53 is, to roughly the same extent, a reason for ruling out the possibility that John is 54.)

(When one or more of the parameters S , P , or C is missing and not supplied by context, the claim that a speech-act is tolerant is to be understood as the claim that the speech-act is tolerant relative to some choice for the missing parameters.)

One final point: it is worth noting that on the present understanding of tolerance, the assumption that a speech-act is tolerant does *not* lead to contradiction. This is in contrast with the (closely related) notion of tolerance set forth in Wright 1976.²

1.2 The practice of assertion

I would like to suggest that there is an illuminating analogy between the toy examples we have just considered and our everyday practice of assertion. The purpose of this section is to spell-out the analogy.

1.2.1 A Stalnakerian description of our assertoric practice

The analogy relies on a particular way of thinking about our assertoric practice, so some stage-setting will be necessary. Following Stalnaker,³ I shall assume that a conversation is usefully modelled as taking place against the background of a *context set*: the set of mutually exclusive possibilities that a conversational participant treats as *live options* for

²Wright says that a predicate 'F is *tolerant* with respect to [a parameter] ϕ if [any object which F characterizes may be changed into one which it does not simply be a sufficient change in respect of ϕ and] there is also some positive degree of change in respect of ϕ insufficient to ever affect the justice with which F applies to a particular case' (pp. 156–7). One can use classical logic to derive a contradiction from the assumption that there is a tolerant predicate in Wright's sense. (Here I take for granted that a predicate applying to different cases with 'equal justice' applies *simpliciter* to either both or neither.)

³See the articles collected in Stalnaker 1999, especially Stalnaker 1979 and Stalnaker 1998. There are slight differences between the picture I present here and Stalnaker's.

the purposes of a given stage of the conversation. For a possibility to be treated as a live option is for it to be compatible with every proposition that the participant treats as a *presupposition* of the relevant stage of the conversation. A proposition is treated as a presupposition by a conversational participant just in case she is disposed to act, for the purposes of the relevant stage of the conversation, as if the proposition was common knowledge.

As the conversation evolves, so does each participant's context set. If Susan sneezes during the course of a conversation she is involved in, participants can be expected to treat the fact that she has sneezed as a presupposition, and their context sets can be expected to evolve accordingly. Similarly, each time an assertion is made, participants can be expected to treat the fact that the assertion has been made as a presupposition, and their context sets can be expected to evolve accordingly. But there is a more specific way in which the context set might evolve as a result of an assertion. If the assertion is accepted, each participant's context set will evolve so as to exclude any possibilities she is disposed to use the assertion to rule out on the basis of her *linguistic competence* with the sentence asserted. This is the assertion's *essential effect*.

Suppose, for example, that Susan asserts 'I live in a blue house' and that your context set is exhausted by possibilities where Susan lives in a blue house and possibilities where Susan lives in a white house. If you accept Susan's assertion, you will presumably be disposed to use the assertion to exclude possibilities of the latter kind on the basis of your linguistic competence. So every possibility of the latter kind will be excluded from your context set as an essential effect of my assertion.

There may be differences between the context sets corresponding to different conversational participants. Such differences can be tolerated as long as they do not interfere with the conversational goals. But if differences become significant enough to threaten conversational goals, participants can be expected to adjust their presuppositions in such a way that discrepancies are reduced.

1.2.2 *The analogy*

You and I are involved in a conversation. At a particular stage in the conversation I am the speaker and you are the listener. It is common knowledge that you wish to know the colour of Susan's house. Nothing special is presupposed about the colour of Susan's house, but I know that she lives in a blue house, and wish to convey this information to you. I assert 'Susan lives in a blue house'.

Your present situation is not unlike your situation in the toy examples, especially examples 3–5:

- In the toy examples, you use a speech-act—the handing-out of a photograph, or the reading out of an adjective, or the assertion of a dying witness—as a ground for ‘retaining’ some items and ‘ruling out’ others. The items in question are integers in the case of examples 3 and 4, and men in the case of example 5. The items you ‘retain’ are those within the range you specify; the items you ‘rule out’ are those that lie outside this range.

In our conversation, you use my speech-act—an assertion—as a ground for adjusting your conversational presuppositions. The result of adjusting your presuppositions is that you ‘retain’ some of the possibilities in the context-set and ‘rule out’ others. The possibilities you ‘rule out’ are those that cease to be treated as live options as a result of the adjustment; the possibilities you ‘retain’ are those that continue to be treated as live after the adjustment.

- In examples 3 and 4, there are prizes and punishments. If John’s age in days is amongst the integers you retain, you receive a prize which is inversely proportional to the number of integers you retain. If John’s age in days is amongst the integers you rule out, you receive a punishment instead of a prize. In example 5, there are ‘prizes’ and ‘punishments’. If the murderer is amongst the men you retain, you make it easier for the police to catch their man (and the help you give the police is inversely proportional to the number of men you rule out). If the murderer is amongst the men you rule out, you make it more likely that the murderer will go free.

In our conversation, there are ‘prizes’ and ‘punishments’. If the actual possibility is amongst the possibilities you retain, you will be in possession of a ‘prize’. For if you trust my assertion and update your beliefs so as to make them incompatible with every possibility which was ruled out from the context-set (but otherwise as conservatively as possible), you will have strengthened your beliefs without making them false. And the fewer possibilities you retain the more your beliefs will be strengthened, so your ‘prize’ will be inversely proportional to the number of possibilities you retain. On the other hand, if the actual possibility is amongst the possibilities you rule out, you receive a ‘punishment’ instead of a prize. For if you trust my

assertion, and update your beliefs so as to make them incompatible with every possibility which was ruled out from the context-set, you are sure to end up with false beliefs.

- You face a semi-principled decision in the course of our conversation, as you did in the toy examples. Although you will certainly treat my assertion as a ground for deciding how to adjust your conversational presuppositions, my assertion will not immediately put you in a position to decide which particular adjustment to carry out: secondary considerations will be necessary. Since our conversation presupposes nothing special about the colour of Susan's house, the context-set will contain possibilities ascribing a range of different colours to Susan's house. And when the range is sufficiently fine-grained, my assertion will be tolerant with respect to the possibilities in the context set. Accordingly, your linguistic competence with the sentence asserted will not be perfectly suited to the task at hand. You are able to use your competence to discriminate between sufficiently dissimilar possibilities (e.g. a possibility whereby Susan's house is a clear case of blue and a possibility whereby Susan's house is a clear case of white). But the task at hand calls for discriminations much finer than that. You are therefore unable to express your linguistic competence with the sentence asserted as a partition of possibilities in the context set.

(As emphasized above, what is at issue here is not whether the *content* of my assertion determines a particular partition of the possibilities in the context set. What is at issue is your *ability* to use my assertion to choose a particular partition, whatever the assertion's content.)

- IMPROVEMENT FROM LIMITATION holds in our conversation, as it did in the toy examples. Suppose our conversation takes place on Susan's street. We can both see that there are only two houses on the street, a blue one and a white one. It is common knowledge that this is so, and that one of the houses is Susan's. Then the context set can be expected to consist of only two possibilities: one according to which Susan lives in the blue house and one according to which Susan lives in the white house. When I assert 'Susan lives in a blue house' you will have no difficulty deciding which possibility to rule out on the basis of your semantic competence.

- FURTHER SECONDARY CONSIDERATIONS holds in our conversation, as it did in the toy examples. If being mistaken about the colour of Susan's house has serious consequences, you can be expected to be more conservative about ruling out possibilities than you otherwise might have been.
- SENSITIVITY TO EXPRESSIVE RESOURCES holds in our conversation, as it did in the toy examples. If, for example, you take me to have a very limited vocabulary, you can be expected to be more conservative about ruling out possibilities than you otherwise might have been.
- Assertions of 'Susan lives in a blue house' display INSTABILITY, just like the speech-acts of the toy examples. Consider two different scenarios. In Scenario A, the context-set consists of possibilities p_{BG} and p_{PG} (according to p_{BG} , Susan lives in a blueish-grey house; according to p_{PG} Susan lives in a pinkish-grey house). In Scenario B, the context-set consists of possibilities p_{BB} and p_{BG} (according to p_{BB} Susan lives in a bright blue house; p_{BG} is as before). In both scenarios I assert 'Susan lives in a blue house'. In Scenario A you retain p_{BG} and rule out p_{PG} ; in Scenario B, you rule out p_{BG} and retain p_{BB} , even though p_{BG} is retained in Scenario A on the basis of the same assertion.

According to a certain conception of linguistic competence, what it is to master a language is to gain cognitive access to *rules* corresponding to speech-acts of the language, and acquire the ability to use speech-acts by applying the corresponding rules. The analogy developed in the present section suggests that this picture is mistaken. Part of what it is to master a language is to acquire the ability to adjust one's conversational presuppositions on the basis of the speech-acts of one's interlocutors. But carrying out the relevant adjustments is not simply a matter of applying a rule: the subject must make a semi-principled decision. (This is *not* to say that speech-acts lack contents; it is only to say that language-mastery is not a matter of gaining access to the relevant contents and then adjusting one's presuppositions as the contents mandate.)

1.2.3 The Principle of Clarity

There are scenarios in which the tolerance of an assertion will interfere with your goals. Suppose Susan is having you over for lunch. You have not been to her home before, and she is giving you directions over the phone. After explaining how to get to her street, she says 'I live in a blue

house'. You walk to Susan's street, and discover that there are only two houses, a blueish-purple one and a blueish-green one. You find that you are unable to use Susan's assertion to come to a confident decision about which door-bell to ring.

In normal cases, however, tolerance will not get in the way. If, for example, you get to Susan's street and discover that there are only two houses, a blue one and a white one, then you will have no problem using Susan's assertion to make a decision about which door-bell to ring. And it is no accident that in normal cases tolerance does not get in the way. For in normal cases the speaker is cooperative, and will do her best to choose her assertion in such a way that it is clear to her audience how she proposes that it be used to carry out the task at hand. So if Susan lives on a street that has a blueish-purple house and blueish-green house, she can normally be counted on not to use 'I live in a blue house' to give you directions.

More generally, well-run conversations are governed by principles of cooperation. This means, among other things, that participants in a well-run conversation will conform to the following principle, and take for granted that others conform as well:

PRINCIPLE OF CLARITY

Make your assertions in such a way that it is clear to your audience how you propose that they be used to modify the conversation's presuppositions.

An immediate consequence of the Principle of Clarity is that one must refrain from making assertions that would be tolerant relative to the context set on which the assertion is to have its essential effect. For instance, Susan's assertion of 'I live in a blue house' will violate the Principle of Clarity if the possibilities in the context set on which the assertion is to have its essential effect transition smoothly from possibilities whereby Susan lives in a clear case of a blue house to possibilities whereby Susan lives in a clear case of a purple house. For it will be unclear to her audience how she proposes that the assertion be used to partition the context-set.

The Principle of Clarity demands that one refrain from making assertions that would be tolerant relative to the context set on which the assertion has its essential effect, but the context set on which the assertion has its essential effect need not be the context set which is in place at the time of the assertion. In cases in which Susan's assertion is tolerant relative to the context set which is in place at the time of the assertion, you will normally conclude that Susan's presuppositions are

different from yours, since you will be taking for granted that Susan conforms to the Principle of Clarity. So you will attempt to change your presuppositions so as to reduce discrepancies with Susan's. In particular, you will attempt to change your presuppositions in such a way that Susan turns out not to violate the Principle of Clarity. But since Susan is being cooperative, she will have only made the relevant assertion if she is confident that your attempt will be successful. So Susan's assertion can be expected to have its essential effect on a context set with respect to which there is no tolerance, even if the assertion is tolerant with respect to the context set which is in place at the time of the assertion. This pragmatic phenomenon is sometimes referred to as *accommodation*, following Lewis 1979.

The phenomenon of accommodation will play a crucial role in what follows, so it is worth considering a couple of examples. Here is the first.⁴ Susan asserts 'I live in a blue house' in a context in which nothing special is presupposed. It is therefore compatible with your presuppositions that Susan's house have any of a broad range of colours. Note, however, that if Susan's house was a borderline case of blue her usage of 'I live in a blue house' would have been somewhat unparadigmatic. But Susan did nothing to prevent her assertion from being interpreted in the most obvious way possible, even though she could have easily done so (e.g. by adding 'well, *sort of* blue'), and even though unparadigmatic usage was easily avoidable (she could have instead said, e.g. 'I live in a blueish house'). So, by taking for granted that Susan is being cooperative, you adjust your presuppositions in such a way that it ceases to be a live option that the colour of Susan's house is in the borderline range. As a result, the assertion will have its essential effect on a context set with respect to which there is no tolerance.

Now consider a second example. Susan asserts 'I live in a blue house' in a context in which it is presupposed that her assertion is intended to help you find her house, and in which nothing special is presupposed about her street. It is compatible with your presuppositions that Susan's street consists of a blueish-purple house and a blueish-green house, and that Susan's street consists of a blueish-grey house and a pinkish-grey house. If her street turns out to consist of a blueish-purple house and a blueish-green house, she will not have been very helpful. So, by taking for granted that Susan is being cooperative, you adjust your presuppositions in such a way that it ceases to be a live option that Susan's street consists of a blueish-purple house and a blueish-green house. On the other hand, you need not adjust your presuppositions in such a way

⁴My treatment of this example draws from discussion of related examples in Dorr 2003.

that it ceases to be live option that Susan's street consists of a blueish-grey house and a pinkish-grey house, since Susan will have been helpful in such a situation, even though she will have used 'I live in a blue house' unparadigmatically. You will continue to make adjustments until you feel confident enough that the assertion will have its essential effect on a context set with respect to which there is no tolerance.

It is important to keep in mind that it is *not* part of the story that the content of Susan's assertion mandates a particular adjustment of your presuppositions. Your decision about how to adjust your presuppositions is semi-principled, and is best thought of on the model of choosing an age-range in the toy examples of section 1.1. Although the decision-making process will certainly be guided by your linguistic competence, other factors will be relevant. If you take Susan's vocabulary to be very restricted, or you think she is too tired to give you a careful answer, your presuppositions might be adjusted in an especially conservative way. More generally, you will have to balance off your desire to presuppose as little as possible (and thereby extract as much information as you can from Susan's assertion) with your desire to minimize the risk of excluding the live option that Susan regards as actual (and thereby minimize the risk of being misled). In typical cases, you will find that there is no single choice that stands out to you as best all things considered. So your final decision can be expected to involve a degree of arbitrariness. If there is a particular choice that is in some sense *mandated* by the content of Susan's assertion, this is not something you have access to in the present context. As more and more presuppositions are made, you will grow more and more confident. But you can be expected to lack a principled criterion for deciding when such confidence is confidence enough.

1.2.4 Instrumentalism

The Stalnakerian picture of conversational pragmatics that I have been presupposing might be thought to face an awkward question. Suppose Susan gives you directions by asserting 'I live in a blue house'. Let p_0 be the proposition that Susan's street consists of a house of some particular shade of blueish-purple and a house of some particular shade of blueish-green, let p_{100} be the proposition that Susan's street consists of a house of some particular shade of paradigmatic blue and a house of some particular shade of paradigmatic green, and for $0 < n < 100$ let the p_n be suitably intermediate propositions. Let it be granted that you presuppose that p_0 is not the case but treat p_{100} as a live option, and let it be granted that for any n between 1 and 100 you will only presuppose

that p_n is not the case if you also presuppose that p_{n-1} is not the case. Which is the largest k such that you presuppose that p_k is not the case?

If one thinks of the Stalnakerian picture non-instrumentally, this will indeed be an awkward question. For choosing any particular k leaves the theorist with an explanatory burden: the task of explaining why she chose that particular k rather than one of its immediate neighbors. And from the perspective of a non-instrumentalist, it is not clear that one would be in a position to supply the needed explanation. I would like to suggest, however, that the Stalnakerian machinery is best thought of instrumentally: it is best thought of as a *tool* for making sense of our linguistic practice. Accordingly, there is no sense to be made of the question whether a given proposition is *really* presupposed at a given stage in a conversation, over and above the question whether describing the conversation in terms of that presupposition is a good way of making sense of the relevant linguistic practice.

From this instrumentalist perspective, the question of which is the largest k such that you presuppose that p_k is not the case only makes sense if it is read as asking about the largest k such that it would be useful to describe the relevant behaviour in terms of the presupposition that p_k is not the case. And, so understood, it can be addressed as follows. The ascription of a presupposition is a way of conveying information about the dispositions of speakers. But speakers' dispositions are usefully described in terms of a presupposition only to the extent that they form a pattern of the right kind. Suppose that, upon discovering that Susan's street consists of a blueish-purple house and a blueish-green house, you are disposed to say: 'Wait a minute! How am I supposed to find Susan's house then?' If you are disposed to do suitably related things in suitably related circumstances, your linguistic practice will be usefully described by saying that you presuppose that p_0 is not the case. But as n grows larger, the relevant patterns will begin to dissipate and it will become less and less helpful to describe your behaviour in terms of p_n . Where to draw the line between patterns that are robust enough to be usefully described in terms presuppositions and those that are not is a *practical issue*, which is up to the theorist. If she chooses to ascribe presuppositions only when the relevant dispositional pattern is highly robust, a description in terms of presuppositions will convey a lot of information, but her account of linguistic practice will cover very little ground. If, on the other hand, she is prepared to ascribe presuppositions to weaker patterns, a description in terms of presuppositions will convey less information, but her account of linguistic practice will cover more ground. In all likelihood, she will lack a principled criterion

for striking a balance between these different desiderata, and her final decision will involve some degree of arbitrary choice.

Someone might worry that the instrumentalist is leaving something out, on the grounds that nothing has been done to ensure that our ascriptions of presuppositions track the presence of suitable *rules* in our cognitive system. Such a complaint is based on the idea that to make a conversational presupposition is to gain cognitive access to a rule, and resolve apply the rule in the right sorts of ways.

The view of linguistic practice I have developed in earlier sections suggests that such a conception of presupposition is mistaken. On the view I have been developing, part of what it is to master a language is to acquire the ability to use one's interlocutor's speech-acts (together with contextual cues) to negotiate one's options in the pursuit of one's goals. When Susan gives you directions by saying 'I live in a blue house', for instance, you are able to use her assertion (together with contextual cues) to make decisions that would be sensible on the assumption that it was common knowledge that Susan's street does not consist of a blueish-purple house and a blueish-green house. But in putting Susan's assertion to use in this way, you are not simply following a rule: you are involved in semi-principled decision-making. Accordingly, to say that you have made the conversational presupposition that *p* is not to say that you have gained cognitive access to a rule corresponding to *p*, and have resolved to follow it in the right sort of way. It is to say that, in light of your linguistic competence, you are disposed to negotiate your options in ways that form a certain kind of pattern, namely: your various decisions would be sensible on the assumption that *p* was common knowledge. How robust must the pattern be before it can be described as a presupposition? According to the instrumentalist one can use whichever conception of robustness one likes, as long as it delivers a useful way of making-sense of linguistic practice.

2. Content

In section 1, I discussed linguistic usage. Nothing was said about the *contents* of our speech-acts. (In particular, nothing was said about the information carried by a speech-act, or about whether a particular modification of the context-set is *mandated* or *permitted* by an assertion's content.) The purpose of this section is to explain how content comes into the picture.

2.1 Theories of assertoric content

As I shall understand it here, a theory of assertoric content is a tool for understanding the evolution of the context set throughout a conversation. It works by ‘predicting’ that the possibilities excluded by an assertion’s essential effect will be those of the possibilities in the context set that are incompatible with the assertion’s content. Accordingly, it is successful to the extent that its predictions match the actual evolution of the context set.

A theory of assertoric content is not our only tool for understanding the evolution of the context set. It can be expected to work as one of several components in a larger explanatory system. Such a system will presumably include a theory of conversational dynamics that specifies norms of conversational cooperation such as the Principle of Clarity (section 1.2.3) and yields an account of conversational implicature (see Grice 1989a). It can also be expected to include a theory of figurative speech and an account of the mechanisms whereby the context set will evolve in cases where the assertion has different contents relative to different live options (see Stalnaker 1979). Although the focus of this section will be on theories of assertoric content, it is important to be clear that the success of a theory of assertoric content should always be assessed on the basis of its role in the larger explanatory system. Accordingly, an assertion is used *correctly* (by the lights of a given theory of assertoric content) just in case its essential effect is as predicted by the larger explanatory system of which one’s theory of assertoric content is a part. (In the special case in which pragmatic considerations play no role, an assertion is used correctly by the lights of a theory of assertoric content just in case the possibilities excluded by the assertion’s essential effect are those of the possibilities in the context set that are incompatible with the content that the assertion is assigned by the theory.)

A crucial component of a theory of assertoric content is a *semantic theory*. As I shall understand it here, a semantic theory is a compositional assignment of propositions to sentence-context pairs. One may think of propositions however one likes for present purposes, as long as a proposition determines a definite partition of the space of possibilities into two exhaustive and exclusive groups: possibilities that are compatible with the proposition, and possibilities that are incompatible with the proposition.

In the next couple of sections I will consider two different ways in which semantic theories might be put to use by a theory of assertoric content.

2.1.1 Globalism

Our assertions are typically tolerant with respect to the entire space of worlds. Because of the Principle of Clarity, this means that they will typically be used *locally*: they will typically have their essential effects on context sets that are highly restricted in the possibilities they contain. It is none the less tempting to suppose that the various instances of local usage fall into patterns that are usefully described in terms *global* contents: contents that determine a definite partition of the entire space of possibilities. Accordingly, it is tempting to think of the task of constructing a theory of assertoric content as the task of finding an assignment of global contents that fits local usage as neatly as possible: it is a matter of piecing together the multiple instances of local usage into unified partitions of the space of possibilities, in much the way that multiple satellite-photographs are pieced together into a unified picture of the surface of the Earth.

The simplest version of this globalist approach proceeds in two steps. The first is to choose a context-insensitive semantic theory (I ignore indexicals to keep things simple.) The second step is to identify the global content of an assertion with the proposition assigned to the sentence asserted by one's preferred semantic theory. The task of constructing a theory of assertoric content therefore boils down to the task of finding an assignment of semantic values to basic lexical items which yields the result that actual instances of usage are reliably correct.

The simple approach suffers from two important limitations. The first is that it does not, by itself, give us a way of predicting when an assertion will be tolerant relative to the context set which is in place at the time of the assertion (and thereby force an accommodation). Say that the semantic value of 'blue' is characterized in terms of hues, saturations and luminosities. To keep things simple, suppose that 'blue' applies to all and only objects with hues between 210 and 270 for some fixed saturation and luminosity,⁵ and that sentences involving 'blue' are assigned global contents accordingly. Now consider a context in which, for every k between 1 and 360, p_k is treated as a live option, where p_k is the possibility that Susan lives in a house with hue k . On the simple globalist approach, the global content of 'Susan lives in a blue house' is compatible with p_k for k within 210–270 and incompatible with p_k for k outside this range. So the simple approach supplies no way of resisting the conclusion that you will use an assertion of 'Susan lives in a blue house' to rule out every p_k with k outside 210–270 from the context set,

⁵ I assume a 0–360 hue scale in which 0 (= 360) is red, 60 is yellow, 120 is green, 180 is cyan, 240 is blue, and 300 is violet.

and keep the rest. And, of course, this is not what will actually happen. In the absence of special contextual cues, an assertion of ‘Susan lives in a blue house’ will be tolerant relative to the p_k . In light of the Principle of Clarity, the assertion should therefore be expected to have its essential effect on an accommodated context set. (See section 1.2.3.)

The difficulty might be addressed by claiming that the contents of sentences partition the space of possibilities in three: possibilities that are (definitely) compatible with the content, possibilities that are (definitely) incompatible with the content, and possibilities that are neither. The suggestion would then be that accommodation is explained by the fact that ‘neither’ possibilities must be excluded from the context set before an assertion can have its essential effect. (Alternatively, the globalist might shift the explanatory burden from semantics to pragmatics, and argue that the practice of assertion is subject to a Margin for Error Principle according to which possibilities that are ‘too close’ to the border between compatible possibilities and incompatible possibilities should be excluded from the context set before an assertion can have its essential effect.) We will return to this sort of response in section 2.2.2.

A second limitation of the simple approach is that, without supplementation, it will only issue reliable predictions about the evolution of the context set in the presence of a stable boundary between possibilities that are ruled out when a particular sentence is asserted and possibilities that are not. But the phenomenon of instability discussed in section 1 makes clear that there is no such boundary. Here is an example (adapted from Graff 2000). In scenario A it is common knowledge that Susan’s street consists of a house painted blueish-grey and a house painted pinkish-grey, and you use my assertion of ‘Susan lives in the blue house’ to rule out possibilities whereby Susan lives in the pinkish-grey house. In scenario B it is common knowledge that Susan’s street consists of a house painted blueish-grey and a house painted bright blue, and you use my assertion of ‘Susan lives in the blue house’ to rule out possibilities whereby Susan lives in the blueish-grey house. So the possibility that Susan lives in the blueish-grey house cannot be stably classified either as a possibility that should be ruled out on the basis of assertions of ‘Susan lives in the blue house’ or as one that should not.

A globalist might address the phenomenon of instability by complicating the simple proposal. It might be suggested that the semantic value of an expression like ‘blue’ is subject to a contextual parameter, with the result that my assertion of ‘Susan lives in a blue house’ has different contents in scenarios A and B. (Alternatively, the globalist might

shift the explanatory burden from semantics to pragmatics, and claim that, even though my assertion *says* the same thing in each of the two scenarios, there is a difference in what is communicated.) We will return this sort of response in section 2.2.1.

2.1.2 A localist alternative

On a globalist theory of assertoric content, one predicts the evolution of the context set by determining which of the possibilities in the context set are compatible with the global content of the sentence asserted. On the localist alternative I would like to consider in this section, one predicts the evolution of the context set in a different sort of way.

The intuitive picture is this. Linguistic communication is a matter of partitioning sets of possibilities in ways that are rendered *salient* by prior linguistic usage in the context of the relevant speech-act. But the information about previous usage that is accessible to speakers is very rough, so successful communication requires that there be a *gap* in the possibilities under consideration: they must fall into classes that are distinct enough from each other for the rough information about prior usage that is accessible to speakers to be enough for rendering some partition uniquely salient in the context of the speech-act. (When there is not enough of a gap, the Principle of Clarity is violated.) The role of gaps in rendering a partition salient yields the result that which partition is rendered salient depends not just on how the expressions in the sentence asserted have been used in the past, but also on the location of the gap. Simple globalists characterize a notion of assertoric content that is not sensitive to the location of the gap: they characterize a notion of content on the basis of previous usage alone. This yields the result that the content of an assertion determines partitions of arbitrary sets of possibilities, but it is not a very good way of predicting the evolution of the context set. Localists, on the other hand, characterize a notion of assertoric content that is highly sensitive to the location of the gap. This will not yield the result that the content of an assertion determines partitions of arbitrary sets of possibilities—assertoric content will turn out to be *local* to a particular set of possibilities—but it does deliver a better way of predicting the evolution of the context set.

The purpose of this section is to suggest a way in which this intuitive idea might be spelled-out. I will proceed in two stages. The first is to identify suitable patterns of prior linguistic usage; the second is to characterize a notion of salience on the basis of these patterns. With respect to the first stage, the localist's crucial observation is that even though no particular semantic theory yields a very satisfactory means for predict-

ing the evolution of the context set, different semantic theories are successful to different degrees, and the differences supply valuable information about patterns of prior linguistic usage.

A first pass at quantifying the success of a semantic theory is this. Say that an assertion is *correct* by the lights of a semantic theory α just in case the possibilities ruled out by the assertion's essential effect are those of the possibilities in the context set that are incompatible with the proposition assigned by α to the sentence asserted. (I assume, for simplicity, that the semantic theories in question are context-invariant.) One can then say that a semantic theory α is *successful* to degree δ just in case δ is the percentage of actual assertions that are correct by the lights of α . Here is an example. Let α_n be a semantic theory assigning 'Susan lives in a blue house' a proposition compatible with all and only possibilities in which Susan lives in a house with hue between 220 and n . Then one would expect the success rate of the α_n with respect to assertions of 'Susan lives in a blue house' to be as depicted in figure 1. Because no provisions are made for accommodation, and because there is no stable boundary between possibilities that assertions of 'Susan lives in a blue house' are used to ruled out and the rest, none of the α_n can be expected to be very successful. But the variation of success rates amongst the α_n conveys substantial information about patterns of prior linguistic usage.

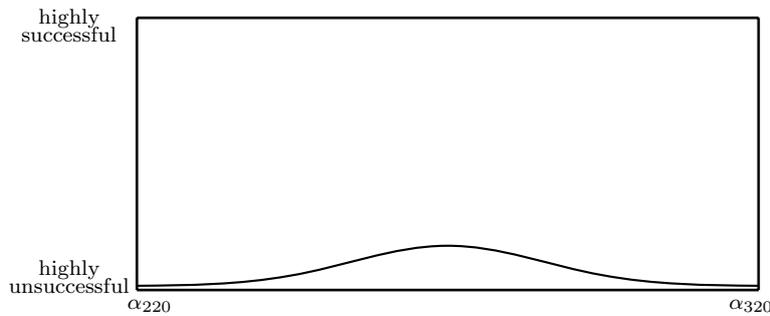


Figure 1: Expected success rate of the α_n with respect to assertions of 'Susan lives in a blue house'

The next step is to use this information to characterize a notion of salience. A straightforward way of doing so is as follows. Say that a *partition* of a context set is a classification of the possibilities in the context set into two classes: possibilities that should be ruled out and possibili-

ties that should not. One can then say that a partition P is rendered *salient* by an assertion just in case the following two conditions are met:

- (1) P is correct by the lights of a significant proportion of those semantic theories that are relatively successful when compared to their peers.⁶
- (2) P is the only partition of C satisfying condition 1 that has not been dismissed on independent grounds.

It is worth illustrating the resulting notion with some examples. In each case, I consider an assertion of ‘Susan lives in a blue house’ and assume that only *non-trivial* partitions are candidates for salience. (Trivial partitions—partitions whereby every possibility is to be ruled out or every possibility is to be retained—can typically be dismissed on the basis of purely pragmatic considerations; see Stalnaker 1979.)

- *An easy case*

Let the context set C consist of possibilities p_B and p_V (according to p_B , Susan lives in a blue house; according to p_V Susan lives in a violet house). Let P be the partition of C whereby only p_V is ruled out.

The intuitive picture is this: there is a large gap between p_B and p_V , and the gap is so positioned that P is rendered uniquely salient by the assertion in light of prior linguistic usage. To see that the proposed definition of salience vindicates this picture, note that although semantic theories that are successful when compared with their peers often disagree about the proposition assigned to ‘Susan lives in a blue house’, there is broad agreement about the correctness of P . The only other non-trivial partition of C is the partition whereby only p_B is ruled out. Such a partition receives little support amongst semantic theories that are relatively successful when compared to their peers. So P is rendered uniquely salient by the assertion.

⁶ *Notation:* partition P of context set C is *correct* (for sentence S) by the lights of semantic theory α just in case an assertion of S whose essential effect consists in modifying C in accordance with P is used correctly by the lights of α ; semantic theories are *peers* if they agree with respect to the semantic values of lexical items not occurring in the sentence asserted.

The situation is illustrated in figure 2. Points within the main box represent semantic theories, with darkness indicating their level of success relative to their peers. The position of a semantic theory on the horizontal axis represents the place at which it locates the cutoff-point between blue and violet; its position on the vertical axis represents the place at which it locates the cutoff-point between saturated colours and greys. (More specifically, the semantic theory with coordinates $\langle i, j \rangle$ assigns ‘Susan lives in a blue house’ a proposition compatible with all and only possibilities in which Susan lives in a house with hue between 220 and i and saturation between j and 100.) Circles represent possibilities in C , with the position of a circle on the horizontal and vertical axes corresponding, respectively, to the hue and saturation of Susan’s house according to the possibility represented.

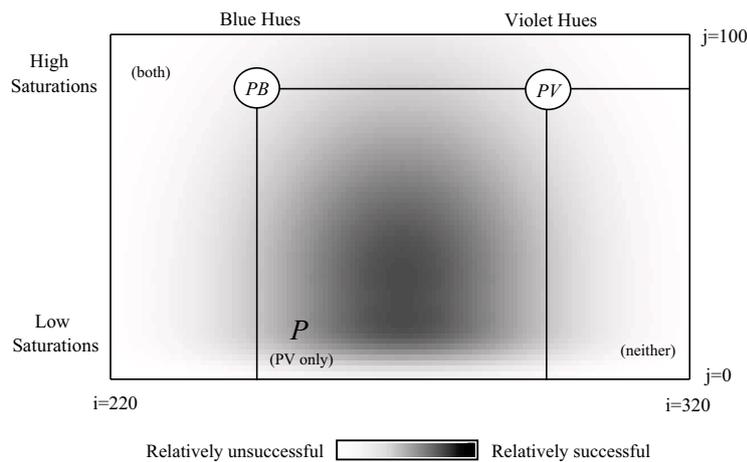


Figure 2: A large and well-placed gap

The straight lines in the diagram do *not* represent partitions of C . Their role is to divide the space of semantic theories into *cells* in such a way that semantic theories placed in the same cell favour the same partition of C . (To determine which partition is favoured by a given semantic theory consider the rectangle whose lower right corner is the point representing the semantic theory and upper left corner is the upper left corner of the diagram. Circles with centres within the rectangle represent possibilities retained by the favoured partition; circles with centres outside the rectangle represent possibilities ruled out.) Accordingly,

semantic theories in the cell labelled ‘(both)’ favour the partition whereby both p_B and p_V get ruled out; semantic theories in the cell labelled ‘(p_V only)’ favour partition P , whereby only p_V is ruled out; and semantic theories in the cell labelled ‘(neither)’ favour the partition whereby neither p_B nor p_V gets ruled out. (Semantic theories favouring the partition whereby only p_B is ruled out cannot be represented in this diagram, but their absence is harmless because they are all highly unsuccessful relative to their peers.) Finally, each cell in the diagram may be thought of as representing a different partition of C —the partition that is favoured by semantic theories within that cell.

Here is how to use the diagram (and those to follow). Single out the cells that contain a significant proportion of the relatively successful semantic theories. If exactly one of them consists of semantic theories favouring a non-trivial partition of C , conclude that that partition is rendered uniquely salient by the assertion. Otherwise, conclude that no partition of C is rendered salient. (In the case of fig. 2, the only cell containing a significant proportion of the relatively successful semantic theories is the cell labelled ‘(p_V only)’, so we may conclude that partition P is rendered uniquely salient by the assertion.)

- *A narrower gap*

We have considered a case in which the gap between possibilities in the context set is large and ideally placed. But a smaller and less well-placed gap would have been sufficient to render a partition uniquely salient. Let C consist of possibilities p_B and p_{BV} (according to p_{BV} Susan lives in a blueish-violet house; p_B is as before). Let P be the partition of C whereby only p_{BV} is ruled out.

The intuitive picture is this. There is a gap between p_B and p_{BV} . It is so positioned that P might have competed for salience with the rival partition P^* whereby every possibility is retained. But P^* can be dismissed on account of its triviality. So P is rendered uniquely salient by the assertion in light of prior linguistic usage.

To see that the proposed definition of salience vindicates this intuitive picture, note that semantic theories that are relatively successful when compared to their peers can be expected to fall into one of two categories. According to those in the laxer category, the proposition assigned to ‘Susan lives in a blue house’ is compatible with both p_B and p_{BV} ; so P^* is favoured. According

to semantic theories in the stricter category the proposition assigned to ‘Susan lives in a blue house’ is compatible with p_B but incompatible with p_{BV} ; so P is favoured. But P^* is trivial. So P is rendered uniquely salient amongst non-trivial partitions.

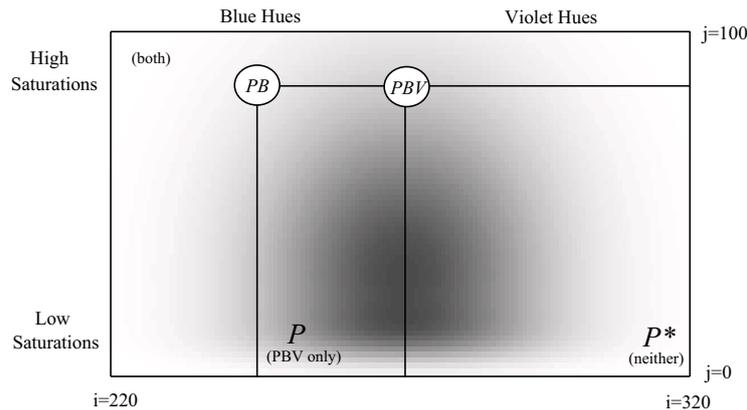


Figure 3: A smaller and less well-placed gap

The situation is illustrated in figure 3: amongst the cells containing a significant proportion of the relatively successful semantic theories, only the cell labelled ‘(p_{BV} only)’ corresponds to a non-trivial partition. So we may conclude that partition P is rendered uniquely salient by the assertion. (Semantic theories favouring the partition whereby only p_B is ruled out cannot be represented in this diagram, but, as before, their absence is harmless because they are all highly unsuccessful relative to their peers.)

- *An ill-placed gap*

I would now like to consider a case in which there is a significant gap amongst possibilities in the context-set, but the gap does not yield the result that a partition is rendered uniquely salient by the assertion. Let the context set C consist of possibilities p_{BV} and p_{BG} (p_{BV} is as before; according to p_{BG} Susan lives in a blueish-grey house). Let P be the partition of C whereby only p_{BV} is ruled out, and P^* be the partition of C whereby only p_{BG} is ruled out. Each of these partitions is favoured by a significant proportion of relatively successful semantic theories. So neither of them is rendered uniquely salient.

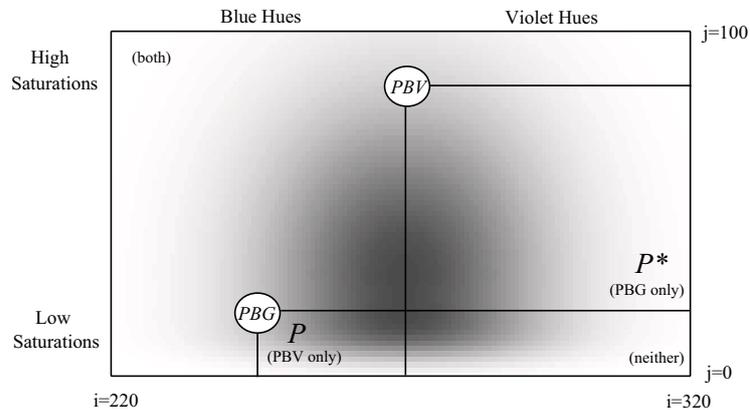


Figure 4: An ill-placed gap

The situation is illustrated in figure 4: amongst the cells containing a significant proportion of the relatively successful semantic theories, more than one corresponds to a non-trivial partition. This means that there are different ways of using the assertion to partition the context set, each of them sensible in light of previous usage. So the Principle of Clarity is violated.

- *Not enough of a gap*

Now consider a case in which there is not enough of a gap between the possibilities in the context set. Let the possibilities in C transition smoothly from possibilities whereby Susan lives in a clear case of a blue house to possibilities whereby Susan lives in a clear case of a violet house. There will be broad disagreement amongst relatively successful semantic theories about which partition is correct. So none of the possible partitions is rendered uniquely salient.

The situation is illustrated by figure 5: there are many cells corresponding to non-trivial partitions and containing significant proportions of the relatively successful theories.

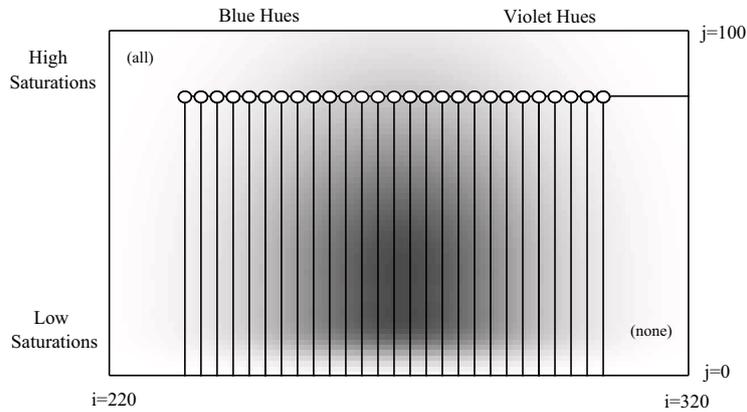


Figure 5: Not enough of a gap

- *Instability*

Finally, it is worth considering an example of instability. In scenario 1 the context set C_1 consists of possibilities p_{BG} and p_{VG} (according to p_{VG} Susan lives in a violetish-grey house; p_{BG} is as before). In scenario 2 the context set C_2 consists of possibilities p_B and p_{BG} (both as before).

The situation is illustrated in figure 6 (see overleaf). In the first diagram, p_{BG} is retained by the sole non-trivial partition with a significant proportion of the relatively successful semantic theories; in the second diagram that same possibility is ruled out by the sole non-trivial partition with a significant proportion of the relatively successful semantic theories. This is explained by the fact that there is a big difference in the location of the gap. (Semantic theories favouring the partition whereby only p_{BG} is ruled out cannot be represented in the first diagram; semantic theories favouring the partition whereby only p_B is ruled out can not be represented in the second diagram. As above, their absence is harmless because they are all highly unsuccessful relative to their peers.)

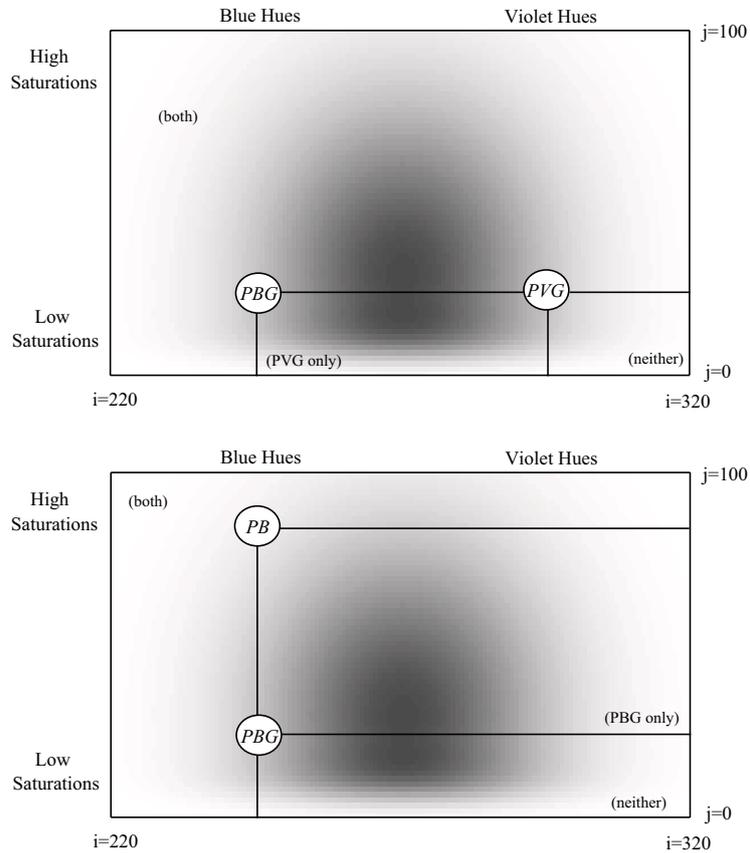


Figure 6: Instability

2.1.3 The localist theory at work

In section 2.1.1 we identified two limitations of the simple globalist account. The first is that simple globalism does not, by itself, supply a means for predicting when an assertion will force accommodation. From our present perspective, on the other hand, there is a straightforward story to tell. Say that an assertion is *defective* just in case it fails to render salient a partition of the context set on which it is to have its essential effect. The localist is in a position to make predictions about accommodation by way of the following principle:

First Principle

Speakers should accommodate to avert defectiveness.

Suppose I assert ‘Susan lives in a blue house’. Nothing special is presupposed about the colour of Susan’s house. Accordingly, the possibilities in the context set which is in place at the time of the assertion are smoothly distributed across the colour spectrum. The localist’s account of salience yields the result that no partition is rendered salient by my assertion because there is not enough of a gap in the possibilities under consideration. So defectiveness threatens, and the localist’s first principle predicts that speakers will revise their presuppositions so that defectiveness is averted.

How exactly should presuppositions be revised? The localist’s first principle is silent on this matter—all it requires is that defectiveness be averted somehow or other. But there is a further principle that immediately suggests itself:

Second Principle

Accommodation should be as conservative as possible.

In the case at hand, the context set to be accommodated consists of possibilities smoothly distributed across the colour spectrum. So the only way of averting defectiveness is by creating a gap. In the absence of special contextual cues, the obvious way of doing so is by making the assumption that the assertion has not been used unparadigmatically, and eliminating possibilities whose compatibility with the content of ‘Susan lives in a blue house’ is a matter of dispute amongst semantic theories that are relatively successful when compared to their peers. The second principle will then recommend eliminating no more of them than is needed to avert defectiveness.

The second principle yields substantial information about the sort of accommodation one should expect. But it will typically not put one in a position to specify a unique revision of the context set. One reason this is so is that different but equally conservative revisions might succeed in averting defectiveness. Another is that the localist’s notion of salience will not always determine whether a particular revision succeeds in averting defectiveness, since salience was characterized in terms of vague expressions such as ‘significant’ and ‘relatively successful’. One should be suspicious of theories of assertoric content *not* featuring this lack of specificity. By issuing overly specific predictions about the evolution of the context set, a theory of assertoric content would indicate the presence of robust patterns of linguistic practice where there are none.

The second limitation we identified in the simple globalist account is that reliable predictions about the evolution of the context set will only be issued in the presence of a stable boundary between possibilities that

are ruled out when a given sentence is asserted and possibilities that are not. From our present perspective, on the other hand, there is a straightforward way of issuing reliable predictions in the absence of stable boundaries. Let the *local content* of an assertion be a function from possibilities in the context set on which the assertion has its essential effect to truth-values: a possibility gets the value ‘F’ if it is ruled out by the partition that is rendered salient by the assertion, and the value ‘T’ otherwise. (Local content is undefined if the assertion is defective.) The localist can then predict the evolution of the context set by way of the following principle:

Third Principle

The possibilities excluded by an assertion’s essential effect should be those of the possibilities in the context set that are assigned the value ‘F’ by the assertion’s local content.

(That this principle delivers the desired predictions can be illustrated by revisiting the examples in section 2.1.2.)

Before bringing this section to a close, I would like to make four additional observations about the localist approach.

- Truth and falsity

The localist is in a position to give straightforward characterisations of assertoric truth and falsity:

If the context set on which a non-defective assertion is to have its essential effect contains the actualized possibility, then the assertion is *true* if the actualized possibility is assigned the value ‘T’ by the assertion’s local content; otherwise it is *false*.

If none of the possibilities in the context set on which a non-defective assertion is to have its essential effect is actualized, then the assertion is said to be an instance of *presupposition failure*.⁷

A defective assertion might be said to be neither true nor false (though, I am not sure there is much that hinges on the matter).

The notion of sentential truth can then be characterized derivatively: a sentence is true relative to a context just in case an as-

⁷ Steve Yablo’s unpublished work on semantic arithmetic might supply a good way of characterizing true and false instances of presupposition failure.

servation of that sentence in that context would be true. (And, of course, the contextual parameter might be left tacit: when a sentence is true relative to a context that is suitably salient, the sentence might be said to be true *simpliciter*.)

A notion of sentential content can also be characterized derivatively: a sentence's (local) content relative to a context is the (local) content of an assertion of that sentence in that context. Relative to an assignment of local contents to sentences one could, if one wanted, say that an assignment of semantic-values to subsentential expressions is adequate just in case it determines an assignment of propositions to sentences that is consistent with the local content of every sentence. But, of course, one should not generally expect to get the result that there is a unique assignment of semantic-values that is counted as adequate.

- Logic

For a sentence to be *logically true* is for its truth to be guaranteed by its semantic structure together with the meanings of the logical terms involved. Here is one way in which the localist might cash out this idea. Say that a semantic theory is *logically adequate* if it assigns the standard semantic-values to the logical terms. (One could, if one wanted, give a substantive account of what makes the value of a logical term 'standard'; one might claim, for example, that the standard values are those amongst the values satisfying a logicity constraint⁸ that most reliably yield successful semantic theories.) Logical truth can then be characterized as follows:

A sentence is logically true just in case it is assigned a proposition compatible with every possibility by any logically adequate semantic theory.

On reasonable assumptions, this yields the result that all and only classical validities are counted as logically true.

The localist can go on to say that an *assertion* is logically true if the sentence asserted is logically true. As long as semantic theories that are relatively successful when compared to their peers turn out to be logically adequate, this will deliver the re-

⁸ An example of a logicity constraint is invariance under permutation—see Tarski 1986, Sher 1991, and McGee 1996.

sult that logically true assertions are guaranteed to be true independently of the context in which they are asserted and independently of the specific patterns of usage of the non-logical terms involved.

An analogous story can be told about logical consequence.

- Pragmatics

In section 2.1 I suggested that a theory of assertoric content should not be thought of as predicting the evolution of the context set on its own. It should be seen as part of a larger explanatory system, which includes a theory of conversational pragmatics. This complication was ignored in the preceding section to simplify the exposition.

A more detailed account of the localist position should bring the additional considerations to bear in at least two places. First, the success of a semantic theory should be assessed against the background of the larger explanatory system. Second, the localist's three principles should be thought of as predicting the evolution of the context set in conjunction with other components of the larger explanatory system.

- Status of the definitions

The definitions of success and salience I introduced on the localist's behalf should not be thought of as constitutive of the proposal. They are meant to illustrate the general shape that a localist account of assertoric content might take. What is constitutive of the localist proposal is that it deliver a way of combining information about past linguistic usage with information about gaps in the context set to determine which partition (if any) should be rendered salient by an assertion. The localist should use whichever notion of salience works best as a tool for understanding the evolution of the context set. Whether or not the specific characterization of salience I have suggested here would maximize empirical adequacy is an empirical question, and calls for an investigation that is beyond the scope of this paper.

2.1.4 Higher-order vagueness

The localist's notion of salience is vague, having been characterized in terms of vague expressions such as 'significant' and 'relatively successful'. This raises the question of whether there is a sharp boundary

between sets of possibilities with respect to which a given assertion renders some partition salient, and sets of possibilities with respect to which no partition is rendered salient. If there is a sharp boundary, then one would like to have an explanation of why the boundary is not at some very proximal but distinct location. If there is no such boundary, then one would like to have some account of the nature of the transition between sets of possibilities with respect to which some partition is rendered salient by a given assertion and sets of possibilities with respect to which no partition is rendered salient by the assertion. The problem of higher-order vagueness, as it arises for the localist, is the problem of addressing this dilemma.

The dilemma might well constitute a source of concern if one thinks of theories of assertoric content non-instrumentally. But here we are thinking of a theory of assertoric content as a *tool* for predicting the evolution of the context set. Accordingly, there is no sense to be made of the question whether a partition is *really* salient, over and above the question whether treating the partition as salient is a good way of making sense of our linguistic practice (and therefore no sense to be made of the question whether an assertion *really* has a particular local content, over and above the question whether ascribing that content to that assertion would be a good way of making sense of linguistic practice).

When matters are thought of in this way, the dilemma's second horn can be addressed as follows. The claim that a certain partition is salient is a device for issuing predictions. In clear cases of salience or clear cases of non-salience, the localist will be in a position to issue reliable predictions. But as one ventures into grey areas, the reliability of the localist's predictions will begin to drop. Where to draw the line between cases of salience or non-salience that are clear enough for predictions to be issued and those that are not depends on one's theoretical aims. If one chooses to issue predictions only when salience or non-salience is very clear, the reliability of one's predictions will increase, but one will often fail to make predictions. If, on the other hand, one is prepared to issue predictions in murkier cases, one's reliability will drop, but predictions will be issued more often. In the end, it is up to the theorist to decide how to proceed, and her decision can be expected to involve some degree of arbitrary choice.

If one thought that understanding an assertion was a matter of gaining cognitive access to a rule corresponding to the assertion's content, one might worry that the instrumentalist is leaving something out, since nothing has been done to ensure that our ascriptions of content track the presence of the relevant rules in our cognitive system. But the

picture of language I have been developing suggests that such a conception of linguistic understanding is mistaken. Understanding an assertion is not a matter of applying a rule; it is a matter of making sensible semi-principled decisions about how to partition the context set in light of past linguistic usage and the location of the gap.

2.1.5 *The Sorites Paradox*

Here is an instance of the Sorites Paradox:

For n between 0 and 200,000, man n has n hairs on his head. (The men are otherwise very much alike, and hair distributions are normal.) It is tempting to think that each of the following is true:

(S₁) Man 0 is bald

(S₂) For each n between 0 and 199,999, if man n is bald, then man $n + 1$ is bald

(S₃) Man 200,000 is not bald

But a contradiction can be derived from (S₁)–(S₃).

From a localist perspective, the contradiction is averted because (S₂) fails to be true. The easiest way to see this is to note that the semantic theories I used to characterize the notion of salience are all *classical*. Since (S₁)–(S₃) are classically inconsistent, any semantic theory whereby the propositions assigned to (S₁) and (S₃) are compatible with the actualized possibility will be a semantic theory whereby the proposition assigned to (S₂) is incompatible with the actualized possibility. But any semantic theory that is relatively successful when compared to its peers can be expected to assign (S₁) and (S₃) propositions that are compatible with the actualized possibility. Enter the proposed characterization of salience, and one gets the result that assertions of (S₂) will be false if non-defective, and that assertions of (S₂)'s negation will be true if non-defective.

The localist and the classical epistemicist both deny that (S₂) is true. But their denials amount to very different things. From the epistemicist's denial one is entitled to conclude that there is a cutoff point for 'bald': an n such that 'bald' applies to man n but not man $n + 1$. This is because the classical epistemicist believes that there is a 'correct' semantic theory: a semantic theory that assigns each expression in the language its 'true' semantic content. From the epistemicist's denial of (S₂) one can therefore conclude that for one of the men in the series, the 'true' semantic content of 'bald' applies to him but not his successor. No

such inference is warranted from the localist's perspective. For the localist, the untruth of (S₂) is no more than an artefact of the use of classical semantic theories, and of the particular way in which the notion of salience has been cashed out. The localist thinks that there is no reason to expect that any particular semantic theory will count as 'correct'. The role of semantic theories in the localist framework is to supply the information about patterns of past linguistic usage that is needed to characterize the notion of salience, which is in turn used to characterize the notion of (local) content. If it turns out that—contrary to expectations—the local contents of sentences involving 'bald' can be pieced together into a single assignment of global contents, one might be able to find a semantic theory that assigns each sentence in the language the relevant global content. Under such extraordinary circumstances the localist might think of the chosen semantic theory as 'correct', and see it as assigning 'bald' its 'true' semantic-value. The localist would then be happy to say that linguistic usage determines a cutoff-point for 'bald'. But nothing short of that will do.

There is a certain resemblance between the reason localists deny that (S₂) is true and the reason classical supervaluationists deny that (S₂) is true. On both proposals, there is a link between a sentence's truth-status and the truth-value it is assigned by the semantic theories in a certain class, and on both proposals the semantic theories in question are classical, with the result that (S₁)–(S₃) cannot all be assigned the value 'true'. The resemblance is imperfect, since there is no notion of admissibility with respect to which a localist could say that an assertion is true if it is true by the lights of every admissible semantic theory. The resemblance is also superficial, since localists and supervaluationists disagree about the nature of the link between a sentence's truth-status and the truth-value it is assigned by semantic-theories in the relevant class. For the supervaluationist, a semantic theory is counted as 'admissible' just in case it respects everything that is semantically determinate about the language. Accordingly, when the supervaluationist says that a sentence is true just in case it is true on every admissible semantic theory she says that a sentence is true if it turns out to be true no matter how the indeterminate features of the language are rendered determinate. This means, in particular, that the supervaluationist's claim that (S₂)'s negation is true warrants the conclusion that there will be a cutoff point for 'bald' no matter how the indeterminate features of the language are rendered determinate. Localists think of the matter very differently. For the localist, it would be a mistake to think that sentences have incomplete global contents which might be completed in more than one way:

they lack global contents altogether. Accordingly, the role of semantic theories in characterizing the notion of truth is not that of articulating possible completions of the language. Their role is to supply the information about past linguistic usage which is needed to characterize salience. So the localist's attitude towards (S₂) does not warrant the conclusion that there will be a cutoff point for 'bald' no matter how the indeterminate features of the language are rendered determinate. It is also worth keeping in mind that there is no guarantee that the localist will end up embracing a notion of salience like the one I have suggested here. As emphasized above, the localist should use whichever notion of salience is best as a tool for understanding the evolution of the context set. Should she base her proposal on non-classical semantic theories, or choose to cash out the notion of salience in a different way, she might well assign a different semantic status to (S₂).

In addressing the Sorites Paradox, it is not enough to tell a story whereby (S₂) fails to be true. One must also explain our inclination to accept (S₂). It seems to me that the localist is unusually well-placed to supply the necessary explanation. We are tempted to think that (S₂) is true because we make a certain kind of mistake. We think of tolerance—which is a feature of our ability to *use* linguistic representations—as a semantic principle governing the *correctness* of our assertions. This leads us from the unobjectionable observation that we are unable to use 'bald' to discriminate between man n and man $n + 1$ to the mistaken conclusion that 'bald' can only be correctly applied to man n if it is also correctly applied to man $n + 1$. It is easy to make this mistake if one is under the grip of a certain conception of language: the idea that there are semantic rules corresponding to sentences, and that language-mastery is a matter of gaining cognitive access to the relevant rules and learning to apply them in the right sorts of ways. For—as emphasized in Wright 1976—this picture makes it natural to suppose that one can uncover the semantic rules governing our language by straightforward reflection on our usage. But the localist thinks of matters very differently. Language-mastery is not a matter of applying semantic rules; it is a matter of making sensible semi-principled decisions about how to partition the context set in light of past linguistic usage and the location of the gap.

One final point. Even though it contains no linguistic vocabulary, an assertion of, say, 'a baker's dozen is 13' can be used to communicate something about language, namely: that 'baker's dozen' is used to pick out collections of 13 objects. (See Stalnaker 1979.) I have been neglecting pragmatic phenomena of this kind to keep things simple, but it is worth

noting that they are relevant to the case at hand. Even though it contains no linguistic vocabulary, an assertion of (S₂) can be used to communicate something about language: that ‘bald’ has no sharp cutoff-points, for instance. Similarly, (S₂)’s negation might be used to communicate the thought that ‘bald’ has a sharp cutoff-point. A localist might therefore think that (S₂)’s negation can be used to communicate something *false* (on the grounds that ‘bald’ has no cutoff point) even though it should be counted as true if non-defective when pragmatic phenomena are ignored. This complicates the terrain, but I hope it does not obscure the localist’s position.

2.2 *Fine-tuning globalism*

2.2.1 *Contextualism*

On the simplest version of the globalist account, each sentence of the language is assigned a global content independently of context. But globalists might be expected to go beyond the simple account, and claim that there can be significant contextual variation in the global contents assigned to sentences. A sufficiently dramatic increase in contextual variation might get globalism to agree with localism: one might get the result that, in general, the global content that the globalist would ascribe a sentence relative to the context of an assertion of that sentence is an extension of the (local) content that the localist would ascribe the assertion.

It is easy enough to achieve this result by brute force, since one can *define* the global content of a sentence relative to a context of assertion as an extension of the (local) content that the localist ascribes to the assertion. But on the brute force approach globalism does no explanatory work—it is simply a different method of bookkeeping. The interesting question is not whether globalism could be made to agree with localism by postulating greater contextual variation; the interesting question is whether globalism is explanatorily advantageous.

A lot of good work has gone into the project of postulating contextual parameters in non-indexical expressions. Such work is useful to globalists and localists alike. It is useful to globalists because a contextualized version of globalism can fit the facts about usage much more neatly than simple globalism. It is useful to localists because localism relies on semantic theories to identify the patterns of local usage that are used to characterize salience, and the more successful a semantic theory the more information it supplies about patterns of local usage. What is at issue between globalists and localists is not whether postulat-

ing contextual parameters is a good idea; what is at issue is whether the postulation of contextual parameters can do away with the localist idea that the content of an assertion depends on the partition of the context set that the assertion renders salient in light of past usage and the location of the gap.

My suspicion is that some version of the localist idea is inevitable, either as an explicit tenet of one's theory of assertoric content, or as the basis of one's explanation of how contextual parameters get filled in particular assertions. But this is not a matter that needs to be settled for present purposes. The point of putting localism on the table is to give an example of a theory of assertoric content that enables us to make sense of a linguistic practice involving tolerant assertion, not to show that globalism is mistaken.

2.2.2 *Three-way partition views*

On the simple version of the globalist account, assertoric contents partition the space of possibilities in two: possibilities that are compatible with the content and possibilities that are not. But some globalists might wish to hold that assertoric contents partition the space of possibilities in *three*: possibilities that are (definitely) compatible with the content, possibilities that are (definitely) incompatible with the content, and possibilities that are neither. Doing so might bring globalism into agreement with localism: one might get the result that tolerant assertions avert defectiveness (in the localist's sense) by excluding possibilities in the 'neither' category from the context set on which the assertion is to have its essential effect.

As before, it might be possible to achieve this result by brute force, since one might be able to *define* 'neither' possibilities on the basis of whether their exclusion from the context set would yield non-defectiveness. But on the brute force approach the three-way partition of possibilities does no explanatory work—it is simply a different method of bookkeeping. The interesting question is not whether globalism could be made to agree with localism by postulating a three-way partition; the interesting question is whether there is some way of motivating the idea that global contents yield a three-way partition independently of the localist thought that non-defective assertion requires enough of a gap in the context set for a (two-way) partition to be rendered salient.

An independent motivation is not out of the question. Suppose it turns out that some particular range of possibilities is consistently excluded from the context sets on which assertions of sentence *S* in contexts of type τ have their essential effect. Then the globalist might

suggest that the possibilities in that range are placed in the ‘neither’ category by the content of *S* relative to contexts of type τ .

In order for this sort of story to be a genuine improvement over localism, the context-types relative to which three-way partitioning global contents are defined must not be characterized on the basis of whether there is enough of a gap in the context set to render a (two-way) partition salient. My suspicion is that such a characterisation would be hard to find, but the matter need not be settled for present purposes. As emphasized above, the point of putting localism on the table is to give an example of a theory of assertoric content that enables us to make sense of a linguistic practice involving tolerant assertion, not to show that globalism is mistaken.

3. Closing remarks

It seems to me that accounts of vagueness in the literature tend to suffer from a methodological weakness. What they *should* do is follow a bottom-up approach: they should start by getting clear about the connection between linguistic usage and semantic theorising, and from there go on to address the question of what form a semantic theory should take. What they do instead is follow a top-down approach: they start with a certain preconception about the form a semantic theory should take, and proceed to fill out the details so as to maximize agreement with linguistic usage.

The preconception in question is that some version or other of globalism is correct: it is taken for granted that context-sentence pairs determine partitions of the entire space of possibilities. The task of the vagueness theorist is then to explain what the relevant partitions are like, and how they get determined by context-sentence pairs. (Beyond this, there is considerable disagreement: whether partitions are two- or many-fold; whether context-sentence pairs determine their partitions uniquely, or whether there is some sort of indeterminacy; whether complex sentences determine their partitions compositionally or, e.g. supervaluationally; whether contextual parameters or pragmatic considerations should play a significant role in explaining how partitions are determined; and so forth.)

The result of adopting the top-down approach is that one appears to face a dilemma: one must choose between countenancing sharp boundaries, and limiting one’s ability to convey useful information about the nature of vague representation.

An example will help illustrate the point. Suppose one thinks that (relative to a fixed context) there are cases in which ‘Harry is bald’ is definitely true, cases in which it is definitely false and cases in which it is neither. Then one has a choice. One option is to claim that, for each possibility, it is determined whether at that possibility ‘Harry is bald’ counts as definitely true, definitely false, or neither. So, in particular, ‘Harry is bald’ determines a sharp boundary between ‘definitely true’ possibilities and the rest. This leaves one with an explanatory task: the task of explaining why ‘Harry is bald’ carves the space of possibilities just as it does, rather than in a slightly different way. From the standpoint of a top-down approach, it is unclear how such an explanation could be given: it is unclear what sorts of reasons could be adduced for thinking that ‘Harry is bald’ would count as definitely true if, for example, Harry had 41,702 hairs on his head, but not if he had 41,703.

The other option is to invoke the *standard move*: one might claim (a) that the metalinguistic expression ‘is definitely true’ is itself vague, and (b) that this has the effect that one’s theory is not committed to a sharp division between possibilities at which ‘Harry is bald’ is definitely true and the rest. To see why the standard move is objectionable, it is useful to begin with an analogy. You wish to know whether Harry has left the party and I say ‘all the bald people have left’. If Harry is a clear case of baldness, then you can use my utterance to learn that Harry has left the party. But if Harry is a borderline case of baldness, then you are unable to use my utterance to learn whether Harry has left the party. In the latter scenario, my utterance is objectionable as a proposal about Harry’s whereabouts. But the complaint is *not* that the proposal is false. Nor is it that the use of vague language in a proposal about Harry’s whereabouts is somehow unacceptable (after all, the same sentence was uttered in the former scenario and there you got the information you needed). What is objectionable is that the vagueness gets in the way: it interferes with your ability to use my utterance to determine whether or not Harry is at the party.

Similarly, the complaint about the about the standard move is *not* that the resulting proposal is false. Nor is it that the use of a vague metalanguage is somehow unacceptable (it is presumably impossible to come up with a *non*-vague metalanguage, after all). The complaint is that the vagueness postulated in the metalanguage gets in the way: it interferes with one’s ability to use the proposal to gain useful information about the nature of the transition between possibilities at which ‘Harry is bald’ counts as definitely true and possibilities in which it does not. The result of claiming that ‘is definitely true’ is vague is *not* that

one's theory sheds light on the nature of the transition. The result is rather that one is unable to use the proposal to learn what the transition is like. In particular, one is unable to use the proposal to learn whether or not 'Harry is bald' determines a sharp partition of the space of possibilities. And by failing to learn whether or not a sharp partition is determined one fails to learn something important about the nature of vague representation. In general, the effect of applying the standard move is that the resulting proposal fails to convey useful information about the most interesting parts of the transition between possibilities at which a sentence counts as having a particular status and possibilities at which it does not, since the blind-spots will be located at just the places where the proposal would be forced to countenance sharp boundaries were it not for the newly postulated vagueness.

The standard move can only succeed in avoiding the problem of accounting for the location of sharp boundaries at the cost of crippling one's ability to shed light on the nature of vague representation. So, from the perspective of the top-down approach, one would appear to be trapped between a rock and a hard place.

Now consider how things look from the perspective of the bottom-up approach. If the discussion in this paper is along the right lines, attention to linguistic usage recommends a localist theory of content, rather than a globalist one. And on a localist theory of content, the globalist's dilemma is suddenly rendered tractable. For when one is only concerned with partitions of on limited sets of possibilities—sets containing a *gap* of the right kind—the location of a the boundary can be easily explained on the basis of our linguistic practice.

There is, of course, the residual problem of accounting for the boundary between sets of possibilities relative to which content is defined and the rest. But this problem loses its bite when a semantic theory is thought of as a *tool* for understanding the evolution of the context set. For then the question of where to locate the relevant boundary is in part a pragmatic decision on the part of the theorist, rather than something that needs to be settled by the world. The decision might depend on the theorist's theoretical aims—for example, on the relative importance of reliability and breadth—but it might also involve some degree of arbitrary choice. (Sometimes one needs to settle things one way or another, even though it is unimportant how one settles them: think of using sets to define ordered-pairs.)

An instrumentalist attitude of this kind is also open to the globalist. But in order for it to do genuine explanatory work, it is crucial that one clarify which aspects of the theory are to be settled by looking at the

world and which ones can be decided on the basis of one's theoretical aims (or by arbitrary choice)—and that one explain how pragmatic decisions affect what one's theory tells us about linguistic practice. The cost of adopting a top-down approach is that one embarks on the task of accounting for vagueness before getting clear about such issues.⁹

Department of Linguistics and Philosophy
Massachusetts Institute of Technology
77 Massachusetts Ave 32-D932
Cambridge, MA 02139
USA
arayo@mit.edu

AGUSTÍN RAYO

References

- Bräuerle, Rainer et al. (eds) 1979: *Semantics from Different Points of View*. Berlin & New York: Springer-Verlag.
- Dorr, Cian 2003: 'Vagueness without Ignorance'. *Philosophical Perspectives*, 17, pp. 83–113.
- Evans, Gareth and J. McDowell (eds) 1976: *Truth and Meaning*. Oxford: Clarendon Press.
- Graff, Delia 2000: 'Shifting Sands: An Interest-Relative Theory of Vagueness'. *Philosophical Topics*, 28, pp. 45–81.
- Grice, Paul 1989a: 'Logic and Conversation'. In Grice 1989b, pp. 22–40.
- 1989b: *Studies in the Ways of Words*. Cambridge, MA: Harvard University Press.
- Keefe, Rosanna and P. Smith (eds) 1996: *Vagueness: A Reader*. Cambridge, MA: MIT Press.
- Laurence, Stephen 1996: 'A Chomskian Alternative to Convention Based Semantics'. *Mind*, 105, pp. 269–301.

⁹ I am deeply indebted to a number of texts, not all of which are referenced in the discussion above. I would especially like to mention Lewis 1969, Wright 1976, Stalnaker 1979, Sainsbury 1990, McGee and McLaughlin 1994, Williamson 1994, Laurence 1996, Graff 2000, and Dorr 2003. Less direct inspiration comes from the later Wittgenstein.

For their many helpful comments, I am grateful to Elizabeth Barnes, Eliza Block, Liz Camp, Jonathan Cohen, Juan Comesaña, Roy Cook, Adam Elga, Michael Glanzberg, Patrick Greenough, Oystein Linnebo, Michael Rescorla, Elliott Sober, Bob Stalnaker, Robert Williams, Tim Williamson, Crispin Wright and an editor and two anonymous readers for *Mind*. I am also grateful to audiences at the universities of Bologna, California at Davis, Glasgow, Monash, Princeton, St Andrews, Warsaw, and Wisconsin at Madison, as well as the Bellingham Summer Philosophy Conference, the Jagiellonian University, MIT, the Ohio State University at Columbus, the Southern California Philosophy Conference, and to seminar participants at *Arché* Philosophical Research Centre for Logic, Language, Metaphysics and Epistemology, MIT, the National Autonomous University of Mexico, and the University of California at San Diego.

- Lewis, David K. 1969: *Convention: A Philosophical Study*. Cambridge, MA: Harvard University Press.
- 1979: 'Score-Keeping in a Language Game'. *The Journal of Philosophical Logic*, 8, pp. 339–59. Reprinted in Bräuerle et al. 1970, and in Lewis 1983.
- 1983: *Philosophical Papers, Volume I*. Oxford: Oxford University Press.
- McGee, Vann 1996: 'Logical Operations'. *The Journal of Philosophical Logic*, 25, pp. 567–80.
- McGee, Vann and B. McLaughlin 1994: 'Distinctions without a difference'. *The Southern Journal of Philosophy*, XXXIII, pp. 203–51.
- Sainsbury, R. M. 1990: 'Concepts without Boundaries'. Inaugural Lecture, given at King's College London. Reprinted in Keefe and Smith 1996, pp. 251–64.
- Sher, Gila 1991: *The Bounds of Logic: A Generalized Viewpoint*. Cambridge, MA: MIT Press.
- Stalnaker, Robert C. 1979: 'Assertion'. *Syntax and Semantics*, 9, pp. 315–22. Reprinted in Stalnaker 1999, pp. 78–95.
- 1998: 'On the Representation of Context'. *Journal of Logic, Language, and Information*, 7, pp. 3–19. Reprinted in Stalnaker 1999, pp. 96–113.
- 1999: *Context and Content*. Oxford: Oxford University Press.
- Tarski, Alfred 1986: 'What are Logical Notions?' *History and Philosophy of Logic*, 7, pp. 143–54.
- Williamson, Timothy 1994: *Vagueness*. London and New York: Routledge.
- Wright, Crispin 1976: 'Language-Mastery and the Sorites Paradox'. In Evans and McDowell 1976, pp. 223–47. Reprinted in Keefe and Smith 1996, pp. 151–73.